Universitext

Ernest E. Shult

Points and Lines

Characterizing the Classical Geometries

Springer

Universitext

For further volumes:
http://www.springer.com/series/223

Ernest E. Shult

Points and Lines

Characterizing the Classical Geometries

Springer

Ernest E. Shult
Kansas State University
Denison Avenue 419
Manhattan KS 66502
USA
ernest_shult@yahoo.com

Editorial board:
Sheldon Axler, San Francisco State University
Vincenzo Capasso, Università degli Studi di Milano
Carles Casacuberta, Universitat de Barcelona
Angus J. MacIntyre, Queen Mary, University of London
Kenneth Ribet, University of California, Berkeley
Claude Sabbah, CNRS, École Polytechnique
Endre Süli, University of Oxford
Wojbor A. Woyczynski, Case Western Reserve University

ISBN 978-3-642-15626-7 e-ISBN 978-3-642-15627-4
DOI 10.1007/978-3-642-15627-4
Springer Heidelberg Dordrecht London New York

Mathematics Subject Classification (2010): 51A45, 51A50, 51A05

© Springer-Verlag Berlin Heidelberg 2011
This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilm or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable to prosecution under the German Copyright Law.
The use of general descriptive names, registered names, trademarks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

Cover design: deblik, Berlin

Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)

Professor Otto Kegel
Professor Jacques Tits
The von Humboldt Foundation

Preface

This book is about characterizing the classical geometries of Lie type by simple axioms on point and lines.

If this book were written for experts only, it would be half as long as it is. Instead it is a teaching book.

I have tried to write this book so that anyone – knowing very little of Incidence Geometry – could actually learn all about it. Ideally, without losing any connecting rapport, the graduate student beginner could be led on a path beginning at an elementary level involving familiar objects (graphs and groups, for example) right up to the research level. Of course that requires three things:

1. A vast amount of definitions, and lengthy explanations of where theorems are leading and why one *should* be led by them.
2. A presentation which, at each stage, seeks the most pedagogically efficient path for one whose only background is that obtained from preceeding stages. As a result, the route taken is somewhat different from standard treatments. (For example in the presentation of Tierlinck's Theory, or in the definition of "building.")
3. A self-contained account. As a result, the student will find almost everything proved in detail.[1] In this sense the book also serves partially as a "source-book," since complete proofs of certain basic classical theorems will now be found in one place – for example, complete proofs of the infinite-rank versions of the

[1] Indeed one can easily list the invoked results which are *not proved* in this book:

(a) Tits' classification of the rank three non-degenerate polar spaces.
(b) The proof of the existence of buildings of any reasonable type (Ronan and Tits: "blueprints").
(c) The theorem of Coxeter classifying the finite Coxeter groups.
(d) Cohen's technical theorem showing that no point residual of a rank three polar space can be the bizarre sort of quadrangle that results from a counterexample involving tentative Grassmann spaces. This may be a technical result, but it is a vital one. His proof depends on (a), so it is a deep theorem. However there is a proof due to Koen Thas using only the fact that any generalized quadrangle that is a residue of a polar space of higher rank must be Moufang [118].

Veblen-Young Theorem and the Tits/Cuypers-Pasini-Johnson classification of polar spaces of rank of at least four, as well as Tits "Local-approach Theorem".

Yes, there are many pages of detailed proofs in this book that would try the patience of any researcher that already knows them. The author hopes that experienced researchers will accustomize themselves to skipping such sections without hesitation. There is much in Parts III and IV that might well interest the expert – whether it be a new presentation of standard theories (Chap. 9 (chamber systems and buildings), Chap. 11 (locally-truncated geometries and sheaves) and Chap. 13 (parapolar spaces)) or new results scattered alongside older ones (mostly in Chaps. 14, 15, 16, 17, and 18)).

Of course it would take several years to teach everything in this book. But over the years many parts of the book have served the author as semester courses in themselves. Permit me to list a few:

1. Covers of graphs, geometries and chamber systems (Chaps. 1, 2, early sections of Chap. 9).
2. A complete classification of projective spaces, using Teirlinck's Theory (Chaps. 4, 5, and 6).
3. A complete teachable classification of polar spaces (Chap. 7).
4. Everything one should know about chamber systems, and the characterizations of buildings (Chap. 9). Perhaps one could include the short Chap. 10 (a fleshed-out account of Tits' "Local Approach Theorem"), if there is time.

It is fair to state what is and is not in this book. Not so long ago geometers described what they are doing under two distinct headings: (1) *Synthetic Geometry* and (2) *Descriptive Geometry*. My impression of the distinction is this: in (1), one wants to characterize objects by axioms, while in (2) one wishes to understand completely the objects which nature so pervasively places before us. The distinction has become fuzzy on several levels.

On the research level, how can anyone show why these imortant classical objects are so pervasive if one cannot recognize the curious properties that bring them to our attention?

On the pedagogical level, how can a teacher introduce concepts in the classroom without rooting the student's intuition in constructs that reduce to solid familiar examples? So both the researcher and the teacher (who are often the same person) must live in both worlds.

The author (an intruding novice from Finite Group Theory) cannot be sure, but guesses that this book comes down heavily on the side of the synthetic point of view. The goal of the book is to characterize things by axioms – and it barely rises above familiar examples derived from graphs, vector spaces, and groups to do so.

It follows that there is another side that this book does not even touch. The objects being characterized are not proved to exist in this book (one of the handful of items listed as not proved in the footnote above), and when they do there is not a survey of the models which display them. One will find nothing about algebraic groups in this book. One cannot even detect a mention of Lie Algebras anywhere in it!

How is a student to learn about triality? These are important subjects that I hope the student-reader of this book will learn elsewhere. And there are sources: to borrow a paradigm from the dieticians, I think of this book and the two recent books of Buekenhout-Cohen [22] and Weiss [145] as forming a sort of "pedagogical food pyramid" providing a basic menu of classical incidence geometry.

There is much more that is not here. I will take a few moments to describe other directions a curious student might think of exploring.

Different fields of geometry ask different questions. In recent times one of the most exciting fields has been finite geometry.

- Finite buildings of rank at least three are not difficult to handle. But those of rank two are extremely difficult. There is a vast history here and a student can get an idea of what is going on by consulting the books of Payne and Thas [98] (for generalized quadrangles), and the book of Van Maldeghan [141] (for all finite generalized polygons) and later on, Tits-Weiss [140] (about infinite Moufang polygons).
- Most of finite geometry, however, has been involved in the study of what I call "strange curves." These would be weird subsets of classical objects with extreme properties. Many of these have led to a strange nexus of interconnected objects. What J.-C. Rota once called a "cryptomorphism" is daily bread here. Let us list a few:

  1. Blocking sets, maximal arcs, and related MDS codes.
  2. Ovoids of planes, PG(3)s, and polar spaces and the constructions of translation planes
  3. A vast interconnection relating BLT-sets, flocks of a cone, Kantor families, q-clans, herds, and constructions of both generalized quadrangles and certain projective planes. (Some of the connections were first discovered by noticing that the same equations arose in two different contexts.) More recently, this has occurred with a new generation of objects. Again, the connections (called the "silver thread") are as mysterious as anything discovered by Ramanujan – another person who paid attention to nature's details.
  4. m-Systems, spreads, etc.

- There has been no attempt to relate the characterizations of Part IV to the workers in continuous geometry who study ruled manifolds. It just seems amazing to me that so many mathematicians and physicists who deal with symmetric spaces and ruled manifolds seem to be unaware of the simple point-line characterizations of these hypersurfaces.

  Instead of simply guessing what Lie group should control a Candidate String Theory, one might ask what axioms such a theory forces upon a light trajectory? Such an approach, if carefully developed, would be directly pointed toward the point-line geometries of Part IV.

  In fact that raises a strange question. Anyone reading this book will realize that locally described objects are at best homomorphic images of objects which are universal covers. Why doesn't one hear this in physics? Is it possible that

our world, rather than being a building determined by the Lie Algebra of some Hamiltonian, is in fact a homomorphic image of such an object? (After all, the Hamiltonian is nothing more than the *Diary Of A Particle* – as seen from the point of view of the particle.) Embracing the *full* centralizer algebra of this Casimir datum might be a leap of faith? Why assume the universe is a simply connected building rather than a locally truncated geometry (Chap. 13)?

Of course that question is also not answered in this book or anywhere else that I know of. I just stuck it in there.

So there is a lot of geometry that is not here – or anywhere.

Teaching, in its own way, is an act of love. It is the ideal love of the truths that a student's mind can understand. Nothing more or less.

This book is dedicated to two mathematicians who had so much to do with this book and to the von Humboldt Foundation (whose devotion to human enlightenment is simply unparalleled).

During my von Humboldt Fellowship in Freiburg 1987–1988, Professor Otto Kegel suggested that I teach a course about buildings. My hand-outs became a set of unpublished notes that had been passed around (in the text they are pretentiously called the *Freiburg Notes* as if everyone knew about them). But after several incarnations as classroom notes for courses, they eventually became the skeleton for this book. Without that opportunity for growth, I would never have written this book which is dedicated to Professor Kegel, as well as one other, Professor Jacques Tits.

To this author, Tits' theory of buildings is certainly a candidate for the greatest theorem of the twentieth century! Almost nothing in this book would be possible without Professor Tits' amazing insight (I confess that he has been my intellectual ideal for many years). So this book is dedicated to him as well.

Manhattan, KS Ernest E. Shult

Contents

Part I Basics

1 Basics About Graphs 3
1.1 The Language of Graphs 3
1.1.1 Connectedness and the Distance Metric 4
1.1.2 Subgraphs 5
1.1.3 Special Types of Graphs and Subgraphs 6
1.1.4 Metrical Properties of Subgraphs 9
1.2 Morphisms of Graphs 14
1.2.1 Basic Definitions 15
1.2.2 Fiberings, Covering Morphisms, and Lifts of Walks 16
1.2.3 Universal $\mathcal{C}$-Covers 18
1.3 $\mathcal{C}$-Homotopy 19
1.3.1 Further Properties of the Closure Operator on Circuits . . . 21
1.3.2 Control of $\mathcal{C}$-Connectedness Through a Subgraph 24
1.3.3 Tits' Condition for Being Simply $\mathcal{C}$-Connected 26
1.4 The Existence of Universal $\mathcal{C}$-Covers 31
1.4.1 $\mathcal{C}$-Connectedness and Subgraphs 33
1.4.2 The Construction 34
1.4.3 Deck Transformations 36
1.5 Exercises for Chapter 1 37
1.5.1 Exercises for Section 1.1 37
1.5.2 Exercises for Section 1.2 40
1.5.3 Exercises for Section 1.3 41

2 Geometries: Basic Concepts 43
2.1 Introduction 43
2.2 Geometries: Definitions and Basic Concepts 43
2.2.1 Basic Definitions 43
2.2.2 Subgeometries 44
2.2.3 Truncations 44
2.2.4 Point-Line Geometries 45

2.2.5 Flags and Chambers . . . 45
2.2.6 Residues . . . 45
2.2.7 The Interplay of Residues and Truncations . . . 46
2.2.8 Shadows . . . 46
2.3 Examples . . . 47
2.4 Morphisms of Geometries . . . 50
2.4.1 Definition . . . 50
2.4.2 Automorphisms . . . 50
2.4.3 Morphisms Defined by a Group of Automorphisms . . . 52
2.4.4 Truncations and Morphisms . . . 53
2.4.5 Residues and Morphisms . . . 53
2.5 Connectedness Properties . . . 54
2.5.1 Residual Connectedness . . . 54
2.6 Exercises for Chapter 2 . . . 57

3 Point-Line Geometries . . . 59
3.1 Introduction . . . 59
3.1.1 On Choosing a Reasonable Definition . . . 59
3.1.2 Our Definition of Point-Line Geometry . . . 60
3.2 The Point-Collinearity Graph . . . 61
3.3 Morphisms and Covers of Point-Line Geometries . . . 61
3.4 Subspaces . . . 62
3.4.1 Generalizations of the Notion of Subspace . . . 65
3.5 Special Types of Point-Line Geometries . . . 66
3.5.1 Partial Linear Spaces and Linear Spaces . . . 66
3.5.2 Gamma Spaces . . . 66
3.6 Local Connectedness in Gamma Spaces . . . 67
3.6.1 The Decomposition of a Gamma Space into Locally Connected Components . . . 68
3.6.2 How Local Characterizations of Gamma Spaces Reduce to the Locally Connected Case . . . 69
3.7 Enriching Geometries . . . 70
3.8 Products of Point-Line Geometries, a Construction . . . 71
3.9 Exercises and Examples for Chapter 3 . . . 72

4 Hyperplanes, Embeddings, and Teirlinck's Theory . . . 79
4.1 Veldkamp Spaces . . . 79
4.1.1 Introduction . . . 79
4.1.2 Geometric Hyperplanes . . . 80
4.1.3 The Veldkamp Space . . . 81
4.2 Teirlinck's Theory . . . 82
4.2.1 The Hypotheses . . . 82
4.2.2 The Exchange Property in Closed Sets . . . 83

4.2.3 To What Extent Does $\mathcal{H}$ Separate Points? ... 84
4.2.4 The $\mathcal{H}$-Closure of Two Inequivalent Points ... 85
4.2.5 The Natural Morphism ... 86
4.2.6 Singular Subspaces are Generalized Projective Spaces ... 87
4.3 The Effect of Teirlinck's Theory on the Veldkamp Space ... 89
4.3.1 The Veblen Axiom ... 89
4.3.2 The First Reason Why a Veldkamp Space Might be a Projective Space ... 91
4.3.3 Another Reason why the Veldkamp Space $\mathcal{V}_{\mathcal{H}}$ Might be a Projective Space ... 93
4.4 The Connections Between Projective Embeddings and Veldkamp Spaces ... 94
4.5 Teirlinck's Theory and Embedability into the Second Veldkamp Space ... 96
4.6 Exercises for Chapter 4 ... 98
4.6.1 Exercises on Geometric Hyperplanes ... 98
4.6.2 Exercises Relating to Veldkamp Spaces ... 100
4.6.3 Exercises Relating to Teirlinck's Axioms ... 101

Part II The Classical Geometries

5 Projective Planes ... 105
5.1 Introduction ... 105
5.2 The Axioms for a Projective Plane ... 106
5.3 Affine Planes ... 109
5.4 The Relation Between Affine and Projective Planes ... 111
5.5 Translation Planes ... 112
5.5.1 Introduction ... 112
5.5.2 Translation Planes from Groups ... 113
5.5.3 Endomorphism Rings ... 115
5.5.4 The Kernel of $(G, \mathcal{S})$ ... 116
5.5.5 A Short Course on the Klein Correspondence ... 118
5.5.6 Ovoids and Spreads ... 119
5.6 Recognizing the Classical Planes ... 120
5.6.1 Perspective Collineations ... 120
5.6.2 The Moufang Case ... 123
5.6.3 The Desarguesian Case ... 124
5.6.4 The Proof of Theorem 5.6.8 ... 126
5.7 Appendix to Chapter 5 ... 127
5.7.1 Introduction ... 127
5.7.2 Planes from Ovoids: A Den of Ubiquity ... 128
5.8 Exercises for Chapter 5 ... 133

6 Projective Spaces 135
6.1 Introduction 135
6.2 Basic Properties of Projective Spaces and Their Partially Ordered Systems of Subspaces 136
6.2.1 The Definition of a Projective Space 136
6.2.2 Matroids 138
6.2.3 The Matroid of a Projective Space 140
6.3 Subplanes of Larger Projective Spaces 141
6.4 Embedding Three-dimensional Projective Spaces 143
6.4.1 Affine Spaces 143
6.4.2 The Classical Affine Space $AG(n, D)$ 145
6.4.3 Completing an Arbitrary Affine Space to its Associated Projective Space 147
6.5 Embedding Three-dimensional Projective Spaces, Some of Whose Planes are Classical 148
6.5.1 Introduction 148
6.5.2 Synthesizing Coordinates: A Basic Lemma 148
6.5.3 Coordinatizing P Minus a Desarguesian Plane 149
6.6 Extending Projective Embeddings of Projective Spaces 153
6.7 The Main Theorem: All Projective Spaces of Projective Dimension at Least Three are Classical 155
6.8 The Fundamental Theorem of Projective Geometry 158
6.9 Three Technical Results Concerning Projective Spaces 160
6.10 Exercises for Chapter 6 163
6.10.1 Exercise on the Definition of Projective Space 163
6.10.2 Axioms of Affine Space 163
6.10.3 Simple Results on Embedded Projective Spaces 165

7 Polar Spaces 167
7.1 Introduction 167
7.2 Generalized Quadrangles 170
7.3 Basics of Polar Spaces 173
7.3.1 Definition 173
7.3.2 Opposite Lines and the Radical 173
7.3.3 The Basic Morphism 174
7.3.4 Basic Properties of Non-degenerate Polar Spaces 175
7.3.5 The Rank of a Polar Space 176
7.4 Oriflame Polar Spaces of Rank Three 178
7.5 Further Properties of Abstract Polar Spaces 182
7.5.1 Geometric Hyperplanes of Polar Spaces 182
7.5.2 Hyperplane Complements in Polar Spaces of Rank at Least Three 184
7.5.3 Automorphisms of Non-degenerate Polar Spaces 186

7.6 The Classical Polar Spaces 192
7.6.1 Morphisms from a Projective Space to Its Dual, or Where Sesquilinear Forms Come From 192
7.6.2 Quasi-Polarities 195
7.6.3 The Relation Between Quasi-Polarities and Reflexive Forms 197
7.6.4 Admissible Pairs and Gramm Matrices 202
7.6.5 PseudoQuadratic Forms 207
7.6.6 The Polar Space of a Pseudoquadratic Form 209
7.7 Polar Spaces Embedded in Projective Spaces 211
7.7.1 Introduction 211
7.7.2 The Geometry of Embedded Polar Spaces 212
7.7.3 The Main Theorem 217
7.8 Polar Spaces Embedded in a Classical Polar Space 218
7.9 Non-degenerate Polar Spaces of Rank at Least Four Are Classical 225
7.9.1 Introduction 225
7.9.2 Teirlinck's Condition in $\mathcal{V}_p$ 226
7.9.3 The Classification of Non-degenerate Polar Spaces of Rank at Least Four 230
7.10 Epilogue 232
7.10.1 Non-embeddable Rank Three Polar Spaces 232
7.10.2 Tits' Classification Theorem 234
7.10.3 The Buekenhout-Johnson-Shult Theorem 236
7.11 Exercises for Chapter 7 237
7.11.1 Exercises Involving Only Basic Properties 237
7.11.2 Exercises on Classical Polar Spaces 239
7.11.3 Exercises from Section 7.9 242
7.12 Appendix to Chapter 7: Known Quadrangles in 2001 243
7.12.1 The Older Examples 243
7.12.2 The Last 15 Years Before 2001: Kantor-Families, the Heisenberg Group and Flocks of a Quadratic Cone . . . 246
7.12.3 Prospects for a Classification 248

8 Near Polygons 251
8.1 Introduction 251
8.1.1 What We Know About Polar Spaces 251
8.1.2 The Definition of Near Polygons 252
8.1.3 Some Non-classical Examples 253
8.1.4 Near Polygons Arising from Chamber Systems of Buildings 257
8.1.5 Near Polygons of Fischer Type 257
8.1.6 Classical Near Polygons: The Dual Polar Spaces 258

8.2 The Existence of Quads in Near Polygons 261
8.2.1 Internal and External Diameter 261
8.2.2 Bounded Diameter Point-Sets in Near Polygons 262
8.2.3 The Existence of Quads 263
8.3 The Convexity of Quads 264
8.4 The Geometry of Quads 266
8.4.1 The Diagram 266
8.4.2 The Relations Between Quads 267
8.4.3 The Relations Among Points and Quads 267
8.5 A Characterization of Dual Polar Spaces 270
8.5.1 The Axioms 270
8.5.2 The Structure of the Residue of a Point 270
8.5.3 The (Global) Convex Subspace Structure 273
8.5.4 The Dual Polar Space Structure in Near Polygons of Finite Diameter 279
8.5.5 Recapitulation 280
8.6 Exercises for Chapter 8 281

Part III Methodology

9 Chamber Systems and Buildings 291
9.1 Introduction 291
9.2 Chamber Systems 293
9.2.1 The Chamber System of a Geometry 293
9.2.2 Abstract Chamber Systems 294
9.2.3 Residually Connected Chamber Systems 303
9.3 Chamber Systems with Strongly Gated Residues 310
9.3.1 Introduction 310
9.3.2 Basic Properties Concerning Strongly Gated Residues . . . 310
9.3.3 Intersections of Strongly Gated Residues 312
9.3.4 2-Simply Connectedness is a Consequence of Strong Gatedness at Low Rank 314
9.3.5 Λ-Homotopy 316
9.3.6 Further Consequences of the Hypothesis (RG_2) 318
9.3.7 Equivalence of Various Gatedness Conditions 323
9.4 Generalized Polygons 323
9.4.1 Panel Homotopy 323
9.4.2 The Chamber System of a Generalized Polygon 326
9.4.3 Generalized n-Gons as Geometries 334
9.4.4 Existence of Generalized Polygons 342
9.5 Diagrams 345
9.5.1 Introduction 345
9.5.2 Rank Two Diagrams 346
9.5.3 Diagram Geometries of Higher Rank 347

9.5.4 Chamber Systems Belonging to a Diagram 348
9.5.5 Diagrams and the Functors Connecting Chamber Systems and Geometries 349
9.5.6 Some Examples Concerning Diagram Geometries and Chamber Systems 349
9.6 Chamber Systems with a Coxeter Diagram 350
9.6.1 Coxeter Groups and Coxeter Systems 350
9.6.2 The Cayley Graph of the Coxeter System (W, R, M), and the Coxeter Chamber Systems 353
9.6.3 Other Properties of Coxeter Chamber Systems 361
9.6.4 Walls, Roots, and Distance in a Coxeter Chamber System 363
9.6.5 Gatedness and Convexity of Residues 367
9.6.6 When Is a Coxeter Group Finite? 370
9.7 Chamber Systems of Type M 371
9.7.1 Introduction 371
9.7.2 The Three Levels of Homotopy 373
9.8 Buildings 377
9.8.1 Introduction 377
9.8.2 The Conditions (G_c) and (P_c) 378
9.9 Apartments 383
9.9.1 The Tits Metric of a Building 383
9.9.2 Strong Isometries and the Standard Apartment Axioms for a Building 386
9.10 Appendix to Chapter 9: Spherical Buildings and (B, N)-Pairs 387
9.10.1 Tits Systems 387
9.10.2 (B, N)-Pairs and Tits Systems 388
9.10.3 Sphericity 389
9.11 Exercises for Chapter 9 391
9.11.1 Exercises on Chamber Systems 391
9.11.2 Exercises on Residual Connectedness 393
9.11.3 A Few Exercises on Gatedness 393
9.11.4 Exercises on Generalized Polygons 394
9.11.5 An Interplay of Examples and Exercises on Diagram Geometries 395
9.11.6 Exercises Concerning Chamber Systems of Type M 397

10 2-Covers of Chamber Systems 399
10.1 Introduction 399
10.2 Kasikova's Lemma on 2-Coverings of Chamber Systems 401
10.3 An Important Property of Buildings 403
10.4 When Are Universal 2-Covers Buildings? 403
10.5 The Structure of an M-Homotopy Class 408
10.6 Tits' "Local Approach" Theorem 412

11 Locally Truncated Diagram Geometries 415
11.1 Locally Truncated Geometries: A Quick Review 415
11.2 The Idea of Sheaves 416
11.2.1 Flags and Residues 416
11.2.2 Sheaves 417
11.2.3 The Chamber System Associated with a Sheaf 417
11.3 Sheaves for Locally Truncated Geometries 419
11.3.1 Modelling Geometries on Truncations 419
11.3.2 Geometric Rigidity 420
11.3.3 Constructing Sheaves 421
11.4 Recovering Γ from $C(\mathcal{F})$ 427
11.5 Residual Connectedness of Γ and $(C(\mathcal{F}))_J$ 429
11.6 Configurations Produced by a 1-Covering of $C(\mathcal{F})$ 431
11.7 Applying the Local Approach Theorem: A Useful Tool for Point-Line Characterizations 433
11.8 A Covering of a Point-Collinearity Graph Derived from a Locally Truncated Geometry 434
11.9 Exercises for Chapter 11 437
11.9.1 Exercises Involving Connectedness Properties 437
11.9.2 Morphisms by Chamber-Semiregular Automorphism Groups Do Not Preserve Residual Connectedness 438

12 Separated Systems of Singular Spaces 441
12.1 Introduction 441
12.1.1 The Basic Context: Paraprojective Spaces 441
12.1.2 Local and Global Hypotheses 442
12.1.3 Separated Systems of Singular Subspaces 443
12.2 Geometries with Two Systems of Subspaces Meeting at Lines 444
12.2.1 Introduction 444
12.2.2 Geometries with Two Systems of Subgeometries 444
12.2.3 A Purely Local View 447
12.3 A Criterion for Global Separation of Singular Subspaces 447
12.3.1 Examples Illustrating Local Separation Without Global Separation 447
12.3.2 A Sufficient Condition for Local Separation to Imply Global Separation 450
12.4 Exercises for Chapter 12 452

13 Cooperstein's Theory of Symplecta and Parapolar Spaces 455
13.1 Introduction 455
13.2 Polar Families 455
13.3 Elementary Properties of Symplecta 461
13.4 Parapolar Spaces: A Stage for Characterizing the Lie Incidence Geometries 462

13.4.1 The Definition of Parapolar Spaces 462
13.4.2 Examples of Parapolar Spaces 465
13.5 Local Connectedness in Parapolar Spaces 466
13.5.1 Introduction 466
13.5.2 Isolating Polar Spaces from Locally Connected Parapolar Spaces 471
13.6 Hanssens' Principle 472
13.7 When Is Either $\Delta_2^*(p)$ or $P_2^*(p)$ a Subspace of a Parapolar Space? 478
13.8 Generalized Parapolar Spaces 482
13.8.1 Introduction 482
13.8.2 Generalized Parapolar Spaces of Nexus 1 486
13.8.3 Generalized Parapolar Spaces of Higher Nexus 487
13.9 Exercises for Chapter 13 488
13.9.1 Exercises Concerning Parapolar Spaces 488
13.9.2 Exercises Concerning an Application of Hanssens' Principle 491
13.9.3 Exercises for Section 13.7 493
13.9.4 Exercises for Section 13.8: On Generalized Parapolar Spaces of Higher Nexus 493
13.9.5 Symplecta Which Cannot Intersect at a Single Point 493

Part IV Applications to Other Lie Incidence Geometries

14 Characterizations of the Classical Grassmann Spaces 497
14.1 The First Characterization 497
14.1.1 Axioms on Two Families of Maximal Singular Subspaces 497
14.2 A Variation for Half-Grassmann Spaces 506
14.2.1 Consequences of Axioms (G1) and (G2) 507
14.2.2 The Effect of the Other Axioms 510
14.3 Sprague's Theory 511
14.3.1 Introduction 511
14.3.2 (G1)-Geometries with the Veblen and Dual-Veblen Axioms, (G2) and (H3) 513
14.3.3 Adding the Finiteness Conditions and the Proof of Theorem 14.3.1 519
14.4 Exercises for Chapter 14 523

15 Characterizing the Classical Strong Parapolar Spaces: The Cohen–Cooperstein Theory Revisited 527
15.1 Prognosis 527
15.2 Introduction 528
15.3 Cohen's Characterization of Grassmann Spaces 530

15.3.1 A Technical Result . . . 530
15.3.2 Cohen's Hypotheses and Their Elementary Consequences . . . 530
15.4 Strong Parapolar Spaces with Constant Symplectic Rank $r > 3$ and Satisfying $(CC)_{r-2}$ . . . 535
15.4.1 Introduction . . . 535
15.4.2 The Case $r = 4$ . . . 536
15.4.3 The Case $r = 5$ . . . 537
15.4.4 The Case $r = 6$ . . . 538
15.4.5 The Case $r > 6$ . . . 540
15.4.6 The First Cohen–Cooperstein Theorem . . . 540
15.5 The Uniformizing Condition . . . 541
15.5.1 Introduction . . . 541
15.5.2 A Basic Lemma . . . 541
15.5.3 The Principle . . . 542
15.5.4 The Uniformizing Principle and Constant Symplectic Rank . . . 543
15.6 Exercises for Chapter 15 . . . 545
15.6.1 General Features . . . 545
15.6.2 Understanding the Exceptional Strong Parapolar Spaces . . . 545
15.6.3 Involutions Acting on Classical Strong Parapolar Spaces . . . 548

16 Characterizing Strong Parapolar Spaces by the Relation Between Points and Certain Maximal Singular Subspaces . . . 553
16.1 Introduction . . . 553
16.2 Parapolar Spaces with the "Clean Hypothesis" . . . 554
16.2.1 Elementary Consequences of the "Clean Hypothesis" . . . 555
16.2.2 The Third Case of the Trichotomy . . . 557
16.3 The Second Case of the Trichotomy . . . 558
16.3.1 The Maximal Singular Subspaces . . . 558
16.3.2 The Symplecta . . . 560
16.3.3 More Maximal Singular Subspaces . . . 561
16.3.4 The Final Result for Case 2 with (H2) . . . 563
16.4 The Final Case 1: $d = 0$ . . . 564
16.4.1 Fibered Symplecta . . . 565
16.4.2 Unfibered Symplecta: M-Projections which Embed a Symplecton . . . 567
16.4.3 The Case that Some Member of $\mathcal{M}$ is a Line . . . 569
16.4.4 Symplecta Disjoint from No Member of $\mathcal{M}$ . . . 570
16.4.5 When All Symplecta are Grids . . . 571
16.4.6 The Final Result for $d = 0$ . . . 573
16.5 Proof of Theorem 16.2.1 . . . 575

16.6 A Corollary of Theorem 16.2.1 575
16.6.1 Uniformity of Point-Residuals 576
16.6.2 The Proof of Corollary 16.6.1 578
16.7 A Further Corollary of Theorem 16.2.1 581
16.8 Appendix to Chapter 16 596
16.8.1 Background 596
16.8.2 Admissible Triples 596
16.8.3 Constructions of the Relevant Parapolar Spaces 598

17 Point-Line Characterizations of the "Long Root Geometries" 603
17.1 Introduction 603
17.1.1 Long Root Geometries 603
17.1.2 The Main Results of this Chapter 605
17.2 A Special Class of Strong Parapolar Spaces 606
17.3 The Proof of Theorem 17.1.1 617
17.3.1 Simple-connectedness of the Point-Collinearity Graph of Γ 617
17.3.2 The Uniform Structure of the Point-Residuals 619
17.3.3 The Case of Finite Singular Rank 619
17.3.4 The Case of Infinite Singular Rank 621
17.4 Exercises for Chapter 17 625
17.5 Appendix to Chapter 17: Long Root Geometries Updated 626

18 The Peculiar Pentagon Property 627
18.1 Introduction 627
18.1.1 The Pentagon Axiom; Its Beginnings 628
18.1.2 Where Are We Going with this Axiom? 628
18.1.3 Introduction 629
18.2 Strong Parapolar Spaces with the Pentagon Property 630
18.2.1 Examples 630
18.2.2 The Role of Grids 630
18.2.3 Why We Assume that Γ Is Not a Polar Space 631
18.2.4 What Extra Axioms Are Needed? 632
18.3 Classifying the Parapolar Spaces Satisfying (PL) and (PL*) 633
18.3.1 An Elementary Consequence of (PL*) 633
18.3.2 The Case that $\mathcal{D}$ is Empty and a Revisitation of Cohen's Theorem 634
18.3.3 The Case that Some Symplecton has Polar Rank at Least Three 637
18.3.4 The Case that $\mathcal{D} \neq \emptyset$ and All Members of $\mathcal{Q}$ Have Polar Rank Two 640
18.3.5 What We Have Proved 645
18.4 The Proof of the Main Theorem 647
18.4.1 The Hypothesis (H) 647

18.4.2 The Axiom (Q*) 650
18.4.3 Local Uniformity 650
18.4.4 The Metasymplectic Case 652
18.4.5 Beyond the Metasymplectic Case: A Related Geometry . . 652
18.4.6 Regarding Γ^* as a Locally Truncated Geometry 653
18.4.7 Enriching Γ^* to a Rank Four Geometry via Hanssen's Principle 654
18.4.8 The Conclusion that Γ Is a Homomorphic Image of a Polar Grassmannian 654
18.4.9 The Simple Connectedness of the Geometry Γ in the Case that $k = 2$ 655
18.5 Final Comments on the Main Theorem 656
18.6 Exercises for Chapter 18 657
18.6.1 Relaxing Hypothesis (PL*) 657
18.6.2 Relaxing (Q): What if the Class $\mathcal{Q}$ is Empty? 659

References 665

Index 671

Part I
Basics

Chapter 1
Basics About Graphs

Abstract A review of basic concepts of simple graphs introduces the important notion of a strongly gated subgraph. Graph morphisms, along with a theory of universal covers, are described. For any collection $\mathcal{C}$ of circuits in a connected graph, there exists a universal $\mathcal{C}$-cover defined by $\mathcal{C}$-homotopy classes of paths emanating from a fixed base vertex. Among criteria for simple $\mathcal{C}$-connectedness, one of these, due to J. Tits, will play a role in Chap. 10.

1.1 The Language of Graphs

Graphs were invented both to simplify and to generalize problems. They simplify problems because they break things down into the essential relations. They generalize problems because a theorem that works for a class of graphs enjoying the hypotheses of the theorem also works for every situation which can be represented by such a graph.

A *simple graph* is simply a set V (called the set of *vertices*) together with a distinguished collection E of pairs of vertices (that is, a collection of 2-subsets of V) which are called *edges* of the graph. So the graph is completely described by specifying the pair (V, E). Example: the vertex set is asserted to be $V = \{0, 1, 2, 3, 4\}$; the edges are declared to be the collection

$$E := \{\{0, 1\}, \{1, 2\}, \{2, 3\}, \{3, 4\}, \{4, 0\}\}$$

of 2-subsets. I think you can see that this is a pentagon. The custom is to draw pictures of these things so that the vertices are rendered as rather heavy dots, while the edges are rendered by continuous arcs connecting the vertices of an edge. (It may happen that the only way to draw this on a blackboard requires that some arcs seem to cross. In that case crossed arcs do not indicate a new vertex; rather it is a defect inherited from trying to draw these things on a two-dimensional blackboard. Just imagine that one arc is behind or in front of another.)

Simple graphs are useful for representing any symmetric relation on a set (such as the property of being "first cousins" on the set of people, or the relation of sharing a

E.E. Shult, *Points and Lines*, Universitext, DOI 10.1007/978-3-642-15627-4_1,

© Springer-Verlag Berlin Heidelberg 2011

common member among a set of committees – here the committees are the vertices, an edge is two committees sharing a common member).

Directed graphs or graphs with multiple edges or loops will not make an appearance in this book.

The reader should be forewarned that in this book the system of natural numbers, denoted by **N**, includes zero as well as the rest of the positive integers.

1.1.1 Connectedness and the Distance Metric

A *walk* is a sequence of vertices

$$w = (v_o, \ldots, v_n),$$

such that $\{v_i, v_{i+1}\}$ is an edge, $i = 0, \ldots, n-1$. The *length of the walk* w is the natural number n. Notice that the length is not the number of vertices x_i encountered during a tour of the walk: rather the length of w is the number of *hops* it would take a frog to go from v_0 to v_n, if each hop were along an edge $\{v_i, v_{i+1}\}$. The beginning vertex v_0 is called the *initial vertex* of the walk $w = (v_0, \ldots, v_n)$, and v_n (where the frog stops) is called the *terminal vertex of the walk*. Then w is said to be a *walk of length n connecting* v_0 (the initial vertex) *to* v_n (the terminal vertex of the walk).

We say that vertex a *is connected to* vertex b if and only if there exists a walk p from a to b. Clearly, if a is connected to b by a walk p, and vertex b is connected to c by a walk q, then, by first traversing walk p and then continuing on through walk q by frog hops, we obtain a walk $p \circ q$ from a to c called the *concatenation* of the walks p and q. So the relation of being connected by a walk is a *transitive relation*.

Note that a is connected to a via the walk (a) of length zero (I know this sounds like one of those tricks that mathematicians pull) so the relation is reflexive (that is, everybody is "connected to themselves").

Finally, it is obvious that the connectedness relation is symmetric. If there is a walk $w = (a = v_0, v_1, \ldots v_n = b)$ from a to b, then there is a backward walk $(b = v_n, v_{n-1}, \ldots v_1, v_0 = a)$, denoted by w^{-1}, from b to a which traverses the vertices of w in reverse order. The walk w^{-1} is called the *inverse of the walk* w.

Thus the relation on the set of vertices of being connected by a walk is reflexive, symmetric, and transitive, and so is an equivalence relation. The equivalence classes under this relation are called the *connected components*. Roughly speaking, if you thought of the graph (V, E) as a bunch of points in space whose edges were represented by connecting strings which somehow can pass through each other in space at non-vertex points so that no interlocking can occur, then the connected components are the pieces that continue to hang together no matter how you shake out this network in 3-space. A graph with just one connected component is said to be *connected*.

A few straggling terms remain to be defined. A *simple backtrack* is a walk of the form $w \circ w^{-1}$. A *subsegment of a walk* $w = (v_0, \ldots v_n)$ is a walk of the form $u = (v_i, v_{i+1}, \ldots, v_j)$ $(0 \leq i \leq j \leq n)$, so that $w = w_0 \circ u \circ w_1$ for some walks w_0, w_1 possibly of length zero. For example, any simple backtrack, of length

at least two, must contain a subsegment with is a backtrack of length two. A *path* is a walk $(v_0, v_1, v_2, \ldots, v_n)$, for which $v_i \neq v_{i+2}$, where i is relevant, so that it has no subsegments of positive length which are simple backtracks. Equivalently, $v_i \neq v_{i+2}, i = 0, \ldots, n-2$. Obviously two vertices are connected by a walk if and only if they are connected by a path. A *circular walk* is a walk $w = (v_0, \ldots, v_n)$ for which $v_0 = v_n$. (Note that it has length n.)

Suppose now that two vertices of a graph Γ are connected by some walk (by definition of finite length). Then among all walks connecting a to b there are those of shortest possible length. These are called *geodesics*. Their common minimal length is called *the distance between two vertices a and b* and is denoted by $d_\Gamma(a, b)$. It should be clear that every geodesic is free of backtracks and so is a path.

Now note that

$$d_\Gamma(a, b) \leq d_\Gamma(b, c) + d_\Gamma(a, c)$$

for any vertices $\{a, b, c\}$. That means this strange function

$$d_\Gamma : V \times V \to \mathbf{N} \tag{1.1}$$

obeys the "triangle inequality."

Just so we know what we are talking about, every vertex is distance zero only from itself. Two vertices a and b which are distinct, but form one of the edges $\{a, b\} \in E$ of our graph, are said to be *adjacent* and are distant one from one another. If two distinct vertices a and b are not adjacent, but there is a third vertex adjacent to both of them, then they are *at distance two from one another*. (It is a convention that if two vertices are in distinct connected components of the graph $\Gamma = (V, E)$ – that is, there is no path whatsoever connecting them – then the the distance between them is said to be infinite. This does not hurt the inequality (1.1) if we take infinity to be larger than any positive integer.)

The maximum of the numbers $\{d_\Gamma(x, y) | x, y \in V\}$ is called the *diameter* of the graph $\Gamma = (V, E)$. It is the length of a longest geodesic path that one could draw in this graph. The vertices and edges of a rectangle or a pentagon form a graph of diameter two; those of an octagon or nonagon have diameter four.

The subset of vertices of distance exactly i from vertex x is denoted by $\Delta_i(x)$, and is called the *radius at distance i from x*. Thus $\Delta_0(x) = \{x\}$ and $\Delta_1(x)$ is the set of all vertices adjacent to x.

We also will have occasion to consider the set $\Delta_i^*(x) = \{v \in V | d(x, v) \leq i\}$ of all vertices of distance *at most i* from vertex x. This set is called the *ball of radius i about x*. One other bit of notation is handy: for each vertex x we write $x^\perp := \{x\} \cup \Delta_1(x)$, the set $\Delta_1^*(x)$ of vertices at distance at most one from x.

1.1.2 Subgraphs

Now we need some notation. Suppose (V, E) is a simple graph, and that X is some specified subset of the set of vertices V. In that case, let E_X denote the collection

of all edges in E both of whose vertices belong to X. For example, if (V, E) is the pentagon on $V = \{0, 1, 2, 3, 4\}$ desribed above and $X = \{0, 1, 2\}$, then E_X consists of just two edges: $\{\{0, 1\}, \{1, 2\}\}$

A graph (X, E') is a *subgraph of* (V, E) if and only if $X \subseteq V$ and $E' \subseteq E_X$. It is called an *induced subgraph* if and only if $E' = E_X$, that is any adjacency among vertices in X is already represented by an edge in E'. Induced subgraphs are completely determined by the set X of vertices which they contain. For this reason, such graphs are often denoted by that set of vertices and, conversely, when a subset X of a vertex set is described as a subgraph, it is the induced graph of X that is intended.

Let $\Sigma_i = (X_i, E_i'), i = 1, 2$ be two subgraphs of the graph $\Gamma = (V, E)$. The *intersection of the subgraphs* is the subgraph $\Sigma_1 \cap \Sigma_2 := (V_1 \cap V_2, E_1' \cap E_2')$ – that is, the graph whose vertices are those vertices in both subgraphs and whose edges are those edges in both subgraphs. Notice that this makes sense since any edge in both graphs has its incident vertices in both vertex sets.

The important observation at this stage is as follows.

Lemma 1.1.1 *The class of induced subgraphs is closed under intersections.*

If $\Gamma' = (X, E')$ is a subgraph of (V, E), then the intersection of all induced subgraphs of Γ containing Γ' is again an induced subgraph of Γ called the *induced closure of* Γ'.

Lemma 1.1.2 *Suppose* $\Gamma' := (X, E')$ *is a subgraph of a graph* $\Gamma = (V, E)$. *Then the induced closure of* Γ' *– the unique smallest induced subgraph of* Γ *containing* Γ' *as a subgraph – is simply the induced graph* (X, E_X).

1.1.3 Special Types of Graphs and Subgraphs

Cliques and Cocliques

A graph of n vertices, any two of which form an edge, is called a *complete graph on n vertices* and is denoted by K_n (n can be infinite here). A subgraph which is complete is called a *clique* (of course, it would then be an induced subgraph).[1] A subset X of vertices for which $E_X = \emptyset$ is called a *coclique* or *an independent set*.[2] Obviously, the class of cliques and the class of cocliques are each closed under subgraph intersection. A clique (coclique) of n vertices is called an *n-clique* (*n-coclique*).

---

[1] In the graph theory of three decades ago, the word "clique" only referred to maximal complete subgraphs. Today, the term refers to any complete subgraph. We currently reserve the term "maximal clique" for what used to be called "clique" in the older literature. Students must be warned when the literature displays multiple definitions.

[2] We tend to eschew the latter term when our graph becomes immersed in contexts where "independent" acquires exterior meanings. I might add that the very useful term "coclique" was coined around 1974 by my dear friend Prof. Jaap Seidel.

Complements of Graphs

If $\Gamma = (V, E)$ is a simple graph, then the *complement of graph* Γ is the graph $\bar{\Gamma} = (V, E^*)$ where E^* is the collection of all 2-cocliques of Γ.

Multipartite Graphs

A graph is said to be *bipartite* if and only if its vertex set V can be partitioned into two sets $V = V_1 + V_2$, so that each edge involves one vertex in V_1 and one vertex in V_2. In other words a graph is bipartite if and only if its vertex set is the disjoint union of two independent sets.

Theorem 1.1.3 *A simple graph is bipartite if and only if it has no circuits of odd length.*

Proof Suppose first that $\Gamma(V, E)$ has no circuits of odd length. Without loss of generality, we may assume that Γ is connected. Then, choosing a vertex x, we may partition the vertices according to their distance from x:

$$V = \{x\} + \Delta_1(x) + \Delta_2(x) + \cdots .$$

Now if there were an edge $\{a, b\}$ inside some $\Delta_d(x)$, then there is a geodesic path $p = (x, x_1, \ldots, x_d = a)$ of length d connecting x to a, and a geodesic path $q = (b = y_0, y_1, \ldots, y_d = x)$ connecting b to x. Then

$$p \circ (a, b) \circ q := (x = x_0, \ldots x_d = a, b = y_0, y_1, \ldots, y_d = x)$$

is a circuit of length $2d+1$, an odd number. Since there are no circuits of odd length, there can be no edge in any of the subsets $\Delta_i(x)$ – that is, they are all independent sets. So if we set

$$V_1 = \{x\} + \Delta_2(x) + \Delta_4(x) + \cdots$$
$$V_2 = \Delta_1(x) + \Delta_3(x) + \Delta_5(x) + \cdots$$

then V_1 and V_2 are independent sets partitioning V, so Γ is bipartite.

On the other hand, suppose $\Gamma = (V, E)$ is given to be bipartite so that there is a partition $V = V_1 + V_2$ of the vertex set into two independent sets. Then it is clear that in following through a circuit, we are alternating vertices between V_1 and V_2 and so, in order to end where it began, a circuit must have even length. □

A *k-multipartite graph* is a graph (V, E) with a k-fold partition of the vertex set into independent components — that means one can divide the vertex set V onto k pairwise disjoint non-empty parts, $V = V_1 + V_2 + \cdots + V_k$ – so that all edges involve only two vertices from distinct parts. A *non-trivial k-vertex coloring* would be an assignment of a color to each vertex from a pool of k colors, such that:

1. The two vertices of any edge bear different colors
2. All k colors were actually used in painting the vertices (this is why we call it 'nontrivial')

Obviously the following is true.

Lemma 1.1.4 *The class of graphs which can be colored in at most k colors is exactly the class of graphs which are ℓ-multipartite for some $\ell \leq k$.*

[At this stage we are actually staring at the class of graphs which comprise the definition of "geometry." The vertices are objects of the geometry and the colors are the types of objects. The difference is that in a geometry this coloring is set for all time: we never need explore the class of all possible colorings. But we are getting a little ahead of ourselves.]

A *complete bipartite graph* is one in which every possible edge is drawn between the two parts. Thus $K_{3,3}$ is a graph on six vertices which have been divided into two independent sets V_1 and V_2 of size three, with edges being one of the nine pairs (x_1, x_2) with $x_i \in V_i$. A *complete multipartite graph* is a graph whose vertices can be partitioned into k non-empty cocliques and whose edge set E consists of all possible pairs of vertices belonging to distinct parts.

Graphs with Regularity Properties

An *isomorphism* of a graph $\Gamma = (V, E)$ to a graph $\Gamma' = (V', E')$ is a bijection $\sigma : V \to V'$ of the vertex sets for which (u, v) is an edge of E if and only if (u^σ, v^σ) is an edge of E'. Isomorphic graphs are essentially identical; we have merely changed the names of the vertices, if you like.[3] An *automorphism* of Γ is just an isomorphism of Γ with itself – that is, a permutation of the vertex set V which stabilizes the set E of edges. Clearly the identity mapping, the inverse mapping of an automorphism, and the composition of two automorphisms are all automorphisms themselves. Thus the automorphisms of Γ form a group under the operation of composition – a group which we denote by $\mathrm{Aut}(\Gamma)$.

Many of the graphs that we derive from geometries admit groups of automorphisms that cause them to be very symmetrical. A vertex v of a simple graph is said to have *degree* $d(v)$ if and only if $d(v)$ is the cardinality of the set of edges which are incident with v. If a simple graph is a *finite graph* (that is, a simple graph with finitely many vertices) then of course the numbers $d(v)$ are non-negative integers for each vertex v. The graph is said to be *regular* if all vertices have the same degree.

Clearly if $G = \mathrm{Aut}(\Gamma)$ is transitive on vertices (equivalently we say Γ is *vertex transitive*), then Γ is regular. Also, for each vertex x we have an induced subgraph $\Gamma(x) := \Delta_1(x)$ on the vertices which are adjacent to, but distinct from, the vertex x – the so-called "neighborhood graph of x".[4] If a graph is vertex-transitive then we even know that all of the subgraphs $\Gamma(x)$ are isomorphic.

In general we say that Γ *is locally* Y if, for each vertex x, the neighborhood graph $\Gamma(x)$ is isomorphic to a fixed graph Y.

Even stronger transitivity properties occur for the point-collinearity graphs of polar spaces, generalized polygons, dual polar spaces, Grassmann spaces, and other

[3] Isomorphisms are special instances of morphisms of graphs discussed fully in the next section.

[4] The custom is to denote this subgraph as $X(x)$ if "X" is the name of the parent graph.

geometries that we shall meet in this book. A graph Γ is said to be *distance transitive* if and only if, for each vertex x, and positive integer i, the subgroup G_x of automorphisms of Γ which fix x, transitively permutes the vertices of $\Delta_i(x)$, the vertices at distance i from x. These graphs are studied very completely in the book *Distance transitive graphs* by A. Brouwer et al. [12]. Most of the geometries just mentioned have point-collinearity graphs in this class.

Of course one may encounter regular properties like this without a transitive group. Suppose, for a finite graph Γ of diameter d, there are constants $(a_1 = 1, a_2, \ldots a_d; b_1, \ldots b_d; c_1, \ldots, c_{d-1}, c_d = 0)$ such that for every vertex x and integer $i \in [1, d]$, each vertex y of $\Delta_i(x)$,

$$\begin{aligned} |y^\perp \cap \Delta_{i-1}(x)| &= a_i \\ |y^\perp \cap \Delta_i(x)| &= b_i \\ |y^\perp \cap \Delta_{i+1}(x)| &= c_i. \end{aligned}$$

Then Γ is said to be *distance regular of diameter* d. Clearly any finite distance transitive graph is distance regular. Moreover, edge counts force

$$\begin{aligned} a_i + b_i + c_i &= |\Delta_1(x)| = d(x), \text{ and} \\ |\Delta_i(x)|a_i &= |\Delta_{i-1}(x)|c_{i-1}, \end{aligned}$$

for $1 \leq i \leq d$.

Distance regularity also forces the adjacency matrix of the graph to satisfy a polynomial equation of degree d. When $d = 2$, the situation is especially tight. In this case the graph is said to be *strongly regular* and the eigenvalues and their multiplicities can be determined by the parameters $k := d(x), \lambda := c_1$ and $\mu := a_2$ alone. This theory is recounted in [12]. Distance regular graphs are instances of coherent configurations, first introduced by D. Higman [74]. These have been studied from another point of view in the excellent book of P-H. Zieschang [150].

1.1.4 Metrical Properties of Subgraphs

Convexity

A subgraph (X, E') of $\Gamma = (V, E)$ is said to be *convex* if and only if, for each geodesic path $g = (v_0, \ldots v_d)$ of Γ connecting two vertices $a = v_0$ and $b = v_d$ of X, all intermediate vertices v_i also belong to X. (Note that this does not mean that edges (v_i, v_{i+1}) of the geodesic are edges of the subgraph. This would be true if the subgraph were an induced subgraph, but one can see that the definition of convexity of subgraph (V, E') makes no reference to the edge set E' at all. Convexity is really a property describing how *a set of vertices* – rather than a subgraph – sits in an ambient graph Γ.)

One easily observes the following:

Lemma 1.1.5 *The intersection* $\cap \Sigma_\sigma$ *over any family* $\{\Sigma_\sigma\}$ *of convex subgraphs is convex.*

Notice that in Lemma 1.1.5, no one is asserting that the members of the family are induced subgraphs. Indeed, given any subgraph, $\Gamma' = (X, E')$ of $\Gamma = (V, E)$, one might well consider the collection of graphs $K(\Gamma')$ which are convex and contain Γ' as a subgraph. The resulting intersection over $K(\Gamma')$ is called the *convex closure of* Γ'. It is the smallest convex subgraph of Γ containing Γ' as a subgraph.

Note that we may now combine this with a previous closure operator to form a new one. The *induced-convex-closure of a sub-graph* $\Gamma' = (X, E')$ *in* Γ is the intersection of all convex induced subgraphs of Γ that contain Γ' as a subgraph. (An important relation of this concept with connectedness is given in Exercise 1.12 (Sect. 1.5.1 at the end of this chapter).[5]

Corollary 1.1.6 *Let* $\Gamma' = (X, E')$ *be a subgraph of* $\Gamma = (V, E)$. *Then the induced-convex-closure of* Γ' *in* Γ *is simply the convex closure* (*in* Γ) *of the induced closure* (*in* Γ) *of* Γ.

Proof This is an elementary exercise (Exercise 1.3 at the end of this chapter). □

There is one further very important property of convex induced subgraphs of a graph.

Lemma 1.1.7 *Suppose* (X, E_X) *is a convex induced subgraph of a graph* $\Gamma = (V, E)$, *and suppose* v *is a vertex not in* X. *Then* $v^\perp \cap X$ *is a clique.*

Proof By way of contradiction assume $v^\perp \cap X$ is not a clique. Then there is a coclique $\{x, y\}$ of Γ contained in $v^\perp \cap X$, and so, since X is an induced subgraph, (x, v, y) is a geodesic path of Γ connecting x to y. By convexity of X, the intermediate vertex v must belong to X, contrary to the choice of v. Thus $v^\perp \cap X$ must be a clique. □

Isometric Embeddings

Given a subgraph $\Sigma = (X, E')$ of a graph $\Gamma = (V, E)$, there are two distinct metrics which can be applied. First, one can restrict the distance function $d_\Gamma : V \times V \to \mathbf{N}$ to $X \times X$. On the other hand, since Σ is a simple graph in its own right, by definition, it has a metric d_Σ which measures the distance between vertices of X in terms of the population of geodesics of the subgraph Σ alone. Always, we have

$$d_\Gamma(x_1, x_2) \leq d_\Sigma(x_1, x_2) \tag{1.2}$$

for all vertices x_1, x_2 of the subgraph Σ.

A subgraph $\Sigma = (X, E')$ is said to be *isometrically embedded* in its ambient graph $\Gamma = (V, E)$ if and only if the two available metrics for Σ coincide – that is, for every pair of vertices $\{x_1, x_2\}$ chosen from X, one has

[5] This result is a vital step in showing that buildings of finite rank are residually connected.

$$d_\Sigma(x_1, x_2) = d_\Gamma(x_1, x_2). \tag{1.3}$$

Example The vertices of the graph $\mathbf{H}_n$ are the 2^n binary n-tuples – that is, sequences of zeroes and ones of length n. Two such vertices are adjacent if and only if they differ in exactly one position of an entry. (This graph is called the *graph of the n-dimensional hypercube*.) It is easy to see that two vertices are at distance k in this graph if and only if their entries differ in exactly k positions. Clearly, these graphs are bipartite. The 8-vertex graph H_3 is called the *graph of the cube* since its vertices can be represented by the eight corners of a cube, two of them being adjacent if and only they represent corners connected by an edge of the cube.

Now there are two sorts of subgraphs (X, E') of H_3 which are hexagons – that is, paths of length six which end where they start. We can describe them, by describing these circular paths. First is

$$C_1 := (000, 001, 011, 111, 101, 100, 000).$$

The second is

$$C_2 := (000, 001, 011, 111, 110, 100, 000).$$

Now one may notice that the second and fifth vertices of C_1, namely 001 and 101 are at distance three in the hexagon, but are adjacent in the cube H_3. On the other hand, the hexagon C_2 is isometrically embedded in H_3.

We observe the following.

Lemma 1.1.8 *1. Any isometrically embedded subgraph is induced.*
2. Any convex induced subgraph is isometrically embedded.

But there are limitations. There are convex subgraphs which are not induced. For example, if $\Gamma = K_n$, the complete graph on n vertices, we can take as a subgraph the graph $\Sigma = (V, \emptyset)$, which contains all n vertices but possesses no edges. We see that a strict reading of the definition forces Σ to be convex, and for $n \geq 2$, it is not induced.

Similarly the dream of possible converses to both parts of the lemma is easily dispelled: (1) there are induced subgraphs which are not isometrically embedded – consider C_1 as a subgraph of H_3 above – and (2) there are isometrically embedded subgraphs which are not convex – for example, C_2 in H_3 above.

Gatedness and Strong Gatedness

In all likelihood, most of the graph-theoretic concepts discussed up to this point are well known to most readers. By way of contrast, the two concepts emblazoned in the title of this subsection are rarely considered – and for good reason: most subgraphs do not enjoy these properties.

However, the classical geometries which play center stage in this course – the *buildings* – seem to surround themselves with graphs which contain a family of

strongly gated subgraphs. Since geometries derived from buildings are precisely the geometries we wish to characterize, we need to understand these concepts.

Again, let $\Sigma = (X, E')$ be a subgraph of a connected graph $\Gamma = (V, E)$. Let y be an arbitrary vertex in V. We say that vertex y is *gated with respect to subgraph* Σ if and only if there exists a vertex g in X such that for every vertex $x \in X$, one has

$$d_\Gamma(y, x) = d_\Gamma(x, g) + d_\Gamma(g, y). \tag{1.4}$$

That means the Γ-distance from y to any vertex of the subgraph Σ can be "measured through the vertex g" – that is, one can find a geodesic from y to any vertex of Σ which contains g. Note that, since Γ is assumed to be connected, when we say that one of its subgraphs is gated with respect to y, all of the distances in Eq. (1.4) are natural numbers – there are no infinity symbols here.

Given y and subgraph Σ, this special vertex g is unique. If there were another vertex g' playing the role of g in Eq. (1.4), then

$$d_\Gamma(y, g') = d_\Gamma(y, g) + d_\Gamma(g, g'),$$
$$d_\Gamma(y, g) = d_\Gamma(y, g') + d_\Gamma(g, g'),$$

and since all distances in these equations are finite, it follows that $d_\Gamma(g, g') = 0$ so $g = g'$.

So g, being uniquely determined by the pair (y, Σ), is given a special name: it is called the *gate of* Σ *with respect to* y.[6]

Now there is a slightly stronger notion: we say that subgraph Σ (of the same connected graph Γ) is *strongly gated with respect to vertex* y if and only if

$$d_\Gamma(y, x) = d_\Gamma(x, g) + d_\Sigma(g, y). \tag{1.5}$$

Note that we are using the internal distance metric in the last term.

Now the following results at once from (1.2).

Lemma 1.1.9 *Strong gatedness of a subgraph with respect to y implies gatedness with respect to y.*

Finally, note that if y is itelf a vertex of Σ, then Σ is gated (but not necessarily strongly gated) with respect to y with y being its own gate g. In particular, if $\Sigma = (X, E')$ is a *spanning subgraph* – that means $X = V$ – then Σ is gated with respect to every vertex y. But in that case (unless it is isometrically embedded) it is not in general strongly gated with respect to every vertex.

Now we universally quantify the above on the vertex variable y. We say that a *subgraph* $\Sigma = (X, E')$ *is* (*strongly*) *gated* if and only if it is (strongly) gated with respect to every vertex y of its parent graph Γ.[7] We have the following.

---

[6] This concept was introduced by Andreas Dress and Rudolph Scharlau in [59].

[7] Again (without the adjective "strong") this definition, and part 1. of Lemma 1.1.10 to which it applies, are due to Dress and Scharlau [59].

Lemma 1.1.10 *1. Any gated subgraph is convex.*
2. Any strongly gated subgraph is convex and isometrically embedded.

Lemma 1.1.11 *Suppose A is a subgraph of a connected graph Γ which is strongly gated with respect to a vertex y, with gate h_A. Suppose B is a subgraph of A which is strongly gated in A with respect to the vertex h_A. Then B is a subgraph of Γ which is strongly gated with respect to vertex y.*

Proof Let h_B be the gate B with respect to vertex h_A. The hypotheses imply that for every vertex u of B,

$$d_A(h_A, u) = d_A(h_A, h_B) + d_B(h_B, u), \tag{1.6}$$
$$d_\Gamma(y, u) = d_\Gamma(y, h_A) + d_A(h_A, u). \tag{1.7}$$

Substitution of the right side of (1.6) for the last summand in (1.7) yields

$$\begin{aligned} d_\Gamma(y, u) &= d_\Gamma(y, h_A) + d_A(h_A, h_B) + d_B(h_B, u) \\ &= d_\Gamma(y, h_B) + d_B(h_B, u), \end{aligned}$$

by (1.7) with $u = h_B$. Thus B is strongly gated with respect to y with gate h_B. □

Theorem 1.1.12 *Let $\{\Sigma_\sigma\}_I$ be a family of subgraphs $\Sigma_\sigma = (V_\sigma, E_\sigma)$ of a connected graph $\Gamma = (V, E)$. Assume each subgraph of the family is strongly gated with respect to Γ, and suppose there is a vertex z lying in all vertex sets V_σ. Then the intersection $\cap_I \Sigma_\sigma$ is itself a strongly gated subgraph.*

Proof Let us write $V_\infty := \cap_I\{V_\sigma\}$ and $E_\infty := \cap_I\{E_\sigma\}$ so that $\Sigma_\infty := (V_\infty, E_\infty)$ is the intersection of all of the subgraphs Σ_σ. By hypothesis V_∞ is non-empty. Being strongly gated, each Σ_σ is a convex induced subgraph. Thus each is isometrically embedded in Γ, and their intersection Σ_∞ is also a convex induced isometrically-embedded subgraph of Γ. It follows that Σ_∞ is strongly gated with respect to any of its own vertices.

By way of contradiction, suppose Σ_∞ is not strongly gated. Then there exists a non-empty set Y of vertices with respect to which Σ_∞ is *not* strongly gated. Since Γ is connected, there is a vertex $y \in Y$ whose distance

$$d_\Gamma(y, \Sigma_\infty) := \min\{d_\Gamma(y, z) | z \in V_\infty\}$$

is minimal. Call this distance d. Then from the last sentence of the previous paragraph, we have $d \geq 1$.

Now choose a vertex $h \in V_\infty$ at which the minimal distance from y is achieved – i.e., $d(y, h) = d$. For each $\sigma \in I$, let h_σ be the gate of Σ_σ with repect to the

vertex y. We claim that h is a vertex of V_∞ nearest h_σ. If x were a vertex of V_∞ with $d(x, h_\sigma) < d(h, h_\sigma)$ then

$$\begin{aligned} d(y, x) &= d(y, h_\sigma) + d(h_\sigma, x) \\ &< d(y, h_\sigma) + d(h_\sigma, h) \\ &= d(y, h), \end{aligned}$$

against $d(y, \Sigma_\infty) = d(y, h)$.

Now suppose for some $\sigma \in I$ that $d(h_\sigma, \Sigma_\infty) < d$. Then by minimality of d, Σ_∞ is strongly gated in Σ_σ with respect to the vertex h_σ. Since h is a vertex of Σ_∞ nearest h_σ, it is in fact the gate of Σ_∞ with respect to the vertex h_σ. We now have all of the hypotheses of Lemma 1.1.11 with Σ_σ and Σ_∞ in the roles of subgraphs A and B, respectively. By the conclusion of that lemma, Σ_∞ is strongly gated with respect to y. That contradicts our choice of y.

Thus, for *every* $\sigma \in I$, we must have $d(h_\sigma, \Sigma_\infty) \geq d$. But then,

$$\begin{aligned} d &= d(y, h) = d(y, h_\sigma) + d(h_\sigma, h) \\ &\geq d(y, h_\sigma) + d, \end{aligned}$$

so

$$d(y, h_\sigma) = 0.$$

Thus $y \in \Sigma_\sigma$ for every $\sigma \in I$. But that means $y \in \Sigma_\infty$ against $d \geq 1$. This final contradiction shows that Y is empty and so the conclusion of the theorem holds. □

1.2 Morphisms of Graphs

As we have said, geometries can be uniquely described by graphs, often in several different ways. When we wish to identify a graph by certain "geometric" axioms, we usually feel better when these axioms are of a "local nature" – that is, they only refer to the structure of the graph that is in some neighborhood of a vertex (with respect to the distance metric). We like such theorems because they make it easier to recognize what geometry one is dealing with and so can become more useful in building a classification theory.

Unfortunately, one sometimes must pay a price for this. Suppose a very nice known geometry possesses a number of local properties P. Conceivably, there are points and lines and other objects at great distance from one another which could be identified without disturbing the purely local properties. In that case how can one ever hope to obtain a significant classification theorem with P as hypothesis? Well, their sometimes *is* a way. If we can show that life is so organized that any geometry with local hypothesis P is a homomorphic image of a *particular* "universal geometry" $U(P)$ and if we can simply characterize $U(P)$ (say, by parameters), then we have more or less accomplished the classification.

It is for this reason that we must develop a good general context for morphisms and develop a general theory about the existence of "universal preimages." Once again, very simple-minded graph theory provides us with both the general context and its easy accessibility to any reader.

1.2.1 Basic Definitions

Suppose $\Gamma = (V, E)$ and $\Gamma' = (V', E')$ are two simple graphs. A *graph homomorphism* (or *morphism of graphs*) is a mapping $f : V \to V'$ such that if $\{x, y\}$ is an edge in E, then either $f(x) = f(y)$ or else $\{f(x), f(y)\}$ is an edge of E'. Clearly, graph homomorphisms may be composed when the domains and images allow this, and these graphs and their morphisms form a category. A morphism is declared to be *surjective* or *injective* if and only if they are surjective or injective in the parent category – that is, as maps on vertex sets. The *image graph* is the graph $f(\Gamma) := (f(V), f(E))$, where

$$f(E) := \{\{f(a), f(b)\} | a, b \in V, f(a) \neq f(b)\}.$$

Given a surjective morphism $f : V \to V'$ of graphs, the collection

$$f^{-1}(x) := \{y \in V | f(y) = x\}$$

of preimages of vertex x, is a non-empty subset of V which we call the *fiber above* x.

We say that the morphism $f : V \to V'$ is *full* if and only if every pair of distinct image vertices in $f(V)$ which are adjacent in Γ' comprise the image of an edge in Γ. Put another way, this means $f(\Gamma)$ is an induced subgraph of Γ'. Obviously, if the codomain of one graph morphism is the domain of another, the two can be composed to obtain a graph morphism between the extremal members of the chain. The composition of full morphisms is full.

A vertex bijective morphism $\phi : \Gamma_1 \to \Gamma_2$ which is full is called an *isomorphism of simple graphs*.[8] Note that because of the fullness requirement and the fact that we are considering here only graphs without multiple edges, an isomorphism induces a bijection on edges. Thus the inverse of an isomorphism is an isomorphism.

An isomorphism of a graph with itself is called an *automorphism* of that graph. For example, the identity mapping of the vertices is always an automorphism. As observed on p. 8, the automorphisms of a graph Γ form a group under the composition operation; we denote this group by the symbol $\mathrm{Aut}(\Gamma)$.

In fact, a very good way to show that a conjectured group with prescribed properties actually exists is to find a graph of which it is the group of automorphisms. For example, the Higman Sims group, MacLaughlin's group, and the group *Suz* – three of the 26 sporadic finite simple groups – were discovered in this way.

---

[8] In the category of graphs with multiple edges, this would not be the desired definition of isomorphism.

Here is a standard construction of a full morphism. Suppose A is a group of automorphisms of a graph $\Gamma = (V, E)$. Then there is a natural partition of the vertex set into A-orbits – specifically, we have a surjective mapping $f : V \to V/A$ which sends each vertex v to the A-orbit, $v^A = \{v^a | a \in A\}$ to which it belongs. The symbol V/A simply denotes the collection of components of this partition – the full set of (non-empty) A-orbits. This way the mapping ϕ is (vertex) surjective.

Now we make a graph $\Gamma/A = (V/A, E(V/A))$ where $E(V/A)$ is the collection of pairs of distinct A-orbits, (x^A, y^A) such that at least one vertex in x^A is adjacent to a vertex in y^A. This definition is designed to make the vertex mapping ϕ a full vertex-surjective morphism of graphs $\phi : \Gamma \to \Gamma/A$. (For full details see Exercise 1.17 in Sect. 1.5.2 at the end of this chapter.)

1.2.2 Fiberings, Covering Morphisms, and Lifts of Walks

Suppose $f : V' \to V$ is the mapping on vertices induced by a graph homomorphism ϕ from $\Psi = (V', E')$ into $\Gamma = (V, E)$. We say that f (or even ϕ) is *locally bijective* if and only if, for any vertex x' in V', the set of edges of Ψ on x' is mapped one-to-one to the set of edges of Γ on the image $f(x')$. In the notation introduced earlier, this means that for any vertex $y \in V'$, the mapping $y^\perp \to f(y)^\perp$ induced by restriction of f is bijective as a mapping on vertices. (It is not usually a morphism of subgraphs.)

Finally, $f : V \to V'$ is called a *fibering* if and only if:

1. f is surjective (on vertices)
2. f is locally bijective

This means that for every edge $(f(x), f(y))$ in E', and any preimage vertex x in the fiber above $f(x)$, there is a *unique* edge (x, y) of E on x, such that $f(y) = y$. This already implies that f is full.

Let us consider a simple example. Suppose Γ is the 8-vertex graph of the cube – that is, the graph H_3 of the previous section. In this graph, each vertex v faces a unique antipodal vertex v' at distance three from it. We then form a mapping f of these vertices onto the four vertices of the complete graph $\Gamma' = (V', E') = K_4$ so that vertices possess the same image if and only if they are in the antipodal distance three relation in Γ. (This is possible since being antipodal is an equivalence relation.)

Now, it is clear that adjacent vertices possess distinct adjacent images, so, indeed, this mapping $f : H_3 \to K_4$ is a surjective morphism of graphs. Moreover, since both graphs have valence three and fibers are cocliques, the morphism is locally bijective. So it is a fibering. So we say that "the cube is a 2-fold cover of K_4."

Of course, any ordinary graph isomorphism is an example of a fibering – one characterized by the fact that all fibers contain a single point.

Let $\phi : \Psi = (V', E') \to \Gamma = (V, E)$ be a fibering with the induced surjection $f : V' \to V$ on the vertices. Consider a walk w in Γ having initial vertex $i(w) = x$ and let x' be any vertex in the fiber above x so that $f(x') = x$. There is then a *unique lift* $\lambda(w)$ of w – that is, a unique walk w' of Ψ starting at x', such that $f(w') = w$. The local bijectivity forces this, for if $w = (x_0, x_1, \ldots, x_n)$ (so $x = x_0$), then there

is a unique preimage (x', x_1') in Ψ of the edge (x, x_1). Since now $f(x_1') = x_1$, local bijectivity of ϕ at x_1' forces a unique preimage (x_1', x_2') of the edge (x_1, x_2) and one continues in this way, lifting each succeeding edge in the walk w at the terminal vertex of the preceding lifted edge.

Of course, if p is a path connecting distinct vertices a and b and a' is an element of the fiber above a, then the unique lift of p at vertex a' is a path p' from a' to a vertex b' in the fiber above b. If no vertex of path p is repeated (that is path p does not cross itself) then also no vertex of the lifted path p' is repeated, and p' and p are isomorphic as subgraphs. So the main difference between p' and p that could be recorded is that p crossed itself at certain spots which failed to show up as crossing points in the lifted path p'. But that immediately focuses attention upon circular walks w' which begin and end at a vertex a. We call these "pointed circuits" since the vertex declared to be the beginning and ending has been distinguished. To be precise, a *pointed circuit* is a walk $w = (x_0, x_1, \ldots, x_k)$ for which $x_0 = x_k$ – it is *pointed at* its common beginning/ending $x_0 = x_k$. We adopt this notation since it still allows us to concatenate them.[9]

A special sort of pointed circuit plays a special role: any closed walk of the form $w = u \circ u^{-1}$ (where u is any walk) is called a *backtrack*. (The reader may recall the definitions of "concatenation" and "inverse walk" from p. 4.) Note that if t is already defined as a backtrack, and v is any walk that ends at the initial vertex of v, then $w = v \circ t \circ v^{-1}$ is a also a backtrack. (Since $t = s \circ s^{-1}$, w has the form $w = u \circ u^{-1}$ where $u = v \circ s$.)

Let $c = (x_0, \ldots, x_n = x_0)$ be a pointed circuit. Then $(x_1, \ldots, x_n, x_1)$ is a *rotation* of c, and all iterates of rotations of c are called the *oriented association class* of c. The iterated rotations of c or $c^{-1} = (x_0, x_{n-1}, \ldots, x_1, x_0)$ together are called the *association class of* c and its members are the *accociates* of c – if we have to put a name to all these notions.

For example, in the covering $H_3 \to K_4$ given as an example just above, every single triangle of K_4 – without exception – lifts to a hexagon of H_3, indeed one which is isometrically embedded.

But there is another rather trivial example of a fibering. We could simply let $\Psi = \Gamma_1 \dot{+} \cdots \Gamma_k$ be a graph whose connected components are just subgraphs $\Gamma_i = (V_i, E_i)$, each isomorphic to $\Gamma = (V, E)$. Then the morphism $\phi : \Psi \to \Gamma$ simply

[9] We would lose this advantage of notation if we paid attention only to the circular structure itself with no distinguished beginning or end – though ultimately, it is the "breaking open" of such an unpointed circuit which describes what goes on in a covering morphism. The notion of a general circuit with no particular beginning/end is probably best expressed by the notion of an "association class" of pointed circuits, (defined in the next few paragraphs).

Indeed one might imagine dispensing with pointed circuits altogether from the following point of view. Let $\Gamma = (V, E)$ be a simple graph. Set $U = \mathbf{GF}(2)V$ be the free vector space over $\mathbf{GF}(2)$ with basis V – that is, the vector space of functions Hom $(V, \mathbf{GF}(2))$. Similarly, we form the free vector space $W := \mathbf{GF}(2)E$ spanned by the edges. Now the mapping which sends each edge $e = (x_1, x_2)$ to the vector $x_1 + x_2$ extends to a transformation $\gamma : U \to W$ whose kernel K is called the *cycle space*. Of course every closed walk w maps to a vector in K

It turns out that this world of adding circuits mod 2 regardless of their orientation – as interesting as it is – simply does not provide a machine that can give us a universal cover.

applies the isomorphism $\phi_i : \Gamma_i \to \Gamma$ to each connected component. Clearly ϕ is a fibering, but not a very interesting one. For this reason, when Γ is a connected graph, we say that $\phi : \Psi \to \Gamma$ is a *covering morphism* if and only if it is a fibering for which the domain-graph Ψ is connected. (By vertex surjectivity it will follow that the graph Γ comprising the range is also connected.)

$\mathcal{C}$-Covers

Now fix a collection $\mathcal{C}$ of pointed circuits of the connected graph $\Gamma = (V, E)$. Let $\phi : \Psi \to \Gamma$, be a cover – that is ϕ is a covering morphism with connected domain Ψ. In this context, as just observed, every walk w beginning at vertex $\phi(x)$ in Γ lifts uniquely to a walk w_x of Ψ with initial vertex x.

We say that ϕ is a *$\mathcal{C}$-cover* if and only if every pointed circuit in Γ which belongs to $\mathcal{C}$ lifts to a closed walk – that is a pointed circuit.

Of course the full collection of pointed circuits of Γ which lift to pointed circuits may be much larger than $\mathcal{C}$. For the moment, let us denote the collection of pointed circuits which lift to pointed circuits of Ψ by the symbol $\bar{\mathcal{C}}$. The following observations are direct consequences of the uniqueness of lifts.

Lemma 1.2.1 *Let $\phi : \Psi \to \Gamma$ be a $\mathcal{C}$-cover.*

1. *At any vertex in the fiber above its initial point, every simple backtrack (including those of length zero) lifts to a simple backtrack.*
2. *Suppose w is a walk in Γ with initial vertex $\bar{x}$ and terminal vertex $\bar{y}$. Let w_x be the lift of w at vertex x in the fiber $\phi^{-1}(\bar{x})$, with terminal vertex y. Then y is in the fiber above $\bar{y}$, and the walk w_x^{-1}, is the unique lift at vertex y of the walk w^{-1} in Γ from $\bar{y}$ to $\bar{x}$.*
   In particular, if c is a pointed circuit in $\bar{\mathcal{C}}$, then so is c^{-1}.
3. *If c_1 and c_2 are pointed circuits in $\bar{\mathcal{C}}$, both pointed at the same vertex, then $c_1 \circ c_2$ is also a pointed circuit. .*
4. *Suppose $c = (y_0, y_1, \ldots, y_m)$ lifts to a pointed circuit $c_x := (x_0, x_1, \ldots, x_m)$ at vertex $x = x_0$ in the fiber over $y_m = y_0$. Then its associate circuit $(x_1, \ldots, x_m, x_1)$ is the unique lift of the pointed circuit $(y_1, \ldots, y_m, y_1)$ at the vertex x_1 in the fiber above y_1. Thus, in general, $\bar{\mathcal{C}}$ is closed under taking associates of pointed circuits.*

Thus one sees that if ϕ is the $\mathcal{C}$-cover of Lemma 1.2.1 then the closure $\bar{\mathcal{C}}$ contains the entire association class of each of its pointed circuits. For this reason it will be convenient to refer to the association classes of pointed circuits, by the simpler term "*circuits*."

1.2.3 Universal $\mathcal{C}$-Covers

A $\mathcal{C}$-cover $\kappa : \bar{\Gamma} \to \Gamma$ is said to be *universal* if and only if, for every $\mathcal{C}$-cover $\phi : \Lambda \to \Gamma$, there exists a covering morphism $\rho : \bar{\Gamma} \to \Lambda$ such that the composition of covering morphisms $\phi \circ \rho$ is κ.

It is easy to see that any two universal $\mathcal{C}$-covers of a graph (should they exist) must be isomorphic.

Finally we say that a graph $\Gamma = (V, E)$ with a prescribed collection of pointed circuits $\mathcal{C}$ is $\mathcal{C}$-*simply connected* if and only if it is a connected graph with the property that any $\mathcal{C}$-covering morphism is an isomorphism. In that case, of course, a universal $\mathcal{C}$-covering exists – namely itself.

In fact the following is true.

Theorem 1.2.2 *Let $\Gamma = (V, E)$ be any connected simple graph, and let $\mathcal{C}$ be any collection of circuits of Γ. Then a universal $\mathcal{C}$-covering for Γ always exists.*

This theorem, though frequently not explicitly recognized, is a very general theory that in some form has been part of folklore for a number of years. The proof has been written explicitly first for 2-covers of chamber systems in [139], and for general systems of pointed circuits $\mathcal{C}$ of graphs in Ernest Shult, Coverings of graphs (1998), and in [4]. An elementary development of the homotopy theory behind this result is presented in the next few sections along with several applications. Theorem 1.2.2 itself is proved in Sect. 1.4 of this chapter.

Why Universal Covers Are Important

We said that a graph Γ was *simply $\mathcal{C}$-connected* if and only if every $\mathcal{C}$-cover is an isomorphism of graphs. This means that such a graph is its own universal $\mathcal{C}$-cover. This is important, for it means that such a graph is a unique isomorphism class of objects. In certain instances, one can show that a graph is simply $\mathcal{C}$-connected only from local data (such as the nature of $x^\perp$ and possibly $x^\perp \cap y^\perp$, when x and y are at distance two). That means the local data uniquely determines the whole graph.

There are three very important contexts where this theory is useful:

1. Local characterizations of graphs. (Certain amalgams of groups can be characterized in this way.) Here $\mathcal{C}$ is all triangles (3-circuits).
2. Diagram geometries. Here $\mathcal{C}$ are the circuits lying within a residue of rank two.
3. Chamber systems. Here the edges bear sets of labels which must be preserved by the morphisms, but otherwise the theory is basically the same. The pertinent collection $\mathcal{C}$ consists of all pointed circuits lying within any residue of rank two.

1.3 $\mathcal{C}$-Homotopy

Elementary Homotopies

Throughout $\mathcal{C}$ is a collection of pointed circuits of a graph Γ.[10] Let $\mathcal{C}^*$ be the set of association classes of pointed circuits of $\mathcal{C}$. We then let $\mathcal{C}^{-1}$ be the set of pointed circuits c^{-1} for all $c \in \mathcal{C}$.

[10] Recall, from the definition on p. 17, that a "pointed circuit" is simply a circular walk; it is not necessaily a *circular path*.

Throughout the remainder of this section we make these three assumptions about $\mathcal{C}$:

(Hypothesis (H))

1. $\mathcal{C}$ contains all (pointed) simple backtracks. In particular, it contains all trivial circuits of length zero.
2. $\mathcal{C}$ is association closed – that is, if $(x_0, x_1, \ldots, x_n)$ is a pointed circuit in $\mathcal{C}$, then so is $(x_1, \ldots, x_n, x_1, x_0)$.
3. (Orientation invariance) If $c \in \mathcal{C}$, then also $c^{-1} \in \mathcal{C}$ – that is, we may reverse orientations of pointed circuits in $\mathcal{C}$. We render this assertion as $\mathcal{C} = \mathcal{C}^{-1}$.

A pair of walks (p, q) are said to be an *elementary $\mathcal{C}$-homotopy* if and only if

$$p = ((x_0, x_1, \ldots, x_m) \circ (a_0, a_1, \ldots, a_r) \circ (y_0, y_1, \ldots, y_n)),$$
$$q = ((x_0, x_1, \ldots, x_m) \circ (b_0, b_1, \ldots, b_s) \circ (y_0, y_1, \ldots, y_n)),$$

where

$$(a_0, a_1, \ldots, a_r) \circ ((b_0, b_1, \ldots, b_s))^{-1}$$

is a (pointed) circuit of $\mathcal{C}$. Note that as $\mathcal{C}$ contains every trivial circuit of length zero, this relation includes the identity relationship – i.e., p is always elementary $\mathcal{C}$-homotopic to itself.

Moreover the relation is symmetric: if p is $\mathcal{C}$-homotopic to q, then q is $\mathcal{C}$-homotopic to p. This is because $\mathcal{C} = \mathcal{C}^{-1}$.

Note that an elementary $\mathcal{C}$-homotopy does not disturb the initial and terminal vertices of the walks.

$\mathcal{C}$-Homotopy and Circuit Closure

Two walks p and q are said to be *$\mathcal{C}$-homotopic* if and only if there exists a finite sequence of elementary $\mathcal{C}$-homotopies

$$(p, p_1), (p_1, p_2), \ldots, (p_{k-1}, p_k = q),$$

the natural transitive extension of the reflexive and symmetric relation of elementary $\mathcal{C}$-homotopy. Thus the relation of being $\mathcal{C}$-homotopic is an equivalence relation.

Since $\mathcal{C}$ is fixed, we may write $p \sim q$ precisely when p and q are $\mathcal{C}$-homotopic walks. There are several immediate observations about this relation.

Lemma 1.3.1 *The following assertions hold:*

1. *The relation $\sim$ is an equivalence relation on walks.*
2. *If $p_1 \sim p_2$ and $q_1 \sim q_2$ and $t(p_1) = i(q_1)$, then*

$$p_1 \circ q_1 \sim p_2 \circ q_2.$$

3. *If $p \sim q$, then $p^{-1} \sim q^{-1}$.*

Proof Part 1 has already been established.

Part 2 is straightforward from the definitions since, if $((a_1, a_2), \ldots (a_{n-1}, a_n))$ is a sequence of elementary $\mathcal{C}$-homotopies, then so is either

$$(a_1 \circ b, a_2 \circ b), \ldots (a_{n-1} \circ b, a_n \circ b)), \text{ or}$$
$$(c \circ a_1, c \circ a_2), \ldots (c \circ a_{n-1}, c \circ a_n)).$$

In Part 3 it suffices to prove that if (p, q) is an elementary $\mathcal{C}$-homotopy, then so is (p^{-1}, q^{-1}). The former assertion means

$$p = a \circ c_1 \circ b \text{ and } q = a \circ c_2 \circ b,$$

where $c_1 \circ c_2^{-1} \in \mathcal{C}$. It follows that

$$p^{-1} = b^{-1} \circ c_1^{-1} \circ a^{-1} \text{ and } q^{-1} = b^{-1} \circ c_2^{-1} \circ a^{-1}.$$

But since $\mathcal{C} = \mathcal{C}^{-1}$, and is closed under taking associates, $c_1^{-1} \circ c_2$ is a circuit in $\mathcal{C}$. Thus (p^{-1}, q^{-1}) is an elementary $\mathcal{C}$-homotopy. The proof is complete. □

We can avoid a rather tangled world if we can describe $\mathcal{C}$-homotopy in terms of the set $\bar{\mathcal{C}}$ of all pointed circuits c of the form

$$c = p \circ q^{-1},$$

where (p, q) ranges over the collection of all C-homotopic pairs of paths.

By definition $\mathcal{C} \subseteq \bar{\mathcal{C}}$. Moreover, any sequence of elementary $\bar{\mathcal{C}}$-homotopies can be refined to a longer sequence of elementary $\mathcal{C}$-homotopies – that is, $\bar{\bar{\mathcal{C}}} = \bar{\mathcal{C}}$. Thus we have the following.

Lemma 1.3.2 *The "bar" operation is a closure operator in the poset of all subsets of the set of pointed circuits.*

1.3.1 Further Properties of the Closure Operator on Circuits

Given $\mathcal{C}$, what sort of pointed circuits belong to $\bar{\mathcal{C}}$? Here it is:

1. Obviously, $\mathcal{C}$ together with all backtracks belong to $\bar{\mathcal{C}}$.
2. If $c \in \bar{\mathcal{C}}$ then also $c^{-1} \in \bar{\mathcal{C}}$ (whether or not $c \in \mathcal{C}$).
3. If $c \in \bar{\mathcal{C}}$ then the entire association class containing c belongs to $\bar{\mathcal{C}}$.
4. Moreover, $c \in \bar{\mathcal{C}}$ if and only if $w \circ c \circ w^{-1} \in \bar{C}$, for any walk w whose terminal vertex is the pointed vertex of c.
5. Consider two closed walks c_1 and c_2 beginning at vertex c. If two of the pointed circuits in the set $\{c_1, c_2, c_1 \circ c_2\}$ belong to $\bar{\mathcal{C}}$, then so does the third.
6. Finally, $\mathcal{C}$-homotopy is $\bar{\mathcal{C}}$-homotopy.

All of these assertions are proved in Sect. 1.4 of this chapter.

We conclude this section with a simple theorem which will allow us to view the elements of $\bar{\mathcal{C}}$ in a slightly different way. A circular walk c is said to be $\mathcal{C}$-*contractible* if and only if there is a finite chain of elementary $\mathcal{C}$-homotopies

$$c \to c_1 \cdots \to c_N = (s), \tag{1.8}$$

reducing c to the circular walk of length zero consisting of its single pointed vertex, s. (Note that each elementary $\mathcal{C}$-homotopy is *for the path along c beginning and ending at s* – that is, the vertex s cannot be moved by such a homotopy.) It is easy to see that any member of $\mathcal{C}$ – in particular any backtrack – is homotopic to its initial vertex (regarded as a circular walk of length zero). For each contractible pointed circuit c, let $E(c)$ denote the minimal length N of a chain of elementary homotopies from c to (s) as in (1.8).

Theorem 1.3.3 *Suppose $\mathcal{C}$ is a fixed collection of circuits as before. The class $\bar{\mathcal{C}}$ coincides with the class of contractible pointed circuits.*

Proof First we show that any member of $\bar{\mathcal{C}}$ is contractible. Suppose $c = p \circ q^{-1}$ where p and q are $\mathcal{C}$-homotopic paths from vertex s to vertex t. Then there is a chain of elementary homotopies

$$p \to p_1 \to \cdots \to p_m = q.$$

Then we have a similar chain of $\mathcal{C}$-homotopies,

$$p \circ q^{-1} \to p_1 \circ q^{-1} \to \cdots \to p_m \circ q^{-1} = q \circ q^{-1},$$

whose last term is a backtrack which is contractible. Thus c is contractible.

Next we must show that every contractible circular walk is a member of $\bar{\mathcal{C}}$ – that is, it has the form $g \circ h^{-1}$ where g and h are $\mathcal{C}$-homotopic walks.

By way of contradiction, suppose

$$c = (s = x_0, x_1, \ldots, x_n = s)$$

is a contractible circular walk which is not a member of $\bar{\mathcal{C}}$, and among these we suppose c is chosen so that $E(c)$ is minimal. Then there is a finite chain of elementary $\mathcal{C}$-homotopies

$$c \to c_1 \to \cdots \to c_N = (s), \text{ where } N = E(c), \tag{1.9}$$

reducing c to a single vertex, s. If $N = 1$, then $c \to (x)$ is by definition an (elementary) $\mathcal{C}$-homotopy, so $c \in \mathcal{C}$, and hence $c \in \bar{\mathcal{C}}$ against our choice of c. Thus $E(c) = N > 1$.

That means the first elementary homotopy of the chain (1.9) is non-trivial, and transforms c to a pointed circuit c_1 by replacing a segment $d = (x_i, \ldots x_j)$ of c by a

new segment $d_1 = (x_i = y_0, y_1, \ldots, y_m = x_j)$, for some pair of indices $\{i, j\}$, with $0 \leq i < j \leq n$. Moreover, since c_1 must be distinct from c, we have $E(c_1) < E(c)$. Thus, by the minimality of $E(c)$, we have $c_1 \in \bar{\mathcal{C}}$, and we are able to write

$$c_1 = (s = x_0, \ldots, x_i = y_0, y_1, \ldots, y_m = x_j, \ldots, x_n = s) \tag{1.10}$$
$$= p \circ q^{-1}, \tag{1.11}$$

where p and q are $\mathcal{C}$-homotopic walks of length at least one beginning at vertex s and having the same terminal vertex u. At this point two cases appear as the terminal vertex u of p (also the initial vertex of q^{-1}) is one of the x_is or not – that is, whether u is an internal vertex of the segment d_1 or not.

Case 1: We suppose u is *not* an internal vertex of the segment d_1. Then u is either one of the vertices in $\{x_1, \ldots, x_i\}$ or one of the vertices in $\{x_j, x_{j+1}, \ldots, x_{m-1}\}$. However, since c is a member of $\bar{\mathcal{C}}$ if and only if c^{-1} is, we may reverse the orientation if necessay to assume that u is in the former set. Then putting

$$p = (s = x_0, \ldots, u),$$
$$b = (u, \ldots, x_i = y_0),$$
$$e = (x_j, x_{j+1}, \ldots, x_n = s),$$

we have the configuration depicted in Case 1 of Fig. 1.1. One now has the following homotopies:

$$p \sim q = e^{-1} \circ d_1^{-1} \circ b^{-1}, \tag{1.12}$$
$$d \sim d_1, \text{ and so} \tag{1.13}$$
$$e^{-1} \circ d_1^{-1} \circ b^{-1} \sim e^{-1} \circ d^{-1} \circ b^{-1}. \tag{1.14}$$

But $c = p \circ b \circ d \circ e$ and so it has the form $g \circ h^{-1}$ where $g = p$ and $h = e^{-1} \circ d^{-1} \circ b^{-1}$. Since $g \sim h$ by the homotopy relations of (1.12) and (1.14), we see that $c \in \bar{\mathcal{C}}$, a contradiction.

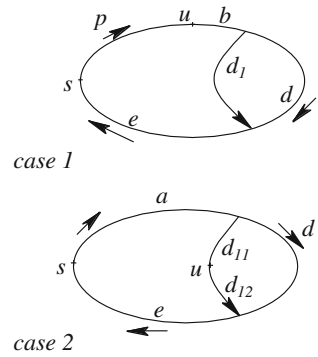

Fig. 1.1 The configurations of the two cases appearing in the proof of Theorem 1.3.3. Labels of walks accompany an *arrowhead* indicating their orientation. The terminal vertex u of p in the factorization $c_1 = p \circ q^{-1}$ is indicated by an *arrow* while the initial/terminal vertex s of c is indicated by a *dot*

Case 2: We suppose u to be one of the internal vertices y_k of d_1 $(0 < k < m)$. We then write $d_1 = d_{11} \circ d_{12}$ where

$$d_{11} := (x_j = y_0, \ldots, y_k = u) \text{ and} \tag{1.15}$$

$$d_{12} := (u = y_k, y_{k+1}, \ldots, y_m). \tag{1.16}$$

We set $a := (s = x_0, \ldots, x_i)$ and let $e = (x_j, x_{j+1}, \ldots, x_n = s)$ as in the previous case. One now has $c = a \circ d \circ e$ and the configuration of Case 2 in Fig. 1.1.

Now

$$p = a \circ d_{11} \sim e^{-1} \circ d_{12}^{-1} = q, \text{ and} \tag{1.17}$$

$$d \sim d_1 = d_{11} \circ d_{12}. \tag{1.18}$$

Thus, using the homotopies of (1.17), and shrinking one backtrack, one obtains

$$\begin{aligned} a \circ d &\sim a \circ d_1 \\ &= a \circ d_{11} \circ d_{12} \\ &\sim e^{-1} \circ d_{12}^{-1} \circ d_{12} \\ &\sim e^{-1}. \end{aligned}$$

Thus c has the form $g \circ h^{-1}$, where $g = a \circ d$ and $h = e^{-1}$, with g homotopic to h by homotopy equivalence of the first and last terms listed just above. So again this is a contradiction and the theorem is proved. □

1.3.2 Control of $\mathcal{C}$-Connectedness Through a Subgraph

Suppose $\mathcal{C}$ is a family of pointed circuits, and let $\bar{\mathcal{C}}$ be its closure, the collection of all pointed circuits $p \circ q^{-1}$ where p and q are $\bar{\mathcal{C}}$-homotopic walks. Suppose (X, E') is a subgraph of a graph $\Gamma = (V, E)$. Recall that this means $E' \subseteq E_X$ where E_X is the set of all edges of E whose incident vertices both lie in X.

Let $\mathcal{C}_X$ be the pointed circuits of $\bar{\mathcal{C}}$ which are subgraphs of (X, E'). (We have to say it this way since (X, E') may not be an induced subgraph. It is not enough to consider circuits whose *vertices* are in X.)

Lemma 1.3.4 *Suppose (X, E') is a subgraph of the connected graph Γ and $\mathcal{C}$, $\bar{\mathcal{C}}$, and $\mathcal{C}_X$ are defined as above. Suppose the subgraph (X, E') has this property:*

(H) *Any walk of Γ connecting two vertices of X is $\mathcal{C}$-homotopic to a path in (X, E') connecting the same vertices.*

Then Γ is simply $\mathcal{C}$-connected if and only if $\mathcal{C}_X \subseteq \bar{\mathcal{C}}$.

Proof Let $c = (x_0, x_1, \ldots, x_n)$, be any circuit pointed at $x_n = x_0$. Choose any vertex x in X, and let p be a path from x to x_0. Now by property 4 of $\bar{\mathcal{C}}$ listed at

the beginning of Sect. 1.3.1, p. 21, c is in $\bar{\mathcal{C}}$ if and only if $c' := p \circ c \circ p^{-1}$ is. But c', being a walk connecting two vertices of X, is by hypothesis C-homotopic to a circuit c_X of (X, E'), pointed at x. Thus $c \in \bar{\mathcal{C}}$ if and only if $c_X \in \bar{\mathcal{C}}$. The conclusion follows from this statement. □

Remark Note that one can't replace $\mathcal{C}_X$ by those circuits in $\mathcal{C}$ which lie in X. The latter set may be empty. Also it is not necessarily true that the circuit-closure of $\mathcal{C}_X$ recovers $\bar{\mathcal{C}}$. Although circuits of Γ are $\mathcal{C}$-homotopic to circuits within subgraph X, starting the other way, with the circuits of $\mathcal{C}_X$, one has no fore-ordained definition of $\mathcal{C}$ with which to reverse these $\mathcal{C}$-homotopies.

If $\mathcal{C}$ is a collection of pointed circuits of the graph (V, E), we say that a subgraph $\Sigma = (X, E')$ *controls $\mathcal{C}$-homotopy in* Γ if and only if condition (H) of Lemma 1.3.4 holds.

In the case that one is concerned with homotopy via backtracks and triangles, there is a useful sufficient condition for the control of homotopy by a subgraph. Here a *triangle* is simply a pointed circuit of length three, and the symbol $\mathcal{T}$ will denote the class of all triangles and backtracks. Consider the following.

Theorem 1.3.5 *Suppose (X, E') is a subgraph of $\Gamma = (V, E)$ with these two conditions:*

(i) For each vertex x not in X, the intersection of the induced subgraph $x^\perp$ with (X, E') is connected
(ii) Every edge lies in a maximal clique of Γ that has a non-empty intersection with X

Then the subgraph (X, E') controls $\mathcal{T}$-homotopy in Γ, where $\mathcal{T}$ is the collection of all triangles and backtracks of Γ.

Proof Suppose $p = (x_0, x_1, \ldots, x_n)$ is a path of minimal length connecting two vertices $x = x_0$ and $y = x_n$ of X subject to *not* being $\mathcal{T}$-homotopic to a path in (X, E'). Then by minimality we may assume *all* intermediate vertices $x_1, \ldots, x_{n-1}$ lie in $V - X$.

Suppose $n = 2$. Then by (i), there is a path q in the intersection of the (induced) subgraph $x_1^\perp$ and (X, E') connecting x with y. Then clearly, q is $\mathcal{T}$-homotopic to $p = (x_0, x_1, x_2)$, and we are done.

So suppose $n > 2$. By (ii) the edge $e = \{x_1, x_2\}$ lies in a maximal clique M in Γ containing a vertex v in $M \cap X$. Now by minimality of the length of p, the path $(v, x_2, \ldots, x_n)$ is $\mathcal{T}$-homotopic to a path r in (X, E'). By (i), there is a path q of (X, E') all of whose vertices lie in $x_1^\perp$, connecting $x = x_0$ and v. Now

$$p \sim (x_0, x_1, v) \circ (v, x_2, \ldots, x_n) \tag{1.19}$$

$$\sim q \circ r, \tag{1.20}$$

a path in (X, E'). Since this contradicts the choice of p, the proof is complete. □

Theorem 1.3.5 has several uses. The hypotheses (i) and (ii) appear especially natural in the context in which (X, E') is the collinearity graph of the points of a geometric hyperplane of a point-line geometry. We shall use it to show that rank two truncations of residually connected geometries are themselves residually connected (Sect. 2.6).

One also has the following

Corollary 1.3.6 *(Weetman [144]) Suppose Γ is a connected graph with this property: for every vertex x and positive integer i, we have:*

1. *For each vertex y in $\Delta_i(x)$, the induced subgraph $y^{\perp} \cap \Delta_{i-1}(x)$ is connected*
2. *For each edge $e = \{x, y\}$ in $\Delta_i(x)$, the set $e^{\perp} \cap \Delta_{i-1}(x)$ is non-empty*

Then Γ is $\mathcal{T}$-simply connected, where $\mathcal{T}$ is the class of all triangles (3-cycles, or equivalently, the class of oriented 3-cliques).

The proof is left as an exercise in Sect. 1.5.3.

1.3.3 Tits' Condition for Being Simply $\mathcal{C}$-Connected

In the previous section of this chapter we said that a graph H was simply $\mathcal{C}$-connected if every $\mathcal{C}$-cover was an isomorphism. In fact, one readily observes the following.

Lemma 1.3.7 *($\mathcal{C}$-connectedness) Let $\mathcal{C}$ be a family of circular walks of a graph H. Then the following conditions are equivalent:*

1. *H is simply $\mathcal{C}$-connected*
2. *If $\psi : G \to H$ is a $\mathcal{C}$-cover, then every circular walk of H lifts to a circular walk*
3. *For any two vertices c and d of H, all walks from c to d are $\mathcal{C}$-homotopic to one another*

In order to state the main theorem of this section in a convenient way, we require one or two definitions. Let $\mathcal{C}$ and H be as in the hypothesis of the preceding lemma (that is, *before* the listing of the conditions). If Y is any induced subgraph of H, we let $\mathcal{C}_Y$ denote the collection of all circular walks of $\bar{\mathcal{C}}$ which lie entirely in Y. Now fix an induced subgraph X. We say that $\mathcal{C}$ *controls the cycle structure of* X if and only if every connected component X_i of the subgraph X is simply $\mathcal{C}_{X_i}$-connected.[11]

Another small point concerns the notion of an "*image of a walk*." We would like the term to refer to a walk. Suppose $f : \Gamma_1 = (V_1, E_1) \to \Gamma_2 = (V_2, E_2)$ is an ordinary morphism of simple graphs. So far we have been describing a walk in Γ_1 by a sequence of vertices $w = (x_0, \ldots, x_n)$, where each vertex in the sequence is

[11] Recall from the definition on p. 24 that $\mathcal{C}_{X_i}$ is the collection of all pointed circuits of $\bar{\mathcal{C}}$ which are also circular walks of the subgraph of X induced on X_i. One notices that if graph H is simply $\mathcal{C}$-connected, then it is grammatically correct to say that "$\mathcal{C}$ controls the cycle strucure of H." But, as we can see, this notion applies to subgraphs which are not connected while the notion of simple $\mathcal{C}_X$-connectedness does not.

distinct from but adjacent to its sucessor (although $x_n = x_0$ is allowed). Now if we apply the morphism f to each vertex of the sequence, we get a new sequence of vertices of Γ_2: $(f(x_0), \ldots, f(x_n))$, which technically may not be a walk since now some of the succesive pairs of vertices are no longer distinct – that is, $f(x_j) = f(x_{j+1})$ for certain values of j. The point is that we can always cast out any vertex that is equal to the one preceding it in the sequence and iterate this process until we obtain a sequence with successive vertices distinct. The result will be a walk

$$f(w) := (y_0, y_1, \ldots, y_m),$$

where:

1. $y_0 = f(x_0)$ and $y_m = f(x_n)$
2. For $i = 0, \ldots, m-1$, $y_{i+1} = f(x_{i+})$ where $i+$ represent the smallest integer greater than i for which $f(x_{i+}) \neq f(x_i)$

Clearly $m \leq n$. Then $f(w)$ is a walk of Γ_2 which we call the *image of the walk* w.

Note that if w is a pointed circuit of (V_1, E_1), then its image $f(w)$ is also a (possibly shorter) pointed circuit of (V_2, E_2). If $\mathcal{C}$ is a collection of pointed circuits of Γ_1, the collection of images $\{f(w)|w \in \mathcal{C}\}$, noting that each $f(w)$ is a pointed circuit of Γ_2, will be denoted by $f(\mathcal{C})$.

Although it is a trivial consequence of the definition of an image of a walk, the following Lemma has important uses.

Lemma 1.3.8 *Suppose $f : \Gamma_1 \to \Gamma_2$ is a morphism of graphs. Let $\mathcal{C}$ be any collection of pointed circuits of G_1.*

1. *Suppose $u \to v$ is an elementary $\mathcal{C}$-homotopy of walks in Γ_1. Then either $f(u)$ and $f(v)$ are the same walk of Γ_2, or else $f(u) \to f(v)$ is an elementary $f(\mathcal{C})$-homotopy of walks in Γ_2.*
2. *If two walks of Γ_1 are $\mathcal{C}$-homotopic, then their images (as defined above) are $f(\mathcal{C})$-homotopic.*

The following criterion for simply $\mathcal{C}$-connectedness of a graph G is in terms of an epimorphic image H of G. The theorem does not appear in this isolated form in the literature. Rather, I have extracted it from the proof of the famous "Tits' Local Approach Theorem" [139] which is considered in Chap. 10 of this book.

Theorem 1.3.9 (J. Tits) *Let $\phi : G \to H$ be a morphism of connected graphs which is a surjection for their vertex sets. For each vertex x of the graph H the symbol G_x will denote the subgraph of G induced on the fiber $\phi^{-1}(x)$ above x. Let $\mathcal{C}$ be a collection of circular walks of G. We assume $\mathcal{C}$ satisfies these properties:*

(i) $\mathcal{C}$ controls the cycle structure of the induced subgraph G_y, for every vertex y in H.

(ii) If (a_1, b_1) and (a_2, b_2) are two edges of G whose images are the same edge in H, i.e.,

$$a = \phi(a_1) = \phi(a_2) \neq b = \phi(b_1) = \phi(b_2),$$

then there exists paths P and Q in G_a and G_b respectively, such that $P \circ (a_2, b_2) \circ Q \circ (b_1, a_1)$ is a circular walk of $\mathcal{C}$.

(iii) The collection of images $\phi(\mathcal{C})$ of the circuits of $\mathcal{C}$ controls the cycle structure of H. (Since H is connected this means H is simply $\phi(\mathcal{C})$-connected.)

Then G is simply $\mathcal{C}$-connected.

Remark Note that the morphism ϕ in the theorem is very different from the covering morphisms discussed earlier in this chapter. For example, if ϕ is uniformly full, condition (ii) implies that each fiber is connected – a real departure from the case of topological morphisms where the fibers contain no edges at all.

In order to prove the theorem we first prove the following lemma.

Lemma 1.3.10 *Assume the graph morphism $\phi : G \to H$ and the collection $\mathcal{C}$ of circular walks of G, satisfies the hypotheses (i) and (ii) of the statement of Theorem 1.3.9.*

Suppose w and w' are two walks of G with the same image walk $\phi(w) = \phi(w')$ in H and let a and b respectively be the initial and terminal vertices of this image-walk. Then there exists a path Y from the initial vertex of walk w to the initial vertex of walk w' and lying entirely within the induced subgraph of the fiber $\phi^{-1}(a)$, and similarly there exists a second walk Z from the terminal vertex of w to the terminal vertex of w' and lying in the subgraph of the fiber $\phi^{-1}(b)$, such that the circular walk

$$w \circ Z \circ (w')^{-1} \circ Y^{-1}$$

is $\mathcal{C}$-contractible.

Proof Let $\phi(w) = \phi(w') = (a = c_0, c_1, \ldots, c_n = b)$. Then we can write the circular walk w in the form

$$w = P_0 \circ (x_1, y_1) \circ P_1 \circ (x_2, y_2) \circ P_2 \circ \cdots \circ P_{n-1} \circ (x_n, y_n) \circ P_n$$

where each path P_i lies in the fiber $\phi^{-1}(c_i)$ and connects y_i to x_{i+1} for $i = 0, \ldots, n$ where y_0 denotes the initial vertex of w (as well as P_0) and x_{n+1} denotes the terminal vertex of w (as well as the walk P_n).

Similarly the walk w', which we presume to have initial vertex y'_0 and terminal vertex x'_{n+1}, can be written in the form

$$w" = P'_0 \circ (x'_1, y'_1) \circ P'_1 \circ (x'_2, y'_2) \circ P'_2 \circ \cdots \circ P'_{n-1} \circ (x'_n, y'_n) \circ P'_n$$

where the walks P'_i connect y'_i to x'_{i+1} and lie within the fiber $\phi^{-1}(c_i)$ for $i = 0, 1, \ldots, n$. Since hypothesis (ii) makes the fibers connected, there exists a path Y_0 connecting the initial vertices of w and w' (that is, it runs from y_0 to y'_0) lying entirely in the fiber $\phi^{-1}(c_0)$. Similarly, there is a path from the terminal vertex x_{n+1} of w to the terminal vertex x'_{n+1} of w' lying entirely inside the fiber $\phi^{-1}(c_n)$.

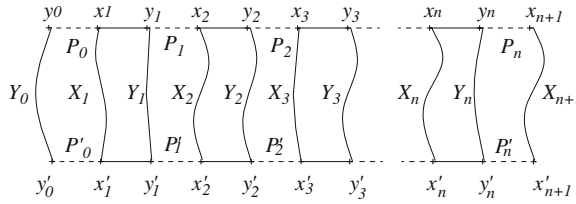

Fig. 1.2 The configuration derived from hypotheses (i) and (ii) of Tits' Homotopy Theorem from two walks of graph G having the same image walk in graph H. The two walks are the the *horizontal trails across the top and bottom of the figure*. The *dashed segments* belong to a common fiber

Now the images of the edges (x_i, y_i) and (x'_i, y'_i) are the edges (c_{i-1}, c_i), for $i = 1, \ldots n$. Using the full strength of hypothesis (ii), there exists a walk X_i from x_i to x'_i lying within fiber $\phi^{-1}(c_{i-1})$, and a walk Y_i in fiber $\phi^{-1}(c_i)$ from y_i to y'_i such that the pointed circuits

$$(x_i, y_i) \circ Y_i \circ (y'_i, x'_i) \circ X_i^{-1} \tag{1.21}$$

are $\mathcal{C}$-contractible for $i = 1, \ldots, n$.

By hypothesis (i), all fibers are $\mathcal{C}$-simply connected. Thus the circular walks

$$P_i \circ X_{i+1} \circ (P'_i)^{-1} \circ Y_i^{-1} \tag{1.22}$$

are also $\mathcal{C}$-contractible for $i = 0, \ldots n$ since each lies within a fiber.

We now have the the network of walks given in Fig. 1.2.

Since all the constituent cycles of (1.21) and (1.22) are $\mathcal{C}$-contractible, it follows that the path w is $\mathcal{C}$-homotopic to the path $Y_0 \circ w' \circ X_{n+1}^{-1}$. Thus the conclusion of the lemma holds. □

Corollary 1.3.11 *Suppose the morphism $\phi : G \to H$ and the class $\mathcal{C}$ satisfy the hypotheses (i) and (ii) of Theorem 1.3.9, as in the previous Lemma 1.3.10. Then $\phi(\bar{\mathcal{C}}) = \phi(\bar{\mathcal{C}})$, – in particular, every $\phi(\mathcal{C})$-contractible circular walk of H is itself an image of a $\mathcal{C}$-contractible circular walk of G.*

Proof First it is straightforward from the definition of the "image of a walk" that if $w_1 \to w_2$ is an elementary $\mathcal{C}$-homotopy of walks of G, then either $\phi(w_1) = \phi(w_2)$ or else $\phi(w_1) \to \phi(w_2)$ is an elementary $\phi(\mathcal{C})$-homotopy of walks in H. It follows from this that the image of a $\mathcal{C}$-contractible circular walk of G is $\phi(\mathcal{C})$-contractible, that is

$$\phi(\bar{\mathcal{C}}) \subseteq \phi\bar{(\mathcal{C})}.$$

To show the reverse containment we must show that every $\phi(\mathcal{C})$-contractible circular walk of H is itself an image of a $\mathcal{C}$-contractible circular walk of G. For this, it suffices to show that the class $\phi(\bar{\mathcal{C}})$ of images of $\mathcal{C}$-contractible circular walks of G is itself closed under elementary $\phi(\mathcal{C})$-homotopy.

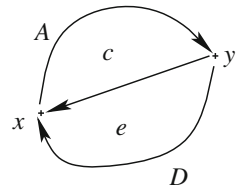

Fig. 1.3 An elementary $\phi(\mathcal{C})$-homotopy in H. The walk D is replacing walk B, represented by the straight arrow from y to x

We suppose then that $c = A \circ B \to A \circ D = f$ is an elementary $\phi(\mathcal{C})$-homotopy in H, where $c \in \phi(\bar{\mathcal{C}})$. This means the deformation $c \to f$ is by a circuit $e = D \circ B^{-1}$ which belongs to $\phi(\mathcal{C})$. We depict these relations of circuits and walks in Fig. 1.3 for reference. We let x be the initial vertex of A which at the same time is the common terminal vertex of B and D (this is because c is a circular walk). Similarly, y designates the terminal vertex of A and at the same time the initial vertices of both B and D. These vertices are indicated in the reference figure. Now by assumption, $c = A \circ B$ is the ϕ-image of a $\mathcal{C}$-contractible circular walk which can be written as $c' = A' \circ B'$ where $\phi(A') = A$ and $\phi(B') = B$. Thus the initial vertex x' of A' (also the terminal vertex of B') is in the fiber $\phi^{-1}(x)$, and the initial vertex y' of B' (also the terminal vertex of A') is in the fiber $\phi^{-1}(x)$.

Similarly, there is a circuit $e' = D' \circ (B'')^{-1}$ of $\mathcal{C}$ with $\phi(B'') = B$ and $\phi(D') = D$. Thus the terminal vertex x'' of B'' belongs to the fiber $\phi^{-1}(x)$ and the initial vertex y'' of B'' belongs to the fiber $\phi^{-1}(y)$. Now by Lemma 1.3.10, there exists a path Y from y' to y'' in fiber $\phi^{-1}(y)$, and path X from x' to x'' in fiber ϕ^{-1}, such that

$$(B')^{-1} \circ Y \circ B'' \circ X^{-1} \in \bar{\mathcal{C}} \tag{1.23}$$

(see Fig. 1.4). Putting Eq. (1.23) together with

$$c' = A' \circ B' \in \bar{\mathcal{C}} \text{ and} \tag{1.24}$$
$$e' = (B'')^{-1} \circ D' \tag{1.25}$$

we see that

$$A' \circ Y \circ D' \circ X^{-1}$$

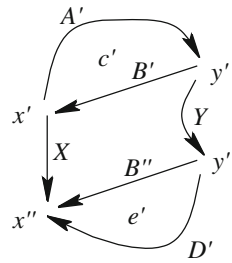

Fig. 1.4 The configuration obtained from ϕ-preimages of circular walks c and e of Fig. 1.3

is a circular $\mathcal{C}$-contractible walk whose image $\phi(f')$ is f, as desired. The proof is complete. □

Proof of Theorem 1.3.9 It suffices to show that any circular walk of G is $\mathcal{C}$-contractible. In addition to hypotheses (i) and (ii) used in the two previous results, we now have hypothesis (iii): the fact that H is $\phi(\mathcal{C})$-simply connected. Suppose c is a circular walk of G. Then its image $\phi(c)$ is a circular walk of H which by hypothesis (iii) is $\phi(\mathcal{C})$-contractible. By Corollary 1.3.11, $\phi(c) = \phi(c_2)$ where c_2 is a $\mathcal{C}$-contractible circuit of G. Thus $\phi(c) = \phi(c_2)$. Then there exists a vertex x in c and a vertex x_2 of c_2 with $\phi(x) = \phi(x_2)$, and we may regard c and c_2 as pointed circuits inducing the same orientation on their common image $\phi(c)$.

Then we can invoke Lemma 1.3.10 to produce two paths X and Y from x_2 to x such that

$$c \circ X^{-1} \circ c_2^{-1} \circ Y \in \bar{\mathcal{C}}. \tag{1.26}$$

But noting that $c_2 \in \bar{\mathcal{C}}$ and $X \circ Y^{-1} \in \bar{\mathcal{C}}$ by hypothesis (i), we obtain the following chain of $\bar{\mathcal{C}}$-homotopies:

$$\begin{aligned} c &\sim c \circ X^{-1} \circ Y \\ &\sim c \circ X^{-1} \circ c_2^{-1} \circ Y \in \bar{\mathcal{C}} \text{ by (1.26)}. \end{aligned}$$

Thus c is $\mathcal{C}$-contractible, as we were to prove. □

1.4 The Existence of Universal $\mathcal{C}$-Covers

Throughout this section, $\Gamma = (V, E)$ is a simple graph that is connected and $\mathcal{C}$ is a fixed collection of pointed circuits of Γ. For each walk w in Γ we let $[w]$ denote the equivalence class of all walks which are $\mathcal{C}$-homotopic to w. Our objective is to construct a universal $\mathcal{C}$-covering $\kappa : \hat{\Gamma} \to \Gamma$. For this purpose it is first useful to prove the results about $\bar{\mathcal{C}}$ described earlier in this chapter (see p. 21).

The following lemmas concern properties of this closure $\bar{\mathcal{C}}$.

Lemma 1.4.1 *The following hold:*

1. *First, $\bar{\mathcal{C}}$ contains all simple backtracks $p \circ p^{-1}$. This includes all circuits of length zero.*
2. *Any simple backtrack is $\mathcal{C}$-homotopic to the circuit of length zero at the initial point of the backtrack.*
3. *Any pointed circuit is in $\bar{\mathcal{C}}$ if and only if it is $\mathcal{C}$-homotopic to the circuit of length zero at its initial point ($\mathcal{C}$-contractibility again).*
4. *$\bar{\mathcal{C}}$ is $\mathcal{C}$-homotopy closed — that is, $a \sim b$ and $b \in \bar{\mathcal{C}}$ together imply $a \in \bar{\mathcal{C}}$. (Here, "$\sim$" indicates the relation of being $\mathcal{C}$-homotopic).*
5. *$\bar{\mathcal{C}}$ is invariant under orientation reversal. That is, if c is a pointed circuit in $\bar{\mathcal{C}}$, then so is c^{-1}.*

Proof 1. Such backtracks are members of $\mathcal{C}$ by hypothesis, and $\mathcal{C} \subseteq \bar{\mathcal{C}}$ by definition of the latter.
2. Suppose $s = p \circ p^{-1}$ where $p = (x_0, x_1, \ldots, x_n)$. Then as $s \in \mathcal{C}$, $s \sim (x_0)$, so $s = s \circ ((x_0))^{-1} \in \bar{\mathcal{C}}$, by definition of $\bar{\mathcal{C}}$.
3. This is an immediate consequence of Theorem 1.3.3.
4. Suppose $a \sim b \in \bar{\mathcal{C}}$, for pointed circuits a and b having initial vertex x. Then by 3, $b \sim (x)$. By transitivity of "$\sim$", $a \sim (x)$. Then $a \in \bar{\mathcal{C}}$, by the converse implication in 3.
5. To say that $c \in \bar{\mathcal{C}}$ means $c = u \circ v^{-1}$ where $u \sim v$. Since $v \sim u$, by definition, $c^{-1} = v \circ u^{-1} \in \bar{\mathcal{C}}$. □

Lemma 1.4.2 *Suppose c_1 and c_2 are two pointed circuits of the graph with the same initial vertex. If two of the three circuits $\{c_1, c_2, c_1 \circ c_2\}$ belong to $\bar{\mathcal{C}}$, then so does the third.*

Proof Suppose first that c_1 and c_2 belong to $\bar{\mathcal{C}}$ with common initial vertex x. Then by Lemma 1.4.1, part 3, we have $c_1 \sim (x) \sim c_2$. Then applying Lemma 1.3.1, part 2, $c_1 \circ c_2 \sim (x) \circ (x) = (x)$. Then Lemma 1.4.2, part 3, also implies $c_1 \circ c_2 \in \bar{\mathcal{C}}$.

Now assume c_1 and $c_1 \circ c_2$ are in $\bar{\mathcal{C}}$. Now, using for the first time Lemma 1.3.1, part 3, we have that c_1^{-1} belongs to $\bar{\mathcal{C}}$. Then by the previous paragraph, $c_1^{-1} \circ (c_1 \circ c_2) \in \bar{\mathcal{C}}$. But the left member is clearly $\mathcal{C}$-homotopic to c_2. So, by Lemma 1.4.1, part 4, $c_2 \in \bar{\mathcal{C}}$. □

Lemma 1.4.3 *The following hold:*

1. *(Transfer of base-point) If $c \in \bar{\mathcal{C}}$, and p is a walk from a vertex x to the initial vertex $y = i(c)$ of c, then $p \circ c \circ p^{-1}$ belongs to $\bar{\mathcal{C}}$.*
2. *(Removal of transfer of basepoint) Suppose $p \circ c \circ p^{-1}$ lies in $\bar{\mathcal{C}}$. Then so does c.*
3. *(Closure under taking associates) If $c = (x_0, x_1, \ldots, x_n)$ is a pointed circuit in $\bar{\mathcal{C}}$, then so is $(x_1, \ldots, x_n = x_0, x_1)$.*

Proof Throughout "$\sim$" indicates the relation of $\mathcal{C}$-homotopy of walks.

1. By Lemma 1.4.1, part 3, $c \sim (y)$. Then by Lemma 1.3.1, part 2, and Lemma 1.4.1, part 3,

$$p \circ c \circ p^{-1} \sim p \circ (y) \circ p^{-1} = p \circ p^{-1} \sim (y).$$

It follows from one more application of Lemma 1.4.1, part 3, that $p \circ c \circ p^{-1} \in \bar{\mathcal{C}}$.
2. Suppose $p \circ c \circ p^{-1}$ is in $\bar{\mathcal{C}}$. Then by transfer of base point, so is

$$a := p^{-1} \circ (p \circ c \circ p^{-1}) \circ p.$$

Then by Lemma 1.3.1, part 2, and Lemma 1.4.1, part 2,

$$a \sim (y) \circ c \circ (y) = c.$$

Then Lemma 1.4.1, part 4, implies $a \in \bar{\mathcal{C}}$.

3. Suppose $c = (x_0, x_1, \ldots, x_n) \in \bar{\mathcal{C}}$. By transfer of basepoint, the circuit

$$d := (x_1, x_0)(x_0, x_1, \ldots, x_n) \circ (x_n = x_0, x_1) \in \bar{\mathcal{C}}.$$

But

$$d = (x_1, x_0, x_1) \circ (x_1, \ldots, x_n, x_1) \sim (x_1, \ldots, x_n, x_1)$$

and so $(x_1, \ldots, x_n, x_1) \in \bar{\mathcal{C}}$, by Lemma 1.4.1, part 4. □

Corollary 1.4.4 *$\mathcal{C}$-homotopy is $\bar{\mathcal{C}}$-homotopy.*

Proof It suffices to show that an elementary $\bar{\mathcal{C}}$-homotopy is a $\mathcal{C}$-homotopy. So suppose (p, q) is an elementary $\bar{\mathcal{C}}$-homotopy. Then $p = a \circ c_1 \circ b$ and $q = a \circ c_2 \circ b$, where $c_1 \circ c_2^{-1} \in \bar{\mathcal{C}}$. But the latter only means that, in turn, $c := c_1 \circ c_2^{-1} = u \circ v^{-1}$ where u and v are $\mathcal{C}$-homotopic walks.

Replacing c by c^{-1} and transposing the roles of u and v if necessary, we may assume that the terminal vertex of c_1 does not come after the terminal vertex t of u. Then we can factor c_2 as $c_2 = v \circ d_2$, where v is a walk proceeding from $x := i(c) = i(c_1)$ to t as a prefix in the concatenation of c_2, and d_2 continues the walk from t to $y := t(c_1) = t(c_2)$. Then by hypothesis, there is a sequence of elementary $\mathcal{C}$-homotopies

$$v \sim v_2 \sim \cdots \sim v_k = u = c_1 \circ d_2^{-1}.$$

Then there are elementary $\mathcal{C}$-homotopies

$$c_2 = v \circ d_2 \sim \cdots \sim u \circ d_2 = c_1 \circ d_2^{-1} \circ d_2,$$

and the last term is in turn elementary $\mathcal{C}$-homotopic to c_1 since $\mathcal{C}$ contains all simple backtracks. Thus c_1 is $\mathcal{C}$-homotopic to c_2. It follows from Lemma 1.3.1, part 2, that $p = a \circ c_1 \circ b$ is $\mathcal{C}$-homotopic to $q = a \circ c_2 \circ b$, as desired. □

1.4.1 $\mathcal{C}$-Connectedness and Subgraphs

Let $X := (X, E')$ and $Y := (Y, E'')$ be subgraphs of the graph Γ. The intersection of the two subgraphs is the subgraph $X \cap Y := (X \cap Y, E' \cap E'')$. If $\mathcal{C}$ is a collection of pointed circuits of Γ, recall that its closure, $\bar{\mathcal{C}}$ has several equivalent definitions:

1. It is the collection of all circuits which lift to circuits under any $\mathcal{C}$-covering (which we now know exists).
2. It is the set of (pointed) circuits c of the form $p \circ q^{-1}$ where p is $\mathcal{C}$-homotopic to q.
3. It is the set of pointed circuits of Γ, which are $\mathcal{C}$-homotopic to a point (Theorem 1.3.3). (This version of the condition was called *$\mathcal{C}$-contractibility*.)

In our notation, $\mathcal{C}_x = \bar{\mathcal{C}} \cap (X, E')$ is all circuits of the subgraph (X, E') which are $\mathcal{C}$-contractibile. We say that (X, E') is *$\mathcal{C}$-simply connected*, if and only if each of its circuits belong to $\bar{\mathcal{C}}$ – that is, they are all $\mathcal{C}$-contractible. (Bear in mind that if one actually wishes to carry out a series of $\mathcal{C}$-homotopies converting a pointed circuit of (X, E') to its initial point, one may very well need to utilize circuits which are not even *in* the subgraph (X, E'). Thus a graph may be $\mathcal{C}$-connected without being $\bar{\mathcal{C}} \cap (X, E')$-connected.)

Lemma 1.4.5 *Suppose $\mathcal{C}$ is a collection of pointed circuits of the graph $\Gamma = (V, E)$. Let $X := (X, E')$ and $Y := (Y, E'')$ be $\mathcal{C}$-simply connected subgraphs. If the intersection $(X \cap Y, E' \cap E'')$ is connected in the ordinary graph-theoretic sense, then the subgraph*

$$(X \cup Y, E' \cup E'')$$

is $\mathcal{C}$-simply connected.

Proof Suppose $c = (x_0, x_1, \ldots, x_n)$ is a pointed circuit of $\delta = (X \cup Y, E' \cup E'')$ which is not $\bar{\mathcal{C}}$. Now the edges $e_i = \{x_i, x_{i+1}\}$ are either in E' or E''. Because of the symmetry of the hypotheses with respect to (X, E') and (Y, E''), we may assume edge e_0 belongs to E'. There is then a minimal index j such that e_i does not belong to E' – otherwise the entire circuit belongs to (X, E') and hence to $\bar{\mathcal{C}}$. Since it *is* a circuit, there is a smallest index k larger than j at which the circuit first returns to (X, E') – that is, $e_k \in E''$, but $e_{k+1} \in E'$. Thus the vertex x_{k+1} is the circuit's first "toehold" upon regaining the coastline of X. Since $(X \cap Y, E' \cap E'')$ is connected, there exists a path p in $(X \cap Y, E' \cap E'')$ from x_{k+1} to x_i. Writing $p_1 = (x_0, \ldots x_{k+1})$ and p_2 for $(x_{k+1}, \ldots x_n = x_0)$, we see that

$$c \sim p_1 \circ p \circ p^{-1} \circ p_2,$$

the concatenation of a circuit in (Y, E'') and a circuit which leaves subgraph (X, E') a fewer number of times t than did c. By induction on t, each factor of the concatenation belongs to $\bar{\mathcal{C}}$, and the proof is complete. □

1.4.2 The Construction

Fix any vertex, say b, in V. Let $\hat{V}$ be the collection of all $\mathcal{C}$-homotopy classes of walks in Γ which begin at vertex b. Recall that if w is such a walk having b as an initial vertex, then all walks in $[w]$ also begin at b and have a common terminal vertex, $\kappa([w]) := t(w)$.

We make $\hat{V}$ into a graph $\hat{\Gamma} = (\hat{V}, \hat{E})$, by declaring $([w_1], [w_2])$ to be an edge of $\hat{E}$ if and only if:

1. $(t_1, t_2) := (\kappa([w_1]), \kappa([w_2]))$ is an edge of Γ
2. The path $w_1 \circ (t_1, t_2)$ is $\mathcal{C}$-homotopic to w_2

One notes that adjacency is well-defined: that is, if $w_i' \in [w_i]$, $i = 1, 2$, then the t_is are unchanged, and the $\mathcal{C}$-homotopy of $w_1 \circ (t_1, t_2)$ and w_2 is contingent upon the $\mathcal{C}$-homotopy of $w_1' \circ (t_1, t_2)$ and w_2'.

Also, the former homotopy holds if and only if $w_2 \circ (t_2, t_1)$ is $\mathcal{C}$-homotopic to w_1. Writing "$\sim$" to indicate $\mathcal{C}$-homotopy, one alway has

$$w_1 \sim w_1 \circ (t_1, t_2) \circ (t_2, t_1),$$

and the latter term is homotopic to $w_2 \circ (t_2, t_1)$ if and only if $w_1 \circ (t_1, t_2) \sim w_2$. Thus adjacency is a symmetric non-reflexive relation on $\hat{V}$, and so, letting $\hat{E}$ denote all adjacent pairs in $\hat{V}$, $\hat{\Gamma}$ becomes a simple graph.

The graph $\hat{\Gamma}$ is connected. This is because any vertex $[w]$ is connected by an obvious walk in $\hat{V}$ to the vertex $[(b)]$, the class of all circuits pointed at b, which are homotopic to the circuit of length zero (at b).

Now $\kappa : \hat{V} \to V$, which reads off the common terminal vertex of each homotopy class of walks starting at b, takes edges to edges. Now let w be any walk from b to z in Γ. By the definition of adjaceny in $\hat{\Gamma}$, all vertices in $\hat{V}$ adjacent to $[w]$ have the form $[w \circ (z, y)]$, where y ranges over all neighbors of z in Γ. Thus κ induces a bijection of edges on $[w]$ to edges on $z = \kappa([w])$. Thus κ is a fibering.

Now suppose $w = (x_0, x_1, \ldots, x_n)$ is any walk in the graph Γ. Let $\hat{x_0} = [u]$ be any vertex of the fiber $\kappa^{-1}(x_0)$. What is the lift $\hat{w}$ of w at $\hat{x_0}$? It is simply

$$\hat{w} := ([u_0], [u_1], \ldots, [u_n]),$$

where $u_i := u \circ (x_0, x_1, \ldots, x_i)$, $i = 0, 1, \ldots, n$.

Next suppose $c = (x_0, x_1, \ldots, x_n)$ is a pointed circuit of Γ belonging to $\mathcal{C}$. Then the lift of c at $[u] \in \kappa^{-1}(x_0)$ is a walk in $\hat{\Gamma}$ beginning at $[u]$ and ending at $[u \circ c]$. But since c is in $\mathcal{C}$, it is homotopic to the pointed circuit (x_0), of length zero (Lemma 1.4.1, part 3). Thus $[u] = [u \circ c]$, so $\hat{w}$ is a circuit of $\hat{\Gamma}$ pointed at $[u]$. It now follows that κ is a $\mathcal{C}$-cover of Γ.

It remains to show that $\kappa : \hat{\Gamma} \to \Gamma$ is a universal cover. For this purpose we consider an arbitrary $\mathcal{C}$-covering, $\phi : \Lambda \to \Gamma$. Let V_L denote the set of vertices of Λ. Now for any walk in Γ, there are two sorts of lifts: those in $\hat{\Gamma}$ and those in Λ, and it will be neccesary to distinguish them notationally.

Fix for all time a vertex x in $\phi^{-1}(b)$. For any walk w in Γ beginning at b, let w^* be the unique lift of that walk in Λ beginning at x. Define the mapping

$$\psi_x : \hat{V} \to V_L$$

by the formula

$$\psi_x([w]) = t(w^*)$$

for all walks w of Γ beginning at vertex b. To see that this map is well-defined, merely note that $w_1 \sim w_2$ implies $w_1 \circ {w_2}^{-1}$ in $\bar{\mathcal{C}}$, so the lift ${w_1}^* \circ ({w_2}^*)^{-1}$ at x is a

pointed circuit of Λ at x. Thus by Lemma 1.4.3, part 3, its associate, $(w_2{}^*)^{-1} \circ w_1^*$ is a circuit pointed at $t(w_1^*)$. Thus, taking terminal vertices, we have $t(w_1^*) = t(w_2^*)$, so ψ_x is well-defined.

Some immediate observations:

1. First the ψ_x-image of a typical edge $([u], [u \circ e])$ (where e is a directed edge – that is, a walk of length one – beginning at the terminal vertex of u) must be $(t(u^*), t(u \circ e)^*)$, an edge of Λ.
2. Second, ψ_x is easily seen to be a fibering. One just composes the restiction of κ to a neighborhood $[u]^\perp$ of $\hat{\Gamma}$ with the inverse of ϕ at the image of that neighborhood.

Thus ψ_x is a fibering. Now it is easy to see that

$$\begin{aligned} \phi(\psi_x([u])) &= \phi(t(u^*)) \\ &= t(u) \\ &= \kappa([u]). \end{aligned} \tag{1.27}$$

Thus

$$\phi \circ \psi_x = \kappa$$

and so κ is a universal covering.

Remark It is important to note that there is nothing unique about ψ_x. It takes the distinguished vertex $[(b)]$ of $\hat{\Gamma}$ to x. Replacing x by some other vertex of the fiber in Λ above b would give a new "ψ".

As one might guess, such ambiguities surrounding the construction of a universal object must give rise to automorphisms of that universal object. Although this sounds like a category-theoretic incantation, it does motivate the next section.

1.4.3 Deck Transformations

Suppose now that in the construction of the previous section, $\phi : \Lambda \to \Gamma$ was also a universal $\mathcal{C}$-covering. Then by the construction, there is a fibering

$$\rho : \Lambda \to \hat{\Gamma},$$

taking distinguished vertex x in Λ to any desired vertex $[c]$ of the fiber in $\hat{\Gamma}$ above b and satisfying

$$\kappa \circ \rho = \phi.$$

Clearly c is a pointed circuit of Γ beginning and ending at b. It follows that

$$\mu := \rho \circ \psi_x : \hat{\Gamma} \to \hat{\Gamma}$$

is an automorphism of $\hat{\Gamma}$ taking each vertex $[u]$ of $\hat{\Gamma}$ to $[c \circ u]$. Clearly, μ preserves all fibers in Λ.

Conversely, for any $\mathcal{C}$-cover $\sigma : \Sigma \to \Gamma$, an automorphism of Σ which stabilizes all σ-fibers is called a *deck transformation*. Non-trivial deck transformations may or may not exist for a given $\mathcal{C}$-covering. But at least the following is true.

Lemma 1.4.6 *Let Γ be a connected graph and let $\mathcal{C}$ be a collection of pointed circuits which contains all backtracks. Let $\sigma : \Sigma \to \Gamma$ be a $\mathcal{C}$-cover, for some collection $\mathcal{C}$ of pointed circuits of Γ, and let D be the group of deck transformations for this fibering. Then D acts semiregularly on each fiber – that is, any deck transformation that fixes a vertex is the identity automorphism.*

Proof Suppose deck transformation d fixes a vertex x. Since distinct neighbors of x in Σ belong to distinct fibers and d stabilizes each fiber, d fixes all neighbors of x. The result now follows from Σ being connected (recall that Σ is connected by the definition of $\mathcal{C}$-cover). □

Coupling this with our observations about universal covers we have the following.

Corollary 1.4.7 *The group D of deck transformations of the universal $\mathcal{C}$-cover $\hat{\Gamma}$ of graph Γ acts regularly on the elements of each fiber.*

1.5 Exercises for Chapter 1

1.5.1 Exercises for Section 1.1

1.1 This exercise concerns bipartite graphs.

1. Show that a connected bipartite graph is bipartite in a unique way. Specifically, if a connected finite graph is bipartite with respect to two partitions of its vertex set, $V = A + B$ and $V = A' + B'$, then (A', B') is either (A, B) or (B, A).
2. Show that if $e = \{x, y\}$ is an edge in a connected bipartite graph (V, E), then there is a partition of the vertices into two sets:

$$D(x) := \{v \in V | d(v, x) \leq d(v, y)\},$$
$$D(y) := \{v \in V | d(v, y) \leq d(v, x)\},$$

that is, into sets of vertices which are closer (with respect to the distance metric) to one vertex of e than the other.

1.2 Show that if a non-empty coclique C is a convex subset of a connected graph Γ, then $|C| = 1$.

1.3 Write out a complete proof of Corollary 1.1.6.

1.4 Let X be a set of n objects and let k be an integer less than n. The *Johnson graph* $J_{n,k}$ is the graph whose vertices are the $\binom{n}{k}$ subsets of X of cardinality k — that is, k-subsets. The edges are pairs of k-subsets $\{A, B\}$ which intersect at a $(k-1)$-subset of X – that is, $|A \cap B| = k-1$. Show that two vertices A and B are at distance ℓ in this graph if and only if $|A \cap B| = k - \ell$. Conclude from this that $J_{n,k}$ has diameter $\min(|n-k|, k)$.

1.5 Graph-theorists often write T_n for $J(n, 2)$ – referring to it as the "triangular graph." Its complementary graph is then $\bar{T}_n$. Show that the graph $\bar{T}_5$ is strongly regular with parameters $(k, \lambda, \mu) = (3, 0, 1)$. (See p. 9 for the definitions of these parameters.) Show that in general, $\bar{T}_n$ is strongly regular for $n > 1$, and compute its parameters (k, λ, μ).

1.6 Let $\Gamma_i = (V_i, E_i)$, be two simple graphs with disjoint vertex sets V_i, $i = 1, 2$. A *product* $\Gamma_1 \times \Gamma_2$ is a graph with vertex set the Cartesian product $(V_1 \times V_2)$ and edge set $(E_1 \times V_2) \cup (V_1 \times E_2)$ where

$$E_1 \times V_2 := \{\{(x, u), (y, u)\} | \{x, y\} \in E_1, u \in V_2\}$$

and

$$V_1 \times E_2 := \{\{(x, u), (x, v)\} | \{u, v\} \in E_2, x \in V_1\}.$$

(Note that both of these are collections of 2-subsets of $V_1 \times V_2$, and so can serve as a sets of edges for a graph on $V_1 \times V_2$.) The *n-by-m grid* is any graph isomorphic to $K_n \times K_m$. Here n and m are any non-zero cardinal numbers and K_k is the complete graph on k vertices. (Sometimes this is called the "lattice graph" and is denoted by $L_{n \times m}$, or L_n when $n = m$.) Show that the 3-by-3-grid, $L_3 = K_3 \times K_3$ is isomorphic to its complement $\bar{L}_3$.

1.7 Show that the Johnson graph $J_{n,k}$ is locally a k-by-$(n-k)$ grid.

1.8 Suppose the graphs Γ_i are locally Y_i, $i = 1, 2$. Explicitly prove that $\Gamma_1 \times \Gamma_2$ is locally $Y_1 \dot{+} Y_2$, which denotes the vertex-disjoint union of the graphs Y_1 and Y_2. Observe that if neither Γ_i is the empty graph, then this product is locally not connected.

1.9 It has always been a fascinating problem to show that if a graph Γ is locally Y, for some fixed Y, or fixed alternatives for Y, then the graph Γ belongs to a known list of graphs. Such a result infers a complete global type from purely local data, and so is always important. Its mathematical fascination stems from the fact that for some choices of Y nothing happens, while for other isomorphism types, the global types are completely determined. Just why this happens is a mystery.

1. For any positive integer n, the symbol K_n denotes the complete graph on n vertices ("complete" means every pair of distinct vertices forms an edge). Show that any graph that is locally K_n is isomorphic to K_{n+1}.
2. Show that any graph which is locally complete bipartite is itself complete tripartite.

Remark There are many nice families of graphs $\mathcal{G}$ which locally do determine their ambient graphs. One very beautiful example is the class of locally Petersen graphs. The *Petersen Graph* is the graph $\bar{T}_5$ whose vertices are the ten 2-subsets that could be chosen from a 5-set $\Omega = [5] := \{1, 2, 3, 4, 5\}$, two of these 2-subsets are adjacent if and only if they share *no* common member. A famous theorem of Jon Hall [64] shows that any graph that is locally Petersen is one of the following three finite graphs: (1) the graph of 2-subsets of a 7-set, two being adjacent if and only if they are disjoint, (2) the commuting graph of a conjugacy class of involutions in the split extension $3 \cdot \mathrm{Sym}(7)$, and (3) a similar thing for the group $P\Gamma L(2, 5^2)$, the semidirect product of the simple group $PSL(2, 5^2)$ of 2-by-2 matrices of determinant one over the field of 5^2 elements, modulo the central group of scalar matrices, extended by the involutory field automorphism. That top factor encompasses the involution class of vertices of this graph.[12]

1.10 Given any simple graph $\Gamma = (V, E)$ there is a symmetric $(0, 1)$-matrix A called the *adjacency matrix* A whose rows and columns are indexed by the vertex-set V, and whose (i, j)-entry $a_{i,j}$ is zero if $i = j$, or if $\{i, j\}$ is not an edge. Show that for any positive integer k, the (i, j)-th entry of A^k is the number of distinct walks of length k from vertex i to vertex j.

1.11 Prove Lemma 1.1.10.

1.12 Suppose $\mathcal{F} = \{\Sigma_\sigma\}$ is a family of convex induced subgraphs of a graph Γ. Show that if each graph Σ_σ is connected, then so is the intersection $\cap \Sigma_\sigma$.

1.13 Show that a connected graph, all of whose edges are gated, is bipartite.

1.14 This has two parts:

1. Let H_n be the graph of the n-hypercube. Show that every 4-circuit is strongly gated.
2. Suppose Γ is the graph of 12 vertices and 30 edges of the dodecahedron. Show that the induced 5-circuit of any face is never a strongly gated subgraph.

1.15 Prove that for every non-trivial subgraph $\Gamma' = (X, E')$ of a graph $\Gamma = (V, E)$, there exists a unique subgraph $(\bar{X}, \bar{E}')$ which is the smallest strongly gated subgraph of Γ which contains Γ' as a subgraph. We can call this the *strongly*

[12] This example was inserted to show how tricky it is. The average graph theorist might not have thought of these group examples, but on the other hand, many irregular examples that graph theorists *do* think of do not easily come to mind for the average finite group theorist. The two tribes need each other.

gated closure of Γ'. Why must it contain the induced-convex closure of Γ' in Γ? Give an example (Γ', Γ) where the strongly gated closure of Γ' in Γ is *not* equal to the induced-convex closure.

1.5.2 Exercises for Section 1.2

1.16 Suppose $H = (H, E_H)$ is a strongly gated subgraph of the connected graph $\Gamma = (V, E)$. Let $\gamma : V \to H$ be the mapping that sends every vertex v to its unique gate in H. Show that $\gamma : \Gamma \to H$ is a morphism of graphs. (This morphism restricts to the identity mapping on H, and for this reason is sometimes called a *retract*.)

1.17 Let A be a subgroup of the automorphism group $\mathrm{Aut}(\Gamma)$ of a connected graph Γ. We construct a new graph $\Gamma/A = (V/A, E(V/A))$ whose vertex set is the collection V/A of A-orbits on the vertex set V. We say that two distinct A-orbits, x^A and y^A, belong to $E(V/A)$ if and only if there is some edge of E bridging these two orbits – equivalently, $x^\perp \cap y^A \neq \emptyset$.[13]

1. Show that the mapping $V \to V/A$, which takes each vertex to the unique A orbit containing it, induces a graph morphism $\phi_A : \Gamma \to \Gamma/A$ that is surjective on both vertices and edges.
2. Show that ϕ_A is a fibering if and only if:
   (a) Γ is connected
   (b) For each non-identity element $\alpha \in A$ and each vertex x, the distance $d_\Gamma(x, x^\alpha)$ is either zero or at least three

1.18 We say that a graph Γ is *locally connected* if and only if, for each vertex x, the "neighborhood graph" $\Gamma(x)$ induced on the vertex set $\Delta_1(x)$ is connected. (One might think of this as a graph with a good system of interlocking triangles.) Show that any arbitrary connected graph Γ contains a family $\mathcal{S} := \{\Gamma_\sigma\}_I$ of induced subgraphs with these properties:

1. Each subgraph Γ_σ is locally connected.
2. Any two distinct members of $\mathcal{S}$ share no edge.
3. Every clique of two or more vertices in Γ lies in a unique member of $\mathcal{S}$.

[Hint: For each vertex x of Γ, let $\{N_\tau(x) | \tau \in I_x\}$ be a complete listing of the connected components of the neighborhood graph $\Gamma(x)$. Let $\hat{V}$ be the set of all pairs $(x, N_\tau(x))$ where x ranges over V, and once x is fixed, τ ranges over

[13] Note that we have written $E(V/A)$ rather than the tempting "E/A" since, strictly speaking, the latter should mean the collection of A-orbits on E. But if (x^A, y^A) is to be an edge, the induced subgraph on $x^A \cup y^A$ might involve more than one edge-orbit. Letting E/A be the edge set would then commit us to multiple edges in Γ/A.

I_x. Declare the two vertices $(x, N_\tau(x))$ and $(y, N_\sigma(y))$ to be adjacent if and only if the following three conditions hold:

1. (x, y) is an edge of Γ
2. $y \in N_\tau(x)$
3. $x \in N_\sigma(y)$

Now let $\hat{E}$ be all adjacent pairs in $\hat{V}$ so that $\hat{\Gamma} = (\hat{V}, \hat{E})$ is now a graph. Observe that there is a vertex- and edge-surjective graph morphism $f : \hat{\Gamma} \to \Gamma$ obtained by "reading off the first coordinate" – that is, sending $(x, N_\sigma(x))$ to x. Now let C be any connected component of $\hat{\Gamma}$. Show that the restriction $f|_C : C \to f(C)$ is a fibering of induced subgraphs. Now, by letting $\mathcal{S}$ be the full collection of images $f(C)$ as C ranges over the connected components of $\hat{\Gamma}$, show that $\mathcal{S}$ satisfies the three required conditions 1–3 above.]

1.5.3 Exercises for Section 1.3

1.19 Prove Corollary 1.3.6. [Hint: Prove each $\Delta_i(x)$ controls $\mathcal{T}$-homotopy in $\Delta_{i+1}(x)$. From this prove all circuits are $\mathcal{T}$-contractible.]

Chapter 2
Geometries: Basic Concepts

Abstract Geometries over a typeset I are defined in terms of multipartite graphs, enabling the language of graphs to describe flags and their residues, truncations, and shadows. Similarly, morphisms of geometries are type-preserving graph morphisms. Truncation is then seen as a functor between categories of geometries. The relation of morphism and residues is more tenuous. The chapter concludes with a discussion of residual connectedness, and exercises.

2.1 Introduction

The central theme of this course concerns characterizations of geometries. For this reason one desires theorems whose hypotheses, while being kept as simple as possible, imply that the geometry to which they are attached is some well-known classical geometry.

Point-line geometries provide the simplest context in which to state the characterizing hypotheses. But the conclusion geometries are often very elaborate structures with many types of objects in addition to points and lines and, as discussed in the previous chapter, sometimes we are only able to characterize the universal covers of these geometries. As a result, we must study geometries and their morphisms in a general uniform way.

2.2 Geometries: Definitions and Basic Concepts

2.2.1 Basic Definitions

A *geometry over I* is a multipartite graph $\Gamma = (V, E)$ whose (non-empty) multipartite components have been labeled by elements of a set I. That is, $\Gamma = (V, E)$ is a graph with a function $\tau : V \to I$ whose fibers $V_i := \tau^{-1}(i)$, $i \in I$, are cocliques of Γ. In the language of geometries, the function τ is called the *type function*, I is called the *set of types*, and the vertices of a fiber V_i are called the *objects of type i*. If $\{x, y\}$ is an edge of Γ, then the vertex x is said to be *incident with vertex y* (or one

E.E. Shult, *Points and Lines*, Universitext, DOI 10.1007/978-3-642-15627-4_2,

© Springer-Verlag Berlin Heidelberg 2011

says *object x is incident with object y*). If we wish to indicate the type function, we sometimes denote a geometry Γ by the symbol $(V, E; \tau)$ rather than just (V, E). If we wish to display the types of objects in the geometry, we may also denote the geometry by a symbol, $(V_1, V_2, \ldots, V_n; *)$, where the '$*$' indicates the incidence relation, or just $(V_1, \ldots, V_n)$ if it is clear that we are talking about a geometry in which the incidence relation is already understood.

Note that the type function $\tau : V \to I$ need not be surjective. Thus if $I \subseteq J$, any geometry over I is automatically also a geometry over J. Of course, in an effort to convey as much information as possible, one usually sets $I = \tau(V)$, but the reason for not insisting on this in general has to do with morphisms of geometries.[1] Often one must consider embeddings of one geometry into another which has more types of objects.

The cardinal number $|\tau(V)|$ is called the *rank* of the geometry. This is the number of types of objects which are actually represented in the geometry.

A geometry is said to be *connected* if and only if it is connected as a graph.

2.2.2 Subgeometries

Now suppose $\Gamma = (V, E; \tau)$ is a geometry over I and that $\Gamma' = (V', E')$ is a *subgraph* of Γ as defined in the previous chapter – that is, $E' \subseteq E_{V'}$ and $V' \subseteq V$. Now it should be clear that when we restrict the type mapping τ to the vertex set V', we obtain another geometry over I, namely $(V', E'; \tau|_{V'})$. We call this the *geometry induced on the subgraph* Γ'. (Note that the rank of a geometry induced on a subgraph might be smaller than the original.)

There are two very important instances of forming subgeometries from subgraphs: truncations and residues.

2.2.3 Truncations

Suppose $\Gamma = (V, E; \tau)$ is a geometry over I and suppose J is a proper subset of I. Let Γ_J be the induced subgraph on the set of vertices V_J whose type belongs to subset J – that is, the set

$$V_J := \tau^{-1}(J) = \cup_{\sigma \in J} V_\sigma .$$

It seems natural to write E_J for E_{V_J}, the edges both of whose incident vertices have a type belonging to the set J. The subgeometry Γ_J is called the *truncation of type J of the geometry* Γ.

---

[1] The reader with experience of such things should be aware that this situation is in sharp contrast with the notion of "chamber systems of type I," where the edge-labelling mapping is required to be surjective.

Note that if J is a subset of $I - \tau(V)$, the truncation does not throw away any objects. Nonetheless, its truncation is asserted to be a geometry over J (not just I, which it also is).

2.2.4 Point-Line Geometries

In the next chapter we shall define the notion of a "point-line geometry." At the very least a point-line geometry is a rank-two geometry in which one of the types has been designated "points" and the other has been designated "lines" – that is, there is a specific distinction between the two types which one might not have paid attention to if we had just said it was a rank two geometry and left it at that. In order to make the distinction clear, we usually denote the types in a particular order. Thus the symbol $(\mathcal{P}, \mathcal{L})$ indicates a rank two geometry in which the first entry, $\mathcal{P}$, designates the set of points and the second entry, $\mathcal{L}$, designates the set of lines.[2]

The reason for this ordering of types is that so many of the concepts surrounding these geometries (e.g., subspaces, convexity, and other concepts drawn from properties of the point-collinearity graph) are not at all symmetric with respect to the two types. Thus $(\mathcal{P}, \mathcal{L})$ and $(\mathcal{L}, \mathcal{P})$ are in general to be regarded as different point-line geometries. The latter is called the *dual geometry* of the former.

Now we see that, given any geometry $\Gamma = (V, E; \tau)$ over I of rank at least three, we can potentially form a number of different point-line geometries (and their duals) by forming a truncation Γ_J to a 2-subset J of the realized type set $\tau(V)$.

2.2.5 Flags and Chambers

A *flag* of a geometry $\Gamma = (V, E; \tau)$ is simply a clique of the multipartite graph (V, E). Note that a flag can contain at most one vertex of each type. Thus the restriction of the type function to a flag F is injective. The subset $\tau(F)$ is called the *type of the flag*.

A flag F whose type is $\tau(V)$, the set of all types achieved by objects of Γ, is called a *flag chamber of* Γ (or simply "chamber of Γ" if the geometry context is clear).

2.2.6 Residues

Let $\Gamma = (V, E; \tau)$ be a geometry over $\tau(V) = I$ (without loss of generality). Let F be a flag of type J. Recall that this means that F is a clique in Γ and that $\tau(F) = J$.

[2] The definition at the beginning of Chap. 3 will ultimately require that each line is incident with at least two points and that distinct lines have distinct "point-shadows," but we need not worry about those conditions at this juncture.

The *residue of the flag* F is the induced subgraph of Γ consisting of all vertices in $F^{\perp} - F$. Note that none of these vertices can possess a type that is in J since then there would be a vertex in $F^{\perp}$ and a vertex in F which were distinct and incident and yet of the same type. Thus, by restricting the type function to the vertex set $F^{\perp} - F$, we obtain a geometry over $I - J$ – the geometry induced on the induced subgraph $F^{\perp} - F$. This geometry is denoted by $\mathrm{Res}_{\Gamma}(F)$.

There are several observations to make here.

First, it could well happen that the flag F is a maximal clique in the graph $\Gamma = (V, E)$. In that case the geometry $\mathrm{Res}_{\Gamma}(F)$ contains no objects at all – it is a rank zero geometry. Nevertheless, we are still allowed to call it a geometry over $I - J$.

The second observation has to do with how we treat information about residues. Assertions about the residues of a non-empty flag F are taken to be *local* information, and hence available as a hypothesis. On the other hand, researchers are basically artists and so are painfully aware of the distinction between hypotheses on residues of very large rank, and hypotheses on residues of rather small rank. Now there is a condition called residual connectedness, which provides a scheme for reducing assertions about high-rank residues in terms of lower rank residues. That notion tends to exclude the anomaly of the preceding paragraph where conceivable flags of relatively low rank could be maximal. We'll thrash this all out.

2.2.7 The Interplay of Residues and Truncations

If F is a flag in a geometry Γ over I and J is a subset of I, we observe the following containment as subgeometries:

$$(\mathrm{Res}_{\Gamma}(F))_J \subseteq \mathrm{Res}_{\Gamma_J}(F_J), \tag{2.1}$$

where F_J is the subset $F \cap \tau^{-1}(J)$, that is, the objects in F whose type is in J. Clearly F_J is a flag (clique) in the truncation Γ_J, so the terms in the containment make sense. Moreover, any object whose type is in J and which belongs to the left side of (2.1) is incident with every member of F_J and so lies in the right side. But the right side can be much larger: There can be elements on the right which belong to ${F_J}^{\perp} - F^{\perp}$.

At least we are forewarned.

2.2.8 Shadows

Now there is a useful notion combining truncations and residues – at least it simplifies the language. Suppose F is a flag in the geometry Γ over I. Suppose i is any type in I. Then the truncation

$$V_i(F) := (\mathrm{Res}_{\Gamma}(F))_{\{i\}}$$

is simply the rank-one geometry of all objects of type i which are equal to, or are simultaneously incident with every element of F – that is, $\tau^{-1}(i) \cap \mathrm{Res}_\Gamma(F)$. This is called the *$i$-shadow of the flag F*.

Let us first grasp the pathologies. It could happen that no element of type i exists in the residue of F. In that case we are dealing with the empty geometry over type set $\{i\}$.

The second pathological point is that it may happen that F itself contains an element of type i – say, y. Then once again the i-shadow is empty.

At best, the i-shadow of a flag is simply a coclique and so it is intrinsically uninteresting. Of course, one reason to examine the i-shadow of a collection of flags is that it provides a way of modeling the truncation $\Gamma_{I-\tau(V)}$ as a geometry of subsets of a given set V_i, the truncation of Γ to type i. This is an important philosophy since rank two truncations of a geometry may be regarded as point-line geometries (in two ways!) and there we reap harvests from regarding lines as sets of points subject to recognizable axioms.

2.3 Examples

Example 1 (The octahedron) An octahedron possesses 6 corners (which we shall call vertices), 12 edges and 8 faces. We form the three sets X_1, X_2, and X_3, consisting of the set of vertices, the set of edges and the set of faces of the octahedron, respectively. Viewing each of these objects as closed sets of points in Euclidean space, we may define "incidence" as proper containment (in some direction) among the point sets making up each vertex, edge, or face. Thus each vertex is incident with four edges and four faces; each edge is incident with just two vertices and two faces; and finally, each face is incident with three vertices and three edges. Thus the incidence relation on the objects of all three types (vertex, edge, or face) forms a tripartite graph which is a geometry of rank three.

Here, every flag lies in a chamber.

Example 2 (Desargues' configuration) Consider the famous Petersen Graph $P = (V, E)$. A simple model for this graph is the dual-triangular graph $\bar{T}_5$ on 10 vertices. Let $\Omega = \{0, 1, 2, n-1\}$ and let V be the set of $\binom{n}{2}$ 2-subsets of Ω. Let T_n be the graph with vertex set V, with two 2-subsets being adjacent if and only if they intersect at a singleton set. Then T_n is the *triangular graph* – in fact it is exactly the Johnson graph $J(n, 2)$ presented in the previous chapter. The complement of this graph is denoted by $\bar{T}_n$: it has the same vertex set except now distinct 2-subsets of Ω are considered to be adjacent if and only if they have an *empty* intersection. Thus the *Petersen graph* $P := \bar{T}_5$, has as vertices the $\binom{5}{2} = 10$ 2-subsets of the five-element set $\Omega = \{0, 1, 2, 3, 4\}$, two vertices being adjacent if and only if (as 2-subsets of Ω) they possess an empty intersection.

The graph P has valence three, no triangles, diameter two, girth five, and just two $\mathrm{Aut}(P)$-congucay classes of cocliques with representatives $C_1 := \{\{0,1\},\{0,2\},\{1,2\}\}$ and $C_2 := \{\{0,x\}|x \in \{1,...,4\}\}$. Let $\mathcal{L}$ be the collection of ten 3-cocliques of the class C_1^G where $G = \mathrm{Sym}(\Omega)$. Finally, we let $\triangledown = (V, \mathcal{L})$ be the rank two geometry of 10 vertices (the "points") and 10 maximal 3-cocliques of P (the "lines", $\mathcal{L}$) where incidence is membership of a vertex in a coclique. This is called *Desargues' configuration*. Each of its points is on three lines; each of its lines has three points. For each point p, the set of points $V - p^{\perp}$ itself constitutes a line of three points.

Example 3 (The Payley designs) Let F be a finite field of q elements where q is an odd prime power such that $q \equiv 3 \bmod 4$. The multiplicative group of non-zero elements of F is cyclic of even order $q-1$ and the set Q of non-zero squares forms a subgroup of index 2 in this multiplicative group. (In the case that q is a prime p, then F is just the ring of integers modulo p, and Q is the set of quadratic residues mod p.) The elements of the field F will be called *points* and denoted by $\mathcal{P}$. The additive translates, $Q + x$, $x \in F$ will be called *lines* and we denote the set of all such translates by the symbol $\mathcal{L}$. A point is incident with such a translate if and only if it is one of its elements. Thus each point is incident with $(q-1)/2$ lines and each line is incident with $(q-1)/2$ points.

The importance of the congruence condition is that it forces any two distinct lines to intersect at $(q-3)/4$ points. (In other congruences there are two distinct intersection sizes.) Such a structure of points and blocks in which there are equally as many points as blocks and in which any two blocks intersect at λ points also enjoys the property that any two distinct points together lie in exactly λ blocks – and so is called a *square design*. A square design with $\lambda = 1$ it is called a *finite projective plane*. There are zillions of these. If $\lambda = 2$ it is called a *finite biplane*. There are just a few of these known.

Example 4 (The classical projective geometries) Let W be a right vector space of finite dimension $n \geq 2$ over a division ring D. Let (W, E) be the graph whose vertices are all the *proper* subspaces of W – that is, 0 and W itself have been excluded – and whose edges are unordered pairs $\{A, B\}$ of proper subspaces of W, one of which is properly contained in the other. The type function $\tau : V \to \{1, \ldots, n-1\} = I$ records the dimension of any proper subspace. As we know, this function is surjective. Then $\Gamma = (V, E : \tau)$ is a geometry over I of rank $n-1$ – *the classical projective geometry of projective dimension* $n-1$ *over* D, also denoted by $PG(W)$ or $PG(n-1, D)$.

Here a chamber is a set $\{U_1, \ldots, U_{n-1}\}$ of subspaces U_i of W of dimension i, with $U_i < U_{i+1}$. Every flag is a subset of such a set.

The truncation of a projective geometry $PG(V)$ to one- and two-dimensional subspaces is called a *classical projective space* – rather than a classical projective *geometry*. It is a point-line geometry $(\mathcal{P}, \mathcal{L})$ whose points are the set $\mathcal{P}$ of all one-dimensional subspaces of V and whose lines are the set $\mathcal{L}$ of all two-dimensional subspaces of V.

Example 5 (A Polar space geometry) Let $V = \mathbf{R}^{(6)}$, the vector space of all 6-tuples with real coefficients. Let $Q : V \to \mathbf{R}$ be the function defined by the formula

$$Q(x_0, x_1, \ldots, x_5) = x_0^2 + x_1^2 + x_2^2 - x_3^2 - x_4^2 - x_5^2.$$

Such a function (defined by a homogeneous quadratic polynomial in the x_i) is called a *quadratic form*. A subspace W of V is said to be *totally singular* if and only if $Q(W) := \{Q(w)|w \in W\} = \{0\}$ — that is, Q vanishes on subspace W (and perhaps elsewhere).

There are three varieties of proper totally singular subspaces of V: those of dimensions 1, 2, and 3, respectively. For example: the set of vectors A of shape (a, b, c, a, b, c), a, b, c real numbers, form a totally singular subspace of dimension 3. Similarly, the set B of vectors of shape (a, b, c, a, c, b) is also a totally singular subspace, and the intersection of these two subspaces

$$A \cap B = \{(a, c, c, a, c, c)|a, c \in \mathbf{R}\}$$

is a two-dimensional totally singular subspace.

Let $\mathcal{S}_i$ be the set of all totally singular subspaces of dimension i, $i = 1, 2, 3$. If we say that two totally singular subspaces are *incident* if and only if one is properly contained in the other, then we obtain a rank three geometry $\Gamma = (\mathcal{S}_1, \mathcal{S}_2, \mathcal{S}_3; *)$.

There are many variations of this example using proper subspaces of a vector space which are totally isotropic or totally singular with respect to a sesquilinear form or a quadratic form on a finite dimensional vector space. All finite ranks are possible. Such geometries are called *classical polar space geometries*. Their truncations to their one- and two-dimensional varieties are called *classical polar spaces* and are point-line geometries.

Example 6 (Coset geometries) Let G be a group and let $\{H_\sigma|\sigma \in I\}$ be a collection of pairwise distinct subgroups of G indexed by the set I. Let $X_\sigma = G/H_\sigma$ be the full set of right cosets of H_σ in G. Let $X = \cup_I X_\sigma$. Two right cosets in X will be declared to be incident if and only if they are distinct and possess a non-empty intersection. Let (X, E) be the graph of this incidence relation. Then it is clear that the sets X_σ are pairwise disjoint cocliques so that the mapping $\tau : X \to I$ which sends all members of X_σ to the indexing subscript σ is a well-defined type function of the geometry $(V, E; \tau)$. This geometry, being completely determined by G and the specified family of subgroups, is given a special denotation: it is the *coset geometry* $(G; \{H_\sigma\}_I)$. In the case where the collection of subgroups is a small finite list, we discard the extra wavy brackets and write $(G; H_1, \ldots, H_n)$, which is then a geometry of rank n.

This construction is more general than it looks, for it can represent any geometry with a sufficiently rich automorphism group. All of the previous five examples are actually coset geometries.

2.4 Morphisms of Geometries

2.4.1 Definition

Let $\Gamma_i = (V_i, E_i : \tau_i)$, $i = 1, 2$ be two geometries over the same type set I. A *morphism of geometries over* I is a graph homomorphism $\phi : \Gamma_1 \to \Gamma_2$ which does not alter the types of the vertices. Thus for each vertex $v \in V_1$, we must have $\tau_2(\phi(v)) = \tau_1(v)$.

The collection of all geometries over I forms a category with respect to the geometry morphisms, for the latter can be composed where it makes sense.

2.4.2 Automorphisms

If the underlying graph homomorphism is an isomorphism of graphs, then the geometry morphism is called an *isomorphism of geometries*. An isomorphism of a geometry with itself is called an *automorphism*, and of course all automorphisms form a group $\mathrm{Aut}(\Gamma)$ preserving the types.

Obviously, an automorphism of the geometry Γ must take any flag of a given type to another flag of the same type. A group of automorphisms is said to be *flag transitive* if and only if it is transitive on the set of flags of any given type, no matter what the typeset. It is said to be *chamber transitive* if and only if it is transitive on the flags of type $\tau(V)$. In a residually connected geometry of finite rank, the two concepts flag transitive and chamber transitive are easily seen to be the same (Exercise 2.2 at the end of this chapter).

Let's take a look at our examples.

In Example 1, it should be clear that the automorphism group of the octahedron geometry is transitive on vertices and that the stabilizer of a vertex faithfully induces the dihedral group of order eight on the subgeometry of edges and faces incident with it (a so-called *residue* of the vertex). So the group has order 48. Now there is a graph induced by the (appropriately named) vertices and edges of the geometry (as we shall see, this is an example of a *truncation*) which is the tri-partite graph $K_{2,2,2}$. It follows easily that the full automorphism group is $Z_2 \times \mathrm{Sym}(4)$.

In Example 2, the stabilizer of a point is $\mathrm{Sym}(3) \times Z_2$, which induces $\mathrm{Sym}(3)$ on the three lines on the point. The observation that $V - x^\perp$ is a line shows that a point-stabilizer here is already a line stabilizer – namely a 3-Sylow-normalizer of $\mathrm{Aut}(\triangledown)$. This shows that there is a natural duality – that is, an isomorphism between a point-line geometry and its dual,

$$\triangledown = (V, \mathcal{L}) \to \triangledown^t := (\mathcal{L}, V).$$

In Example 3, we have a point-line geometry $(\mathcal{P}, \mathcal{L})$ whose points are the elements of the field $F = \mathbf{G}F(q)$, $q \equiv 3 \bmod 4$, and whose lines are the q translates $Q + x$, where x ranges over F. Obviously addition of y to all points is a permutation of the points which takes translate $Q + x$ to $Q + (x + y)$ and so

produces a permutation of the lines which preserves incidence. This produces a group T of translations which is isomorphic to $(\mathbf{G}F(q), +)$, the additive group of the field. Similarly, right multiplication of all points (field elements) by a non-zero square s also permutes the points in a manner that transforms translate $Q + x$ to $(Q + x)s = Q + xs$. We call this multiplative group Q (with only a small abuse of notation). Thus the group of transformations

$$\{a_{y,s} : x \to xs + y, s \in Q\}$$

is a group TQ of automorphisms of order $q(q-1)/2$ which is flag-transitive. (The notation TQ is justified since the translation group T is a normal subgroup of this group, with complement Q.)

We can utilize these automorphisms to show that any two distinct lines intersect at a set of a fixed cardinality. Suppose x and y are distinct elements of F, and consider the pair of lines $(Q + x, Q + y)$. We can translate this pair (as well as its intersection) to the pair $(Q, Q + (y - x))$. By hypothesis, $y - x$ is non-zero. Since $q \equiv 3$ modulo 4, -1 is a representative non-zero non-square – that is $F^{\#}$ is the disjoint union, of the two cosets Q and $(-1)Q = -Q$. So there exists a square s such that $(x - y)s = \pm 1$. Thus right multiplication by s brings the pair $(Q, Q + (y - x))$ to the pair $(Q, Q + 1)$ or to the pair $(Q, Q - 1)$. But translation by $+1$ brings the latter to the former. (In effect, we have shown that our group of translations and square multiplications is 2-homogeneous on lines.) It follows that all intersections of distinct lines are sets of the same size in one orbit under the automorphism group.

There are two subcases where the group of automorphisms is actually larger than the group TQ.

The case $q = 7$. We can take the points $\mathcal{P}$ to be the set $\{0, 1, \ldots, 6\}$ representing the residue classes mod 7. Then the set of quadratic residues is $Q = \{1, 2, 4\}$. Then the seven translates are listed in Fig. 2.1b, and can be arranged as a point-line geometry by the so-called *Fano configuration* which appears in Fig. 2.1a. Here the seven points are labeled, and the lines are the six straight lines of the figure as well as the triplet of points which make up the central circle.

The reader who has not done it before in an earlier life, should try to compute the order of the group of automorphisms. The order of the full automorphism group fixing a point can be calculated from the left figure, and has index $|T|$ because of T is transitive. This object is called the *projective plane of order* 2.

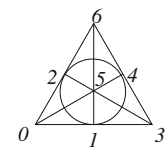

(a)

$B_0 = \{1,2,4\}$
$B_1 = \{2,3,5\}$
$B_2 = \{3,4,6\}$
$B_3 = \{4,5,0\}$
$B_4 = \{5,6,1\}$
$B_5 = \{6,0,2\}$
$B_6 = \{0.1.3\}$

(b)

Fig. 2.1 The Fano configuration (**a**), and the blocks of the Payley design (**b**) for $q = 7$. Here B_i is $Q + i$ of the text

The case $q = 11$. The geometry $(\mathcal{P}, \mathcal{L})$ has 11 points and 11 "lines". Any two lines intersect in exactly 2 points, so it is a biplane.

Again, one can show that the stabilizer of a point is the alternating group Alt(5), acting transitively on the remaining ten points as it does on the vertices of Petersen's graph. But this takes a little work. As a result $\mathrm{Aut}((\mathcal{P}, \mathcal{L}))$ is a group of order 660. (It is in fact the simple group $PSL(2, 11)$ in an unusual doubly transitive action on 11 letters.)

In Example 4, it should be clear that the group $GL(V)$ (of all bijective linear transformations of vector space V onto itself) is a flag-transitive subgroup of the automorphism group of the projective geometry $PG(V)$. (Recall that if V has finite dimension n over the division ring D, one can write $GL(n, D)$ for $GL(V)$ and $PG(n-1, D)$ for $PG(V)$.)

Let Q be the quadratic form on $V = \mathbf{R}^{(6)}$ of Example 5. The subgroup of $GL(V)$ consisting of transformations $t : V \to V$ which *preserve the form* – that is, $Q(t(v)) = Q(v)$ for all vectors v – form a subgroup $GO(V, Q)$, called *the general orthogonal group with respect to the quadratic form* Q. This group acts flag-transitively on the (orthogonal) polar space geometry of Example 5.

In Example 6 concerning the coset geometry $(G; \{H_\sigma\}_I)$, right multiplication by any group element g of G permutes the elements of $X_\sigma = G/H_\sigma$, thus permutes their intersections and so preserves incidence as well as the type function, and so induces an automorphism $\alpha(g)$. It easily follows that

$$\alpha : G \to \mathrm{Aut}(G; \{H_\sigma\}_I)$$

is a group homomorphism. Its kernel is the intersection of the groups N_σ where N_σ is the *core* of the group H_σ – that is, the group $N_\sigma := \cap_{x \in G} x^{-1} H_\sigma x$, the kernel of the action of G on X_σ.

2.4.3 Morphisms Defined by a Group of Automorphisms

At this point we should mention a very canonical example of a morphism of geometries, one already fully discussed in the more general context of graphs.

Suppose $\Gamma = (V, E; \tau)$ is a geometry over I. Suppose A is a subgroup of $\mathrm{Aut}(\Gamma)$. Let V/A be the collection of A-orbits on the objects V of the geometry. Since A preserves types, all objects within an A-orbit, v^A, possess the same type which can be attached to this orbit. This gives us a function $\tau/A : V/A \to I$. Now say that an A-orbit x^A is *incident* with an A-orbit y^A if and only if some vertex in x^A is adjacent (that is, incident) with some vertex in y^A. As we have seen in the previous section, this construction produces a *full* homomorphism of the underlying graphs, and since the partition of the objects into A-orbits only refines the partition into τ-fibres, we have a homomorphism of geometries.

Example 1 provides a very nice instance of this. Each vertex of the octahedron is associated with a unique antipodal vertex at distance two from it in the subgraph of vertices and edges. Exchanging each vertex with its antipodal partner preserves

edges and triangles, and is an automorphism z of order two which is central in the automorphism group. Moreover, for each edge and triangle there is a unique opposite image under z. Let $A = \langle z \rangle$, a group of automorphisms of order two. Then $\Gamma/A = (V/A, E/A)$ is a geometry over I with just three vertices, six edges, and four triangles. This geometry can be modeled on the tetrahedron with the same edges and faces, but now "vertices" are pairs of skewed edges of the tetrahedron.

Here is another example. Let V be the the vector space $\mathbf{C}^{(n)}$ of all n-tuples of complex numbers, with the usual addition and scalar multiplication. Let $Q : V \to \mathbf{C}$ be the "sum of squares" quadratic form – this is, for any n-tuple $\vec{x} = (x_1, \ldots, x_n)$, $Q(\vec{x}) = \sum x_i^2$. A subspace U of V is said to be *a Q-singular vector subspace* if and only if it "vanishes" under the quadratic form, that is,

$$Q(U) = \{Q(\vec{u}|u \in U\} = \{0\}.$$

There exist Q-singular subspaces of V for every dimension less than or equal to $n/2$. If we think of the non-zero Q-singular subspaces of V as the objects of a geometry with the dimension as a type function, we obtain a geometry Γ of rank $\lfloor n/2 \rfloor$. For example, when $n = 11$, we obtain a geometry with five types of objects.

Now consider the mapping which takes each vector $\vec{x} = (x_1, \ldots, x_n)$ to the vector of complex conjugate coordinates $(\bar{x_1}, \ldots, \bar{x_n})$. This operator is easily seen to preserve linear dependence (and independence) of sets of vectors, and preserves the singularness (or non-singularness) of any vector. It follows that it induces an automorphism σ of the geometry Γ of order two. Setting $A = \langle \sigma \rangle$, we obtain a homomorphism Γ. In Exercise 2.4 at the end of this chapter, the reader is asked to prove that σ fixes no object of the geometry – that is, *sigma acts freely on the geometry* Γ.

2.4.4 Truncations and Morphisms

One now notices the following: if $\rho : \Gamma \to \Sigma$ is a morphism connecting two geometries over the same type set I, then the restriction of ρ to vertices whose type belongs to J induces a morphism of the corresponding truncations of these geometries: $\rho_J : \Gamma_J \to \Sigma_J$. In the language of categories, that means that truncation to type J is a *functor* from the category of geometries over I to the category of geometries over J.

2.4.5 Residues and Morphisms

Suppose F is a flag in the geometry Γ and $\phi : \Gamma \to \Gamma'$ is a morphism among geometries of type I. Then there is an induced geometry morphism

$$\mathrm{Res}_\Gamma(F) \to \mathrm{Res}_{\Gamma'}(\phi(F)).$$

In most cases this is not a surjection. (Of course it is injective or full if ϕ is.)

2.5 Connectedness Properties

We have mentioned elsewhere that a geometry is essentially a multipartite graph, and that the geometry is "connected" if and only if it is connected in the graph-theoretic sense.

2.5.1 Residual Connectedness

Let Γ be a geometry over a finite type set I. Recall that there is then a (possibly non-surjective) type map $\tau : V \to I$ and that no two vertices of the same type are adjacent. Recall that the cliques of Γ are called flags, and that $\tau(F)$ is called the type of the flag F. Moreover, by restriction of the type function, the residue of a flag F of type J is the induced subgraph

$$\mathrm{Res}_\Gamma(F) := F^\perp \cap (V - F),$$

which is a geometry over $I - J$.

A flag of type I is called a *chamber* (or if necessary, a "flag chamber"). Two chambers are said to be *adjacent chambers* if and only if they differ in only one object. We now form the graph $C(\Gamma)$ whose vertices are all flag chambers of Γ, with adjacency as just defined.[3] A geometry Γ is said to be *chamber connected* if and only if $C(\Gamma)$ is connected – that is, the graph of the adjacency relation on chambers is connected. A geometry with no chambers is chamber connected.

Whether all types are realized or not, the *rank* of $\Gamma = (V, \tau, *)$ is defined to be $|\tau(V)|$ and is denoted by $\mathrm{rnk}(\Gamma)$ (see p. 44). The *corank* of a flag F of type J is $|\tau(V) - J|$.

We say that such a geometry is *residually connected* if and only if the following two conditions hold:

(RC1) *Every flag of corank one has a non-empty residue.*
(RC2) *Every flag of corank two or more has a connected non-empty residue.*

This notion has immediate consequences.

Lemma 2.5.1 *Suppose Γ is a residually connected geometry. Then:*

1. *Every residue of Γ is itself residually connected*
2. *Every flag lies in a chamber*
3. *If Γ has finite rank, then every residue of Γ is chamber connected*

[3] In fact one can insert more information into this graph. We say that two flag chambers F and G are *i-adjacent* if and only if they differ only by an object of type i. We can then include this information by attaching the label "i" to the edge (F, G) of $C(\Gamma)$. We now have a new thing: a graph with edge labels which is denoted by $C(\Gamma)$, and is called the *chamber system derived from the geometry* Γ. We shall have much more to say about the relations between the categories of chamber systems and the categories of geometries in a later chapter. Residual connecteness will have a crucial role to play there.

Proof 1. Let F_0 be a flag in residue $R := \mathrm{Res}_\Gamma(F_1)$. Then $F := F_0 + F_1$ is a flag of Γ and by hypothesis

$$\mathrm{Res}_\Gamma(F) = \mathrm{Res}_R(F_0)$$

is connected if rank two or more, and non-empty if rank one. Thus R is residually connected.

2. Any flag lies in a maximal flag (in fact by a simple application of Zorn's Lemma, any clique in a graph lies in a maximal clique). Let F be a maximal flag. Then $\mathrm{Res}_\Gamma(F)$ is empty. If F has positive corank, then by (RC1) and (RC2), $\mathrm{Res}_\Gamma(F)$ is non-empty, a contradiction. Thus F has corank zero – that is, it is type I, and so is a flag-chamber. Part 2 is proved. (Note that the proof did not require finite rank of Γ.)

3. Since the empty flag has a connected residual, Γ is a connected geometry. As a result the graph-theoretic distance between any two objects of the geometry is finite. Suppose Γ has been chosen of minimal rank such that there are pairs of chambers in distinct connected components of the chamber adjacency graph, $C(\Gamma)$. Among such pairs choose a pair (X, Y) so that

$$d = d(X, Y) = \min\{d(x, y) | (x, y) \in X \times Y\}$$

is minimal.

If $X \cap Y = F$ is a non-empty flag, then, because

$$\mathrm{rnk}(\mathrm{Res}_\Gamma(F)) < \mathrm{rnk}(\Gamma),$$

$X_0 := X - F$ and $Y_0 := Y - F$ are chambers of $\mathrm{Res}_\Gamma(F)$ connected by a chain $(X_0, X_1, \ldots, X_m = Y_0)$ of chambers, each adjacent to its successor. Then setting $F_i = F \cup X_i$, we obtain such a chain $(X = F_0, F_1, \ldots F_m = Y)$ connecting X and Y, a contradiction. Thus $X \cap Y = \emptyset$.

Now there exists a geodesic path $(x = x_0, x_1, \ldots, x_d = y)$ connecting a point x in X with a vertex y in Y. By 2. there is a chamber Z containing $\{x, x_1\}$. Then F is in the same chamber-adjacency component as X since $X \cup Z$ is non-empty. It is also in the same component as Y since $d(Z, Y) < d$ and d was minimal. Thus X and Y are in the same component, contrary to the way they were chosen. The proof is complete. □

The *truncation* of type J of the geometry Γ is the induced graph with vertex set $V_J = \tau^{-1}(J)$, regarded as a geometry Γ_J over J under the restriction of the type mapping. In short, one simply ignores all vertices whose type is not in J.

Theorem 2.5.2 *Let Γ be a residually connected geometry of finite rank and let Γ_J be its truncation of type J, a subset of I of cardinality at least 2. Let $\mathcal{T}$ be the collection of all triangles of the graph Γ.*

Then any path in Γ *connecting two vertices of* Γ_J *is* $\mathcal{T}$*-homotopic to a path in* Γ_J *connecting the two vertices. That is, the truncation* Γ_J *controls* $\mathcal{T}$*-homotopy in* Γ.

Proof We assume Γ residually connected and I chosen of minimal rank so that the theorem fails. It suffices to show that the truncation Γ_J, as a subgraph of Γ satisfies the two conditions (1) and (2) in the hypothesis of Theorem 1.3.5 in Chap. 1.

Let M be a maximal clique of Γ. Then as Γ is residually connected, M is a chamber, and so contains vertices of type J. So condition (2) holds – that is every maximal clique meets Γ_J.

Suppose x is a vertex of Γ whose type is not in J. Then by Lemma 2.5.1 part 1, $\mathrm{Res}_\Gamma(x)$ is residually connected of smaller rank than Γ. By induction, *its* truncation of type J — that is $x^\perp \cap \Gamma_J$ – being of rank at least two, controls $\mathcal{T}$-homotopy in $\mathrm{Res}_\Gamma(x)$, and so at the very least, is a connected graph. Thus as $x^\perp \cap \Gamma_J$ is always connected for x exterior to the truncation, condition (1) of Theorem 1.3.5 holds. □

Corollary 2.5.3 *If* Γ *is a residually connected geometry of finite rank, then any proper truncation of rank at least two is also residually connected.*

Remark A "proper truncation" is one in which J meets $\tau(V)$ at a proper subset. That is, some actual objects of Γ are missing in the truncation. The reason that condition is there is that for truncations Γ_J, where J contains $\tau(V)$, nothing is lost – Γ_J *is* Γ, but with a reduced type set. (For example this is true when $J = I = \tau(V)$.) In that case nothing new is asserted by the theorem, since the hypothesis that Γ is residually connected immediately appears as part of the conclusion. Thus throwing out this case costs us nothing. On the other hand, it allows us to state a converse which (as one can see) would be patently trivial otherwise.

Proof Any flag F_0 in Γ_J lies in a chamber F of Γ and $F \cap \Gamma_J$ is a chamber of Γ_J containing F_0. So (RC1) holds for $\mathrm{Res}_\Gamma(F_0)$. So assume $\mathrm{rnk}(\mathrm{Res}_\Gamma(F_0)) \geq 2$. Now by Lemma 2.5.1, $\mathrm{Res}_\Gamma(F_0)$ is residually connected, so by Theorem 2.5.2 its truncation $(\mathrm{Res}_\Gamma(F_0))_J$ is connected. But as

$$(\mathrm{Res}_\Gamma(F_0))_J = \mathrm{Res}_\Gamma(F_0) \cap \Gamma_J = \mathrm{Res}_{\Gamma_J}(F_0),$$

(RC2) holds and we are done. □

Remark In a paper [20] on "basic diagrams" F. Buekenhout defined a geometry to be *strongly connected* if and only if every truncation of rank two was connected. He proved that a residually connected geometry was always "strongly connected." The reader is asked to observe that that theorem is contained in Corollary 2.5.3, for certainly every rank two residue, being residually connected, is connected. The corollary is in fact a stronger statement.

The converse of Corollary 2.5.3 (as we have been careful to phrase it) is in fact not true. Consider the following rank three geometry.

Counterexample to the converse of Corollary 2.5.3. Consider a regular pentagon embedded in the plane and draw all diagonals to form a complete graph K_5 embedded there (see Fig. 2.2). Let V be the 5 vertices (corners) of the pentagon, labelled

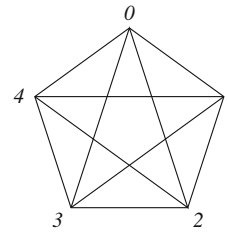

Fig. 2.2 The special rank three geometry derived from the pentagram

in clockwise order: (0, 1, 2, 3, 4). Let E_1 be the "outside edges" – indicated by their extremal vertices as $\{\{i, i+1\} | i \in \mathbf{Z}/(5)\}$. Let E_2 be the diagonal edges — that would be the set of edges $\{\{i, i+2\} | i \in \mathbf{Z}/(5)\}$ (see Fig. 2.2). We must define incidence. A vertex in V is incident with an edge if and only if it is one of its extremal vertices. That defines incidence in $V \times E_1$ and $V \times E_2$. We must next define incidence in $E_1 \times E_2$ We say that a pair of edges $(e, f) \in E_1 \times E_2$ are *incident* if and only if (1) they meet at a vertex v, and (2) their interior angle about v is as small as possible.[4] Thus the "exterior edge" $\{0, 1\}$ of E_1 is declared to be incident to the two edges $\{0, 2\}$ and $\{1, 4\}$ of E_2, but is *not* incident with the edges $\{1, 3\}$, $\{0, 3\}$, or $\{2, 4\}$ which it does not even intersect.

Now when we view the residual of the edge $e = \{0, 1\}$ in the geometry $\Gamma := (V, E_1, E_2; *)$ we obtain the subgeometry consisting of the two vertices 0 and 1 and the two edges $\{0, 2\}$ and $\{1, 4\}$. As a subgeometry, it is easy to see that this is not connected. It follows that Γ is not residually connected. Nonetheless, when we look at the three rank-two truncations $((V, E_1; *)$, $(V, E_2; *)$, or $(E_1, E_2; *))$ all of them are the incidence system of the vertices and edges of a pentagon – as a bipartite graph, it is a simple 10-circuit. In any case, all of these truncations are connected. We thus have a rank three geometry in which every residue of rank two is connected (i.e., Γ is strongly connected) yet the geometry itself is not residually connected.

2.6 Exercises for Chapter 2

2.1 Suppose $\Gamma = (V, I; \tau)$ is a geometry over I of rank $|\tau(V)|$ at least two. Following F. Buekenhout, say that Γ is *strongly connected* if and only if every rank-two truncation of Γ is a connected geometry. Write out a simple proof of the following Theorem using Corollary 2.5.3.

Theorem 2.6.1 (F. Buekenhout [20]) *For a geometry Γ of rank at least two, residual connectedness implies strong connectedness.*

Does this theorem hold for infinite rank? [Rexamine the proof of Theorem 2.5.2 in this light.]

[4] That happy phrase saves one from actually having to measure the angles formed between sides and diagonals.

2.2 Suppose $\Gamma = (V, I; \tau)$ is a geometry over I in which every type in I is *realized* – that is, $\tau(V) = I$. Show that if Γ is chamber transitive, then $\mathrm{Aut}(\Gamma)$ transitively permutes the flags of finite cotype which lie in some flag-chamber. [Hint: This is obviously an induction proof.] As a corollary, conclude that the automorphism group of a residually connected chamber transitive geometry of finite rank is transitive on all flags of a given type.

2.3 Suppose Γ is a geometry over $I := \tau(V)$. If Γ is residually connected, show that Γ need not be chamber connected. [This was proved in part 2 of Lemma 2.5.1 for the case that the rank of Γ was finite. So the problem here is to find a counterexample of infinite rank.]

2.4 For $n>1$, let Γ be the geometry of subspaces of $V = \mathbf{C}^{(n)}$ which are singular with respect to the sum of squares form $Q : V \to \mathbf{C}$. Let σ be the automorphism of Γ which is induced by the mapping which takes the complex conjugate of each coordinate of a vector $\vec{x}$ in $\mathbf{C}^{(n)}$. Show that if p is a 1-subspace generated by a singular vector, then p and p^σ cannot both be in a common singular subspace. In particular, σ acts freely on the geometry.

2.5 Suppose Γ is a geometry $(V, E; \tau)$ over $I = \tau(V)$ with these properties: (1) every flag of Γ lies in a flag chamber; (2) Γ is chamber-connected. Show that Γ is a connected geometry. [Do not assume I is finite.]

Chapter 3
Point-Line Geometries

Abstract Point-line geometries are just rank two geometries, and so inherit the concepts of morphism and cover from the last chapter. The symmetry between the two types is broken by the concept of a subspace, which treats points differently from lines. A new graph, the point-collinearity graph, is useful in describing geometric properties. Singular spaces, partial linear spaces, linear spaces, and gamma spaces, all of which appear among the Lie incidence geometries, are introduced at this point. Of special importance is the notion of a locally connected component of a gamma space. The exercises present a number of classical examples and questions concerning product geometries.

3.1 Introduction

3.1.1 On Choosing a Reasonable Definition

This chapter concerns geometries $\Gamma = (\mathcal{P}, \mathcal{L})$ over a type set of size two in which (for purposes of discussion) one of the two types has been designated as "points" (denoted by $\mathcal{P}$), the other as "lines"(denoted by $\mathcal{L}$). This means that Γ is a simple bipartite graph, $(\mathcal{P} \cup \mathcal{L}, E)$ with $\mathcal{P}$ and $\mathcal{L}$ being the two fibers of a type function $\tau : V \to \{\text{points}, \text{lines}\}$.

It is important to note that we are talking about geometries *over* the type set {points, lines}. They need not actually be rank two geometries since either of the two sets $\mathcal{P}$ or $\mathcal{L}$ could be empty.

This entire book is ultimately about geometries of points and lines with simple extra axioms that force them to be classical Lie incidence geometries. We have noted from examples of geometries of points and lines in Chap. 2 that most of the properties and axioms involving points and lines do not say the same thing when we transpose the words "points" and "lines" in the statements of these properties. In fact the whole basic theory which characterizes classical geometries is built from *asymmetric* hypotheses on points and lines, the concept of "subspace" being a prime example.

E.E. Shult, *Points and Lines*, Universitext, DOI 10.1007/978-3-642-15627-4_3,
© Springer-Verlag Berlin Heidelberg 2011

In our definition of "point-line geometry," which will appear below, we will attach two further modest properties to a geometry of points and lines. But which axioms should we add?

As remarked in our discussion of residual connectedness in the previous chapter, there has been a practice of building into the basic definitions hypotheses universally enjoyed by the geometries one's theory hopes to characterize. After all, such definitions do not restrict the class we are trying to describe, and at the same time they seem to "streamline" the presentation of proofs.

But there is a difference between simplicity at the pedagogical level of "presentation," where the teacher explains how the assumptions imply the conclusion as neatly as possible, and simplicity relative to the landscape of discovery. In the latter world, no hypothesis is invoked unless it is really needed in the proof. That way future generations of students are invited to extend a theorem without constantly having to alter the definition of each "X" to a "generalized X."

Another reason for being very conservative about inculcating unnecessary properties into the basic definitions is that if we allow such inherent assumptions, the class of objects being defined may no longer be closed under morphisms, products, constructions derived from intersections with shadows of subobjects, or other pertinent instruments of its natural category. It is for this reason that we should surely want the definition of point-line geometry to allow the sets $\mathcal{P}$ and $\mathcal{L}$ to be empty.

3.1.2 Our Definition of Point-Line Geometry

On the other hand one can go too far in the other direction of not making assumptions. Certainly if nearly every proof that one would make requires a certain assumption, then that assumption should be part of the definition. For this book we adopt the following.

Definition A *point-line geometry* is a triple $(\mathcal{P}, \mathcal{L}, *)$ where $\mathcal{P}$ and $\mathcal{L}$ are disjoint sets (whose elements are called "points" and "lines," respectively) and $*$ is a symmetric relation (called "incidence") on $\mathcal{P} \times \mathcal{L}$, subject to these two axioms:

1. (Lines are firm.) Every line is incident with at least two points.
2. (There are no repeated lines.) If two lines are distinct, then there is a point incident with one of the lines, which is not incident with the other line.

If L is a line of $\Gamma = (\mathcal{P}, \mathcal{L})$, the set of points incident with L is called the *point shadow of* L and is denoted by $sh_{\mathcal{P}}(L)$ or simply $sh(L)$ if the point set is understood. The second axiom for a point-line geometry asserts that distinct lines (if there are any) must have distinct point shadows.

To indicate that Γ is such a point-line geometry, we often suppress the name of the incidence relation (the "$*$") in the the notation, and simply write $\Gamma = (\mathcal{P}, \mathcal{L})$. As mentioned before, we have not begun with the "bipartite graph" description, in

order to prepare the way for the asymmetry of the two types in our definitions.[1] However, if $\mathcal{P}$ and $\mathcal{L}$ are both non-empty, the bipartite graph (sometimes called *the incidence graph of* Γ) is an important point of view of point-line geometries with many uses – for example, in describing "dual geometries" when this is allowed.

3.2 The Point-Collinearity Graph

Before venturing further, let us define the *point-collinearity graph*. Two points of $\mathcal{P}$ are said to be *collinear* if and only if there exists a line L in $\mathcal{L}$ incident with both points. If every point is incident with at least one line (this must occur if the geometry is connected) then point-collinearity is a reflexive relation. The *point-collinearity graph* is the graph whose vertex set is the set $\mathcal{P}$ of points, and whose edges are unordered pairs of distinct points which are incident with a common line. We denote this graph by the symbol $\Delta = (\mathcal{P}, \sim)$ where "$\sim$" denotes the collinearity relation restricted to pairs of distinct points. (The "Δ" is there to distinguish it completely from the other graph Γ, the bipartite one.)

Now if x is a point of $\mathcal{P}$ in a point-line geometry $(\mathcal{P}, \mathcal{L})$, we let $x^\perp$ have the exact meaning that symbol would have for the point-collinearity graph $\Delta = (\mathcal{P}, \sim)$ – that is, $x^\perp$ is the point x together with the set of all points y which are collinear with x.

3.3 Morphisms and Covers of Point-Line Geometries

This brief section simply recasts the notions of morphisms and covers of geometries to point-line geometries. There is a slight wrinkle. Of course we could simply define a morphism $\phi(\mathcal{P}, \mathcal{L}) \to (\mathcal{P}_1, \mathcal{L}_1)$ as we would for any rank-two geometry. The problem is that the image geometry $(\phi(\mathcal{P}), \phi(\mathcal{L}))$ may not also be a point-line geometry. So instead of defining a morphism of point-line geometries as a type of geometry morphism, we describe it in terms of a morphism of the point collinearity graphs.

A *morphism of point-line geometries* $\phi : \Gamma = (\mathcal{P}, \mathcal{L}) \to (\mathcal{P}_1, \mathcal{L}_1) = \Gamma_1$ is a morphism $\Delta \to \Delta_1$ of the point-collinearity graphs of Γ and Γ_1, respectively, with this property:

1. *For any line $L \in \mathcal{L}$, the restriction of ϕ to its point shadow $sh_{\mathcal{P}}(L)$ is either a bijective mapping of L onto the point-shadow $sh_{\mathcal{P}_1}(L_1)$ of some line $L_1 \in \mathcal{L}_1$, or else $\phi(sh_{\mathcal{P}}(L))$ is a single point of $\mathcal{P}_1$.*

Thus if we set $\mathcal{L}_\phi$ to be the collection of subsets $\phi(sh_{\mathcal{P}}(L))$ of $\mathcal{P}_1$ which contain at least two points, then the condition described above ensures that $\phi(\Gamma) := (\phi(\mathcal{P}), \mathcal{L}_\phi)$ is a point-line geometry. (Of course, here $\phi(\mathcal{P}) \subseteq \mathcal{P}_1$, $\mathcal{L}_\phi \subseteq \mathcal{L}_1$ and the $\mathcal{P}_1$-point shadows of elements of $\mathcal{L}_\phi$ are contained in $\phi(\mathcal{P})$.)

[1] There is an even more mundane reason: in the thinking of many combinatorialists, the two vertex-components of a bipartite graph must be non-empty.

Here is a simple example. The nine vertices and six maximal cliques of the three-by-three grid form a geometry of points and lines. But by shrinking one parallel class of lines to single points, one obtains a morphism of the grid-geometry to the geometry consisting of a single line with three points.

So far, everything in this section resembles legislation more than it does mathematics. But at least it gives us a language. For example we say that the morphism ϕ is *full* if it is full as a morphism of point-collinearity graphs. Thus if u and v are distinct points of $\phi(\mathcal{P})$ which are $\mathcal{L}_1$-collinear then some point of the fiber $\phi^{-1}(u)$ is $\mathcal{L}$-collinear with a point of the other fiber $\phi^{-1}(v)$ (see p. 15).

Similarly, $\phi : \Delta \to \Delta_1$ is is called a *fibering of point-line geometries* if ϕ is a fibering as a graph morphism. (Note that in this case no line of $\mathcal{L}$ is ever mapped to a single point of $\mathcal{P}_1$.)

But things work without so much definitional artifice in the reverse direction. Let $\Gamma = (\mathcal{P}, \mathcal{L})$ be a point-line geometry; let $\Delta = (\mathcal{P}, \sim)$ be its point-collinearity graph; and let $\mathcal{T}$ be the collection of all 3-circuits – that is, triangles – of the point-collinearity graph Δ. Now suppose

$$\kappa : \Delta' = (\mathcal{P}', \sim') \to \Delta = (\mathcal{P}, \sim)$$

is a $\mathcal{T}$ covering of Δ in the sense of Chap. 1. (We spare the reader from looking up the definitions by recalling that this means κ is a full vertex-surjective morphism between connected graphs Δ and Δ_1 which induces an isomorphism of induced neighborhood graphs at every vertex. In particular, the vertex-fibers partition the vertices of Δ into cocliques.)

The point is that the Δ' of the previous paragraph is not just a graph. It has a well defined line structure. Since κ is a $\mathcal{T}$-covering, every clique of the graph Δ_1 lifts to a unique clique at any point in the fiber of one of its points. Thus in a canonical way we obtain a special family $\mathcal{L}'$ of cliques in the graph Δ'. Then $\Gamma' := (\mathcal{P}', \mathcal{L}')$ is itself a point-line geometry, and κ becomes a fibering $\Gamma' \to \Gamma$ of point-line geometries – simply called a *covering of the point-line geometry* Γ.

Of course, drawing further on Chap. 1, we have a *universal $\mathcal{T}$-covering*

$$\tau : \hat{\Delta} \to \Delta$$

which can also be refined to a morphism of point-line geometries

$$\tau : \hat{\Gamma} = (\hat{\mathcal{P}}, \hat{\mathcal{L}}) \to \Gamma = (\mathcal{P}, \mathcal{L})$$

which we call the *universal cover of the point-line geometry* Γ.

3.4 Subspaces

We now come to the basic concept which forces an asymmetry between points and lines. We say that a subset S of $\mathcal{P}$ is a *subspace* if and only if for every line L, its point-shadow $\mathrm{sh}(L)$ meets S at either the empty set (then L is said to be *exterior to*

S), meets S at a single point (then L is said to be *tangent to* S), or the point-shadow of the line, $\mathrm{sh}(L)$ lies entirely within X (when L could be called an *internal line*). Since the point-line geometry $(\mathcal{P}, \mathcal{L})$ has no repeated lines, each line can be identified with its point-shadow. This allows a simpler formulation of the definition of subspace: we could say:

> *A set of points S is a subspace if and only if, for every line L, either $|L \cap S| \leq 1$ or $L \subseteq S$.*

A subspace S of a point-line geometry can be regarded as a point-line geometry $(S, I(S))$ in its own right, where $I(S)$ denotes the collection of internal lines, that is, lines whose point-shadows lie entirely inside S. It follows that any coclique C of the point-collinearity graph Δ is a subspace – indeed a subspace without lines (since we have required that all lines contain at least two points). Note that the extreme subsets of $\mathcal{P}$, namely the empty set and $\mathcal{P}$ itself, are both subspaces of Γ.

A subspace S of $(\mathcal{P}, \mathcal{L})$ is said to be *proper* if and only if S is a proper subset of $\mathcal{P}$.

The next observation is that the intersection of any collection of subspaces is a subspace. This should be clear from the definition of subspace.

Suppose $X \subseteq \mathcal{P}$ is an arbitrary subset of the point set $\mathcal{P}$. The intersection of all subspaces of $\Gamma = (\mathcal{P}, \mathcal{L})$ which contain X is called *the subspace of* Γ *generated by* X and is denoted by $\langle X \rangle_\Gamma$. Similarly, if S_1 and S_2 are subspaces, there is a unique smallest subspace $\langle S_1 \cup S_2 \rangle_\Gamma$ containing both of them. Thus the subspaces of a point-line geometry form a lattice. In general such a lattice is not a nice one – rarely distributive or even modular.

In certain cases it is useful to view the subspace generated by a set of points in a different way.

Let $X \subseteq \mathcal{P}$, where $\Gamma = (\mathcal{P}, \mathcal{L})$ is a point-line geometry. Similarly, let $Y \subseteq \mathcal{L}$ be any collection of lines. Let $\mathcal{L}(X)$ be the set of lines which are incident with at least two distinct points of X. (These lines are often called "secant lines." Note that when X is a subspace of Γ, then $\mathcal{L}(X)$ is just the set of internal lines $I(X)$ defined above.) Let $\mathcal{P}(Y)$ be the set of points incident with at least one of the lines of Y.[2] In fact, the definitions assert that X is a subspace if and only if $\mathcal{P}(\mathcal{L}(X)) \subseteq X$. We define a new operator "$\mathcal{PL}$" by the equation $(\mathcal{PL})(X) := X \cup \mathcal{P}(\mathcal{L}(X))$, for any set of points X.

Then

$$X, (\mathcal{PL})(X), (\mathcal{PL})((\mathcal{PL})(X)), \ldots, (\mathcal{PL})^{(n)}(X), \ldots$$

is an increasing sequence of sets of points whose union we denote as $\langle X \rangle_{\mathcal{P}}$. Similarly,

$$\mathcal{L}(X), \mathcal{L}(\mathcal{P}(\mathcal{L}(X))), \ldots \mathcal{L}((\mathcal{PL})^{(n)})(X), \ldots$$

[2] Note that, despite the notation, $\mathcal{P}(Y)$ is *not* the "dual" of the operator "$\mathcal{L}$" – the latter would be the set of points incident with at least two distinct lines of Y.

is an increasing sequence of lines whose union we denote by $\langle X\rangle_{\mathcal{L}}$. (Note that by definition $\mathcal{L}((\mathcal{PL})(X))$ contains $\mathcal{L}(X)$.) Then the following applies.

Lemma 3.4.1 *1. For any $X \subseteq \mathcal{P}$, $\langle X\rangle_{\mathcal{P}} = \langle X\rangle_{\Gamma}$, the subspace generated by X.*
2. The internal lines of the resulting subspaces are all lines constructed along the way: that is,

$$I(\langle X\rangle_{\Gamma}) = \langle X\rangle_{\mathcal{L}}.$$

Proof Part 1. Since $(\mathcal{PL})(\langle X\rangle_{\mathcal{P}}) = \langle X\rangle_{\mathcal{P}}$, the right side is a subspace of Γ containing X and so contains the intersection of all such subspaces, namely, $\langle X\rangle_{\Gamma}$.

On the other hand, if S is any subspace containing $(\mathcal{LP})^{(n)}(X)$ (including the case that $X = (\mathcal{PL})^{(0)}(X)$), then it contains $(\mathcal{PL})^{(n+1)}$ and so contains $\langle X\rangle_{\mathcal{P}}$. Since S is arbitrary, $\langle X\rangle_{\mathcal{P}}$ lies in the intersection of all such S, namely $\langle X\rangle_{\Gamma}$.

The proof of part 2 is left as Exercise 3.1 in the final Sect. 3.9 of this chapter. □

Remark We have noted that if S is already a subspace of $\Gamma = (\mathcal{P}, \mathcal{L})$, then $I(S) = \mathcal{L}(S)$. In that case, with two symbols to choose from, we tend to use the notation $\mathcal{L}(S)$ rather than $I(S)$.[3]

A subspace S of a point-line geometry Γ is said to be a *singular subspace* if and only if any two distinct points of S are collinear – that is, the collinearity graph of $(S, \mathcal{L}(S))$ is a clique in the full point-collinearity graph of Γ. Γ itself is called a *singular space* if and only if it is a singular subspace of itself – that is, the full point-collinearity graph Δ is a complete graph. Once again, we observe that the collection of all singular subspaces is closed under taking intersections.

There is still another way to attribute adjectives to subspaces, a way which depends only on how the point-collinearity graph of a subspace is embedded in the full point-collinearity graph $\Delta = (\mathcal{P}, \sim)$. Suppose S is a subspace. Since $\mathcal{L}(S) \subseteq \mathcal{L}$, two points of S are collinear in $\Sigma = (S, \mathcal{L}(S))$ if and only if they are collinear in Γ. Thus the following applies.

Lemma 3.4.2 *Suppose S is a subspace of the point-line geometry $\Gamma = (\mathcal{P}, \mathcal{L})$. Then the point-collinearity graph $\Delta(\Sigma)$ of the geometry $\Sigma = (S, \mathcal{L}(S))$ is an induced subgraph of the point-collinearity graph Δ of Γ.*

We say that a subspace S is a *convex subspace* if and only if its induced point-collinearity subgraph is convex. This means that if $(x_0, x_1, \ldots, x_n)$ is a shortest possible sequence of points with x_i collinear with x_{i+1} connecting two given points x_0 and x_n of S, then all the intervening points x_i also belong to S. Note that all singular subspaces are convex, according to this definition.

From Lemma 1.1.5 in Chap. 1, the following is derived.

[3] The reasons for this choice have to do with competing meanings demanding the use of the symbol "I." Besides "internal," there are notions like "incidence," "index sets of types," etc.

Lemma 3.4.3 *Convex subspaces are closed under the intersection operation.*

Convex subspaces have an important property.

Lemma 3.4.4 *Suppose S is a convex subspace of the point-line geometry $\Gamma = (\mathcal{P}, \mathcal{L})$. If $x \in \mathcal{P} - S$, then any two points of $x^\perp \cap S$ are collinear. In particular, if $x^\perp$ is a subspace of Γ, then $x^\perp \cap S$ is a singular subspace of Γ.*

Proof Noting that Lemma 3.4.2 makes S an induced subgraph of the point-collinearity graph $\Delta = (\mathcal{P}, \sim)$, the result follows immediately from Lemma 1.1.7, p. 10, of Chap. 1. □

An even stronger property relating a subspace to its ambient space is the following. We shall say that a *subspace S is strongly gated* if and only if the subgraph of the point-collinearity graph induced on the vertex set S is strongly gated. Of course, since we are talking about induced subgraphs, the subgraph is strongly gated if and only if it is gated. Thus, equivalently, we can say that a *subspace S is gated* if and only if it is strongly gated. Again, Theorem 1.1.3 of Chap. 1 shows us that intersections of strongly gated subspaces are strongly gated. This includes the empty subspace.

The hypothesis that a class of subspaces is gated is quite strong. We shall later meet *near polygons* for which lines are gated, and *dual polar spaces* which contain a class of convex generalized quadrangles which are gated.

3.4.1 Generalizations of the Notion of Subspace

In this brief section, we consider two generalizations of the notion of subspace. We have seen that a subspace S of a point-line geometry $\Gamma = (\mathcal{P}, \mathcal{L})$, is a subset S of the point set which meets any line at 0, 1 or all of its points, and that we may regard S as the points of a point-line geometry $\Sigma = (S, \mathcal{L}(S))$ whose collinearity graph $\Delta(\Sigma)$ is an induced subgraph of $\Delta(\Gamma)$, the point-collinearity graph of Γ.

On the other hand, suppose $\mathcal{X} \subseteq \mathcal{L}$ and let $\mathcal{P}(\mathcal{X})$ be the set of all points incident with at least one line in $\mathcal{X}$ – that is, the union of the lines of $\mathcal{X}$ regarded as point sets. Then, by restriction of the incidence relation of Γ, we obtain a point-line geometry $S = (\mathcal{P}(\mathcal{X}), \mathcal{X})$ which we call *an embedded subspace of* Γ.[4] Like ordinary subspaces, the shadow of each line of the embedded subspace is the full point-shadow of the line, as a line of Γ. The main difference is that the point-collinearity graph of an embedded subspace need not be an induced subgraph of $\Delta(\Gamma)$.

The concept of embedded subspace is very important, especially when Γ is a classical projective space (see Chap. 7).

Finally we should mention point-line geometries $(\mathcal{P}_0, \mathcal{L}_0)$ where $\mathcal{P}_0$ and $\mathcal{L}_0$ are respectively sets of points and lines of $\Gamma = (\mathcal{P}, \mathcal{L})$, such that for each line $L \in \mathcal{L}_0$,

[4] Note that in $(\mathcal{P}(\mathcal{X}), \mathcal{X})$, it is quite possible for a point to lie in a unique line of $\mathcal{X}$. This was not forbidden in our definition of point-line geometry.

$|L \cap \mathcal{P}_0| \geq 2$. We call such geometries *partial subspaces of* Γ. In general, the point shadow of any of its lines, might be smaller than the point shadow of that same line in Γ, and again, its point-collinearity graph might not be an induced subgraph of $\Delta(\Gamma)$. However, if $\mathcal{L}_0 := \{L \in \mathcal{L} || L \cap \mathcal{P}_0| \geq 2\}$ – that is *all* lines of $\mathcal{L}$ which intersect the set $\mathcal{P}_0$ in at least two points are used – then the point-collinearity graph *is* an induced subgraph and $(\mathcal{P}_0, \mathcal{L}_0)$. Such a *geometry of secant lines of* $\mathcal{P}_0$ is also called a *lax embedding* (see [132]). (For those familiar with the notions, a Baer subplane of a projective plane would be an example of a partial subspace which is also a geometry of secant lines.)

The generalizations of this section are arranged in this way:

$$\text{subspaces} \subseteq \text{embedded subspaces} \subseteq \text{partial subspaces}.$$

3.5 Special Types of Point-Line Geometries

3.5.1 Partial Linear Spaces and Linear Spaces

A point-line geometry is called a *partial linear space* if and only if any two distinct points are collinear with at most one line. A partial linear space which is a singular space is simply called a *linear space*. Thus for a linear space, every pair of distinct points are together incident with exactly one line in common.

Let's review these concepts in our examples:

- The geometry of vertices and edges of the octohedron geometry is a partial linear space with thin lines, but is not a linear space.
- The points and lines of Desargues' configuration is a partial linear space which is not a singular space.
- The vertices and edges of the complete graph K_n is a linear space (V, E) with thin lines.
- Let $\Gamma = (\mathcal{P}, \mathcal{L})$ be dual of the point-line geometry of Example 3 of the previous chapter, p. 48. Thus the points of the geometry are the $\mathbf{GF}(q)$ translates $Q + x$ of the set Q of non-zero quadratic residues of the field $\mathbf{GF}(q)$, $q \equiv 3 \bmod 4$ and lines are the field elements. Then for $q > 7$, Γ is a singular space which is not a linear space. If $q = 7$, Γ is a particular linear space called a "projective plane of order 2".

3.5.2 Gamma Spaces

A point-line geometry $\Gamma = (\mathcal{P}, \mathcal{L})$ is called a *gamma space* if and only if, for every line L and point p we have either:

(i) $p^\perp \cap sh(L) = \emptyset$
(ii) $|p^\perp \cap sh(L)| = 1$, or
(iii) $L \subseteq p^\perp$

That means that if p is collinear with at least two distinct points of the line L, then it is collinear with all of the points in the point-shadow of L.

Note that if p is incident with L, then $p^\perp \supseteq L$ (since p is now collinear with itself) so conclusion (iii) holds. All of the Lie incidence geometries of this book are gamma spaces.[5]

Theorem 3.5.1 *The following statements are equivalent for a point-line geometry* $\Gamma = (\mathcal{P}, \mathcal{L})$:

1. Γ *is a gamma space*
2. *For every point* p, *the set* $p^\perp$ *is a subspace of* Γ
3. *For any clique* X *of the point-collinearity graph of* Γ, *the subspace that it generates,* $\langle X \rangle_\Gamma$, *is a singular subspace of* Γ

Proof That 1 implies 2 is trivial. That 3 implies 1 is fairly trivial. Consider, if 1 failed, there would exist a thick line L and a point p collinear with at least two distinct points a and b of L, while not being collinear with a third point c of L. Then L and p lie in the subspace S generated by the 3-clique $\{a, b, p\}$. By assumption 3, S is a clique. But that is impossible since it contains the non-collinear pair of points $\{p, c\}$.

So the problem is whether 2 implies 3. Suppose X is a clique as hypothesized. It is sufficient to prove that if Z is a clique of points, then so is $(\mathcal{PL})(Z)$. Any point x of the latter set lies in some line L_x which meets Z in at least two distinct points – say a and b. Thus if y is any further point of $(\mathcal{PL})(Z)$, then there are two distinct points y_i of $L_y \cap Z$ each of which is collinear with the two points a and b (since Z is a clique) and so are collinear with x by part 2. Similarly, as x is now collinear with the two points y_i, the statement of part 2 implies that x is collinear with y. Since x and y were arbitrary points of $(\mathcal{PL})(Z)$, the latter must be a clique. The proof is complete. □

Remark Note the argument that 2 implies 3 would not be easy to deduce if we did not have a second "atomic" characterization $\langle X \rangle_\mathcal{P}$ of the subspace $\langle X \rangle_\Gamma$ generated by a set of points (see Lemma 3.4.1).

3.6 Local Connectedness in Gamma Spaces

Now suppose $\Gamma = (\mathcal{P}, \mathcal{L})$ is a gamma space. Recall that this property asserts that $p^\perp$ is a subspace of Γ for each point $p \in \mathcal{P}$. For any line L, $L^\perp$ denotes the set of points which are either incident with L or are otherwise collinear with every point of L. For each point p let $\mathcal{L}_p$ be the collection of all lines of $\mathcal{L}$ which are incident with point p. By the gamma space hypothesis, if L_1, L_2 are distinct lines in $\mathcal{L}_p$ there are two possibilities: (1) $L_1{}^\perp \cap L_2 = L_2{}^\perp \cap L_1 = \{p\}$ and (2) $L_1 \subseteq L_2{}^\perp$ and

[5] That is not to say that there are not very interesting geometries (such as partial geometries or $(0, \alpha)$-geometries) which happen not to be gamma spaces.

$L_2 \subseteq L_1{}^{\perp}$. Let us write $L_1 \sim L_2$ if and only if (2) holds. Then at each point p we obtain a simple graph $\Lambda_p := (\mathcal{L}_p, \sim)$, which we call the *tangent graph at p*. We say that the gamma space $\Gamma = (\mathcal{P}, \mathcal{L})$ is *locally connected*, if and only each tangent graph Λ_p is connected.

3.6.1 The Decomposition of a Gamma Space into Locally Connected Components

Again, let $\Gamma = (\mathcal{P}, \mathcal{L})$ be a gamma space. Let $\hat{\mathcal{P}}$ be the set of all pairs (p, C) where C is a connected component of the graph Λ_p. We can now define a line structure on $\hat{\mathcal{P}}$ as follows. Select a line $L \in \mathcal{L}$ with point-shadow $sh(L)$. Define the following subset of $\hat{\mathcal{P}}$:

$$\hat{L} := \{(p, C_p) \in \hat{\mathcal{P}} | p \in sh(L), L \in C_p\}.$$

(Note that in this definition, the line L is fixed, and, as p ranges over $sh(L)$, C_p is determined by p and the fixed L.) Next set $\hat{\mathcal{L}} := \{\hat{L} | L \in \mathcal{L}\}$.

Now observe that if $\hat{L_1} = \hat{L_2}$ for lines L_1, L_2 of $\mathcal{L}$, then $L_1 = L_2$. This is because the first equation forces $sh(L_1) = sh(L_2)$ and this forces $L_1 = L_2$ since any gamma space is by definition a point-line geometry and so does not possess repeated lines. Thus the mapping defined by $L \to \hat{L}$ induces a bijection $\mathcal{L} \to \hat{\mathcal{L}}$.

Then we have a new geometry $\hat{\Gamma} := (\hat{\mathcal{P}}, \hat{\mathcal{L}})$. Note that in this new geometry, a point (p, C) is collinear with a point (q, D) in $\hat{\Gamma}$ if and only if p and q are distinct points of $\mathcal{P}$ and there exists at least one line $L \in \mathcal{L}_p \cap \mathcal{L}_q$ such that $L \in C \cap D$. From this one may verify that $\hat{G}$ is also a gamma space.

Now we define two mappings: (1) $f : \hat{P} \to \mathcal{P}$ which sends (p, C) to the point p, and (2) $\ell : \hat{\mathcal{L}} \to \mathcal{L}$ which sends line $\hat{L}$ to its unique L. Then as $(p, C) \in \hat{L}$ implies $p \in sh(L)$, we see that $(f, \ell) : \hat{\mathcal{P}} \times \hat{\mathcal{L}} \to \mathcal{P} \times \mathcal{L}$ is a morphism of geometries which induces a bijection $\hat{L} \to L$ when restricted to any line $\hat{L}$.

Now let $\hat{\mathcal{P}} = \sum_{\sigma \in X} \hat{\mathcal{P}}_\sigma$ be the partition of the vertices of the point-collinearity graph $\hat{\Delta}$ of $\hat{\Gamma}$ into connected components. Then the point-shadow of each line $\hat{L}$ lies in a unique component – say $\hat{\mathcal{P}}_\sigma$. Thus the geometry $\hat{\Gamma}$ is a disjoint union of connected geometries $\hat{\Gamma}_\sigma = (\hat{\mathcal{P}}_\sigma, \hat{\mathcal{L}}_\sigma)$.

From the definition of C as a connected component of the graph Λ_p, for two lines L and N of C, there is sequence of lines $(L = L_0, L_1, \ldots, L_k = N)$ with $L_i \subseteq L_{i+1}{}^{\perp}$. It is then easy to show that lines $\hat{L_i}$ and $\hat{L_{i+1}}$ lie together in a larger singular subspace of $\hat{\Gamma}$, and so lines $\hat{L}$ and $\hat{N}$ are connected in the tangent $\hat{\Lambda}_{(p,C)}$ of $\hat{\Gamma}$, since any two lines $\hat{L}$ and $\hat{N}$ on point $(p, C) \in \hat{\mathcal{P}}$ come from lines L and N on p in the same connected component C of Λ_p.

So we see that two lines L and N on point p lie in the same connected component of Λ_p if and only if their lifts $\hat{L}$ and $\hat{N}$ meet at a common point $(p, C) \in \hat{\mathcal{P}}$ and so lie together in a unique connected (and automatically locally-connected) component of $\hat{\Gamma}$.

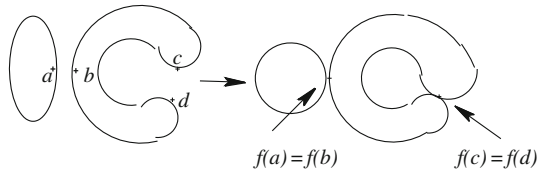

Fig. 3.1 An example illustrating how the canonical morphism $\phi : \hat{\Gamma} \to \Gamma$ pastes together connected components of $\hat{\Gamma}$ to produce a geometry which is not locally connected. Note that two points of the same component can become identified if they are not collinear

Thus we have the following.

Theorem 3.6.1 *Let $\Gamma = (\mathcal{P}, \mathcal{L})$ be a gamma space, and construct the covering morphism $\phi = (f, \ell) : \hat{\Gamma} = (\hat{\mathcal{P}}, \hat{\mathcal{L}}) \to (\mathcal{P}, \mathcal{L})$:*

1. *Each connected component $\hat{\Gamma}_\sigma$ of the geometry $\hat{\Gamma} = (\hat{\mathcal{P}}, \hat{\mathcal{L}})$ is a locally connected gamma space.*
2. *The mapping $\ell : \hat{\mathcal{L}} \to \mathcal{L}$ is a bijection.*
3. *The point-fibers $f^{-1}(p)$ are sets of pairwise non-collinear points of $\hat{\mathcal{P}}$.*
4. *The images $S_\sigma = f(\hat{\mathcal{P}}_\sigma)$ are a canonically determined set of subspaces of Γ, each of which is maximal with respect to the property of being a locally connected subspace.* (We call the S_σ the *locally connected components*.)

Corollary 3.6.2 (Uniqueness of the locally-connected components.) *Every gamma space possesses a unique system of connected locally-connected subspaces with the property that each line is contained in a unique member of the system.*

In effect, the original gamma space Γ is constructed by identifying isolated points from a disjoint union of connected locally connected geometries. This is illustrated in Fig. 3.1.

Theorem 3.6.1 seems to have been rediscovered many times, and should probably be attributed to folklore.[6]

3.6.2 How Local Characterizations of Gamma Spaces Reduce to the Locally Connected Case

Suppose one has a hypothesis (P) which declares that above each point p of a point-line geometry, there exist certain relationships among hypothesized subspaces on

[6] The unpublished Theorem 3.6.1 and the corollary of this section on local connectedness were proved independently by Peter Johnson, J.I. Hall, Rolf Gramlich and certainly others, including the author, though the results are so simple one wonders why one even mentions this fact. On the other hand, one doesn't want to make the important obvious theorems unmentionable on the grounds that they shame the intellectual reputation of those who utter them. A lot of important stuff can be forgotten that way.

a point p. Often such hypotheses can be translated into what I would call "purely local hypotheses" involving only the lines and singular subspaces of a gamma space incident with any given point.

In some characterization theorems, local connectedness is hypothesized outright.[7] Sometimes these local hypotheses – as we have called them – imply that the geometry is locally connected in a more indirect way. But occasionally such a local hypothesis does not imply local connectedness at all. The student can test this by seeing whether two locally connected geometries with property (P), which have been pasted together at a point, still have property (P).

In view of Theorem 3.6.1 we have these principles:

1. Any characterization theorem based on local hypotheses that do not include the hypothesis of being locally connected should have a conclusion set which is closed under the pasting construction described above in this section
2. Any proposed theorem which fails in this way – while literally untrue as stated – can be easily rectified by changing the conclusion "Γ satisfies C" (which implicity assumed local connectedness) to read "the locally connected components of Γ satisfy C"
3. It follows that one can always reduce a characterization theorem to the case that the geometry is locally connected if one is willing to alter the conclusion as in the previous item

3.7 Enriching Geometries

This section concerns the following: Given a point-line geometry, $\Gamma = (\mathcal{P}, \mathcal{L})$, we can sometimes produce a geometry of higher rank, say

$$\Gamma' := (\mathcal{P}, \mathcal{L}, S_3, S_4, \ldots, S_k),$$

where the S_i are certain classes of subspaces of Γ. In this way we can produce a geometry of rank k which we call an *enrichment of* Γ from our original point-line geometry $\Gamma = (\mathcal{P}, \mathcal{L})$ of rank two.[8] This process of enrichment is pretty much the reverse of the process of truncation introduced in the previous chapter. Indeed, with the notation just introduced, we have that the truncation to points and lines of Γ' presented above is the original point-line geometry.

Many examples of enrichment are actually provided by special hypotheses on the points and lines. Here are some examples that will appear in this book:

1. For singular subspaces of gamma spaces, perhaps the following two examples are the most important:

[7] See for example the theorems in [68, 72].

[8] This very useful term "enrichment" is due to Pasini (see [95]).

(a) As we have defined the term so far, a linear space is called a *projective space* if and only if the subspaces generated by two intersecting lines is a projective plane with thick lines.[9] If lines are not maximal singular subspaces, new (necessarily singular) subspaces exist and each is also a projective space of a smaller dimension. Thus, looking only at those singular subspaces of finite dimension, we obtain an enriched geometry Γ where S_d is the set of subspaces of projective dimension $d-1$. Projective spaces are studied in Chap. 5.

(b) Polar spaces are partial linear spaces with a special axiom even simpler than Veblen's axiom. If lines are not maximal singular subspaces, a polar space contains a system of singular subspaces – namely the singular subspaces of finite rank – with respect to which it becomes an enriched geometry of rank at least three. These are studied in Chap. 7.

2. A simple axiom on near polygons produces a family of "quads" $\mathcal{Q}$ (which are convex subspaces which are generalized quadrangles), thus producing an enriched geometry $\Gamma' := (\mathcal{P}, \mathcal{L}, \mathcal{Q})$ of rank three. If the quads are strongly gated subspaces (defined earlier in this chapter), a classical theorem of Cameron shows that they are derived from polar spaces. This theory is developed in Chap. 9.
3. Quite a number of Lie incidence geometries obtained from buildings satisfy a point-line hypothesis defining a "parapolar space." These examples all possess a family Σ called a "*polar family*" consisting of point-pairs $\{x, y\}$ which are at mutual distance 2 in the point-collinearity graph, possessing two properties. An argument called "Cooperstein's theory" proves that the convex closure of any such pair is a convex subspace called a "*symplecton*." At this stage, the geometry has been enriched by the addition of the symplecta and any singular subspace properly containing a line which lies in such a symplecton. (Cooperstein's theory – adapted to polar families – is given in detail in Chap. 13.)

So one can see that enrichment of a geometry plays a fundamental role in classifying the original point-line geometry.

3.8 Products of Point-Line Geometries, a Construction

Product geometries are important since they arise in many contexts, for example, as point-residuals of Grassmann spaces.

Suppose $\Gamma_i = (\mathcal{P}_i, \mathcal{L}_i)$, $i = 1, 2$, are two point-line geometries. We wish to construct a new geometry $\Gamma_1 \times \Gamma_2 = (\mathcal{P}, \mathcal{L})$ with point set $\mathcal{P} = \mathcal{P}_1 \times \mathcal{P}_2$ and line set $\mathcal{L} = (\mathcal{L}_1 \times \mathcal{P}_2) \cup (\mathcal{P}_1 \times \mathcal{L}_2)$.

A point $p = (a, b) \in \mathcal{P}_1 \times \mathcal{P}_2$ is incident with a line $L = (L_1, p_2) \in \mathcal{L}_1 \times \mathcal{P}_2$ if and only if $b = p_2$ and a is incident with L_1 in Γ_1. Similarly, $p = (a, b)$ is incident

[9] Later, on p. 89, Lemma 4.3.1 will tell us that a linear space is a projective space if and only if it satisfies Veblen's axiom. This is the standard definition.

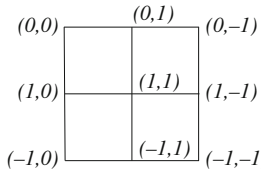

Fig. 3.2 The 3-by-3 grid as the product of two lines

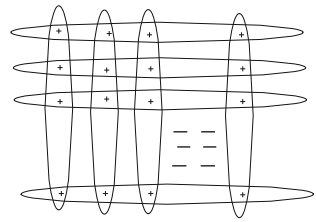

Fig. 3.3 A general picture of a product geometry

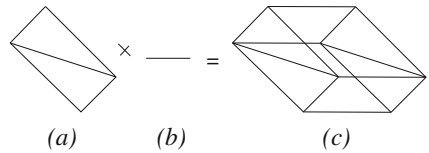

Fig. 3.4 The depiction in (**c**) of a simple product of the geometries in (**a**) and (**b**)

with a line $N = (p_1, N_2) \in \mathcal{P}_1 \times \mathcal{L}_2$ if and only if $a = p_1$ and b is incident with N_2 in the geometry Γ_2.

Let us look at a very simple example. Suppose $\Gamma_1 = \Gamma_2 = \Gamma$ consists of just three points $\mathcal{P}_i = \mathcal{P} = \{0, 1, -1\}$ – that is, the elements of $\mathbf{Z}/(3)$ – and there is just one line L which is incident with all three points. Then $\Gamma_1 \times \Gamma_2$ is the 3-*by*-3 *grid*. That is, it has the nine points of $\mathcal{P} \times \mathcal{P}$ and the six lines (L, a) or (a, L), as a ranges over $\mathcal{P}$ with incidence depicted as in Fig. 3.2.

In general, $\mathcal{L}$ consists of two classes of lines: the "vertical" lines of the form $L \times \{p\} \in \mathcal{L}_1 \times \mathcal{P}_2$, and the "horizontal" lines of the form $\{r\} \times N \in \mathcal{P}_1 \times \mathcal{L}_2$. Thus the points of $\Gamma_1 \times \Gamma_2$ are partitioned into "vertical" subspaces, all isomorphic to Γ_1, as well as "horizontal" subspaces isomorphic to Γ_2, as crudely depicted in Fig. 3.3.

For example, suppose Γ_1 is the thin geometry of vertices and edges of K_4 minus one edge (the graph of Fig. 3.4a). Suppose also that Γ_2 is just two vertices and an edge (Fig. 3.4b). Then the product geometry is the graph in Fig. 3.4c, under the convention that vertices are points and edges are lines.

3.9 Exercises and Examples for Chapter 3

There are several ways to form a point-line geometry from a geometry of rank $n > 2$.

First, one can simply set $(\mathcal{P}, \mathcal{L}) = (V_i, V_j)$, the truncation of a geometry $(V_1, \ldots V_n; *)$ to a 2-set $\{i, j\}$ of the type set $I = \{1, \ldots, n\}$. Here are three examples derived from a classical projective geometry:

Fix an $n+1$-dimensional vector space V over a division ring D and let Γ be the classical projective geometry $PG(V)$ of all proper subspaces of V. Then Γ is a geometry of rank n. (Except for our first example, n is finite.)

Example 1 (A Classical Projective Space.) Let $\mathcal{P}$ be all one-dimensional subspaces of V and let $\mathcal{L}$ be all two-dimensional subspaces of the same space. Then with the usual incidence, $(\mathcal{P}, \mathcal{L})$ is a well defined point-line geometry. This geometry is called a *classical projective space* (or sometimes, a *Desarguesian projective space*), even when V has infinite dimension.

Example 2 Suppose $\mathcal{P}$ is again all 1-spaces of the $n+1$-space V, but now let $\mathcal{L}$ be the set of all n-subspaces of V (these are called the *projective hyperplanes of* $PG(V)$). Then $(\mathcal{P}, \mathcal{L})$ is a nice point-line geometry. If D is a finite field $\mathbf{GF}(q)$, then

$$|\mathcal{P}| = |\mathcal{L}| = 1 + q + \cdots + q^n.$$

Example 3 Suppose $\dim V = 5$. Let $\mathcal{P}$ and $\mathcal{L}$ be respectively, the sets of all 2- and 3-subspaces of V, with the usual incidence.[10]

3.1 Prove part 2 of Lemma 3.4.1.

3.2 In the above three examples:

1. Which are singular spaces?
2. Which are linear spaces? (Recall that as $n > 2$, $\dim V \geq 4$.)
3. Which are gamma spaces?

3.3 Let W be any k-dimensional subspace of V, where $n \geq k \geq 3$. For $\ell \leq k$, let $P_\ell(W)$ be the collection of all ℓ-subspaces of V which are contained in the subspace W.

1. In Example 1, $\mathcal{P}_1$ is a subset of the set of points $\mathcal{P} = \mathcal{P}_1(V)$, the "projective points." Show that $\mathcal{P}_1(W)$ is a subspace of the projective space $(\mathcal{P}, \mathcal{L})$.
2. In Example 3, suppose W is a 4-subspace of the 5-subspace V. Now $\mathcal{P}_2(W)$, by our definition, is the collection of all 2-subspaces of V which lie in W. Therefore it is a subset of the set of points of Example 3. Show that it is a subspace of that space.

Example 4 (Attenuated spaces.) Let W be any k-dimensional subspace of the vector space V, where $n \geq k \geq 3$. Let $\mathcal{P}^W$ and $\mathcal{L}^W$ be the collections of all 1- and 2-subspaces of V which intersect the subspace W at the 0-subspace. Then the point-line geometry $A(V : W) = (\mathcal{P}^W, \mathcal{L}^W)$ is called the *attenuated space of V with kernel W*. It is an embedded subspace of the projective space $\mathbf{P}(V)$, but is not a subspace of it.

[10] I know of no special name for this geometry. In Chap. 14, we will call such spaces "Half-Grassmann spaces."

Example 5 (Some classical polar spaces over a field.) Let V be an n-dimensional vector space over a field F, and let $f : V \times V \to F$ be a *reflexive sesquilinear form*. That means, for any vectors $a, b, c \in V$, and scalar α in F, there is an involution σ of Aut(F) such that

$$f(a+b,c) = f(a,c) + f(b,c),$$
$$f(a,b+c) = f(a,b) + f(a,c),$$
$$f(\alpha a, c) = \alpha f(a,c),$$
$$f(a,\alpha c) = f(a,c)\alpha^{\sigma},$$
$$f(a,b) = 0 \text{ if and only if } f(b,a) = 0.$$

We say that *vector a is perpendicular to a vector b with respect to a reflexive sesquilinear form f* if and only if $f(a,b) = 0$, or equivalently, $f(b,a) = 0$. For any set of vectors X, we let $X^{\perp} = \{v \in V | f(v,x) = 0, \text{for all } x \in X\}$, the set of vectors simultaneously perpendicular to every vector in set X. Thus

$$X^{\perp} = \cap_{x \in X} x^{\perp}.$$

A subspace A of V is said to be *totally singular with respect to f* if and only if $f(A,A) = \{0\}$ – that is, $A \subseteq A^{\perp}$.

The set Rad $(V) := \{a \in V | f(a,v) = 0 \text{ for all } v \in V\}$ is a vector subspace of V called the *radical of the form f*. We say the form f is *non-degenerate* if and only if Rad $(V) = \langle 0 \rangle$.

Now suppose f is a non-degenerate sesquilinear form on a finite-dimensional vector space V. It is an easy theorem that all maximal totally singular subspaces of V possess the same dimension, called the *Witt index of the form f*, and that this index must not exceed one-half the dimension of V. Let A_i be the collection of all totally singular subspaces of dimension exactly i. Then $\Gamma = (A_1, \ldots, A_r; \subset \text{ or } \supset)$, is a geometry of rank r which we call *the polar geometry of* (V, f). Its truncation to $J = \{1, 2\}$ is a point-line geometry called *the classical polar space associated with* (V, f). The truncation to $K = \{r, r-1\}$ (in that order) is a point-line geometry called a *classical dual polar space*.

3.4 Show the following:

1. For any non-incident point-line pair $(p, L) \in \mathcal{P} \times \mathcal{L}$, of a polar space $(\mathcal{P}, \mathcal{L})$, p is collinear to exactly one, or to all of the points of L.
2. For such a non-incident pair (p, L) in a dual polar space, p is collinear with at most one point of L.
3. In the case the Witt index r is two, the polar spaces and dual polar spaces are point-line geometries which are duals of each other. In this case, for each non-incident point-line pair (p, L), p is collinear with exactly one point of L (these are examples of *classical generalized quadrangles*).

A second way to form a point-line geometry from a geometry is the following. Let $I = \tau(V)$ be the collection of all realized types of a geometry Γ. For each subset J of I let $\mathcal{F}_J$ be the collection of all flags F of type J – that is, $\tau(F) = J$. (If Γ is residually connected, then $\mathcal{F}_J$ is non-empty if J is.) Now let J and K be two distinct subsets of I. We can now define a point-line geometry $(\mathcal{P}, \mathcal{L}) = (\mathcal{F}_J, \mathcal{F}_K)$ by asserting that a pair of flags $(F, G) \in \mathcal{F}_J \times \mathcal{F}_K$ are *incident* if and only if $F \cup G$ is itself a flag of type $J \cup K$. (Note that incidence means that the union of the two cliques F and G is also a clique – equivalently $F \subseteq G^\perp$ or $G \subseteq F^\perp$.)

Note that this construction generalizes that of rank-two truncation. If $J = \{i\}$ and $K = \{j\}$, then $\mathcal{F}_J$ is just all objects of type i, and $\mathcal{F}_K$ is just all objects of type j, and the geometry $(\mathcal{F}_J, \mathcal{F}_K)$ is just the truncation $\Gamma_{\{i,j\}}$ in this case. So we get something new when at least one of the sets J or K has cardinality exceeding one.

Example 6 (The Grassmann spaces $A_{n,k}(D)$.) Again let $\Gamma = (V_1, \ldots, V_n; *)$ be the geometry $PG(V)$ of all proper subspaces of an $(n+1)$-dimension vector space V. Fix an integer k with $1 < k < n$. A $(k-1, k+1)$-*flag* is a pair of incident subspaces (A, C) (that is, $A < C$) with $\dim A = k-1$ and $\dim B = k+1$ – that is a flag of type $K := \{k-1, k+1\}$ in the projective geometry $PG(V)$. Let $\mathcal{P} = \mathcal{P}_k(V)$, the set of all k-subspaces of V and let $\mathcal{L}$ be the set $\mathcal{F}_{\{k-1,k+1\}}$ of all $(k-1, k+1)$-flags of $PG(V)$. We say that a "point" B is *incident* with a "line" (A, C), a $(k-1, k+1)$-flag, if and only if all these spaces form a chain $A < B < C$. That means, of course, that $\{A, B, C\}$ is a flag of Γ of type $\{k-1, k, k+1\}$ so this follows the construction above. Such a point line geometry $(\mathcal{P}, \mathcal{L})$ is called a *Grassmann space*[11] or *Grassmannian*. Since the vector space V is determined by the parameters n and D, and our geometry by k, this Grassmannian is sometimes denoted by $A_{n,k}(D)$.

3.5 Let $(\mathcal{P}, \mathcal{L})$ be the Grassmann space $A_{n,k}$ of Example 6. Suppose $1 < k < n$:

1. Show that $(\mathcal{P}, \mathcal{L})$ is a partial linear space
2. Show that $(\mathcal{P}, \mathcal{L})$ is a gamma space
3. Show that if A and B are points (that is, k-subspaces of V), then the distance between A and B in the point-collinearity graph Δ is the dimension of the factor space $A/(A \cap B)$ – i.e., the *codimension* of $A \cap B$ in A

Now let's try something different.

Example 7 Suppose Ω is a presented set of six letters $\omega = \{0, 1, \ldots 5\}$. Let $\mathcal{P}$ be the 15 2-subsets of Ω and let $\mathcal{L}$ be the 15 partitions of Ω into three subsets of size 2. We say that a 2-subset is incident with one of the partitions if and only if the unordered pair is a component of the partition. In this case each 'line' is incident with exactly three points and each point is incident with exactly three lines.

Example 8 Another variation on the same theme. Now suppose $\Omega = \{0, 1, \ldots, 8\}$, a set of nine letters. Our points $\mathcal{P}$ will be the set of all $\binom{9}{3} = 84$ 3-subsets (which we

[11] In honor of the German mathematician and high-school teacher who nearly invented them.

call 'triplets') chosen from Ω. The lines are the 280 partitions of Ω into three 3-sets. As before, a point is incident with a line if and only if it is one of the component 3-sets of the partition.

3.6 Show the following:

1. The point-line geometries of either Example 6 or 7 are partial linear spaces: moreover, in them, no point is collinear with all remaining points.
2. In Example 7, for each point p not incident with a line L, p is collinear with exactly one point of L. (This means Example 7 is a generalized quadrangle of order (2, 2).)
3. The point-line geometry of Example 7 is a gamma space.
4. Fix a point $p = \{a, b, c\}$, one of the triplets in Example 8. Let X be the set of 45 triplets which meet p at a single letter of Ω. Then X is a subspace of $(\mathcal{P}, \mathcal{L})$.

The next few problems involve generation.

3.7 Show that the generalized quadrangle of order 2 in Example 7 can be generated by five of its points; but cannot be generated by four points.

3.8 Suppose Γ_1 and Γ_2 are point-disjoint point-line geometries which can be generated by g_1 and g_2 points respectively. Show that the product geometry $\Gamma_1 \times \Gamma_2$ can be generated by $g_1 g_2$ points. If g_1 and g_2 are *minimal* numbers of generators, can $\Gamma_1 \times \Gamma_2$ be generated by fewer than $g_1 g_2$ generators?

On p. 65 we introduced the concept of an *embedded subspace* as something more general than a subspace. This notion can be transported to a categorical context in the following way. Let $\Gamma_i = (\mathcal{P}_i, \mathcal{L}_i)$, $i = 1, 2$ be two point-line geometries. Suppose $e : \Gamma_1 \to \Gamma_2$ is a morphism of these point-line geometries and the property that line shadows are bijectively mapped to line shadows of their image lines. That means there are really two mappings: $e_1 : \mathcal{P}_1 \to \mathcal{P}_2$ on points and $e_2 : \mathcal{L}_1 \to \mathcal{L}_2$ on lines preserving incidence. We say the morphism e is *point-injective (surjective)* if and only if e_1 is injective (surjective). *Line-injective* and *line-surjective* have similar meanings with e_2 replacing e_1.

An *embedding* $e : \Gamma_1 \to \Gamma_2$ is a morphism of point-line geometries with these properties:

1. e is point-injective
2. For every line $L \in \mathcal{L}_1$, we have

$$e_1(sh(L)) = sh(e_2(L))$$

That is, the images of all points incident with a line L comprise the full point-shadow (in Γ_2) of the image of that line

3. The set of image points spans Γ_2, – that is,

$$\langle e_1(\mathcal{P}_1) \rangle_{\Gamma_2} = \mathcal{P}_2.$$

Then the image geometry $e(G_1) := (e_1(\mathcal{P}_1), e_2(\mathcal{L}_1))$ is an embedded subspace of Γ_2 in the sense defined on p. 65. It need not be an actual subspace.

Example 9 Form the vector-space $V = (\mathbf{GF}(2))^{(2n)}$, $n \geq 2$ and let $f : V \times V \to \mathbf{GF}(2)$ be the ordinary "dot" product. Then f is a bilinear form on V (that is, it obeys the axioms of a reflexive sequilinear form in Example 5 above, with the $\sigma = 1_F$, the identity automorphism of F). The *weight* of a vector v is the number of "1"s that appear in the n-tuple expression for V.

Let $j = (1, 1, \ldots, 1)$, the unique vector of weight $2n$. The inner product $f(j, v)$ simply records the weight of vector v taken mod 2. It follows that the collection of all even weight vectors forms a hyperlane $j^\perp$ of V. Then, as $f(j, x) = 0$ for all $x \in j^\perp$, we see that

$$s(x, y) := f(x + \langle j \rangle, y + \langle j \rangle) = f(x, y)$$

for all $x, y \in j^\perp \times j^\perp$. From this formula f induces another bilinear form on the factor space $W = j^\perp / \langle j \rangle$. Now, it is clear that this form obeys the rule

$$s(w, w) = 0, \text{ for all} w \in W. \tag{3.1}$$

Such a form is called an *alternating* (or *symplectic*) form.

Let $\Omega := \{1, 2, \ldots, 2n\}$. Now any vector w can be identified with a pair consisting of a subset $\Omega(w)$ of Ω and its complement in Ω, of even weights k and $2n - k$, respectively.

3.9 Show the following:

1. $s(u, w) = 0$ if and only if $|\Omega(u) \cap \Omega(w)|$ is even.
2. If $s(r, w) = 0$ for all $w \in W$, then $r = 0$.

3.10 Let (V, s) be the symplectic form just described for $n = 3$. Show that each non-zero vector w can be represented so that $\Omega(w)$ is a unique 2-subset of Ω. Conclude that $\dim W = 4$, and that the polar space of 1-spaces and totally isotropic 2-subspaces is isomorphic to the generalized quadrangle Γ of Example 7, p. 75.

Show that this gives us an embedding $e : \Gamma \to \mathbf{P}(W)$ whose image is not a subspace of the projective geometry.

3.11 Suppose $e : \Gamma_1 \to \Gamma_2$ is an embedding of point-line geometries. Show that if Γ_1 is a linear space, then $e(\Gamma_1)$ is actually a subspace of Γ_2. Conclude from the spanning property that e is an isomorphism of point-line geometries. In particular, Γ_2 is a linear space.

3.12 Suppose A and B are two locally connected point-line geometries, each containing at least two points. Is $A \times B$ locally connected?

Chapter 4
Hyperplanes, Embeddings, and Teirlinck's Theory

Abstract L. Teirlinck's theory characterized projective spaces among singular spaces by an axiom on geometric hyperplanes. That theory is recast here so that it applies to any point-line geometry with a suitable subcollection $\mathcal{H}$ of geometric hyperplanes satisfying the Teirlinck axiom. If $\mathcal{H}$ can separate points, then all singular subspaces are projective spaces. This axiom affects the associated Veldkamp space $\mathcal{V}_{\mathcal{H}}$, and has a lot to do with the ability to embed a point-line space into a projective space.

4.1 Veldkamp Spaces

4.1.1 Introduction

It is a sad fact that the appreciation of some great insights comes late. Such is the case with the notion of *geometric hyperplanes* due to F. D. Veldkamp [143].[1]

[1] Of course they were not known by that name at that time. Despite Veldkamp's demonstration of the usefulness of this concept in characterizing embedded polar spaces (as "polar spaces" were defined at that time), the notion was basically ignored until the rediscovery of its usefulness 20 years later. For example, about 30 years after Veldkamp's work, J. Hall and the author, simply called the hyperplanes "big subspaces" [65]. The term "geometric hyperplane" was introduced by Mark Ronan in the 1980s [102]. Even then this caused a certain amount of friction, since we often call certain subspaces of affine spaces "hyperplanes," which are not at all geometric hyperplanes in the sense of Ronan-Veldkamp. In order to preserve a vital distinction, some preferred to call the latter "projective hyperplanes" to indicate that they were different from affine hyperplanes in just the way projective hyperplanes of projective spaces differed from their affine analogs. But, as Sect. 4.4 shows, geometric hyperplanes are intimately bound up with embeddings of geometries into classical projective spaces. Therefore there is a great linguistic need to distinguish a Veldkamp-type hyperplane (in the *domain* of the embedding morphism) from the projective hyperplane of the projective space (comprising the *range* of the embedding morphism). We can't just call both of them "projective hyperplanes" and leave the reader to sort it out. Words should be the servant of those using them, not the other way round. In this book, we adopt Ronan's very neutral term "geometric hyperplane." The words indicate its source: a subspace defined solely by the domain geometry. The word "projective hyperplane" will be restricted to the very special case of geometric hyperplanes of projective spaces, as the words suggest – and nothing more.

E.E. Shult, *Points and Lines*, Universitext, DOI 10.1007/978-3-642-15627-4_4,
© Springer-Verlag Berlin Heidelberg 2011

From the point of view of merely describing point-line geometries, it is not all that useful. But from the point of view of creating the machinery which would allow us to classify geometries from first principles – that is, from certain axioms on points and lines, – it was the insight of Veldkamp that was the key concept.

Suppose you wish to characterize a geometry Γ by abstract axioms. You would probably have a head start if you could classify the *embedable* geometries – that is, the geometries that can be embedded into classical projective space. Let us make this a little more precise.

A *projective embedding of a point-line geometry* $\Gamma = (\mathcal{P}, \mathcal{L})$ is an object-injective point-line geometry morphism $\phi : \Gamma \to P_{1,2}(V) := \mathbf{P}(V)$ taking points of $\mathcal{P}$ to 1-spaces of the vector space V, and the lines of $\mathcal{L}$ to 2-spaces of V so that the point shadows of lines of Γ are surjectively mapped onto the point-shadows of the image projective lines, $\phi(L), L \in \mathcal{L}$. In effect, it is a geometry isomorphism from a point-line geometry to an embedded subgeometry of a projective space. For convenience, it is often assumed that the image points $\phi(\mathcal{P})$ span $\mathbf{P}(V)$.

To say that a point-line geometry is embedable is quite strong. Obviously many subconfigurations of points and lines can't exist inside a projective space as a subgeometry in the special sense that the lines of the embedded subgeometry "fill out" the actual projective lines.

Suppose, as a pioneer in the field, you realized this. Wouldn't it be valuable to show that many classical geometries are embedable in this particular way? In this "garden of forked paths"[2] the correct choice was taken very early. In looking at how classical polar space geometries are embedded by their natural descriptions, Veldkamp saw that there was a way that this description could be characterized intrinsically, knowing only the incidence of points and lines. Thus was born the notion that we now call a *Veldkamp space* a point-line geometry.

4.1.2 Geometric Hyperplanes

A *geometric hyperplane* of a point-line geometry Γ is a *proper* subspace S of Γ with with no external lines. Thus, for every line L, the point-shadow of L meets the proper subspace S non-trivially.

The examples of point-line geometries presented in Chaps. 2 and 3 already provide numerous geometric hyperplanes. However, to show that a prescribed set of points is a geometric hyperplane is really an exercise, in each instance. At this point the student is invited to peruse such a collection of exercises in Sect. 4.6.1 at the end of this chapter.

[2] To borrow a charming image from the writer Borjes.

4.1.3 The Veldkamp Space

Given $\Gamma = (\mathcal{P}, \mathcal{L})$ let $\mathcal{V}$ be the collection of all geometric hyperplanes of Γ. We begin with two very fundamental lemmas.

Lemma 4.1.1 *A geometric hyperplane H of a point-line geometry $\Gamma = (\mathcal{P}, \mathcal{L})$ is a maximal subspace if and only if the collinearity graph induced on $\mathcal{P} - H$ is connected.*

Proof This is left for the student as Exercise 4.8 of Sect. 4.6.2. □

We say that *Veldkamp points exist* if and only if every geometric hyperplane is a maximal subspace.

We say that *Veldkamp lines exist* if and only if, for any three geometric hyperplanes A, B and C, with A not contained in C,

$$A \cap B \subseteq C \text{ implies } C \cap B \subseteq A.$$

In general, we say that *Veldkamp $r-1$-spaces exist* if and only if the following two assertions both hold:

1. Veldkamp $r-2$-subspaces exist (or Veldkamp lines exist, if $r = 4$)
2. If $A_1, \ldots A_r$ are geometric hyperplanes with the property that $A_1 \cap \cdots \cap A_{r-1} \subseteq A_r$ and $A_1 \cap \cdots \cap A_{r-2}$ is not contained in A_r, then

$$A_1 \cap \cdots \cap A_{r-2} \cap A_r \subseteq A_{r-1}$$

Now suppose Veldkamp lines exist for $\Gamma = (\mathcal{P}, \mathcal{L})$. Suppose A and B are distinct geometric hyperplanes of Γ. Then for any two distinct hyperplanes X and Y containing $A \cap B$, one has $A \cap B = X \cap Y$. Thus the collection $L(A, B)$ of all geometric hyperplanes containing $A \cap B$ is the same collection $L(X, Y)$ for all distinct pairs $X, Y \in L(A, B)$. In that case $L(A, B)$ is called the *Veldkamp line determined by A and B*.

Let $\mathcal{V}$ be the collection of all geometric hyperplanes of Γ, and let $\mathcal{V}_2$, be the collection of all intersections $H \cap K$ of distinct hyperplanes H and K of Γ. We say that an intersection $A \cap B$ in $\mathcal{V}_2$ is *incident* with a geometric hyperplane X if and only if $A \cap B \subset X$. With respect to this definition of incidence one obtains a rank two geometry $(\mathcal{V}, \mathcal{V}_2)$ which we denote $\mathcal{V}(\Gamma)$.

Lemma 4.1.2 *Suppose Veldkamp lines exist for the point-line geometry $\Gamma = (\mathcal{P}, \mathcal{L})$. Then the incidence geometry $\mathcal{V}(\Gamma) = (\mathcal{V}, \mathcal{V}_2)$ is a linear space. (In this case the space $\mathcal{V}(\Gamma)$ is called the Veldkamp space of Γ).*

Remark Note that there is no asserted Veldkamp space unless Veldkamp lines already exist.

It will be useful in further work to be able to verify that Veldkamp $(r-1)$-spaces exist for higher values of r and for particular collections of geometric hyperplanes.

For this purpose we generalize Lemma 4.1.1. First we note a convenient way to render the basic condition.

Lemma 4.1.3 *Let $\mathcal{H}$ be a collection of geometric hyperplanes of a point-line geometry $\Gamma = (\mathcal{P}, \mathcal{L})$. Veldkamp $(r-1)$-spaces exist for $\mathcal{H}$ if and only if, for every $s \leq r$, and hyperplanes $A, A_1, \ldots, A_s$ of $\mathcal{H}$, it is true that $A_1 \cap \cdots A_{s-1} \cap A$ is either $A_1 \cap \cdots \cap A_s$ or $A_1 \cap \cdots \cap A_{s-1}$.*

Proof This is an exercise for the student (again, Exercise 4.8, Sect. 4.6.2). □

Lemma 4.1.4 *Let $\mathcal{H}$ be any set of geometric hyperplanes of a geometry $\Gamma = (\mathcal{P}, \mathcal{L})$ and fix a positive interger r. Suppose, for any subset $\{A_1, \ldots, A_s\}$ of $\mathcal{H}$, where $s \leq r$, that*

$$A_1 \cap \cdots \cap A_{s-1} - A_1 \cap \cdots \cap A_s$$

has a connected collinearity graph (this includes the case that the difference of these sets is empty).

Then Velkamp $(r-1)$-spaces exist for $\mathcal{H}$.

Proof We need only verify the condition given in Lemma 4.1.3. Suppose, for hyperplanes in $\mathcal{H}$, $A_1 \cap \cdots \cap A_s \subseteq A$. Then we have

$$A_1 \cap \cdots \cap A_s \subseteq A_1 \cap \cdots \cap A_{s-1} \cap A \subseteq A_1 \cap \cdots \cap A_{s-1}.$$

The last set is a subspace S of Γ and the first set $H := A_1 \cap \cdots \cap A_s$ is either equal to S or is a geometric hyperplane of it. If $H = S$ all three sets are equal. Otherwise, H is a hyperplane of S and by hypothesis, $S - H$ has a connected collinearity graph. It follows from Lemma 4.1.1 that H is a maximal subspace of S. Thus, in either case $A_1 \cap \cdots \cap A_{s-1} \cap A$ is either H or S. But this is exactly the condition for Veldkamp $(r-1)$-spaces to exist for $\mathcal{H}$ in Lemma 4.1.3. The proof is complete. □

If Veldkamp lines exist for $\mathcal{H}$, then the incidence system

$$\mathcal{V}_{\mathcal{H}} = (\mathcal{H}, \mathcal{H}_2),$$

where $\mathcal{H}_2$ denotes the collection of all intersections $H_1 \cap H_2$ of distinct members of $\mathcal{H}$, is a linear space called the *Veldkamp space relative to the collection $\mathcal{H}$*.

4.2 Teirlinck's Theory

4.2.1 The Hypotheses

Fix a point-line geometry $\Gamma = (\mathcal{P}, \mathcal{L})$. In this section we suppose that $\mathcal{H}$ is a collection of geometric hyperplanes of Γ subject to these hypotheses:

(H1) *$\mathcal{H}$ contains at least one non-empty geometric hyperplane and the intersection of all the geometric hyperplanes in $\mathcal{H}$ is the empty set.*

(H2) (Teirlinck's condition for $\mathcal{H}$.) *Suppose H_1 and H_2 are distinct members of $\mathcal{H}$, and p is any point. Then there exists an element $H_3 \in \mathcal{H}$ containing $(H_1 \cap H_2) \cup \{p\}$.*

(H3) (Veldkamp lines exist for $\mathcal{H}$.) *Suppose A, B, C are elements of $\mathcal{H}$ with $A \cap B \subseteq C$, but A not contained in C. Then $A \cap C \subseteq B$. Also, no member of $\mathcal{H}$ properly contains another.*

Remarks 1. The requirement that $\mathcal{H}$ possesses at least one non-empty hyperplane is a matter of convenience for it excludes only a trivial case. If the empty set were a geometric hyperplane, the geometry Γ could have no lines – the whole geometry would just be a set of points.

2. Note that the axioms make sense even if the members of $\mathcal{H}$ are not geometric hyperplanes. Therefore we will try to prove as much as we can without using the fact that they are hyperplanes: and we shall indicate the places in which we do make use of this assumption.

The collection $\mathcal{H}$ defines a closure operator $\tau : 2^{\mathcal{P}} \to 2^{\mathcal{P}}$ on the poset of subsets of $\mathcal{P}$ as follows. For any subset X of $\mathcal{P}$, $\tau(X)$ will be $\mathcal{P}$ itself if X lies in no member of $\mathcal{H}$, and will be the intersection of all members of $\mathcal{H}$ which contain X otherwise. It is clear from this definition that

$$X \subseteq Y \subseteq \mathcal{P} \text{ implies } \tau(X) \subseteq \tau(Y), \text{ and}$$
$$\tau^2 = \tau.$$

In general, we call $\tau(X)$ the *$\mathcal{H}$-closure of X*, and say that X is *$\mathcal{H}$-closed* if and only if $X = \tau(X)$.

4.2.2 The Exchange Property in Closed Sets

We first observe the following.

Theorem 4.2.1 (Teirlinck) (The exchange property.) *Suppose $\mathcal{H}$ satisfies the standing hypotheses (H1), (H2), (H3), and suppose A is $\mathcal{H}$-closed, and b is a point in $\tau(A \cup \{a\})$ which is not in A. Then $a \in \tau(A \cup \{b\})$.*

Proof Suppose $a \notin \tau(A \cup \{b\})$. Then there is an element H_2 of $\mathcal{H}$ containing $\tau(A \cup \{b\})$ but not a. Since $b \notin A$, there is an element H_1 containing A but not b. Now by axiom (H2), there is an element H_3 of $\mathcal{H}$ containing a and $H_1 \cap H_2$. Now if $H_2 \subseteq H_3$, we must have $H_2 = H_3$ (by the last part of (H3)), which would force $a \in H_2$, contrary to the choice of H_2. Thus H_2 is not a subset of H_3, so (H3) makes $H_1 \cap H_2 \subseteq H_3$ imply $H_3 \cap H_2 \subseteq H_1$. That is impossible, as the left side contains b while the right side does not. The proof is complete. □

Lemma 4.2.2 *Every point lies in some member of $\mathcal{H}$.*

Proof Fix a point p. We wish to show that p lies in a member of $\mathcal{H}$. If $\mathcal{H}$ contains two distinct members, either one of them contains p or else (H2) provides a third

member which does. If $\mathcal{H}$ contains one member, by (H1) that member is the empty set, an excluded case. □

4.2.3 To What Extent Does $\mathcal{H}$ Separate Points?

Lemma 4.2.3 *Suppose x and y are distinct points of $\mathcal{P}$. If there is an element of $\mathcal{H}$ containing x but not y, then there is an element of $\mathcal{H}$ containing y but not x.*

Proof Suppose $H_1 \in \mathcal{H}$ contains x but not y. By $(H1)$ there is an $H_2 \in \mathcal{H}$ not containing x. If H_2 contains y we are done. Otherwise, there is an $H_3 \in \mathcal{H}$ containing y and $H_1 \cap H_2$. Now by Lemma 4.1.1, $H_1 \cap H_2 = H_1 \cap H_3 = H_2 \cap H_3$. But if H_3 also contained x, then $H_1 \cap H_3$ would contain x. Since this intersection is in H_2 which does not contain x, we should have a contradiction. Thus x is not in H_3, and H_3 is the desired element of $\mathcal{H}$. □

Now consider the following relation on the set of points, $\mathcal{P}$. We write $x \equiv y$ if and only if $y \in \tau(x)$ – that is, if and only if every element of $\mathcal{H}$ which contains x also contains y. Transparently, this relation is reflexive and transitive.

Lemma 4.2.4 *The relation "$\equiv$" is symmetric.*

Proof Suppose $x \equiv y$ but that some element of $\mathcal{H}$ contains y but not x. Then by the previous lemma, there is an element of $\mathcal{H}$ containing x but not y, and that is impossible by hypothesis. □

It follows that the relation "$\equiv$" is an equivalence relation. For each point $p \in \mathcal{P}$, let $[p]$ denote the equivalence class with respect to "$\equiv$" which contains point p. Thus a sort of "granular structure" is imposed on the $\mathcal{H}$-closed sets:

(*) *Each $\mathcal{H}$-closed set A is the union of its equivalence classes $[a]$, as a ranges over A.*

We now have the following.

Lemma 4.2.5 *For each point p, $[p]$ is a subspace of $\Gamma = (\mathcal{P}, \mathcal{L})$.*

Proof Suppose L is a line intersecting $[p]$ at two distinct points, x and y. Then any element of $\mathcal{H}$ which contains x also contains y, and, since it is a subspace of Γ, contains L. Thus $L \subseteq [p]$. The proof is complete. □

Remark Note that this is the first time we have used the fact that elements of $\mathcal{H}$ are subspaces. From this point onward, we make full use of the fact that elements of $\mathcal{H}$ are geometric hyperplanes – that is, $\mathcal{H} \subseteq \mathcal{V}$.

Lemma 4.2.6 *Suppose A is $\mathcal{H}$-closed and x is a point not in A. Then A is a geometric hyperplane of the subspace $\tau(A \cup \{x\})$.*

Proof Since x is not in A, which is $\mathcal{H}$-closed, there is an element H_1 of $\mathcal{H}$ containing A but not x. Thus since H_1 is a geometric hyperplane of Γ, $K := H_1 \cap \tau(A \cup \{x\})$ is a geometric hyperplane of $\tau(A \cup \{x\})$ as well as being an $\mathcal{H}$-closed subspace. But if there were a point $y \in K - A$, then by the Exchange Theorem (Theorem 4.2.1), $\tau(A \cup \{y\}) = \tau(A \cup \{x\})$. But since K is a proper $\mathcal{H}$-closed subspace of $\tau(A \cup \{x\})$ containing y, that is impossible. Thus no such y exists, and we see that $A = K$ is now a geometric hyperplane of $\tau(A \cup \{x\})$. □

Corollary 4.2.7 *The equivalence classes* $[p]$ *are always cocliques of the point-collinearity graph of* Γ.

Proof Suppose x and y were distinct points of $[p]$ on some line L. Then $L \subseteq [p]$ since the latter is a subspace. Now by (H1) there is a hyperplane $H_1 \in \mathcal{H}$ not containing x. Then $H_1 \cap L$ is a single point z. But this contradicts $[z] = [x]$. Thus no distinct collinear pairs of points exist in $[p]$. □

4.2.4 The $\mathcal{H}$-Closure of Two Inequivalent Points

In the following we shall denote the $\mathcal{H}$-closure of two distinct points x and y by $\tau(x, y)$, rather than $\tau(\{x, y\})$.

Theorem 4.2.8 *1. Set $S = \tau(x, y)$ for two inequivalent points x and y. Then for any pair of distinct points u and v in S for which $[u] \neq [v]$, one has $S = \tau(u, v)$.*
2. If $S = \tau(x, y)$ for two inequivalent points x and y, then either S is a coclique, or $S = \cup_{u \in L}[u]$ for some line L.

Proof 1. It suffices to show that if $v \in \tau(x, y) - [x]$ then $\tau(x, y) = \tau(x, v)$. But this follows easily from the Exchange Theorem (Theorem 4.2.1).

2. Suppose a and b are distinct collinear points in $S = \tau(x, y) = \tau(a, b)$. Since S is a subspace of Γ, any line L on a and b lies in S.

   Suppose $z \in S - L$ but that z belongs to no class $[u]$ for $u \in L$. In particular, $[z] \neq [a]$ so there is a hyperplane H from $\mathcal{H}$ which contains z but does not contain a. Then $H \cap L$ contains a second point y', which by hypothesis does not belong to $[z]$. Then H must contain $\tau(z, y')$. But by part 1 of this proof, $\tau(z, y') = \tau(x, y) = S$, so H contains S. That is impossible since H does not contain the point a. Thus no such point z exists in S and so S is the union of the classes $[u]$ as u ranges over the points of L. □

In general, $(\mathcal{P}, \mathcal{L})$ need not be a partial linear space, although (as we shall see) its singular subspaces must be partial linear spaces. Consider the following:

Example 1 Let $\mathcal{P} = H_1 + H_2 + \cdots + H_n, n > 1$ be a partition of a set. A subset T of $\mathcal{P}$ is called a *transversal* if it contains exactly one element from each component of the partition H_i. Now let $\mathcal{L}$ be any non-empty collection of transversals whatsoever. Set $\mathcal{H} = \{H_i | 1 \leq i \leq n\}$. Then $(\mathcal{P}, \mathcal{L})$ is a point-line geometry, and each H_i is a geometric hyperplane of this geometry. Now, since the intersection of any two distinct hyperplanes of $\mathcal{H}$ is empty, all of the axioms (H1), (H2), and (H3) are in

place. Clearly we can make examples where $(\mathcal{P}, \mathcal{L})$ is not a partial linear space (one could take nearly any rank n residually connected geometry, let the partition be that induced by the type function, and let $\mathcal{L}$ be the collection of all flag-chambers).

However, the next section will show that $(\mathcal{P}, \mathcal{L})$ has a nice homomorphic image which is indeed a partial linear space.

4.2.5 The Natural Morphism

Let $\bar{\mathcal{P}} = \{[v] \mid v \in \mathcal{P}\}$, and let $\phi : \mathcal{P} \to \bar{\mathcal{P}}$ be the mapping which takes each point p to the equivalence class $[p]$ which contains it. Then by Corollary 4.2.7, the restriction of ϕ to any singular subspace is injective. In particular this is true of lines. Let $\bar{\mathcal{L}} = \{\phi(L) | L \in \mathcal{L}\}$. We now have a morphism of geometries

$$\Gamma = (\mathcal{P}, \mathcal{L}) \to \bar{\Gamma} := (\bar{\mathcal{P}}, \bar{\mathcal{L}}),$$

where the incidence on the right is the obvious one: $[u]$ is incident with $\bar{L}$ if $[u]$ is one of the classes in $\bar{L}$.

For each $\mathcal{H}$-closed set A of Γ, let $\bar{A} = \phi(A)$. The "granular" phenomenon in (*) shows that

(**) *if $x \notin A$ then $\phi(x) \notin \bar{A}$.*

In particular, for any geometric hyperplane $H \in \mathcal{H}$, $\bar{H} := \phi(H)$ is a geometric hyperplane of $\bar{\Gamma}$. We set $\bar{\mathcal{H}} := \{\bar{H} | H \in \mathcal{H}\}$, and note that because of (**) the mapping $\mathcal{H} \to \bar{\mathcal{H}}$ induced by ϕ is bijective.

Using the observation (**) above, it is now elementary to see the following.

Lemma 4.2.9 $(\bar{\Gamma}, \bar{\mathcal{H}})$ *satisfies the three axioms* (H1), (H2), (H3) *of Teirlinck's theory. In addition it satisfies the following.*

(H4) ($\bar{\mathcal{H}}$ separates points.) *For any two distinct points $\bar{p}$ and $\bar{q}$ of $\bar{\mathcal{P}}$, there is a hyperplane $\bar{H}$ in $\bar{\mathcal{H}}$, containing $\bar{p}$ but not $\bar{q}$.*

For each subset $\bar{X}$ of $\bar{\mathcal{P}}$, we let $\bar{\tau}(\bar{X})$ be the intersection of all elements of $\bar{\mathcal{H}}$ which contain $\bar{X}$ (with the usual convention that this is $\bar{\mathcal{P}}$ when no such elements of $\bar{\mathcal{H}}$ exist) – that is, the $\bar{\mathcal{H}}$-closure of $\bar{X}$. Then it is an easy exercise to show the following.

Lemma 4.2.10 *For all $X \subseteq \mathcal{P}$, one has $\phi(\tau(X)) = \bar{\tau}(\phi(X))$.*

Theorem 4.2.11 *For any two distinct points $\bar{p}$ and $\bar{q}$ of $\bar{\mathcal{P}}$, we have the following:*

1. *$\bar{\tau}(\bar{p}, \bar{q}) = \bar{\tau}(\bar{u}, \bar{v})$ for any two distinct points $\bar{u}$ and $\bar{v}$ in $\bar{\tau}(\bar{p}, \bar{q})$.*
2. *If $\bar{p}$ is collinear with $\bar{q} \neq \bar{p}$, then $\bar{\tau}(\bar{p}, \bar{q})$ is any line on $\bar{p}$ and $\bar{q}$ – so there is only one such line of $\bar{\mathcal{L}}$.*
3. *$\bar{\Gamma}$ is a partial linear space.*

Proof The first two statements are immediate from Theorem 4.2.8 and Lemma 4.2.10. The third statement follows from the second. □

4.2.6 Singular Subspaces are Generalized Projective Spaces

Now we come to a consideration of singular subspaces. In fact singular spaces have a very particular structure. A *generalized projective plane* is a firm linear space in which any two lines intersect at a point. If all lines are thick, these are called *projective planes*. If lines are not thick, for each integer n greater than one, there is only one isomorphism type, consisting of one "long" line whose complement is a single point.

Theorem 4.2.12 *Suppose S is an $\mathcal{H}$-closed singular subspace of Γ:*

1. *Any subspace of S which is generated by finitely many elements (as a subspace of Γ) is $\mathcal{H}$-closed*
2. *S is a linear space in which any three distinct points of S generate a line or a generalized projective plane*

Proof 1. Suppose a singular subspace S contains two distinct points. Then it contains a line. If, moreover, it can be generated as an $\mathcal{H}$-space by two elements, then, by Theorem 4.2.8, part 2, it is itself a line. Thus all lines of S are $\mathcal{H}$-closed.

Let us suppose $A = \tau(\{a_1, \ldots, a_n\})$, where the a_i lie in S and where we may suppose that the set $\{a_1, \ldots, a_n\}$ is minimal with respect to having its $\mathcal{H}$-closure equal to A. We must show that A is simply the (necessarily singular) subspace of Γ generated by $\{a_1, \ldots, a_n\}$. From the previous paragraph we may suppose $n > 2$, and apply induction on n. Thus we are permitted to suppose, that $H := \tau(\{a_1, \ldots, a_{n-1}\}) = \langle a_1 \ldots a_{n-1} \rangle_\Gamma$. But by Lemma 4.2.6, H is a geometric hyperplane of A. But since A is singular, that means $A - H$ has a connected induced point-collinearity graph, and so H is actually a maximal subspace of A by Lemma 4.1.1. It follows that

$$A = \langle H, a_n \rangle_\Gamma = \langle a_1, \ldots, a_n \rangle_\Gamma,$$

as required. So the first result holds by induction.

2. Now S is an arbitrary singular subspace of Γ, not necessarily $\mathcal{H}$-closed. Suppose, $\{a, b, c\}$ is a 3-subset of S which lies in no line. Set $A_0 := \tau(\{a, b, c\})$ and let $S_0 = A_0 \cap S$. Then for any line L in S_0, we see $\tau(L) = \cup_{u \in L}[u]$ is a geometric hyperplane of A_0 (Lemma 4.2.6). Thus $L = \tau(L) \cap S$ is either equal to S_0 or is a geometric hyperplane of S_0. Thus if S_0 is not itself a line, any of its lines L is a geometric hyperplane of S_0. In that case any two distinct lines of S_0 intersect at a one-point set. Thus S_0 is a generalized projective plane. □

Remark A linear space in which any three distinct points generate (as a subspace) either a line or a generalized projective plane is called a *generalized projective*

space. If lines are thick, it is simply called a *projective space.* Generalized projectives spaces can have their points partitioned into point-sets consisting of (1) a single point, or (2) a projective space (one with thick lines – including a single thick line all by itself), all edges bridging two of these components, themselves comprising a thin line. Thus everything is determined by these components. They will be studied in some detail in the next two chapters.

Lemma 4.2.13 *Suppose* Γ *is a singular space and* $\mathcal{H}$ *is a class of hyperplanes satisfying axioms (H1) and (H2). Then axiom (H3) must also hold.*

Proof Assume Γ is a singular space, and that (H1) and (H2) hold. First, for any $H \in \mathcal{H}$, $\mathcal{P} - H$ has a complete collinearity graph, so, by Lemma 4.1.1, H is a maximal subspace of Γ. Thus no hyperplane is properly contained in another.

Suppose A, B, and C are geometric hyperplanes of Γ with $A \cap B \subseteq C$ but such that A is not contained in C. Now $A \cap B$ and $A \cap C$ are both hyperplanes of A. But since $A - B$ has a connected collinearity graph, by Lemma 4.1.1 $A \cap B$ is a maximal subspace of B. Since no truly intermediate subspace such as $A \cap C$ is allowed between A and B, either $A \subseteq C$ or else $A \cap B = A \cap C$, which, as can be seen, is enough to prove (H3). □

Corollary 4.2.14 (Teirlinck's characterization of projective spaces [128].) *Suppose* Γ *is a singular space satisfying (H1) and (H2), relative to some collection* $\mathcal{H}$ *of geometric hyperplanes of* Γ. *Then* Γ *is a generalized projective space. Of course if* Γ *has only thick lines, it is a projective space.*

Proof By Lemma 4.2.13, Theorem 4.2.12 applies, showing that all subspaces generated by two intersecting lines are generalized projective planes. That's enough. □

Remarks This is a very useful theorem. In Chap. 7, we will use it to show that non-degenerate polar spaces have all their singular subspaces projective and so these polar spaces are partial linear spaces – all in one stroke.[3]

Teirlinck [128] used the hypotheses (H1) and (H2) entirely in the context of singular spaces (which he called "2-covers"). There he proved many theorems involving planarity, counting subspaces of various ranks, etc., beyond what we need to do here. His great insight was discovering the power of axiom (H2).[4] This section has represented the author's attempt to recast the most superficial parts of his theory in the wider context of general point-line geometries.

---

[3] This observation is due to N. Percsy and F. Buekenhaut, independently.

[4] I beg the reader to forgive a personal note. Francis Buekenhout once told me that most of incidence geometry was not proving theorems so much as discovering the right axioms. Pick the wrong ones and one proves only artificial-looking trivialities. Pick the right axioms and (1) it will be easy to check whether they are in place, (2) the axioms will be useful because they are realized in a lot of places, and (3) the consequences of the axioms will seem to be at a miraculous logical distance away. He suggested this was an art that good geometers learn. His comments impressed me early in my career, and for me made this sort of geometry one of the most esthetically attractive parts of mathematics.

4.3 The Effect of Teirlinck's Theory on the Veldkamp Space

This section offers two results implying that a Veldkamp space is a projective space: each enjoys a separate section. In order to accomplish this, we first intervene with a few remarks about how to recognize a generalized projective space.

4.3.1 The Veblen Axiom

On p. 87 we defined a *generalized projective plane* to be a linear space in which any two lines intersected at a point. We also defined a projective space as a linear space in which the subspace generated by any two lines which intersect at a point is a projective plane (again, see p. 88). To characterize such a space, one would have to test every pair of lines that could appear in the subspace generated by two distinct intersecting lines to see if they intersect. It is not so convenient to see how these lines are to be chosen from an unknown space generated by two intersecting lines. It is for this reason that it is convenient to invoke a special axiom called the Veblen axiom.

The Veblen axiom (V) *Suppose* $\Gamma = (\mathcal{P}, \mathcal{L})$ *is a partial linear space. If* L_1 *and* L_2 *are distinct lines intersecting in a point* p *and* N_1 *and* N_2 *are lines intersecting both* L_i *at points distinct from* p*, then* N_1 *and* N_2 *themselves intersect at a point.*

Lemma 4.3.1 *Suppose* $\Gamma = (\mathcal{P}, \mathcal{L})$ *is a linear space generated by two intersecting lines. Moreover, suppose that* Γ *satisfies the Veblen axiom (V). Then any two lines of* Γ *possess a non-empty intersection. In particular,* Γ *is a generalized projective plane.*

Proof This proof fully utilizes the description of subspaces generated by a prescribed set of points expounded in Sect. 3.4, especially in Lemma 3.4.1. The reader may recall that such a subspace is built up from a generating set X by taking set $\mathcal{L}(X)$ of secant and internal lines with respect to the set X, then adjoining to X all points incident with these lines – the set $\mathcal{P}(\mathcal{L}(X))$. Iterating this process, we obtain new generations of lines $\mathcal{L}(\mathcal{P}(\mathcal{L}(X)))$. We propose to show how this works when the Veblen axiom (V) is in charge.

Recall that we begin with a set of two distinct lines $\mathcal{L}_0 = \{L_1, L_2\}$, intersecting at the point p and begin with point set $X_0 = L_1 \cup L_2$. We then form the set of secant lines $\mathcal{L}_1 := \mathcal{L}(X_0)$ and the set X_1 of points they contain, and repeat these operations to obtain ascending chains of point- and line-sets:

$$X_0 \subseteq X_1 \subseteq X_2 \subseteq \cdots, \tag{4.1}$$

$$\mathcal{L}_0 \subseteq \mathcal{L}_1 \subseteq \mathcal{L}_2 \subseteq \cdots. \tag{4.2}$$

Suppose at some stage we reach a generation X_i of points which we will describe as "*very old points*." We then call the secant lines $\mathcal{L}_{i+1} := \mathcal{L}(X_i)$ "old lines," and

we call the set $X_{i+1} := \mathcal{P}(\mathcal{L}(X_i)$ of points they contain "*old points*," a set which clearly contains the *very old points*. Finally, by taking secants of this set we obtain a set $\mathcal{L}(X_{i+2})$ which may contain some new lines. Definitions prove nothing, but at least the words confine us to these rules:

1. Any line incident with two *very old points* is an *old line*
2. Any old line is incident with at least two very old points
3. Any point incident with an old line is an *old point* – possibly a *very old point*

A line that is incident with two distinct *old points*, but which is not itself an old line, will be called a *new line*. Clearly a *new line* cannot be incident with two distinct *very old points*.

We wish to prove that all lines in this generated subspace intersect. We can do this by induction on the index that records the successive generations of lines that evolve from the process just described. The Veblen axiom itself takes care of the first two generations of lines, $\mathcal{L}_0$ and $\mathcal{L}_1 = \mathcal{L}(X)$. So by induction we may assume the following:

(*Old*) *Any two* old lines *intersect non-trivially.*

Using axiom (V) we intend to show two things:

Step 1. *Any new line intersects each old line non-trivially.*
Step 2. *Any two new lines intersect each other.*

To prove Step 1 suppose N is a *new line* and M is any *old line*. By definition, N contains two distinct old points – say u and v. Suppose there exists a *very old point* p not on either M or N. Then by assumption (Old), the lines pu and pv must both meet the line M. Applying the Veblen axiom (V) to the configuration $\{N, M, pu, pv\}$ one must conclude that M and N intersect non-trivially, as desired.

So we must assume every *very old point* either lies on line M or is the unique such point on line N. This is impossible if Γ is being generated by the array of *very old points* of the two generating thick lines. Thus one of the generating lines is not thick, and in that case it easily follows that Γ is just the degenerate plane with a unique thick line for the reason that the process of producing succesive generations of points doesn't even reach a second generation in that case.

To prove Step 2 we consider two distinct new lines N and M. We must show that they intersect. By definition of "new line," the line N carrries two distinct old points $\{a, b\}$. By definition these two points a and b must live on old lines A and B respectively, and each of these lines must carry at least two very old points. As old lines, A and B must intersect at a point x. Further, by Step 1, the new line M must intersect each of the old lines A and B at points.

If x is not in M, then (A, B, M, N) are in the configuration of the Veblen axiom with A and B as the initial intersecting lines. Thus N and M intersect by (V).

So we must assume $x \in M$. We mentioned that A and B each carry at least two very old points. But if a and b were both very old, then N would be an old line

contrary to its choice as a new line. Thus one of these two points – say b – is not a very old point. Thus the line B carries a very old point b' not equal to either x or b. Also, A carries a very old point a' distinct from x (though $a' = a$ is possible). Since $\Gamma = \langle X_0 \rangle$ is a linear space, there is an old line C on a' and b'. Now C intersects both N and M by Step 1, and meets B at b'. Thus the lines (B, C, N, M) have the configuration of the Veblen axiom with B and C as the initial intersecting lines, forcing N and M to intersect by (V).

Thus we have shown that if every pair of lines of $\mathcal{L}_{i+1}$ intersect, then the same is true of $\mathcal{L}_{i+2}$. Thus $\Gamma = \langle X_0 \rangle$ is a linear space in which any two lines intersect, and so is a generalized projective plane. □

Corollary 4.3.2 *Suppose $\Gamma = (\mathcal{P}, \mathcal{L})$ is a linear space. Then Γ satisfies the Veblen axiom (V) if and only if it is a generalized projective space.*

Proof Suppose Γ is a generalized projective space. Then by definition any two lines in a subspace spanned by two intersecting lines must intersect. Thus one has the Veblen axiom for such subspaces, and so (V) (which posits two such intersecting line L_1 and L_2 and focuses on two lines N_1 and N_2 in the subspace that they generate) must hold.

Conversely, suppose Γ is a linear space satisfying axiom (V). Let P be a subspace generated by two intersecting lines. Then axiom (V) is valid for the lines within the subspace P. By Lemma 4.3.1, P is a generalized projective plane. Thus Γ satisfies the definition of a generalized projective space. □

4.3.2 The First Reason Why a Veldkamp Space Might be a Projective Space

Here we wish to show that if a collection $\mathcal{H}$ of geometric hyperplanes for which Veldkamp planes exist satisfies Tierlinck's condition then the corresponding Veldkamp space is a generalized projective space.

Theorem 4.3.3 *Assume $\Gamma = (\mathcal{P}, \mathcal{L})$ is a connected point-line geometry and $\mathcal{H}$ is a collection of geometric hyperplanes with these properties:*

1. *Teirlinck's condition holds for $\mathcal{H}$ (condition (H2) of Teirlinck's theory)*
2. *Velkamp planes exist for $\mathcal{H}$ – that is:*
   - *(a) No member of $\mathcal{H}$ is properly contained in another*
   - *(b) For A, B, C in $\mathcal{H}$, $A \cap B \subseteq C$, and $A \not\subseteq C$, one has $A \cap C \subseteq B$*
   - *(c) For $A, B, C, D \in \mathcal{H}$, $A \cap B \not\subseteq D$, the statement $A \cap B \cap C \subseteq D$ implies $A \cap B \cap D \subseteq C$*

Then the Velkamp space $\mathcal{V}_{\mathcal{H}}$ is a generalized projective space.

Proof We need only verify the Veblen axiom for two intersecting thick lines of $\mathcal{V}_{\mathcal{H}}$. We draw the configuration for this axiom in Fig. 4.1, naming the Veldkamp lines and points.

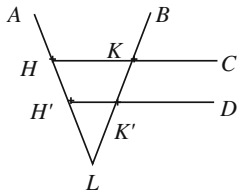

Fig. 4.1 The configuration for the Veblen axiom in the Veldkamp space $\mathcal{V}_{\mathcal{H}}$. Here the Veldkamp lines are $A = H \cap L$, $B = K \cap L$, $C = H \cap K$, and $D = H' \cap K'$

Then

$$H \cap L = H' \cap L = H \cap H', \tag{4.3}$$

$$K \cap L = K' \cap L = K \cap K', \tag{4.4}$$

and $H \cap H \neq H' \cap K'$.

Without loss of generality we may assume there is a point p in $(H' \cap K') - (H \cap K)$. Note that p does not lie in L, for if it did, then it would lie in

$$H' \cap (K' \cap L) = H' \cap (L \cap K) = (H \cap L) \cap K \subseteq H \cap K,$$

contrary to the choice of p.

Now by the Teirlinck condition, there is a geometric hyperplane R containing p and $H \cap K$. We claim that $H' \cap K' \subseteq R$.

Suppose not. Then

$$H' \cap K' \nsubseteq R,$$

but by Eq. (4.3)

$$H' \cap K' \cap L = H \cap K \cap L \subseteq R.$$

So, since Veldkamp planes exist,

$$H' \cap K' \cap R \subseteq L.$$

But that is impossible, as the point p is in the left side, but is not in the right side.

Thus the hyperplane R is a Veldkamp point incident with both the Veldkamp lines $H \cap K$ and $H' \cap K'$. So the Veblen axiom holds. □

Remark The above proof is due to A. Pasini, in a personal communication to the author. (The fact that he was using it for polar spaces of rank at least four did not detract from the transportability of his argument. The polar space context could be replaced by the hypothesis that Veldkamp planes exist without changing his proof at all.)

4.3.3 Another Reason why the Veldkamp Space $\mathcal{V}_{\mathcal{H}}$ Might be a Projective Space

Theorem 4.3.4 *Assume $\Gamma = (\mathcal{P}, \mathcal{L})$ is a point-line geometry with a non-empty set $\mathcal{L}$ of thick lines. Suppose $\mathcal{H}$ is a family of geometric hyperplanes satisfying axioms (H1)–(H3) of Teirlinck's theory. Let $\mathcal{V}_{\mathcal{H}} = (\mathcal{H}, \mathcal{H}_2)$ be the attendant Veldkamp space, where $\mathcal{H}_2$ is the set of all proper intersections $A \cap B$, $A, B \in \mathcal{H}$. Then $\mathcal{V}_{\mathcal{H}}$ is a linear space. Suppose, in addition, the following.*

(H5) *For any two points p and q with $p \notin [q]$, and any $H \in \mathcal{H}$,*

$$H \cap \tau(p, q) \neq \emptyset.$$

Then $\mathcal{V}_{\mathcal{H}}$ is a generalized projective space. (Of course if all Veldkamp lines from $\mathcal{H}$ are thick, it is a projective space.)

Proof For each point p let $\mathcal{H}_p := \{H \in \mathcal{H} | p \in H\}$. Lemma 4.2.2 shows this set is never empty. Also, by (H1), $\mathcal{H}_p$ is a proper subset of $\mathcal{H}$. Obviously if p lies in distinct elements A and B of $\mathcal{H}$, it lies in any member of $\mathcal{H}$ containing $A \cap B$. Thus $\mathcal{H}_p$ is a proper subspace of the Veldkamp space for $\mathcal{H}$, $\mathcal{V}_{\mathcal{H}}$.

We proceed by a series of short steps.

Step 1. *For each point p, $\mathcal{H}_p$ is a geometric hyperplane of $\mathcal{V}_{\mathcal{H}}$.*

This is an immediate consequence of Lemma 4.5.1.

Now let $\mathcal{H}^* = \{\mathcal{H}_p | p \in \mathcal{P}\}$, a subcollection of hyperplanes of $\mathcal{V}_{\mathcal{H}}$.

Step 2. $\cap_{p \in \mathcal{P}} \mathcal{H}_p = \emptyset$.

Any hyperplane in the indicated intersection would have to contain all points and so could not be a proper subspace.

Step 3. *Given distinct points p and q of Γ, and any geometric hyperplane $H \in \mathcal{V}_{\mathcal{H}}$, there exists a point r such that:*

1. $\mathcal{H}_r \supseteq \mathcal{H}_p \cap \mathcal{H}_q$
2. $H \in \mathcal{H}_r$

Now by (H5), there is a point $r \in \tau(p, q) \cap H$. But then 1 and 2 are immediate.

We can now complete the proof of the theorem. By the three steps above, $\mathcal{V}_{\mathcal{H}}$ is a linear space having a system of hyperplanes $\mathcal{H}^*$ satisfying (H1) and (H2). Thus $\mathcal{V}_{\mathcal{H}}$ is a generalized projective space by Teirlinck's characterization theorem (Corollary 4.2.14). □

Remark Unfortunately condition (H5) does not hold generally.

Much of the work for conclusion 2 of this Theorem appears in [113].

4.4 The Connections Between Projective Embeddings and Veldkamp Spaces

Let $\Gamma = (\mathcal{P}, \mathcal{L})$ be a point-line geometry. Let V be a vector space, and let $\mathbf{P}(V)$ be its associated *classical projective space*. (Recall that $\mathbf{P}(V) = (P_1(V), P_2(V))$ is the point-line geometry of one- and two-dimensional subspaces of vector space V (the projective points and lines)).[5]

A *projective embedding of the point-line geometry* $\Gamma = (\mathcal{P}, \mathcal{L})$ is a point-injective mapping $e_1 : \mathcal{P} \to P_1(V)$ of geometric points onto a spanning set of 1-subspaces of a vector space V, and a mapping $e_2 : \mathcal{L} \to P_2(V)$ such that for each line $L \in \mathcal{L}$, the image of its point-shadow, $e_1(sh(L))$, is all 1-spaces of the line $e_2(L) \in P_2(V)$. We denote such an embedding by the symbol $e : \Gamma \to \mathbf{P}(V)$.

Now let $\mathbf{H}$ be the projective point-shadow of some vector-space hyperplane of V (recall that such a set of points was a geometric hyperplane of the projective space called a *projective hyperplane* – see Exercise 4.1. at the end of this chapter). It is an easy exercise to see that

$$e_1{}^{-1}(\mathbf{H} \cap e(\mathcal{P})),$$

the set of all points which map into $\mathbf{H}$, is a geometric hyperplane of $(\mathcal{P}, \mathcal{L})$. We unambiguously denote this geometric hyperplane by the symbol $e^{-1}(\mathbf{H})$. We say that a geometric hyperplane M of Γ *arises from the projective embedding* e if and only if $M = e^{-1}(\mathbf{H})$ for some projective hyperplane $\mathbf{H}$.

We let $\mathcal{V}_e$ denote the collection of all geometric hyperplanes of Γ which arise from the embedding e.

Theorem 4.4.1 *Suppose* $e : \Gamma \to \mathbf{P}(V)$ *is a projective embedding of the point-line geometry* $\Gamma = (\mathcal{P}, \mathcal{L})$. *Suppose Veldkamp* $(r-1)$*-spaces exist for* Γ (remember this is an entirely intrinsic property of Γ: it does not depend on e). *Then the following statements hold:*

1. *Every vector subspace* U *of codimension at most* r *in* V *is spanned by the image points* $e(\mathcal{P})$ *which lie within it.*
2. *If* r *is at least* 2 *(so Veldkamp lines exist),* $\mathcal{V}_e$ *is a subspace of the full linear Veldkamp space* $\mathcal{V}$.
3. *There is an incidence-preserving bijection between the subspaces of* V *of codimension at most* r *and the flats of* $\mathcal{V}_e$ *of dimension at most* r. *(This map converts dimension to codimension.)*
4. *If Veldkamp lines exist,* $\mathcal{V}_e$ *is a projective space.*

[5] Following the convention of Chaps. 2 and 3, unless the dimension is so low that it is unavoidable, the *projective space* $\mathbf{P}(V)$ will *never* mean the entire *projective geometry* of proper subspaces of V at all possible dimensions – the latter is denoted $PG(V)$.

Proof 1. For each subspace U of V, set $S(U) := \{p \in \mathcal{P} | e(p) \in U\}$. Clearly $S(U)$ is a subspace of Γ and if U has codimension 1, then $S(U)$ is a geometric hyperplane of Γ. We proceed by induction on r.

Assume $r = 1$, so Veldkamp points exist. Then all geometric hyperplanes are maximal subspaces. Let $\mathbf{H}$ be a projective hyperplane of $\mathbf{P}(V)$, that is, $\mathbf{H}$ is the set of one-dimensional vector subspaces of a vector subspace of codimension 1 in V, and suppose $\mathbf{K} := \langle e(\mathcal{P}) \cap \mathbf{H} \rangle_V$ has codimension at least 2 in V. Then there exists a point $p \in \mathcal{P} - S(\mathbf{K})$ and so there is a hyperplane $\mathbf{H}'$ containing $\langle \mathbf{K}, e(p) \rangle_V$. Then $S(\mathbf{H})$ is properly contained in $S(\mathbf{H}')$ while both are geometric hyperplanes. This contradicts $S(\mathbf{H})$ maximal.

Now assume the result holds for all values of r less that k, and suppose Veldkamp $(r-1)$-spaces exist for Γ. Suppose by way of contradiction that $\mathbf{U}$ is a vector subspace of codimension k in V such that $\mathbf{K} = \langle e(\mathcal{P}) \cap \mathbf{U} \rangle_V$ has codimension greater than k. Again choose $p \in \mathcal{P} - S(\mathbf{U})$. Then $\mathbf{K}$ is a vector-space hyperplane $\langle \mathbf{K}, e(p) \rangle_V$ and so there is a subspace $\mathbf{U}'$ of codimension k in V containing $\langle \mathbf{K}, e(p) \rangle_V$ and meeting $\mathbf{U}$ at a vector-space hyperplane (that is, $\mathbf{U} \cap \mathbf{U}'$ has codimension $k+1$). It follows that $\mathbf{U}$ and $\mathbf{U}'$ can be expressed as the intersection of almost the same k independent hyperplanes:

$$\mathbf{U} = \mathbf{H}_1 \cap \cdots \cap \mathbf{H}_k, \tag{4.5}$$

$$\mathbf{U}' = \mathbf{H}_1 \cap \cdots \cap \mathbf{H}_{k-1} \cap \mathbf{H}, \tag{4.6}$$

where $\mathbf{H}$ and the $\mathbf{H}_i$ are all vector-space hyperplanes of V.

Now set $S(\mathbf{H}_i) = A_i$ and $S(\mathbf{H}) = A$, all of them being geometric hyperplanes. Then

$$S(\mathbf{K}) = A_1 \cap \cdots \cap A_k \subseteq A,$$

and

$$S(\mathbf{K}) \subseteq A_1 \cap \cdots \cap A_{k-1} \cap A \subseteq A_1 \cap \cdots \cap A_{k-1}. \tag{4.7}$$

The first containment here is proper since the middle term contains point p. Since Veldkamp $(r-1)$-spaces exist, we must have the middle term equal to the extreme right. But then, by induction on k, the embedded image of this $(k-1)$-fold intersection on the right spans a subspace of codimension at most $k-1$. On the other hand, by Eq. (4.6), the embedded images of the middle term of Eq. (4.7) are trapped in $\mathbf{U}'$ of codimension k. This is a contradiction.

Thus $\mathbf{K} = \mathbf{U}$, as desired.

2. Suppose A and B are two distinct "points" of $\mathcal{V}_e$, that is $A = S(\mathbf{H}_A)$ and $B = S(\mathbf{H}_B)$ for two projective hyperplanes $\mathbf{H}_A$ and $\mathbf{H}_B$ of $\mathbf{P}(\mathbf{V})$ meeting at subspace $\mathbf{K}$ at codimension 2. Now suppose C is any geometric hyperplane of Γ which happens to contain $A \cap B$ and suppose C is distinct from both A and B. Since Veldkamp lines exist, we see that $A \cap B = A \cap C = B \cap C$ and so C cannot

be contained in $A \cup B$. Thus C contains a point p in neither A nor B. Then $\mathbf{H} := \langle \mathbf{K}, e(p) \rangle_V$ is a vector-space hyperplane contained in $\langle e(C) \rangle_V$, giving us the containment

$$A \cap B \subset S(\mathbf{H}) \subseteq C.$$

The first containment is proper so the existence of Velkamp points forces $C = S(\mathbf{H})$. Thus $C \in \mathcal{V}_e$. This proves that $\mathcal{V}_e$ is a subspace of the Veldkamp space.

3. For $s \leq r$ any subspace $\mathbf{U}$ of codimension s in V is the intersection of s independent vector-space hyperplanes, and so $S(\mathbf{U})$ is an intersection of s geometric hyperplanes which, because Veldkamp $(r-1)$ spaces exist, cannot be expressed as an intersection over a proper subcollection of these hyperplanes. Conversely, any subspace expressible as an intersection of s geometric hyperplanes with s minimal has its image spanning a codimension s subspace of V. Thus the two correspondences are inverses of each other.
4. It follows from part 3 that $\mathcal{V}_e$, as a Veldkamp subspace, is isomorphic to the incidence geometry of codimension-1 and codimension-2 subspaces of V. This, of course is the classical projective space derived from the dual space V^*.

The proof is complete. □

4.5 Teirlinck's Theory and Embedability into the Second Veldkamp Space

We begin with a vital matter

Lemma 4.5.1 *Assume the axioms (H1) and (H2) of Teirlinck's theory for the collection $\mathcal{H}$ of geometric hyperplanes of the geometry $\Gamma = (\mathcal{P}, \mathcal{L})$. Let $\mathcal{V}_\mathcal{H}$ be the associated Veldkamp space $(\mathcal{H}, \mathcal{H}_2)$. For each point $p \in \mathcal{P}$ let $V(p)$ be the set of all hyperplanes in $\mathcal{H}$ which contain the point p. Then $V(p)$ is a geometric hyperplane of the linear space $\mathcal{V}_\mathcal{H}(\Gamma) := \mathcal{V}_\mathcal{H}$.*

Proof Obviously, if two hyperplanes of $\mathcal{H}$ contain point p, then so do all hyperplanes of $\mathcal{H}$ which contain their intersection, so $V(p)$ is a subspace of $\mathcal{V}_\mathcal{H}$. Given any Veldkamp line defined by an intersection $A \cap B$, Teirlinck's condition (H2) shows that at least one member of each Veldkamp line of $\mathcal{V}_\mathcal{H}$ contains p. The proof is complete. □

Now $\mathcal{V}_\mathcal{H}$ is a linear space. *Its* Veldkamp space is called the *second Veldkamp space of* Γ *and* $\mathcal{H}$ and is denoted $\mathcal{V}(\mathcal{V}_\mathcal{H})$ or simply $\mathcal{V}^*_\mathcal{H}$.

Example 2 If you are lucky and $\mathcal{V}_\mathcal{H}$ is a projective space, then $\mathcal{V}^*_\mathcal{H}$ is the ordinary dual of a projective space and so is itself a projective space. But if $\mathcal{V}_\mathcal{H}$ is something-weird – say a strange Steiner triple system– then its Veldkamp space might not have many elements. Indeed, an affine space (to be discussed in Chap. 6) has no geometric hyperplanes whatsoever (see the first footnote of this chapter).

However, since we have the axioms of Teirlinck's theory at work here, Lemma 4.5.1 provides insurance that we have plenty of points to this second Veldkamp space of Γ and $\mathcal{H}$. It is time to look at the line structure of this space.

Since the original Veldkamp space $\mathcal{V}_{\mathcal{H}}$ is a linear space, Theorem 4.2.1 shows that Veldkamp r-spaces exists for $\mathcal{V}_{\mathcal{H}}$ for arbitrarily large r. This shows that the second Veldkamp space is itself a linear space, whose subspaces possess a "semi-modular-lattice-like structure" called a "matroid" (details of which appear in Chap. 6). In effect we are saying that the second Veldkamp space has a very strong structure. The problem is that usually the second Velkamp space is a very nice space, but it can be a very sparse structure that bears little relation to the original geometry Γ. However, in the presence of Teirlinck's theory, there is a very close connection between Γ and the second Veldkamp space– specifically when we replace (H1) by the point-seperation axiom (H4).

Theorem 4.5.2 (The embedding theorem of Teirlinck's theory) *Suppose $\mathcal{H}$ is a family of geometric hyperplanes of the point-line geometry $\Gamma = (\mathcal{P}, \mathcal{L})$ for which Veldkamp lines exist and Teirlincks' condition (Teir) hold – that is hypotheses (H3) and (H2) of this section are in force. Hypothesis (H1), that the elements of $\mathcal{H}$ eventually intersect at the empty set, is replaced by the stronger hypothesis (H4) asserting that the elements of $\mathcal{H}$ separate points. (In the language of the previous section, every point is $\mathcal{H}$-closed.)*

Then there is an injection

$$e_1 : \mathcal{P} \to \mathcal{V}^*_{\mathcal{H}}$$

whose image on the point-shadow of any line L is a full Veldkamp line of the space $\mathcal{V}^*_{\mathcal{H}}$. That is, in the language introduced earlier, there is an embedding $e : \Gamma \to S \leq V^*_{\mathcal{H}}$ of Γ onto a subgeometry S of the second Veldkamp space (another linear space).

Proof Our proposed mapping $e : \mathcal{P} \to \mathcal{V}^*_{\mathcal{H}}$ is that which sends each point p to the hyperplane $V(p)$ of the first Velkamp space $\mathcal{V}_{\mathcal{H}}$ (i.e., a "point" of the second Veldkamp space $\mathcal{V}^*_{\mathcal{H}}$). We must show that the induced mapping on any line of Γ, surjectively maps onto a Veldkamp line of $\mathcal{V}^*_{\mathcal{H}}$.

So suppose L is a line of the geometry Γ and a and b are distinct points of this line. Then $V(a) \cap V(b)$ is precisely the set of all hyperplanes of Γ which contain the line L (this uses nothing more than the fact that hyperplanes are subspaces). Now suppose $\mathcal{X}$ is a collection of hyperplanes of Γ which includes all hyperplanes of Γ containing L, but itself forms a geometric hyperplane of the linear space $\mathcal{V}_{\mathcal{H}}$. Thus

$$V(a) \cap V(b) \subseteq \mathcal{X}, \text{ a hyperplane of } \mathcal{V}_{\mathcal{H}}.$$

Since $\mathcal{V}_{\mathcal{H}}$ is a linear space, Veldkamp lines exist for it, and so $V(a) \cap V(b)$ cannot itself be a hyperplane. Thus $V(a) \cap V(b)$ is properly contained in $\mathcal{X}$. So there is a hyperplane X of Γ in $\mathcal{X}$ not in $V(a) \cap V(b)$. That means X does not contain

line L. But being a hyperplane of Γ it must intersect L at a unique point z. Thus $V(a) \cap V(b) \cup \{X\}$ lies in $V(z)$.

Now $V(a) \cap V(b)$ is a geometric hyperplane of $V(z)$, a hyperplane subspace of a linear space, and so is a maximal subspace of $\mathcal{V}_{\mathcal{H}}$. Since Veldkamp lines exist for the second Veldkamp space, $\mathcal{X} \cap V(a) = V(a) \cap V(b)$ and so $V(a) \cap V(b)$ is also a maximal subspace of $\mathcal{X}$. Thus the subspace $\mathcal{Y} := \langle X, V(a) \cap V(b) \rangle_{\mathcal{V}^*}$ of the second Veldkamp space generated by $V(a) \cap V(b)$ and X, properly contains $V(a) \cap V(b)$, but is in both $V(z)$ and $\mathcal{X}$. But Veldkamp lines exist for $\mathcal{V}_{\mathcal{H}}$, and so $V(z) = \mathcal{Y} = \mathcal{X}$. The proof is complete. □

Corollary 4.5.3 *Assume the hypotheses (H2), (H3), and (H4) of Theorem 4.5.2. Suppose for some reason that the Veldkamp space $\mathcal{V}_{\mathcal{H}}$ was a projective space. Then there is a projective embedding $e : \Gamma \to \mathcal{V}^{**}$, the subspace of $\mathcal{V}^*_{\mathcal{H}}$ spanned by the $V(p)$ as p ranges over the points $p \in \mathcal{P}$.*

Proof Immediate from the preceding Theorem 4.5.2. All one needed was an embedding into a projective space. □

4.6 Exercises for Chapter 4

4.6.1 Exercises on Geometric Hyperplanes

4.1 Consider a classical projective space $P = (\mathcal{P}, \mathcal{L})$ obtained from a vector space V of finite dimension n over a division ring D. Recall from Example 1 of Chap. 3, p. 73, that then the points $\mathcal{P}$ are the 1-subspaces of V, and the lines $\mathcal{L}$ are the 2-subspaces of V – there are no further hypothesized objects. Now let H be a vector subspace of V of codimension 1. The point-line truncation of its residue in the full projective geometry of spaces of all types would be the 1-subspaces and 2-subspaces which lie in H. Show that this residue is clearly a (point-line) subspace of the projective space characterized by the fact that the point-shadow of every 2-space meets it in at least one 1-space. That means that $sh_{\mathcal{P}}(H)$ – the set of projective points of H – is a geometric hyperplane of the projective space. Since these are the only sort of geometric hyperplanes that exist in a projective space, we call them *projective hyperplanes*.

4.2 Consider the polar space of rank three of Example 5 of Chap. 2, p. 49. The points are the singular 1-spaces with respect to the form $Q : \mathbf{R}^{(6)} \to \mathbf{R}$: the lines are the singular 2-subspaces. Then this geometry is easily seen to be a gamma space. Fix a totally singular 1-space $x = \langle v \rangle$. Then show that $x^{\perp}$ (the vertices of the point-collinearity graph of Γ which are equal to or adjacent to x) is a geometric hyperplane of $(\mathcal{P}, \mathcal{L})$.

4.3 Consider the generalized quadrangle $(\mathcal{P}, \mathcal{L})$ of order $(2, 2)$ in Example 7, Chap. 3, p. 75. The 15 points are the set $\mathcal{P}$ of all 2-subsets of a fixed 6-set Ω.

The lines are the 15 $(2-2-2)$-partitions of Ω. Consider the collection O of all 2-subsets of $\Omega = \{1, 2, \ldots, 6\}$ which contain the letter "1." Then O is a set of $1 + 2^2 = 5$ pairwise non-collinear points.

1. Show that O is a maximal coclique of the point-collinear graph $(\mathcal{L}, \sim)$.
2. Show that O is a geometric hyperplane of $(\mathcal{P}, \mathcal{L})$. (It follows that every point not in O is collinear with exactly three points of O.)

4.4 Also the quadrangle of the preceding exercise contains subspaces which are grids. For example, one may partition the six letters of Ω into two components of size 3 – for example

$$\Omega = A + B = \{1, 2, 3\} + \{4, 5, 6\}.$$

Then we let S be the nine 2-sets $(a, b) \in A \times B$.

1. Show that the set S can have its points are arranged in Fig. 4.2 to exhibit its structure of internal lines.
2. Show that S is a geometric hyperplane. [Hint: Show that the internal and tangent lines already exhaust all 15 lines.]

4.5 Let us generalize Example 3 of Chap. 3. Suppose V is a right vector space over a division ring D, of dimension $n + 1 > d + 1 > 2$, for some integer d. Let $\Gamma = (\mathcal{P}, \mathcal{L})$ be the truncation of $PG(V)$ to (V_d, V_{d+1}), the d- and $(d + 1)$-dimensional vector subspaces of V. (We call this a "half Grassmann space.") Let H be any subspace of dimension n – that is, a (vector space) hyperplane of V. Let S be the set of all d-subspaces of V which lie in H. Then the student can show the following:

1. $(\mathcal{P}, \mathcal{L}) = (V_d, V_{d+1})$ is a partial linear space
2. S is a geometric hyperplane of $(\mathcal{P}, \mathcal{L})$

4.6 Let Γ be the classical polar space of (V, f) where V is a vector space over a division ring D and $f : V \times V \to D$ is a non-degenerate reflexive sesquilinear form. (See Example 5 of Chap. 3, p. 74.) A subspace U of V is said to be *totally isotropic with respect to the form* f if and only if $f(U, U) = 0$. The points of Γ are the isotropic 1-subspaces of V, and the lines of Γ are the totally f-isotropic 2-subspaces of V.

1. Given two polar points U and W (these are isotropic 1-subspaces), show that U is collinear with W in Γ if and only if $f(U, W) = 0$.

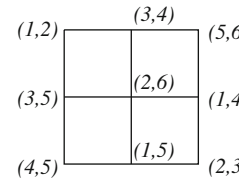

Fig. 4.2 A grid subspace of Example 7, Chap. 3, p. 75

2. If U is a polar point, show that $U^{\perp}$, the collection of all polar points collinear with U including U itself, is a geometric hyperplane of U.
3. Suppose H is a vector space hyperplane of V – that is, a vector subspace of codimension 1. Let $\Gamma(H)$ be the collection of all polar points which, in their guise as isotropic 1-subspaces of V, are vector subspaces of H. Show that $\Gamma(H)$ is a geometric hyperplane of Γ.

4.7 Let Γ be the classical Grassmann space $A_{n,d}(D)$ of Example 6 of Chap. 3, p. 75. The points of Γ are the d-dimensional subspaces of an $(n+1)$-dimensional vector space V over a division ring D. The lines are the $(d-1, d+1)$-dimensional flags (A, B). Fix a d-dimension vector subspace X of V. Let $H(X)$ be the collection of all d-subspaces of V which intersect X non-trivially – that is, the vector-space intersection with X is not the 0-subspace. Then $H(X)$ is a subset of $\mathcal{P}$, the points of Γ.

1. Show that $H(X)$ is a subspace of V.
2. Show that $H(X)$ is a geometric hyperplane of V.
3. Recall from Exercise 5 at the end of Chap. 3, that the distance between two points A and B of Γ in the collinearity graph $\Delta = (\mathcal{P}, \sim)$ is the codimension of $A \cap B$ in either A or B. Conclude that if $n+1 \geq 2d$ then as a graph, Δ has diameter d.
4. Conclude that if $n+1 \geq d$, then $H(X) = \Delta_{d-1}(X)$, the sphere of radius $d-1$ about vertex X in the graph Δ.

Remark The reader will encounter many other instances in which the proper spheres of maximal radius $\Delta_{d-1}(x)$ about a vertex x in the point-collinearity graph of a geometry turn out to be geometric hyperplanes. This happens in all polar spaces, dual polar spaces, the even-parity half-spin geometries and in all of the hexagonic ("long root") geometries of Chap. 17 of this book.

4.6.2 Exercises Relating to Veldkamp Spaces

4.8 Prove Lemmas 4.1.1 and 4.1.3 in detail.

4.9 This problem is intended as a challenge to the beginner student. Consider once more the generalized quadrangle of Example 7 of Chap. 3, p. 75, whose points were the 15 2-subsets of $\Omega = \{1, 2, 3, 4, 5, 6\}$, and whose lines were the 15 partitions of the elements of Ω into three disjoint pairs. Prove the following statements:

1. There are 15 hyperplanes of shape $p^{\perp}$, where p is a point.
2. There are 6 cocliques of size five in the collinearity graph (these are called *ovoids*) and these must be hyperplanes.
3. There are 10 subspace grids (known to be hyperplanes from Exercise 4).

4. All hyperplanes of this quadrangle are one of these three types (the more difficult part of the exercise), so there are $2^5 - 1 = 15 + 6 + 10$ geometric hyperplanes in all.

4.6.3 Exercises Relating to Teirlinck's Axioms

4.10 Prove Lemma 4.2.9.

4.11 Suppose $\mathcal{H}$ is a family of geometric hyperplanes of a point-line geometry $\Gamma = (\mathcal{P}, \mathcal{L})$ satisfying the axioms (H1)–(H3). Show that if Γ is a gamma space, then Γ is a partial linear space. [Hint: In a gamma space, two lines intersecting at two distinct points must lie in a common singular subspace.]

4.12 Again suppose $(\Gamma, \mathcal{H})$ satisfies the hypotheses (H1)–(H3) of Teirlinck's theory. Let $\phi : \Gamma \to \bar{\Gamma}$ be the morphism discussed in this section. Suppose Γ satisfies this property:

(H5) If $a, b \in p^{\perp} - \{p\}$ then either $a = b$ or $[a] \neq [b]$.

Show that the morphism between the point-collinearity graphs Γ and of $\bar{\Gamma}$ induced by the geometry morphism ϕ, is locally bijective.

4.13 Show that the extra hypothesis (H5) of the previous exercise implies Γ is a partial linear space.

Part II
The Classical Geometries

Chapter 5
Projective Planes

Abstract Translation planes, the Klein correspondence, and collineations of projective planes prepare one for the goal of this chapter: a proof of the fact that planes admitting every possible collineation are classical (that is Desarguesian). As a corollary, any projective plane properly embedded in a larger projective space is classical. The treatment is coordinate-free.

5.1 Introduction

From a strictly axiomatic point of view, the simplest rank two geometry seems to be the projective plane. It forms one of the basic "building blocks" of most of the Lie incidence geometries – the central focus of this book.

The known projective planes come in three types: (1) the canonical class, the Desarguesian (or classical) projective planes which are parameterized by the class of division rings, (2) the strictly larger class of Moufang planes, and (3) the bewildering jungle of non-Moufang projective planes. In this book, we shall be interested in those planes which play a role in higher rank classical geometries – in particular as they are presented in the projective spaces and the non-degenerate polar spaces. A projective plane that can be embedded as a proper subspace of a larger projective subspace will be seen to be classical (Chap. 6). A projective plane that is a maximal singular subspace of a rank three polar space is necessarily Moufang (Chap. 8).

Moufang planes are parameterized by alternative division rings, which have been classified up to division rings.[1]

There is convincing evidence that the remaining planes cannot be classified by a finite set of parameters. There has never been a question of classifying projective planes having infinitely many points. Such an endeavor appears hopeless since there are so-called "free constructions" which allow one to embed[2] weakly an arbitrary

[1] The Bruck-Kleinfeld-Skornyakov theorem.

[2] A weak embedding is quite different from the embeddings introduced in the previous chapter. In a weak embedding, the embedded image points of a point shadow of a line of the domain geometry need not be a full point-shadow of the line into which this set is embedded. But nevertheless this image of a line shadow must lie in some unique projective line.

E.E. Shult, *Points and Lines*, Universitext, DOI 10.1007/978-3-642-15627-4_5,

© Springer-Verlag Berlin Heidelberg 2011

partial linear geometry $(\mathcal{P}, \mathcal{L})$ into a projective plane. This not an embedding as a subspace: it is an iterative process in which new lines and new points are added at each stage along with new incidences so that (unless $(\mathcal{P}, \mathcal{L})$ is already embedded in a finite plane) the shadows of preexisting points and lines progressively grow larger and larger. Perhaps the only interesting thing about that side of the subject is how "free" it really is. One can fix it so that all sorts of desired properties are carried along with the construction – for example, one can even arrange the free construction so that the group of automorphisms of the final plane is doubly transitive on points!

On the other hand, things might be different in the finite world. Indeed, there was at one time a feeling that it might be possible to sift one's way through the world of finite non-Desarguesian planes toward an actual classification. There were two reasons encouraging this optimism. (1) First, there was the general argument that in the finite case, a plane is a very tight structure – at least it is very difficult to be a projective plane. You have to have $n + 1$ pairwise orthogonal latin squares and there are theorems, such as the Bruck-Ryser theorem [17], which rule out many values of n as being able to be the order of a finite projective plane. (2) A second reason was the observation that very sharp consequences follow from the existence of certain automorphisms of finite planes.[3] The proposed Lenz-Barlotti classification of planes used as a paradigm, the various ways a group of automorphisms could act on a plane, and the scheme remains today as "archeological" evidence that there indeed was once a guarded optimism that a classification of finite planes was possible (see [57, 76]). The optimism for a classification of all finite planes has since faded. Even translation planes now seem unreachable, and at present it is difficult to detect an enthusiasm for estimating the ubiquity of finite planes whose automorphism group has order 1.

5.2 The Axioms for a Projective Plane

A *projective plane* is a point-line geometry $(\mathcal{P}, \mathcal{L})$ satisfying these axioms:

(PP1) Any two distinct points are incident with exactly one line
(PP2) Any two distinct lines are incident with exactly one point
(PP3) There exist at least two distinct lines each of which is incident with at least three points

Let us review these axioms in the language of Chap. 3 on point-line geometries. Axiom (PP1) says that $(\mathcal{P}, \mathcal{L})$ is a *linear space*. Together, axioms (PP1) and (PP2)

[3] For translation planes, this optimism may have been encouraged by the theorems of Wedderburn and Artin-Zorn asserting that finite alternative division rings were fields.

assert that both $(\mathcal{P}, \mathcal{L})$ and its dual $(\mathcal{L}, \mathcal{P})$ are linear spaces.[4] Finally, (PP3) provides the existence of at least two thick lines.

NOTATION: As with any partial linear space, we will adopt the convention of denoting the unique line incident with two distinct collinear points x and y by the symbol xy.

We claim that if $(\mathcal{P}, \mathcal{L})$ is a projective plane, then so is the dual structure $(\mathcal{L}, \mathcal{P})$. In trying to view $(\mathcal{L}, \mathcal{P})$ as a point-line geometry, every statement about $(\mathcal{P}, \mathcal{L})$, can be applied to $(\mathcal{L}, \mathcal{P})$ to yield a new statement in which the words "point" and "line" have been transposed, and the words "incident" or "incidence" are left alone. (Of course, since we are dealing with a partial linear space one might identify a line with its point shadow, so that incidence might be expressed by saying that "a line contains a point." In the dual statement, the word "contains" would then become the phrase "is contained in.")

From this point of view, one sees that the "dual" of statement (PP1) is (PP2) and the "dual" of (PP2) is (PP1). The dual of statement (PP3) is the assertion that there are two distinct points, x and y, say, each of which is incident with at least three distinct lines. Assuming both (PP1) and (PP2), this dual statement means that there are at least two distinct lines A and B on point y which do not contain point x. Now each of the three lines on x intersect (as point-shadows) three distinct points of A. (For if two of these three points were equal, (PP2) would fail.) Similarly, line B contains at least three points; and the existence of thick lines A and B implies (PP3).

Writing P' for the dual of proposition P, we have shown the implications

$$\begin{gathered}(PP1)' \Leftrightarrow (PP2),\\ (PP2)' \Leftrightarrow (PP1),\\ (PP1) + (PP2) + (PP3)' \Leftrightarrow (PP1) + (PP2) + (PP3).\end{gathered}$$

Thus the following is reached.

Lemma 5.2.1 *A point-line geometry $(\mathcal{P}, \mathcal{L})$ is a projective plane if and only if its dual point-line geometry $(\mathcal{L}, \mathcal{P})$ is also a projective plane.*

Any plane with either a finite number of points or a finite number of lines is called a *finite plane*.

Lemma 5.2.2 *For every finite plane $\Pi = (\mathcal{P}, \mathcal{L})$, there exists an integer n (called the order of the plane) such that:*

1. *Every point lies on $n + 1$ lines*
2. *Every line contains $n + 1$ points*
3. *There are in all exactly $n^2 + n + 1$ points*
4. *There are in all exactly $n^2 + n + 1$ lines*

[4] Recall that a point-line geometry is *firm* if and only if every line is incident with at least two points, and each point is incident with at least two lines.) A firm point-line geometry with these two axioms, (PP1) and (PP2), is called a *generalized projective plane*. (This concept was formally introduced in Chap. 4, p. 87.)

Proof This is fairly elementary. Pick a point p and let $n+1$ be the finite number of lines on p. Since every point is on at least two lines, and each line has at least two points, there is, first of all, a point z distinct from p and a line on z not containing p. Let L be any such line not on p. Then each line on p meets L at a unique point y of L; conversely, any point y of L (by (PP1)) determines a unique line py containing $\{p, y\}$. Thus there is a bijection between the $n+1$ lines on p and the points of the line L. Specifically, there is a bijection of shadows

$$sh_{\mathcal{L}}(p) \to sh_{\mathcal{P}}(L) \tag{5.1}$$

for every non-incident point-line pair (p, L). Thus, to show that two lines A and B are incident with the same number of points, we need only produce a point p not in $A \cup B$ (strictly speaking, $sh_{\mathcal{P}}(A) \cup sh_{\mathcal{P}}(B)$, if we identify a line and its unique point-shadow). If that were impossible, then A and B would have to be the only two lines on their point of intersection $A \cap B = \{c\}$. Then by (5.1) every line not on c has only two points. But in that case A and B are the only two lines of $\mathcal{L}$ possessing more than two points and both must do so. In that case, there are distinct points a_1, a_2 in $A - \{c\}$ and distinct points b_1, b_2 in $B - \{c\}$ and lines $a_i b_i$ are thin with point-shadows $\{a_i, b_i\}$, $i = 1, 2$. Thus line $a_1 b_1$ does not intersect $a_2 b_2$ at a point, as required by (PP2).

Thus, we now know that, for any two lines, there is a point not in the union of their point-shadows. It follows from the bijection (5.1) that any two lines possess the same number of points. By duality, any two points lie on the same number of lines (we have assigned this number to be $n+1$) and so each line contains $n+1$ points. At this stage (1) and (2) hold.

Now, each point p is on $n+1$ lines each of which carries n further points. By (PP2) none of the n further points of one line on p can be one of the n further points contributed by a second line on p. Thus

$$|\mathcal{P}| = 1 + n(n+1) = 1 + n + n^2.$$

The argument for lines is the dual of this one. □

Example 1 (The Desarguesian Projective Planes.) This example was mentioned in Chap. 3, p. 73. Let V be a three-dimensional right vector space over a division ring K. Let $\mathcal{P}$ and $\mathcal{L}$ be the collections of all one-dimensional and two-dimensional vector subspaces of V, respectively. Then with incidence being proper containment among subspaces, $(\mathcal{P}, \mathcal{L})$ is one of the Desarguesian projective spaces discussed as an example in the previous chapter. However, since any two two-dimensional subspaces of V must intersect at a unique one-dimensional subspace, we have a linear space satisfying (PP2), and so it is a projective plane. This is the "classical" example.

Example 2 (Difference sets.) Let G be any multiplicative group, and let D be a subset of G. We are about to write an equation whose right side is a multiset of group elements. That is, it is an accountant's list of how many group elements of

each kind we have on hand. (These inventories, connected by plus signs might be thought of as elements of the group ring $\mathbf{Z}G$ with integer coefficients.) Let D^{-1} be the set of inverses of the elements of D. We say that D is a *difference set* if and only if

$$DD^{-1} = |D|1 + G^{\#},$$

where $G^{\#}$ is each non-identity element of G listed exactly once, and the first term on the right indicates that the identity element occurs exactly $|D|$ times, as dd^{-1} for each $d \in D$. Thus each non-identity element of G is uniquely expressible as $d_1 d_2^{-1}$, for some d_1 and d_2 in D.[5] Now we produce a point-line geometry $(\mathcal{P}, \mathcal{L}) = (G, \mathcal{L})$ where $\mathcal{L}$ is the collection of all right translates $\{Dg | g \in G\}$.

We leave to the student the proof that if D is a difference set of at least three elements in a group G, then $(G, \mathcal{L})$ satisfies the three axioms of a projective plane.[6]

We have already seen in Chap. 3 that $D = \{1, 2, 4\}$ is a difference set in the group $\mathbf{Z}/(7)$ and defines a projective plane of order 2, in the above sense for his additive group. Here are some more examples of difference sets:

1. $D = \{3, 6, 7, 12, 14\}$ for integers mod 21
2. $D = \{1, 2, 4, 8, 16, 32, 64, 55, 37\}$ for integers mod 73

No one so far has found a difference set for a group of order $n^2 + n + 1$ where n is not a prime power (though $(18)^2 + 18 + 1 = 7^3$ sounds suggestive). We know this for the larger (sociological) reason that no one has so far discovered a single finite projective plane whose order is not a prime power.

5.3 Affine Planes

An *affine plane* is a firm point-line geometry $A = (\mathcal{P}, \mathcal{L})$ satisfying these axioms:

(AP1) $(\mathcal{P}, \mathcal{L})$ is a linear space.
(AP2) Given a line L and any point p, not incident with L, there exists exactly one line on p which meets L at the empty set.

Let us say that two lines of an affine plane are *parallel* if and only if they intersect at the empty set or are the same identical line. It is then easy to see the following.

---

[5] Please note that if G were an abelian group, then we could naturally write this as an additive group, and our multiset equation would read

$$D + D = |D|0 + G^{\#}.$$

But now the "+" on the right should not be confused with group addition.

[6] In fact, at this point the student should be able to do the first four exercises at the end of this chapter.

Lemma 5.3.1 *The relation of being parallel is an equivalence relation on lines.*

Proof It is clear from the definition that the parallel relationship on lines is reflexive and symmetric. It remains only to show that it is transitive.

Suppose A, B, and C are lines with $A \cap B = \emptyset = B \cap C$ while $A \cap C = \{p\}$. Then p is a point, not on B incident with two lines which meet B at the empty set. This contradicts axiom (AP2), completing the proof. □

Thus the set of lines $\mathcal{L}$ partitions into equivalence classes under the parallel relation; these classes are called the *parallel classes*. For each line L, let $[L]$ denote the unique parallel class containing line L.

Example 3 (The Desarguesian affine plane.) Let $\mathcal{P} = F \times F$ where F is any division ring. Set $\mathcal{L}_\infty := \{\{a\} \times F | a \in F\}$, and for every ordered pair of scalars $(a, b) \in \mathcal{P}$, let $L_{a,b} := \{(x, y) \in \mathcal{P} | y = ax + b\}$ (the "graph of the equation of a line" in the language of "Freshman Algebra"). Set $\mathcal{L} = \mathcal{L}_\infty \cup \{L_{a,b} | (a, b) \in F \times F\}$. A line L is defined to be incident with a point p if and only if p is a member of L.

Example 4 (An equation-free description of the same example.) Let V be a two-dimensional right vector space over a division ring D. The elements of V are our points, that is, $V = \mathcal{P}$. Let V_1 be the collection of all 1-subspaces of V. Then the lines will be the collection of all cosets of these 1-subspaces, that is, $\mathcal{L} = \{W + v | W \in V_1, v \in V\}$.

Note that, in this example, the full collection V/W of cosets of a 1-subspace W, forms a parallel class of lines. The group $(V, +)$ induces a regular group of automorphisms of this plane stabilizing every parallel class.

Example 5 (The Moulton planes.) Now let $\mathbf{R}$ be the real field, and set $\mathcal{P} = \mathbf{R}^{(2)}$. Fix a real non-zero number r. The parallel classes of lines are as follows. First there is the class of vertical lines $\mathcal{L}_\infty$ – that is, the additive cosets in $V = \mathbf{R}^{(2)}$ of the y-axis $L = \{(0, y) | y \in \mathbf{R}\}$. Then for each real number a, there is a parallel-class consisting of the additive translates of the set

$$L_a = \{(y, ay) | y \geq 0\} \cup \{(y, ray) | y < 0\}.$$

These lines look like this. In the upper half-plane the line looks like a classic straight line of slope a (allowing $a = \infty$ in order to include the vertical lines in this discussion), but they continue below the x-intercept as a line of slope ra. That is, these lines are "bent" with a constant index of refraction upon entering the lower half-plane from above. It doesn't take much checking to see that this is indeed an affine plane. I think the reader might guess that this example depends in a particular way upon the fact that R is totally ordered. So don't expect anything like this for finite planes.

Example 6 (Some translation affine planes.) Now let V be a vector space over a field F of even dimension $2n$. Suppose we happen to have a collection $\mathcal{S}$ of n-dimensional spaces with these properties:

1. The n-spaces in $\mathcal{S}$ pairwise intersect at the 0-subspace. (That is, for every two members A and B of $\mathcal{S}$, one has $V = A \oplus B$.)
2. For every non-zero vector v in V, there exists a subspace $W \in \mathcal{S}$ containing it.

Such a collection of subspaces is called a *spread of half-dimensional subspaces of* V.[7] Now, given such a spread $\mathcal{S}$ of half-spaces of an even-dimensional space V, we can construct an affine plane as follows. Let $\mathcal{P}$ be the vectors of the vector space V – these are the points. We describe the lines as subsets of $\mathcal{P}$ so that incidence will need no separate description. A *line* is any additive coset $A + v$ of a subspace A belonging to the spread $\mathcal{S}$.

Let's check the axioms. This will be made easier once we observe that the additive group $(V, +)$ acts as a point-transitive group of automorphisms of this incidence system $(\mathcal{P}, \mathcal{L})$. In this way, all assertions about a point can be tested when that point is the zero vector.

Is $(\mathcal{P}, \mathcal{L})$ a linear space? We want to know if two distinct vectors are incident with a unique line. From the previous paragraph we may assume that one of the vectors is the zero vector 0. Then the result follows from the fact that the subspaces of $\mathcal{S}$ partition the non-zero vectors of V.

What about the parallel axiom? We imagine that we have a line, W, which we may suppose (via translation) to contain the 0-vector and so be a member of $\mathcal{S}$ and an additional vector v outside W. Obviously the coset $W + v$ is a line parallel to W. Suppose there were a second such line $S + v$ for $S \in \mathcal{S}$ on v. Since $S \neq W$, by hypothesis $V = S \oplus W$, so we can write $v = w + s$ where $w \in W$. Then $S + v = S + w$. Then, translating by w, we see that $(S + v) \cap W = \emptyset$ implies $S \cap W = \emptyset$, which is absurd since they both contain the zero vector.

5.4 The Relation Between Affine and Projective Planes

Let $\Pi = (\mathcal{P}, \mathcal{L})$ be any projective plane and fix a line $L \in \mathcal{L}$. As usual, we let the symbol $sh_{\mathcal{P}}(L)$ denote the *point-shadow of the line* L — that is, the full set of points which are incident with L. Now let $\mathcal{P}' = \mathcal{P} - sh_{\mathcal{P}}(L)$ and let $\mathcal{L}' := \mathcal{L} - \{L\}$. Then, by restricting the incidence relation to $\mathcal{P}' \times \mathcal{L}'$, we obtain a derived point-line geometry: $A := (\mathcal{P}', \mathcal{L}')$.

Now A is a linear space. Suppose $N \in \mathcal{L}'$ and p is a point of A not on N. Now consider any line M on p. If $M \cap N \cap \mathcal{P}' = \emptyset$, then the unique point of Π incident with both these lines is in L. This means that in Π, $M \cap N = L \cap N = q$ and so $M = pq$, a uniquely determined line on p. Conversely, if $\{q\} = L \cap N$, then pq is a line on p which in A is parallel to N. Thus the following is reached.

[7] There is a more general notion, in which a class of d-subspaces of V pairwise intersect at the 0-subspace, and whose union is all of V – called a *d-dimensional spread*. If such a spread exists for a finite vector space, one can show that d divides n. Does this divisibility property hold for the case of vector spaces of finite dimension n over infinite fields and division rings?

Lemma 5.4.1 *If L is a line of a projective plane, then the incidence relation induced on all the points not on L, and all lines distinct from L, is an affine plane.*

We shall always denote the affine plane obtained by removing line L from the projective plane π by the symbol $Af_\pi(L)$, or by $Af(L)$ if the plane π is understood.

This process is completely reversible. Suppose $A = (\mathcal{P}, \mathcal{L})$ is an affine plane. Let L_∞ be the set of all parallel classes of lines. We now invent a new incidence structure, $\Pi = (\mathcal{P}', \mathcal{L}')$ as follows:

1. $\mathcal{P}'$ consists of two sorts of points:
   (a) $\mathcal{P}'_1 = \mathcal{P}$, the old points of A (the "affine" points)
   (b) $\mathcal{P}'_2 :=$ the elements of L_∞, the parallel classes of A[8]
2. The lines are also of two types:
   (a) $\mathcal{L}'_1 := \mathcal{L}$, the old lines of A
   (b) $\mathcal{L}'_2$ consisting of the single new line L_∞

So far $\mathcal{P}' = \mathcal{P}'_1 \cup \mathcal{P}'_2$ and $\mathcal{L}' = \mathcal{L}'_1 \cup \mathcal{L}'_2$. We must define incidence.

1. An affine point p is incident with a line of $\mathcal{L}'_1$ if and only if they formed an incident pair in A.
2. A "point at infinity" – that is, a parallel class $[L]$ – is incident with all of the lines of A which belong to that parallel class.
3. A "point at infinity," $[L]$ is always incident with the "line at infinity," L_∞.

The incidence structure $\mathbf{P}(A) := (\mathcal{P}', \mathcal{L}')$ produced in this way is called the *projective plane determined by the affine plane A*.

Lemma 5.4.2 *Given an affine plane A, the enlarged incidence system $\mathbf{P}(A)$ is a projective plane.*

5.5 Translation Planes

5.5.1 Introduction

Historically the algebraic structures in this category have arisen via a coordinating structure for an affine plane A. For translation planes that structure turns out to be a quasifield. This creature is an additive group $(L, +)$ admitting a second binary operation "$\circ$," for which left and right multiplication by non-zero elements are surjective on the non-zero elements $N = L - \{0\}$ and which possesses a two-sided identity (that is, $(N, \circ)$) is a loop); moreover, "$\circ$" must be right distributive with respect to addition. Historically this led to the study of several species of quasifields. We list the most prominent ones:

[8] In this setting, these points are traditionally called the "points at infinity," more to confuse people than anything.

- A quasifield for which the multiplication "$\circ$" is associative is called a *nearfield*.
- A quasifield with the left distributive law is called a *semifield*.
- A *skew field* (also called *division ring* in this book) is the intersection of the above two classes: both distributive laws hold and multiplication is associative.
- An *alternative division ring* is a semifield with a very weak version of the associative law: for any elements a and b of the semifield, one has:

$$a(ab) = (aa)b,$$
$$(ab)b = a(bb).$$

There is a lot of interesting algebra here, and an account of how these coordinate systems arise from properties of the associated projective planes can be found in the classic book *Projective planes* by Hughes and Piper [76] and in the beautiful book *Projective Geometry* by Casse [28].

The two cases of interest for us are the planes coordinated by division rings (the classical, or Desarguesian planes), and the planes coordinated by alternative division rings (the Moufang planes). The alternative division rings have been classified up to division rings (the Bruck-Kleinfeld-Skornyakov theorem). As mentioned at the beginning of this chapter, these two cases arise for projective planes that are embedded in larger projective spaces or are maximal singular subspaces of rank three polar spaces, respectively.

We discuss the general quasifield case (which represents the general translation affine plane) only for the purpose of convincing the reader that these probably cannot be classified. To avoid a lot of messy equations our approach uses a little abstract algebra to reach the setting of spreads.

5.5.2 Translation Planes from Groups

Let us begin with a group G with a very special family of subgroups $\mathcal{S}$ having these properties:

(T1) Every element of G lies in some member of $\mathcal{S}$.
(T2) The intersection of any two distinct members of $\mathcal{S}$ is the identity subgroup.
(T3) If X and Y are any two distinct members of $\mathcal{S}$, then $G = XY$.[9]

The first two conditions simply require that the family of subgroups $\mathcal{S}$ induces a partition of the set of non-identity elements of G.

[9] In group theory, one says that a group G "factorizes as $G = XY$" if and only if every element of G can be written as a product xy of two elements, where x is in the set X and y is in the set Y. If the sets X and Y are subgroups of G, then $G = YX$, as well. Moreover, if X and Y are subgroups, and $X \cap Y = 1$, as assumed in (T2), then each element $g \in G$ can be written as a product of an element of X and an element of Y as $g = xy$ in a unique way.

From this structure one may construct a point-line geometry $\mathcal{A}(G, \mathcal{S}) := (\mathcal{P}, \mathcal{L})$ as follows. Let $\mathcal{P}$ be the set of elements of G, and let $\mathcal{L}$ be the collection of all right cosets in G of the subgroups in $\mathcal{S}$ – that is $\mathcal{L} = \{Xg | (X, g) \in \mathcal{S} \times G\}$. We say that g is *incident* with a line Xa if and only if g is a member of the coset Xa.

Lemma 5.5.1 *Let $\mathcal{A} = \mathcal{A}(G, \mathcal{S})$, the point-line geometry just constructed.*

1. *$\mathcal{A}$ is a linear space.*
2. *Given a point g not in line Xa, the coset Xg is the unique line through g which intersects Xa at the empty set.*
3. *$\mathcal{A}$ is an affine plane.*

Proof Part 1. Suppose x and y are distinct elements of G, that is, distinct points of $\mathcal{P}$. By (T1) xy^{-1} lies in a unique subgroup Z of $\mathcal{S}$. Then x and y both belong to the line Zy. Suppose, conversely, both x and y were points of a line Ub, a coset of some subgroup U of $\mathcal{S}$. Then xy^{-1} would lie in U, whence $U = Z$ and $Ub = Zx = Zy$ is the *only* line containing both x and y.

Part 2. Suppose g is a point not in the line Xa. Then Xg is a line containing g which intersects Xa trivially. Suppose, by way of contradiction, that Zg is a second line containing g that does not intersect Xa. Then the subgroup Z is distinct from X, so that $G = ZX$ and we may write $ga^{-1} = zx_1$ where $z \in Z$ and $x_1 \in X$. Then $Zga^{-1} = Zx_1$, which intersects X at x_1. Thus Zg intersects Xa at x_1a, contradicting the supposition. Thus Xg is the only line through g which is "parallel" to Xa.

Part 3. Parts 1 and 2 have established the axioms (AP1) and (AP2) which define an affine plane. The proof is complete. □

At this point we formally define an *affine translation plane* to be any plane of the form $\mathcal{A} = \mathcal{A}(G, \mathcal{S})$, where the pair $(G, \mathcal{S})$ satisfies the axioms (T1)–(T3), heading this subsection.

We now see that a typical parallel class of lines in the affine plane $\mathcal{A} = \mathcal{A}(G, \mathcal{S})$ is a full set G/X of right cosets of one of the subgroups X of $\mathcal{S}$. We next observe that G itself acts on $\mathcal{A}$ by right multiplication. That is, for each element g of G the mappings

$$x \to xg \text{ for all } x \in G$$

and

$$Xa \to Xag \text{ for all cosets } Xa, X \in \mathcal{S}$$

define an automorphism α_g of $\mathcal{A}$ which stabilizes each parallel class of lines, and, except for the identity element, acts freely on points.

Lemma 5.5.2 (Baer.) *Suppose $\mathcal{S}$ is a family of subgroups of the group G, satisfying the three axioms (T1)–(T3) above. Then G is an abelian group. All non-identity elements of G have infinite order, or else they all have prime order, for the same prime p.*

Proof Let $\mathcal{A} = \mathcal{A}(G, \mathcal{S})$ be the affine plane constructed as in Lemma 5.5.1 above, and let G act on $\mathcal{A}$ by right multiplication so that G is a point-regular group of automorphisms of $\mathcal{A}$ which stabilizes each parallel class.

Step 1. If g is a non-identity element of G, and the elements x and xg (regarded as points of $\mathcal{A}$) lie on a line Uh of $\mathcal{L}$, then α_g stabilizes Uh – that is $Uhg = Uh$.

Since $x \in Uh$ we see that the point xg is in Uhg as well as Uh. Since Uh and Uhg are parallel, we must have $Uh = Uhg$. Thus Uh is an α_g-invariant line.

Step 2. We claim that if X and Y are distinct subgroups of $\mathcal{S}$, then $[X, Y] = \{1\}$, that is, every element of X commutes with every element of Y.

Choose non-identity elements x and y in X and Y, respectively. We then have

$$\{1, x\} \subseteq X, \{1, y\} \subseteq Y, \{y, xy\} \subseteq Xy, \text{ and } \{x, yx\} \subseteq Yx.$$

Now by Step 1 the line on y and yx is α_x-invariant. Since α_x fixes no point, this line cannot intersect the α_x-invariant line X and so is parallel to X. Thus, the line on y and yx is Xy. Similarly, x and xy are conjugate under the right action of y, and so the line through these two points is parallel to Y and must be the coset Yx. So we have shown

$$yx \in Xy \text{ and } xy \in Yx.$$

But from the inclusions above, $yx \in Yx$ and $xy \in Xy$, so yx and xy both lie in the intersection $Xy \cap Yx$, and so are the same point.

Step 3. G is abelian.

Fix $X \in \mathcal{S}$. By Step 2, for any $Y \in \mathcal{S} - \{X\}$, Y lies in the centralizer $C_G(X)$. This means $G = \langle G - X \rangle = C_G(X)$, so each subgroup X is in the center of G. Since X was arbitrarily chosen, G is generated by elements of its center, and so is abelian.

Step 4. All non-identity elements of G possess the same order, which is either infinite or a prime number p.

Suppose G contains a non-identity element of finite order. Then, by raising that element to the appropriate power, we obtain a non-identity element g of some prime order p. Let X be the unique subgroup of $\mathcal{S}$ containing g. Select an arbitrary element y not in X and choose Y so $y \in Y \in \mathcal{S}$. Then xy lies in a third member of $\mathcal{S}$, say Z, which is distinct from X and Y. Then $(xy)^p = x^p y^p = y^p$, since G is abelian. Comparing the first and last members of the equation it is clear that this element lies in Z and at the same time lies in Y, and so is the identity element. Thus $y^p = 1$ for all $y \in G - X$. Applying the same argument with the roles of X and Y transposed we see that every non-identity element of X also has order p. The proof is complete. □

5.5.3 Endomorphism Rings

A little time out for algebra. Let A be an arbitrary abelian group (written multiplicatively). An *endomorphism* of A is simply a group homomorphism $A \to A$.

If α and β are endomorphisms of A we define $(\alpha \circ \beta)$ as the composition of the two endomorphisms (with α applied first — that is the mapping which sends the group element g to $(g^{\alpha})^{\beta}$). One defines a "sum" $\alpha + \beta$ as the mapping which sends group element g to $g^{\alpha} g^{\beta}$ (where the juxtaposition of the two elements indicates the group operation). Since G is abelian, it is easy to see that $\alpha + \beta$ is also a group endomorphism. In fact, with respect to the operations "$\circ$" and "$+$", the set $End(A)$ of all endomorphisms of A satisfy the axioms of a ring. This ring $(End(A), \circ, +)$ is called the *endomorphism ring of* A.

The identity automorphism $id : A \to A$ is the multiplicative identity element (the "one" in the ring of endomorphisms) and the trivial endomorphism $\theta : A \to \{1\}$ is the additive identity (the "zero" of the ring of endomorphisms).

5.5.4 The Kernel of $(G, \mathcal{S})$

Once again consider a pair $(G, \mathcal{S})$ satisfying the axioms (T1), (T2), and (T3). Since G is abelian, we may form the endomorphism ring $E := End(G)$, as described in the previous subsection. The *kernel* of $(G, \mathcal{S})$ is the subset of E consisting of all endomorhisms of G which stabilize each subgroup belonging to $\mathcal{S}$.

Theorem 5.5.3 *Let $(G, \mathcal{S})$, E, and K be as defined in the previous paragraph. Then the following statements hold:*

1. *The kernel K is a subring of the endomorphism ring E.*
2. *In fact K is a division ring.*[10]

Proof Suppose α and β are endormorphisms in K – that is X^{α} and X^{β} are subgroups of X for each $X \in \mathcal{S}$. Then, for any such X,

$$X^{\alpha\circ\beta} = (X^{\alpha})^{\beta} \subseteq (X^{\beta}) \subseteq X,$$
$$X^{\alpha+\beta} = \{x^{\alpha} x^{\beta} | x \in X\} \subseteq X.$$

Of course the identity automorphism also stabilizes each member of $\mathcal{S}$, and so K is a subring of E. Thus conclusion 1 holds.

To prove statement 2, we must show that every endomorphism in K which is not the "zero morphism" $\theta : A \to \{1\}$ has a multiplicative inverse – that is to say, it is an automorphism of G. So, consider an endomorphism $\alpha \in K$ distinct from θ.

[10] A warning to the reader: In the book *Projective Planes* by Hughes and Piper [76], the word "division ring" means something more general – a quasifield with the right distributive law. Here we follow the standard terminology of abstract algebra (see Jacobson [79]): "division ring" is synonymous with "skew field."

First we show that α is injective. Suppose otherwise. Then as α is a group homomorphism, its kernel,[11] $\ker\alpha = \{g \in G | g^\alpha = 1 \in G\}$ contains a non-identity element g which, of course, must lie in a unique subgroup U of $\mathcal{S}$. Choose any $h \in G - U$. Then h lies in a subgroup V of $\mathcal{S}$, while gh lies in a subgroup W of $\mathcal{S}$, the three subgroups, U, V, and W being pairwise distinct. Then $(gh)^\alpha = h^\alpha$ lies in $W \cap U = 1$. Thus $h^\alpha = 1$ so $h \in \ker\alpha$. Considering the wide range of h, we must have

$$\ker\alpha \supseteq \langle G - U\rangle = G.$$

But that would imply $\alpha = \theta$ contrary to our choice of θ. Thus α is indeed injective.

To show that α is an automorphism of G, it remains only to show that in fact α is surjective. For any element g of G, we must show that there exists a group element g_1 such that $g_1^\alpha = g$. Since $1^\alpha = 1$, we may assume g is a non-identity element of G, so g lies in a unique subgroup U of $\mathcal{S}$. Select any $h \in G - U$, so h lies in another subgroup $V \in \mathcal{S}$ distinct from U.

If $h^\alpha g = 1$, then $h^\alpha = g^{-1}$ lies in $U \cap V = \{1\}$, so $h \in \ker\alpha$ forcing $\alpha = \theta$ by the conclusion of the previous paragraph. Since α was chosen not to be θ at the outset, this is impossible. Thus we must have $h^\alpha g \neq 1$, so $h^\alpha g$ lies in a unique subgroup $W \in \mathcal{S}$ which is distinct from U. Thus the line Uh of the affine plane $\mathcal{A} = \mathcal{A}(G, \mathcal{S})$ is not parallel to W, and so there exists a unique affine point x in $Uh \cap W$. Then x (as an element of Uh) must have the form $x = g_1 h$, for some element g_1 in U. Now

$$x^\alpha = g_1^\alpha h^\alpha \in Uh^\alpha \cap W^\alpha = Uh^\alpha \cap W.$$

But also $h^\alpha g = gh^\alpha$ lies in W as well as Uh^α (note the use of the commutative law for G). Since only one element lies in the intersection $Uh^\alpha \cap W$, we have

$$x^\alpha = g_1^\alpha h^\alpha = h^\alpha g = gh^\alpha.$$

It follows that $g_1^\alpha = g$, as desired. Thus α is surjective.

Thus α is an automorphism of G which induces an automorphism on each subgroup of $\mathcal{S}$. It follows that its inverse mapping α^{-1} does the same thing. Thus $\alpha^{-1} \in K$. The proof is complete. □

A multiplicatively written abelian group is said to be *divisible*[12] if, for every positive integer k and non-identity group element g, there exists a group element h such that $h^k = g$.

---

[11] This is the kernel of a homomorphism, not to be confused with the kernel K of the configuration $(G, \mathcal{S})$. The history that produced these two meanings for one word could not have the foresight that one day they might be in conflict in the same proof.

[12] See Marshall Hall's book *Theory of groups*, p. 197, where it is defined for additively written groups.

Corollary 5.5.4 *If every element of G has infinite order, then the abelian group G is a divisible group.*

Proof One notes that for each positive integer k the "k-th power map" which sends each element of G to its k-th power is an endomorphism of G stabilizing each subgroup of $\mathcal{S}$. But by the previous theorem such a mapping is surjective. □

At this stage we see that the abelian group G is a right vector space V over the division ring K, and that $\mathcal{S}$ is a *spread* of K-subspaces, with the property that $V = A \oplus B$, for any two distinct members A and B of $\mathcal{S}$. If $C \in \mathcal{S} - \{A, B\}$, then $V = A \oplus C = B \oplus C$. Thus we have a chain of isomorphisms of K-modules:

$$A \simeq V/C \simeq B.$$

In particular, all of the subspaces in $\mathcal{S}$ possess the same dimension over K. If this dimension is a finite number n, then $n = 2d$ where d is the dimension of the subspaces is $\mathcal{S}$. The reader met these in Example 6, p. 111.

When $d = 1$, we have the classical (Desarguesian) translation affine plane. But even in the very next case where $d = 2$, an unending bonanza of non-Desarguesian affine planes seems to appear. When K is a field, and $n = 4$, the Klein correspondence provides an easier way of studying the spreads.

5.5.5 A Short Course on the Klein Correspondence

Suppose V is a vector space of dimension 4 over a field F. Since multiplication in F is commutative, one may form the wedge-space $W := V \wedge V$, that is the subspace of homogeneous degree 2 elements of the exterior algebra. The student does not actually need to understand all this double-talk: What is happening is that there is a graded algebra

$$E(V) = F \oplus V \oplus V \wedge V \oplus (V \wedge V \wedge V) \oplus (V \wedge V \wedge V \wedge V),$$

where $a \wedge a = 0$ for all degree one vectors $a \in V$. Everything else is a consequence of the simple distributive multiplication operation $\wedge$. Thus the dimensions of the homogeneous summands above, reading from left to right, are 1, 4, 6, 4, 1. If V has a basis $X = \{x_1, \ldots, x_4\}$, then the 6-space $W = V \wedge V$ has a basis $B(X) = \{e_{ij} := x_i \wedge x_j | 1 \leq i < j \leq 4\}$.

A vector w in $W = V \wedge V$ is said to be *pure* if and only if it has the form $u \wedge v$ for two vectors, u and v of the original four-dimensional vector space V. (Thus, relative to the basis $\{e_{ij} | 1 \leq i < j \leq 4\}$ given above, we see that each basis element e_{ij} is pure, while a vector of the form $e_{12} + e_{34}$ is not pure.)

Now any vector in $W = V \wedge V$ can be written uniquely as

$$v = \sum\nolimits_{1 \leq i < j \leq 4} \alpha_{ij} e_{ij},$$

with respect to the basis $B(X)$.

Thus we may define a quadratic form $Q_X : V \to F$ by the rule that

$$Q(\sum\nolimits_{1\le i<j\le 4}\alpha_{ij}e_{ij}) = \alpha_{12}\alpha_{34} - \alpha_{13}\alpha_{24} + \alpha_{14}\alpha_{23}.$$

Now we come to the truly remarkable part of this quadratic form. You can see that it is entirely defined in terms of the coefficients α_{ij} which in turn appear there because of the choice of basis $B(X)$, in turn determined by a basis X of the 4-space V. Suppose we had chosen an entirely different basis Y of V. Of course that determines an entirely new basis $B(Y)$ of $W = V \wedge V$, and (employing the formula above on the new coefficients) an entirely new quadratic form – or so one would expect. That is exactly the miracle: the two forms will always be the same up to a scalar multiple – specifically $Q_X(v) = 0$ if and only if $Q_Y(v) = 0$ – they vanish on the same variety.

But even this is an immediate consequence of the following

Theorem 5.5.5 *For any basis X, the quadratic form Q_X vanishes only on the pure vectors lying in $W = V \wedge V$.*

It is a relatively simple consequence of the fact that the pure vectors span $W = V \wedge V$, that any non-zero quadratic form Q which vanishes on all the pure vectors is a scalar multiple of Q_X and so vanishes *only* on the pure vectors. We call any one of these non-zero forms the *Klein quadric* (the literature also refers to the common projective variety they define – the collection of all pure 1-subspaces – by this same name).

Now the miraculously universal 1-space of quadratic forms $\{cQ = Q_X | c \in F - \{0\}\}$ (X any basis of V whatsoever) has a number of obvious properties. Suppose $U = \langle a, b\rangle = \langle c, d\rangle$ is a 2-space in V. Then $a \wedge b$ is a scalar multiple of $c \wedge d$ (the scalar in question being the determinant of the transition matrix that passes the first basis to the second). It follows that there is a one-to-one-correspondence (commonly called the *Klein correspondence*):

$$\kappa : \text{2-spaces of } V \to \text{singular 1-subspaces of } W.$$

where singularity is measured with respect to the quadratic form Q_X, where X is any arbitrary basis of V.

But the Klein correspondence has an additional property:

Any pair of distinct 2-subspaces of V intersect at a 1-space if and only if the κ-images are perpendicular singular 1-spaces of the orthogonal space (W, Q), Q being any member of the class of the pairwise proportional forms Q_X, X any basis of V.

5.5.6 Ovoids and Spreads

Continuing the discussion when K is a finite field, $GF(q)$, the Klein correspondence tells us that any spread of 2-spaces in V corresponds to a collection of $1 + q^2$

pairwise non-orthogonal singular 1-subspaces of (W, Q) where $W = V \wedge V$ and Q is the Klein quadric of the previous subsection. It is easy to show that this is the maximal size of a coclique in the graph of singular 1-spaces of (W, Q) under the perpendicular relation. Such a set is called an *ovoid of the polar space of type* $O^+(6, q)$. (The polar space of type $O^+(6, q)$ is also called the *quadric* $Q(5, q)$, where the "5" refers to the projective dimension of the projective space in which the quadric is embedded as an algebraic variety.)

What we are saying is this: for every ovoid in the polar space of type $O^+(6, q)$, there is a spread of 2-spaces of the 4-space V, and a well-defined translation plane. If two ovoids of (W, Q) are conjugate under the classical group $O^+(6, q)$ then the resulting translation planes are certainly isomorphic. The problem is that they might be isomorphic anyway – but that problem is not very great. If two such non-conjugate ovoids led to the same projective plane they would have to do so with respect to residuals with respect two distinct lines, and by that time, with two full translation groups available, the plane must be a finite Moufang plane and hence Desarguesian by the Artin-Zorn theorem.[13] An appendix to this chapter contains a story intended to convince the reader that there are probably unclassifiably many translation planes. It exploits the connection between ovoids and translation planes just described.

5.6 Recognizing the Classical Planes

The next few chapters will produce instances in which one wishes to insist that a plane is a classical Desarguesian plane, or (in the case of a plane arising in a rank three polar space) something pretty close to it, called "Moufang." So at least some method is needed to characterize the classical or near classical planes.

The classical projective planes (those obtained from three-dimensional vector spaces) can be characterized in two ways: (1) by a configurational property called "Desargues' configuration," and (2) by the existence of all possible perspective collineations. From the point of view of this book, we shall be interested in projective planes embedded in higher rank projective or polar spaces. For that reason, we are not interested in contributing to an endless zoo of non-Moufang planes. The classically embedded planes, as we shall see, possess many perspectivities; accordingly we shall follow method (2) in identifying these planes.

5.6.1 Perspective Collineations

For the sake of discussion, $\Pi = (\mathcal{P}, \mathcal{L})$ is a projective plane. The (type-preserving) automorphisms of Π form a group $G = \mathrm{Aut}(\Pi)$. For any automorphism $g \in G$ and subset $X \subseteq \mathcal{P} \cup \mathcal{L}$, the symbol X^g denotes the image of the subset X under the action of g.

[13] For a nice proof of the Artin-Zorn theorem [151], see Hughes and Piper [76], pp. 149–153.

Lemma 5.6.1 *Suppose $g \in G$ fixes every point of a line L. Then there is a point, p fixed by g for which g stabilizes every line on p. If g is a non-identity element, then the pointwise fixed line L and the line-wise fixed point p are unique for g.*

Proof We may assume that g is a non-identity element.

If g fixes a point p not on L, then it fixes every line on p since any such line meets L at a point q fixed by g and so has the form pq.

So assume g fixes no point not on L. Now choose any point x not on L and form the line $N := xx^g$, which meets L at p. Then $N = xp = x^g p = N^g$ is stabilized by g. Now take any point $y \in \mathcal{P} - (N \cup L)$. Then similarly $M := yy^g$ is a line fixed by g and so $M \cap N$ is a fixed point, and so lies on L. That is, $M \cap N = \{p\}$. Since y was arbitrary, it follows that all lines on p are fixed.

For the last assertion, assume g is a non-identity element of G having two fixed points x and y and stabilizing every line on each point. Since every point z not on xy is in the intersection of two such lines, every such point is fixed. If u is in $xy - \{x, y\}$, then u lies on a thick line N containing two points not in xy and so is stablized by g. Then u is the unique point on the two g-invariant lines N and xy, and so it, too, is fixed. Thus g is the identity element, a contradiction.

The uniqueness of a line of fixed points for a non-identity element, is just the dual argument. □

An automorphism g satisfying the hypothesis of this lemma is called a *perspective collineation*. If it is not an identity element, the unique line of fixed points is called its *axis*; the unique point p, all of whose lines are stabilized by g is called the *center* of g. If the center p is incident with the axis L, then g is said to be an *elation*. Otherwise, it is called a *homology*.

Lemma 5.6.2 *A perspective collineation which is not the identity fixes only the points of its axis and its center.*

Proof If such a collineation fixed a point r, not on its axis L, then all lines through r meet the axis L, and so are stabilized by g. Thus r is the unique center of g, namely p. □

Lemma 5.6.3 *Suppose α and β are non-identity elations with a common axis L, but distinct centers p and q, respectively. Then $\alpha\beta$ is also a non-identity elation with axis L and center distinct from both p and q.*

Proof The element $\alpha\beta$ cannot be the identity since then $\alpha = \beta^{-1}$ and β would have the same center. On the other hand, $\alpha\beta$ is a perspective collineation since it fixes every point on L, and so, by Lemma 5.6.1 possesses a unique center r. Thus to show that $\alpha\beta$ is an elation, we need only show that r is a point on L.

Suppose r is not on L. Then $r = r^{\alpha\beta}$, so $r^\alpha = r^{\beta^{-1}}$. Then $\{p, r, r^\alpha\}$ is in a line N, and $\{q, r, r^{\beta^{-1}}\}$ is contained in a line M – just from the definition of p and q as centers. But since these lines already meet at r and r^α, these are the same line – a line $N = M$ meeting L at both p and q, contradicting (PP1). Thus r is on L and $\alpha\beta$ is an elation.

Now, if the center r of $\alpha\beta$ were equal to p, the center of α, then $\beta = \alpha^{-1} \cdot (\alpha\beta)$ would have center p, against the hypothesis that this center is q distinct from p. Thus $r \neq p$. Similarly, r is not q.

All parts of the lemma have been proved. □

We now see the following:

Corollary 5.6.4 *The set of all elations with axis L (including the identity element) form a group* (which we denote by $E(L)$).

The elations with center p and axis L form a subgroup $E(p, L)$ of $E(L)$.

Let (p, L) be a flag of the plane Π. We say that Π is (p, L)-*transitive* if and only if, for any two distinct points x and y not on L whose line xy meets L at p, there exists an element $\alpha \in E(p, L)$ taking x to y. By Lemma 5.6.2 such an α is easily seen to be unique. In other words, (p, L)-transitivity means that every conceivable elation that *could* be in $E(p, L)$ is actually there.

We say that Π is a *translation plane with respect to the axis L* if and only if the elation group $E(L)$ acts transitively on the set of points not on L.

In fact, the following is reached.

Lemma 5.6.5 *The following two statements hold:*

1. *If Π is a translation plane with respect to axis L, then Π is (p, L)-transitive for each point p of L.*
2. *If Π is (p, L)- and (q, L)-transitive for distinct centers p and q in L, then Π is a translation plane with axis L. Moreover, there is a factorization $E(L) = E(p, L) \cdot E(q, L)$.*

Proof 1. If the elation group $E(L)$ is transitive on $\mathcal{P} - L$, then for any two distinct points x and y not in L, there is a unique element $e \in E(L)$ taking x to y. But the line $N := xy$ meets L at a point p which is fixed by e. Thus

$$N^e = (xy)^e = (xp)^e = x^e p = yp = xy$$

is a line stabilized by e. We claim that $E \in E(p, L)$. Since every non-identity element of $E(L)$ is an elation, e is an elation with center $q \in L$. If q were not equal to p, then one would see e move the line qx to the distinct line qy, which is impossible if q is to be the center of e. Thus $q = p$ and $e \in E(p, L)$ as claimed. Since x and y can be chosen as any distinct points of N, we see that $E(p, L)$ is transitive on $N - \{p\}$. Further, as x and y can be chosen so that xy intercepts any point p on the axis L, we see that Π is (p, L) transitive for each point p of L.

2. Suppose Π is (p, L)- and (q, L)-transitive for distinct centers p and q in L. Let x and y be distinct points not in L. We claim that there exists an elation taking x to y. If $xp = yp$, this follows from (p, L)-transitivity. A similar conclusion holds if $xq = yq$. Thus we may assume that y is not in the line xp. Now the lines xp and qy must intersect at a point z. Then by hypothesis there are elations e_1 and e_2, taking x to z and taking z to y, respectively. Then $e = e_1 \cdot e_2$ is an elation taking x to y.

Since x and y were arbitrarily chosen, we see that $E(L)$ is transitive on $\mathcal{P} - L$, and that we have $(E(L) = E(p, L) \cdot E(q, L)$. □

Corollary 5.6.6 *If* Π *is a translation plane with axis* L*, then the elation group* $E(L)$ *is abelian, and every non-identity elation of* $E(L)$ *has the same order, which is either infinite or a prime number.*

Proof Set $G := E(L)$ and let $\mathcal{S} := \{E(p.L)|p \in L\}$. Then the preceding lemmas show that $(G, \mathcal{S})$ satisfies the axioms (T1), (T2), and (T3) of Sect. 5.5.2 The result then follows from Lemma 5.5.2. □

5.6.2 The Moufang Case

A plane that is a translation plane with respect to every line L, is called a *Moufang plane*.[14] If $G = \text{Aut}(\Pi)$, being Moufang entails at least these facts:

1. G is doubly transitive on points.
2. G is transitive on non-incident point-line pairs (i.e., G is *antiflag transitive*).
3. G has its elation groups $E(L)$ all conjugate in G.

Of course Moufang planes have many other properties.[15]

---

[14] After Ruth Moufang, who first studied this situation for real planes.

[15] For any line L, the affine plane $Af(L)$ is coordinated by a certain non-associative quasifield C (this means points of $Af(L)$ are the elements of $C \times C$, with lines defined by equations of the form $x = c \in C$ or $y = xm + b$.) In particular the non-associative quasifield is an additive group with a multiplication satisfying:

1. Both distributive laws.
2. The existence of a multiplicative identity, for which every non-zero element has a two-sided inverse.
3. The law: $a^{-1}(ab) = b$, for all $a, b \in C$. (In fact all of these together imply the *left alternative law*: $(aa)b = a(ab)$.)

(Every such system C is called an *alternative division ring*.)

It is not difficult to see that any classical (Desarguesian) projective plane (whose derived affine planes are given in Examples 3 and 4 preceding) has enough elations to make it Moufang. On the other hand, there *are* Moufang planes which are not classical. However, they are all infinite. The reason is as follows.

Theorem 5.6.7 (Artin-Zorn theorem.) *Every finite alternative division ring is a field.*

Remark Because of the famous Wedderburn theorem that any finite division ring is a field, the proof of the Artin-Zorn theorem is reduced to demonstrating the associative law. Still, it is not an easy task (for a proof, see Sect. 20.6 of Marshall Hall's classic book *The Theory of Groups* [66]).

What are these alternative division rings? A very important theorem of Bruck and Kleinfeld [16] proves that they are Cayley-Dickson algebras. I do not expect this statement to be particularly enlightening to the general reader, but its importance is that these are algebras having a basis with prescribed rules of multiplication that generalize (to other fields) the famous algebra of the Octonions.

5.6.3 The Desarguesian Case

Definition of the Basic Transitivity Properties

Now suppose (p, L) is a non-incident point-line pair – that is, an antiflag. The group of all perspective collineations with center p and axis L is a group of homologies which we denote $H(p, L)$. As before, we say Π is (p, L)-*transitive* if and only if for any two points x and y not in $\{p\} \cup L$, together on a line on p, there exists a homology γ in $H(p, L)$, taking x to y. This means the group $H(p, L)$ acts on $\mathcal{P} - (\{p\} \cup L)$ with orbits being the sets $M - (\{p\} \cup L)$ as M ranges over the lines of Π on p.

We say that *all possible perspective collineations exist* for the projective plane Π if and only if Π is Moufang (i.e., all possible elations exist), and, for each antiflag (p, L), Π is (p, L)-transitive (i.e., all possible homologies exist). What is the effect of the presence of all the additional homologies?

Planes with All Possible Perspective Collineations Are Classical

The goal of this section is to prove the following.

Theorem 5.6.8 *Suppose Π is a projective plane having all possible perspective collineations. Then Π is the classical plane (V_1, V_2) of 1-spaces and 2-spaces of a three-dimensional vector space V over a division ring D.*

We shall prove this theorem using only elementary concepts of module theory. In this way, the reader can be spared the traditional lengthy route through Desargues' configurations, coordination, alternative rings, etc.[16]

Assume now all possible perspective collineations exist for the projective plane Π and let L be a line of Π. This plane Π and line L will remain fixed for the rest of the discussion.

Since Π is Moufang, Lemmas 5.6.2, 5.6.3, and 5.5.2, and Corollary 5.6.4, imply the following.

Lemma 5.6.9 *$E(L)$ is an abelian group of elations inducing a point-regular group of automorphisms of the derived affine plane $Af(L)$ which leaves invariant each parallel class of $Af(L)$. For any two distinct points p and q of L we have*

$$E(L) = E(p, L) \times E(q, L),$$

(the internal direct product).

---

[16] Our choice of development does not mean that the subject of classifying Moufang planes through coordination by alternative division rings does not have an import role in its own right. Indeed, such a program is vital to the construction and classification of rank three polar spaces in Tits' book [138] (see the final section of Chap. 7) and to the recent classification of the Moufang generalized polygons by Tits and Weiss [140]. It is just that we can achieve our immediate goal (the Veblen Young theorem) without it.

Now the affine plane $Af(L)$ is modeled as follows: its points are the elements of the abelian group $E(L)$ with some designated point (say "0") being the identity element. For each point $p \in L$, the cosets of the subgroup $E(p, L)$ form a parallel class of lines. Thus letting $\mathcal{E}$ be all such cosets, our model of $Af(L)$ is the rank two incidence system $(E(L), \mathcal{E})$.

Fix a point "$p0$" not on line L corresponding to the element "0" in the model of $Af(L)$ described in the previous paragraph. Then, by assumption, the group $H(p0, L)$, of all homologies with center $p0$ and axis L also induces a group of automorphisms of the affine plane $Af(L)$ whose point-orbits there consist of $\{p0 = 0\}$ and the set $B - \{0\}$ where B ranges over the lines of $Af(L)$ which are on point 0.

Now set

$$G_0 := \langle \{E(p, p0) | p \in L\} \cup H(0, L) \rangle.$$

Then G_0 is a subgroup of $\mathrm{Aut}(\Pi)$ stabilizing L and fixing the point $p0$.

Lemma 5.6.10 *The subgroup of G_0 generated by the elation groups $E(p, p0)$ acts doubly transitively on L and so acts doubly transitively on the lines of $Af(L)$ on the point 0.*

Proof It suffices to show that the indicated subgroup is doubly transitive on L. But because of our Moufang assumption, each group $E(p, p0)$ transitively permutes the points of $L - \{p\}$. Since L has at least three points, the proof is complete. □

Lemma 5.6.11 *G_0 normalizes $E(L)$.*

Proof Let x be either an elation of $E(p, p0)$ for some point p in L, or a homology of $H(0, L)$, and let e be an elation of $E(L)$. Then $x^{-1}ex$ fixes L pointwise, and has all of its fixed points in L. It follows from Corollary 5.6.4 that $x^{-1}ex$ is an element of $E(L)$. Since such x generate G_0 we see that G_0 normalizes $E(L)$. □

At this point we see that we have a semidirect product of groups $E(L) : G_0$, and that G_0 acts on the affine plane $Af(L)$ exactly as it acts on its model $(E(L), \mathcal{E})$ by group-theoretic conjugation.

Lemma 5.6.12 *For each point p of L, $H(0, L)$ is centralized by each element of $E(p, p0)$. It follows from the definition of G_0 that $G_0 = C_0H$ where $C_0 := \mathbf{C}_{G_0}(H)$,*

Proof Choose $h \in H(0, L)$ and $f \in E(p, p0)$. Since f stabilizes L and h fixes L pointwise, the commutator $[x, f]$ fixes L pointwise. Similarly, since h stabilizes the line $p0$ which is the axis of the elation f, the commutator $[x, f]$ also fixes $p0$ pointwise. Since the commutator $[e, x]$ fixes all points on both lines L and $p0$, it is the identity. The proof is complete. □

Lemma 5.6.13 *$E(L)$ is an irreducible $\mathbf{Z}C_0$-module.*

Proof Suppose E_1 were a proper (non-identity) C_0-invariant subgroup of $E(L)$. Let $X = 0^{E_1}$, the E_1-orbit on $Af(L)$ containing the point 0. Since $X^{C_0} = 0^{E_1C_0} = 0^{C_0E_1} = 0^{E_1} = X$, X is a C_0-invariant set.

Since E_1 contains a non-identity element and acts semi regularly on $\mathcal{P} - L$, we see that X must contain a point a distinct from 0. Set $N = a0$, the line of Π through a and 0, set $\{p\} := N \cap L$, and set $N_1 := N \cap X$. Then N_1 contains 0 and a and possibly further elements of X. Since X is C_0-invariant, and C_0 is transitive on the lines through 0 (Lemma 5.6.10), all lines on 0 meet X in at least two points. In particular X is larger than N_1. Choose a point b in $X - N_1$ and form the line $M = bp$ of Π. Now $E(p, N)$ is a subgroup of C_0, fixing $N = p0$ pointwise, but acting transitively on the set $M - \{p\}$, a consequence of our Moufang assumption. Thus $M - \{p\} = b^{E(p,M)}$ is an entire affine line of $Af(L)$ lying inside X and $M - \{p\}$ is parallel to $N - \{p\}$ in $Af(L)$. Since E_1 acts regularly on X there is an element $f \in E_1$ such that $0 = b^f$. Since $N - \{p\}$ is the line through 0 parallel to $M - \{p\}$ and f preserved all parallel classes of Af we see that $N = M^f$. But now as $M \subseteq X$ and f stabilizes X, we conclude that $N - \{p\}$ is in X. Thus the entire affine line $N - \{p\}$ lies in X.

Since any line through 0 meets X in at least two points, any of them can play the role of N in the previous paragraph. Thus every line of $Af(L)$ on point L is in X and we conclude that $X = \mathcal{P} - L$ and so $E_1 = E(L)$ follows from the regular action. This contradicts the assumption that E_1 was a proper subgroup of $E(L)$, and completes the proof that $E(L)$ is an irreducible $\mathbf{Z}C_0$-module. This completes the proof of Lemma 5.6.13. □

5.6.4 The Proof of Theorem 5.6.8

We continue this discussion to obtain a proof of Theorem 5.6.8. Now, by Schur's Lemma $K := \mathrm{Hom}_{C_0}(E(L), E(L))$ is a division ring. Thus $E(L)$ is a K-vector space (K acting on the left) – more precisely, an irreducible KC_0-module (with this ring acting on the right). The left action of K and the right action of C_0 on $E(L)$ commute.

Now our assumption that all possible homologies exist shows us that H acts on $Af(L)$ fixing 0 and each affine line through zero, and acting transitively on the remaining points of each of those lines. In terms of the model $(E(L), \mathcal{E})$ for this plane, it means that H stabilizes each subgroup $E(p, p0)$ and acts transitively on its non-identity elements. Thus, as a K-space, each one-dimensional subspace of $E(L)$ contains a full subgroup $E(p, p0)$ for some appropriate point p of L. But since $E(L)$ is generated as an additive group by just two such subgroups (Lemma 5.6.9) we see that if a one-dimensional K-subspace containing $E(p, p0)$ also contained an element not in $E(p, p0)$, then $E(L)$ would have K-dimension one.

In the case where $E(L)$ has dimension 1 as a left K-vector space, the multiplicative group of non-zero elements of K (denoted K^*) acts transitively on the elements of $E(L)^\#$ (and hence on the points of $Af(L) - \{0\}$) from the left. (Note we have not yet shown that K preserves the lines of the plane A, so its elements are *not necessarily automorphisms of that plane*.) Fix a point x distinct from 0 on an affine line $N = p0 - \{p\}$, $p \in L$. Then from the transitivity of K^*, we can find a "scalar" α which takes x to a point αx not on line N. Next choose any non-identity element

e from the elation group $E(p, p0)$. By Lemma 5.6.12, e is in C_0 so the right action of e commutes with the left action of K^*. Thus

$$(\alpha x)^e = \alpha(x^e). \tag{5.2}$$

But since e is an elation with center $p \in L$ and axis $p0$, it acts semiregularly on the points of $L - \{p\}$, and hence acts semiregularly on the lines of $Af(L)$ on point 0 which are distinct from N. Since αx is not on N, the affine line on 0 and αx is different from the line on 0 and $(\alpha x)^e$. In particular $\alpha x \neq (\alpha x)^e$.

Since N is the axis of the elation e, $x^e = x$, the right side of Eq. (5.2) becomes αx, contradicting Eq. (5.2). Thus $E(p, p0)$ is itself a one-dimensional K-subspace. Since this holds for every $p \in L$, the 1-subspaces of the K-space V are precisely the lines through 0 in the model $(E(L), \mathcal{E})$ of $Af(L)$, all remaining lines being translates of these. Thus $Af(L)$ is isomorphic to the classical affine plane formed from the vectors of a two-dimensional vector space, and the cosets of its 1-spaces . It now follows from Exercise 5.3 that Π is a classical projective plane, and Theorem 5.6.8 is proved.

5.7 Appendix to Chapter 5

5.7.1 Introduction

Since classical projective planes are characterized by three-dimensional vector spaces, a discussion of the ubiquity of isomorphism classes of projective planes centers on non-Desarguesian planes.

For planes of infinite order you can just about have anything, as was explained in our discussion of free constructions. The non-Desarguesian projective planes of finite order fall into two basic classes:

1. Those planes π for which $G = \mathrm{Aut}(\pi)$ does not stabilize a line or a point
2. Those planes in which G must fix a line or a point

The author strongly recommends the excellent survey article on projective planes by Beutelspacher in the *Handbook of Incidence Geometry* [7], p. 110 in which the first class above are called *proper planes*. I suppose the second class might be called "improper."[17]

[17] By law, any projective plane whose automorphism group has order one belongs to this second class. I really don't know of any plane whatsoever of this sort. My fantasy is that their are jillions of them. So why don't we see them? One suspects that there is a sort of "anthropomorphic principle" at work here. Statistics about what is being discovered in science cannot to be a fair sample of the universe that is out there, but rather an artifact of the pitiful tool-box we possess for making the discoveries. At the time of writing, the methods for discovering planets of other stars favor very

Now there really are infinite families of proper finite planes. Beutelspacher gives a list. The two most astounding cases were the *Hughes planes* (essentially the first infinite family of proper non-Desarguesian planes known), and at the other end of the discovery-time-scale the *Figuera Planes* discovered only a few years ago.

In the second class are the planes whose full automorphism group fixes a line or a point (or both). By duality, one may consider the case that the full autormorphism group G fixes a line L. In that case G acts on the affine plane $A(L)$. Conversely, if we start with a group G of automorphisms of the affine plane A, then G acts on the projective plane $\Pi(A)$. So, in a sense, this case is about affine planes.

But as noted above, this class already includes all planes (and hence affine planes) whose automomorphism group has order one – a class we can't yet approach.

The really unending fountain of planes seems to be in the area of translation planes. My argument is that there are so many planes in this zone that a classification of them is really inconceivable for a long time to come. I mean centuries – or worse: it is like classifying the prime numbers.

Does that mean that the classification of projective planes is dead? Well, if not, you need a classification theorem that says that a projective plane is either a Desarguesian plane, belongs to an author's list, or is a translation plane (a class undivided by further speciation). Anyone can see that even if such a theorem were possible, it is far outside the courageous Lenz-Barlotti scheme. The future author's list must include the untouched class of planes whose only automorphism is the identity. The outlook is simply dismal.

5.7.2 Planes from Ovoids: A Den of Ubiquity

Experiments in Santa Cruz

Experiments employing ovoids to construct spreads of $PG(3, q)$ were begun at a miniconference in U. C. Santa Cruz in 1984. Mason and Ostrom had investigated line spreads of $PG(3, q)$ admitting the groups 2Alt(6) (a non-split central extension of the alternating group of six letters by Z_2) and $Q_8 * D_8 : \mathrm{Sym}(5)$ (the semidirect product of the normal extraspecial group $Q_8 * D_8$ – a central product of quarternion and dihedral groups of order 8 – with the symmetric group of five letters (the latter is permuting the five subgroups of type $Z_2 \times Z_2$ in $Q_8 * D_8$). This group does not live in $SL(3, q)$, and so the translation planes arising from such a spread must be

---

heavy planets much closer to the star than is Mercury to our Sun. That does not mean the universe is typically that way and no physicist has a problem here – they *know* the statistics are screwy.

The same is true of non-Desarguesian projective planes. Almost all of the known finite non-Desarguesian projective planes still exhibit a rich group of automorphisms. They seem to have groups which make them translation planes or their duals, or have rich groups which preserve a subplane. But these groups are there for the simple reason that they played a role in the very construction of these planes. Without them, one would not have found the planes. Thus we should not make any inferences about the density of planes not derived from such groups: they, like the planets, may be there. We just don't know how to see them.

non-Desarguesian. The new paradigm was to view a spread of $PG(3, q)$ as an ovoid of the quadric $Q^+(5, q)$ – that is, the polar space of type $O^+(6, q)$, to identify it by the name of its classical group.[18] Many new planes admitting the Mason-Ostrom group made their appearance at once ([86]).[19] Much later, Eric Moorhouse proved that ovoids admitting $Q_8 * D_8$: Sym(5) exist for all primes p congruent to one modulo four [88]. There was no great mathematics here: it was just a sociological event. We just knew there were very many planes.

During the conference Bill Kantor noticed that groups such as $Q_8 * D_8$: Sym(5) and 2Alt(6), used to insure that the spread is non-Desarguesian, were subgroups of the *Weyl group of* E_8, not just subgroups of $2 \cdot O^+(6, q)$.

We should have taken the suggestion and moved up two dimensions, for we knew the following. If $\mathcal{O}$ is an ovoid of the orthogonal geometry of type $O^+(8, q)$ – that is, $\mathcal{O}$ is a collection of $1 + q^3$ pairwise non-perpendicular singular 1-spaces of an 8-space over $\mathbf{GF}(q)$ with a non-degenerate quadratic form for which there are totally singular subpaces of dimension four (the Witt index) – and if s is a singular vector outside $\mathcal{O}$, then $s^\perp \cap \mathcal{O}$ consists of $1 + q^2$ singular 1-spaces which (taken mod $\langle s \rangle$) forms an ovoid $\mathcal{O}_s$ of $s^\perp/\langle s \rangle$, an orthogonal geometry of type $O^+(6, q)$. The latter, as we have seen, corresponds to a translation plane. Obviously if s and t are singular 1-subspaces outside $\mathcal{O}$ which are conjugate in the stabilizer of the ovoid O in $O^+(8, q)$, the translation planes derived from $\mathcal{O}_s$ and $\mathcal{O}_t$ are isomorphic. But these planes could easily be isomorphic anyway. That is the subtlety.

At least an ovoid of the orthogonal geometry of $O^+(8, q)$ would provide a "master system" organizing a family of translations planes. That was well-known at the time of the Santa Cruz miniconference; we just didn't know how to persue it.

Kantor's Ovoids and Two "Sporadic" Ovoids

Right after the Santa Cruz miniconference, Bill Kantor found three systems of ovoids of $O^+(8, q)$ admitting the respective groups, $U(3, q^2)$, $q \equiv 0, 2 \bmod 3$, the Ree groups $R(q)$ where q is an odd power of three, and $SL(2, q^3)$ where q is a power of 2 [83]. In each of these cases the listed groups act doubly transitively on the eight-dimensional Kantor ovoid.

Soon afterward, two further ovoids appeared. Cooperstein's ovoid is constructed as follows. Let X be an orthonormal basis for a 10-space V over $\mathbf{GF}(5)$. Then all vectors which are sums of five elements of X are singular and are perpendicular to the all-one vector $j = (1^{10})$, with respect to the usual "dot product" form B : $V \times V \to \mathbf{GF}(5)$. Since j is itself a singular vector, the factor space $W = j^\perp/\langle j \rangle$ aquires a quadratic form (W, Q) of type $O^+(8, 5)$. Now let U be the collection of

[18] The group-theorists like the group-name; the pervasive tradition of Italian geometry prefers the former name, that of an algebraic variety. Because of the research efforts of both groups have thoroughly comingled in the literature of the last five decades, the literate reader will just have to get used to this dual language. There is no choice.

[19] In fact, the method was so accessible that, upon explaining it to a Santa Cruz graduate student (Cabaniss) late one afternoon, he produced a new plane of order 11^2 in just an hour or two.

$\binom{10}{5} = 252$ vectors in V with five entries $+1$ and five entries -1 (that is, of shape $(1^5, (-1)^5)$) and observe the following:

1. The vectors of U are singular and lie in $j^\perp$.
2. If $u, v \in U$, then $B(u, v) = 0$ if and only if $v = \pm u$. So the vectors of U span a collection $\bar{U}$ of 126 singular 1-spaces.
3. No two of the 1-spaces of $\bar{U}$ are congruent mod $\langle j \rangle$.

Since $\langle j \rangle$ is perpendicular to every element of $j^\perp$, the form B induces a bilinear form $\bar{B}$ on the factor space $W := j^\perp / \langle j \rangle$, defined by

$$\bar{B}(u + \langle j \rangle, v + \langle j \rangle) = B(u, v).$$

Then the natural homomorphism $\phi : j^\perp \to W$ takes the set $\bar{U}$ to a collection $\mathcal{O} = \phi(\bar{U})$ of $126 = 1 + 5^3$ pairwise non-perpendicular elements in the quadric of type $O^+(8, 5)$. Clearly Sym(10) acts on this structure by permuting the elements of the orthonormal basis X. This group acts in just three orbits on the singular 1-spaces of W outside $\mathcal{O}$. As a result, one has three translation planes: a Desarguesian plane, a semifield plane, and the Walker plane [42].

Similarly, one can let $V = (\mathbf{Z}/7)^{(8)}$ be the 8-space over $\mathbf{GF}(7)$ with an orthonormal basis B. (In this way, the bilinear form is the ordinary "dot product" among 8-tuples.) Let S be the collection of 1-spaces generated by the set of vectors of shapes[20]

$$\begin{aligned} ((\pm 1)^7 : 0) &\quad \text{(128 vectors in 64 1-subspaces)} \\ ((\pm 3)^3, 0^4 : 1) &\quad \text{(280 vectors spanning as many 1-spaces)} \end{aligned}$$

Clearly, with respect to the "sum of squares form," all of the vectors listed are singular (that is the sum of the squares of their coordinates is 0). More important is the fact that the inner product of any two of these vectors is non-zero (not difficult to check). Now, counting these 1-spaces reveals the other miracle:

$$280 + 64 = 344 = 7^3 + 1.$$

That means that our set S is an ovoid of $O^+(8, 7)$. This group induces the group $E(2^6) : \mathrm{Sym}(7)$, or $W(E_7)$ on the ovoid, with 22 orbits on singular points of the ovoid [109]. The result: 22 candidates for translation planes. But how many of these actually give rise to distinct translation planes?

[20] In this convenient notation of Conway, the entries listed before the colon occur in any order in the first seven coordinates, and the entry after the colon is the eighth coordinate.

The Conway-Kleidman-Wilson-Moorehouse Constructions

What was needed was two results: (1) a method to produce zillions of other ovoids of $O^+(8, q)$ (Kantor having provided three infinite families) and (2) a sharp method of distinguishing them. It should fall upon a very great mathematician who had never before worked in projective planes, to provide essentially both of these items, perhaps without realizing the effect this would have upon the incumbents of the projective plane world. The mathematician in question is John Conway.

I will explain 1, but just say a few words about 2 (every bit as great an achievement as 1).

The E_8 Root Lattice

Consider an eight-dimensional Euclidean space (by definition over the real numbers). Then there is a nice orthonomal basis $B = \{e_1, \ldots, e_8\}$ where the "inner product" is defined by $B(e_i, e_j) = \delta_{ij}$, the Kronecker "delta funtion." There is a rather nice "norm function" $N : V \to \mathbf{R}^+$ in which

$$N\left(\sum\nolimits_{i=0}^{8} \alpha_i e_i\right) = \sum\nolimits_{i=0}^{8} \alpha_i^2.$$

Of course this is a positive definite quadratic form.

Now this vector space contains a remarkable subset Λ_2 whose 240 elements are

$$\begin{aligned} &\{\pm e_i \pm e_j\} \quad \text{as } (i, j) \text{ ranges over the 2-subsets of } \{1, \ldots 8\}, \text{ and} \\ (1/2)&\sum\nolimits_{i=1}^{8} \epsilon_i e_i \quad \text{such that } \epsilon_i = \pm 1 \text{ and } \textstyle\sum_i \epsilon_i \text{ is even.} \end{aligned}$$

Let Λ be the abelian subgroup generated by this set. Λ is a spanning discrete subgroup of the vector space E (false of most subsets), and is called a *lattice*, although it has nothing whatsoever to do with the class of posets by that name. This particular lattice is called the E_8 *root lattice*. It is an "even lattice," which means that all norms of its vectors are even integers, while the Grammian of a generating set is invertible in $Z^{(8\times 8)}$.

Let Λ_n be the collection of all vectors of norm $2n$ – that is, the set $\{v \in \Lambda | (v \cdot v) = 2n\}$. Then the set Λ is completely partitioned as

$$\Lambda = \Lambda_0 + \Lambda_1 + \Lambda_2 + \cdots .$$

Now each of the sets Λ_n lies on a sphere of finite radius about zero in the ambient Euclidean space, and thus is a compact discrete subset of this space. Thus $|\Lambda_N| := c_n$ is a finite integer for $n = 0, 1, \ldots$. The series

$$\sum\nolimits_{0}^{\infty} c_n z^n$$

is called the *Eisenstein series of the lattice*. In this case and in the case of most even lattices, the Eisenstein series defines a function in a domain of the complex field, which satisfies certain remarkable functional equations. These functions are called *modular forms* and the proof that these Eisenstein series satisfy the requisite equations uses Poisson inversion. (There is a very good account of this theory in Sects. 2 and 3 of Chap. VII of the book "A course in arithmetic" by J.-P. Serre [107].) One feature of this theory is that if you know a suitable finite number of the coefficients c_n, of the function, you in fact know all of them. It is for these reasons that in the case of the E_8 root lattice, we obtain a truly amazing fact:

$$c_n = 240\sigma_3(n),$$

where, for any non-negative integer n, $\sigma_3(0) = 0$ and otherwise $\sigma_3(n)$ is the sum of the cubes of the divisors of n.

The Constructions

Conway saw what the rest of us did not. If p were an odd prime, then Λ_p is a set of $240(1 + p^3)$ vectors which one might use to produce $1 + p^3$ 1-spaces of an ovoid in $\Lambda/p\Lambda$ – an orthogonal space of type $\Omega^+(8, p)$. How do we get rid of the 240? For that, we must understand that $\Lambda/2\Lambda$ possesses a quadratic form Q which (1) vanishes precisely on all cosets $v + 2\Lambda$ for which $N(v) \equiv 0 \bmod 4$, and (2) whose associated bilinear form (actually a symplectic form over $\mathbf{GF}(2)$), is the inner product on Λ read modulo 2. This space contains a set N of 120 non-singular vectors and a set S of 135 non-zero singular vectors. It follows that the morphism

$$\phi : \Lambda \to (\Lambda/2\Lambda)$$

takes Λ_p onto N. A typical fiber $\phi^{-1}(w)$, $w \in N$, consists of exactly $1 + p^3$ antipodal vector pairs, $\{v, -v\}$, which, taken modulo $p\Lambda$, produces a collection O of $1 + p^3$ singular 1-spaces of the $O^+(8, p)$ geometry of $\Lambda/p\Lambda$. It is then not difficult to show that O, defined in this way, is an ovoid. Conway and his colleagues P. Kleidmann and R. A. Wilson, went on to produce a second recipe for extracting such an ovoid O from Λ_p with any one of the three orbits of the Weyl group $W(E_8)$ on Λ_3 replacing Λ_2 [40]. In all cases, the ovoid O admits a subgroup G of the Weyl group stabilizing a vector in Λ_2 or one of the three orbits on Λ_3, where G (mod any center) is $W(E_7)$, Alt(9), or $W(D_7)$. The ovoid of Cooperstein, and the ovoid over $\mathbf{GF}(7)$ given above are obtained from the second recipe.[21]

Now as p gets larger and larger, the group G acts on more and more orbits on the singular 1-spaces of the ovoid, and thus yield more and more six-dimensional ovoid sections (those giving us translation planes).

[21] Now one sees why Kantor's remark was indeed prophetic. The groups one inherits here *are* subgroups of the Weyl group of type E_8.

These methods have been greatly expanded in recent work of Moorehouse [87, 88]. The result is still millions of projective planes.

The hard part is to tell if the planes are non-isomorphic. For the non-Desarguesian ones that amounts to an isomorphisms between ovoids of $O^+(6, q)$. An invariant of a six-dimensional ovoid that is highly sensitive was discovered by Conway: he calls it a "fingerprint." This is the second of the two accomplishments of Conway mentioned earlier in this subsection and can be found in the work of Charnes [29, 30] which shows just how sensitive this invariant really is. With it Charnes showed that isomorphisms do occur among sections – for example, he showed that the 22 sections of the ovoid of $O^+(8, 7)$ given earlier, give rise only to 15 distinct planes of order 7^2.

One thing is clear: the Conway-Kleidman-Wilson construction (despite some duplications) produces new translation planes by the truckload.

5.8 Exercises for Chapter 5

5.1 Prove Lemma 5.4.2 by checking the axioms (PP1)–(PP3). [For the last recall that A is always firm, and observe that every line of $\mathbf{P}(A)$ contains at least one more point than a line of A.]

5.2 Following the recipe above, formally construct the projective plane of order 2 from the affine plane of order 2. (The latter is essentially the incidence system of vertices and edges of the complete graph K_4 with four vertices.)

5.3 Suppose A is the classical affine plane of Example 4, the incidence system of vectors and cosets of 1-spaces of a two-dimensional K-space W over some division ring K). Give a complete proof that the related projective plane $\mathbf{P}(A)$ is isomorphic to the classical projective plane (V_1, V_2) of 1- and 2-subspaces of a three-dimensional K-space V, as described in Example 1. [You will need to produce the isomorphism with $\mathbf{P}(A)$ as domain.]

5.4 Prove that if D is a difference set for the group D, $(G, \mathcal{L})$ (where $\mathcal{L}$ is described as on p. 109) is a projective plane.

Chapter 6
Projective Spaces

Abstract The partially-ordered system of subspaces of a projective space form a matroid, thus endowing every projective space with a "projective dimension" (one less than the associated matroid rank). The goal of the chapter is a complete proof of the Veblen–Young theorem, which asserts that every projective space whose projective dimension exceeds two is the classical geometry of one- and two-dimensional subspaces of a (possibly infinite-dimensional) right vector space over a division ring. The so-called "Fundamental Theorem of Projective Geometry" asserts that any embedding of a classical projective space into another is induced by a semilinear transformation of the underlying vector spaces. It is proved here, allowing infinite dimensional spaces. The chapter concludes with three technical results that are vital to the classification of the polar spaces of the next chapter.

6.1 Introduction

In this chapter, we shall introduce projective spaces through axioms as simple as those for projective planes – in fact they are weaker axioms, for projective spaces properly include all projective planes.

Chapter 5 revealed a distinction between "nice"planes (the classical or Desarguesian planes) and the "pretty bad" planes (the non-Desarguesian planes) which in Chap. 5 turned out to be plentiful in both the finite and infinite cases.

However, projective spaces which are not planes are very very nice: they are exactly the classical projective spaces of Example 1 of Sect. 3.9, p. 73. This assertion is usually called the Veblen-Young theorem.[1]

[1] The author was once exhorted to call it the Hilbert-von Staudt-Veblen-Young theorem to give full credit to the inventors of the methodology of the Veblen-Young theorem. The new arrivals to the nomenclature of this theorem tended to show a preference for coordinatizing by algebras over the real numbers – a fact that did not detract in the least from the transportability of their methods. No theorem stands in isolation of the work of predecessors – particularly this one. So there is always a potential conflict between the goal of crediting theorems for the purpose of doing justice to the preceding contributions and the goal of having students know which theorem you are talking about. With the latter view in ascendency for this author, I will revert to calling the principle theorem of this chapter the Veblen-Young theorem. No insult to history is intended.

E.E. Shult, *Points and Lines*, Universitext, DOI 10.1007/978-3-642-15627-4_6,
© Springer-Verlag Berlin Heidelberg 2011

The objective of this chapter is to give an account of this theorem. The reader should bear in mind that we must be prepared to have classical projective spaces $\mathbf{P}(V)$ in our conclusions, even when V has infinite dimension.

To achieve this, we must establish a series of results for projective spaces Γ (abstractly defined by axioms below) which are not themselves planes. Here is the list of desired results:

1. All of the projective planes of Γ are Desarguesian and isomorphic to one another.
2. The lattice of subspaces of Γ is a matroid, and hence is a ranked poset.
3. (A result about embeddings in $PG(3, D)$.) Suppose $f : \pi \to f(\pi) \leq \mathbf{P}(V)$ is an embedding of a projective plane π into the projective space $\mathbf{P}(V)$ obtained from a right D-vector space V for which $\dim V = 4$. Let $PG(V)$ be the full rank three geometry of proper subspaces of V. Suppose S is a rank three projective geometry containing π as a plane. Then the embedding f extends to an isomorphism $e : S \to PG(V)$. (This result uses the fact that planes are Desarguesian.)
4. (The previous result in arbitrary dimension.) Similarly, if H is a subspace of the projective space Γ, and $K = \langle H, p\rangle_\Gamma$ is the subspace of Γ generated by H and an external point p, then any embedding $e : H \to \mathbf{P}(W)$ can be extended to an embedding $e' : \langle H, p\rangle \to P(W')$ where W is a vector space hyperplane of W'. (This argument has to work when H and W are infinite-dimensional.)
5. We then show that there is an embedding $e : \Gamma \to \mathbf{P}(V)$ by the "well-ordering principle," followed by a trans-finite induction argument. As a consequence, Γ will be isomorphic to $\mathbf{P}(V)$.

That is basically our proof of the Veblen-Young theorem.

6.2 Basic Properties of Projective Spaces and Their Partially Ordered Systems of Subspaces

6.2.1 The Definition of a Projective Space

A *projective space* is a point-line geometry $\Gamma = (\mathcal{P}, \mathcal{L})$ satisying these axioms:

(PS1) Γ is a linear space.
(PS2) (The Veblen-Young axiom.) Whenever two distinct lines L_1 and L_2 intersect at a point p and M_1 and M_2 are two lines not containing p which non-trivially intersect both L_i, $i = 1, 2$, then M_1 and M_2 must intersect each other at a common point.
(PS3) All lines are thick.

Remarks 1. Suppose $\Gamma = (\mathcal{P}, \mathcal{L})$ consists of only a single thick line. Then all three axioms hold (the first two vacuously), and so this is a projective space.
2. Otherwise there is more than one line, and since axiom (PS1) prevents one line from being properly contained in another, for each line, there is a point not on that line, and so, appealing to (PS1) once more, there is another line intersecting it.

3. A firm geometry, which does not satisfy (PS3) but satisfies (PS1) and (PS2) (we call this a *generalized projective space*) is not difficult to understand. Let us say p and q are *strongly collinear* if and only if they are connected by a thick line. The relation of being "strongly collinear or equal" is easily seen to be an equivalence relation whose equivalence classes are either (a) single points, (b) a line, or (c) a larger projective space in the sense axiomatized above. All pairs of points in different equivalence classes are connected by a thin line. So it suffices to understand these spaces with thick lines as axiomatized above.

Suppose now that Γ is a projective space. If S is a subspace of Γ then S is either empty, a single point, a line, or another projective space (using its internal lines as its lines).

Theorem 6.2.1 *Suppose S is a proper subspace of the projective space $P = (\mathcal{P}, \mathcal{L})$ and suppose p is a point not in S :*

1. *The subspace $\langle S, p\rangle_P$ generated by S and p, consists precisely of all points which lie on some line joining p to a point of S.*
2. *S is a geometric hyperplane of $\langle S, p\rangle$.*
3. *If point r is in $\langle S, p\rangle_P$ but not in S, then $\langle S, r\rangle_P = \langle S, p\rangle_P$.*

Proof 1. Let S_1 be the set of all points which lie on a line joining p to a point of S. (Be aware that by definition $S \subseteq S_1$.) Choose any two distinct points x and y of S_1. We wish to show that the line $L = xy$ lies in S_1. If both of these points are in S, there is nothing to prove. If L is on p, and $x \neq p$, then $L = xp$ meets S since $x \in S_1$ and P is a linear space. If one of x or y lies in S, (say x), and y does not, we can replace x by $x' \in L - \{x, y\}$. Thus we can assume that x and y are in "general position" – by which we mean:

1. $\{x, y\} \subseteq S_1 - S$
2. The lines xp and yp are distinct

Now let $\{a\} = px \cap S$, $\{b\} = py \cap S$ and let H be the unique line of S on a and b (see Fig. 6.1).

Then H and L intersect ap and bp, but not at p. By (PS2), $H \cap L$ is a single point ℓ. Now choose any point r in $L - \{x, y\}$. Then rp and H meet ap and L, but not at x, so by (PS2), $rp \cap H$ is non-empty. Thus r is on a line with p meeting H and so belongs to S_1. Thus, because of the arbitrary choice of r, $L \subseteq S_1$, as required.

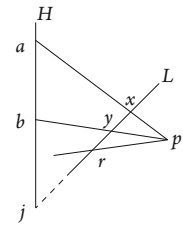

Fig. 6.1 The subspace generated by a point and a subspace not on it

2. We have just seen, that any line L meets subspace S, so S is a geometric hyperplane of S_1.
3. Since $S_1 - S$ has a connected collinearity graph, 2 implies S is a maximal subspace of S_1, from which assertion 3 follows. □

The theorem provides another way to phrase the projective space axioms, as follows.

Corollary 6.2.2 $\Gamma = (\mathcal{P}, \mathcal{L})$ *is a projective space if and only if:*

1. Γ *is a linear space*
2. *Any two intersecting lines generate a projective plane*

Proof This proof is left as an exercise. □

Theorem 6.2.1 exerts a very tight control over the structure of the poset of subspaces of a projective space. In fact this will lead us to the notion of the *dimension* of a subspace of a projective space. A good way to understand this is through the notion of a *matroid*.

6.2.2 Matroids

Perhaps it is a bad time to bring up dependence relations and dimension. But there is a universal notion about these things, which under general circumstances yields a notion of *rank* or *dimension*.

Suppose X is a set and $\mathcal{F}(X)$ denotes the collection of all finite subsets of X. "Dependence" is a relation in $X \times \mathcal{F}(X)$ subject to these three axioms:

(D1) (A reflexive property of dependence.) Each element x depends upon every finite set which contains it.
(D2) (Transitivity of dependence.) If x depends upon finite set B and every element of B depends upon the common finite set C, then x depends on C.
(D3) (The exchange axiom.) Suppose y depends upon the finite set S, but for some $x \in S$, does not depend on $S - \{x\}$. Then x depends upon $(S - \{x\}) \cup \{y\}$. That is, x may be exchanged for y.

Any such dependence relation on $X \times \mathcal{F}(X)$ is called a *matroid on* X.

Example 1 Suppose $X = V$, the elements of a vector space V over X. Then any vector x is said to *depend* on a finite set of vectors, U, if and only if x can be expressed as a linear combination of the elements of U – that is, $x \in \langle U \rangle_V$. It is easy to check that all three axioms (D1) – (D3) hold for "linear dependence" in this case.

Example 2 Let K be a field extension of a field F. An element y is said to *depend algebraically (over F) on a set* $\{a_1, a_2, \ldots a_n\}$ of elements of K if and only if $F(a_1, \ldots, a_n, y)$ is a finite extension of $F(a_1, \ldots a_n)$.

Given a dependence theory, or matroid, $D \subseteq X \times \mathcal{F}(X)$, we can extend this to a relation $\hat{D} \subseteq X \times 2^X$ between points and arbitrary subsets of X, by asserting that x *depends on set* A if and only if, x depends (in the old sense) on some finite subset A_x of A. Using this new definition, a subset Y of X is *dependent* if and only if for some $y \in Y$, y depends on $Y - \{y\}$. A set is *independent* if and only if it is not dependent.

Let $\mathcal{I}$ be the collection of all independent subsets of X. It can be shown that this collection possesses the following properties:

(I1) *The family $\mathcal{I}$ is closed under taking subsets.*
(I2) (Another kind of exchange axiom.) *If $A, B \in \mathcal{I}$ with $|A| > |B|$, then there is an element $a \in A - B$ such that $\{a\} \cup B \in \mathcal{I}$.*
(I3) *If every finite subset of a set A belongs to $\mathcal{I}$, then A belongs to $\mathcal{I}$.*

In fact, the three axioms (I1)–(I3) can be taken as defining a matroid (see the classic books of Oxley or Welsh [94, 146] but for a succinct account the author recommends the excellent book of Aigner [3]).

For any subset Y of X, the *flat generated by* Y is the set of all elements of X which depend on Y. The mapping

$$\sigma : 2^X \to 2^X,$$

which sends each subset Y to the flat $\sigma(Y)$ which it generates, is a *closure operator* on the Boolean poset of all subsets of X – that is:

1. $Y \subseteq \sigma(Y)$
2. It is monotone, that is, $Z \subseteq Y$ implies $\sigma(Z) \subseteq \sigma(Y)$
3. It is an idempotent mapping – i.e., $\sigma^2 = \sigma$

Now there are completely different axiomatic formulations of a matroid, in terms of

- Axioms about "circuits" which are "almost independent" sets of X
- Axioms about flats
- Axioms about a closure operator (the Steinitz-Maclane axiom)

A subset Y is called a *spanning set* if and only if every element x of X depends on a finite subset Y_x of Y – that is, $X = \sigma(Y)$.

It is easy to show the following.

Lemma 6.2.3 *Every maximal independent set is a minimal spanning set, and conversely.*

We call a maximal independent set a *basis* of the matroid. In Example 1, this corresponds to the ordinary notion of a basis for a vector space. In the case of algebraic dependence this notion corresponds to the notion of a transcendence basis of a field extension. The main result for matroids is the following.

Theorem 6.2.4 (The existence of dimension or rank.) *In a matroid, any two bases possess the same cardinality.*

This common cardinality is called the *rank of the matroid*.

Also, for each subset Y of X, the restriction of the dependence relation D to $Y \times \mathcal{F}(Y)$ defines a matroid based on the set Y, and the rank of this matroid is called the *rank of* Y and is denoted $\rho(Y)$. Thus rank is a function

$$\rho : 2^X \rightarrow \text{ cardinal numbers},$$

with special properties. In fact one can prove:

(R1) *For each subset* Y *of* X, $0 \leq \rho(Y) \leq |Y|$.
(R2) (Monotonicity.) *If* $Z \leq Y \leq X$, *then* $\rho(Z) \leq \rho(Y)$.
(R3) (The submodular inequality.) *If* A *and* B *are subsets of* X, *then*

$$\rho(A \cup B) + \rho(A \cap B) \leq \rho(A) + \rho(B).$$

As the reader might suspect, (R1)–(R3) is also an axiomatization of a matroid. Here one lets $\mathcal{I}$ be all subsets Z with the property
(*) *For every finite subset* A *of* Z, *we have* $\rho(A) = |A|$.

[At this point, readers previously unfamiliar with matroids might have a look at Exercises 6.2–6.4 at the end of this chapter.]

6.2.3 The Matroid of a Projective Space

Let $P = (\mathcal{P}, \mathcal{L})$ be any linear space. We will say that a point p *depends on a finite set of points* Y if and only if p is in the subspace generated by Y – that is $p \in \langle Y \rangle_P$. With respect to this special dependence relation on points, we insist that when $(\mathcal{P}, \mathcal{L})$ is a projective space, we get a matroid. First axiom (D1) holds because A is always contained in the subspace that it generates. Suppose now p depends on A and every one of the elements of A depends on the elements of B. This says that

$$p \in \langle A \rangle \text{ and } A \subseteq \langle B \rangle.$$

From the second part, $\langle A \rangle \subseteq \langle B \rangle$, so p indeed depends on B. Thus (D2) holds.

But now the third axiom (D3) is a complete paraphrase of conclusion 3 of Theorem 6.2.1. Thus the following applies.

Lemma 6.2.5 *With respect to the relation of a point depending on a finite set if it lies in the subspace generated by that set, the points of a projective space form a matroid.*

Remark Note that the subspaces of the projective spaces are the flats (or closed sets).

Let $\Gamma = (\mathcal{P}, \mathcal{L})$ be a projective space. The *rank of the projective space* Γ is defined to be the rank of the matroid defined just above in Lemma 6.2.5 and the discussion preceding it.

Warning The *rank of a projective space* is not to be confused with the *dimension of a projective space* (also called its *projective dimension*). A projective space $\Gamma = (\mathcal{P}, \mathcal{L})$ has finite (projective) *dimension* $r - 1$ if and only if it has (matroid) *rank* r. Thus a projective plane is two-dimensional while its matroid of subspaces is rank three. However, if the matroid rank is an *infinite cardinal number* r, one cannot distinguish r and $r - 1$, and so we also call that cardinal number the *(projective) dimension* of the projective space.

The reader is also warned that the literature is somewhat ambiguous about the use of this term "rank" for projective spaces. Many authors like to call the projective dimension of a projective space its *rank*, perhaps because it has fewer syllables than the former perfectly serviceable word. Since we already have a notion of dimension of a projective space, we have reserved "projective rank" to mean the matroid rank for this book.

A single projective point has rank one (and dimension zero). A line is a subspace with a basis of two elements, and so has rank two (and dimension one). The rank three subspaces are all projective planes and have dimension two.

In the case of the projective space $\mathbf{P}(V)$, obtained from a right vector space V over a division ring D, one has that the projective rank of this space is the vector-space-dimension, $\dim V$, while its *dimension* as a projective space is one less than that number. As previously noted, if $\dim V \geq 3$, the symbol $\mathbf{P}(V)$ is a point-line geometry. But if $\mathbf{P}(V)$ has finite dimension n (that is, $\dim V = n + 1$), then the symbol $PG(n, D)$ is the rank n geometry of all *proper subspaces* of V, typed by their vector space dimensions. Thus the *geometric rank* of $PG(n, D)$ is in fact the projective dimension of $\mathbf{P}(V)$ – the integer n.

What about the rank function?

The submodular inequality (R3) enjoyed by the rank function will be useful. From it, one must conclude that any two distinct projective planes of a projective space of dimension three (that is, one having a rank four matroid of subspaces) must intersect in a line. We shall use this fact freely in the next section.

6.3 Subplanes of Larger Projective Spaces

Suppose $P = (\mathcal{P}, \mathcal{L})$ is a projective space properly containing a projective plane. Then the rank of P is at least four, and every projective plane is a proper subspace since, by the last section, we have a rank function. This means every line lies in at least two subplanes. We shall use this fact.

Fix a plane P_3. Choose any line L of P_3. Then L lies on a second plane P_3', and from the rank function we see that both of them are geometric hyperplanes in the

rank four space P_4 that they generate. Fix a point p in P_3 not on line $L := P_3 \cap P_3'$. We wish to show two things:

1. P_3 is isomorphic to P_3' by an isomorphism that fixes the line of intersection pointwise.
2. P_3 is (p, L)-transitive (consult p. 124 for the definition of this property). Basically it means that all possible homologies are available.

First choose a point $z \in P_4 - P_3 \cup P_3'$. For each point x of P_3, form the line xz and let $\rho_z(x)$ be the point at which this line meets plane P_3'. (Recall that since P_3' is a geometric hyperplane of P_4, such a point must exist.) Then $\rho_z : P_3 \to P_3'$ is a bijection fixing the line of intersection L pointwise. We need to show that this mapping carries lines onto lines, so that it will be an isomorphism.

Obviously L is carried onto L. Suppose N is another line of P_3. Then there is a plane $B = \langle N, z\rangle$ which must intersect the other plane at a line which is forced to be $\rho_z(N)$. Thus ρ_z indeed takes lines to lines.

Now we address the alleged (p, L)-transitivity. Suppose x, y, and p are three distinct points belonging to a line N of P_3, with none of these points lying in L. Then as just noted, $B = \langle N, z\rangle_P$ is a plane meeting P_3' in a line $M := \rho_z(N)$. So $\rho_z(x)$ is a point of M. Then $y\rho_z(x)$ is another line of plane B meeting line pz at a point t. One notes that since x and y are distinct, t cannot lie in P_3'. Similarly, since $\rho_t(y) = \rho_z(x)$ and $\rho_t((x)$ are distinct, t cannot lie in P_3. That is, t lies in $P_4 - P_3 \cup P_3$, and so can be used to define a second isomorphism $\rho_t : P_3 \to P_3'$.

Now form the mapping

$$\alpha := (\rho_t)^{-1} \circ \rho_z.$$

Then α is an automorphism, of the plane P_3, fixing p and L pointwise, and moving x to y. Thus all (p, L) homologies exist in this case.

The situation for elations is similar. One picks p on line L, and chooses two distinct points x and y of $N - \{p\}$, where N is a line of P_3 on p. Just pick any line K on p lying in neither plane P_3 or P_3' (possible since thickness of lines does not allow P_4 to be the set-theoretic union of the points of just two planes). Then N and K generate a projective plane subspace C which meets P_3' at a line M. Choose any point f on $M - L$. Then xf meets K at a point z, and yf meets K at a point t. Then using the exact same formula for α given in the previous paragraph, one can see that α is now an elation, with axis L and center p, moving x to y. Thus we have (p, L)-transitivity for elations as well.

Since all possible perspective collineations exist for plane P_2, by Theorem 5.6.8 of the last section of the previous chapter, P_2 is isomorphic to a classical $PG(2, D)$ for some division ring D. Of course, so is P_2', for the same D. Now in at most three steps we can string together any two planes by a sequence of planes, each of which meets its successor at a line. By the transitivity of the isomorphism relation, all of the planes are classical $PG(2, D)$'s.

We record all of this as follows.

Theorem 6.3.1 *The following statements are true:*

1. *Any two projective subplanes of a projective space P of projective dimension at least three are isomorphic.*
2. *If two of them intersect at a line, they are in fact isomorphic by an isomorphism which fixes the line of intersection pointwise.*
3. *Each plane is isomorphic to the classical one, $PG(2, D)$, for the same division ring D.*

6.4 Embedding Three-dimensional Projective Spaces

6.4.1 Affine Spaces

An *affine space* is a firm point-line geometry $A = (\mathcal{P}, \mathcal{L})$ satisfying these axioms:

(AS-1) A is a linear space.
(AS-2) The subspace generated by any two intersecting lines is (with its internal lines) an affine plane.
(AS-3) There is an equivalence relation (which we will call the "parallel" relation) on the set $\mathcal{L}$ of all lines, which, when restricted to the lines in any affine plane subspace, becomes the parallel relation as defined for affine planes.

Remarks

1. Obviously, any subspace that happens to be an affine plane, arises as in (AS-2) since it contains two distinct intersecting lines.
2. Another way to express axiom (AS-3) is this. Let us say that two lines are "plane parallel" if they are either equal or are contained together as non-intersecting lines in some affine plane subspace of A. Then the transitive extension of the "plane parallel" relation is an equivalence relation which we will call "extended plane parallel." Of course this large equivalence class might include intersecting lines. It could even include all lines. Now axiom (AS-3) just asserts that no two lines in the same "extended plane parallel" class can intersect. (If they did, they would not be parallel in the plane they generate.) Note that (AS-3) does *not* say that the (non-extended) "plane parallel" relationship is itself is an equivalence relation. However that will soon emerge (see Lemma 6.4.2 below).
3. Is axiom (AS-3) a consequence of (AS-1) and (AS-2)? Or is it possible that (AS-3) could be false, while the first two axioms are true? There is an example due to Marshall Hall in which all lines have three points, satisfying (AS-1) and (AS-2) but not (AS-3). The points are the unique class of involutions t^G in a certain group G of order $2 \cdot 3^5$. The lines are the triplets of involutions which lie together in a common subgroup isomorphic to Sym(3). In the group G, the involutions satisfy these properties:

   (a) *If $r, s \in t^G$, then $r = s$ or rs has order three.*

(b) *If r, s, u are three distinct members of t^G, with u not in $\langle r, s\rangle \simeq Sym(3)$, then $\langle u, r, s\rangle$ is the semidirect product of $Z_3 \times Z_3$ by Z_2 where conjugation by the generator of the latter group inverts every element of the former.*

The properties (a) and (b) yield (AS-1) and (AS-3) for this space. But (AS-3) fails.

However, the role of the prime three is critical in this example. There is a theorem of Buekenhout (one of the great theorems of modern incidence geometry) [18] as follows.

Theorem 6.4.1 (F. Buekenhout.) *Let $(\mathcal{P}, \mathcal{L})$ satisfy axioms (AS-1) and (AS-2). Suppose $(\mathcal{P}, \mathcal{L})$ has at least two affine subplanes, and that every line has at least four points. Then $(\mathcal{P}, \mathcal{L})$ is an affine space – that is, (AS-3) holds.*

4. Finally there is another possible question. Suppose axioms (AS-1) and (AS-3) hold? Must (AS-2) hold? The exercises of Sect. 6.10.2 at the end of this chapter is offered to suggest that the answer is "no." (These exercises are interesting because of their connections with partial spreads and cocliques in half-spin geometries.)

Suppose now that $(\mathcal{P}, \mathcal{L})$ is an affine space (so all three axioms hold). The first observation is this.

Lemma 6.4.2 *The relation of being "plane parallel" is already an equivalence relation. Any two lines belonging to the same parallel class as prescribed in axiom (AS-3) generate an affine plane and are parallel lines of that plane.*

Proof As in the remarks preceding this lemma, let "extended plane parallel" be the transitive extension of the "plane parallel" relation. Then this equivalence relation refines the "parallel" relation hypothesized in (AS-3).

Now fix a line L, and let $[L]$ and $[[L]]$ respectively be the "extended plane parallel" – and "parallel" equivalence classes containing L. It suffices to show that any line in $[[L]]$ is already "plane parallel" to L.

Suppose N were a line in $[[L]]$. Choose point p in L and point q in N. Then L and pq are two intersecting lines which by hypothesis generate an affine subplane A. Then there is a line L' in A on point q which is "plane parallel" to L in A. Then we see that N and L' are now intersecting lines of $[[L]]$ and so must coincide by (AS-3). Thus $N = L'$ is "plane parallel" with L. □

There are two very important examples of affine spaces.

Example 1 (The classical affine spaces $AG(V)$.) Let V be a right vector space over a division ring D. Then the points are the vectors of V, and the lines are the additive cosets of the one-dimensional subspaces. (These are fully described from the point of view of coordinates in the next section.)

Example 2 (Affine spaces derived from projective spaces.) Suppose $P = (\mathcal{P}, \mathcal{L})$ is a projective space and that H is a hyperplane – that is, a subspace of rank $n - 1$. We form a new space, $A_H(P) = (\mathcal{P}', \mathcal{L}')$ whose points are the points of $\mathcal{P} - H$, and

whose lines are the sets $L - (L \cap H)$ where L is any line of $\mathcal{L}$ not internal to H. (Of course, since H is a hyperplane, such lines meet H at a single point.)

Now we have to verify the axioms. Clearly A is a linear space – no problem. If lines L_1' and L_2' meet at a point of A, they come from intersecting lines L_1 and L_2, respectively of $\mathcal{L}$ which generate a projective plane π of P which meets H at a line M. Then $A_{12} = \Pi - M$ is an affine plane in A generated by the L_i', so (AS-2) holds. For (AS-3) we must come up with the equivalence relation of parallel – but that is easy. We say the two lines of $\mathcal{L}'$ are parallel if and only if the lines of P from which they were derived meet H at the same point. It is immediately from this definition and the definition of how affine planes are derived from projective planes that this definition extends the "plane parallel" relation.

What we don't know is whether the projective space P discussed in this example is a classical space or whether the derived affine space $A_H(P)$ is a classical affine space, or can you have mixtures?

The point of the next section is that you can't have any "mixing up": if P is a classical projective space, so is $A_H(P)$, for any hyperplane H. Conversely, if $A_H(P)$ is classical for some projective space-hyperplane pair (P, H), then the projective space has to be classical.

However we only need these results when the space P has finite rank.

6.4.2 The Classical Affine Space AG(n, D)

In the previous section we defined the classical affine spaces $AG(V)$ where V was a right vector space over division ring D. When V has finite dimension n, we indicate this space by the determining symbol $AG(n, D)$.

In this case, one may model this space through a system of "coordinates." From this point of view, the points are n-tuples in $D(n) := D \times \cdots \times D$ (with n Cartesian factors), and the lines are the sets of vectors of the form

$$\{(x_1, \ldots, x_n) | x_i = b_i t + c_i, i = 1, \ldots, n, t \in D\}$$

for fixed constants (b_i, c_i), $i = 1, \ldots n$, with not all $b_i = 0$. In other words, lines are the vector sets defined by a complete system of linear parametric equations.

Just as was the case for planes, there is a close relationship between classical affine spaces and classical projective spaces.

First we pass from classical projective spaces $PG(n, D)$ to the affine space $AG(n, D)$ by removing a hyperplane. Let $V = D^{(n+1)}$, the vector space of $(n+1)$-tuples over D. Then we may represent points $PG(n, D)$ by the so-called *homogeneous coordinates* in which the non-zero $(n+1)$-tuples

$$(a_1, \ldots, a_{n+1}) \text{ and } (aa_1, \ldots, aa_{n+1})$$

are taken to represent the same point $p = \langle(a_1, \ldots, a_{n+1})\rangle$. Then those points whose homogeneous $(n+1)$-th coordinate is "0" form a hyperplane H and the remaining points can then be unambiguously assigned coordinates

$$(a_1, \ldots, a_n, 1).$$

Now a line L not in H meets H at a point $p = \langle v\rangle$ where $v = (v_1, \ldots, v_n, 0)$ and contains a point $\langle u\rangle$ where $u = (u_1, \ldots, u_n, 1)$, and consists of all 1-spaces in the 2-space

$$\langle u, v\rangle = \{ua + vb | a, b \in D\}.$$

Thus each such point of $L - \{p\}$, has a unique homogeneous coordinate of the form

$$v + ut = (v_1 + u_1 t, \ldots), t \in D,$$

and so consists of those points with last homogeneous coordinate 1 and the first n homogeneous coordinates defined by a system

$$x_i = v_i + u_i t, t \in D, i = 1, \ldots, n$$

of linear parametric equations. We have thus shown that the mapping

$$e^{-1} : \mathcal{P} - H \to A$$

which takes

$$(a_1, \ldots a_n, 1) \to (a_1, \ldots, a_n) \in D^{(n)}$$

is a bijection of points taking lines outside H to lines of A.

But we have also shown the opposite. The classical affine space $AG(n, D)$ can be completed to a classical $PG(n, D)$ projective space as follows. Let $e : A \to PG(n, D) = P(D^{(n+1)})$ which sends each affine point $(a_1, \ldots, a_n)$ to the point with homogeneous coordinates $(a_1, \ldots, a_n, 1)$. Adjoin the lines and points of the projective space $H = P(V_0)$ where V_0 is the subspace of $D^{(n+1)}$ of vectors with last homogeneous coordinate zero. To each line L of A defined by the parametric system

$$x_i = b_i t + c_i, i = 1, \ldots, n$$

we adjoin the point p_L of H with homogeneous coordinates $(c_1, \ldots, c_n, 0)$ to obtain a line $L^* = e(L) \cup \{p_L\}$ of $P(V)$.

In this way the classical affine space $AG(n, D)$ completes to a classical projective space $P(A) \simeq PG(n, D)$ – precisely derived as $PG(n, D)$ minus one of its hyperplanes.

We think of the lines $L_i = \{(0, \ldots, a_j, 0, \ldots, 0)\}$ whose points have all but the i-th coordinate equal to zero, as the "axis lines" of the coordinatization. The group of automorphisms of a classical affine space is a semidirect product of a normal group of translations isomorphic to $(V, +)$ which acts transitively on points, and the group $\mathrm{Aut}(V) \simeq \Gamma L(n, D)$ which acts transitively on any collection n lines on the zero vector which span V. As a result, there are many ways to coordinatize a classical affine space. Thus the following applies.

Lemma 6.4.3 *A classical affine space of finite rank can be coordinatized as n-tuples over a division ring D so that:*

1. *Any preassigned affine point can be labled as the "origin" – and assigned the "zero" n-tuple $(0, \ldots, 0) \in D^{(n)}$.*
2. *Any preassigned spanning collection of n lines on the origin can be taken as the "axes" – that is, those lines whose coordinate points have at most one entry non-zero.*
3. *Every line of A consists of a coset $\langle v \rangle + u$ of a one-dimensional subspace of $V = D^{(n)}$ – i.e., the coordinates are linear functions of a variable $t \in D$.*

6.4.3 Completing an Arbitrary Affine Space to its Associated Projective Space

Without equations the process generalizes to an arbitrary affine space A. We can construct a projective space and a hyperplane H, whose removal yields A. Lemma 6.4.2 shows that the points of the hyperplane H *must* correspond to the parallel classes of A, and that the lines of H are those collections of parallel classes which make an appearence in some affine subplane of A. Everything is determined!

Thus the following applies.

Theorem 6.4.4 *Every affine space A is derived by removing a hyperplane from some uniquely determined projective space P. Suppose $A = A_H(P)$ is an affine space derived by removing a hyperplane H from a projective space P of finite rank n. Then P is classical if and only if A is classical.*

Note that H seems to play a subdued role in this theorem. Does H being classical force A and H to be classical? The answer is "no" if H is simply a line. Then P can be a non-Desarguesian plane. However if H is a classical $P(W)$ which has projective dimension larger than 1, then we know that the parallel classes and affine subplanes of A must form a classical projective space. Is this enough to make A or P classical? The following section begins to address this question.

6.5 Embedding Three-dimensional Projective Spaces, Some of Whose Planes are Classical

6.5.1 Introduction

This section is the most technical of the chapter: the rest is relatively straightforward. As a sort of road map, let us lay out what we intend to accomplish.

We suppose that P is a projective space of projective dimension 3 – i.e., its matroid of subspaces has matroid rank four.[2]

As we have seen in Sect. 6.3, each plane of P is a classical $PG(2, D)$ and so is embeddable. Our objective is really quite simple. We wish to show that if we have an embedding (isomorphism, actually) $f : \Pi \to PG(2, D)$, where Π is any plane of P, then f can be extended to an isomorphism $e : P \to PG(3, D)$ (where, in a natural way, the range $f(\Pi)$ is seen as a subspace of $PG(3, D)$.

Of course this extension e is not unique. However it *is* unique if one specifies the fate (under e) of two points outside Π.

6.5.2 Synthesizing Coordinates: A Basic Lemma

Lemma 6.5.1 *Let $(\mathcal{P}, \mathcal{L})$ be the point-line geometry consisting of two classical affine planes A_2 and A_3 over D, sharing a common line L_1. If there exists an isomorphism*

$$\rho : A_2 \to A_3,$$

which is the identity mapping on the line L_1, then it is possible to give coordinates $(x, y) \in D \times D$ to the points of A_1 and coordinates $(x, y)' \in D \times D$ to the points of A_2 which are "compatible coordinatizations" in this sense:

1. *Points of L_1 have coordinates $(x, 0) = (x, 0)', x \in D$.*
2. *Lines of A_2 are given by linear equations in the coordinates (x, y).*
3. *Lines of A_3 are given by linear equations in its coordinates $(x, y)'$.*

Proof First as A_2 is classical, it possesses a system of coordinates for the points so that points of L_1 receive coordinates $(x, 0)$, and any desired line L_1 of A_2 on the point $(0, 0)$ has all its points (x, y) satisfying a linear equation, or is a "vertical" line whose points have coordinates $(0, x)$, $x \in D$. Now we coordinatize A_3 by giving the point $\rho(x, y)$ the coordinates $(x, y)'$. Note that since ρ is an isomorphism, it

[2] At this point we cannot recklessly declare that P is a truncation to points and lines of a $PG(3, D)$. Such a rash act would presume that such a projective space was already a classical space derived from a four-dimensional vector space. That is the very fact we are trying to prove here, so we must be careful not to be betrayed by our own notation.

must take a line of A_2 to a line of A_3. Thus an image line involves a set of points $(x, y)'$ with x and y satisfying a linear equation. Thus assertions 1–3 hold. □

We put this lemma to work at once.

6.5.3 Coordinatizing P Minus a Desarguesian Plane

For the rest of this section we assume that $P = (\mathcal{P}, \mathcal{L})$ is a projective space of projective dimension 3. We shall suppose that P contains as many classical planes as we may require.

Our goal is to show that P is a classical $PG(3, D)$. We shall do this by choosing a subplane H, and showing that the associated affine space $A := A_H$ is classical. The latter assertion is proved by producing a coordinatization

$$A \to D \times D \times D,$$

where D is the division ring that is the kernel of all subplanes of P, and then proving that all lines of A have a parametric representation

$$\{(a_1(t), a_2(t), a_3(t)) | t \in D\}$$

where the functions a_1, a_2, a_3 are all linear functions of t – i.e., $a_i(t) = d_i t + e_i$ for constants $d_i, e_i \in D$.

Let H be a projective subplane (hyperplane) of P, and let $A := A_H$ be the associated affine space whose points are those of $P - H$. Since there will be a need to speak of projective lines of P as well as lines of the affine space A, we adopt the following convention. If L is a line of the affine space A, then L^* will denote the unique line of the projective space P, such that $L^* - H = L$ (as a point set). As for any linear space, we write pq for the unique line containing points p and q.

Let us fix three independent points f_1, f_2, f_3, of the plane H and fix for all time a point o outside H. Let L_i^* be the projective line $of_i, i = 1, 2, 3$. We form the projective plane subspaces $P_{ij} = \langle o, f_i, f_j \rangle$ together with the associated the affine planes $A_{ij} := P_{ij} - H$, as is illustrated in Fig. 6.2. By choosing a point z in $f_2 f_3 - \{f_2, f_3\}$ we obtain a projection of planes $\rho_z : P_{12} \to P_{13}$ which takes the line $H \cap P_{12} = f_1 f_2$ to the line $H \cap P_{13} = f_1 f_3$, while fixing the line $L_1^* := P_{12} \cap P_{13} = f_1 o$ pointwise. So ρ_z induces an isomorphism of the affine planes:

$$\rho_z' : A_{12} \to A_{13},$$

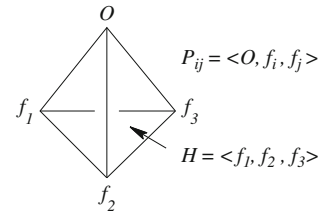

Fig. 6.2 The basic situation

which fixes the affine line $L_1 := L_1^* - \{f_1\}$ pointwise. So we have the following:

(*) *By Lemma 6.4.3 we can assign a system of coordinates* $\{(x, y, 0)|(x, y) \in D^{(2)}\}$ *to the points of* A_{12} *so that point* o *is the "origin"* $(0, 0, 0)$*, and lines* $L_1 = \{(t, 0, 0)|t \in D\}$ *and* $L_2 = \{(0, t, 0)|t \in D\}$*) are the axes and all lines are described by linear equations in the first two coordinates. By Lemma 6.5.1 the isomorphism* ρ_z *provides us with a compatible system of coordinates* $\{(x, 0, z)|(x, z) \in D^{(2)}\}$ *for the points of* A_{13} *such that lines in* A_{13} *are given as coordinates whose first and third coordinates satisfy a linear equation. In this coordinate system lines* L_1 *and* L_3 *become the axes.*

Having established a system of compatible coordinates for the two affine planes A_{12} and A_{13}, we have no further need of the point z.

Referring to Fig. 6.2, one can construct further projections,

$$\pi_{12} : P - \{f_3\} \to P_{12},$$
$$\pi_{13} : P - \{f_2\} \to P_{13},$$

where π_{12} sends each point $x \in P - \{f_3\}$ to the unique point of $xf_3 \cap P_{12}$, the *projection from* f_3 *onto* P_{12}. Similarly, π_{13} sends each point $x \in P - \{f_2\}$ to the unique point of $xf_2 \cap P_{13}$, *the projection from* f_2 *onto* P_{13}. (These maps act like retracts, that is π_{12} is the identity on P_{12}, and a similar statement for π_{13}.) Note that π_{12} and π_{13} induce similar retracts of the associated affine spaces:

$$\pi_{12} : A \to A_{12},$$
$$\pi_{13} : A \to A_{13}.$$

We are now ready to coordinatize all of A. Suppose x is an arbitrary point of $P - H$. Then the projection from f_3 sends x to a point $\pi_{12}(x)$ with coordinates $(a_x, b_x, 0)$ in A_{12}. Similarly, the projection from f_2 sends $x \in P - H$ to a point $\pi_{13}(x)$ in A_{13} with coordinates $(a'_x, 0, c_x)$. We claim that $a_x = a'_x$.

First form the projective plane $R = \langle x, f_2, f_3 \rangle_P$. This plane meets P_{12} at the line N^* through f_2 and $\pi_{12}(x)$. So far, this much is represented in Fig. 6.3a. But as a drawing in the affine plane A_{12}, we see that $N := N^* \cap A_{12}$ is a line parallel to the axis $L_2 := f_2 o \cap A_{12}$. The "L_1-coordinate" of the point of $R \cap L_1$ is $(a_x, 0, 0)$ where

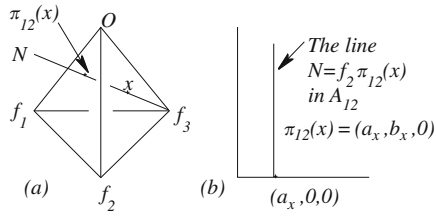

Fig. 6.3 How projections give consistent coordinates

a_x is the first coordinate of $\pi_{12}(x) = (a_x, b_x, 0)$. (This is represented as a drawing in A_{12} in Fig. 6.3b.) Similarly, if $\pi_{13}(x) = (a'_x, 0, c_x)$, we see that the same plane R intersects the plane P_{13} at a line M^*, through f_3 and $\pi_{13}(x)$, and $M := M^* \cap A_{13}$ is a line of A_{13} parallel to the L_3-axis $of_3 - \{f_3\}$ of A_{13}. It follows as before that $M \cap L_1$ is the point $(a'_x, 0, 0)$. But when one intersects a projective plane (such as R) with a non-incident line (such as L_1^*), at best one gets a unique point: thus $N \cap L_1 = R \cap L_1 = M \cap L_1$, so $a'_x = a_x$.

So our claim is proved. Because of this, we can unambiguously assign coordinates to all points of A. We have already assigned coordinates to all points of $A_{12} \cup A_{13}$. For $x \in A - (A_{12} \cup A_{13})$, one has unique constants $a_x, b_x, c_x \in D, b_x c_x \neq 0$ such that $\pi_{12}(x) = (a_x, b_x, 0)$ and $\pi_{13}(x) = (a_x, 0, c_x)$, and so we can assign to the point x the coordinates $\kappa(x) := (a_x, b_x, c_x) \in D \times D \times D$. Thus we have a coordinatizing bijection:

$$\kappa : \mathcal{P} - H = A \to D \times D \times D.$$

The new coordinates are related to those of the two affine planes $A_{1,j}$:

(**) *The projection mapping* $\pi_{1j} : A \to A_{1j}$ *induces this map:*

$$(a, b, c) \to \begin{cases} (a, b, 0), & \text{if } j = 2 \\ (a, 0, c), & \text{if } j = 3 \end{cases}.$$

Now let L be any line of the affine space $A = P - H$ which does not lie in either of the two subplanes A_{12} or A_{13}. Then $L = L^* - H$ for a projective line L^* of P. Let h be the unique point of $L^* \cap H$.

Consider first the case that $h \neq f_3$. Then the projection of L from f_3 – that is $\pi_{12}(L)$ – is a line of A_{12} (see Fig. 6.4), namely the intersection of the plane $\langle f_3, L^* \rangle$ of P, with the plane P_{12} with the point $\pi_{12}(h)$, removed. Then by choice of the coordinate system for the classical plane A_{12}, the line $\pi_{12}(L)$ is given by a linear equation – i.e., there exists linear functions $a_i : D \to D, i = 1, 2$, such that

$$\pi_{12}(L) = \{(a_1(t), a_2(t), 0) | t \in D\}.$$

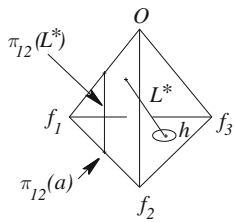

Fig. 6.4 Why these lines are described by linear equations

Note that if h, the unique point of $L^* \cap H$, lies on the line $f_2 f_3$, but is not f_3, then the projection $\pi_{12}(L)$ is parallel to the axis L_2, so that the function a_1 is a constant function.

Similarly, if $h = f_1$, then a_2 is a constant function while a_1 is a non-constant linear function.

On the other hand, if $h = f_3$, then the projection $\pi_{12}(L)$ is no longer a line, but is some point – say (a, b). In that case, both of the functions a_i, $i = 1, 2$, are constant functions.

Similarly, projecting this same affine line $L = L^* - \{h\}$ onto the projective plane A_{13} from the point f_2, we obtain a line $\pi_{13}(L)$, provided, $h \neq f_2$, and the points of that projected line have coordinates $\{(a_1(t), 0, a_2(t) | t \in D\}$, for linear functions a_i, $i = 1, 3$, not both of which are constant functions. Finally, if $h = f_2$, then $\pi_{13}(L)$ is a single point, and so both a_1 and a_3 are constant (and hence linear) functions. But of course, h cannot simultaneously be equal to f_2 and f_3, and so it is not possible for all three of the linear functions a_i, $i = 1, 2, 3$, to be constant functions.

Thus by (**), we have the following.

Lemma 6.5.2 *Under the coordinate system given above for the affine space A, every line of A is described by a system of linear functions. More precisely, given line L of the affine space A, there exist three linear functions $a_i : D \to D$, $i = 1, 2, 3$, not all three of them constant functions, such that each point of L has coordinates $(a_1(t), a_2(t), a_3(t))$ for some $t \in D$.*

It follows that A is classical.

Corollary 6.5.3 *Suppose P is a three-dimensional projective space, with four projective subplanes, H, π_{12}, π_{13}, and π_{23} which together intersect at the empty set. If the three planes π_{ij} are all classical planes, then the affine 3-space $A = P - H$ is the classical $AG(3, D)$.*

It follows from Theorem 6.4.4 that P itself is a classical $PG(3, D)$.

Corollary 6.5.4 *Suppose P is a three-dimensional projective space, H is a projective subplane of P and o is a point not in H. Suppose e is an injective mapping*

$$e : H \cup \{o\} \to PG(3, D)$$

which, when restricted to H, induces an embedding onto a subplane of $PG(3, D)$.

Then e extends to an embedding $e' : P \to PG(3, D)$. (By Exercise 6.5, p. 77, e' is an isomorphism of point-line geometries.)

Also if e has its domain enlarged to accommodate one more point m outside of H and distinct from o, with the proviso that $e(m)$ lies on the line of $PG(3, D)$ spanned by $e(o)$ and $e(mo \cap H)$, then e can uniquely be extended to an embedding e' of P in $PG(3, D)$.

Proof We are given an embedding $e|_H$ of H onto a subplane $PG(2, D)$, which by Exercise 6.7, p. 77, is an isomorphism. By Theorem 6.3.1 every projective subplane of P is now isomorphic to $PG(2, D)$. Select three independent points f_1, f_2, f_3 of H and form the subplanes $\pi_{ij} = \langle f_i, f_j, o \rangle$, all of which are classical

$PG(2, D)$s. Now we can coordinatize the image space $PG(3, D)$ with homogeneous coordinates so that points of $e(H)$ are represented by 1-spaces generated by vectors $(x, y, z, 0)$, points outside $e(H)$ are represented in homogeneous coordinates vectors $(x, y, z, 1)$, and $e(f_1), e(f_2), e(f_3), e(o)$ are respectively represented by the vectors of a fundamental basis of $D^{(4)}$. Then the completion of the embedding amounts to describing coordinates for the points not in H so that lines match. Corollary 6.5.3 shows this can be done.

Next we address the last remark. Suppose e' is not sending m to the prescribed point on the line $e(om)$. Then we may precompose e' with a collineation α of P which fixes H and o pointwise, and extends a homology of a subplane π on om with axis $\pi \cap H$ and center o. That m can move to the right preimage of e' follows from the (p, L)-transitivity of the Desarguesian plane π.

Finally, the question of uniqueness. Suppose there were two such embeddings agreeing on $H \cup \{o, m\}$. Since these embeddings are isomorphisms, the composition of one with the inverse of the other would give an automorphism of P fixing this set. The reader can see that such an automorphism must be the identity. □

6.6 Extending Projective Embeddings of Projective Spaces

Corollary 6.5.4 of the previous section demonstrated that any isomorphism of a subplane $H = \pi$ of a projective space P of projective dimension three into a classical plane $\mathbf{P}(W) = PG(2, G)$ could be extended to an isomorphism $P \to \mathbf{P}(W \oplus \ell)$ so as to take any desired point $p \in P - H$ to the 1-space ℓ. In this section, we generalize this corollary, allowing H to possess arbitrary projective dimension greater than 2. The proof will make repeated use of the Corollary.

Theorem 6.6.1 *(Extending classical embeddings of hyperplanes.) Suppose H is a proper subspace of projective rank at least three in a projective space P, and let p be any point of P not in H. Then H is a hyperplane of the subspace $H' := \langle H, p\rangle_P$ generated by H and the point p.*

Suppose there exists a projective embedding

$$e : H \to P(W)$$

where W is a right vector space over the division ring D. Let $W' = W \oplus \ell$, where ℓ is any 1-space over D.

Then there exists an embedding

$$e' : \langle H, p\rangle_P \to P(W'),$$

sending point p to ℓ and extending e.

Proof First set $H' := \langle H, p\rangle_P$ and choose a point $h \in H$, and a point $q \in ph - \{p, h\}$. Next, let $W' = W \oplus \ell$ where ℓ is some one-dimensional D-space and choose a 1-subspace m of $\langle e(h), \ell\rangle_{W'}$ distinct from both $e(h)$ and ℓ.

Let

$$e' : H \cup \{p, q\} \to \mathbf{P}(W')$$

be the injection which restricts to the isomorphism e on H and sends p to ℓ and q to m. This much is fixed for the rest of the proof.

For each line L of H on point h we are going to define a bijection

$$e^{(L)} : H' := \langle H, p\rangle_P \to \mathbf{P}(W')$$

which extends e', and which restricts to an injective morphism on the plane $P_L := \langle p, L\rangle_P$ and on the projective rank-three spaces $R_{L,x} := \langle P_L, x\rangle$ for each $x \in H' - (H \cup P_L)$.

First let us choose an arbitrary point x in $H' - (H \cup P_L)$. (This is possible since otherwise $H = L$, contrary to assumption.) Let $h(x)$ be the unique point of $px \cap H$. Then $\langle L, h(x)\rangle_P$ is a subplane $H_{L,x}$ of H, and $R_{L,x} := \langle P_L, x\rangle_P$ is a subspace of projective rank three containing this plane. Then the restriction of the injection e to $H_{L,x}$ is an injective morphism

$$H_{L,x} \to \mathbf{P}_{x,L} := \mathbf{P}(\langle e(L), e(h(x))\rangle_W \oplus \ell) \simeq PG(3, D)$$

and so by Corollary 6.5.4 extends uniquely to an isomorphism

$$e_{L,x} : R_{L,x} \to \mathbf{P}_{x,L}$$

taking p to ℓ and q to m.

Then $e_{L,x}$ in turn restricts to an isomorphism

$$\psi_x : P_L \to \mathbf{P}(\langle e(L)\rangle_W \oplus \ell)) := \mathbf{P}_L.$$

Now if we replace x by some other point $y \in H' - (H \cup P_L)$, we obtain another rank three projective space $R_{L,y}$, along with isomorphisms $e_{L,y} : R_{L,y} \to \mathbf{P}_{y,L}$ and $\psi_y : P_L \to \mathbf{P}_L$.

We claim that the two restrictions ψ_x and ψ_y coincide on the plane P_L. Since they have the same range $\mathbf{P}_L$, $\psi_x \circ \psi_y^{-1}$ is an automorphism of the plane P_L fixing line L pointwise (for $e_{L,x}$ and $e_{L,y}$ behave as e on L) and fixing distinct points p and q not on L. This automorphism is clearly the identity and the claim is proved.

This fact allows us to define a point-bijection $e^{(L)} : H' \to \mathbf{P}(W')$ which restricts to the point-line isomorphism $e : H \to \mathbf{P}(W)$. One sets

$$e^{(L)}(x) = \begin{cases} e(x) & \text{if } x \in H \\ \psi_y(x) & \text{if } x \in P_L - H \text{ and } y \in H' - (H \cup P_L). \\ e_{L,x}(x) & \text{if } x \in H' - (H \cup P_L) \end{cases}$$

This is well-defined since, as just seen, ψ_y is the same for all y.

At this point it is not clear that $e^{(L)} : H' \to \mathbf{P}(W')$ is an isomorphism. One would first need to know that it is an embedding, and proving that reduces to showing that $e^{(L)}$ takes lines of H' not in H to lines of $\mathbf{P}(W')$.

This part of the argument will be made easier by first establishing that the bijection $e(L) : H' \to \mathbf{P}(W')$ does not depend on the line L at all.

Suppose N is another line of H on point h. Then $e^{(L)}$ and $e^{(N)}$ agree on $H \cup \{p, q\}$.

Suppose x is any point of $R := \langle L, N, p\rangle$. Then $R_{L,x} = R_{N,x} = R$ and both $e_{L,x}$ and $e_{N,x}$ are isomorphisms $R \to \mathbf{P}_{x,L} = \mathbf{P}_{x,N}$ agreeing on the plane $R \cap H$ as well as on $\{p, q\}$. So by Corollary 6.5.4 they agree.

Next suppose x is a point of $H' - (H \cup R)$. Now

$$T := R_{L,x} \cap R_{N,x} = \langle h, p, h(x)\rangle,$$

a plane, and the restrictions of $e_{L,x}$ and $e_{N,x}$ to T have the same range $\mathbf{P}(\langle e(h), e(h(x))\rangle_W \oplus \ell))$, and the same values on all points of $hh(x) \cup \{p, q\}$. So they agree on this plane. But then at x we have

$$e^{(L)}(x) = e_{L,x}(x) = e_{N,x}(x) = e^{(N)}(x).$$

Thus $e^{(L)}$ and $e^{(N)}$ are the same bijection $H' \to \mathbf{P}(W')$.

Now, to complete the proof, suppose E is any line of H' not in H. If $E = ph$, any $e^{(L)}$ maps E to $\mathbf{P}(e(h) \oplus \ell)$, a projective line. If $h \in H$ and $E \neq ph$, then for any point $x \in E - H$, $e^{(L)}$ restricts to an isomorphism $e_{L,x} \to \mathbf{P}_{x,L}$ taking E to a line of $\mathbf{P}(W')$. So we may suppose that E meets H at a point $h(E) \neq h$, and let N be the line $hh(E)$. Now set $R := R_{N,x} = \langle P_N, E\rangle_P$, where x is any point of $E - H$. Then, as above, $e_{N,x}$ is a restriction of $e^{(N)}$ (in both domain and range). Since $e^{(L)} = e^{(N)}$, the desired result follows. Thus $e^{(L)} = e^{(N)}$ gives us an isomorphism $H' \to \mathbf{P}(W')$ extending e and the proof is complete. □

6.7 The Main Theorem: All Projective Spaces of Projective Dimension at Least Three are Classical

Let P be any projective space of projective dimension at least three. By Theorem 6.3.1 all planes are classical $PG(2, D)$s for a common division ring D. If P has projective dimension three, then by Corollary 6.5.4 P is a classical $PG(3, D)$; so for the remainder of this section we assume that P has projective dimension at least four.

Since any projective embedding of P is an isomorphism of P with a classical projective space $\mathbf{P}(V)$ (Exercise 6.6, p. 77), it suffices to produce such an embedding for P.

Now recall that P is a matroid with its subspaces as flats, and so has a (matroid) basis B – that is, B is a minimal spanning set for P.

We will now invoke the *well-ordering principle* (WO)[3] to impose a well-ordering on the set B.

This will cause some funny language in what follows, for which the author apologizes.

On the one hand B and each of its elements can be regarded as ordinal numbers, and each member b can be thought of as the collection $\{x < b | x \in B\}$ of all its predecessors. If that collection has a maximal member a then we write $b = a + 1$. If that collection has no maximal member we say that b is a *limit ordinal*, and simply identify b itself with this collection of its predecessors.

On the other hand elements of b are points of P in the basis set B, so subsets of B are invoked to span subspaces. In that guise, they don't look like ordinal numbers.

The two ways of speaking are both utilized in the following definition. For each element $b \in B$ we define a subspace S_b of P as follows:

$$S_b := \begin{cases} \langle \{x \in B | x < b\} \rangle_P & \text{if } b \text{ is a limit ordinal} \\ \langle \{x \in B | x \leq b\} \rangle_P & \\ = \langle S_a, b \rangle_P & \text{if } b = a + 1 \end{cases} .$$

One then notes the following:

1. If $x \leq y$ in the total ordering on B, then $S_x \subseteq S_y$.
2. If b is a limit ordinal, then S_b is the set-theoretic union over the "tower" (or totally containment-ordered system) of subspaces $\{S_x | x < b\}$.
3. If $b = a + 1$, then S_a is a hyperplane of S_b.
4. The collection of subspaces $\{S_x | x \in B\} \cup \{P\}$ is a well-ordered set with respect to the containment relation.

Let V be a right D-vector space which has a basis $\{v_b | b \in B\}$ indexed by the elements of the basis B. Then we can produce a well-ordered chain of vector subspaces, just as we did for subspaces of P. Thus for each basis element (also an ordinal number) b in B, we write

$$V_b = \begin{cases} \langle \{v_x | x < b\} \rangle_V & \text{if } b \text{ is a limit ordinal} \\ \langle \{v_x | x \leq b\} \rangle_V & \\ = V_a \oplus \langle v_b \rangle_V & \text{if } b = a + 1 \end{cases} .$$

Then each of the four assertions just listed above hold with S_b replaced by V_b. In particular, with respect to the vector-subspace relation, the system

$$\{V_x | x \in B\} \cup \{V\}$$

is a well-ordered set.

---

[3] See Devlin [58, pp. 58–59] for the well-ordering principle (WO) and its ZF-equivalence to the *axiom of choice*.

Now it is time to talk about embeddings. Being well-ordered, B contains three initial elements; call them b_1, b_2, and b_3, and let π be the projective plane they span. Now there is also a 3-subspace V_3 of V spanned by the first three basis vectors v_1, v_2, and v_3 that would have been indexed by the first three b_i were we able to tolerate subscripts of subscripts.

It has been remarked already that the plane π is isomophic to a classical $PG(2, D) = \mathbf{P}(W_3)$. Since (1) there is a vector space isomorphism $W_3 \to V_3$, (2) $Aut_D(V_3)$ contains $PGL(V_3)$ and (3) $GL(V_3)$ is transitive on bases of V_3, we can compose mappings to obtain an embedding

$$e_3 : \langle b_1, b_2, b_3 \rangle_P = \pi \to \mathbf{P}(V_3),$$

which takes b_i to $\langle v_i \rangle_V$, for $i = 1, 2, 3$. (Just to complete the notation to facilitate generic assertions later, we set $V_1 := \langle v_1 \rangle_V$, $V_2 := \langle v_1, v_2 \rangle_V$ and write e_c for the embbeddings $e_c : S_c \to \mathbf{P}(V_c)$ which are restrictions of e_3 as c ranges over b_1, b_2, and b_3.)

Now let us call an element $b \in B \cup \{B\}$ – regarded here as an ordinal number – *embedding coherent* (or CE for brevity) if and only if the following holds.

There exists a system of projective embeddings of point-line geometries

$$\{e_x : S_x \to \mathbf{P}(V_x) | x < b\},$$

such that the following assertions hold:

1. For each basis element $c < b$ of the matroid basis for S_x, $e_x(c) = \langle v_x \rangle_V$. (The embeddings take the standard basis elements of S_x (a subset of B) to the right 1-spaces of V.)
2. (Extending e_3.) e_x concides with e_i when restricted to S_i, $i = 1, 2, 3$.
3. (Coherence of the embeddings.) If $x \leq y < b$ is a chain in the totally ordered set B, the restriction of the embedding e_y of S_y to its subspace S_x coincides with e_x.

Clearly, by this definition, the first three elements b_1, b_2, and b_3 of B are CE (with respect to the system of morphisms $\{e_1, e_2, e_3\}$).

Now let X be the collection of all elements (ordinal numbers) of B which are not CE and let d be its minimal member. Then for each $x < d$, x is CE.

Suppose this minimal d is a limit ordinal. Then

$$S_d = \cup\{S_x | x < d\} \text{ and } \mathbf{P}(V_d) = \cup\{\mathbf{P}(V_x) | x < d\}.$$

Now there is a system of coherent embeddings $e_x : S_x \to \mathbf{P}(V_x)$ (extending e_3, if $x \geq b_3$) such that for any $y \leq x$, $y \in B$, e_x restricted to the subspace S_y of S_x is e_y.

Now S_d is the union of the subspaces $\{S_x | x < d\}$ and V_d is similarly the set-theoretic union of the vector subspaces V_d (and similarly for their corresponding projective spaces). It is thus possible to define an embedding $e_d : S_d \to \mathbf{P}(V_d)$ extending all the previous e_xs. Each point p (or any line L) of S_d lies in some S_x for $x < d$. Set $e_d(p) := e_x(p)$ (or $e_d(L) := e_x(L)$). The coherence shows that the

definition of $e_x(p)$ (or $e_x(L)$) does not depend on the particular choice of $x < d$ as long as p (or L) lies in S_x. The coherence also shows that e_d preserves incidence of points and lines. Since the point-images are a spanning set of S_d (from defining S_d as a union above), one is forced to conclude that $e_d : S_d \to \mathbf{P}(V_d)$ is a genuine projective embedding of S_d that extends all the preceding embeddings e_x. It follows that d is CE. But that contradicts the choice of d as a member of X.

Next suppose d is not a limit ordinal. Then $d = a + 1$. That means that S_a is a hyperplane of S_d and there already exists an embedding $e_a : S_a \to \mathbf{P}(V_a)$ behaving correctly on the basis for S_a and extending all predecessor-embeddings $\{e_x | x < a\}$. You should not be surprised that everything is set up for an application of Theorem 6.6.1. We may assume that $a \geq 3$, set $H = S_a$, $p = d$ (d is a point in basis B on this visitation), set $W = V_a$ and $\ell = \langle v_d \rangle_V$ and $W' = V_d = V_a \oplus \ell$. With this choir in place one applies Theorem 6.6.1 to conclude that there is an embedding e_d extending e_a and sending d to $\ell = \langle v_d \rangle$. Again it follows that d is CE, contrary to the choice of d.

Since X is a proper subset of B (b_1 is not in it), it must be empty. That means B (itself viewed as the maximal ordinal number completing the well-orderd set B when the latter has no maximal member) is CE. In particular, there is an embedding

$$e : P := S_B \to \mathbf{P}(V_B) = \mathbf{P}(V).$$

By Exercise 6.6, p. 77, e is an isomorphism and so we have at last the following theorem.

Theorem 6.7.1 (Veblen-Young.) *Suppose P is an arbitrary projective space which is not empty, a single point, a line, or a projective plane. Then P is a classical projective space $P(V)$ obtained from some vector space V over a division ring D.*

6.8 The Fundamental Theorem of Projective Geometry

Suppose V_i is a right vector space over the division ring D_i, $i = 1, 2$. Suppose $\sigma : D_1 \to D_2$ is an isomorphism of division rings. A mapping $f : V_1 \to V_2$ is said to be σ-*semilinear* if and only if

$$f(u + v) = f(u) + f(v),$$
$$f(v\alpha) = f(v)\alpha^{\sigma}$$

for all vectors u, v of V_1 and all scalars α of D_1.

Let $S(V_i)$ be the lattice of all subspaces of V_i (including 0 and the whole space). Then any σ-semilinear transformation $t : V_1 \to V_2$ takes subspaces of V_1 to subspaces of V_2 and so induces a morphism $S(V_1) \to S(V_2)$ of lattices, and even a morphism of projective spaces $\mathbf{P}(V_1) \to \mathbf{P}(V_2)$ whose kernel is the subspace $\ker(t) := \{v \in V_1 | t(v) = 0\}$.

The issue of the "fundamental theorem" is to see to what extent these statements can be reversed. One can see that if $t : V_1 \to V_2$ is a σ-semilinear mapping, and γ

is a fixed non-zero scalar, then the mapping $f\gamma : V_1 \to V_2$ defined by $(\gamma f)(v) = f(v)\gamma$ is a σ'-semilinear form (where $\sigma'(\lambda) = \gamma^{-1}\lambda^{\sigma}\gamma$) and induces that same mapping of subspaces $S(V_1) \to S(V_2)$ as does f. In this case we say that $f\gamma$ is *proportional to* f.

Theorem 6.8.1 (The Fundamental Theorem of Projective Geometry.) *Let V_1 and V_2 be finite-dimensional right vector spaces over division rings D_1 and D_2, respectively. We assume* $\dim V_1 \geq 3$. *Let $L_i = S(V_i)$ be the lattice of all subspaces of V_i, $i = 1, 2$. If there is an isomorphism between these lattices,*

$$\zeta : L_1 \to L_2,$$

then there is an isomorphism $\sigma : D_1 \to D_2$, $\dim V_1 = \dim V_2$ *(over their respective division rings), and ζ is induced by a bijective σ-semilinear transformation $\eta : V_1 \to V_2$. Moreover η is unique up to proportionality.*

Proof See Jacobson's classic book *Basic Algebra I* [79, pp 448–451], where the presented proof is Artin's. □

An adaptation closer to the spirit of this chapter would be as follows.

Corollary 6.8.2 *Suppose $P := \mathbf{P}(V) = (\mathcal{P}, \mathcal{L})$ is a projective space of finite rank at least three. Suppose $e : P \to \mathbf{P}(W)$ is an embedding of point-line geometries. Then e is an isomorphism which is induced by a bijective σ-semilinear transformation $t : V \to W$ where $\sigma : D_1 \to D_2$ is an isomorphism from the coordinatizing ring of V_1 to the corresponding ring for V_2.*

Proof $e(P)$ is a spanning subspace of $\mathbf{P}(\mathbf{V})$ which is a linear space. This forces $e(P) = \mathbf{P}(W)$, so e is an isomorphism. Then, as the points and lines of P define the matroid of all subspaces of P, e induces an isomorphism of the lattices of subspaces $\mathbf{P}(V)$ and $\mathbf{P}(W)$. By the "fundamental theorem" preceding, this mapping is induced by a bijective semilinear transformation $t : V \to W$. Upon truncating to points and lines, e is induced by t. □

Theorem 6.8.3 (The infinite version of the Fundamental Theorem of Projective Geometry.) *Let $e : P_1 \to P_2$ be an embedding of one projective space into another. Assume P_1 has rank at least three, so $P_i = \mathbf{P}(V_i)$, $i = 1, 2$. Then the isomorphism e is induced by a bijective σ-semilinear transformation $t : V_1 \to V_2$ (where σ is as in the proceeding corollary).*

Proof For each non-zero subspace U of V_1 let $\mathcal{P}(U)$ be the full collection of its incident 1-subspaces. Then its image under e is a projective subspace of $\mathbf{P}(V_2)$, and is in fact the "flat" F in the matroid of $\mathbf{P}(V_2)$ spanned by $e(\mathcal{P}(U))$. It is easy to see that the vectors contained in the 1-spaces in F, are closed under scalar multiplication and addition, and so, form a vector space $e(U)$ such that $F = \mathbf{P}(e(U))$. (In fact, $e(\mathcal{P}(U) = \mathbf{P}(e(U))$ exactly.)

Let $\mathcal{U}$ be the collection of all subspaces of V_1 of dimension at least three. Now for each $U \in \mathcal{U}$, there is a bijective semilinear transformation $\eta_U : U \to e(U)$

which induces the embedding $e|_{\mathcal{P}(U)} : \mathcal{P}(U) \to e(\mathcal{P}(U))$, and which is unique up to proportionality.

Now fix a 3-subspace U_0 of V_1, and select a particular bijective σ-semilinear transformation $\eta_0 : U_0 \to e(U_0)$ inducing $e|_{\mathcal{P}(U_0)}$. Let $\mathcal{U}_0$ be the members of U which contain U_0. For each $W \in \mathcal{U}_0$, we have a η_W restricted to U_0 inducing the same mapping $\mathbf{P}(U_0) \to e(\mathbf{P}(U_0)$ as η_0. Then we can scale the semilinear $\eta_W : W \to e(W)$ so that $\eta_W|_{U_0} = \eta_0$. Then this newly-defined η_W must be σ-semilinear.

Now for every vector $v \in V_1$, set $\hat{t}(v) := \eta_{U_0+\langle v\rangle}(v)$. Now it is easy to see that $\hat{t}$ is a σ-semilinear transformation $V_1 \to V_2$ which is injective and satisfies

$$\langle \hat{t}(v)\rangle = e(\langle v\rangle)$$

for every non-zero vector v in V_1. The proof is complete. □

The last comment is to adapt this to general morphisms of projective spaces. Morphisms of projective spaces are not truly morphisms in the sense of morphisms of point-line geometries. One recalls from Chap. 4, that a morphism of projective spaces, is really a mapping $f : \mathbf{P}_1 - K \to \mathbf{P}_2$ with these properties:

1. If $p \vee K = r \vee K$ for a projective point p not in K, then $e(p) = e(q)$ is a point of $\mathbf{P}_2$.
2. If L and N are projective lines meeting K trivially, with $L \vee K = N \vee K$, then $e(L) = e(N)$, and it is a full projective line of $\mathbf{P}_2$.

The subspace K is called the *kernel* of the morphism.

Of course, if $\mathbf{P}_1$ and $\mathbf{P}_2$ are classical polar spaces, we can write $\mathbf{P}_i = \mathbf{P}(V_i)$ where V_i is a right vector space over division ring D_i, $i = 1, 2$. Then a morphism f as above is uniquely determined as an embedding $e : \mathbf{P}(V_1/K) \to \mathbf{P}(W)$ where W is a D_2-subspace of V_2. The D_1-dimension of V_1/K is the D_2-dimension of W and is called the *index* of the morphism f. Now we can restate this in a general way as follows.

Corollary 6.8.4 (The fundamental theorem of projective geometry, once again.) *Let $f : \mathbf{P}_1 \to \mathbf{P}_2$ be a morphism of projective spaces of index at least three. Then there are division rings D_i, and vector spaces V_i over D_i such that $\mathbf{P}_i = \mathbf{P}(V_i)$, $i = 1, 2$. More importantly for f, there exists an isomorphism $\sigma : D_1 \to D_2$ of the division rings, and a σ-semilinear transformation*

$$t : V_1 \to V_2,$$

inducing the morphism f – that is, $t(p) = f(p)$ for every 1-space p of V_1.

6.9 Three Technical Results Concerning Projective Spaces

This seems to be the appropriate place in which to present a few special technical results concerning projective spaces that are vital for the analysis of polar spaces in the next Chap. 7. The reader wishing to cover only, the main theory of projective spaces might be advised to defer this section until later.

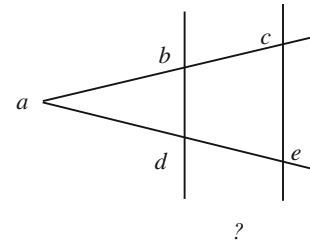

Fig. 6.5 The Veblen configuration

Theorem 6.9.1 (Cuypers-Johnson-Pasini [46].) *Suppose* $\Lambda = (\mathcal{P}, \mathcal{L})$ *is a linear space with thick lines and suppose* S *is a proper subspace of* Λ. *Assume the Veblen axiom holds for every configuration of four lines, not all of which lie in the subspace* S. *Then the Veblen axiom holds generally, and* Λ *is a projective space.*

Proof We suppose abc and ade are collinear triples of distinct points in subspace S. We must show that the lines bd and ce intersect (Fig. 6.5).

Choose point $f \in \mathcal{P} - S$, and choose a point b' in line bf distinct from both b and f. Notice that b' is not in S. Then lines ab' and cf meet two lines ac and bf, and, as bf is not in S, they intersect at a point c' not in S (Fig. 6.6a).

Next, lines db' and ec' meet the two lines ae and ab' on a, and as ab' is not in S, they intersect at a point g (Figure 6.6b).

If $g = f$, then $b'g = b'f = bd$, so $b' \in bd \subseteq S$, which is untrue. So $g \neq f$. We can then consider the line fg.

Now fg and ce meet $c'f$ and $c'g$ (neither in S) and so intersect at a point x (Fig. 6.6c). Similarly fg and bd meet lines $b'f$ and $b'g$ (neither in S) and so meet in a point y.

Now as $ce \cup bd \subseteq S$, we see that $\{x\} = fg \cap S = \{y\}$. Thus cd and bd intersect at $x = y$, as required. □

Here is the second technical result of this section.

Theorem 6.9.2 *Let* $\mathcal{P} = S + T + C$ *be a partition of the points of a projective space* $\mathbf{P}$ *with this property:*

(*) Every projective line which carries one point of S *and a point of* C *in fact carries exactly one point of* S *and all of the remaining points of the line belong to* C.

Then the subspace $\langle S \rangle_{\mathbf{P}}$ *generated by* S *lies in* $S \cup T$.

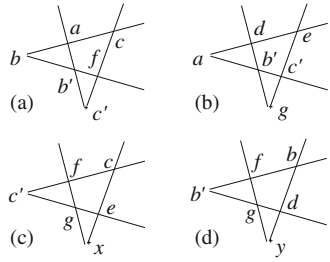

Fig. 6.6 Applications of the Veblen axiom in the proof of Theorem 6.9.1

Proof Immediate consequences of (*) are these:

(**) If a line carries a point of S and a point of T, then it lies in $S \cup T$.
(***) If a line carries a point of T and a point of C then the line is contained in $T \cup C$.

In these two cases, there is no information telling us how the points of the line are distributed among S and T or among C and T.

It also follows from (*) that if a line is spanned by elements of S, then that line lies in $S \cup T$. More generally, suppose A is a subspace of $\mathbf{P}$ with $A \subseteq S \cup T$ and $s \in S - A$. We claim that the subspace $B := \langle A, s \rangle_{\mathbf{P}}$ lies in $S \cup T$. Suppose not. Then $B - A$ must contain a point c in C. Since A is a hyperplane of B, the line sc meets A at a point a. But $a \in S \cup T$, by hypothesis, while by (*), every point of $sc - \{s\}$ must belong to C, so point a belongs to C, a contradiction. So the claim holds. Thus:

If A is a subspace of $\mathbf{P}$ lying in $S \cup T$, and $s \in S$, then the subspace generated by A and s lies in $S \cup T$.

This principle provides an easy induction argument showing that any subspace of $\mathbf{P}$ generated by finitely many elements of S lies in $S \cup T$.

Now suppose c were an element of $\langle S \rangle_{\mathbf{P}} \cap C$. Then c depends on a finite subset of S and so lies in a subspace generated by finitely many elements of S. But we have just seen that that is impossible. Thus no such c exists and the proof of the theorem is complete. □

Our third result continues our investigation of a projective space $\mathbf{P}$ with a point-partition $\mathcal{P} = S + T + C$ satisfying the condition (*) of p. 161.

We say a subspace X of $\mathbf{P}$ is *pure* if and only if

1. $X \subseteq S + T$, and
2. For every subspace Y containing X as a hyperplane, $Y - X$ is either contained in $S + T$ or is contained in C.

For example, the hypothesis (*) tells us that any point s of S is a pure point, since, for any line L on s, either $L - \{s\} \subseteq S + T$ or $L - \{s\} \subseteq C$.

We first observe that the class of pure subspaces of $\mathbf{P}$ is closed under taking subspaces – i.e., the class of pure subspaces forms an order ideal in the poset (actually a lattice) of all subspaces of P. We now have the following.

Theorem 6.9.3 *The set U of all pure points of $\mathbf{P}$ forms a subspace of $\mathbf{P}$. Thus U is the unique maximal pure subspace of P.*

Proof Let U be the set of all pure points of $\mathbf{P}$. Suppose u and v are two distinct points of U and let $L = uv$ be the projective line determined by these points. Suppose N is a line of $\mathbf{P}$ meeting L at a point z, and suppose $N - z$ contained a point $c \in C$ as well as a point $t \in S + T$ (that is, z is *not* pure). Form the plane

$\pi := \langle L, N\rangle$. Since u is pure, all points on the line ut lie in $S + T$. Since each line of π on v must meet a point of ut, and since v is pure, all of these lines lie in $S + T$. But that is impossible as vc is one of these lines. Since N is an arbitrary line meeting L, we see that any point of L (such as z) is pure, and that the line L itself is a pure subspace.

Thus U is a subspace of $\mathbf{P}$. If $U = \mathcal{P}$ then $\mathcal{P} = S + T$, and every subspace is pure. If, on the other hand, U is a proper subspace of $\mathbf{P}$, we may consider a subspace Y containing U as a hyperplane. Choose $y \in Y - U$. Any point x of $Y - U$ is either equal to y or lies on a line joining y to a pure point u in U. Thus $x \in C$ if and only if y is in C. Thus all points of $Y - U$ lie in $S + T$ or lie in C according as $y \in S + T$ or $y \in C$. We conclude that U is itself a pure subspace. Then so are all of its subspaces. Conversely any pure subspace lies within U since U contains all their pure points. Thus U is the unique maximal pure subspace of $\mathbf{P}$ and the proof is complete. □

Corollary 6.9.4 *Suppose $C \neq \emptyset$ and suppose $\langle S\rangle_{\mathbf{P}}$ is a hyperplane of $\mathbf{P}$. Then $\langle S\rangle_{\mathbf{P}} = S + T$ and the hyperplane complement $\mathcal{P} - \langle S\rangle_{\mathbf{P}}$ has all its points in C.*

Proof Since the points of S are pure, the hyperplane $\langle S\rangle_{\mathbf{P}}$ is a pure subspace. Since $C \neq \emptyset$ and $\langle S\rangle_{\mathbf{P}} \cap C = \emptyset$, one must have $\mathcal{P} - \langle S\rangle_{\mathbf{P}} = C$. □

6.10 Exercises for Chapter 6

6.10.1 Exercise on the Definition of Projective Space

6.1 Prove Corollary 6.2.2 on p. 138.

6.10.2 Axioms of Affine Space

6.2 Suppose V is a six-dimensional vector space over a finite field $\mathbf{GF}(q)$. Let $\mathcal{T}$ be a collection of $1 + q^3$ three-dimensional subspaces which pairwise intersect at $\{0\}$ and whose set-theoretic union is V. In terms of the associated projective geometry $PG(V)$, one may say that $\mathcal{T}$ is a partition of the projective points of $PG(V)$ by a subcollection of its projective planes. The reader will recall that in the language of Chap. 5, p. 111, $\mathcal{T}$ is called a *spread* of V into three-dimensional vector subspaces. [Such spreads exist. For example one may regard V as a two-dimensional vector space over the field $\mathbf{GF}(q^3)$ and take $\mathcal{T}$ to be the collection of one-dimensional $\mathbf{GF}(q^3)$-subspaces. There are certainly others.]

Let H be any projective hyperplane of $PG(V)$. Show that H must contain exactly one of the component planes of the spread $\mathcal{T}$. (In this way, $\mathcal{T}$ partitions the hyperplanes of $PG(V)$ and so also induces a spread of the dual projective space.)

6.3 Let $W \simeq PG(4,q)$ be a fixed hyperplane of $PG(V) \simeq PG(5,q)$ of the previous exercise, and let π be the unique projective plane component of $\mathcal{T}$ contained in W. Let $\mathcal{S}$ denote the collection of intersections of W with the planes of $\mathcal{T}$ which are distinct from π. Then $\mathcal{S}$ is a collection of q^3 (projective) lines of W whose point-shadows partition the points of $W - \pi$. Suppose now that $H \simeq PG(3,q)$ is a hyperplane of W. Show that either:

1. H contains π and contains no line of $\mathcal{S}$, or
2. H meets π at a line L_H, and H contains exactly q members of $\mathcal{S}$. (In this case one notes that the lines of $\mathcal{S}$ which are contained in H together with the line L_H form a spread of H.)

[Hint: One possible approach. Suppose H is a hyperplane which meets π at a line. Let a be the number of lines of $\mathcal{S}$ which are contained in H and let b be the number of lines of $\mathcal{S}$ which meet H at a point. This gives rise to two equations in a and b.]

6.4 Continuing the notation of the previous exercise, let $\mathcal{H}$ denote the collection of $q^2(1+q+q^2)$ hyperplanes of W which do not contain π. We define a point-line geometry $A = (\mathcal{S}, \mathcal{H})$. We say that a "Point" $S \in \mathcal{S}$ is *incident* with a "Line" $H \in \mathcal{H}$ if and only if they are incident as objects (a projective line and a hyperplane) of $PG(W)$.
Prove each of the following statements:

1. A is a linear space with q^3 Points and $q^2(1+q+q^2)$ Lines.
2. Each Line is incident with q points and each Point is incident with q^2+q+1 Lines.
3. For each projective line N of the plane π, there are exactly q^2 members of $\mathcal{H}$ which intersect π at N. Letting $\mathcal{H}_N$ denote this subcollection of hyperplanes of $\mathcal{H}$ that contain N, show that $\mathcal{H}_N$ is a collection of Lines of the space A which partition the Points of A. (Simply as an intuitive aid, we think of $\mathcal{H}_N$ as a "parallel class" of Lines.)
4. If N and M are distinct projective lines of π, $\mathcal{H}_N \cap \mathcal{H}_M = \emptyset$. Thus there is a partition

$$\mathcal{H} = \{\mathcal{H}_N | N \text{ ranging over the projective lines of } \pi\}$$

of the Lines of A into q^2+q+1 "parallel classes".

6.5 The standard model of the classical affine geometry $AG(3,q)$ takes as its points P_{AG} the projective points of a $PG(3,q)$ which do lie in a fixed hyperplane U (which is a projective plane). Its lines L_{AG} and planes Π are the projective lines and projective planes of the $PG(3,q)$ which do not lie in U. (See Examples 1 and 2 on p. 145.) Show that $AG(3,q)$, has these properties:

1. It is a linear space with q^3 points, and $q^2(1+q+q^2)$ lines each of size q.

2. The set of lines may be partitioned into $(1 + q + q^2)$ classes, $L_{AG}^{(p)}$ (the "classical parallel classes"), indexed by points p of U, such that the q^2 lines of any class partition the points.
3. There is a partition $\Pi = \{\Pi_N\}$ of the set of planes into component subsets indexed by the projective lines N of the plane U, there being q planes, each an $AG(2, q)$, in each class, no two of them sharing a point.
4. Any two distinct lines of a line-parallel class $L_{AG}^{(p)}$ generate a unique plane of Π.

6.6 The reader has by now noticed that the linear space $A = (\mathcal{S}, \mathcal{H})$ of Exercise 6.4 satisfies properties 1 and 2 of the previous exercise. This raises the question whether $A \simeq AG(3, q)$? The student is asked to examine the possibility that A contains any subspaces isomorphic $AG(2, q)$ at all. Describe the difficulties in detail.

6.10.3 Simple Results on Embedded Projective Spaces

6.7 Suppose $e : P \to \mathbf{P}(V)$ is an embedding of a general projective space P into a spanning set of the classical projective space $P(V)$, where $\dim V = 4$. Show that e is an isomorphism.

6.8 Prove that the previous assertion is true when V is allowed to assume arbitrary dimension.

Chapter 7
Polar Spaces

Abstract The goal of this chapter is a complete teachable proof from first principles of the classification of all polar spaces of polar rank at least four. Polar spaces are defined by the well-known "one or all" axiom on points and lines. Each such space is uniquely associated with a non-degenerate polar space whose singular subspaces are projective spaces. If a maximal singular subspace of the latter has finite rank, this number is called the polar rank of the space. The classical objects on the scene are the polar spaces defined by sesquilinear forms and pseudoquadratic forms. The classification of polar spaces of rank at least four proceeds in two steps: (1) showing that a polar space that can be embedded in a projective space must be classical and (2) showing that a non-degenerate polar space of rank at least four is in fact embeddable. The technical results from the end of the preceding chapter are used in both steps, while a homotopy-proof of Tits' unpublished result that polar spaces of rank at least three are characterized by the cone over a point facilitates the proof (due to Cuypers, Johnson, and Pasini) of (2).

7.1 Introduction

Before pondering an abstract definition, let us begin with some examples. A *generalized quadrangle* is a partial linear space, $\Gamma = (\mathcal{P}, \mathcal{L})$, with the property that, given any non-incident point-line pair (p, L), there is a unique point on L collinear to p.

The "thin" version of this notion is the very simple geometry of the vertices and edges of a square.

But there are many examples with all lines thick. To start with, one could start with a point p, and impose upon it a "bouquet" of lines of various sizes (at least two), partitioning the remaining points. We call this a "star." It has two atypical properties: (1) there is a point p collinear with all remaining points, and (2) all points distinct from p live on only one line.

Are there examples without either of these two features? In Chap. 3, Example 7, p. 75, we gave an example of such a geometry. Its points were the 15 2-subsets of a set $\Omega = \{1, 2, \ldots, 6\}$. Its lines were the $(2, 2, 2)$-partitions of Ω, with the obvious incidence. For example $L = \{(12), (34), (56)\}$ is a typical line. There are then 15

E.E. Shult, *Points and Lines*, Universitext, DOI 10.1007/978-3-642-15627-4_7,

© Springer-Verlag Berlin Heidelberg 2011

lines, each line has exactly three points, and each point is on exactly three lines.[1] This is a generalized quadrangle.

Example 1 Fix a set $\Omega = \{1, \ldots, 7\}$ of seven neutral letters. We construct a geometry as follows: The *points* ($\mathcal{P}$) are the 35 3-subsets of Ω. The *lines* ($\mathcal{L}$) are the triplets of 3-subsets which pairwise intersect at a fixed common letter ℓ. For example, {(123), (145), (167)} is a line. Thus two points are *collinear* if and only if they are represented by 3-sets which intersect at a single letter ℓ. The third point of the line they are on is uniquely determined: it is the triplet consisting of the letter of intersection together with the two letters not covered by the union of the 3-sets representing the given first two points. Thus $\Gamma = (\mathcal{P}, \mathcal{L})$ is a partial linear space.

Now consider a typical line: $L = \{(123), (145), (167)\}$. Suppose (abc) represents a 3-set $p = \{a, b, c\}$ not equal to one of the 3-sets of L. Then if $1 \in \{a, b, c\}$, one sees that p meets two of the triplets of L at two points each and meets the third at just the letter 1. Thus p is collinear with exactly one point of L. On the other hand, if $\{a, b, c\}$ does not contain the letter 1, there are two possibilities describing how the set $\{a, b, c\}$ intersects the three 3-sets comprising the points of L: (1) p meets one of these triplets at two letters of Ω, another triplet at one letter and the last triplet at the empty set, or (2) the 3-set (abc) meets each 3-set of L at one point each. In case (1) (abc)) is again collinear with exactly one member of L, while in case (2), p is collinear with all points of L.

So in Example 1, we seem to have this property:

(P) (The "all or one" property.) $\Gamma = (\mathcal{P}, \mathcal{L})$ *is a point-line geometry such that for any point-line pair* (p, L)*,* p *is collinear with all or exactly one of the points of* L*.*

A point-line geometry $(\mathcal{P}, \mathcal{L})$ satisfying axiom (P) will always be called a *polar space*

This turns out to be an unbelievably strong property despite some trivial cases. First, if $\mathcal{L}$ is empty, we just have a bunch of points, no two collinear. However, if there is even one line, the geometry is connected. But this just opens the door to another trivial example: a singular subspace, where, for each antiflag (p, L), p is collinear with every point of L. Second, it is possible to reduce the study of these geometries to the case that no point is collinear with all remaining points (as was the case with the "star" example above).

Finally, before taking one step further we record the following fundamental fact which hardly needs proof.

Lemma 7.1.1 *Any subspace of a polar space is also a polar space.*

[1] To those students familiar with graph theory, this geometry would be the incidence system of edges and 1-factors of the complete graph K_6.

Suppose we are beyond the sorts of trivialities encountered in the previous paragraph: namely, no pair of distinct points is collinear, or *every* such pair is a collinear pair or there is a point collinear with all remaining points.. Then one can show that $(\mathcal{P}, \mathcal{L})$ is a partial linear space with projective singular subspaces.[2] Much more monumental is the "near classification": one gets (1) the case of no lines, (2) a generalized quadrangle, (3) a difficult intermediary rank case (rank three polar spaces), or (4) classical objects obtained from sequilinear and pseudoquadratic forms over a vector space. Only the generalized quadrangles in this list are unclassified.

The generalized quadrangles seem to be intermediate between the totally unclassifiable projective planes, and the truly sparse population of generalized hexagons and octagons. Of course there are once again the free constructions of generalized quadrangles, but even in the finite case there are about 16 infinite families of generalized quadrangles (only 6 of which can even reasonably be called classical). This focuses the attention on polar spaces of rank three or higher. Although it is fair to say that this case succumbs to a complete classification in rank three or higher (as was the case of projective spaces of dimension three or higher), the proof does not allow such a simple bifurcation.

At rank three there are some non-classical polar spaces – but they have been classified.[3] Nonetheless, there is a spot in Part IV of this book in which this classification of rank three polar spaces is invoked – at least to rule out a very peculiar generalized quadrangle. This appears as a technical lemma due to Arjeh Cohen that is absolutely vital to all that follows it: the last four chapters of this book.[4] Although we are forced to refer to Tit's famous book [138] for the actual classification of polar spaces of rank three, an informal – almost recreational – sketch of the argument is given in Sect. 7.9, just to give an idea of it.

In contrast to this, the polar spaces of rank four or more can be approached from quite elementary first principles, and we will do that.

The history of this subject has had quite a number of developments. First Veldkamp [143] dealt with a number of axioms which amounted to the embeddable polar spaces, then deduced that they must be classical. Of course, missed in this were the strange rank three examples. J. Tits produced a simpler and more general set of

---

[2] Still today the author finds that a startling conclusion from only the "one or all" axiom.

[3] However the the arguments reaching them employ the classification of reductive algebraic groups of rank three. Since there is simply no way that a lecturer can insert a "short course on algebraic groups" as a mere "side bar," we will not be able to present a "ground up" development of the rank three polar spaces. I do not apologize for this. Any thinking mathematician should be humbled by the set of mathematical facts that we do not understand. Even the miniscule portion supported by theorems that have been proved by someone cannot always be demonstrated ab initio in two hours in any classroom, for the simple reason that the methods for proving them are quite beyond the course.

[4] During the writing of this book K. Thas found a proof of Cohen's lemma that did not require the classification of the non-degenerate polar spaces of rank three.

axioms, which included the non-embeddable examples, and classified all of these. In other words he and Veldkamp did most of the hard work. Later Buekenhout and Shult [23] simplified these axioms in what is their ultimate form (modulo a few remarks about infinite rank and thin lines), but even these author's were merely reducing everything to Tits' axioms. No part of Tits' original theory was actually replaced: it was just used.

Another development was clarifying the extent to which the classification could be made to apply to polar spaces of infinite rank, an effort more or less completed by Peter Johnson [80]. In that work, Johnson also made it clear once and for all how the second Veldkamp space of Chap. 4 was being used to produce a projective embedding – on that score, our arguments follow his. .

Finally, H. Cuypers, P. Johnson, and A. Pasini [46, 47] evolved a presentation of polar spaces all of whose steps are at a short logical distance from first principles – that is, it could be taught to students in Kansas.

So the "guts" of the chapter is what I call the Cuypers-Johnson-Pasini (CJP) theory, expounded in Sect. 7.4 which shows that such a polar space has an embedding as absolute points of a projective space with a polarity. The embeddability is still a version of Veldkamp's orginal theory using a space made out of geometric hyperplanes but this time for a "cone" $p^\perp$.

In order to make the CJP-theory work, we present a proof of an unpublished theorem of Tits which says that for non-degenerate polar spaces of rank at least four, say, P and Q containing points p and q, respectively, any isomorphism $p^\perp \to q^\perp$ can be extended to an isomorphism $P \to Q$.

Finally, the author apologizes for inserting small bits of theory here and there to complete his own understanding of the existing theorems in the literature.[5]

7.2 Generalized Quadrangles

As mentioned in the introduction, polar spaces include the generalized quadrangles, and these low rank polar spaces stand outside the classification theory just as the projective planes stand outside the Hilbert/von-Staudt/Veblen-Young classification of projective spaces.

So, before leaving generalized quadrangles altogether, we should say something about the known examples of quadrangles. Again there are free constructions for the infinite case, and quite a few non-classical examples in the finite case, though, so far, the adjective '‘ubiquitous" does not seem to be justified, as it was for projective planes.

[5] Two examples are notable. In that literature one cannot find the invocation of Theorem 6.9.2 to prove the basic Lemma 7.7.3, nor an appeal to Theorem 7.6.5 to prove the key Theorem 7.7.15. I think both results are indispensable. My aim here is not to review the literature, but simply to prevent a student from being impaled on a logical gap in the supposedly self-contained presentation of this chapter.

A *generalized quadrangle* is a point-line geometry $\Gamma = (\mathcal{P}, \mathcal{L})$ with this property:

(GQ) For any antiflag $(p, L) \in \mathcal{P} \times \mathcal{L}$, one has $|p^\perp \cap L| = 1$.

Like any point-line geometry, it is required that all lines have at least two points, and that there are no "repeated lines" (distinct lines with the same point shadows). However it is not required that each point be on at least two lines. A simple example is the "*star*," consisting of a "bouquet of lines" – that is, a point p incident with all lines, each of which contains at least one other point apart from p and which pairwise intersect at p. In this case every point distinct from p lies on a unique line.

Now it is clear that two distinct lines cannot intersect in two or more points, for in that case there would be a point in one that is not in the other ("no repeated lines"), and that point would be collinear to at least two points of the other line, against (GQ). We may thus conclude the following.

Lemma 7.2.1 *A generalized quadrangle is always a partial linear space.*

We also observe the following.

Lemma 7.2.2 *In a generalized quadrangle there is at most one point p with $p^\perp = \mathcal{P}$, and in the case where that happens, Γ is a star.*

We say that the generalized quadrangle is *non-degenerate* if it is not a star.

The next observation is that the axiom (GQ) is self dual. Basically it is a property of all antiflags, and the axiom reads the same when everything is "dualized" – the incidence geometry $(\mathcal{P}, \mathcal{L})$ is replaced by $(\mathcal{L}, \mathcal{P})$. The minor problem is our requirement that a point-line geometry have at least two points (see Chap. 3, p. 61). This property is not always realized when we dualize a rank two geometry. For example, the dual of a star, would be a set L of points in which the whole set L and all its 1-point subsets are declared to be the "lines." However dualizing the star turns out to be the only case among generalized quadrangles where an attempt to dualize runs into the legislative difficulty that the dual is techinically not a point-line geometry. It will be a consequence of a theorem of the next section that if a point of a polar space lies on only one line, then there must exist points p with $p^\perp = \mathcal{P}$. As just remarked, that can only happen for a generalized quadrangle when it is a star.

Thus we may conclude the following.

Two lines L and N of a generalized quadrangle, are said to be *opposite* if they do not intersect (we will generalize this notion later for polar spaces). It easily follows from (GQ) that the mapping which takes each point of a line to the point it is collinear with in an opposite line is a bijection. Thus if the opposite relation on lines is connected, then all lines have the same cardinality.

In the next section we will show the following.

Lemma 7.2.3 *If the generalized quadrangle is not a star, then every line has a line opposite it.*

As a corollary we see the following.

Lemma 7.2.4 *If the generalized quadrangle $(\mathcal{P}, \mathcal{L})$ is not a star, then every point lies on at least two lines. In this case the dual geometry $\Gamma^* = (\mathcal{L}, \mathcal{P})$ is also a generalized quadrangle.*

Then one can easily show the following.

Theorem 7.2.5 (A trichotomy with two trivial cases.) *For a non-degenerate generalized quadrangle, at least one of the following statements hold:*

1. *All lines are thin, and Γ is the incidence structure of the vertices and edges of a complete bipartite graph.*
2. *$(\mathcal{P}, \mathcal{L})$ is a grid, that is, there are two parallel classes of lines partitioning the points: each line in one class meets each line in the other; each point is on just two lines, one from each class.*
3. *Every point is on at least three lines, at least one of which is thick. In this case, there are cardinal numbers s and t such that:*

   (a) Every line has $1 + s$ points
   (b) Each point is on $1 + t$ lines

The first two cases are dual to each other. They are completely determined up to isomorphism by two parameters giving the cardinalities of the two components of the bipartite graph in case 1, or the cardinalities of the the two parallel classes in Case 2. So the interesting case is the third one. Here Γ is said to be *a generalized quadrangle of order* (s, t), where s and t are the two distinguished parameters. A long-standing problem is this: can one have a generalized quadrangle of order (s, t) when one of the parameters is finite and the other is infinite?

When both are finite, there are constraints: one must have (1) s or t equal to 1 (which puts you in one of the trivial cases of the trichotomy), or else (2) $\sqrt{t} \leq s \leq t^2$. (This is a key result of D. Higman [73]. For general results, see also the "bible" in this area: *Finite Generalized Quadrangles* by S. Payne and J. Thas [98], and for commentaries and updatings of this bible, see the book of H. van Maldegham [141], and the extensive papers of J. A. Thas written after the former book.) When s and t are both finite, we can count the lines and points. We have

$$|\mathcal{P}| = (1 + s)(1 + st), \tag{7.1}$$

$$|\mathcal{L}| = (1 + t)(1 + st). \tag{7.2}$$

We require no further results about generalized quadrangles in order to present the theory that follows. At best these quadrangles reappear as residues.

Still, the many examples of quadrangles can serve the dual purpose of reinforcing the student's intuitive understanding of the axiomatic definition, and the other objective of filling in the blanks (as far as known at the time of writing) of the classification prospects for these objects. For this reason we have preserved the examples known at the time of writing in an appendix to this chapter – just as we did in Chap. 5 for non-classical projective planes.

7.3 Basics of Polar Spaces

7.3.1 Definition

A *polar space* is a point-line geometry with thick lines satisfying the following.

(P) *For any antiflag* $(p, L) \in \mathcal{P} \times \mathcal{L}$, $p^\perp \cap L$ *is either a single point or all of* L.

Looking back to Chap. 3 on point-line geometries, the reader should see that this is a very special sort of gamma space. Thus we find the following.

Lemma 7.3.1 *For any polar space,* $\Gamma = (\mathcal{P}, \mathcal{L})$*, the following assertions hold:*

1. *If a point is collinear with two distinct points of a line, it is collinear with all such points.*
2. *The subspace generated by any set of points forming a clique in the point-collinearity graph is a singular subspace.*
3. *Maximal singular subspaces exist and are the the maximal cliques of the point-collinearity graph.*

As usual, we write $x^\perp$ to denote the set of points which are collinear with point x together with x itself. In any gamma space, $x^\perp$ is a subspace, but in the case of a polar space, it is either all of $\mathcal{P}$ or is a geometric hyperplane of Γ.

7.3.2 Opposite Lines and the Radical

We set $\mathrm{Rad}(\Gamma) = \{p \in \mathcal{P} | p^\perp = \mathcal{P}\}$. These are the points which are collinear with all remaining points other than themselves. From axiom (P), the set $\mathrm{Rad}(\Gamma)$ is a subspace of Γ.

A line N is *opposite* a line L if and only if is a bijection $\phi : N \to L$ such that, for each point x in N, $\{\phi(x)\} = x^\perp \cap L$, the unique point of L collinear with x. Clearly this is a symmetric relation between lines.

Lemma 7.3.2 *1. Suppose* a *and* b *are distinct points on a line* N*, and* L *is a second line with* $a^\perp \cap L$ *and* $b^\perp \cap L$ *being distinct single point subsets of* L*. Then* N *is opposite* L *and* $(L \cup N) \cap Rad\ (\Gamma) = \emptyset$*.*

2. *If a line* L *does not intersect the radical Rad* (Γ)*, there is a line opposite* L *— in fact, on any point* w *not in* $L^\perp$*, there is a line on* w *opposite* L*.*

Proof 1. Suppose N and L are lines, and that a and b are distinct points of N for which $a^\perp \cap L = \{a'\}$ and $b^\perp \cap L = \{b'\}$, $a' \neq b'$. Now for any point $x \in N$, $x^\perp \cap L = L$ or is a single point. But if $L \subseteq x^\perp$, then a', which is not collinear with b would be collinear only with x in N, forcing $x = a$. But similarly, the fact that b' is not collinear with a would force $x = b$. We can't have both equations as a and b are distinct. Thus each point x of N is collinear with a unique point x' of L.

Also, we cannot have $x' = y'$ for distinct points x and y of N, for then x' would be collinear with all points of N, forcing $x' = a' = b'$, against a' and b' distinct.

Thus the mapping $\phi : N \to L$ which takes each point x of N to the unique point x' of L with which it is collinear, is an injective mapping. But it is surjective as well, since for any point c in L, axiom (P) implies the existence of a point $u \in c^\perp \cap N$ for which $c = \phi(u)$. Thus ϕ is a bijection.

Finally, if $r \in L \cap \mathrm{Rad}(\Gamma)$, then $r^\perp$ would contain N contrary to the existence of the bijection ϕ. So $L \cap \mathrm{Rad}(\Gamma) = \emptyset$. The same holds with N replacing L, by the symmetry of the hypotheses between N and L.

2. Fix line L and suppose $L \cap \mathrm{Rad}(\Gamma) = \emptyset$. For each point u in L, let $\Delta(u) := \{x \in \mathcal{P} | x^\perp \cap L = \{u\}\}$. Then by (P), we have a disjoint union:

$$\mathcal{P} = L^\perp + \cup_{u \in L} \Delta(u) \ .$$

Since Rad $(\Gamma) \cap L = \emptyset$, there are at least two distinct points u and v on L for which both $\Delta(u)$ and $\Delta(v)$ are non-empty. Choosing $u_1 \in \Delta(v)$, we see that there must be a line M on u_1 and v. Without loss of generality, we may suppose our point w lies in $\Delta(u)$ and so is collinear with u but not v. Then by (P), $w^\perp \cap M$ is a point v_2 of $\Delta(v)$. Finally let N be a line on w and v_2. We see that $w^\perp \cap L = \{u\}$ and ${v_2}^\perp \cap L = \{v\}$, with $v \neq u$ and $w \neq v_2$. Thus by part 1, N is opposite L. □

7.3.3 The Basic Morphism

Consider now the following relation on points of a polar space $\Gamma = (\mathcal{P}, \mathcal{L})$: we write $x \simeq y$ if and only if $x^\perp = y^\perp$. Then the relation "$\simeq$" is manifestly an equivalence relation, and we write $\bar{x}$ for the equivalence class containing point x. Now suppose $x \simeq y$ for distinct points x and y. Then there is a line L on x and y. Now if L intersected the radical $R := \mathrm{Rad}(\Gamma)$ trivially, by part 2 of Lemma 7.3.2, there would be a line N opposite L, and in that case there would be a point on N collinear to x but not y, an impossibility. Thus every line on x and y intersects the radical R.

Conversely, if x and y are distinct points lying on a line that intersects the radical R, the reader should be able to prove that $x^\perp = y^\perp$. Thus we have the following.

Lemma 7.3.3 *For any two points x and y of the polar space Γ, $x^\perp = y^\perp$ if and only if $x = y$ or x and y lie on a line meeting the radical Rad(Γ) non-trivially.*

Any polar space for which $\mathrm{Rad}(\Gamma) = \emptyset$ is said to be a *non-degenerate polar space*. So this lemma spawns the immediate corollary below.

Corollary 7.3.4 *In a non-degenerate polar space, $x^\perp = y^\perp$ implies $x = y$.*

Now let $\Gamma' := (\mathcal{P}'.\mathcal{L}')$, the geometry of points outside $R = \mathrm{Rad}(\Gamma)$ and all lines exterior to R (that is, they do not intersect R). Now we can define a point-line morphism $\Gamma' \to \bar{\Gamma} = (\bar{\mathcal{P}}, \bar{\mathcal{L}})$ by sending each point x not in R to its equivalence class $\bar{x}$, and sending each line L not meeting R to the collection of equivalence

classes of the points on L. Note that as L is a line not meeting the radical, the "bar" mapping is injective when restricted to L.

Now one can use the fact that a point of $\mathcal{P}'$ is collinear to a point of a $\simeq$-class if and only if it is collinear with all points of the class to conclude the following.

Lemma 7.3.5 *Given $\Gamma = (\mathcal{P}, \mathcal{L})$, a polar space, the "bar image" $\Gamma/Rad(\Gamma) := (\bar{\mathcal{P}}, \bar{\mathcal{L}})$ is a polar space (the non-degenerate section of Γ) with the extra property, that Rad $((\bar{\mathcal{P}}, \bar{\mathcal{L}}))$ is empty.*

Given a non-degenerate polar space $\bar{\Gamma}$, it is possible to adjoin a singular subspace S for each point $\bar{p}$ of $\bar{\Gamma}$ to assign a singular space A_p containing S as a subspace, and lines connecting these subspaces, to recover a polar space $(\mathcal{P}, \mathcal{L})$ with radical S. The trouble is that this can be done with many arrangements of lines for the same non-degenerate section $\bar{\Gamma}$. In fact the cardinalities of lines can vary all over the place.

7.3.4 Basic Properties of Non-degenerate Polar Spaces

In this section Γ is a non-degenerate polar space. If $\mathcal{L}$ is empty, then Γ is just a set of pairwise non-collinear points. In this case, if it has at least two points, we say that Γ is a *rank one polar space*, since it is a rank one geometry. If it has just one point, it is a rank zero geometry, since its geometry of proper subspaces is empty.

We shall prove the following.

Theorem 7.3.6 *A non-degenerate polar space is a partial linear space all of whose singular subspaces are projective spaces.*[6]

Let $\Gamma = (\mathcal{P}, \mathcal{L})$ be a non-degenerate polar space. Let M be a maximal singular subspace. For each point $x \in \mathcal{P} - M$, $H_x := x^{\perp} \cap M$ is a proper subspace M (otherwise $M \cup \{x\}$ is a larger clique against the maximality of M among singular subspaces) and by (P), H_x must non-trivially intersect each line of M. Thus H_x is a geometric hyperplane of M. Let $\mathcal{H} = \{H_x | x \in \mathcal{P} - M\}$ denote the full collection of such hyperplanes.

Lemma 7.3.7 *The system of hyperplanes $\mathcal{H}$ satisfies the following:*

1. $\cap\{H | H \in \mathcal{H}\} = \emptyset$.
2. *Teirlinck's condition: if H_1 and H_2 are distinct members of $\mathcal{H}$, and p is a point of M not in $H_1 \cup H_2$, then there is a member of $\mathcal{H}$ containing $\{p\} \cup (H_1 \cap H_2)$.*[7]

Proof 1. Obviously if r were a point in the intersection of all H_x, $x \in \mathcal{P} - M$, then $r^{\perp} = \mathcal{P}$, so $r \in \text{Rad}(\Gamma)$, which is empty. So the first part holds.

---

[6] Of course the converse statement does not hold.

[7] See condition (H2) of Sect. 4.2, p. 83.

2. Now suppose H_x and H_y are distinct members of $\mathcal{H}$, and let p be a point of $M - H_x \cap H_y$. Now choose a point $u \in H_x - H_y$ (the reader should be able to explain why this is possible) and let N be a line containing u and p. Now it is clear that $N \cap (H_x \cap H_y) = \emptyset$, for otherwise N would contain two points of H_x, and so would lie in H_x, against the fact that it contains p.

Form the subspace $S = (H_x \cap H_y)^\perp$. Then S is a polar space containing N, x, and y and contains $H_x \cap H_y$ in its radical. Suppose now $r \in \text{Rad}(S)$. Then $r \in M^\perp$, and since M is a maximal clique of the collinearity graph of Γ, we must have $r \in M$. On the other hand, $x \in x^\perp \cap y^\perp$, so $r \in H_x \cap H_y$. Thus $H_x \cap H_y = \text{Rad}(S)$.

Now that means $N \cap \text{Rad}\,(S) = \emptyset$ and so by Lemma 7.3.2, part 2, S contains a line L opposite N, and L, in turn, contains a point z with $z^\perp \cap N = \{p\}$. Now it is clear that H_z belongs to $\mathcal{H}$ and contains $\{p\} \cup (H_x \cap H_y)$. The proof is complete. □

It now follows from Corollary 4.2.14 of Chap. 4 that any maximal singular subspace M of Γ is a projective space.

To prove Theorem 7.3.6 it remains only to address the partial linearity. But if L_1 and L_2 were distinct lines meeting in at least two points, $L_1 \cup L_2$ would be a clique, by (P), and so would lie in a maximal clique which is a maximal singular space which, as just observed, is a linear space containing L_1 and L_2, a contradiction. Thus any two distinct lines intersect in at most one point.

Remark This very neat proof of Theorem 7.3.6 is due independantly to F. Buekenhout (personal communication) and N. Percsy [99]. The author first learned of it from Peter Johnson.

7.3.5 The Rank of a Polar Space

In some sense, we propose to attach to each polar space P, a number called its *polar rank* (or just *rank* if it is clear that this adjective applies to a polar space). One would like this to be the matroid rank of the poset of all maximal singular subspaces of the associated non-degenerate polar space "$P/\text{Rad}(P)$" (that is, one more than the projective dimension of one of its maximal singular subspaces) and the lemma below shows that that certainly makes sense if all these maximal singular subspaces of the specified non-degenerate polar space have the same rank.

Lemma 7.3.8 *Suppose P is a non-degenerate polar space. Suppose M is a maximal singular subspace of Γ of finite projective rank r. Then all maximal singular subspaces of Γ possess the same projective rank.*

Proof Suppose N is a maximal singular subspace with $M \cap N = \emptyset$. Then, as shown, M and N are projective spaces. The Veldkamp space of a projective space N has as points and lines all hyperplanes an all codimension-two subspaces, and is usually called the the *dual space of* N, and is denoted by N^*. As observed in the previous section, the mapping $M \to N^*$ which takes a point x of M to a hyperplane $x^\perp \cap N$ of N, also takes lines L of M to codimension 2 subspaces of N, and so induces an embedding $M \to N^*$ as projective spaces. This forces the projective rank of M to

be bounded by that of N^* which, being finite, is also that of N. By the symmetry of the hypothesized relation between M and N, they have the same projective rank. □

Thus, for a non-degenerate polar space having a maximal singular subspace which is a projective space of finite dimension r, all its maximal singular subspaces have the common projective dimension r; we call $r + 1$ the *polar rank* of the polar space. We extend this notion to possibly degenerate polar spaces P by defining the *polar rank* of P to be that of the non-degenerate section $P/\mathrm{Rad}(P)$. Thus, when P possesses a maximal singular subspace in which the radical has finite codimension, we define the *polar rank* of P to be the common finite projective dimension of the maximal singular subspaces of its non-degenerate section $P/\mathrm{Rad}(P)$.

Just to touch base, a polar space of rank zero is equal to its radical. A non-degenerate polar space of rank one is just a non-empty collection of at least two points, no two being collinear – in fact there are no lines. A polar space of rank at least 2 is just one that contains a line L disjoint from its radical – in particular, a non-degenerate polar space of rank two is just a generalized quadrangle which is not a star. Similarly, a polar space of rank at least 3 possesses a projective plane A sharing no point with the radical, and in particular, any non-generate polar space which possesses a projective plane as a maximal singular subspace must have polar rank three.

However, the conclusion of Lemma 7.3.8 fails miserably when polar space P contains a maximal singular subspace which contains the radical at infinite codimension below it. (In the example which follows shortly one obtains a non-degenerate polar space having two sorts of maximal singular subspaces – one having countably infinite projective rank, and the other having uncountable projective rank.) In that case we say that P has *infinite polar rank*, without specifying any particular cardinal number.

Note that P and $P/\mathrm{Rad}(P)$ always possess the same rank (a non-negative integer or the label "infinite rank"). Now consider the following.

Example 2 Fix a division ring D. Let Π be the set of all sequences $b = \{\beta_i \in D | i \in \mathbf{N}\}$. Similarly, let Σ be all sequences $a = \{\alpha_i | i \in \mathbf{N}\}$ of elements of D for which only finitely many of the α_i are non-zero. Then Σ and Π are right vector spaces over D, called the *direct sum* and *direct product*, respectively, of countably many copies of D. We formally define the direct sum $V = \Sigma \oplus \Pi$ as right D-vector spaces. Now let σ be a fixed antiautomorphism of D of order two – that is $\sigma^2 = 1_D$.[8] We first define a function

$$f : \Sigma \times \Pi \to D,$$

by the rule that if $a = \{\alpha_i\} \in \Sigma$ and $b = \{\beta_i\} \in \Pi$, then

[8] Note that antiautomorphisms are automorphisms when D is a field. Such an antiautomorphism of order two is often called an "*involution of D*" (although elements of order two in the multiplicative group D^* of non-zero elements of D would appear to be just as deserving of that name).

$$f(a,b) = \sum\nolimits_{i\in\mathbf{N}} \alpha_1^{\sigma}\beta_i.$$

This sum is finite since only finitely many of the α_is are non-zero. Now we can define a so-called σ-Hermitian form

$$h : V \times V \to D$$

by the rule that for any $(u, v) \in V \times V$, with $u = a_1 + b_1$ and $v = a_2 + b_2$, where $a_i \in \Sigma$ and $b_i \in \Pi$, $i = 1, 2$, we have

$$\begin{aligned} h(u,v) &= h(a_1 + b_1, a_2 + b_2) \\ &:= f(a_2, b_1)^{\sigma} + f(a_1, b_2). \end{aligned}$$

In Exercise 7.9 the student is asked to prove:

1. That h is additive in each of its arguments
2. That for all scalars $\gamma, \delta \in D$,

$$h(u\gamma, v\delta) = \gamma^{\sigma} h(u,v)\delta$$

3. That

$$h(u,v) = h(v,u)^{\sigma}$$

for all vectors u, v in V

A subspace U of V is said to be *h-isotropic* if $h(U, U) = \{0\}$. It is also left to Exercise 7.9 for the student to show the following:

1. If $\mathcal{P}$ and $\mathcal{L}$ are the h-isotropic one- and two-dimensional subspaces of V, then $\Gamma = (\mathcal{P}, \mathcal{L})$ is a non-degenerate polar space.
2. The subspaces Σ and Π are both maximal singular subspaces. (Note that Σ has countably infinite dimension, while Π has uncountable dimension.)

Remark Note that our definition of rank is a little crude. It does not single out particular infinite cardinals. The rank is either a non-negative integer or is simply declared to be "infinite." However, this is enough to suit our purposes.

At this point, students may test their intuition for the polar space definitions by attempting Exercises 7.1–7.11 at the end of this chapter.

7.4 Oriflame Polar Spaces of Rank Three

In Theorem 7.5.13 of Sect. 7.5 we will prove Tits' theorem showing that non-degenerate polar spaces of rank at least three are completely determined by any "cone," $p^{\perp}$, the hyperplane of all points collinear with p. In order to prove that

theorem, it is essential to have a lemma asserting that if two planes A and B of such a space meet at a line L, then there is an isomorphism of planes $A \to B$ fixing L pointwise and carrying any prescribed point of $A - L$ to a designated point of $B - L$. (A corollary would be that all planes of non-degenerate polar spaces of rank at least three are Moufang.) However, the generic proof of that lemma breaks down when the polar rank is exactly three and there is a line L which lies in only two planes. We propose to handle that situation in this section.

It happens that this is an early case in which the student can see a complete classification of a class of polar spaces using only elementary principles that we have developed so far.

Suppose now that S is a non-degenerate polar space of polar rank exactly 3. Then by the results of the previous section, maximal singular subspaces of S are projective planes and every line lies in at least two of them.

The non-degenerate rank three polar space S is said to be *oriflame* if and only if it contains a line L which lies in exactly two planes.[9] This is a strong condition. Consider the following.

Lemma 7.4.1 *If S is oriflame, then every line of S lies in exactly two planes.*

Proof Suppose L lies in exactly two planes, and let N be any line meeting L at a point p. Let

$$\mathrm{Res}_S(p) := (\mathcal{L}_p, \Pi_p)$$

be the point-line geometry whose "points" and "lines" are the sets of lines and planes of S which are incident with the point p. (Such a geometry is called a *point residual* and is itself a polar space.) Now for any plane A on p and line M on p not in A, $M^\perp \cap A$ is the unique line of A on p which shares a plane with M. This is another way of saying that $\mathrm{Res}_S(p)$ is a generalized quadrangle. It has thick "lines" since the planes of S have thick lines, and so have at least three lines on each of their points. But since it is non-degenerate and has a "point" L on exactly two "lines" this quadrangle must be a grid, and so all its "points" are on exactly two "lines." In particular line N is in exactly two planes of S.

But for any line M of S which does not meet L at a point, the polar space axiom (P) implies the existence of a line N meeting both L and M at points. By the previous paragraph, N lies in just two planes, and, repeating the argument, M also lies in just two planes. The proof is complete. □

For the rest of this section S is a non-degenerate polar space of rank three which is oriflame.

Lemma 7.4.2 *Suppose A_1, A_2, and A_3 are planes with $A_i \cap A_{i+1} = \{p_i\}$, a point, for $i = 1, 2$. Then either $A_1 = A_3$ or $A_1 \cap A_3$ is a single point.*

---

[9] Interpreted as ascending chains of flags, one realizes the whole thing is a pleasant pun on the word "flag," for "oriflame" historically refers to the French battle flag whose shape was a non-convex pentagonal figure ending in just two points on its right end – or perhaps the left end, depending upon which way the wind is blowing.

Proof We may assume $A_1 \neq A_3$.

Case 1. Suppose $p_1 = p_2$. Then $p_1 \in A_1 \cap A_3$. Suppose $A_1 \cap A_3$ were a line N on p_1. Then $N^\perp \cap A_2$ is a line M on p_1 in A_2, and then $C = \langle N, M\rangle_S$ will be a third plane on N, a contradiction. Thus $A_1 \cap A_3 = \{p_1\}$ in this case.

Case 2. Assume $p_1 \neq p_2$. Then ${p_1}^\perp \cap A_3$ is a line of A_3 on p_2, and ${p_2}^\perp \cap A_1$ is a line of A_1 on p_1. The two planes $\langle p_1 p_2, {p_1}^\perp \cap A_3\rangle_S$ and $\langle p_1 p_2, {p_2}^\perp \cap A_1\rangle_S$ are distinct from A_2 and so must concide. Then the lines ${p_1}^\perp \cap A_3$ and ${p_2}^\perp \cap A_1$ of this common plane must intersect at a point p_3 in $A_1 \cap A_3$. But if $A_1 \cap A_2$ were a line, that line would have to be the line ${p_2}^\perp \cap A_1$, already containing p_1, and that would give $A_2 \cap A_3 \supseteq p_1 p_2$, contrary to our hypothesis. Thus $A_1 \cap A_3 = \{p_3\}$. □

At this stage, the relation among planes of being equal or intersecting in exactly one point is an equivalence relation which we shall denote by "$\sim$". Let us fix an equivalence class $\mathcal{C}$ with respect to this relation. Note that, since each line is on two planes, $\mathcal{C}$ is a proper subset of the set of all planes.

Lemma 7.4.3 *The planes of S fall into just two equivalence classes $\mathcal{C}$ and $\mathcal{C}'$ with respect to the relation "$\sim$," where $A \sim B$ if and only if $A = B$ or $A \cap B$ is a single point.*

Proof Fix $A \in \mathcal{C}$. If B_1 and B_2 intersect A at distinct lines N_1 and N_2 of A, then $B_1 \cap B_2$ cannot contain a line since otherwise the "perp" of this line would contain all of $\langle N_1, N_2\rangle_S = A$ against A being a maximal singular subspace. Thus $B_1 \cap B_2$ is a single point. Thus all planes meeting A at a line belong to a single class $\mathcal{C}'$. Now suppose D is any further plane not in $\mathcal{C}$. Then $D \cap A = \emptyset$. Choose a point $d \in D$, and form the planes $\langle d, d^\perp \cap A\rangle_S = B$. Then $B \cap D = \{d\}$ and $D \sim B \in \mathcal{C}'$. Thus all planes not in class $\mathcal{C}$ are in class $\mathcal{C}'$. The proof is complete. □

Now form the rank two geometry

$$\mathbf{P} = (\mathcal{C}, \mathcal{P})$$

with the ordinary incidence of S between planes and points. We wish to think of $\mathbf{P}$ as a geometry of "points" ($\mathcal{C}$) and "lines" ($\mathcal{P}$). Note that since every line of S lies in a plane of each class, one must conclude that each point p is in at least three planes of $\mathcal{C}$ (recall that $\mathrm{Res}_S(p)$ was a grid with thick lines). Since any two planes in $\mathcal{C}$ intersect in exactly one point, we see that $\mathbf{P}$ is a linear space with thick lines.

We set out to verify the Veblen axiom for this linear space. The hypothesis for that axiom on "points" and "lines" is laid out in Fig. 7.1. The three "lines" r, p,

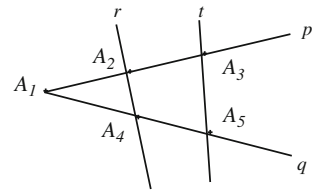

Fig. 7.1 The set-up for the Veblen axiom in the linear space $\mathbf{P}$

and q are points of S which pairwise belong to distinct planes (A_1, A_2, A_4) of $\mathcal{C}$, and so together generate a plane B_r in $\mathcal{C}'$. Similarly, "lines" t, p, and q, as points of S, generate plane B_t of $\mathcal{C}'$. Since B_r and B_t are two planes of $\mathcal{C}'$ on line pq, these planes must coincide. This forces point r to be collinear with point t in S, and since they are distinct, they lie together in a unique plane A_6 in $\mathcal{C}$. That means that the "lines" r and t of Fig. 7.1 meet at the "point" A_6 (not portrayed in Fig. 7.1), so the Veblen axiom holds.

Thus the following applies.

Lemma 7.4.4 $\mathbf{P} = (\mathcal{C}, \mathcal{P})$ *is a projective space.*

For each polar point $p \in \mathcal{P}$ let the symbol $\mathcal{C}(p)$ be the set of planes in $\mathcal{C}$ which contain point p. Thus $\mathcal{C}(p)$ is the "point"-shadow of the "line" p in the projective space $\mathbf{P}$.

Now fix a plane B in $\mathcal{C}'$ and let $N(B)$ be the collection of all planes of $\mathcal{C}$ which meet B at a line. Note that if A_1 and A_2 are distinct members of $N(B)$, then $A_1 \cap B \cap A_2 = A_1 \cap A_2 = \{p\} \in B$ and so $\mathcal{C}(p) \subseteq N(B)$. Thus $N(B)$ is a subspace of $\mathbf{P}$.

In fact, by our oriflame assumption, we have a bijection

$$\text{lines of } B \rightarrow N(B),$$

and another bijection

$$\text{points of } B \rightarrow \text{ "lines" of subspace } N(B) \text{ of } \mathbf{P},$$

with incidences being preserved. Thus the following applies.

Lemma 7.4.5 $N(B)$ *is a plane of the projective subspace* $\mathbf{P}$ *isomorphic to the dual* B^* *of the plane* B *of* S.

Next we note that if p is any point of S, then either p belongs to B or $\langle p, p^{\perp} \cap B \rangle_S$ is a unique plane A of $N(B)$, and in the latter case $\mathcal{C}(p) \cap N(B) = \{A\}$. Thus the following is true.

Lemma 7.4.6 *The plane* $N(B)$ *of the projective space* $\mathbf{P}$ *is a geometric hyperplane of* $\mathbf{P}$. *Thus* $\mathbf{P}$ *is the classical projective space* $PG(3, D)$ *for an appropriate division ring* D.

We can now identify the points of S with the projective lines of $\mathbf{P}$, and identify the two classes of planes $\mathcal{C}$ and $\mathcal{C}'$ with the points and planes of $\mathbf{P}$. Since each line of S is the set of points in two unique planes, one in $\mathcal{C}$ and one in $\mathcal{C}'$, we see that each line of S is the collection of lines of $\mathbf{P}$ which are incident with a point-plane flag of $\mathbf{P}$.

The geometry Γ whose points are the lines of $PG(3, D)$ and whose lines are the point-plane flags of $PG(3, D)$ is called the *Grassmannian of type* $A_{3,2}(D)$. So we have shown the following.

Theorem 7.4.7 *A non-degenerate polar space of rank three which is oriflame is isomorphic to a Grassmannian of type* $A_{3,2}(D)$. *All its planes are Desarguesian and isomorphic.*

Remark The last comment in the statement of Theorem 7.4.7 follows from the fact that each plane B in $\mathcal{C}'$ is a plane of $PG(3, D)$ and that each plane of $\mathcal{C}$ is isomorphic to the dual of any plane of $\mathcal{C}'$ which it meets trivially.

7.5 Further Properties of Abstract Polar Spaces

Before being able to approach the classification of non-degenerate polar spaces of rank at least four, still further properties are needed. They still follow from first principles in an elementary way, occasionally assisted by Exercises 7.1–7.9. This section is sprinkled with a few comments intended to anticipate how these properties will be needed later on.

The first few results involve geometric hyperplanes of polar spaces of polar rank at least three. The next establishes the $\mathcal{T}$-simply-connectedness of the point-collinearity graph induced on a hyperplane complement in a polar space of polar rank at least three. (Here $\mathcal{T}$ is the collection of all circuits of length three.) That result is then used in a discussion of lifting automorphisms on a subspace of a polar space to the whole space and in Tits' theorem, showing that polar spaces of rank four or more are characterized by the cone $p^{\perp}$.

7.5.1 Geometric Hyperplanes of Polar Spaces

Rank and the Veldkamp Conditions

An oft-proved lemma states that if H is a geometric hyperplane of a non-degenerate generalized quadrangle with thick lines, then the subgraph of the point-collinearity graph induced on the complement $\mathcal{P} - H$ is connected (see, for example [38]). This means all geometric hyperplanes are maximal subspaces (Lemma 4.1.1), which in turn can be rephrased as the assertion that Veldkamp points exist. An easy induction shows the following.

Lemma 7.5.1 *The following statements hold:*

1. *Suppose* Γ *is a polar space with thick lines of rank at least* 2. *Then Veldkamp points exist.*
2. *Suppose the polar space* Γ *with thick lines has rank at least* 3. *Then Veldkamp lines exist.*
3. *In general, if* Γ *is a polar space of rank* r, *then Veldkamp* $(r-1)$*-spaces exist for* S.

Proof Part 1. Let H be any geometric hyperplane. It suffices to show that the point-collinearity graph of $\mathcal{P} - H$ is connected. By way of contradiction choose x and y

in distinct connected components of $\mathcal{P} - H$. Then neither x nor y lies in the radical, and $x^\perp \cap y^\perp \subseteq H$. Also, $x^\perp \cap y^\perp$ is not a clique, since Γ has rank at least 1. So there is a non-collinear pair of points (u, v) in $x^\perp \cap y^\perp$. Let A and B be lines on $\{x, v\}$ and $\{x, u\}$ and let C be a line containing $\{v, y\}$. Then $H \cap A = \{u\}$ and $H \cap C = \{v\}$. Since C is a thick line, there is a point y' in $C - \{y, v\}$, belonging to the same connected component as y. By axiom (P), y' is collinear with a point w of line A. But if w were the point u, then $u^\perp$ would contain y and y' and hence would contain all points of C – in particular v, contrary to our choice. Thus w must be a point of $A - H$, and now we have a path (x, w, y', y) connecting x to y.

Part 2. We need only show that if A_1 and A_2 are distinct geometric hyperplanes of a polar space Γ, then $A_1 - (A_1 \cap A_2)$ has a connected collinearity graph. Now it is an easy observation that if G is a polar space of rank r, then any geometric hyperplane A is also a polar space with rank at least $r - 1$. So, if Γ has rank at least three, then $A_1 \cap A_2$ is a geometric hyperplane of a polar space of rank at least 2. So the indicated set is connected by part 1.

Part 3. The general induction proof is left to the student (Exercise 7.11). □

Remark In fact we could prove that if Γ had rank at least 4 then Veldkamp planes would exist. If we could also prove that Teirlinck's condition (Teir)[10] held, we would then know that the Veldkamp space $\mathcal{V}(\Gamma)$ was a projective space as at the end of Chap. 4. However the world of possible geometric hyperplanes of an arbitrary polar space is rather unknowable, and proving (Teir) is a very complicated process (it occupies at least 19 steps in one manuscript). Fortunately, most of the theory surrounding (Teir) involves a *family of hyperplanes*, not *all* geometric hyperplanes. The fact is, that one has a better chance of proving (Teir) for a tightly controlled family of hyperplanes $\mathcal{H}$ rather than for all hyperplanes. We did that in the Buekenhout/Percsy argument of the previous section when considering maximal singular subspaces of a non-degenerate polar space. We had no assurance that the $x^\perp \cap M$ reproduced *all* geometric hyperplanes of the singular space M – nor did we care. We just had a system of hyperplanes that worked.

One of the goals of this chapter is to provide at least a self-contained classification of the non-degenerate polar spaces of rank at least 4. It will turn out that (Teir) is much easier to prove for a naturally-occurring family $\mathcal{H}$ of hyperplanes of a cone $p^\perp$ rather than for the complete Veldkamp space. Just as in the case of a maximal singular subspace, the *environment* of the embedded subspace itself provides the nice family of hyperplanes satisfying (Teir).

Thanks to a vital unpublished theorem of Tits' (Theorem 7.5.13 ahead), at this rank, classifying the cone classifies the non-degenerate polar space.

So we shall not really need to know anything about Veldkamp planes existing when the rank is four even though these (or at least their non-degenerate sections) comprise the target polar spaces of the classification theory presented here.

[10] That is, the conditions (H1) and (H2) of p. 83, Sect. 4.2.

7.5.2 Hyperplane Complements in Polar Spaces of Rank at Least Three

For the next two lemmas and their corollary, $P := (\mathcal{P}, \mathcal{L})$ is a non-degenerate polar space of rank at least two having thick lines.

Lemma 7.5.2 *Let H be a geometric hyperplane of P. Then the point-collinearity graph induced on $\mathcal{P} - H$ has diameter at most* 3. *In particular, since it is connected, H is a maximal subspace of P.*

Proof This statement is equivalent to the assertion that the collinearity graph induced on $\mathcal{P} - H$ is connected. The argument appears in the proof of part 1 of Lemma 7.5.1 of the previous section. □

Lemma 7.5.3 *Suppose H_1 and H_2 are distinct geometric hyperplanes of P. Let $\mathcal{L}_{12}$ be the collection of lines of P which connect a point of $H_1 - H_2$ with a point of $H_2 - H_1$.*

Then $\mathcal{L}_{12}$ is non-empty.

Proof Assume $\mathcal{L}_{12}$ is empty. Because Veldkamp points exist, neither of the H_i is contained in the other, so both sets $H_1 - H_2$ and $H_2 - H_1$ are non-empty. Consider a point $x \in H_1 - H_2$. Then as $\mathcal{L}_{12}$ is empty, $x^\perp$ is contained in H_1 and by maximality of the hyperplane subspaces $x^\perp = H_1$. In particular x is not in Rad(P). These statements hold for arbitrary $x \in H_1 - H_2$.

Choose a particular $x_1 \in H_1 - H_2$. Since P has rank at least 2, $x_1^\perp \cap H_2$ contains two non-collinear points u and v. A line L on x_1 and u carries a third point x_2, not in H_2. Since also $H_1 = x_2^\perp$, x_2 is collinear with v. Then v is collinear with the two points x_1 and x_2 of L and so by axiom (P) is collinear with u, a contradiction.

Thus $\mathcal{L}_{12}$ is non-empty. □

Corollary 7.5.4[11] *P is never the union of two hyperplanes.*

Proof Suppose H_1 and H_2 were the two hyperplanes whose union was $\mathcal{P}$. Let the non-empty $\mathcal{L}_{12}$ be defined as in Lemma 7.5.3. Any thick line $\mathcal{L}_{12}$ must bear a point in neither H_1 nor H_2. The proof is complete. □

Let Δ be any graph, and let $\mathcal{T}$ be the class of all triangles (3-cliques) of Δ. Recall that the phrase "simply connected" (without any explicitly specified collection of cycles) will always means $\mathcal{T}$-simply connected (see Chap. 1 for definitions of these terms).

Theorem 7.5.5 *Assume P is a polar space of rank at least three with thick lines. Let H be any geometric hyperplane of P. Let Γ be the point-line collinearity graph induced on the vertex set $\mathcal{P} - H$. Then Γ is simply connected.*

---

[11] As far as the author is aware, this minor but useful corollary is due to Jonathan Hall, first appearing in [65].

Proof It suffices to show that any circuit of Γ is $\mathcal{T}$-contractible. Suppose $c = (x_0, \ldots, x_n = x_0)$ is a non-contractible circuit of minimal length n. Then $n \geq 4$ and if $|i - j| \geq 2$, then x_i is not collinear with x_j. If $P - H$ contained a point r in the radical of P, then we would have edges in Γ connecting r to any vertex of the circuit c, and c would be contractible. So we can assume Rad $(P) \subseteq H$.

The proof is divided into two cases, $n > 4$ and $n = 4$.

Case 1. $n > 4$. We draw a portion of the circuit in Fig. 7.2, where $x_4 \neq x_0$.

Let L be a line on x_4 and x_3. It r were a point of Rad(P) on L, then r would be distinct from x_3 since it belongs to H. But in that case x_2 would be collinear with both x_3 and r and hence collinear with x_4, by axiom (P). That is against the minimality of c.

Thus the line L does not meet the radical. Since, under the basic morphism $\beta : P \to P/\text{Rad}(P)$ onto a non-degenerate polar space, L maps into a plane, there exists a plane A on L in P which meets the radical trivially.

Set $Y = x_0^{\perp} \cap x_2^{\perp}$, a polar space of rank at least 2 having thick lines. Suppose, for some point $z \in Y - H$ we could find a point $u \in z^{\perp} \cap (A - H)$. Then we would have the circuit decomposition of Fig. 7.3.

Then c would decompose into circuits of smaller length against the assumption that c was non-contractible of minimal length at least five. Thus we must always have:

$$\textit{For each point } z \in Y - H,\ z^{\perp} \cap A = A \cap H. \tag{7.3}$$

Let N denote the line $A \cap H$. So, for each point h in N, we see that $h^{\perp}$ contains $Y - H$. Now if $h^{\perp} \cap Y \neq Y$ we see that Y is a polar space covered by the two geometric hyperplanes $h^{\perp} \cap Y$ and $H \cap Y$ against Corollary 7.5.4. Thus we have

$$h^{\perp} \supseteq Y.$$

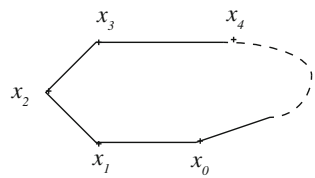

Fig. 7.2 The initial circuit for the case $n > 4$

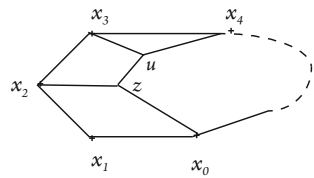

Fig. 7.3 A decomposition for $u \in A - H, z \in u^{\perp} \cap Y - H$. Note that $u = x_3$ or x_4 is possible

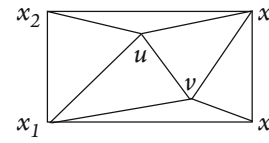

Fig. 7.4 A decomposition for the case $n = 4$

So $N \subseteq Y^{\perp}$. Now x_2 is collinear with a point h_0 of N by axiom (P). Thus $h_0^{\perp}$ contains $\langle x_2, Y\rangle = x_2^{\perp}$. This means any line M on h_0 and x_2 cannot contain an opposite line, and so by Lemma 7.3.2 this line M must intersect the radical of P. But we have seen that the radical of P is in hyperplane H, and so this radical point is $M \cap H = h_0 \in A$. But this contradicts our choice of A as a plane sharing no point with the radical.

Case 2. $n = 4$. The circuit is $c = (x_0, x_1, x_2, x_3, x_0)$. Then $Y = x_0^{\perp} \cap x_2^{\perp}$ is a polar space of rank at least two, and so, by Lemma 7.5.2, $x_0 \cap x_2^{\perp} - H$ is connected, and so contains a path (x, u, v, x_3) or a shorter one in which $u = v$. This yields the triangulation of c given in Fig. 7.4. This completes the proof.

Corollary 7.5.6 *The point-collinearity graph of any affine polar space of rank at least three (in the sense of Cohen and Shult* [38]*) is simply connected.*

Corollary 7.5.7 *Suppose $\Delta = (\mathcal{P}, \mathcal{L})$ is a non-degenerate polar space of rank at least three with thick lines. If $e : \Delta \to \mathbf{P}(V)$ is any embedding, then every geometric hyperplane of Δ arises from the embedding $\hat{e}$, the relatively universal hull of e.* (See p. 94 for relevant definitions.)

7.5.3 Automorphisms of Non-degenerate Polar Spaces

The purpose of this section is to indicate why non-degenerate polar spaces tend to possess many automorphisms and in particular to prove Tits' theorem on extending automorphisms from a cone. The latter is needed for our final classification theorem.

Using this fact, it will be easier to display all of the desired phenomena for non-degenerate polar spaces. As usual, all lines are thick, and the rank of the polar space is at least three.

Automorphisms on Planes

Lemma 7.5.8 *Let S be a non-degenerate polar space of rank at least* 3*. Suppose A and B are planes meeting at a line L and a and b are points of $A - L$ and $B - L$, respectively. Then there is an isomorphism of planes*

$$\alpha_{a,b,L} : A \to B,$$

fixing line L pointwise and taking a to b.

Proof First, the result is easily seen to be true if the planes A and B are isomorphic and Desarguesian. Thus, by Theorem 7.4.7, we may assume that S is not oriflame of rank three. Then by Lemma 7.4.1 we can assume that every line of S lies in at least three planes. We shall use this fact.

If the two planes already lie in a larger singular subspace of S, they are planes of some $PG(3, D)$ and the desired isomorphism $\alpha_{a,b,L}$ exists through a perspectivity as in Chap. 6. So we may assume that point a is not S-collinear with point b.

Now by Exercise 9.10 , at the end of this chapter, the subspace $a^\perp \cap b^\perp$ is a non-degenerate polar space of rank at least two, containing the line L. By Lemma 7.3.2, there is a line N which is opposite L in $a^\perp \cap b^\perp$.

We claim that the line N lies in a plane D which is not in $a^\perp$ nor in $b^\perp$.

First assume S has rank exactly three. Since the line N lies in at least three planes, it is in a plane D distinct from the two planes $\langle a, N\rangle$ and $\langle b, N\rangle$. Since D is a maximal singular space, D is not in $a^\perp$ nor $b^\perp$.

On the other hand, assume S has rank at least four. Then the polar space $N^\perp/N$ is non-degenerate of rank at least two. Moreover $(\{a\} \cup N)^\perp/N$ and $(\{b\} \cup N)^\perp/N$ are two geometric hyperplanes of this polar space. By Corollary 7.5.4, there is a "point" D/N not in either hyperplane. That means D is a plane containing the line N, not in $a^\perp$ or in $b^\perp$, as desired.

Now if a' were a point of $D^\perp \cap A$, then by our choice of D, $a \neq a'$. But then the line aa' of the projective plane A must intersect its other line L at a point t in $N^\perp$, against N being opposite L. Thus $D^\perp \cap A = \emptyset$, and similarly $D^\perp \cap B = \emptyset$, – that is, D is a plane opposite both A and B.

Now axiom (P) forces isomorphisms

$$\alpha_{A,D} : A \to D^*,\ \alpha_{B,D} : B \to D^*$$

of A and B to the dual plane D^* of D, with a and b mapping to line N, and points of L mapping bijectively to the lines of D on $\{z\} := L^\perp \cap D$. It is easy to see that the composition

$$\alpha_{a,b,L} := (\alpha_{B,D})^{-1} \circ \alpha_{A,D}$$

has the desired properties. □

Corollary 7.5.9 *(Tits) Let S be a non-degenerate polar space of rank at least* 3. *Then the planes of S are Moufang.*

Proof If S has rank at least four, the planes lie in larger projective (singular) subspaces and so are Desarguesian, and hence Moufang. So we may assume S has rank three exactly, so planes are maximal singular subspaces.

Let A be any plane of S. Fix one of its lines L and choose two distinct points a_1 and a_2 of $A - L$ and let $c = a_1a_2 \cap L$. We must show that a_1 can be moved to a_2 by a perspectivity of A, with axis L.

Since $L^\perp$ is not a clique, there is a plane B on L distinct from A. Choose $b \in B - L$. Now the composition

$$(\alpha_{a_2,b,L})^{-1} \circ \alpha_{a_1,b,L}$$

of the morphism of the previous Lemma is the desired perspectivity. □

Automorphisms of the Cone

Theorem 7.5.10 *Let S be a non-degenerate polar space of rank at least three. Fix a point p and let H and H' be distinct geometric hyperplanes of the cone $C = p^{\perp}$, neither of which contains the point p. Then there is a unique automorphism α of C (as a point-line geometry) taking H to H', stabilizing every line on p, and fixing pointwise those lines on p which meet $H \cap H'$.*

Proof Consider any plane A of S on point p. Then A meets H and H' at lines which either coincide or else meet at a unique point q of $A \cap H \cap H'$. In the former case let α_A be the identity automorphism of A. In the latter case, since A is Moufang, there is a unique elation α_A with axis pq and center p taking line $A \cap H$ to $A \cap H'$. Thus for each plane A, there is a unique elation α satisfying the required conditions, at least for the plane A.

We need to show that these plane mappings agree on intersections of these planes. If that is true, then the desired automorphism α of C is defined.

Suppose A and B are two planes of C on p which meet at a line L. We must show that α_A and α_B agree on L. If L carries a point of $H \cap H'$, both maps fix L pointwise and so agree. Thus we must assume there are unique points $q_A := A \cap H \cap H'$ and $q_B := B \cap H \cap H'$ in A and B respectively, which are not on line L.

Now, by Lemma 7.5.8, there is an isomorphism of planes $\gamma : A \to B$ fixing line L pointwise, and mapping q_A to q_B. Then $\gamma \circ \alpha_A \circ \gamma^{-1}$ is an automorphism of plane B fixing line pq_B pointwise, fixing every line of B on p, mapping $q = H \cap L$ to $q' = L \cap H'$. But α_B is an elation accomplishing the same feat, and so by the uniqueness of such an elation, we have

$$\gamma \circ \alpha_A \circ \gamma^{-1} = \alpha_B.$$

Now since γ fixes L pointwise we have these equalities upon restriction to L:

$$\alpha_A|_L = \gamma \circ \alpha_A \circ \gamma^{-1}|_L = \alpha_B|_L.$$

So α_A and α_B agree on the intersection. The proof is complete. □

Extending Point-Line Morphisms of the Cone

The first two lemmas of this section sound a little technical, but, together with the triangulation results of the previous section, they do the whole job.

Lemma 7.5.11 *Suppose S is a non-degenerate polar space of rank at least three. Let q be a point not collinear to p, let a be a point of $p^{\perp} \cap q^{\perp}$, and observe that every member of the ascending chain*

$$a^\perp \cap p^\perp \cap q^\perp \subset p^\perp \cap q^\perp \subset p^\perp$$

is a geometric hyperplane of its successor.

Now let H be any hyperplane of $p^\perp$ which contains $a^\perp \cap p^\perp \cap q^\perp$. Then there exists a unique point x on the line aq such that $H = p^\perp \cap x^\perp$.

Proof First the hyperplane H of the cone $p^\perp$ cannot lie in $a^\perp \cap p^\perp \cap q^\perp$ itself, for the latter lies properly in $p^\perp \cap q^\perp$, which is a non-degenerate hyperplane of $p^\perp$, and that would contradict the Veldkamp-points principle for $p^\perp$ (which must hold since $p^\perp$ has polar rank at least two). If $H = a^\perp \cap p^\perp$, then we may take $x = a$ in the theorem. So we may assume that H contains a point u not in $a^\perp \cap p^\perp$. Then u is not collinear with a. Thus u is collinear with a unique point x on the line qa, $x \neq a$. Then by partial linearity of S the lines xa and qa coincide, so

$$p^\perp \cap a^\perp \cap q = p^\perp \cap a^\perp \cap x^\perp.$$

Now

$$p^\perp \cap a^\perp \cap x \subset p^\perp \cap x^\perp \subset p^\perp$$

is also a chain with each member a hyperplane (and hence a maximal subspace) of its successor in the chain. There is also a chain

$$p^\perp \cap a^\perp \cap x^\perp \subset \langle u, a^\perp \cap p^\perp \cap q^\perp \rangle \subseteq p^\perp \cap x^\perp \tag{7.4}$$

with the first containment proper. It follows that the last containment is an equality. But we notice that the middle term of Eq. (7.4) is contained in H, so the hyperplane H of $p^\perp$ contains $x^\perp \cap p^\perp$, another hyperplane of $p^\perp$. Thus $H = p^\perp \cap x^\perp$ as we were required to prove. □

Lemma 7.5.12 *Suppose p_i is a point of a non-degenerate polar space S_i, of rank at least three, and q_i is a point not collinear with p_i, in S_i, $i = 1, 2$. Let Δ_i be the point-collinearity graph of S_i. Suppose*

$$f : {p_1}^\perp \cup q_1 \to {p_2}^\perp \cup q_2$$

is an isomorphism of the induced subgraphs of Δ_1 and Δ_2 defined by the range and domain vertex sets, which restricts to a point-line isomorphism $C_1 \to C_2$ on the cones $C_i = {p_i}^\perp$. (Note that the edges connecting the point q_1 to the cone $C_1 := {p_1}^\perp$ must be mapped to edges connecting q_2 to the cone $C_2 := {p_2}^\perp$ in S_2.) Then:

1. (The extension f^q of f from q_1.) *The mapping f can be extended to an isomorphism*

$$f^q : {p_1}^\perp \cup {q_1}^\perp \to {p_2}^\perp \cup {q_2}^\perp,$$

again, as an induced subgraph.

2. (Symmetry of extensions.) *If r_1 is a point of $S_1 - C_1$ collinear with q_1, the extension $(f^q)^r$ of f^q from r_1 (like f itself) takes q_1 to q_2.*
3. (Extensions f^q and $(f^q)^r$ from collinear q_1 and r_1 agree on the intersection of their domains: $C_1 \cup ({r_1}^\perp \cap {q_1}^\perp)$.) *If (q_1, r_1, s_1) is a triangle in $\mathcal{P} - {p_1}^\perp$, then $f^q(s_1) = (f^q)^r(s_1)$ – that is, if one were to apply part 1 to extend the restriction of f^q which induces the mapping*

$$ {p_1}^\perp \cup r_1 \to {p_2}^\perp \cup {r_2}^\perp $$

to a mapping (of induced graphs)

$$ (f^q)^r : {p_1}^\perp \cup {r_1}^\perp \to {p_2}^\perp \cup f^q(r_2)^\perp, $$

we should discover that the latter map, applied to s_1 yields $f^q(s_1)$.

Proof First how is f^q to be defined? We only need to know its value on any new point r_1 which could appear on a thick line joining point q_1 to a point a_1 in ${p_1}^\perp$. In that case, $H_{r_1} := {r_1}^\perp \cap {p_1}^\perp$ is a geometric hyperplane of ${p_1}^\perp$ which contains ${p_1}^\perp \cap {a_1}^\perp \cap {q_1}^\perp$. Taking f-images, and using the fact that f restriced to C_1 takes lines to lines, we learn that $f(H_{r_1})$ is a hyperplane of $f(p_2)^\perp = {p_2}^\perp$ which contains ${p_2}^\perp \cap f(a_1)^\perp \cap f(q_1)^\perp = C_2 \cap f(a_1)^\perp \cap {q_2}^\perp$. By the previous Lemma 7.5.11, there is a unique point r_2 on the line $f(a_1)f(q_1) = a_2q_2$ such that $f(H_{r_1}) = {r_2}^\perp \cap {p_2}^\perp$. This unique point r_2 is declared to be the image of r_1 under f^q.

f^q is now a well-defined bijective mapping with the stated domain and range. We must show that it preserves adjacencies in the induced subgraph. But that brings us to the second part of this proof.

If s_1 and r_1 are points on the line q_1a_1 of S_1, then, by construction, their images are both on the line, $q_2f(a_1)$ of S_2 and so are collinear.

Next suppose r_1 and s_1 are two collinear points of ${q_1}^\perp - {p_1}^\perp$ which do not lie on a line on q_1. This means that q_1r_1 is a line meeting ${p_1}^\perp$ at a point a_1, that q_1s_1 is a line meeting ${p_1}^\perp$ at a point b_1, and that r_1s_1 is a line meeting ${p_1}^\perp$ at a point c_1. Then the three points a_1, b_1, c_1 lie on a line L of ${p_1}^\perp$, namely the intersection of the singular plane $\langle q_1, r_1, s_1\rangle_{S_1}$ with ${p_1}^\perp$. The f-images of these points are denoted a_2, b_2, and c_2, respectively.

Then $f^q(r_1)$ and $f^q(s_1)$ lie on lines q_2a_2 and q_2b_2, respectively and so lie in the singular plane of S_2 generated by the 3-clique $\{a_2, b_2, q_2\}$ of S_2, and so are collinear. (At this point, part 1 is proved.)

Now note that by Part 1, the extension $(f^q)^r$, must take q_1 to the unique point q' on the line $a_2(f^q)^r(r_1) = a_2f^q(r_1)$ whose "perp" meets $C_2 \cap f^q(r_1)^\perp$ where ${a_2}^\perp$ meets it. But q_2 is a point on that line with that property. Thus from the uniqueness, $q' = q_2$ and part 2 is proved.

Now, by part 1, the restriction of f^q to $C_1 \cup \{r_1\}$ can be extended to

$$ (f^q)^r : C_1 \cup {r_1}^\perp \to C_2 \cup f^q(r_1)^\perp. \tag{7.5} $$

Let

$$y_q := f^q(s_1), \text{ and}$$
$$y_r := (f^q)^r(s_1),$$

that is, the images of s_1 under the two mappings f^q and $(f^q)^r$. Now let us look at the definition of these points.

Set $H := f(C_1 \cap {s_1}^\perp)$, a hyperplane of the cone C_2, not on p_2. Then:

1. y_q is the unique point on line b_2q_2 such that

$$ {y_2}^\perp \cap C_2 = H \text{ and } H \cap {q_2}^\perp = H \cap {b_2}^\perp. \tag{7.6}$$

2. Similarly, y_r is the unique point on line $c_2 f^q(r_1)$ with

$$ {y_r}^\perp \cap C_2 = H \text{ and } H \cap {q_2}^\perp = H \cap {c_2}^\perp. \tag{7.7}$$

Now, by way of contradiction, assume y_q is not equal to y_r. The two points are collinear for they lie in the planes of S_2 generated by $\{q_2, a_2, b_2\}$ and $\{q_2, a_2, c_2\}$, which are the same plane, since a_1, b_2, and c_2 are incident with a common line. So there is a line N on y_q and y_r which meets C_2 at a point y. But then, from the left sides of Eqs. (7.6) and (7.7), $y^\perp$ contains H and p_1, and so lies in $\mathrm{Rad}(C_2) = \{p_1\}$. That is impossible, as y_q is not in C_2.

Thus $y_q = y_r$, and the proof of part 3 is complete. □

Identifying Non-degenerate Polar Spaces by Their Cones

Theorem 7.5.13 (Tits' theorem.) *Suppose S_1 and S_2 are two non-degenerate polar spaces of rank at least three, with respective points p_1 and p_2. Let $C_i := {p_i}^\perp \cap S_i$, a cone in each space. Suppose there is an isomorphism $C_1 \to C_2$ as point-line geometries. Then there is an isomorphism $\hat{f} : S_1 \to S_2$. Such an isomorphism is uniquely determined by its restriction to C_1 and its value at q_1 where q_1 is any point of $S_1 - C_1$.*

Proof We are given an isomorphism β of subspaces $C_1 \to C_2$. Now choose any point q_1 in $S_1 - C_1$, form $H_1 := {q_1}^\perp \cap C_1$, a hyperplane of C_1 not on p_1. Similarly chose q_2 in $S_2 - C_2$ and form $H_2 := C_2 \cap {q_2}^\perp$. Then as β is a point-line isomorphism, $H' = \beta(H_1)$ is also a hyperplane of cone C_2 not on p_2. Now by Theorem 7.5.10, there is a point-line automorphism $\alpha : C_2 \to C_2$ carrying H' to H_2. Thus the composition $f := \alpha \circ \beta$ is an isomorphism $C_1 \to C_2$ of point-line geometries, taking $H_1 = {q_1}^\perp \cap C_1$ to $H_2 = {q_2}^\perp \cap C_2$, and so can be extended to an isomorphism of induced point-collinearity subgraphs

$$f : C_1 \cup \{q_1\} \to C_2 \cup \{q_2\}$$

by setting $f(q_1) = q_2$.

At this point, we have the hypotheses of Lemma 7.5.12, part 1, and so this mapping can be extended to the union of the cones $C_1 \cup q_1{}^{\perp}$. In fact, for any path $P(q_1, y_1) := (q_1 = q_{11}, q_{12}, \ldots, q_{1n} = y_1)$ in $S_1 - C_1$, we can extend f to the "double cone" $C_1 \cup q_{11}{}^{\perp}$, restrict this to $C_1 \cup \{q_{12}\}$, apply part 1 once more to extend this to a mapping f_2 on the "double cone" $C_1 \cup q_{12}{}^{\perp}$, restrict to $C_1 \cup \{q_{13}\}$, and iterate the extension procedure – until ultimately we have assigned a value $f_n(q_{1n})$. If we repeat this procedure using some other path $P'(q_1, y_1)$ from q_1 to y_1 in $S_1 - C_1$, do we get the same value assigned to y_1? Now part 3 of Lemma 7.5.12 says that if one path follows two sides of a triangle, and the other is the third edge of that triangle, the assigned values are the same. This means that if one of the paths is deformed by a homotopy of triangles, the ultimate assigned value of the mapping at y_1 iterating over the deformed path is the same as the value obtained from the original path.

Now Theorem 7.5.5 tells us that $S_1 - C_1$ is $\mathcal{T}$-simply connected, and so path $P'(q_1, y_1)$ is homotopic to $P(q_1, y_1)$. That means the procedure of iterating the double cone extensions, starting with the assignment $f(q_1) = q_2$, forces a unique assignment $f(y_1)$ for any point y_1 in $S_1 - C_1$, giving a mapping

$$f : S_1 \to S_2,$$

which is a graph isomorphism when restricted to any double cone $C_1 \cup y^{\perp}$, $y \in S_1 - C_1$. It follows that f is a graph homomorphism which is locally an isomorphism – that is, a $\mathcal{T}$-covering. Since it is an isomorphism when restricted to C_1, and $S_2 - C_2$ is simply connected, by Theorem 7.5.5, f is bijective and so is a graph isomorphism. Finally, f is then an isomorphism of point-line spaces since lines of S_i are sets $e^{\perp\perp}$, where e is an edge of the point-collinearity graph of the non-degenerate space S_i. The proof is complete. □

Remark Obviously, this theorem fails for polar spaces of rank two. The Kantor-family/flock constructions of generalized quadrangles mentioned in the appendix to this chapter show that for suitable prime powers q, there are non-isomorphic generalized quadrangles of order (q^2, q). Yet for these, there is always an isomorphism of their cones over a point.

As far as the author knows, Tits' Theorem on characterization of polar spaces by their cones has not been published. There is an account of it in Chap. 19 of the recent book "*Diagram Geometries*" by F. Bukenhout and A. Cohen. The above proof differs substantially from that one.

7.6 The Classical Polar Spaces

7.6.1 Morphisms from a Projective Space to Its Dual, or Where Sesquilinear Forms Come From

The first title of this section portrays a general and natural situation. For any projective space **P** there is always a dual space $\mathbf{P}^*$ whose points are the hyperplanes of **P**,

and whose lines are the second-maximal subspaces of **P**. The dual of a projective space **P** is simply its Veldkamp space, pure and simple.

There is not much to say about duals at the level of projective planes: the dual of a plane is another plane. So assume **P** is a projective space of rank at least 4 (projective dimension at least 3), where all planes are Desarguesian and self-dual. Thus we consider a classical projective space $\mathbf{P}(V)$ where V is a right vector space over a division ring D.

In this case there is another description of the dual $\mathbf{P}(V)^*$. Let V^* be the set Hom (V, D), of D-linear morphisms (D being regarded as a right vector space over itself) whose elements are called *functionals*. This set becomes a *left* vector space over D where addition and scalar multiplication are defined by

$$(f_1 + f_2)(v) = f_1(v) + f_2(v),$$
$$(\alpha f)(v) = \alpha f(v).$$

The reader should note that the definition insures $(\alpha\beta)f = \alpha(\beta f)$, as required for a left vector space. Every 1-space of functionals possesses a kernel which is a hyperplane of V, and the collection of all 1-spaces of functionals whose kernels contain a prescribed codimension-2 space, correspond to a 2-subspace of V^*. (If you think about it, this is just the Veldkamp lines condition.) So we have realized that

$$\mathbf{P}(V^*) = \mathbf{P}(V)^*. \tag{7.8}$$

Now a word about the division ring D. There is another ring D^{opp}, called the *opposite ring* defined as follows: $(D^{\mathrm{opp}}, +) = (D, +)$ as additive groups.[12] Multiplication in $(D^{\mathrm{opp}}, \circ)$ is simply defined by the rule that $\alpha \circ \beta = \beta\alpha$ (where juxtaposition on the right indicates multiplication in D), for all $\alpha, \beta \in D$. In other words, addition is the same, but the multiplication table for D^{opp} is the "transpose" of the multiplication table for D. It should be clear that D^{opp} is a division ring.

The reason we bring this up is that any left D-vector space W can be instantly converted into a right D^{opp}-space by the simple device of writing

$$\alpha v = v\alpha. \tag{7.9}$$

(Although the symbols do not suggest it, (7.9) is a non-trivial equation. The "α" on the right side is a member of the new ring D^{opp}, while that on the left is in the old division ring D. The symbols just don't explicitly reveal the worlds in which they operate.)

Now if V is a right vector space over D, then, as we have observed above, V^* is a left vector space over D. So it can be converted to a right vector space over D^{opp}. Note from Eq. (7.9) that any left D-subspace of V^* is also a right D^{opp}-subspace

[12] This is equality, not a mere isomorphism.

of V^*. That means that the projective spaces they define are the same, unambiguously denoted $\mathbf{P}(V^*) = \mathbf{P}(V)^*$.

But this gives us a slight advantage in considering the hypothesis below:

(SQ): There is a morphism of projective spaces

$$f : \mathbf{P}(V) \to \mathbf{P}(V)^*.$$

The right side is the projective space of a *right* vector space over D^{opp} and so if f has index at least three, the fundamental theorem of projective geometry, as enshrined in Corollary 6.8.4, can be invoked to show that:

> *There exists an isomorphism $\sigma : D \to D^{opp}$ and a σ-semilinear transformation $t : V \to V^*$ inducing the morphism f.*

What is an isomorphism $\sigma : D \to D^{\text{opp}}$? It takes $\alpha\beta$ to $\sigma(\alpha) \circ \sigma(\beta) = \sigma(\beta)\sigma(\alpha)$ – that is, it is an *anti-automorphism of D*.

Of course any αt proportional to transformation t also induces f, and so there is an entire proportionality class associated with f.

Now suppose f is induced by the σ-semilinear transformation t. For each $u \in V$, $t(u)$ is a functional. So t defines a mapping

$$\begin{aligned} sq(t) : V \times V &\to D \text{ where} \\ sq(t)(u, v) &:= t(u)(v), \end{aligned}$$

that is, the functional $t(u)$ applied to vector u.

Then the new function $h := sq(t)$ possesses these properties:

$$h(u_1 + u_2, v) = h(u_1, v) + h(u_2, v), \tag{7.10}$$

$$h(u, v_1 + v_2) = h(u, v_1) + h(u, v_2), \tag{7.11}$$

$$h(u, v\alpha) = h(u, v)\alpha, \text{ and} \tag{7.12}$$

$$h(u\alpha, v) = \alpha^\sigma h(u, v). \tag{7.13}$$

Any function $h : V \times V \to D$ satisfying the Eq. (7.10)–(7.13) is called a *σ-sesquilinear form*. The scalar multiple αh (where t is replaced by αt) is a σ'-sesquilinear form where $\sigma' = i \circ \sigma$ and where i is the appropriate inner automorphism of D. Such a multiple αh is said to be *proportional* to the form h.

The subspace

$$V^L := \{u \in V | h(u, v) = 0 \text{ for all } v \text{ in V}\}$$

is called the *left radical of h*. Clearly the kernel of t is the left radical of $h = sq(t)$.

Let's reverse all this. Now suppose h is a σ-sesquilinear form. Then for each vector u in V, the function $h_u : V \to D$ which sends arbitrary vector v to

$h(u, v)$ is clearly a functional of V. In fact, $t_h : V \to V^*$ which sends u to the functional h_u is a σ-semilinear transformation $V \to V^*$ inducing a morphism $f(h) : \mathbf{P}(V) \to \mathbf{P}(V)^*$. It is easy to check that if h' is proportional to h then $f(h) = f(h')$. Conversely, if $\dim V \geq 3$ then $f(h) = f(h')$ implies that h is proportional to h'.

So we have the following.

Theorem 7.6.1 *Let $\mathbf{P}(V)$ be a classical projective space. Then there is a bijection between the projective space morphisms*

$$\mathbf{P}(V) \to \mathbf{P}(V)^*$$

whose kernels have codimension at least three and the proportionality classes of sesquilinear forms whose left radical has codimension at least three in V.

7.6.2 Quasi-Polarities

Let V be a right vector space over a division ring D. Let $S(V)$ be the poset of all subspaces of V.[13] As usual, $P(V)$ is the projective geometry of 1- and 2-subspaces of V, the (projective) points $\mathcal{P}(V)$ and lines $\mathcal{L}(V)$. A *hyperplane of V* is a subspace of codimension 1 in V (its points forming a geometric hyperplane of $P(V)$.)

A function

$$\pi : S(V) \to S(V)$$

is called a *quasi-polarity*[14] if and only if:

1. For each point p (1-subspace of V), $\pi(p)$ is either a hyperplane of V or V itself.
2. If U_1 and U_2 are subspaces of V, then $U_1 \leq \pi(U_2)$ if and only if $U_2 \leq \pi(U_1)$.
3. π is inclusion-reversing.

A quasi-polarity π always defines a relation $\perp = \perp_\pi$ on the set $\mathcal{P}(V)$ of projective points of $P(V)$, by writing $p \perp q$ if and only if $p \in \pi(q)$. The second condition for a quasi-polarity shows that this is a symmetric relation.

On the other hand, we could have begun with a symmetric relation $\perp$ on the projective points $\mathcal{P}(V)$ of $P(V)$ with the property that for each point p the set

---

[13] When V has finite dimension, $S(V)$ differs from the geometry $PG(V)$ (besides being a poset rather than a geometry) by including the 0-subspace and the entire subspace, V. In infinite dimension, $PG(V)$ is confined to the geometry of non-zero finite-dimensional subspaces since otherwise one would confront subspaces of the same dimension one properly contained in the other – making a rank function impossible to define. However for $S(V)$, all subspaces are included, and still form a poset under the inclusion relation, regardless of dimension.

[14] This term is due to Tits [138]: see also the excellent survey article of A. Cohen in *The Handbook of Incidence Geometry*, (F. Buekenhout ed.) [35].

$x^\perp := \{q \in \mathcal{P}(V) | p \perp q\}$ is either a hyperplane of $P(V)$ or the set $\mathcal{P}(V)$ of all projective points. We may also regard this symbol $p^\perp$ as indicating a subspace of V – either a hyperplane of V or V itself. For each vector subspace U of V let $\mathcal{P}(U)$ be the set of 1-spaces of V contained in U. Then we set

$$\pi(U) = \cap_{u \in \mathcal{P}(U)} u^\perp,$$

an intersection of subspaces of V. By convention $\pi(0) = V$. In this way we obtain a mapping $S(V) \to S(V)$ satisfying the three conditions of a quasi-polarity.

There is a third way to regard a quasi-polarity.

Lemma 7.6.2 *Let $P(V)$ be a classical projective space.*

1. *Let $\pi : S(V) \to S(V)$ be a quasi-polarity. Then, truncating both range and domain to spaces of small dimension and codimension, π induces a morphism of projective spaces, $\mu : \mathbf{P}(V) \to \mathbf{P}(V)^*$, with the following "reflexive property": that for any two points, p and q of $\mathbf{P}(V)$,*

$$p \in \mu(q) \text{ if and only if } q \in \mu(p).$$

2. *Suppose, conversely, that $\mu : \mathbf{P}(V) \to \mathbf{P}(V)^*$ is a morphism of projective spaces subject to the "reflexive property" just listed. Then μ extends uniquely to a quasi-polarity from which it was induced.*

Proof Suppose L is a line of $\mathbf{P}(V)$ with distinct points a and b. If $x \in \pi(a) \cap \pi(b)$, then the line L is contained in $\pi(x)$, so $x \in \pi(c)$ for all points c of L. This means that π produces a morphism μ from the projective space of points and lines of $\mathbf{P}(V)$ into its Veldkamp space $\mathbf{P}(V)^*$ whose kernel is the set $\mathrm{Rad}(\pi)$ of all points r for which $\pi(r) = V$. The "reflexive property" follows from the symmetry property of a quasi-polarity.

Conversely, if $\mu : \mathbf{P}(V) \to \mathbf{P}(V)^*$ is a morphism with kernel R, then for any subset X of $\mathbf{P}(V)$ one can set $\pi(X) = \cap_{x \in X} \mu(X)$, and so extend this to obtain an inclusion-reversing mapping $S(V) \to S(V)$. By the hypothesized "reflexive property," this mapping enjoys the symmetric property of a quasi-polarity π which restricts to the morphism μ. The proof is complete. □

Suppose π is a quasi-polarity. A subspace U of V is said to be *absolute with respect to* π if and only if $U \leq \pi(U)$. Note that if $W \leq U \leq \pi(U)$ then $\pi(U) \leq \pi(W)$ so W is also absolute. Thus the collection $\mathcal{A}(\pi)$ of all subspaces of V which are absolute with respect to π form a subposet of $S(V)$ which is closed under taking subspaces.[15] If we "truncate" this poset to absolute 1-spaces and 2-subspaces of V (that is, absolute projective points and lines) we obtain a point-line geometry $S_\pi = (\mathcal{P}_\pi, \mathcal{L}_\pi)$ which is already embedded (in the sense of Chap. 4) in the subspace

[15] In the language of posets, $\mathcal{A}(\pi)$ is an *order ideal*.

of the projective space $P(V)$ which it spans. Note that two points p and q of $\mathcal{P}_\pi$ are collinear in $\mathcal{S}_\pi$ if and only if $p \perp q$.

Now let L be an absolute line with respect to π, and let p be any absolute point. Then $\pi(p)$ is either a hyperplane of $P(V)$ or is $\mathcal{P}(V)$. In either case, it contains one or all of the points of L. Thus p is collinear to one or all of the points of L, so the axiom (P) holds. Thus

$\mathcal{S}_\pi$ is a polar space.

We call $\mathcal{S}_\pi$ the *(classical) polar space defined by the quasi-polarity* π.

Now one will note that if π is a quasi-polarity of V, the collection of subspaces U with $\pi(U) = V$ is closed under taking joins and so this collection possesses a unique maximal subspace R of V which we call the *radical of the quasi-polarity of* π and we denote this space $R = \mathrm{Rad}(\pi)$. The quasi-polarity is said to be *non-degenerate* if and only if Rad (π) is the 0-subspace of V. A non-degenerate quasi-polarity is called a *polarity*. In this case π is injective on the finite-dimensional subspaces.

Suppose R is the radical of the quasi-polarity π of V. Then every subspace U for which U/R is one-dimensional, $\pi(U)$ is always a hyperplane of V (it is never V itself). Thus we can define an induced quasi-polarity $\bar{\pi} : S(V/R) \to S(V/R)$ where $\bar{\pi}(U/R) = \pi(U)/R$ for every vector subspace U/R of V/R. It it easy to see that $\bar{\pi}$ is a polarity of $S(V/R)$.

Of course, for every subspace U of V, a quasi-polarity $\pi : \mathcal{S}(V) \to \mathcal{S}(V)$ induces a quasi-polarity $\pi_U : \mathcal{S}(U) \to \mathcal{S}(U)$, where, for each point p of $\mathcal{P}(U)$, one sets $\pi_U(p) = \pi(p) \cap U$.

In general, as noted, a quasi-polarity π of V may be such that the absolute points do not span V. If they do, that is, $\langle \mathcal{P}_\pi \rangle_V = \langle U \in \mathcal{A}_\pi \rangle_V = V$, we say π is a *trace-valued quasi-polarity*. In this case we have an actual embedding $\mathcal{S}_\pi \to P(V)$ of geometries. Of course, when π is not trace valued, we can always form the quasi-polarity π_W on $W := \langle \mathcal{P}_\pi \rangle_V$, as in the previous paragraph. Then π_W will be trace-valued.

Remark Note that as we have defined things, the identity mapping $0 \to 0$ is actually a non-degenerate quasi-polarity of the 0 vector space. Another trivial example is the constant mapping $\pi : \mathcal{S}(V) \to \mathcal{S}(V)$ with constant value V. Here $S_\pi = P(V)$.

In the next section we look at the algebraic side of the theory of quasi-polarities. This will give us honest-to-goodness examples of quasipolarities (unlike the insipid specimens of the previous paragraph) and the classical polar spaces they define. Also, this change in perspective will allow us to introduce still further classical polar spaces which are not defined by a quasi-polarity.

7.6.3 The Relation Between Quasi-Polarities and Reflexive Forms

Suppose $\pi : S(V) \to S(V)$ is a quasi-polarity. As detailed in the previous section, we may also identify π with a unique projective space morphism $\mathbf{P}(V) \to \mathbf{P}(V)^*$

with the "reflexive property." But from the section before the last, we showed that such morphisms of index at least 3 corresponded to proportionality classes of sesquilinear forms. But now, not any such class will do. It must be one which provides the "reflexive property" for its associated morphisms $\mathbf{P}(V) \to \mathbf{P}(V)^*$.

A σ-sesquilinear form f is said to be *reflexive* if and only if, for any vectors u and v of V, $f(u, v) = 0$ if and only if $f(v, u) = 0$. Fix a vector $u \in V$. The mapping $V \to D$ which sends vector v to $f(u, v)$ is a functional of the vector space V and its kernel $u^{\perp} := \{v \in V | f(u, v) = 0\}$ must be a hyperplane of V or V itself. If f is reflexive, and for each subspace U of V we set

$$\pi_f(U) := \{v \in V | f(U, v) = 0\},$$

then $\pi_f : \mathcal{S}(V) \to \mathcal{S}(V)$ is a quasi-polarity of V.

A reflexive σ-sesquilinear form f is said to *represent a quasi-polarity* π if and only if $\pi = \pi_f$.

A subspace U is *totally isotropic with respect to the reflexive σ-sesquilinear form* if and only if $f(U, U) = \{0\}$, or equivalently, $U \subseteq \pi_f(U)$ – that is, if and only if U is absolute with respect to π_f. In this case the radical, $\text{Rad}(\pi_f) = \pi_f(V)$ is called the *radical of the form* f and is denoted $\text{Rad}((V, f))$ just to remind one how this quasi-polarity was formed. If this subspace is 0, we say that the *form f is non-degenerate* (just as we did for its derived quasi-polarity in the previous section).

Theorem 7.6.3 *1. Suppose π is a quasi-polarity of a right vector space V over division ring D for which $V/Rad(\pi)$ has dimension at least three. Then $\pi = \pi_f$ for some reflexive σ-sesquilinear form f.*

2. Conversely, if f is a reflexive σ-sesquilinear form, then f defines a unique quasi-polarity π_f whose radical coincides with the radical of f.

Proof The second statement has been thoroughly discussed.

The first statement actually uses some previous theorems. If π is a quasi-polarity, it can be regarded as a projective space morphism

$$\pi : \mathbf{P}(V) \to \mathbf{P}(V)^*$$

and so, by Theorem 7.6.1 and the hypothesis on the index, corresponds to a proportionality class of sesquilinear forms. Moreover, these forms are forced to be reflexive forms, since they represent a π having the "reflexive property." □

Remark The condition on the dimension of $V/Rad(\pi)$ in part 1 is absolutely necessary. Suppose, $\mathbf{P}_1$ is the full collection of 1-spaces of a two-dimensional vector space V over a division ring D. Let t be any involution in the (possibly infinite) symmetric group $\text{Sym}(\mathbf{P}_1)$. We may then define a polarity $\pi(t) : \mathcal{S}(V) \to \mathcal{S}(V)$, as follows:

1. Set $\pi(t)(0) = V$, and set $\pi(t)(V) = 0$.
2. For each 1-subspace A, set $\pi(t)(A) := A^t$.

Then one can easily see that since t is an involution, $U \subseteq \pi(t)(W)$ if and only if $W \subseteq \pi(t)(U)$, for any two subspaces U and W of the 2-space V. Thus $\pi(t)$ is indeed a polarity.

Now one can see that if D is the field $\mathbf{Q}$ of rational numbers, then σ is the identity, $|D|$ is countably infinite, and, considering the limited data one needs to define a two-dimensional form, there are only countably many proportionality classes of reflexive σ-sesquilinear forms. On the other hand there are *uncountably* many involutions in $\mathrm{Sym}(\mathbf{P}_1)$, each giving a unique polarity. So we can choose an involution t, so that $\pi(t)$ is not represented by any sesquilinear form.

Reflexive Sesquilinear Forms are (σ, ϵ)-Hermitian Forms

When the radical does not dominate things, reflexive sesquilinear forms are reflexive because of a particular formula connecting $f(u, v)$ and $f(v, u)$.

Theorem 7.6.4 *Let f be a reflexive, σ-sesquilinear form on V, and assume that $\mathrm{codim}_V(\mathrm{Rad}(f)) \geq 2$. Then there exists ϵ in D such that*

$$f(y, x) = f(x, y)^\sigma \epsilon \text{ for all } x, y \in V. \tag{7.14}$$

Moreover, we then have

$$\epsilon^\sigma = \epsilon^{-1}, \text{ and} \tag{7.15}$$

$$a^{\sigma^2} = \epsilon a \epsilon^{-1} \text{ for all } a \in D. \tag{7.16}$$

Proof Let $x \in V - \mathrm{Rad}(f)$, choose $y \in V$ with $f(x, y) = 1$, and set $\epsilon = f(y, x)$. For any $z \in V$ we can write $z = ya + w$ with $w \in x^\perp$ and $a \in D$. Then $f(x, z) = f(x, ya) = a$, while $f(z, x) = f(ya, x) = a^\sigma f(y, x) = f(x, z)^\sigma \epsilon$. That is,

$$f(z, x) = f(x, z)^\sigma \epsilon \text{ for all } z \in V. \tag{7.17}$$

We may write $\epsilon = \epsilon_x$ to emphasize that ϵ depends only on x. Observe that for any $t \in x + y^\perp$ we have $f(t, y) = 1$ and $f(y, t) = \epsilon$. This yields

$$\epsilon_x = \epsilon_t \text{ if } t \in x + y^\perp. \tag{7.18}$$

But if we replace y by an element y' in $y + x^\perp$, the new y' enjoys the same properties:

$$f(x, y') = 1 \text{ and } f(y', x) = \epsilon,$$

which, for y implied (7.18). Thus we may strengthen that equation to read:

(*) For each $y' \in y + x^\perp$ and each $t \in x + (y')^\perp$, one has $\epsilon_x = \epsilon_t$.

Now suppose by way of contradiction that there is a vector u in $V - \mathrm{Rad}(f)$ such that $\epsilon_u \neq \epsilon_x$. Choose $v \in V$ so that $f(u, v) = 1$. If one could find a vector w, say, in $(x + (y')^{\perp}) \cap (u + v^{\perp})$, then by (*) and by Eq. (7.18) (with (u, v) in the role of (x, y)), one could conclude $\epsilon_x = \epsilon_w = \epsilon_u$, a contradiction. Thus, for all $y' \in y + x^{\perp}$, we must have $(x + (y')^{\perp}) \cap (u + v^{\perp})$ empty, and as this is the intersection of two cosets of hyperplanes, this can only happen if they are cosets of the same hyperplane. This means

$$v^{\perp} = (y')^{\perp} \text{ for all } y' \in y + x^{\perp}. \tag{7.19}$$

But Eq. (7.19) declares that any element of $v^{\perp}$ is perpendicular to all elements of $y + x^{\perp}$ and hence to the entire subspace they generate. The latter subspace contains all *differences* of elements of $y + x^{\perp}$ (which form the hyperplane $x^{\perp}$ itself) as well as $y \notin x^{\perp}$, and so is all of V. Thus $v^{\perp} \leq V^{\perp} = \mathrm{Rad}(f)$ contrary to the assumption that $\mathrm{codim}_V \mathrm{Rad}(f) \geq 2$.

Thus Eq. (7.14) of the statement of the theorem holds.

Now $1 = f(x, y) = f(y, x)^{\sigma}\epsilon = \epsilon^{\sigma}\epsilon$, so $\epsilon^{\sigma} = \epsilon^{-1}$. Also for any $a \in D$, $a = f(x, ya)$, $a^{\sigma}\epsilon = f(ya, x)$, and $(a^{\sigma}\epsilon)^{\sigma}\epsilon = f(x, ya) = a$. So

$$\epsilon^{\sigma} a^{\sigma^2} \epsilon = \epsilon^{-1} a^{\sigma^2} \epsilon = a \text{ and } a^{\sigma^2} = \epsilon a \epsilon^{-1}.$$

The proof is complete. □

In order to describe the next theorem we need to define a certain derivation of one graph from another one. Suppose $\Gamma = (X, E)$ is a simple graph. The graph $\Delta(\Gamma)$ has the same vertex set X as Γ, but in this graph two distinct vertices are adjacent if and only if they are not adjacent in G and there is a third vertex not adjacent to either of them in Γ – that is, the adjacent vertices of Δ are distinct and connected by a path of length two in the complement of Γ.

Theorem 7.6.5 (Characterizing Hermitian forms.) *Let V be a right vector space over the division ring D. Suppose S is a spanning set of vectors of the vector space V which is closed under scalar multiplication. Let f be a σ-sesquilinear form on V and suppose S "supports reflexivity of f" in the sense that for all vectors $s \in S$,*

$$f(s, x) = 0 \text{ if and only if } f(x, s) = 0,$$

for all vectors x in V. In partcular, there is a graph $\Gamma = (\bar{S}, \perp)$ whose vertices are the 1-spaces of V generated by the non-zero vectors of S, two such 1-spaces a and b being adjacent if and only they are distinct and are "perpendicular": that is $f(a, b) = 0 = f(b, a)$.

If the graph $\Delta(\Gamma)$ is connected, then f is a $\{\sigma, \epsilon\}$-Hermitian form.

Proof Choose $\bar{s} = \langle s \rangle_V$ in $\bar{S}$, the vertex set of both Γ and $\Delta := \Delta(\Gamma)$. If $\bar{s}$ is the only vertex of $\bar{S}$, then V is one-dimensional, and f is reflexive. So suppose $\bar{t} = \langle t \rangle$ is a vertex of Δ adjacent to $\bar{s}$. Then there exists a non-zero vector $y \in S$ so that $\bar{y} \notin \pi(s) \cup \pi(t)$.

Now we simply mimic the very beginning of the proof of the previous theorem. Clearly, there exists a vector y in $\bar{y}$, such that $f(x, y) = 1$. Then set $\epsilon_s := f(y, s)$. Then we conclude as before that

$$f(s, z) = f(z, x)^\sigma \epsilon_s$$

for all $z \in V$. Then we can choose $t \in \bar{t}$ so that $f(t, y) = 1$, and deduce as above that

$$f(t, y) = f(y, t)^\sigma \epsilon_t, \text{ and that } \epsilon_s = \epsilon_t.$$

Since the graph Δ is connected, there is a uniform ϵ such that

$$f(s, x) = f(x, s)^\sigma \epsilon \text{ for all vectors } x \in V.$$

Now consider two arbitrary vectors x and y of V. Since S spans V we may write

$$x = \sum\nolimits_{i \in I} \alpha_i s_i,$$
$$y = \sum\nolimits_{j \in J} \beta_j s_j.$$

It is then a computational exercise to verify

$$f(x, y) = f(y, x)^\sigma \epsilon.$$

(Note that this computation would not be possible without a uniform value of ϵ.)

Corollary 7.6.6 *Assume all of the hypotheses of the preceding theorem, except that the statement that "S supports reflexivity" is replaced by the weaker assertion that "S supports reflexivity relative to a spanning set of vectors X" – that is, for all vectors* $(s, x) \in S \times X$,

$$f(s, x) = 0 \textit{ if and only if } f(x, s) = 0.$$

Then f *is a* $\{\sigma, \epsilon\}$*-Hermitian form.*

Proof From the connectivity of the graph Δ one can infer the existence of a uniform constant ϵ such that

$$f(x, s) = f(s, x)^\sigma \epsilon,$$

for all vectors $x \in X$ and all vectors $s \in S$. We need only show that X can be replaced by all of V, and apply the preceding theorem. Now an arbitrary vector y in V has the form $y = \sum x_i \alpha_i$, where the x_i are in X. Then for any vector s in S,

$$\begin{aligned} f(y,s) &= f\left(\sum x_i\alpha_i, s\right) \\ &= \sum \alpha_i^{\sigma} f(x_i, s) \\ &= \sum \alpha_i^{\sigma} f(s, x_i)^{\sigma}\epsilon \\ &= \sum (f(s, x_i)\alpha_i)^{\sigma}\epsilon \\ &= f(s, y)^{\sigma}\epsilon. \end{aligned}$$

Thus S supports reflexivity of f with respect to all vectors of V and we now have the conditions of the preceding theorem. □

Remark This theorem associates with each reflexive sesquilinear form, a σ-Hermitian form (take $S = \mathbf{P}(V)$ in the theorem). But it does a lot more. Suppose you are missing an argument that the reflexivity condition holds for f for large sets of pairs of vectors of V. One can still infer that if things look good for an embedded spanning set, the form is still genuinely reflexive.

Note that this part of the theorem – particularly in the form of its general Corollary 7.6.6 – is *absolutely critical* for the classification of embeddable polar spaces recounted later in this chapter. Specifically it is required in showing that an unknown non-degenerate polar space of rank at least three, which has been embedded as a spanning set in some projective space, is really embedded in an ambient classical polar space defined by a σ-Hermitian form.

7.6.4 Admissible Pairs and Gramm Matrices

A pair (σ, ϵ) consisting of an antiautomorphism σ of D and a scalar ϵ of D bearing the relations of Eqs. (7.15) and (7.16), is called an *admissible pair for* D. A σ-sesquilinear form satisfying the Eq. (7.14) is called a (σ, ϵ)-*Hermitian form*.

A non-zero scalar multiple $f' = \lambda f$, $\lambda \neq 0$, of a σ-sesquilinear form f is easily seen to be a σ'-sesquilinear form where, for any scalar α in D, $\alpha^{\sigma'} = \lambda\alpha^{\sigma}\lambda^{-1}$ so σ' is still an anti-automorphism of D. The two forms, f and λf are said to be *proportional*. "Being proportional" is an equivalence relation on sesquilinear forms whose equivalence classes are called *proportionality classes*.

If one of f or λf ($\lambda \neq 0$) is reflexive, so is the other, and in that case both f and λf induce the same quasi-polarity π and so define the same classical polar space S_π.

However, the admissible pair for λf is now (σ', ϵ') where, as above, σ' is the antiautomorphism obtained by first applying σ and then the inner automorphism $\alpha \to \lambda\alpha\lambda^{-1}$ and $\epsilon' = \lambda\lambda^{-\sigma}\epsilon$.

At this stage, provided the corresponding radicals have codimension at least two, we have a pretty good correspondence between these three sets:

- Quasi-polarities
- Proportionality classes of reflexive σ-sesquilinear forms
- Proportionality classes of (σ, ϵ)-Hermitian forms

Many examples exist. Suppose $B = \{b_i\}_{i \in I}$ is a basis for vector space V indexed by I. Then every vector x is uniquely expressible as a finite sum $\sum_{i \in Supp(x)} x_i b_i$ where x_i is a non-zero coefficient. The subset $Supp(x)$ of I is called the *support of vector* x. Now suppose $F = (F(i, j))$ is an $I \times I$ matrix (that is, a function $F : I \times I \to D$) for which $F(i, j) = F(j, i)^\sigma \epsilon$. We can then use this matrix F to define a form $f : V \times V \to V$ by setting

$$f(x, y) = \sum\nolimits_{i,j \in Supp(x) \cap Supp(y)} x_i^\sigma F(ij) y_j. \tag{7.20}$$

Then f is a σ-sesquilinear form on V.

Conversely, if f is already a σ-sesquilinear form on V, then the values of $f(x, y)$ are given by the formula presented in Eq. (7.20), where $F(i, j) := f(b_i, b_j)$, the so-called *Gramm matrix for* f.[16]

Special reflexive Forms

A $(\sigma, 1)$-Hermitian form is called σ*-Hermitian* or simply *Hermitian*. Note that in this case $\sigma^2 = \mathrm{id}$, the identity automorphism.

A special case plays an exceptional role in some of the theorems. We say that a (σ, ϵ)-Hermitian form f is an *alternating* or a *symplectic form* if and only if every vector is isotropic – that is, $f(x, x) = 0$ for all x. Here, $\sigma = \mathrm{id}$, so D is a field, and $\epsilon = -1$. In characteristic distinct from 2, $(\sigma, \epsilon) = (\mathrm{id}, -1)$ also implies f is alternating, but this implication fails in characteristic 2.

Another special case is when $\sigma = \mathrm{id}$ and $\epsilon = 1$. In this case, $f(x, y) = f(y, x)$ and so the form is called *symmetric*. Now we note that in characteristic two, alternating forms are symmetric, but some symmetric forms are not alternating.

If one is willing to replace a form by a scalar multiple (a change which does not affect the associated quasi-polarity), a general (σ, ϵ)-Hermitian form can be made nicer. First, we require part 1 of the following.

Lemma 7.6.7 (A property of admissible pairs.) *Suppose* (σ, ϵ) *is an admissible pair associated with a non-zero* (σ, ϵ)*-form* $f : V \times V \to D$*:*

1. *If* f *is not alternating, there exists a scalar* $\lambda \in D$ *such that* $\lambda^\sigma \lambda^{-1} = \epsilon$.
2. *If* f *is not an alternating form, there exists a scalar* μ *such that* $\mu^\sigma \mu^{-1} = -\epsilon$.

Proof 1. If $\epsilon = 1$, let $\lambda = 1$. So suppose $\epsilon \neq 1$. Now if f is not symplectic, there is a vector x such that $\lambda^{-1} = f(x, x)$ is not zero. Then $\lambda^{-\sigma} \epsilon = \lambda^{-1}$ so $\epsilon = \lambda^\sigma \lambda^{-1}$, as desired in part 1.

[16] Though it is hardly of interest, note that unlike the usual matrix theory which insists on totally ordered bases and totally ordered index sets, the formulae above are free of that, and hence free of Zorn's lemma in that regard. It is important that the sums involved in Eq. (7.20) are all finite, due to the finiteness of the supports of the vectors.

2. Now let $Z := Z(D)$, the center of D. Then form the field extension $K := Z(\epsilon)$ and note that σ acts on K since it must stabilize Z and inverts ϵ.

If σ fixes ϵ we see that $\epsilon^{-1} = \epsilon$ so ϵ is an involution. Then $(1+\epsilon)(1-\epsilon) = 0$ forces $\epsilon = \pm 1$. If D has characteristic 2 we take $\mu = 1$. So we may assume D has characteristic distinct from 2 and that $\epsilon = -1 \neq 1$. Then $\epsilon \in Z(D)$ and so σ has order 2. Then σ acts on D as a vector space over (D) with two eigenspaces V_{-1} and V_{+1} whose direct sum is D. If $\sigma = 1$, f is alternating against our hypothesis. So we may assume there is a non-zero scalar ζ such that $\zeta^\sigma = -\zeta$. Now, since $\epsilon = -1$ one can see that $\mu = \zeta^2$ does the job.

So we may assume σ acts non-trivially on K and on ϵ. Then (as the characteristic can be assumed different than 2) K (as a vector space over (D)) decomposes into two eigenspaces $K_{-1} \oplus K_{+1}$, and we now have an element $\zeta \in K$ commuting with ϵ with $\zeta^\sigma \zeta^{-1} = -1$. Now if λ is chosen as in part 1 so that $\lambda^\sigma \lambda^{-1} = \epsilon$, then, because ζ was carefully chosen to commute with ϵ, $\mu = \lambda\zeta$, satisfies the requirements of λ in the statement of part 2. □

Now we have the following.

Corollary 7.6.8 *Any (σ, ϵ)-Hermitan form that is not an alternating form is proportional to a σ-Hermitian form.*

Proof By part 1 of Lemma 7.6.7, if f is not alternating, we can find λ so that $\epsilon' = \lambda\lambda^{-\sigma}\epsilon = 1$. Then λf is σ-Hermitian. □

Special Subgroups of $(D, +)$ Defined by an Admissible Pair

Certain vital properties of reflexive forms, and their relations to pseudoquadratic forms, come down to basic algebraic properties of the admissible pair and the division ring D itself. We isolate the pertinant subspaces of $(D, +)$ at this point.

Let (σ, ϵ) be an admissible pair for the division ring D. Let $\mathrm{tr}_{(\sigma,\epsilon)} : D \to D$ be the mapping which takes each scalar α to $\alpha + \alpha^\sigma \epsilon$. We are also interested in another closely related mapping, $d_{(\sigma,\epsilon)} : D \to D$ defined by $d_{(\sigma,\epsilon)} = \mathrm{tr}_{(\sigma,-\epsilon)}$. Both of these mappings are clearly endomorphisms of the additive group $(D, +)$. Let $T_{\sigma,\epsilon}$ and $D_{\sigma,\epsilon}$ be the images of $\mathrm{tr}_{(\sigma,\epsilon)}$ and $d_{(\sigma,\epsilon)}$, respectively. Finally, we set $D^{\sigma,\epsilon} := \ker \mathrm{tr}_{(\sigma,\epsilon)} = \{\alpha \in D | \alpha^\sigma \epsilon = -\alpha\}$ and $T^{\sigma,\epsilon} := \ker d_{(\sigma,\epsilon)} = \{\alpha \in D | \alpha^\sigma \epsilon = \alpha\}$.

Lemma 7.6.9 *The following statements hold:*

1. *If $\alpha \in T_{\sigma,\epsilon} = \mathrm{tr}_{(\sigma,\epsilon)}(D)$, then $\alpha^\sigma = \alpha\epsilon^{-1}$. So $T_{\sigma,\epsilon} \subseteq T^{\sigma,\epsilon}$. Similarly, if $\alpha \in D_{\sigma,\epsilon} := d_{(\sigma,\epsilon)}(D)$, then $\alpha^\sigma = -\alpha\epsilon^{-1}$. So $D_{\sigma,\epsilon} \subseteq D^{\sigma,\epsilon}$.*
2. *$(\mathrm{tr}_{(\sigma,\epsilon)})^2 = 2\mathrm{tr}_{(\sigma,\epsilon)}$.*
3. *$(d_{(\sigma,\epsilon)})^2 = 2D_{(\sigma,\epsilon)}$.*
4. *The four subgroups $T_{\sigma,\epsilon}$, $T^{\sigma,\epsilon}$, $D_{\sigma,\epsilon}$, and $D^{\sigma,\epsilon}$ are D^*-modules with respect to the action $\alpha \to \alpha * \lambda := \lambda^\sigma \alpha \lambda$. That is, if A is one of these four groups, then $A = \lambda^\sigma A \lambda$ for each non-zero $\lambda \in D$.*

5. *The quotient group $D^{\sigma,\epsilon}/D_{\sigma,\epsilon}$ is a right D-module (that is, it is a right D-vector space with respect to the multiplication induced by "$*$").*
6. *If the characteristic of D is not 2, then $D_{\sigma,\epsilon} = D^{\sigma,\epsilon}$ and $T_{\sigma,\epsilon} = T^{\sigma,\epsilon}$.*

Proof The first four parts are direct consequences of the definitions of the subgroups and admissibility relations and are left as exercises which should familiarize the student with relations involving antiautomorphisms of non-commutative division rings.

For part 5, one must show that for any two scalars λ_1 and λ_2 of D, that

$$(\alpha + D_{\sigma,\epsilon}) * (\lambda_1 + \lambda_2) = \alpha * \lambda_1 + \alpha * \lambda_2 + D_{\sigma,\epsilon}.$$

This reduces to showing that if $\alpha \in D^{\sigma,\epsilon}$, then $\lambda_1^\sigma \alpha \lambda_2 + \lambda_2^\sigma \alpha \lambda_1$ is always an element of $D_{\sigma,\epsilon}$. But

$$\begin{aligned}
\lambda_1^\sigma \alpha \lambda_2 + \lambda_2^\sigma \alpha \lambda_1 &= -\lambda_1^\sigma(-\alpha\epsilon^{-1}) \cdot \epsilon \lambda_2^\sigma \epsilon^{-1} \cdot \epsilon + \lambda_2^\sigma \alpha \lambda_1 \\
&= -(\lambda_1^\sigma \alpha^\sigma \lambda_2^{\sigma^2})\epsilon + \lambda_2^\sigma \alpha \lambda_1 \\
&= -(\lambda_2^\sigma \alpha \lambda_1)^\sigma \cdot \epsilon + \lambda_2^\sigma \alpha \lambda_1 \\
&\in D_{\sigma,\epsilon}.
\end{aligned}$$

For part 5, we note that if $\alpha \in D^{\sigma,\epsilon}$, then $\alpha^\sigma \epsilon = -\alpha$ and one can then write $\alpha = (\alpha - \alpha^\sigma \epsilon)/2$. The proof for $T_{\sigma,\epsilon}$ is similar: $-\epsilon$ is replaced by ϵ. □

Trace-Valued Forms

At first sight, this does not appear to be an important subject, since upon resriction to an appropriate subspace any (σ, ϵ)-Hermitian form becomes a so-called "trace-valued" form. But they do assume an importance since a number of important algebraic properties are equivalent to the geometric property. We begin with one of the algebraic properties.

Suppose g is a not-necessarily reflexive σ-sesquilinear form. Can we make a reflexive σ-sesquilinear form out of it? Yes we can: for the form $f := g + g^{T\sigma}\epsilon$ (that is, the form with values $f(x, y) = g(x, y) + g(y, x)^\sigma \epsilon$) satisfies

$$\begin{aligned}
f(y, x)^\sigma &= g(y, x)^\sigma + \epsilon^\sigma g(x, y)^{\sigma^2} \\
&= g(y, x)^\sigma + \epsilon^{-1} \epsilon g(x, y) \epsilon^{-1} \\
&= \left(g(y, x)^\sigma \epsilon + g(y, x)\right) \epsilon^{-1} \\
&= (f(x, y))\epsilon^{-1},
\end{aligned}$$

so, $f(y, x)^\sigma \epsilon = f(x, y)$.

Perhaps not all reflexive σ-sesquilinear forms are obtained by this device. So, to distinguish these worlds, we say that a reflexive σ-sesquilinear form f is *trace-valued* if and only if $f(x, y) = g(x, y) + g(y, x)^\sigma \epsilon$ for some σ-sesquilinear form g.

Theorem 7.6.10 (The trace-valued condition.) *The following conditions are equivalent for a (σ, ϵ)-Hermitian form f on vector space V:*

1. *f is trace-valued, that is, $f(x, y) = g(x, y) + g(y, x)^\sigma \epsilon$ for some σ-sesquilinear form g and for all vectors x, y of V.*
2. *For all $x \in V$, $f(x, x) \in T_{\sigma,\epsilon} := \{\alpha + \alpha^\sigma \cdot \epsilon | \alpha \in D\}$.*
3. *There is a basis $\{e_i\}_{i \in I}$ of V such that $f(e_i, e_i) \in T_{\sigma,\epsilon}$.*

Proof The implications $1 \Rightarrow 2 \Rightarrow 3$ are immediate. To show 3 implies 1, totally order I, choose $g_i \in D$ so that $f(e_i, e_i) = g_i + g_i^\sigma \epsilon$ (possible by condition 3). Then we can uniquely determine a σ-sesquilinear form g by setting

$$g(e_i, e_j) = \begin{cases} f(e_i, e_j) & \text{if } i < j, \\ g_i & \text{if i = j,} \\ 0 & \text{if } i > j, \end{cases}$$

noting that g satisfies the conditions required from 1. □

Theorem 7.6.11 (The geometric condition for being trace-valued.) *Suppose f is a reflexive (σ, ϵ)-Hermitian form on V for which Rad $(V) = V^\perp$ does not contain all totally isotropic subspaces. Then V is spanned by its totally isotropic subspaces if and only if the form f is trace-valued.*

Proof First, if V is spanned by its singular spaces, then there is a basis of singular vectors – that is vectors e_i such that $f(e_i, e_i) = 0$, and so condition 3 of Theorem 7.6.10 holds in a trivial way.

We will complete the proof by showing that condition 2 of Theorem 7.6.10 implies V being spanned by its totally isotropic subspaces.

By hypothesis, there exists a vector u not in $V^\perp$ for which $f(u, u) = 0$. Let L be any two-dimensional subspace of V containing u. We claim that L is spanned by its isotropic 1-spaces, and to accomplish that we need only show that there is one other besides $\langle u \rangle$.

So we can assume that $L = \langle x, u \rangle$ where $f(x, x)$ is non-zero. Since f is trace-valued, $f(x, x) = \mu + \mu^\sigma \epsilon$, for some scalar μ in D. We are still free to scale u (that is, replace it by a scalar multiple) so that $f(x, u) = -\mu$. Now we compute

$$\begin{aligned} f(x + u, x + u) &= f(x, x) + f(x, u) + f(u, x) + f(u, u) \\ &= f(x, x) + f(x, u) + f(x, u)^\sigma \epsilon \\ &= (\mu + \mu^\sigma \epsilon) - \mu - \mu^\sigma \epsilon = 0, \end{aligned}$$

so $\langle x + u \rangle$ is a second totally isotropic 1-space of L.

Since V is clearly spanned by all the 2-subspaces containing u, the result follows. □

Corollary 7.6.12 *Let σ, ϵ be an admissible pair for the division ring D. The following statements are equivalent:*

1. *Every (σ, ϵ)-Hermitian form is trace-valued.*
2. *$T_{\sigma,\epsilon} = T^{\sigma,\epsilon}$.*

As noted, the condition holds if D has characteristic distinct from 2.

Proof 1 implies 2. Suppose $\gamma \in T^{\sigma,\epsilon}$. Then we can define a (σ, ϵ)-Hermitian form on a 1-space $\langle e \rangle$ by setting $f(e, e) = \gamma$. Then as f is tracevalued, $\gamma \in T_{\sigma,\epsilon}$ by condition 2 of Theorem 7.6.10.

2 implies 1. Suppose f is a (σ, ϵ)-Hermitian form. Then automatically, $f(x, x) \in T^{\sigma,\epsilon}$. But since $T^{\sigma,\epsilon} = T_{\sigma,\epsilon}$ we see that condition 2 of Theorem 7.6.10 holds, and so f is trace-valued.

The last statement about characteristic distinct from 2 is in Lemma 7.6.9, part 6. □

7.6.5 PseudoQuadratic Forms

Introduction

In the previous section we considered polar spaces that were obtained from quasi-polarities. Except for a trivial rank one case in the projective line, these all turned out to be polar spaces $\mathcal{S}_f$ for some (σ, ϵ)-Hermitian form f. If these were the only classical polar spaces to contend with, life would be much simpler. However, even these classical polar spaces sometimes carry within them further polar subspaces that cannot be described in this way.

The Definition

We begin with the following theorem.

Theorem 7.6.13 (The defining properties of pseudoquadratic forms.) *Let (σ, ϵ) be an admissible pair for the division ring D. The symbol $D/D_{\sigma,\epsilon}$ will refer to the quotient group of $(D, +)$ and $(D_{\sigma,\epsilon}, +)$ as additive groups. The following properties of a mapping $q : V \to D/D_{\sigma,\epsilon}$ are equivalent:*

(i) There exists a σ-sesquilinear form $g : V \times V \to D$ such that

$$q(x) = g(x, x) + D_{\sigma,\epsilon} \tag{7.21}$$

for all vectors $x \in V$.

(ii) For any $(x, \lambda) \in V \times D$, $q(x\lambda) = \lambda^\sigma q(x)\lambda$, and there is a trace-valued (σ, ϵ)-form $f : V \times V \to D$ such that

$$q(x + y) = q(x) + q(y) + f(x, y) + D_{\sigma,\epsilon}, \tag{7.22}$$

for all $x, y \in V$.

Moreover, if g satifies (7.21), then the form f in (7.22) can be chosen so that

$$f(x, y) = g(x, y) + g(y, x)^\sigma \epsilon.$$

Proof Assume (i). Then

$$\begin{aligned} q(x+y)-q(x)-q(y) &= g(x+y,x+y)-g(x,x)-g(y,y)+D_{\sigma,\epsilon} \\ &= g(x,y)+g(y,x)+D_{\sigma,\epsilon} \\ &= g(x,y)+g(y,x)^{\sigma}\epsilon+(g(y,x)-g(y,x)^{\sigma}\epsilon)+D_{\sigma,\epsilon} \\ &= g(x,y)+g(y,x)^{\sigma}\epsilon+D_{\sigma,\epsilon}. \end{aligned}$$

So (ii) holds with f as given by $f(x,y) := g(x,y)+g(y,x)^{\sigma}\epsilon$.

Conversely, suppose g is given as in (ii). Then let $(e_i)_{i\in I}$ be a basis for V and set $q(e_i) := g_i + D_{\sigma,\epsilon}$.

By Zorn's lemma we can totally order the index set I, and use this ordering to define a σ-sesquilinear form g by

$$g(e_i,e_j) = \begin{cases} f(e_i,e_j) & \text{if } i<j, \\ g_i & \text{if } i=j, \\ 0 & \text{if } i>j. \end{cases}$$

Then Eq. (7.22) shows g satisfies (7.21) in (i). The proof is complete. □

A function q as in part (i) of Theorem 7.6.13, is called a (σ,ϵ)-*quadratic form* or *pseudoquadratic form*. The form f which then appears in part (ii) of the theorem is called the *reflexive sesquilinar form associated with* q. This form is uniquely determined by q, for if there were two different forms satisying Eq. (7.22), say f and f', then

$$f(x,y)-f'(x,y) \in D_{\sigma,\epsilon} \tag{7.23}$$

for all vectors x and y, and for at least one pair of vectors, x_0 and y_0, the difference on left side is non-zero. But then, replacing vector y_0 by any scalar multiple $y_0\lambda$ in (7.23), we see that $D = D_{\sigma,\epsilon}D \subseteq D_{\sigma,\epsilon}$. But in that case $\lambda^{\sigma}\epsilon = \lambda$ for every λ. Setting $\lambda = 1$ one gets $\epsilon = 1$, and $\sigma = 1$, from which $D_{\sigma,\epsilon} = 0$, now an absurdity.

The unique (σ,ϵ)-Hermitian form associated with q is denoted $\beta(q)$.

Some Proportional Special Cases

As with the case of sesquilinear forms, we can obtain nice forms proportional to any given pseudoquadratic form.

A $(\sigma,1)$-quadratic form is called a σ-*quadratic form*, and the even-more-special case of a (id, 1)-quadratic form is just a *quadratic form*.

Suppose q is a (σ,ϵ)-quadratic form. Fix a non-zero scalar λ. Let σ' be the composition $i_\lambda \circ \sigma$, where i_λ is conjugation $\tau \to \lambda\tau\lambda^{-1}$, an inner automorphism of D. Similarly set $\epsilon' := \lambda\lambda^{-\sigma}\epsilon$. Then the student may easily calculate that the form λq is a (σ',ϵ')-quadratic form. In this case we say that the form λq is *proportional to* q.

A simple application of Lemma 7.6.7, reveals the following.

Lemma 7.6.14 *Every pseudoquadratic form is proportional to a σ-quadratic form for suitable σ.*

7.6.6 The Polar Space of a Pseudoquadratic Form

Fix a (σ, ϵ)-quadratic form, q and let $f := \beta(q)$ be its associated (σ, ϵ)-Hermitian form. We say that a subspace U of V is *totally singular with respect to* q, if and only if $q(u) = 0$ for all vectors u of U. An immediate consequence of Eq. (7.22) is the following.

Lemma 7.6.15 *Every subspace that is totally singular with respect to q is totally isotropic with respect to its associated form $f = \beta q$. Moreover, if $\{X_i | i \in I\}$ is a family of totally singular subspaces of q which are pairwise orthogonal with respect to f — that is, $X_i \subseteq \pi_f(X_j)$ for all distinct pairs (i, j), then the sum $\sum_I X_i$ is also totally singular.*

Now let p and r be any 1-spaces of vector space V that are totally singular with respect to the pseudoquadratic form q. Lemma 7.6.15 tells us that $p \in \pi_f(r)$ if and only if p and r span a totally singular 2-subspace of V. Let $\mathcal{P}_q$ and $\mathcal{L}_q$ be the collections of all totally singular 1- and 2-subspaces of V and let $\mathcal{S}_q = (\mathcal{P}_q, \mathcal{L}_q)$ be the point-line geometry obtained from the usual incidence by the containment relation. The remarks just made above show that $\mathcal{S}_q$ is a subspace of the polar space $\mathcal{S}_f$ defined in the previous section, and so is itself is a polar space.

We call $\mathcal{S}_q$ the *polar space of the pseudoquadratic form q*. Obviously, if $\lambda \neq 0$, q and λq vanish on exactly the same set of vectors, and so a space is totally singular with respect to q if and only if it is totally singular with respect to λq. In particular $\mathcal{S}_q = \mathcal{S}_{\lambda q}$, the same polar space.

The Relationship of the Polar Spaces for q and $\beta(q)$

One might first consider this question at the level of the forms rather than the spaces.

Theorem 7.6.16 *Let $Q_{\sigma,\epsilon}$ be the collection of all (σ, ϵ)-quadratic forms, and let $S_{\sigma,\epsilon}$ be the collection of all trace-valued (σ, ϵ)-Hermitian forms. The following statements hold:*

1. *The "sesquilinearization mapping"*

$$\beta : Q_{\sigma,\epsilon} \to S_{\sigma,\epsilon}$$

   is surjective.
2. *Any pseudoquadratic form in* $\ker \beta$ *takes its values in $D^{\sigma,\epsilon}/D_{\sigma,\epsilon}$.*
3. *β is bijective if and only if $D^{\sigma,\epsilon} = D_{\sigma,\epsilon}$ (or equivalently, every (σ, ϵ)-Hermitian form is trace-valued).*

Proof We just sketch this. Take any $f \in S_{\sigma,\epsilon}$. Since it is trace-valued, it has shape $f = g + (g^T)^\sigma \epsilon$ where g is σ-sesquilinear. But now that g is available, we can use it to define q by $q(x) = g(x, x) + D_{\sigma,\epsilon}$. Then a direct calculation gives $\beta(q) = f$. So we have the surjectivity.

The statement about $\ker \beta$ is straight from the definitions.

If $\beta(q) = \beta(q')$, then (noting that $Q_{\sigma,\epsilon}$ is an additive group, though not a vector space), $q - q' \in \ker \beta$. The third statement follows directly from this observation. □

Let's see what this means for the associated polar spaces.

Theorem 7.6.17 *Let q be a (σ, ϵ)-quadratic form on V, and let $f = \beta(q)$. If f is alternating and char $D \neq 2$, then $q = 0$, so $\mathcal{S}_q = P(V)$ as point-line spaces.*

So assume f is not alternating. Then

(i) *$\mathcal{S}_q$ is a subspace of $\mathcal{S}_f$.*
(ii) *If $\mathcal{P}_q$ spans $P(V)$, then $Rad(\mathcal{S}_q) = \mathcal{P}_q \cap Rad(\mathcal{S}_f)$.*
(iii) *If $D_{\sigma,\epsilon} = D^{\sigma,\epsilon}$, then $\mathcal{S}_q = \mathcal{S}_f$.*

Proof We proved (i) above.

(ii) Suppose $P(V)$ is spanned by its totally singular 1-spaces. By (i), $\mathrm{Rad}(\mathcal{S}_f) \cap \mathcal{P}_q \subseteq \mathrm{Rad}(\mathcal{S}_q)$. Now assume x is a point of $\mathrm{Rad}(\mathcal{S}_f)$. Then $\mathcal{P}_q \subseteq \pi_f(x)$. But as $\mathcal{P}_q$ spans $P(V)$, we see $\pi_f(x) = P(V)$, so $x \in \mathrm{Rad}(\mathcal{S}_f)$.

(iii) By (i), it suffices to show that $\mathcal{S}_f \subseteq \mathcal{S}_q$. Suppose $D_{\sigma,\epsilon} = D^{\sigma,\epsilon}$. Take any $x \in \mathcal{S}_f$. Then $0 = f(x, x) = g(x, x) + g(x, x)^\sigma \epsilon$, for g given as in part (i) of Theorem 7.6.13. So $g(x, x) \in D^{\sigma,\epsilon} = D_{\sigma,\epsilon}$, and, by Eq. (7.21), we have $q(x) = 0$ so $x \in \mathcal{S}_q$. □

The Augmented Pseudoquadratic Form

The importance of this section is to continue the comparison of the polar spaces of σ-Hermitian and σ-quadratic forms by showing that if f is not a symplectic form in odd characteristic, then in fact $\mathcal{S}_f$ is isomorphic to some $\mathcal{S}_q$ for a suitable σ-quadratic form q.

We have mentioned before that the $D^{\sigma,\epsilon}/D_{\sigma,\epsilon}$ is a vector space where "right scalar multiplication" is defined by

$$\alpha * \lambda + D_{\sigma,\epsilon} = \lambda^\sigma \alpha \lambda + D_{\sigma,\epsilon}.$$

Lemma 7.6.18 *The identity function on $D^{\sigma,\epsilon}/D_{\sigma,\epsilon}$ is a (σ, ϵ)-quadratic form $q_{\sigma,\epsilon}$ whose associated reflexive sesquilinear form f is zero.*

Now let q be a non-zero σ-quadratic form on vector space V (so $\epsilon = 1$ here), and let f be its associated σ-Hermitian form. We now enlarge the vector space V to

$$\hat{V} := V \oplus (D^{\sigma,1}/D_{\sigma,1}),$$

and we define the function $\hat{q} = q \oplus q_{\sigma,1}$ by

$$\hat{q}(v, \alpha + D_{\sigma,1}) = q(v) + \alpha + D_{\sigma,1}, \tag{7.24}$$

where $(v, \alpha + D_{\sigma,1})$ is any vector of the direct sum $V \oplus (D^{\sigma,1}/D_{\sigma,1}) = \hat{V}$.

We call $\hat{q}$ the *augmented form of q*.

Theorem 7.6.19 *$\hat{q}$ is a σ-quadratic form on $\hat{V}$. Then, of course, $S_{\hat{q}}$ is a polar space naturally embedded in $\mathbf{P}(\hat{V})$. Now S_f is a polar space embedded in V, and the linear mapping $\hat{V} \to V$ defined by projection on the first summand in (7.24), induces an isomorphism of the polar spaces $\mathcal{S}_{\hat{q}} \to \mathcal{S}_f$.*

Proof This is straightforward. □

At this stage we can say the following.

Corollary 7.6.20 *A classical polar space has one of these two presentations: (1) it is S_f where f is an alternating form over a vector space in characteristic not 2, or (2) it is a polar space S_q defined by a σ-quadratic form.*

7.7 Polar Spaces Embedded in Projective Spaces

7.7.1 Introduction

Through most of this section S is a polar space embedded in a projective space $\mathbf{P}(V) = (\mathcal{P}(V), \mathcal{L}(V))$, the geometry of 1- and 2-subspaces of a right vector space V over a division ring. Thus $S = (\mathcal{P}, \mathcal{L})$ is a polar space where $\mathcal{P} \subseteq \mathcal{P}(V)$ and $\mathcal{L} \subseteq \mathcal{L}(V)$, each projective point incident with a line of $\mathcal{L}$ belongs to $\mathcal{P}$, and V is spanned by $\mathcal{P}$. The embedding is denoted $e : S \hookrightarrow \mathbf{P}(V)$ (a new notation to indicate that e is defined by the inclusion relation). Then $\mathcal{V}_e$, which denotes the set of all geometric hyperplanes of S which arise from this embedding, is simply the set of geometric hyperplanes H of the form $\mathbf{H} \cap \mathcal{P}$, where $\mathbf{H}$ is a projective hyperplane of $\mathbf{P}(V)$.

We have a bit of a problem. For any subset X of the point-set $\mathcal{P}$, we must distinguish between the *subspace of S* spanned by X and the *subspace of $P(V)$* spanned by X. We always denote the former (always a subset of $\mathcal{P}$) by $\langle X \rangle_S$ while the latter (which may include many projective points not in $\mathcal{P}$) is a subspace of $\mathbf{P}(V)$ which we denote with square brackets $[X]$.

Recall that the *rank* of a non-degenerate polar space is defined to be one plus the projective dimension of a (necessarily projective) maximal singular subspace, and that we have extended this definition to all polar spaces S by setting the rank of S to be that of its non-degenerate section $\bar{S} = S/\mathrm{Rad}(S)$.

[For example, if a polar space is just a line and its points, then its rank is zero, since the non-degenerate section is actually empty. The same is true for any singular space. A star (a "bouquet" of lines on a point as defined for quadrangles), or a "bouquet" of planes on a line, or $PG(3)$s on a plane, and so forth, are all *rank one* polar spaces.[17] There is still nothing interesting going on. Rank at least 2 means there are lines which miss the radical. Rank at least $n \geq 2$ implies that every singular

[17] Whoever invented this descriptive and beautiful term "bouquet" should get a prize.

space of rank $r \leq n$ which misses the radical lies in a projective space of rank n which also does so.]

We can then record an item already proved.

Lemma 7.7.1 *Let r be a non-negative integer. If S has rank at least $r+1$, the subspaces of $\mathbf{P}(V)$ of codimension r or less are spanned by the points of $\mathcal{P}$ embedded in them. In particular if S has rank at least three, this is true of subspaces of codimensions 1 and 2 in V, and so there is an isomorphism*

$$\rho : \mathbf{P}(V) \to \mathcal{V}_e,$$

of projective spaces.

Proof By Lemma 7.5.1 Veldkamp r-spaces exist for S. Then Theorem 4.4.1 part 1 (p. 94, Chap. 4) implies the first statement. The statement about the Veldkamp space being projective follows from part 2 of Lemma 7.5.1. □

7.7.2 The Geometry of Embedded Polar Spaces

Lemma 7.7.2 *Suppose π is a plane in $\mathbf{P}(V)$ carrying at least one point of S. Then one of the following four possibilities must hold:*

1. $\pi \subseteq S$.
2. $\pi \cap S$ *is a pencil of at least two lines of π on a unique point z of the plane.*
3. $\pi \cap S$ *is a single line.*
4. $\pi \cap S$ *is a coclique of points of S; there are no lines of S in π.*

Now the lines $\mathcal{L}(V)$ of the projective space are divided into four classes: (1) *external lines* L, for which $L \cap \mathcal{P} = \emptyset$, (2) *secant lines* L which are not contained in S but contain at least two points of S, (3) lines which intersect the set of embedded polar space points $\mathcal{P}$ at a single point, and (4) the set $\mathcal{L}$, the (embedded) lines of S. The last two classes are lumped together and are called *tangent* lines.

Lemma 7.7.3 *Recall that $[p^\perp]$ is the subspace of $\mathbf{P}(V)$ spanned by the points of $\mathcal{P}$ which are collinear (in S) with p. Every (projective) line of $[p^\perp]$ which lies on p is a tangent line.*

Proof Let $\bar{P}$ be the rank two geometry of all lines and planes of $\mathbf{P}(V)$ on the point p. Then $\bar{P}$ is a projective space. Let $\bar{S}$ be the lines of S which are on p, let $\bar{T}$ be the tangent lines which intersect S exactly at p, and let $\bar{C}$ be the collection of all secant lines on point p. These three sets partition the "points" of the projective space $\bar{P}$. Suppose now that A is a plane spanned by a secant line pq ($q \in S - \{p\}$) on p and a line L belonging to S. Then by the polar space axiom, q is S-collinear with a point r of line L. This point r cannot be p, since pq is a secant line. Thus $A \cap S$ consists of some lines on point r as in case (2) of Lemma 7.7.2. But it cannot consist of all possible lines on r since the secant line pq carries at least one point not in S. Thus we see that with only the exception of line L, all other lines of A on p are

secant lines. This means that the partition $\bar{P} = \bar{S} + \bar{T} + \bar{C}$ satisfies condition (*) of Theorem 6.9.2 of Chap. 6.

The conclusion of that lemma is that the subspace of $\bar{P}$ spanned by $\bar{S}$ lies in $\bar{S} \cup \bar{T}$. But that is precisely the conclusion of this lemma. □

Lemma 7.7.4 *Assume S has polar rank at least 2. For any point $p \in \mathcal{P} - Rad(S)$, $[p^\perp]$ is a hyperplane of $\mathbf{P}(V)$ whose intersection with $\mathcal{P}$ is $p^\perp$.*

Proof Since p is not in Rad(S), there is a point q not S-collinear with p, so pq is a secant line. By the preceding lemma, $[p^\perp]$ is a proper subspace of $\mathbf{P}(V)$ and so lies in a *projective* hyperplane, H. Then $H \cap S$ is a geometric hyperplane of S containing another geometric hyperplane $p^\perp$. Since S has rank at least two, Veldkamp lines exist for S and so these geometric hyperplanes coincide – that is, $p^\perp = H \cap S$.

Also, since S is an embedded space with Veldkamp lines, $H = [H \cap S]$ and the lemma is proved.

Lemma 7.7.5 *Every tangent line meeting S at a point p lies in the space $[p^\perp]$.*

Proof If the point p in question lies in the radical Rad(S), then $p^\perp = S$ spans all of $\mathbf{P}$ and there is nothing to prove.

So let us fix $p \in S - \text{Rad}(S)$ and let $\bar{P}$ be the projective space consisting of the lines and planes of $\mathbf{P}$ that are incident with p. As in the proof of Lemma 7.7.3, the lines of $\mathbf{P}$ on point p (that is, the "points" of $\bar{P}$) are partitioned into three sets: $\bar{S}$, $\bar{T}$, and $\bar{C}$, respectively denoting lines on p that belong to S, that meet S exactly at p, or that are secant lines. As previously noted, condition (*) of Theorem 6.9.2 of Chap. 6 must hold for this partition. By Lemma 7.7.4 $[p^\perp]$ is a hyperplane of $\mathbf{P}$ and so $\langle \bar{S} \rangle_{\bar{P}}$ is a hyperplane of $\bar{P}$. Moreover, as p is not in the radical, secant lines on p must exist, so $\bar{C}$ is non-empty. It follows from Corollary 6.9.4 that $\langle \bar{S} \rangle_{\bar{P}} = \bar{S} + \bar{T}$ exactly. This means $\bar{T} \subseteq \langle \bar{S} \rangle_{\bar{P}}$, which is our conclusion. □

Now let x be any projective point of $\mathbf{P}(V)$. The set $H_x := \{p \in \mathcal{P} | x \in [p^\perp]\}$ is called the *collar of the point x*. If p is in H_x, then the projective line xp is a tangent line (Lemma 7.7.3). Note that if $x \in \mathcal{P}$, the statement $p \in H_x$ is equivalent to $xp \in \mathcal{L}$ or $p \in x^\perp$, so $H_x = x^\perp$ in that case.

If $H_x = \mathcal{P}$, x is said to be *defective point*. From the previous paragraph, the defective points which happen to lie in $\mathcal{P}$ are just the points of the radical Rad (S). We shall let D denote the set of all defective points of $\mathbf{P}(V)$.

Lemma 7.7.6 *If x is not defective, then H_x is a geometric hyperplane of S.*

Proof Assume x is not defective, so H_x is a proper subset of $\mathcal{P}$. Consider any line $L \in \mathcal{L}$. We must show that either exactly one point of L lies in H_x or that $L \subseteq H_x$. If $x \in S$ then $H_x = x^\perp$ is already a hyperplane of S and there is nothing to prove. So we may assume $x \in \mathbf{P} - S$. In particular $x \notin L$ so we can form the projective plane $A := [\{x\} \cup L]$.

If Case 1 of Lemma 7.7.2 holds for A, then $x \in S$, against the reduction of the previous paragraph. So this case can be dismissed.

In Case 2, where the points of $S \cap A$ are those incident with two or more lines on a common point r, we have that xr is the only tangent line to L, all the rest are

secant lines. By Lemma 7.7.5, $x \in [r^\perp]$ (that is, $r \in H_x$ and by Lemma 7.7.3 for each point v of $L - \{r\}$, the secant line xv does not lie in $[v^\perp]$ (that is, v is not in H_x). Thus $|L \cap H_x| = 1$ as desired.

In Case 3 of Lemma 7.7.2, $S \cap A = L$, and all lines of A on x are tangent lines. Appealing to Lemma 7.7.5, $x \in [y^\perp]$, for each point y of L. This means $L \subseteq H_x$ in this case.

Case 4 doesn't arise since A contains line L.

Thus H_x meets each line of $\mathcal{L}$ in one or all of its points, and so it is a geometric hyperplane of S. The proof is complete. □

At this stage we seem to have on hand a map from the non-defective points of $P(V)$ into the points of the Veldkamp space $\mathcal{V} = \mathcal{V}(S)$ which takes any projective point x to its collar H_x. Our next aim is to show that this mapping defines a linear space morphism τ.

Theorem 7.7.7 *Let $\mathcal{V}(S)$ be the Velkamp space of the embedded polar space S of polar rank at least three, so that $\mathcal{V}(S)$ is a linear space. Let $\mathcal{V}'$ be the subspace of $\mathcal{V}(S)$ that is generated by the hyperplanes $\{p^\perp | p \in \mathcal{P}\}$. The mapping τ which sends each projective point x to its collar H_x is a surjective morphism of projective spaces*

$$\tau : P(V) \to \mathcal{V}'$$

whose kernel is the subspace D of defective points of $\mathbf{P}(V)$.

We shall prove this theorem in a series of steps. We must first establish how this mapping "tau" respects lines of the projective space. This endeavor is complicated by the need to manage the unavoidable defective points,[18] but can still be achieved in a few lemmas.

As usual, we maintain the assumption that S has polar rank at least 2.

Lemma 7.7.8 *Suppose L is a tangent line on the point $p \in S$, and suppose d is a defective point of $L - \{p\}$. Then for any point $x \in L - \{d\}$, one has $H_x = p^\perp$.*

Proof Suppose $H_x \neq p^\perp$ for some point $x \in L - \{d\}$. Since S has polar rank at least 2, geometric hyperplanes are maximal subspaces of S (Lemma 7.5.1). Thus there must exist a point $s \in p^\perp - H_x$. Then $N = ps$ is a line of S in $p^\perp$, and since s is not in H_x, x is not in $[p^\perp]$ and so is not in N. Thus $[\{p, x, s\}] = \pi$ is a projective plane containing line $N = ps$ of S. But then all lines of π on defective point d are tangents meeting N, so N is the only line of S in π – and we are in Case 3 of Lemma 7.7.2. But that is impossible since, by choice of s, the line xs is a secant line carrying a further point r of $S - N$.

Thus we must have $H_x = p^\perp$, as desired. □

There is a sort of converse to the preceding lemma.

[18] Even if S is non-degenerate, defective points may be present.

Lemma 7.7.9 *Suppose $L = up$ is a tangent line meeting S at point p. Suppose further, that $H_u = p^\perp$. (Note that by definition of L, u is distinct from p.) Then:*

1. *For each point $v \in L$, $H_v \supseteq H_u$.*
2. *There is a defective point in line L.*

Proof Part 1. Choose an arbitrary point v in line L. We need only show $H_v \supseteq p^\perp$. Let N be any line of S on p and form the plane $\pi := [N \cup \{u\}]$. Then $N \subseteq p^\perp = H_u$, so all lines of π on u meet N and are tangents, forcing $\pi \cap S = N$. Then $N \subseteq H_v$. Since N was arbitrary, $p^\perp \subseteq H_v$. We are done.

Part 2. If $p \in \mathrm{Rad}(S) \subseteq D$ we are done. So suppose q is a point of $S - p^\perp$. By Lemma 7.7.4, $[q^\perp]$ is a projective hyperplane of $\mathbf{P}(V)$, and so meets L in some point v distinct from p. Now H_v contains $p^\perp$ (by part 1) as well as q (since qv is a tangent line by Lemma 7.7.3). Since geometric hyperplanes are maximal subspaces of S, it follows that $H_v = S$ and so v is the desired defective point. The proof is complete. □

Lemma 7.7.10 *Suppose L is any line of $\mathbf{P}(V)$, and x, y, and z are three distinct points of L. Then*

$$H_x \cap H_y \subseteq H_z.$$

Proof If false, there exists a point $p \in H_x \cap H_y - H_z$. If L is a line of S, the conclusion reads $x^\perp \cap y^\perp \subseteq z^\perp$ which follows from the polar space property (P). If $p \in L$ then $xp = L$ is a tangent line simply because $p \in H_x$. But then $L = zp$ being tangent forces $p \in H_z$ contrary to our choice. Thus we must have that L is not in S and does not contain the point p.

We can then form the plane $\pi := [L \cup \{p\}]$. By choice of p, xp and yp are tangent lines while zp is a secant line containing a point $r \in zp \cap S - p^\perp$. Now if both tangents xp and yp were in $[p^\perp]$, then the entire line L would lie in $[p^\perp]$, whereupon Lemma 7.7.3 would force zp to be a tangent line, contrary to its being secant. Thus at least one of the two tangent lines – say xp – does not lie in $[p^\perp]$. But that contradicts Lemma 7.7.5. So the Lemma holds. □

Corollary 7.7.11 *Again, suppose L is any line, and that for distinct points x and y on L, $H_x = H_y$. Then $L \cap D \neq \emptyset$.*

Proof Now suppose $H_x = H_y$. We claim that L contains a defective point. To this end we may assume x is not defective. Then there is a point $p \in S - H_x$. Then the line L must meet the projective hyperplane $[p^\perp]$ in at least one point, say, z. Then by Lemma 7.7.10, H_z contains $H_x \cap H_y = H_x$ as well as the extra point p, and so $H_z = S$ since H_x is a maximal subspace of S. □

Corollary 7.7.12 *The set of defective points D is a subspace of $\mathbf{P}(V)$.*

Proof Suppose x and y were two distinct defective points of a line L. Then by Lemma 7.7.10, $H_z = S$ for any further point z of L. □

Corollary 7.7.13 *For any point $p \in S$, one has $D \subseteq [p^\perp]$.*

Proof Suppose $d \in D - [p^\perp]$ for some point p in S. Then $p \notin \mathrm{Rad}(S)$ so that $[p^\perp]$ is a projective hyperplane. Then there is a point $q \in S - p^\perp$, and we may form the projective plane $\pi := [d, p, q]$.

Now $\pi \cap [p^\perp]$ is a line N on p, but not containing d by hypothesis. Then N meets the line dq of π at a point y. Now yp is a tangent line by Lemma 7.7.3. But as d is a defective point of the dq, $H_y = q^\perp$ (a consequence of Lemma 7.7.10) and so does not contain p. That is, $yp = N$ is not a tangent line after all – a contradiction.

Thus, no such point d exists, and the proof of the corollary is complete. □

Lemma 7.7.14 *Assume S has polar rank at least three so that Veldkamp lines exist (Lemma 7.5.1). Suppose L is a projective line of $\mathbf{P}(V)$ with no defective points. Then the collection $\tau(L) := \{H_v | v \in L\}$ is a full Veldkamp line.*

Proof Suppose x and y are distinct points of projective line L and let H be any geometric hyperplane of S which contains $H_x \cap H_y$. We want to show that H is an image of τ, and for that purpose we may assume H is distinct from both H_x and H_y. Now H is a polar space of rank at least two, and so by Corollary 7.5.4 is not the union of its two hyperplanes $H \cap H_x$ and $H \cap H_y$. So there is a point $p \in H - (H_x \cup H_y)$. If p is defective, then $[p^\perp] = [S] = \mathbf{P}(V)$. If p is not defective, then $[p^\perp]$ is a projective hyperplane (Lemma 7.7.4). In either case, the projective line L on x and y meets $[p^\perp]$ at a point z. Then by Lemma 7.7.10, H_z contains $H_x \cap H_y$ and p, as does H. Then, as Veldkamp lines exist, $H = H_z = \tau(z)$, and the argument is complete:

$$\{H_x | x \in L\} = \tau(L)$$

is a full Veldkamp line. □

Proof of Theorem 7.7.7 We have a mapping $t : \mathbf{P}(V) \to \mathcal{V} \cup \{S\}$ which takes each projective point x to its collar H_x, and assumes the value "S" only on the points of subspace D. Let x and y be distinct projective points not in D. By Lemma 7.7.10 and Corollary 7.7.11, $H_x = H_y$ if and only if the projective line $L = xy$ contains a point of D. On the other hand, if $L \cap D = \emptyset$, Lemmas 7.7.10 and 7.7.14 show that $t(L)$ is a line of the Veldkamp space. Since Corollary 7.7.12 implies D is a projective subspace, we have all the ingredients of a projective space morphism

$$\tau : \mathbf{P}(V) \to \mathcal{V},$$

with kernel D.

The image $\mathcal{V}'$ is then a subspace of the linear Veldkamp space, one that is isomorphic to the projective space $\mathbf{P}(V/D)$. The fact that $\mathcal{V}'$ (here defined as the subspace of the linear space $\mathcal{V}(S)$ spanned by all the hyperplanes $\{H_x | x \in \mathbf{P}(V) - D\}$) is actually spanned by the subcollection $\{s^\perp | s \in S\}$ follows from the fact that any non-defective point not in S lies on some secant line L, and the Veldkamp line which is its image carries the two hyperplanes $s^\perp$ and $t^\perp$, where s and t are distinct points of $L \cap S$.

7.7.3 The Main Theorem

As before, S is an abstract polar space embedded into a projective space $\mathbf{P}(V)$. We shall assume that S has polar rank at least three. The goal of this section is achieved by the following theorem.

Theorem 7.7.15 *There exists a quasi-polarity π of the projective space $\mathbf{P}(V)$ whose associated classical polar space S_π of absolute points (those with $x \in \pi(x)$) and absolute lines (those projective lines with point-shadows in S_π) such that S is a subspace of S_π as point-line geometries. Precisely:*

1. *$S \subseteq S_\pi$ as point sets.*
2. *Every line of S is a line of S_π (as full point shadows).*
3. *Two points of S are S_π-collinear only if they are already collinear in S.*

Thus we have embeddings

$$S \hookrightarrow S_\pi \hookrightarrow \mathbf{P}(V).$$

Proof We have established that τ is a surjective morphism $P(V) \to \mathcal{V}'$ with kernel D, so $\mathcal{V}' \simeq \mathbf{P}(V/D)$. In particular $\mathcal{V}'$ is a projective space. Since S spans $\mathbf{P}(V)$, $\tau(S - D)$ spans $\mathcal{V}'$.

Now for each non-defective point p of S, $[p^\perp]$ is a hyperplane of $\mathbf{P}(V)$ (Lemma 7.7.4) and so $p^\perp$ arises from the given original embedding $e : S \to \mathbf{P}$. Thus $\tau(S - D)$ is a subset of $\mathcal{V}_e$, the collection of all geometric hyperplanes which arise from the embedding e. Since $\mathcal{V}_e$ is a subspace of $\mathcal{V}$ (Lemma 7.5.1 and Theorem 4.4.1 of Chap. 4), we see that $\mathcal{V}' \leq \mathcal{V}_e$. □

Remark It follows that $[H_x]$ is a hyperplane of $\mathbf{P}(V)$ for each non-defective point x of $\mathbf{P}(V)$ – a fact we had not proved up to now.

Since Veldkamp lines exist we have a projective embedding $\rho : \mathcal{V}_e \to \mathbf{P}(V)^*$. Now let's piece this altogether: we have a chain of morphisms:

$$\mathbf{P}(V) \xrightarrow{\tau} \mathcal{V}' \hookrightarrow \mathcal{V}_e \xrightarrow{\rho} \mathbf{P}(V)^*,$$

whose composition π takes each non-defective point x to the projective hyperplane $[H_x]$ and which is a morphism of projective spaces with kernel D, as we have defined such morphisms.

Now we note that for each pair of points $(r, s) \in S \times (S - \mathrm{Rad}(S))$ the statements $r \in H_s$ and $s \in H_r$ are equivalent, for they mean only that s and r are collinear as points of S. Now it is easy to see the following.

> *Let G be the point collinearity graph of $S - Rad(S)$ where S is a polar space of rank at least 3. Then the distance-two graph of the complement, which we call $\Delta(G)$, is connected.*[19]

[19] The reader might verify that $\Delta(G)$ is not connected when G is a grid – a rank two polar space.

Now let x be any point of $\mathbf{P}(V) - D$ and let p be a point of $S - \mathrm{Rad}(S)$.

Suppose, first, that $x \in [p^\perp]$. Then by Lemma 7.7.3, px is a tangent line, and so by definition, $p \in H_x$. We record this as

$$x \in \pi(p) := [p^\perp] \text{ implies } p \in \pi(x) := [H_x].$$

Next suppose $p \in [H_x]$. Since x is not defective, $[H_x]$ is a hyperplane of $\mathbf{P}(V)$. We claim that $[H_x] \cap S = H_x$. Certainly the left side contains the right. Suppose there were a point $s \in S - H_x$ in this hyperplane. Then the subspace $M := \langle s, H_x \rangle_S$ also lies in this projective hyperplane. But, as the geometric hyperplane H_x is a maximal subspace of S, M is in fact S. But in that case, the projective hyperplane $[H_x]$ would contain a spanning set of $\mathbf{P}(V)$, an absurdity. Thus we can conclude, that the point p of S chosen in $[H_x]$ actually lies in H_x. Thus px is a tangent line. Then by Lemma 7.7.5, $px \subseteq [p^\perp]$, in which case $x \in [p^\perp]$. We record this as

$$p \in \pi(x) := [H_x] \text{ implies } x \in \pi(p) := [p^\perp].$$

It now follows from Corollary 7.6.6, p. 201, (with $\mathbf{P}(V) - D$ in the role of X, and $S - \mathrm{Rad}\,(S)$ in the role of S) that π is a quasi-polarity defined by a (σ, ϵ)-Hermitian form $h : V \times V \to D$.

Now, as observed, for two points x and y of S, $x \in x^\perp \subseteq [x^\perp] = \pi(x)$ and x is collinear with y if and only if $x \in H_y = y^\perp$ or equivalently $y \in H_x = x^\perp$. Thus collinearity in S occurs if and only if it represents collinearity in S_π. Thus two points of S are collinear if and only if they are collinear in S_π.

All parts of Theorem 7.7.15 have been proved.

7.8 Polar Spaces Embedded in a Classical Polar Space

Suppose S is a polar space of rank at least 3 which is a subspace of a classical polar space S_π of absolute points of some quasi-polarity $\pi : \mathbf{P}(V) \to \mathbf{P}(V)^*$. Then any embedding of S_π provides an embedding of S into the subspace its images span.

We have see from Sect. 7.4 that there are basically just two possibilities that capture all possible S_π:

1. $S_\pi \simeq S_f$ for an alternating form f in characteristic not equal to 2.
2. $S_\pi \simeq S_q$ for some σ-quadratic form.

In either case, there is a natural ambient space $\mathbf{P}(W)$ supporting one of these forms. Let $\mathbf{P}(V)$ be the subspace of $\mathbf{P}(W)$ spanned by points of S. We may then restrict the forms f or q to the vector space V and without loss of generality, take S_f or S_q to be the classical subspace of absolute or singular points of this restriction. One now has the following assumptions.

Hypothesis

1. *S is a polar space of rank at least 3 which (as a point-line geometry) is a subspace of a classical polar space $S' \simeq S_f$ or S_q defined [in Hypothesis 1.] by forms on vector space V. (Here S_f denotes the isotropic points of an alternating form f in characteristic not 2 and S_q is the singular points of a σ-quadratic form.)*
2. *The subspace S is embedded in $\mathbf{P}(V)$, that is, S spans $\mathbf{P}(V)$ (or $[S] = \mathbf{P}(V)$ in the notation of Sect. 7.5).*

Associated with S' is a quasi-polarity π and we now set S_π to be the polar space of its absolute points. Note that $S_\pi = \mathbf{P}(V)$ in the the alternating case, but that conceivably one might have a chain $S \subseteq S' \subseteq S_\pi$ that is proper in the case $S' \simeq S_q$.

Since S and S' are embedded polar spaces, we may import all of the available theory about such spaces from Sect. 7.5 (including the notation in which square brackets denote subspaces of the ambient projective space generated by the argument they enclose). Now as embedded polar spaces, S and S' themselves induce quasi-polarities of $\mathbf{P}(V)$ whose value at any non-defective point s is the hyperplane $[s^\perp]$ consisting of the set-theoretic union of all projective lines on s which meet S' (or S) at one point or lie entirely within S' (S). (Recall from Sect. 7.5 that both sorts of lines were called *tangent lines*.) Since both these spaces have rank at least three (for S is a subspace of S' containing a pair of opposite planes) these two projective hyperplanes, $[s^\perp]$ and $[s^\perp \cap S]$ coincide for any point s within S. But since S' is a classical space, this hyperplane is precisely $\pi(s)$ for the quasi-polarity π associated with S'. Thus we have the following.

Lemma 7.8.1 *Let s be a point of S. Then $\pi(s) = [s^\perp := \pi(s) \cap S'] = [s^\perp \cap S]$ and is the union of all tangent lines of S' on s as well as the union of all tangent lines of S on s ("tangent" being used in the sense of Sect. 7.5). In particular, any line of $\mathbf{P}(V)$ which meets S at a single point can meet S' at some further point only if it is one of the lines of the polar space S'.*

An immediate consequence is the following.

Corollary 7.8.2 *Suppose x is a point in $S' - S$ and L is a projective line of $\mathbf{P}(V)$ on x and on a point $s \in S$. Then the following conditions are equivalent:*

1. *$L \cap S = \{s\}$.*
2. *$s \in \pi(x)$. (Equivalently, $x \perp s$ or x and s are collinear in S'.)*
3. *L is a line of S'.*

Corollary 7.8.3 *$Rad(S') \cap S = Rad(S)$.*

Theorem 7.8.4 *Under the stated hypothesis, $S = S'$.*

An immediate consequence of this theorem is the following.

Corollary 7.8.5 *Any polar space of polar rank at least three, which possesses an embedding into a projective space, is a classical polar space – i.e., one defined by an alternating form or a σ-quadratic form.*

Proof of Theorem 7.8.4 By way of contradiction, suppose there exists a point x in $S' - S$. Since S has rank at least one, there is a point m in $S - \text{Rad}\,(S)$. Then any point of $xm - \{x, m\}$ lies in $S' - (S \cup \text{Rad}\,(S'))$. So without loss of generality, we assume x lies in $S' - (S \cup \text{Rad}\,(S'))$.

Then by Corollary 7.8.2, there is a projective line M connecting x to a point of S which is not a tangent of S. Thus M is a secant line containing two points $[e_1]$ and $[e_2]$ of S. Then the $[e_i]$ are not collinear in S, and so do not lie in the radical of S. Since S has rank at least two, there is a line L on $[e_1]$ not meeting Rad (S). Then by Lemma 7.3.2, there exists a line L' on $[e_2]$ which is opposite L.

Thus, $W = [L, L']$ is a four-dimensional subspace of V which is non-degenerate with respect to the forms f or q, and contains our secant line M. Let us enlarge the set $\{e_i\}$ already introduced to a basis $B = \{e_1, e_2, e_3, e_4\}$ of W, where $L = [e_1, e_3]$, $L' = [e_2, e_4]$, and as before $M = [e_1, e_2]$. It will be convenient to use coordinates: we write

$$(\alpha_1, \alpha_2, \alpha_3, \alpha_4) := \sum\nolimits_{i=1} \alpha_1 e_i.$$

As noted, there are two basic cases for S': (1) $S' = S_f$ for a symplectic form f over a field D of characteristic not 2, and (2) $S' = S_q$ where q is a σ-quadratic form. The first case is very easy.

Case (1). We may take B to be a hyperbolic basis, so the form f is given by

$$f((\alpha_1, \ldots, \alpha_4), (\beta_1, \ldots, \beta_4)) = \alpha_1\beta_2 - \alpha_2\beta_1 + \alpha_3\beta_4 - \alpha_4\beta_3.$$

Since $x = [p] \leq [e_1, e_2]$ we may write $p = (1, \alpha, 0, 0)$.

Now note that

$$t_1 := (1, 0, 1, 0) \perp (0, -1, 0, 1) := r_1,$$
$$t_{-1} := (1, 0, -1, 0) \perp (0, 1, 0, 1) := r_{-1},$$

where the left members span distinct points of $L = [e_1, e_3]$ and the right-hand members span distinct points of $L' = [e_2, e_4]$. We thus obtain transverse lines $N_\delta = [t_\delta, r_\delta]$, $\delta = \pm 1$. Now note that the vectors

$$v_1 = (1, \alpha, \alpha, -1) \text{ and}$$
$$v_2 = (1, \alpha, -\alpha, 1)$$

lie in $N_1 \cap p^\perp$ and $\in N_{-1} \cap p^\perp$, respectively, and that

$$(v_1 + v_2)/2 = p.$$

Since $[v_1, p]$ is the line of S on v_1 and v_2, $x = [p]$ must be in S, a contradiction.

Case (2) is very similar: in fact the cases are separated only to clarify the calculations. Here q is a σ-quadratic form whose restriction to W is given by

$$q(\alpha_1, \alpha_2, \alpha_3, \alpha_4) = \alpha_1^\sigma \alpha_2 + \alpha_3^\sigma \alpha_4 + D_{\sigma,1}. \tag{7.25}$$

Its associated σ-Hermitian form on W is calculated as

$$f((\alpha_1, \ldots, \alpha_4), (\beta_1, \ldots, \beta_4)) = \alpha_1^\sigma \beta_2 + \alpha_2^\sigma \beta_1 + \alpha_3^\sigma \beta_4 + \alpha_2^\sigma \beta_3. \tag{7.26}$$

(Note that the "epsilon" that might have been expected in the second and fourth terms of the right-hand side of (7.26) does not make its appearance because $\epsilon = 1$.)

Now, we can put $x = [p]$ where $p = (1, \alpha, 0, 0)$. Then, as $x \in S_q$, we must have $\alpha \in D_{\sigma,1}$. (Note that $\alpha \neq 0$ since $x \neq [e_1] \in S$.) Thus there is a scalar θ such that

$$\alpha = \theta - \theta^\sigma.$$

Note that

$$p^\perp = \{(\gamma, \alpha\gamma, *, *)\}$$

where the "$*$" indicates an arbitrary entry from D.

We consider anew the opposite lines

$$L = \{(*, 0, *, 0)\} \text{ and } L' = \{(0, *, 0, *)\}.$$

Now we have

$$e_3 = (0, 0, 1, 0) \perp (0, 1, 0, 0) = e_2,$$

and for each scalar $\nu \in D$ we have

$$t_\nu := (1, 0, \nu, 0) \perp (0, -\nu^\sigma, 0, 1) := r_\nu.$$

Then $N_\nu := [t_\nu, r_\nu]$ is a line of S connecting a point of L to the unique point of L' perpendicular to it.

Now the line $N_{\beta^{-1}\alpha}$ contains the vector

$$\begin{aligned} v_\beta &:= t_{\beta^{-1}\alpha} + r_{\beta^{-1}\alpha}\beta^\sigma \\ &= (1, 0, \beta^{-1}\alpha, 0) + (0, -\alpha^\sigma\beta^{-\sigma}, 0, 1)\beta^\sigma \\ &= (1, -\alpha^\sigma, \beta^{-1}\alpha, \beta^\sigma) \\ &= (1, \alpha, \beta^{-1}\alpha, \beta^\sigma), \text{ (since } \alpha^\sigma = -\alpha \text{),} \end{aligned}$$

which is in $p^\perp$. So $[v_\beta] = N_{\beta^{-1}\alpha} \cap p^\perp$.

Now $\alpha = \theta - \theta^\sigma \neq 0$. So, if we set

$$\gamma_1 = \alpha^{-1}\theta \text{ and } \gamma_2 = -\alpha^{-1}\theta^\sigma$$

then

$$\gamma_1 + \gamma_2 = 1.$$

Let β_1 and β_2 be distinct non-zero scalars from D, and form the linear combination

$$v_{\beta_1} \cdot \gamma_1 + v_{\beta_2} \cdot \gamma_2 = (1, \alpha, {\beta_1}^{-1}\alpha\gamma_1 + \beta_2^{-1}\alpha\gamma_2, \beta_1^\sigma\gamma_1 + \beta_2^\sigma\gamma_2).$$

To get the last two coordinates equal to zero, we require

$$\begin{aligned} &\beta_1^{-1}\alpha \cdot \alpha^{-1}\theta - \beta_2^{-1}\alpha \cdot \alpha^{-1}\theta^\sigma = 0, \quad \text{or} \\ &\beta_1^{-1}\theta - \beta_2^{-1}\theta^\sigma = 0 \text{ and} \qquad (7.27) \\ &\beta_1^\sigma\alpha^{-1}\theta - \beta_2^\sigma\alpha^{-1}\theta^\sigma = 0. \qquad (7.28) \end{aligned}$$

From (7.27) we have

$$\beta_1^{-1}\theta = \beta_2^{-1}\theta^\sigma.$$

Then $\beta_2\beta_1^{-1} = \theta^\sigma \cdot \theta^{-1}$, so

$$\beta_2 = \theta^\sigma\theta^{-1}\beta_1.$$

Substituting in (7.28), we get

$$\begin{aligned} \beta_1^\sigma\alpha^{-1}\theta &= (\theta^\sigma\theta^{-1}\beta_1)^\sigma\alpha^{-1}\theta^\sigma \\ &= \beta_1^\sigma\theta^{-\sigma}\theta\alpha^{-1}\theta^\sigma. \end{aligned}$$

But then the β_1^σ cancels, giving

$$\alpha^{-1}\theta = \theta^{-\sigma}(\theta\alpha^{-1})\theta^\sigma.$$

Taking inverses of both sides yields

$$\begin{aligned} \theta^{-1}\alpha &= \theta^{-\sigma}(\alpha\theta^{-1})\theta^\sigma, \\ \theta^{-1}(\theta - \theta^\sigma) &= \theta^{-\sigma}(\theta - \theta^\sigma) \cdot \theta^{-1}\theta^\sigma, \\ 1 - \theta^{-1}\theta^\sigma &= \theta^{-\sigma}\theta \cdot \theta^{-1} \cdot \theta^\sigma - \theta^{-1}\theta^\sigma \\ &= 1 - \theta^{-1}\theta^\sigma, \end{aligned}$$

an identity! This means that once one chooses β_1 and sets $\beta_2 = \theta^\sigma \theta^{-1} \beta_1$, then both Eqs. (7.27) and (7.28) hold. (Note that $\alpha \neq 0$ implies $\theta^\sigma \neq \theta$, so β_1 and β_2 are necessarily distinct.) Then for such a pair β_1, β_2, we have

$$p = v_{\beta_1}\gamma_1 + v_{\beta_2}\gamma_2 \subseteq p^\perp,$$

so $x \in [v_{\beta_1}, v_{\beta_2}]$, a line of S. This again, contradicts the choice of x and completes the proof.

An Alternative Version of Theorem 7.8.4

Nothing in this minor section is actually required for what follows. It only offers an alternative development of the previous section.[20] The student should view this section as optional reading – a bayou to the side of the main flow of the idea-stream of this chapter.

The previous section has established the famous Lemma 8.10 of Tits' book, essentially following his proof (except for a corrected algebraic "typo"). It is worth mentioning that an entirely different approach to Theorem 7.8.4 was advanced by Francis Buekenhout and his colleagues. I will try to sketch these arguments. We can suppose that we are in this situation: The projective space $\mathbf{P}(V)$ is spanned by an unkown embedded non-degenerate polar space S of rank at least 3. From what we have shown, there is a polarity ρ of $\mathbf{P}(V)$ whose absolute points and lines form a classical polar space C such that S is a subspace of C. Now, if we wish, the embedding of C is a morphic image of an embedding $e : C \to \mathbf{P}(W)$ so the C is either a symplectic polar space over a field F of odd characteristic or is the set of all singular points of a σ-quadratic form. Of course we can then restrict the relevant forms to the subspace spanned by the images of the points of the subspace S under this new embedding of C. Thus without loss (and a slight redefining of C and V and the form) we may assume:

(x) 1. *C is the classical polar space defined by a σ-quadratic form q of vector space V or by a symplectic form in odd characteristic.*
2. *S is a subspace of C such that*
   (a) *S is a non-degenerate polar space of rank at least* 2.
   (b) *S spans the entire space $\mathbf{P}(V)$.*

That's it. In both cases there is a sesquilinear form f that is either the symplectic form in question, or is the sesquilinear form associated with the σ-quadratic form. Thus C is in turn a subspace of S_f, the classical polar space of absolute points with respect to the polarity derived from f – that is, the points $\langle x \rangle \in \mathbf{P}(V)$ such that $f(x, x) = 0$.

---

[20] The author is grateful to Peter Johnson for correcting the author's misunderstanding about this alternative.

The form f is not just random: it is derived from the quasi-polarity which restricts the mapping sending any point p of S (or C_f, for that matter) to the hyperplane H_p spanned by all lines of $\mathbf{P}(V)$ on p which are tangent with respect to S (or C_f). This tells us that S is very much the "child" of the polar space C_f. The lines S are lines of C_f. But more important, any projective line of $\mathbf{P}(V)$ which meets S at a single point s (and is thus tangent to S), is a line of $s^\perp$ – that is, if it contains just one more point of C_f, it is already a line of C_f. (The "$\perp$" symbol is being used with respect to the sesquilinear form f.)

Now we ask ourselves what would happen if there were a projective line L which met the unknown polar space S in at least two points, x and y, but was not a line contained in S. Then, for every third point v in $L - \{x, y\}$ there must exist an f-perspectivity ρ_v with center v. At this point, we must catch up on a few definitions appearing in Exercises 7.21–7.25 at the end of this chapter. One feature of all perspectivities with center v is that they fix all points of $\mathbf{P}(v^\perp)$. Thus ρ_v stabilizes the hyperplane $v^\perp$ pointwise, and so fixes the geometric hyperplane $v^\perp \cap S$ of S, pointwise. That means that if there exists any sort of ρ_v-invariant subset U of S not contained in the maximal subspace $v^\perp \cap S$, then $S = \langle U, v^\perp \cap S\rangle$ is certainly ρ-invariant. But if $S = S_v^\rho$, then $L \cap S$ is invariant under ρ_v for any $v \in L - S$.

Now if there is no projective line L of $\mathbf{P}(V)$ not lying in S, but picking up two points of S, then the point-set of S is itself a subspace of V. Since it is a subspace of C, the spanning requirement would force $S = C$. So such a line L must exist. Then (as L is not a line of S), we have the following.

Lemma 7.8.6 *For every perspectivity ρ with center p, $p \in L - S$, $L \cap S$ is ρ-invariant if and only if S is ρ invariant.*

Now suppose one could prove the following proposition.

Proposition 7.8.7 *Suppose* $\dim V = 2$ *and that f (or q) is a non-degenerate σ-Hermition form (or that q is a non-degenerate σ-quadratic form) such that $D^{\sigma,\epsilon} = D_{\sigma,\epsilon}$. Then there exists no proper subset S_L of $L \cap J$ (where J are the isotropic points of L with respect to f (the singular points of L with respect to q) which is invariant under every f- (q-) perspectivity of $\mathbf{P}(V)$ with center on L.*

This is a calculation in only a two-dimensional subspace of V – the one yielding the projective line L – whereas the Tits' argument presented above uses $PG(3, D)$. Either way they are calulations that depend on $D^{\sigma,\epsilon} = D_{\sigma,\epsilon}$. Still, it is a very nice idea.

Armed with this proposition, no "funny lines" like this exist and $\langle S\rangle = \mathbf{P}(V)$ finishes everything.

Remark Now there is nothing stopping the point v from belonging to S_q, in which case ρ_v is an f-transvection (or a q-transvection). The calculations in the *Handbook of Incidence Geometry* (p. 674 following Lemma 3.29) are slightly misleading in this respect: if $\langle v\rangle$ was restricted to $C - S$, we should only be able to use the above proposition to conclude that $(C-S) \subseteq S^\perp$ (exactly what is argued in the Handbook).

If that condition were true, all perspectivities in question are transvections, and then counterexamples to the proposition listed above (such as in Exercise 7.23) would exist. True enough, it was never stated in the Handbook that the center should belong to S, but the argument displayed in the proof is for that special case. So the reader should beware!

7.9 Non-degenerate Polar Spaces of Rank at Least Four Are Classical

7.9.1 Introduction

The purpose of this section is to prove the assertion which comprises its title.

We suppose throughout this section that S is a non-degenerate polar space of rank at least 4. Fix once and for all a specific point p of S. Let $C = p^\perp$ be the cone of all points collinear with p.

For each point x not in C, set $H_x := x^\perp \cap C$. Then H_x is a geometric hyperplane of C and itself is a non-degenerate polar space of rank at least 3. Note the following:

Lemma 7.9.1 *If G is a geometric hyperplane of H_x, then the subspace $\langle p, G\rangle_S$ is the union of all lines on p which meet G and is a geometic hyperplane of the cone $C = p^\perp$. In particular, if for points x, y in $S - C$, $H_x \neq H_y$, then $H^{xy} := \langle p, H_x \cap H_y\rangle_S$ is a geometric hyperplane of C and is a polar space of rank at least two.*

Proof This is an easy exercise. □

Now let $\mathcal{H}_p$ be the collection of all subspaces of C which are of the form H_x or $H^{xy}, x, y \in S-C$. Let $\mathcal{V}_p$ be the Veldkamp space defined internally by the set $\mathcal{H}_p$ – that is, a point-line geometry whose points are the elements of $\mathcal{H}_p$ and whose lines are intersections $H_1 \cap H_2$ of distinct pairs chosen from $\mathcal{H}_p$. Since this is a partial subspace of the full Veldkamp space $\mathcal{V}(C)$, for which Veldkamp lines exist, $\mathcal{V}_p$ is a linear space.

We intend to show a number of things:

(1) $\mathcal{V}_p$ is actually a subspace of $\mathcal{V}(C)$ – that is, lines of $\mathcal{V}_p$ are full lines of $\mathcal{V}(C)$.
(2) $\mathcal{V}_p$ is a projective space.

This will lead to an embedding of C into a subspace of the dual $\mathcal{V}_p^*$, which by the previous section is embedded as a classical polar space of some $\mathbf{P}(W)$ with a one-dimensional radical. One can then extend the classical form on W to a non-degenerate form on $W \oplus \langle t\rangle$ defining a classical polar space S'. Since C is now a cone of both S and S', by Tits' Theorem 7.5.13, S is isomorphic to S'. Since S' is classical, we are done.

7.9.2 Teirlinck's Condition in $\mathcal{V}_p$

Proving Teirlinck's condition is the most technically difficult part of the proof, but it will prove items 1 and 2 above, and will provide our desired embedding of C into a projective space.

One might ask why one can't prove Teirlinck's condition for the Veldkamp space $\mathcal{V}(S)$ of the *entire* polar space, thus implying that the latter is a projective space and then obtain an embedding of S into a subspace of its dual, from which one can conclude from the previous section that S is classical?

The answer is that it is even harder to prove Teirlinck's condition for a large collection of hyperplanes like $\mathcal{V}(S)$ than it is for $\mathcal{V}_p$, where we have a controllable collection of hyperplanes.[21] That is why we use the cone.

We begin with the following.

Lemma 7.9.2 *Suppose $x, y \in S - C$ and $r \in H_x - (H_x \cap H_y)$. Then there is a point $s \in H_x$ not collinear with r such that*

$$H_x \cap r^\perp \cap s^\perp \subseteq H_x \cap H_y.$$

Proof Let $L = xr$, the line on x and r. Since y is not collinear with r, $y^\perp \cap L = \{t\} \in S - C$. Now the line $M = yt$ meets C at a point $s' \in H_y$. Finally the line ps' meets H_x at a point s. (All of this is drawn in Fig. 7.5). Now if $r \sim s$ then $ps \subseteq r^\perp$, so $s' \in r^\perp$. But since $t \in r^\perp$ and t is distinct from s', y would be collinear with r, a contradiction. Thus s is not collinear with r.

Now suppose $z \in H_x \cap r^\perp \cap s^\perp$. Then $z \sim t$, and since $z \sim s$ implies $z \sim s'$, we see that $z \sim y$, so $z \in H_x \cap H_y$. □

Remark The curious thing about Lemma 7.9.2 is that its hypothesis and conclusion only involve H_x, its hyperplane $H_x \cap H_y$, and certain points in H_x. It doesn't refer to anything outside. This would make one think (replacing H_x by an arbitrary rank three polar space S_0) that if H were a hyperplane of polar space S_0 and $r \in S_0 - H$, then there would be a point s in $S_0 - H$ with $r^\perp \cap s^\perp \subseteq H$. However this is not true in general (see Exercises 7.27 and 7.28 at the end of this chapter).

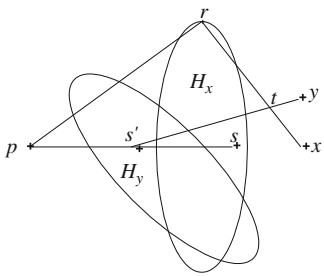

Fig. 7.5 The configuration of Lemma 7.9.2

[21] In other words, when selecting a set of hyperplanes $\mathcal{H}$ with which to apply Teirlinck's theory, *"less is better"* is generally a good maxim.

Lemma 7.9.3 *Suppose $H_x \neq H_y$, and let q be any point of C. Then there is a hyperplane H belonging to $\mathcal{H}_p$ which contains both q and $H_x \cap H_y$.*

Proof If $q = p$ take $H = H^{xy}$. Also if $q \in H_x \cup H_y$ we can take H to be H_x or H_y. Thus we may assume $q \in C - (H_x \cup H_y \cup \{p\})$.

Now line pq meets H_x at a point r. Then by Lemma 7.9.2 there is a point $s \in H_x$ not collinear to r such that

$$H_x \cap r^\perp \cap s^\perp \subseteq H_x \cap H_y.$$

Now if $H_x \cap H_y \subseteq r^\perp \cap s^\perp$, then one would have

$$H^{xy} \subseteq p^\perp \cap r^\perp \cap s^\perp \subseteq p^\perp \cap r^\perp \subseteq C.$$

The last inclusion in this chain is proper since s is in C but not $r^\perp$. Also the preceding inclusion is also proper since r is in the third term of the chain, but not in the second term. But that contradicts the fact that H^{xy} is a maximal subspace of C. Thus $H_x \cap H_y - (r^\perp \cap s^\perp)$ is not empty, and so $H_x \cap r^\perp \cap s^\perp$ is a proper subspace of $H_x \cap H_y$.

In fact it is a geometric hyperplane of it for this reason: H_x is a polar space of rank at least three, and so it possesses Veldkamp lines. Then $H_x \cap H_y$ is one of the Velkamp points of $\mathcal{V}(H_x)$ – one that contains the Veldkamp line $H_x \cap r^\perp \cap s^\perp$. So the latter is a hyperplane of $H_x \cap H_y$. [Indeed, $H_x \cap H_y \cap r^\perp = H_x \cap H_y \cap s^\perp = H_x \cap r^\perp \cap s^\perp$.]

So there is a point t in $H_x \cap H_y$ which is not in $r^\perp \cap s^\perp$, and from the maximality of hyperplanes,

$$H_x \cap H_y = \langle H_x \cap r^\perp \cap s^\perp, t \rangle_S.$$

There is a unique point u in $sx \cap q^\perp$, and there is a point z in $t^\perp \cap qu$ (see Fig. 7.6). Neither u nor z lie in C; for if $u \in C$, then $u = s$ and $z \in C$ and $u \in qu \cap C = \{q\}$, and that would force $t \in (pq)^\perp \cap (sx)^\perp \subseteq r^\perp \cap s^\perp$ against the fact that $t \in H_x - \cap r^\perp \cap s^\perp$. (The situation is depicted in Fig. 7.6.)

Now observe that any element of $H_x \cap r^\perp \cap s^\perp$ lies in $(xs)^\perp$ and hence $u^\perp$, and lies in $(pr)^\perp$ and hence $q^\perp$. But as z is on line qu, we have

$$H_x \cap r^\perp \cap s^\perp \subseteq z^\perp.$$

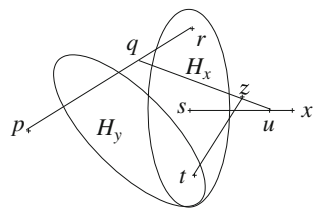

Fig. 7.6 The configuration for Lemma 7.9.3

But $z^\perp$ also contains t and so contains $H_x \cap H_y = \langle H_x \cap r^\perp \cap s^\perp, t\rangle$, as well as q. Thus the hyperplane H_z contains both $H_x \cap H_y$ and q, as desired, completing the proof. □

The next lemma asserts that if a Veldkamp line of $\mathcal{V}_p$ contains a Veldkamp point which is a hyperplane of C not on point p, then it contains at least two such hyperplanes.

Lemma 7.9.4 *Let H be any member of $\mathcal{H}_p$ which contains the point p. Then for each point x of $S - C$, there is a point y, with*

$$H = H^{xy} \text{ and } H \cap H_x = H \cap H_y.$$

Proof Let r and s be two points such that $H = H^{rs}$ (all hyperplanes in $\mathcal{H}_p$ which contain p have this form). Then $H \cap H_r = H \cap H_s$.

We are given x and wish to find y. Since $S - C$ is connected it suffices to prove this for the case that x is collinear with r. We can also assume $H_x \neq H_s$, otherwise we take $y = s$.

Now line xr meets C at point q and line pq meets H_s at a point t. If $H_r \cap H_s \subseteq H_x$, we can set $y = r$, and we are done. So we may assume that there exists a point $z \in (H_r \cap H_s) - (H_r \cap H_x \cap H_s H_x)$, and let z' be the point of H_x on line pz. Then z is not z'. Also, if z' were in H_r then the entire line $z'z = pz$ would be contained in H_r, against r being outside the cone $p^\perp = C$. Finally z' is collinear with a point y on the line ts. Now if this point y were in C, then $y = t$ and z' would be collinear with t, p, q, x, and hence r, and that would force $z = z' = pz \cap r^\perp$, against our choice of z. Also we see that if $H_x = H_y$ then $q^\perp$ would contain y and t and hence would contain s, forcing $q = pq \cap H_s = t$. In that case we would have $H_r \cap H_x = p^\perp \cap q^\perp = H_s \cap H_y$ and, comparing first and last terms, would lie in $H_r \cap H_s$. But then $H_r \cap H_x = H_r \cap H_s$ contrary to our assumption. Thus we see that y is not in C and that $H_x \neq H_y$, giving the configuration depicted in Fig. 7.7.

Now

$$H_r \cap H_s = \langle H_r \cap H_s \cap H_x, z'\rangle \subseteq (H \cap H_x) \cap H_x \cap H_y)$$

and adjoining p to both sides produces the same hyperplane $H^{rs} = H^{xy}$. □

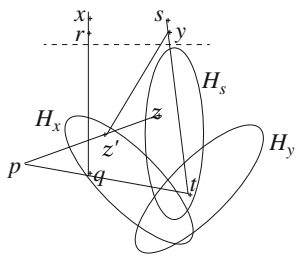

Fig. 7.7 The configuration for Lemma 7.9.4. The *zone below and to the left of the dashed line* belongs to $C = p^\perp$

Lemma 7.9.5 (Teirlinck's condition for the collection $\mathcal{H}_p$.) *Let H_1 and H_2 be any two hyperplanes of $\mathcal{H}_p$, and let q be any point of the cone $C = p^{\perp}$. Then there is a hyperplane H containing q and $H_1 \cap H_2$.*

Proof One can assume q is not in $H_1 \cup H_2$. Also, by Lemmas 7.9.3 and 7.9.4 we can assume that both H_1 and H_2 contain p. Now fix $x \in S - C$. By Lemma 7.9.4, there is a point $y \in S - C$ such that

$$H_1 = H^{xy} \text{ and } H_1 \cap H_x = H_1 \cap H_y.$$

Now, given this point y, we can apply Lemma 7.9.4 once more to produce a point $z \in S - C$ such that

$$H_2 = H^{yz}, H_2 \cap H_y = H_2 \cap H_z.$$

Now

$$\begin{aligned} H_1 \cap H_2 &= \langle p, H_x \cap H_y \rangle \cap \langle p, H_y \cap H_z \rangle \\ &\geq H^{xyz} := \langle p, H_x \cap H_y \cap H_z \rangle. \end{aligned}$$

But as $H_x \cap H_y \cap H_z$ is a hyperplane of $H_x \cap H_y$, we see that H^{xyz} is a hyperplane of $H^{xy} = H_1$. Similarly, H^{xyz} is a hyperplane of $H^{yz} = H_2$. Since C has rank at least three, we have

$$H^{xyz} := \langle p, H_x \cap H_y \cap H_z \rangle = H_1 \cap H_2.$$

Now if the line pq intersects $H_x \cap H_y$, we may reach the conclusion by setting $H = H^{xy}$. Similar conclusions hold if pq intersects $H_y \cap H_x$ or $H_x \cap H_z$. Thus we may assume that the line pq intersects H_x, H_y, and H_z at three distinct points q_x, q_y, and q_z, respectively.

Now by Lemma 7.9.4, there is a hyperplane $H_3 \in \mathcal{H}_p$, containing $H_x \cap H_y$ and q_z. Now if H_3 contained p, then line $pq_z = pq$ would meet $H_x \cap H_y$, contrary to our assumption. Thus $H_3 = H_w$ for some $w \in S - C$. Now $H_x \cap H_w = H_x \cap H_y$. But if $H_w = H_x$ then the equation of the previous sentence would imply $H_x = H_y$ giving $q_x = q_y$, a contradiction. Thus $H_w \neq H_x$.

Now set

$$H := H^{wz} = \langle p, H_w \cap H_z \rangle.$$

The right side contains q_z and contains

$$\langle p, H_w \cap H_x \cap H_z \rangle = \langle p, H_x \cap H_y \cap H_z \rangle = H_1 \cap H_2.$$

Thus we have produced the desired H and the proof is complete. □

Theorem 7.9.6 *The Veldkamp space $\mathcal{V}_p$ on the set $\mathcal{H}_p$ is a projective space.*

Proof Since C has rank at least three, Veldkamp lines exist, and so the full Veldkamp space $\mathcal{V}(C)$ is a linear space. Because Teirlinck's condition holds for $\mathcal{H}_p$ (Lemma 7.9.5), any line of $\mathcal{V}_p$ is a full line of $\mathcal{V}(C)$, and so the former is a subspace of the latter. Thus $\mathcal{V}_p$ is a linear space with thick lines. We need only show that it satisfies the Veblen axiom.

Suppose A, B, G and A, D, E are two triples of elements of $\mathcal{H}_p$ which respectively belong to two distinct lines of $\mathcal{V}_p$. Thus $B \cap D \cap E = B \cap G \cap E$ is a geometric hyperplane of $B \cap D$ and of $G \cap E$. Let q be a point $B \cap D - (B \cap D \cap E)$. By Lemma 7.9.5 there is a hyperplane F in $\mathcal{H}_p$ containing $G \cap E$ and point q. If $B \cap D$ is generated by $B \cap D \cap E$ and q, then F contains $B \cap D$ and F is the desired intersection of Veldkamp lines in the Veblen configuration. So we may assume that $B \cap B \cap E$ is not maximal in $B \cap D$. Then by Lemma 7.5.1 on p. 182, we must assume that $B \cap D$ has rank less than 2, so the rank of $p^\perp$ is at most three (which means that S is rank exactly 4). Then neither B nor D can have rank three, and so both of them are elements of $\mathcal{H}_p$ of the form H^{xy} – that is, they both contain p.

By the exact same argument, with G replacing B and E replacing D, we can similarly conclude that both G and E contain point p. Thus we see that the Veblen axiom holds unless the entire Veblen configuration lies in the subspace of $\mathcal{V}_p$ consisting of all elements of $\mathcal{H}_p$ which contain the point p.

By Theorem 6.9.1 at the end of Chap. 6, $\mathcal{V}_p$ is a projective space. □

7.9.3 The Classification of Non-degenerate Polar Spaces of Rank at Least Four

This section is the focus of a large number of previously prepared theorems, and so will be rather short.

Suppose S is a non-degenerate polar space of rank at least four. Choose a point p, and form the cone $C = p^\perp$ and let $\mathcal{H}_p$, $\mathcal{V}_p$ be defined as in the previous section. Suppose x and y are two distinct points of C. Then as S is non-degenerate, one must have $x^\perp \neq y^\perp$ as geometric hyperplanes of S. So there exists a point r in S such that $r^\perp$ contains x but not y.

If $r \notin C$ this means $H_r \in \mathcal{H}_p$ contains x but not y. If, on the other hand r cannot be chosen this way, we must have

$$x^\perp - y^\perp \subseteq p^\perp.$$

In that case $x \neq p$, so x is not deep in $x^\perp$, that is, some thick line on x does not lie in $p^\perp$. Since y must be collinear with at least two points of such a line, it follows that y is collinear with x. Now the point-residual $\mathrm{Res}_S(x)$ is a non-degenerate polar space of rank at least three which is the union of two geometry hyperplanes, consisting of the lines on x lying in $(px)^\perp$ and the lines on x lying in $(xy)^\perp$, respectively. That is impossible by Corollary 7.5.4.

Thus we have shown that given x and y in C there is a member of $\mathcal{H}_p$ containing x but not containing y.

Thus the collection $\mathcal{H}_p$ is a set of hyperplanes of C which satisfies these three properties:

1. $\mathcal{H}_p$ satisfies condition Teir (Lemma 7.9.5).
2. Veldkamp lines exist for $\mathcal{H}_p$ (this comes back to C being rank at least three).
3. $\mathcal{H}_p$ seperates points (as just argued above).

Thus by Corollary 4.5.3 of Chap. 4, p. 98, there is a projective embedding

$$C \to U \leq \mathcal{V}_p^*$$

of C into a subspace U of the second (or dual) Veldkamp space $\mathcal{V}_p^*$ spanned by the set $\{V(c)|c \in C\}$. (Here, as in Corollary 4.5.3, the symbol $V(c)$ denotes all hyperplanes of $\mathcal{H}_p$ which contain point c.)

By the hard-earned results of the previous section, C is embedded as the isotropic points of a symplectic form f in characteristic not 2, or the singular points of a σ-quadratic form q with associated σ-Hermitian form h on some vector space V so that the radical of C is embedded as a one-dimensional space comprising the intersection of the embedded C with the radical of f or q. That is, we have a vector space V, equipped with a σ-Hermitian or σ-quadratic form f or q, such that

$$V = R \perp N$$

where "$\perp$" is taken with respect to f or h, and R is the one-dimensional radical (with respect to either f or q) and N is non-degenerate, and C is exactly the classical polar space S_f or S_q of isotropic or singular 1-spaces with respect to these forms.

It is then possible to enlarge this vector space and extend the forms to forms f' and q' on

$$W = V \oplus \langle s \rangle = N \perp \langle r, s \rangle,$$

where the forms f' or q' are non-degenerate. Let S' be the non-degenerate classical polar space defined by the forms f' and q' on the space W.

Then we have the following situation. On the one hand, $C = p^\perp \cap S$ for the polar space S. On the other hand, $C = p^\perp \cap S'$, for the classical non-degenerate polar space $S' = S_{f'}$ or $S_{q'}$. By the critical Theorem of Tits (Theorem 7.5.13), we see that S and S' are isomorphic as polar spaces. Thus the following applies.

Theorem 7.9.7 *Any non-degenerate polar space of rank at least four is a classical polar space.*

We probably should not depart from this chapter without reminding the reader of another of its results.

Theorem 7.9.8 *An embeddable non-degenerate polar space of rank at least three is classical.*

Proof This is just the non-degenerate version of Corollary 7.8.5 of Sect. 7.5, (basically contained in the "main theorem" in Sect. 7.5.3). □

7.10 Epilogue

7.10.1 Non-embeddable Rank Three Polar Spaces

Up to this point we have achieved a classification theorem giving unique algebraic descriptions of embeddable rank three or rank four (embeddable or not) non-degenerate polar spaces. Aside from the apparently non-classifiable generalized quadrangles, the only polar spaces which remain unanalyzed are the *non-embeddable rank three polar spaces*. Very strange examples occur here. For example, certain analogs of compact forms of Lie Algebras exist over all sorts of fields F (algebraic number fields, for example). There are compact forms of building geometries of type E_7 which have only three varieties of objects forming a polar space. Yet Jacques Tits classified all these.

At this stage we are unable to provide the student with a completely self-contained account of his classification of non-embeddable rank three polar spaces for several reasons. (1) We should need a lengthy classification of the Cayley-division algebras and a complete proof (using coordinates) that any Moufang plane is coordinatized by such a gadget. The former was achieved by the Bruck-Kleinfeld Theorem [16], and a presentation of the latter is given in the book of Pickert [100, Sect. 7.1]. (2) The classification depends upon the fact that for chamber systems of spherical buildings an isomorphism between the union of all residues of rank two containing a chamber in one building, and a like subgraph in another, uniquely extends to an isomorphism between the buildings. Clearly, uniqueness is the key to classifying anything. But we are at a disadvantage for we have not yet introduced chamber systems or even defined buildings. (3) Finally, the classification depends critically on a structure called a "conformal space."

Nonetheless, if the reader is willing to accept the results in (1), on faith, and can be persuaded to believe that enough local isomorphisms symphonically agreeing on overlapping subdomains of a tight structure like a polar space extend to an isomorphism (a substitute for (2)), it is possible at least to give a student a "feel" about how Tits achieved this classification. The proof will depend upon a subtle virtually invisible substructure, the conformal space, which we will casually peruse below.

Moufang Planes and Cayley Division Algebras

The starting point is that the planes of any non-degenerate polar space of rank three are Moufang: that means that for any point-line flag (p, L) of a plane π, and any pair of distinct points a, b of $\pi - L$ whose line meets L at c, there exists an elation e of π, fixing L pointwise, and taking a to b. The line L is called the *axis* of e, and the point c is called the *center of e* for very good reason: it follows that every line through c is stabilized by e. We went through this in Chap. 5.

The Moufang condition forces the coordinatizing ring to be a particular algebra K over a field k – called a *Cayley-Dickson division algebra* – enjoying a number of identities. The Bruck-Kleinfeld theorem asserts that a Cayley-Dickson division algebra which is not a division ring is a certain eight-dimensional non-associative

division algebra called a *Cayley-Dickson algebra* defined by structure constants – an analog of the quarternion and Cayley algebras only over a general field.

Although there are one or two free parameters describing the multiplication table of a Cayley-Dickson algebra, it is always true that they must possess a "norm" function $n : K \to k$ such that $n(ab) = n(a)n(b)$ and n is a quadratic form with respect to the center. There is a great history of mathematics classifying such norms over specific fields, but they seemed to have assumed a new generic role in classifying buildings.

Now, following Pickert [100], any Moufang plane is coordinatized by K so that the set of points is given by

$$\mathcal{P} = K \times K \cup K \cup \infty.$$

The lines are of three types:

$$\begin{aligned} L_{a,b} &= \{(x, y) \in K \times K | x + by + a = 0\} \cup \{-b\}, a, b \in K, \\ L_a &= \{(x, -a) | x \in K\} \cup \infty, \\ L_\infty &= K \cup \{\infty\}. \end{aligned}$$

This plane is Moufang and has a flag-transitive group that is doubly transitive on points.

Conformal Structures

The analysis begins with a study of so-called *conformal structures*. Basically a conformal structure is an embedded rank one polar space defined by a certain form.

Let k be a field. Let K be a k-space of dimension d at least 2, and let $n : K \to k$ be an anisotropic quadratic form over K. Thus, for $x \in K$, $n(x) \neq 0$ unless $x = 0$.

Next we form the $d + 2$-dimensional k-space $V = K \oplus k^{(2)}$ and extend the quadratic form n to $Q : V \to k$ defined by

$$Q(x_0, x_1, x_2) = n(x_0) - x_1 x_2.$$

We denote the projective space $\mathbf{P}(V)$ associated with V by the sympbol P. Then associated with the quadratic form Q is a polarity $\mu : P \to P^*$. (Recall that a polarity is a quasi-polarity whose iteration is the identity mapping on the set $S(P)$ of subspaces of P.)

With some abuse of notation, we also let Q denote the projective variety in P defined to be the projective points which are singular with respect to the quadratic form Q. This variety is called a "quadric."

Now the subspaces of P intersect the quadric Q at all sorts of sets. The points of $P - Q$, the exterior points, for example, meet Q in the empty set. For each point s on the quadric Q, the polar of it, $s^\perp := \mu(p)$ is a hyperplane of P intersecting Q only at the point s and no further points.

We let Γ denote all the non-empty intersections of subspaces of P with Q. This produces a sort of geometry (Q, Γ) which we call a "conformal space".[22] Thus it has all the points of Q (defined by intersecting Q with either the subspaces s or $s^{\perp}$, $s \in Q$), pairs of points of Q (intersections with certain lines), and larger sets. A *circle*, for example, is the intersection of Q with a plane spanned by three points of the quadric. Similarly, a *hypersphere* is the intersection of the quadric with a hyperplane which is spanned by points of the quadric.

It is important to realize that the conformal space is not just the incidence structure Γ alone. It is given that Q is embedded in some projective space and that the "objects" of Γ are intersections of Q with subspaces of the ambient projective space that are spanned by points of the quadric. Then it is not difficult to see that the ambient space and quadric are essentially unique – that is, they can be "recovered" (Lemma 8.18 of [138]).

The reader may be more familiar with this example. If $d = 2$, then $\dim(V) = 4$, and Q is the familiar elliptic quadric. The conformal space consists of points, pairs of points, and ovals of planes. (The latter are both circles and hyperspheres in this example). Thus we obtain a rank three geometry called a (Miquelian) *inversive plane*.[23]

7.10.2 Tits' Classification Theorem

Theorem 7.10.1 *Each non-degenerate non-embeddable polar space of rank three is canonically describable in terms of a unique Cayely division algebra.*

Moreover, there exists one such polar space for each Cayley division algebra. In other words, these polar spaces are parameterized by the Cayley division algebras.

The argument starts by first noting (as in Corollary 7.5.9, p. 187) that each plane is Moufang, and so has the form $\Pi(K)$ as given above. We can assume that K is not a field (otherwise the polar space is embeddable), so K is a Cayley Dickson algebra over the field k and so comes equipped with the anisotropic norm form $n : K \to k$. Extending n to the larger form Q as above we get a conformal space $\Gamma(n, K)$ as described in the preceding section – only this time it has extra nice properties (conditions (i)–(iii) of Proposition 9.6, p. 184 of [138]), since K is now a particular algebra.

In the plane $\Pi(K)$ this conformal space can be attached to a line and (since the elations preserve this structure) it lives on any line of the plane $\Pi(L)$. (Actually this structure can be imposed on a line in two ways, so there is a sort of "orientation" involved.) Now, take a point p of the plane not on line L. Then there is a bijection between the lines on p and the points on line L, and this lifts the conformal structure on L to one on the lines on p, a "line" of the dual plane. (It is actually a conformal

[22] Of course, to make it a real geometry, in the sense of this book, we would have to figure out just how to define the "type" function on it.

[23] Sometimes called a Möbius plane.

structure with the opposite orientation. This orientation has to do with the involution (antiautomorphism of order two) of the Cayley-Dickson algebra K.)

Now $\Pi(K)$ was just some plane of the rank three polar space S lying on some point p. So, in the cone $p^\perp$ in S, the plane Π looks like a "line" of the residual generalized quadrangle T^* defined by the lines and planes of S on p. But now, as we have seen, a "line" of this triangle has inherited a special conformal structure.

Now comes a key result. If you have two generalized quadrangles T and T' each of whose lines carry *some* conformal structure with the "extra nice property" alluded to, and if there is a conformal isomorphism from a particular line of one quadrangle to a line of the other, then this isomorphism extends to an isomorphism $T \to T'$ of the quadrangles. (What a result! If such a quadrangle has a line-transitive automorphism group, the conformal structures are the same on each line.)

Now suppose S and S' are two non-degenerate rank three polar spaces each of which happens to possess a plane isomorphic to $\Pi(K)$ where K is a Cayley division algebra. Then all their planes are isomorphic. But now from the development of the previous paragraph, all their residual quadrangles are isomorphic. Finally, since the residue of a line is a product geometry (actually just a big bipartite graph), one can utilize appropriate restrictions of the preceding residue-isomorphisms to obtain isomorphisms of the line-residues.

But in fact, because of the way the isomorphisms were derived from one another, one can pick a flag $F = (p, L, \pi)$ in S (a point, line, and plane, respectively) and a flag $F' = (p', L', \pi')$ in S' and patch together the residue isomorphisms to get an incidence preserving bijection $\phi : S_2(F) \to S'_2(F')$ where $S_2(F)$ (or $S'_2(F')$ is the set of points lines or planes of S (or S') which are incident with at least one object of the flag F (or F').

The point is that this mapping ϕ extends to an isomorphism of the rank-three geometries involved – in particular to an isomorphism $S \to S'$ of the polar spaces. So one gets at most one polar space for each Cayley division algebra K.

There are hyperbolic lines here.

A Remark: Muehlherr's Theorem

Before leaving the subject of non-embeddable rank-three polar spaces one might note that they crop up in a peculiar way. About 15 years ago Professor Bernhard Muehlherr [89] proved the following remarkable result.

Theorem 7.10.2 *Suppose Γ is a non-embeddable rank three polar space. Then there exists a field F and an $E_{7,7}$ Lie incidence geometry admitting a type-preserving involution t fixing only some points, lines, and symplecta (the lines pointwise) such that the fixed-point geometry is the non-embeddable rank three polar space Γ.*

The parapolar polar space $E_{7,7}$ mentioned in this theorem makes its debut in Chap. 15, of Part IV of this book.[24]

[24] It is not clear that Muehlherr's discovery could be exploited to give an alternative proof of Theorem 7.10.1. As it stands, it actually exploits it.

7.10.3 The Buekenhout-Johnson-Shult Theorem

The last part of the argument of the previous section involved the extension of an incidence-preserving mapping

$$\phi : S_2(F) \to S_2'(F')$$

to an isomorphism $S \to S'$ (viewed as full rank three geometries because of the inclusion of the planes). Now the reader should notice that the cones over p and p' are already made isomorphic by the inclusion. If the rank had been rank four or more we could have concluded from Tits' theorem about cones (Theorem 7.5.13) that S and S' are isomorphic as polar spaces (and so isomorphic as enriched geometries since all singular subspaces are uniquely determined by their point-shadows).

So you see, something extra has to be done in the case of polar rank three.

The next theorem can be seen this way. By Tits' theorem (Theorem 7.5.13, p. 191) isomorphism of cones implies isomorphism of polar spaces if we are speaking of cones of polar spaces of sufficient rank on both the domain and range. What would happen if we did not know in advance that the ambient spaces S and S' of the domain and range were polar spaces at all; that we only knew that the "cones" involved were cones of a non-degenerate polar space of rank at least three and that for the ambient space S' of the range, this was true no matter what cone $x'^{\perp}$ one looked at. Then S', the ambient space of the range, is also a polar space.

This theorem is elementary when compared to the classification of rank three polar spaces just discussed.[25] Indeed it does not even depend upon the actual classification of polar spaces of rank at least four given in this chapter. Rather, it only uses the basic properties of polar spaces. It is inserted here only because it will be required in our discussion of symplecta in Chap. 13.

Theorem 7.10.3 (Buekenhout [19], Johnson and Shult [81].) *Assume that Γ is a (connected) projective gamma space with this property:*

(LP) For each point p the residue $Res(p) = (\mathcal{L}_p, \Pi_p)$ of all lines and planes incident with point p, is a non-degenerate polar space of rank at least two (but otherwise of arbitrary rank).

Then $(\mathcal{P}, \mathcal{L})$ is itself a non-degenerate polar space.

Proof Recall that a projective gamma space is one whose singular subspaces are projective spaces. Note that since the polar rank of the residual is at least two, every pair of distinct collinear points lie together in a projective plane which is a subspace of Γ. It follows that Γ is a partial linear space.

Fix an arbitrary point p. Since we are in a gamma space, $p^{\perp}$ is a subspace of Γ. We claim that Γ is itself a polar space. Suppose (x, L) is a non-incident point-line pair in $p^{\perp}$. We must show that $x^{\perp} \cap L$ is non-empty. This is true if L contains p

[25] I am actually embarrassed to place them in the same section.

or if $x = p$. So we assume $x \neq p$ and p is not incident with L. Then $A = \langle p, L\rangle$ is a plane on p, and there is a line L on $\{x, p\}$, representing a non-incident "line" and "point" of the point-residual $(\mathcal{L}_p, \Pi_p)$. Since the latter is a polar space, there is a line N on point p such that N is a line of A, and N shares a plane B with line L. Since A is a projective plane, the lines L and N meet at a common point q. Since x is collinear with every remaining point of B, it is collinear with q. Thus $x^\perp \cap L$ is non-empty.

From now on, we will repeatedly use the fact that, for every point $r \in \mathcal{P}$, $r^\perp$ is always a polar space of polar rank at least two.

To show that the gamma space Γ is itself a polar space, we need only show that $x^\perp \cap L$ is non-empty for every point-line pair (x, L). By way of contradiction, assume this is false, so that there is a non-empty collection $\mathcal{A}$ of pairs (x, L) for which $x^\perp \cap L$ is empty. Let $d(x, L)$ be the minimum distance (as measured in the point-collinearity graph Δ of Γ) of x from a point of L.[26] Then among all pairs $(x, L) \in \mathcal{A}$, we choose (x, L) so that $d(x, L)$ assumes the minimum possible value d. By definition of $\mathcal{A}$, one has $d \geq 2$. Now choose y on L so that $d(x, y) = d$, and consider a geodesic path in the graph Δ $(x, u_1, \ldots, u_d = y)$. By minimality of d, $(u_1, L) \notin \mathcal{A}$, so ${u_1}^\perp \cap L$ is non-empty, and so contains y. Thus $d = 2$ and our geodesic is (x, u_1, y).

Now, since the polar rank of ${u_1}^\perp$ is at least two, there exists a plane A_1 on the line yu_1, and $x^\perp \cap A_1$ contains a line M on u_1. Choose any point $z \in L - \{y\}$. Since $y^\perp$ is a polar space, there is a line N contained in $z^\perp \cap A_1$ and containing point y. Since A_1 is a projective plane, the lines N and M must intersect at a point v. Now the polar space $v^\perp$ contains the line vx and the plane $A_2 := \langle v, L\rangle$, so $x^\perp \cap A_2$ contains a line K. Then K and L, being two lines of a common projective plane, must intersect at a point w. Now $w \in x^\perp \cap L$ against our choice of $(x, L) \in \mathcal{A}$. Thus $\mathcal{A}$ is empty, and we see that Γ is a polar space. □

Remark The non-degeneracy only appears for convenience in stating the result. The version in [81] allows radicals.

7.11 Exercises for Chapter 7

7.11.1 Exercises Involving Only Basic Properties

7.1 Recall that if X is a set of points in a point-line geometry $\Gamma = (\mathcal{P}, \mathcal{L})$, the symbol $X^\perp$ is the set of points collinear to or equal with each point of X – that is, $X^\perp = \cap_{x \in X} x^\perp$ in our notation. In any gamma space $X^\perp$ is always a subspace. Show the following:

(1) If A and B are sets of points, $A \subseteq B^\perp$ if and only if $B \subseteq A^\perp$.
(2) In a gamma space, if $X \subseteq A^\perp$ and $X \subseteq B^\perp$, then $X \subseteq \langle A, B\rangle_\Gamma{}^\perp$.

[26] Recall that by hypothesis this graph is connected.

7.2 Let $\mathcal{M}$ be the collection of all maximal singular subspaces of a polar space $\Gamma = (\mathcal{P}, \mathcal{L})$:

(1) Show that every maximal singular subspace contains the radical $\mathrm{Rad}(\Gamma)$.
(2) Show that $\mathrm{Rad}(\Gamma) = \cap\{M \in \mathcal{M}\}$.

7.3 Show that if M is a maximal singular subspace of a polar space $\Gamma = (\mathcal{P}, \mathcal{L})$, and x is a point in $\mathcal{P} - M$, then

(1) $x^\perp \cap M$ is a geometric hyperplane of M.
(2) $\langle x, x^\perp \cap M\rangle_\Gamma$ is also a maximal singular subspace of Γ.

7.4 Suppose $M_1 \cap M_2 = A \neq \emptyset$, for maximal singular subspaces $M_i \in \mathcal{M}$. Suppose x is a point not in $A^\perp$ – that is, x is not collinear with all points of A. Form the subspace $M_0 := \langle x, x^\perp \cap M_1\rangle$. Show that $M_0 \cap M_2 \subset A$.
[Hint: Any point of this intersection lies in $A^\perp$ and $x^\perp \cap M^\perp$. Use Exercise 7.1.]

7.5 Suppose A is a singular subspace of a non-degenerate polar space Γ. Then of course the subspace $A^\perp$ is a polar space in its own right. Show that $A = \mathrm{Rad}\,(A^\perp)$. Conclude that $A^\perp/\mathrm{Rad}\,(A^\perp)$ is non-degenerate.

7.6 Suppose a and b are two non-collinear points of a polar space S. Show that

$$\mathrm{Rad}\,(a^\perp \cap b^\perp) \subseteq \mathrm{Rad}\,(S).$$

7.7 Suppose A and B are maximal singular subspaces of a non-degenerate polar space such that $A \cap B$ has finite projective rank. Show that there exists a maximal singular subspace C such that $C \cap A = \emptyset$ – that is, A and C are "opposite."
[Hint: Use Exercise 7.4 in an induction argument.]

7.8 Let S be a non-degenerate polar space. We say that two singular subspaces A and B are "opposite" if $A^\perp \cap B = B^\perp \cap A = \emptyset$.

(1) Must opposite singular spaces have the same rank?
(2) Show that if A is any singular subspace of Γ of finite rank, then there exists a subspace B which is opposite A.
[Hint: Devise an induction proof.]

7.9 Prove the assertions about the form h requested in Example 1, p. 178.

7.10 Suppose $P = (\mathcal{P}, \mathcal{L})$ is a non-degenerate polar space of polar rank at least three. Suppose $\{a, b\}$ is a pair of non-collinear points. Show that the subspace $a^\perp \cap b^\perp$ (together with its internal lines) is also a non-degenerate polar space of rank one less than that of P. [Slight hint: Note that your proof of non-degeneracy cannot employ any assumptions of finite polar rank. Induction is not the way to go here.]

7.11 Give an induction proof showing that, for every integer r larger than two, Veldkamp $(r-1)$-spaces exist for any polar space of rank at least r. (This was part 3 of Lemma 7.5.1.)

7.11.2 Exercises on Classical Polar Spaces

7.12 Is any (every) alternating form trace-valued?

7.13 What do you conclude about the admissible pair (σ, ϵ) if $D^{\sigma,\epsilon} = D$? Answer the same question for $T^{\sigma,\epsilon} = D$, $D_{\sigma,\epsilon} = 0$ and $T_{\sigma,\epsilon} = 0$).

7.14 Let D be the field of rational functions $F(x, y)$ in two indeterminates x and y. Let σ be the unique field automorphism which exchanges x and y. Set $\epsilon = 1$, so that σ, ϵ is an admissible pair. Show that $T_{\sigma,\epsilon} = T^{\sigma,\epsilon}$.

7.15 Suppose f is a (σ, ϵ)-Hermitian form. Define $f\lambda : V \times V \to D$ by $(f\lambda)(u, v) := f(u, v)\lambda$. Show that $f\lambda$ is a (σ', ϵ')-Hermitian form for an appropriate choice of σ' and ϵ'. With the formulae for σ' and ϵ' in hand, verify from these first principles the admissibility relations.

7.16 Given, $\lambda, \sigma, \epsilon$ as in the previous problem, show that $D_{\sigma,\epsilon}\lambda = D_{\sigma',\epsilon'}$.

7.17 Suppose V is a vector space with countable basis $B = \{e_i | i \in \mathbf{N}\}$. Write $V = \langle e_0 \rangle \oplus W$ where W is the space spanned by all e_i, i a positive integer. We define an alternating form f by setting

$$f(e_i, e_j) = 1 \text{ if } i \text{ is odd and } j = i + 1,$$
$$f(e_i, e_j) = 0 \text{ for } 1 \leq i < j \text{ otherwise,}$$
$$f(e_0, e_i) = 1, \text{ for all positive } i.$$

The rest of the values of f on $B \times B$ are determined by the relations $f(e_i, e_i) = 0$ and $f(e_i, e_j) = -f(e_j, e_i)$ required of an alternating form. Using the symbol "$\perp$" to indicate the quasi-polarity determined by f, show that $W^{\perp} = 0$. (In the case of a non-degenerate alternating form on a finite-dimensional vector space, the "perp" of a proper subspace is never zero. The lesson of this example is that for a quasi-polarity π on the subspaces of V, and a subset X, $\pi(X) = \{0\}$ does not generally imply that X spans V.)

For the next few problems suppose q is a non-zero (σ, ϵ)-quadratic form with associated (σ, ϵ)-form f on a right vector space V over D. Or, f might simply be a form that is not derived from any (σ, ϵ)-quadratic form at all, as in the case of an alternating form in characteristic not 2. In that case the assertions regarding f are still to be proved – those involving the mythical q can be ignored. So don't use the existence of q to prove anything about f unless possibly when q is hypothesized in the problem.

7.18 Show that if $V = D \times D$, that the function q defined by

$$q(x) = x_1^\sigma x_2 + D_{\sigma,\epsilon}, \quad x = (x_1, x_2) \in V,$$

is a (σ, ϵ)-quadratic form whose associated (σ, ϵ)-Hermitian form is defined by

$$f(x, y) = x_1^\sigma y_2 + x_2^\sigma \epsilon y_1.$$

(The student might well ask where the ϵ in the second formula came from since it didn't appear to be in the formula for q? Well, it was there: it was in the definition of $D_{\sigma,\epsilon}$.)

7.19 Continuing the last problem, show that a vector $(a, 1)$ is f-isotropic if and only if $a \in D^{\sigma,\epsilon}$. Show that it is singular with respect to q if and only if $\alpha \in D_{\sigma,\epsilon}$. (This may begin to give one some idea of the roles of these two special subgroups of $(D, +)$.)

7.20 (Transportability of forms.) Suppose $t : V \to W$ is a bijective linear transformation. Define functions $ft : W \times W \to W$ and $qt : W \to W$ by the rules

$$ft(w_1, w_2) := f(t^{-1}(w_1), t^{-1}(w_2)),$$
$$qt(w) = q(t^{-1}(w))$$

for all $w, w_1, w_2 \in W$. Show that qt is a (σ, ϵ)-quadratic form whose associated (σ, ϵ)-Hermitian form is ft.

7.21 (Transvections.) Suppose z is an isotropic vector of V with respect to f. Choose any $\gamma \in D$. Define $t(z, \gamma) : V \to V$ to be the mapping defined by the transformations

$$v \to v - z \cdot \gamma \cdot f(z, x). \tag{7.29}$$

(1) Show that $t(z, \gamma)$ is a linear transformation which fixes $z^\perp$ vectorwise.
(2) Show that if f is not the zero form, and z is isotropic, then $t = t(z, \gamma)$ is an *f-isometry* – that is, for all $y \in V$, $f(x, y) = f(t(x), t(y))$ – if and only if $\gamma^{-1} \in D^{\sigma,\epsilon}$. (In this case t is called an *f-transvection centered at* $\langle z \rangle$.)
(3) We say that a linear transformation $\rho : V \to V$ is a *q-isometry* if and only if $q(\rho(v)) = q(v)$ for all vectors v. Show the following:
   (a) Any q-isometry is an f-isometry.
   (b) If $q(z) = 0$ then $t(z, \gamma)$ is a q-isometry if and only $\gamma^{-1} \in D_{\sigma,\epsilon}$. (In this case, $t(z, \gamma)$ is called a *q-transvection*.)

7.22 Suppose $V = D \oplus D$ with q and f as given in Exercise 7.7. Let z be a vector in V. Let $T(z) = \{t(z, \gamma) | \gamma \in D\}$, the set of all transvections with center $\langle z \rangle$. Show the following:

1. If z is isotropic, $t(z, \gamma_1) \circ t(z, \gamma_2) = t(z, \gamma_1 + \gamma_2)$ and so $T(z)$ is an abelian group isomorphic to $(D, +)$. Show that it induces an action on the projective line $P(V)$ which acts faithfully and regularly on $P(V) - \langle z \rangle$. Conclude that the group $G = SL(V)$ generated by the set $\{T(z) | z \in V, f(z, z) = 0\}$ is doubly transitive on $P(V)$.
2. Suppose z is totally isotropic with respect to f. Let $T_f(z)$ be the set of f-isometries in $T(z)$. Then $T_f(z)$ is a group isomorphic to the additive group $(D^{\sigma, \epsilon^{-1}}, +)$ and this group regularly permutes the elements of $S_f - \langle z \rangle$. Thus $G_f = \langle T_f(z) \rangle | f(z, z) = 0 \rangle$ is doubly transitive on S_f.
3. Suppose z is a totally singular vector, that is, $q(z) = 0$, where q is the given (σ, ϵ)-quadratic form. The subset of the projective line consisting of the totally singular 1-spaces, is called the *quadric defined by q* and we will denote it by Q. Let $T_q(z)$ be the members of $T(z)$ which are q-isometries. Show that $T_q(z)$ is isomorphic to the additive group $D_{\sigma, \epsilon^{-1}}$ and acts semiregularly and faithfully on the set $Q - \langle z \rangle$. Conclude that $G_q = \langle T_q(z) | q(z) = 0 \rangle$ acts doubly transitively on the quadric Q.

7.23 A phenomenon which comes up in (unkown) polar spaces X embedded in a classical polar space $S = S_f$ or S_q is that if the subspace X has rank at least two (as rank is defined for a polar space), then X is *f- or q-transvection closed*. That means that if $x \in S - X$ and p and r are two points of $X - x^{\perp}$ (the perp taken with respect to f), which are moved, one to the other, by an element $t \in T_f(x)$ $(T_q(x), x \in Q)$, then t stabilizes S. This problem concerns an f-transvection(q-transvection)-closed subset X of the S_f (S_q) for the projective line $\mathbf{P}(D \oplus D)$.

Show that if D has characteristic 2, every such transvection merely transposes p and r (conceivably p and r could be equal, but if you have even been reading the previous exercise you would know this can't happen for a non-identity transvection). Conclude that in this characteristic and dimension, every 2-subset of S is a transvection-closed subset of S on this non-degenerate projective line. Produce a simple example over $GF(2^2)$.

7.24 (f-Homologies where f is a Hermitian form.) Suppose f is a (σ, ϵ)-Hermitian form $V \times V \to D$ which is non-degenerate.

1. Show that if $\beta = f(x, x)$, then $\beta \in D^{\sigma, \epsilon}$.
2. Assume $\beta = f(x, x)$ is non-zero. Show that for any scalar $\alpha \in D$, $\alpha^{\sigma} \beta \alpha \in D^{\sigma, \epsilon}$. Must it be true that $D^{\sigma, \epsilon} = \{\alpha^{\sigma} \beta \alpha | \alpha \in D\}$?
3. Set $W = x^{\perp} := \{w | f(x, w) = 0\}$. Show that W is a subspace of V. Must it be true that $V = \langle x \rangle \oplus W$? [Hint: This has nothing to do with the

number of cosets of $(D^{\sigma,\epsilon}, +)$ in $(D, +)$. Choose any vector v, suppose $f(x, v) = \gamma$, and choose $\alpha = \beta^{-1}\gamma$; show that $f(x, x\alpha - v) = 0$.]

4. Show that if $f(x, x) = \beta \neq 0$, then there exists an f-isometry $t : V \to V$ such that $t(\langle x \rangle) = \langle x \rangle$ and t fixes every vector of $W = x^{\perp}$. Write out the formula for $t(u)$, where $u = x\alpha + w$, $w \in W$, entirely in terms of $f(x, x) = \beta$, and $f(x, u)$.

(An f-isometry of this last type, fixing a *non-isotropic* 1-space and fixing its perpendicular complement is called an *f-homology with center* $\langle x \rangle$. The center of such a homology is not necessarily unique.)

7.25 (q-Homologies where q is a σ-quadratic form.) Consider a σ-quadratic form ($\epsilon = 1$), with associated σ-Hermitian form f. If $q(x)$ and $f(x, x)$ are both non-zero, decide the analog of the previous problem: if $W = x^{\perp}$ so $V = \langle x \rangle \perp W$, does there exist a q-isometry fixing $\langle x \rangle$ as well as every 1-subspace of W? If so, give a formula for the generic q-isometry with this property, and identify the group of all such transformations (called *q-homologies with center* $\langle x \rangle$) with a subgroup of D^*, the multiplicative group of non-zero elements of D. (A q-homology or q-transvection with center $\langle x \rangle$ is called a *q-perspectivity with center* $\langle x \rangle$. Similarly, an f-homology or f-transvection with center $\langle x \rangle$ is called an *f-perspectivity*.)

7.11.3 Exercises from Section 7.9

7.26 Prove Lemma 7.9.1.

7.27 Suppose S_0 is a non-degenerate polar space of rank at least three and let H be a geometric hyperplane of S_0. Fix a point r in $S_0 - H$, and suppose s is a point of $S_0 - r^{\perp}$ such that $r^{\perp} \cap s^{\perp} \subseteq H$.

1. Show that if $s \in H$, then $H = x^{\perp}$. [Show first that s must be a deep point of H.]
2. Show that if the point collinearity graph induced on $S_0 - H$ has diameter three, then s must lie in H so the conclusion of the previous part of this exercise applies.

7.28 Suppose S_0 is the classical polar space of rank three or more defined by a non-degenerate quadratic form q on a vector space V over a field of characteristic not two.

1. Show that any non-degenerate two-dimensional subspace of V contains either zero or exactly two singular 1-spaces.
2. If U is a non-degenerate 2-subspace of V, show that the restriction of the quadratic form to

$$U^{\perp} := \{v \in V \; b(u, v) = 0, \text{ for all } u \in U\}$$

(where b is the symmetric bilinear form associated with q) defines a polar space of rank at least two.

3. Use the preceding finding to deduce that the collinearity graph induced on $S_0 - p^{\perp}$ has diameter two.

7.12 Appendix to Chapter 7: Known Quadrangles in 2001

7.12.1 The Older Examples

We have already met the generalized quadrangle of order $(2, 2)$ several times in previous chapters. Its point-collinearity graph is $\bar{T}_6$ and its automorphism group is Sym(6). We now offer more examples of non-degenerate quadrangles.

Example 3 A generalized quadrangle of order $(2, 4)$. Let $V = \mathbf{GF}(4)^{(3)}$, the vector space of 3-tuples over the field of four elements. Let $\mathcal{P}$ be the non-zero vectors of even weight – that is 3-tuples (a_1, a_2, a_3) with exactly one of the coordinates equal to zero. There are 27 such vectors. A *line* is defined to be any triplet of vectors (u, v, w) chosen from $\mathcal{P}$ for which $w = u + v$. (Note that since the characteristic is 2, the equation is equivalent to $v = u + w$ or $u = w + v$. This is a symmetric ternary relation on the vectors, and for any such triplet (u, v, w) satisfying the relation, the vectors u, v, and w must be pairwise distinct.) There are 45 such lines. Each line has three points and each point is on five lines.

See if you can verify the axiom (GQ).

[There is a story in this example. Suppose we consider in our Euclidean space a polynomial function $f(x, y, z) = 0$. The solutions $(x, y, z) \in \mathbf{R}^{(3)}$ form a so-called *algebraic surface*. It is said to *possess no singular points* (algebraic geometers call them *double points*) if it does not cross itself.[27] Here is an example I have often encountered. If you wake up in the morning disoriented (as I often do), there is still the terrain of covers over your bed. Covers, as a rule, do not intersect themselves (although that possibility briefly crosses my mind when they are terribly entangled). Now, God forbid, you awake one morning to discover that the polynomial $f(x, y, z)$ describing your bedcovers is an irreducible cubic polynomial in $\mathbf{Q}[x, y, z]$. Then, believe it or not, this surface contains exactly 27 straight lines. Each line intersects just ten others, and among these ten, they are paired into five pairs according to whether they intersect each other or not. This produces a system of 27 lines and 45 triangles of lines, which has the incidence structure of the generalized quadrangle presented in this example.[28]]

---

[27] At such a point one could find derivatives pointed along two distinct tangent planes, which means the Jacobian is zero when evaluated there.

[28] Obviously the 27 intersection points lie in a compact ball that can be scaled to fit in my bedroom. As a student, I have conducted numerous experiments in the early morning, shrewdly adjusting

Example 4 (Classical examples.) Let f be a non-degenerate reflexive σ-sesquilinear form on a vector space V whose maximal isotropic subspaces are all two-dimensional. (All of these mysterious words are explained in Sect. 7.6.) Let the points and lines of $\Gamma = (\mathcal{P}, \mathcal{L})$ be the systems of totally isotropic one- and two-dimensional subspaces of V with the usual inclusion relation as incidence – that is, these are subspaces of these dimensions in which any two vectors are perpendicular to each other. (Obviously it is already not a Euclidean space or even a Lorentz space.) Then two distinct points p and q (these are 1-spaces) are collinear if and only if they are *perpendicular* – that is, $f(p, q) = 0$ – for in that case, the subspace $U = \langle p, q \rangle_V$ that they generate is totally isotropic – i.e.,

$$f(U, U) := \{f(u, v) | (u, v) \in U \times U\} = 0 := \{\text{the zero vector}\}$$

– and so U is a line of the quadrangle. Given any such line U and point x not in U, $x^\perp$ is a codimension-1 subspace which cannot contain U (otherwise $\langle p, U \rangle$ would be totally isotropic against the maximality of U) and so $p^\perp \cap U$ is an isotropic 1-space. Thus $(\mathcal{P}, \mathcal{L})$ is a generalized quadrangle. If the student desires a very specific example, one may consider the vector space $V = \mathbf{R}^{(n)}$, $n \geq 4$, where, for $u = (a_0, a_1, \ldots, a_{n-1})$ and $v = (b_0, b_1, \ldots, b_{n-1})$, f is defined by

$$f(u, v) = (\sum\nolimits_{i=0}^{n-2} a_i b_i) - (a_{n-1} b_{n-2} + a_{n-1} b_{n-1}).$$

Then $(\mathcal{P}, \mathcal{L})$ is an infinite generalized quadrangle.

There is another version of this sort where the form f is replaced by a generalization (for division rings) of a quadratic form in characteristic two.

Example 5 (Tits' constructions, $T_2(O)$, $T_3(O)$)

(Ovals in $PG(2, q)$.) Suppose $\Pi = (\mathcal{P}, \mathcal{L})$ is a projective plane of finite order $n \geq 2$. A 2-*arc* is a subset A of $\mathcal{P}$ subject to the property that every line of Π meets A in at most two points. A line meeting A in two points is called a *secant line*; a line meeting A at just one point is called a *tangent line* and a line that does not intersect A at all is called an *exterior line*. Such a 2-arc is called an *oval* if and only if each of its points lies on a unique tangent line.

A famous theorem of Segre says that if Π is a Desarguesian plane $PG(2, q)$ over a finite field $\mathbf{GF}(q)$ where q is odd, then any oval is the set on which an irreducible quadratic form vanishes – that is a *conic*. (Perhaps the most accessible proof of Segre's Theorem can be found on p. 139 of Peter Cameron's beautiful book *Combinatorics* [26].) This means that in odd characteristic, all ovals are essentially alike.

If n is even, many ovals exist which are not conics.

---

my bed covers while wondering if I should attend class, but I have never directly observed this phenomenon of the 27 lines.

(Ovoids of $PG(3,q)$.) An *ovoid of* $PG(3,q)$ is a set O of $1+q^2$ points of $PG(3,q)$, each point of which lives on a unique tangent plane – that is, a $\Pi \simeq PG(2,q)$ meeting O at a single point. A famous theorem of Barlotti [6] gives the analog of Segre's theorem for ovoids. If q is odd, then any ovoid is the variety defined by an elliptic quadric – the zeroes of a class of homogeneous polynomials in four variables, and so are all alike.

We paraphrase the construction as presented in [131] pp. 6 and 7. Let $d = 2$ (respectively, $d = 3$) and let O be an oval (respectively an ovoid) of $PG(d,q)$. Further let $PG(d,q) := H$ be embedded as a hyperplane in $PG(d+1,q) := P$. The points $\mathcal{P}$ are of three types:

1. The points of $P - H$
2. The hyperplanes X of P for which $|X \cap O| = 1$
3. The symbol ∞

The lines $\mathcal{L}$ are of two types:

(a) The lines of P which are not in H and meet O (necessarily at a unique point)
(b) The points of O

Incidence is defined as follows. A point of type 1 is incident only with the lines of type (a) (the natural incidence of P). A point of type 2 is incident with all lines of type (a) contained in it and the unique element of O it contains. The point (∞) is incident with all lines of type (b) and no line of type (a). The incidence structure $(\mathcal{P}, \mathcal{L})$ is a generalized quadrangle $T_2(O)$ of order (q,q), when $d = 2$, and is otherwise a generalized quadrangle $T_3(O)$ of order (q, q^2), when $d = 3$.

These quadrangles are classical if and only if O is a conic or an elliptic quadric when $d = 2$ or 3, respectively (see Sects. 3.2.2 and 3.2.4 of the book of Payne and Thas [98]). Thus when q is odd, the theorems of Segre and Barlotti prevent any strange ovals or ovoids from appearing.

But if q is even, the door is open for strange ovals and ovoids. Many ovals which are not conics are known (see [32]). As for the ovoids, one of the most famous is the *Suziki–Tits ovoids*. A little over 40 years ago, Michio Suzuki discovered a new infinite class of simple finite groups, which were defined by so-called "structure equations" (that is, generators and relations similar to those for the Zassenhaus groups). He then proved that these relations could be realized by matrices, indicating that these groups were special subgroups of the classical symplectic groups, $Sp(4,q)$, where q is an odd power of 2 at least 8 [124]. Tits [137] discovered that this group stabilized a new sort of ovoid – hence the Suzuki–Tits ovoid. This and the elliptic quadric are the only two ovoids known. Perhaps there are more (an avenue also being explored by Cherwitzo/O'Keefe and Pentilla). But it is difficult to be such an ovoid: for a beautiful theorem of Matt Brown [15] shows that if the ovoid has even one hyperplane section that is a conic, then the ovoid is an elliptic quadric. Nonetheless, at least one new infinite family of quadrangles results.

Example 6 (Tits' construction $T_2^*(H)$.) If the order n is even, it can be shown that all tangent lines of an oval A meet at a common point n called the *nucleus* of A.

This point can then be adjoined to A to form a set $H := A \cup \{n\}$ called a *hyperoval*. If H is a hyperoval, all lines of Π meet H in 0 or two points: all the tangents have disappeared. Since planes of even order carry plenty of mysterious ovals, they carry many mysterious hyperovals, and only recently has work of Cherwitzo, O'Keefe, Pentilla [31, 92, 93] and others begun to unfold what lies here.

Suppose H is now some hyperoval of a Desarguesian projective plane Π of order $q = 2^r \geq 4$. We can then embed Π as a hyperplane in a $PG(3, q)$. Now let $\mathcal{P}'$ be the points outside this hyperplane – that is, the points of $PG(3, q) - \Pi$. The lines $\mathcal{L}'$ are the lines L of $PG(3, q)$ not lying in Π which meet Π at a point of the hyperoval H. Believe it or not, $(\mathcal{P}', \mathcal{L}')$ is a generalized quadrangle of order $(q - 1, q + 1)$ called $T_2^*(H)$.[29]

Let us model a very specific example. The Desarguesian projective plane over $\mathbf{GF}(4)$ contains 168 hyperovals. Pick one hyperoval and let $L_1, \ldots, L_6$ be the 1-spaces of $V = \mathbf{GF}(4)^{(3)}$ representing its projective points in $PG(2, 4)$. Let $\mathcal{P} = V$, the vectors, and let $\mathcal{L}$ be the full collection of vector translates, $\{L_i + v | v \in V, i = 1, \ldots, 6\}$. Then $(\mathcal{P}, \mathcal{L})$ is a generalized quadrangle of order $(3, 5)$. It has 64 points and 96 lines.

That was the status for a number or years. Then Stanley Payne [96] made a simple observation that had been overlooked by everyone.

Example 7 (Payne's "contraction" construction.) Let $\Gamma = (\mathcal{P}, \mathcal{L})$ be a quadrangle of order (s, s). A point p is said to be *regular* if, whenever q is not collinear with p, the set of points of $\{p, q\}^{\perp\perp}$ is a coclique of $1 + q$ points – a so-called "hyperbolic line." (For example in the quadrangle $W(3, q)$ associated with the the classical group $Sp(4, q)$, every point is regular.) Now fix such a regular point p and set $\mathcal{P}' := \mathcal{P} - p^\perp$, and let the set of lines be of two types: (1) the lines of $\mathcal{L}$ not in $p^\perp$ and (2) the set of hyperbolic lines on p (these partition $\mathcal{P}'$ as one parallel class). Then $(\mathcal{P}', \mathcal{L}')$ is a generalized quadrangle of order $(s - 1, s + 1)$. The class of quadrangles that arise from this construction clearly extends beyond the $T_2^*(H)$ and the Hall-Arends/Szekeres examples since s can be odd.

7.12.2 The Last 15 Years Before 2001: Kantor-Families, the Heisenberg Group and Flocks of a Quadratic Cone

Nearly 15 years went by. What happened next involves highly special techniques and can't really be presented before the reader has an illuminating example. I will simply tell this story as an outside observer. Kantor (who was really hunting for a generalized hexagon) investigated various deformations and constructions of the classical generalized hexagon of type $G_2(q)$. (After all, many new planes were discovered by deforming the definitions of incidence in Baer subplanes and discernible

[29] The letter "T" stands for Jacques Tits, the "2" is the dimension of Π. The "$T_2^*(H)$-examples" were discovered by Sekeres and Arends [2] and by Marshall Hall [67].

reguli, so why not in similar subsets of hexagons?) Instead, he found a new quadrangle $K(q)$ of order (q^2, q) constructed from the $G_2(q)$ hexagon when $q \equiv 2 \mod 3$. Hirschfeld and Thas attempted to understand this from the point of view of normal curves $\{(1, t, t^2, t^3) | t \in \mathbf{GF}(q)\}^{(4)}$ (thus putting the hexagon into the background). Clearly the idea was that if only q would be in the correct congruence, the quadrangle was still a pretty classical generic object.

But then at the Como Conference "Groups and Geometry" (the famous one in which Tits sketched his classification of the affine buildings of rank at least four) Kantor [82] astounded everyone by presenting two entirely new families of generalized quadrangles – not only that, he had the courage to present a paradigm for getting them which included his recent $K(q)$s as well (that is, he actually presented three infinite families). The presentation in Como involved a family of subgroups of the finite "Heisenberg group" (to me a special p-group of order q^5 having center of order q) [82]. As Kantor observed, the required conditions on this family of subgroups could be defined for any group and, in this abstract setting, they came to be known as a *Kantor family*. Of course, as you might guess, it is probably a very special group that admits a Kantor family.[30] Evidence of this was adduced by Frohardt and his student Chen.

Actually there are two Kantor families, some big groups $\mathcal{F}^*$ and some smaller ones $\mathcal{F}$. In the Heisenburg case the big family, mod the center, is a system $\mathcal{C}$ of isotropic 2-subspaces of the symplectic 4-space induced on the Frattini factor (the form is defined by the commutator mapping). Stanley Payne derived a set of conditions making this possible, conditions described by equations. For a moment no one thought too much about these equations. (One might have thought that the idea of modern algebra was to *replace* coordinatizations and calculations by soft general principles which were just as good. This history shows that that is not always a good lesson.) Here is what happened:

1. The equations describing the Kantor families for the Heisenberg group were soon generalized (Stanley Payne's q-clans). The equations were good: a solution meant a generalized quadrangle.
2. They were good in characteristic 2, but no (non-classical) solutions to the equations were known.
3. Then Jef Thas recognized that for no apparent reason these equations were exactly those occurring in an entirely different subject – flocks of a quadratic cone. The mystery mounts.

Note the benefits of these observations. Item 1 opened the door to new quadrangles, as did item 2 in characteristic 2, without the benefit of a Kantor family. Jeff Thas visited the author at that time, excited that the equations which ensure that something is a flock of a cone were the same as those appearing in the Kantor–Payne equations. The interesting thing was that even if you did not know the reason why

[30] I suppose this statement might also benefit from some explanation, but you are not going to get it.

these two sets of equations were the same, the solution implied by one background implied a solution interpreted in the other context. Thus there was a correspondence

$$\text{flocks of a cone} \rightarrow \text{ Kantor families of the Heisenburg group}$$

which might not have been discovered without being able to see the two sets of equations side-by-side.

Almost at once one inherited several new families of generalized quadrangles, just because many classes of flocks of a quadratic cone were already known. For example, the Thas–Fisher flocks were known, but at that time did not correspond to any known quadrangles. But the mere coincidence of equations in two different fields did not really tell us *why* it worked. Finally Norbert Knarr straightened that out for odd characteristic: there was a geometric construction of the correspondence given above. A few years later J. Thas discovered a geometric explanation for it in even characteristic.

There are two lessons here: (1) It is not always good to have contempt for equations. Sometimes one can discover things through equations that would not have been noticed otherwise. (2) If a geometric construction holds with q odd, perhaps its algebraized version coughs up conclusions of the construction with q even, without the construction itself being in place.

To finish the story of quadrangles, each flock construction produced a quadrangle. But it is not that simple: given a flock (a certain partition of a quadratic cone by conics) there is also a derived flock, also giving a quadrangle.

On the other hand it was discovered (W. Kantor, N. Johnson) that there is a correspondence between certain projective planes and the flock-quadrangle story, so the non-unique correspondences are built on a triangle whose vertices are these three categories: certain projective planes, flocks of a quadratic cone in $PG(3, q)$, and the Heisenberg-group-constructions of quadrangles.

The net results are well-described by the survey articles by Stanley Payne *A garden of quadrangles* [97] and the article on generalized polygons by J. Thas in the *Handbook for Incidence Geometry* [129].

7.12.3 Prospects for a Classification

The 2001 census lists fewer than 20 infinite families of generalized quadrangles. There are very few sporadic examples (isolated single freaks belonging to no describable infinite family at all)! So the situation is not as bad as it was for the projective planes.

Should one be as bold as the finite group theorists and guess that these examples are all? I should say no. By the time the most optimistic finite-group-theorist (Gorenstein) stated his program, there were thousands of pages of pure methodology in place. That is not the case with the quadrangles. Yes, we can axiomatically characterize the classical ones. Can we characterize even the quadrangles of order $(q - 1, q + 1)$? With q odd, there is only the Payne construction, and

each of these quadrangles possess this unique parallel class. Must that be there? Is it characterizing? We have a long way to go. Probably we have discovered what quadrangles we have merely because the pitiful methods we possess made them visible. In that respect it is somewhat like looking at the current list of extra-solar planets! Aside from the $(q-1, q+1)$ class or its dual, all of the non-classical quadrangles have been obtained from the flocks-quadrangle nexus or a potential ovoid in $PG(3, 2^{2n+1})$, and all of these are of order (q^2, q). Are there any new quadrangles of order (q^2, q^2)? Potential cases like this exist (see [117]).

Most of this discussion – like the parallel narrative concerning projective planes – concerned finite planes. The infinite case is blessed (or perhaps cursed would be a better word) by many free constructions.

But, as we noticed in the case of planes, Moufang planes, whatever their order, are classified by either a division ring, or a Cayley division algebra – that is, they are bonafide algebraic objects. The correct (but not as pleasant) concept for planes was "Moufang" – not Desarguesian. The same holds for generalized quadrangles. There is a natural hypothesis of "being Moufang" (referring to a minimally rich group of automorphisms). Finite Moufang quadrangles are classical (in fact a result of Thas, Payne, and van Maldegham asserts that in the finite case one gets everything from a weaker "half Moufang" condition [130].[31])

In the infinite case many other Moufang quadrangles make their appearance. This result has been in the works for years and has at last been solved in the recent monumental treatise by Tits and Weiss [140]. In fact their work goes well beyond generalized quadrangles. Many years ago Tits defined the notion of a generalized n-gon (introduced in Chap. 9) as a cornerstone in defining buildings. According to this notion, projective planes turn out to be generalized 3-gons, and non-degenerate generalized quadrangles (as the name suggests) are generalized 4-gons. Finite generalized n-gons with thick lines exist only for $n = 3, 4, 6, 8, 12, 16$ and in the last two cases can have only two lines on each point [62]. The Moufang condition can be uniformly defined for all of these generalized polygons, and the great theorem of Tits and Weiss classifies every single one of them.

[31] In Chap. 11 of the book *Translation Generalized Quadrangles* by Thas et al. [131], one may find many weakened versions of the Moufang condition that still imply it. Work of Tent and Van Maldegham ([134]) has extended some of these the implications to infinite quadrangles.

Chapter 8
Near Polygons

Abstract Near polygons are point-line geometries with the property that, given any point and any line, that line contains a unique point nearest the given point. Under a mild condition, one can show that any two points at distance two lie in a unique convex subquadrangle called a "quad." Cameron's theorem states that if a near polygon has quads and each quad is strongly gated, then the geometry is a dual polar space. The chapter offers a complete proof.

8.1 Introduction

8.1.1 What We Know About Polar Spaces

This is a chapter on near polygons. The main theorem of this subject is Cameron's Theorem [24].

Theorem 8.1.1 *Any near polygon of finite point-diameter with thick lines and gated quads is a dual polar space.*

All of the undefined terms appearing in this theorem will be fully explicated below. The theory centers on the interplay between the distance metric of the point-collinearity graph and lines themselves, developing from this an elaborate poset of subspaces whose top members form a polar space, and then applying the classification of polar spaces as given in the following theorem.

Theorem 8.1.2 *Suppose* $\Gamma = (\mathcal{P}, \mathcal{L})$ *is a non-degenerate polar space of rank at least three. Then* Γ *is one of the following:*

1. *A classical polar space of rank at least* 3.
2. *A non-embeddable polar space of rank exactly three.*

The previous chapter was basically devoted to the proof of just two major propositions: (1) if a non-degenerate polar space of rank at least three is embeddable, then it is classical, and (2) if a non-degenerate polar space has rank four, it is in fact embeddable a priori. This statement completely forces the dichotomy of the preceding theorem.

E.E. Shult, *Points and Lines*, Universitext, DOI 10.1007/978-3-642-15627-4_8,

© Springer-Verlag Berlin Heidelberg 2011

As described in the last section of Chap. 7, Tits obtained a complete classification of polar spaces of the last category.

8.1.2 The Definition of Near Polygons

Every time we encounter a point-line geometry $\Gamma = (\mathcal{P}, \mathcal{L})$, we also encounter a point-collinearity graph $\Delta = (\mathcal{P}, \sim)$ in which a distance function $d : \mathcal{P} \times \mathcal{P} \to \mathbf{N}$ is defined. Consider the following basic axiom:

(NP) *Given any point-line pair $(p, L) \in \mathcal{P} \times \mathcal{L}$, there exists a unique point on line L nearest p.*

Of course, if the point p is incident with the line L (so the p is actually one of its points), then p has distance 0 from itself and distance one from every other point on that line. In general, axiom (NP) is saying that for each point-line pair (p, L), the non-decreasing list of distances from p to a point of L reads $\{d, d + 1, d + 1, d + 1, \ldots\}$ — that is, only one point of L is at distance d from p, the rest are at distance $d + 1$. Of course the words "unique" and "nearest" in the axiom would not operate if the distance d were infinite. Thus we may take axiom (NP) to imply that the point-collinearity graph is connected (or equivalently that the line coaxiality-graph is connected) – either that or the uninteresting case that there are no lines whatsoever, technically a near polygon, but one for which the author apologizes.

Of course if all lines have just two points (all lines are thin), then we are talking about the vertices and edges of a bipartite graph. Extending from this special case we see that, in general, we are discussing a geometry in which every line is a strongly gated subspace.

Now note that if a point p has distance one from two distinct points of a line L, then p must have distance zero from some point of L – that is, it belongs to L. Thus we observe the following.

Lemma 8.1.3 *If $\Gamma = (\mathcal{P}, \mathcal{L})$ is a near polygon, then the following statements hold:*

1. *The maximal cliques of the point-collinearity graph $\Delta = (\mathcal{P}, \sim)$ are exactly the point-shadows of the lines.*
2. *Γ is a partial linear space.*
3. *(Recognizing d-opposite lines.) Suppose x and y are distinct points of line L and u and v are distinct points of line N. If*

$$d(x, u) = d(y, v) = d, \text{ and}$$
$$d(x, v) = d(y, u) = d + 1,$$

then there is a bijection $\epsilon : L \to N$ from the point shadow of L to the point-shadow of N, such that corresponding points are at distance d and all other pairs of points of $L \times N$ not of the form $(x, \epsilon(x))$ are at distance $d + 1$.

We say that pairs of lines L and N satisfying the conclusion of part 3 of the Lemma, are *d-opposite lines*. Note that with this notation the parameter d denotes the *minimum* distance of point-pairs in $L \times N$. For example, non-intersecting lines of a generalized quadrangle are 1-opposite lines.

8.1.3 Some Non-classical Examples

Clearly the spaces with no lines may be said to have already been classified. Otherwise, the near polygon Γ is connected. In that case the point-collinearity graph $\Delta = (\mathcal{P}, \sim)$ is connected. If this graph possesses a *finite diameter n* – that is, a least upper bound n on the possible distances (in Δ) between distinct points – the near polygon is called a *near n-gon*.

We have already met a class of examples of near polygons. If n is an even integer, then a generalized n-gon is in fact a near polygon. The reason for this is that when n is even, the definition of "generalized n-gon" has another formulation, as follows.

Definition If n is an even positive integer, a point-line geometry $\Gamma = (\mathcal{P}, \mathcal{L})$ having a point-collinearity graph of diameter $n/2$ is called a *generalized n-gon* if and only if, for every non-incident point-line pair (p, L), there is a unique point q on L nearest p and there is a *unique* path on the point-collinearity graph from p to q. The generalized polygon Γ is said to be *non-degenerate* if and only if, any pair of points at the maximal distance $n/2$ are connected by at least two geodesic paths.

From this definition, it should be transparent that any generalized n-gon, as defined above for n even, is already a near polygon. However, there exist many more examples.

Rather obviously a near 1-gon is just a line. What is a near 2-gon?

In this case, there must exist points at distance two, and so there exist non-incident point-line pairs (p, L). For each of these pairs (p, L), p is collinear with exactly one point of L. The reader should recognize this property as the key axiom for generalized quadrangles. So we record the following.

Lemma 8.1.4 *Any generalized 2-gon is a (possibly degenerate) generalized quadrangle.*

No one can dispute that a "*star*" – that is, a point p, together with a collection of lines (and their points) which pairwise meet at $\{p\}$ – is a near 2-gon. But it is a degenerate generalized quadrangle.

For $n > 2$ the story is quite different. There are many near polygons which are not generalized polygons.

Let us being with a very easy example.

Near Polygons of Hamming Type

Fix an integer n greater than 2, and an "alphabet" A (actually any set). We set $\mathcal{P} = A^{(n)}$, the collection of all n-tuples with entries from the alphabet A. (These are just sequences of length n of elements of A, but in the suggestive language which

envelopes this situation, these n-tuples are thought of as "words" of length n in the "alphabet" A). These words are the points.

A "line" is the set of all "mispellings" of a word that could result if just one of its letters was replaced by anything in the alphabet A. For example, if letters $a_1, \ldots, a_{i-1}, a_{i+1}, a_{i+2}, \ldots, a_n$ are prescribed in advance, the set of points

$$\{(a_1, \ldots, a_{i-1}, x, a_{i+1}, \ldots, a_n)\} | x \in A\},$$

is a line. It is said to belong to the set $\mathcal{L}_i$, because this line has a free alphabetic variable only in the i-th position. In this way, the full set of lines enjoys a partition $\mathcal{L} = \mathcal{L}_1 + \mathcal{L}_2 + \cdots + \mathcal{L}_n$ into so-called "parallel classes", $\mathcal{L}_i$.

The distance between points (words) in the point-line geometry $\Gamma = (\mathcal{P}, \mathcal{L})$ is the minimal number of "spelling corrections" it takes to convert one word into the other. Put another way, the distance between two words is equal to the number of positions in which the letter-entries offered by the two words differ. It follows that $(\mathcal{P}, \mathcal{L})$ is in fact a near polygon. Suppose $p = (b_1, \ldots b_n)$ is a point, and L is the line whose j-th coordinates are fixed for all j distinct from i, while the i-th coordinate x freely roams over the alphabet A. Then the nearest point on L is found by setting $x = b_i$, since only then is the distance from p reduced by one from all the other values.

The student might want to look at the 27-point example obtained when $n = 3 = |A|$, and check that this is not a generalized hexagon.

The Near Hexagon of M_{24}

This example occurs because of a remarkable coincidence concerning multiply transistive groups and error-correcting codes. Since our only objective is to *describe* the example, I will not prove most of the statements.

Let $V = k^{(n)}$ be the vector space of all n-tuples over a field k. One thinks of this as an n-dimensional vector space with a specified basis B with many of the concepts depending on this basis. The *weight* of a vector $v = (a_1, \ldots, a_n) \in V$ is the number of a_i which are non-zero. We denote the weight of v by $\mathrm{wt}(v)$.

This situation can be described in the language of the previous section. The vectors of V are words of length n in the alphabet $A = k$. The weight of a vector is just its distance from the zero vector in the Hamming near polygon (or *Hamming scheme*). Note that with respect to this scheme, we have

$$d(u + v, w + v) = d(u, w)$$

for all vectors $(u, w, v) \in V \times V \times V$. Thus vector-addition is a group $(V, +)$ of isometries.

The theory of error-correcting codes fully utilizes this metaphor of words and mispellings. Here the object is to find a pairwise high-distance subset of points (called the code) that can be easily dealt with. Usually "easily dealt with" means the subset is a linear subspace.

In the case that $k = \mathbf{G}F(2)$ the situation is even more special. Let (x, y) denote the ordinary "dot product" $\sum_i a_i b_j$ of vectors

$$x = (a_1, \ldots, a_n) \text{ and } y = (b_1, \ldots, b_n).$$

Then

$$\text{wt}(u + v) = \text{wt}(u) + \text{wt}(v) - (u, v).$$

This is not generally true of other fields.

Now there is a remarkable error-correcting code C (called the Golay binary code) sitting in $\mathbf{G}F(2)^{(24)}$ as a 12-dimensional subspace. It contains the unique zero-vector (of weight zero), the unique "all-one" vector $j = 1^{24}$, of weight 24, 759 vectors w of weight 8, another 759 vectors of weight 16, and, in the middle, 2576 vectors of weight 12 – making a total of $1 + 759 + 2576 + 759 + 1 = 2^{12}$ vectors in the subspace C of V. What makes it a good code is not only the high minimum non-zero weight (eight), but the fact that it sits in the total Hamming code in a way noticeable to every other vector. Given any vector v, in V, there is a unique vector in C nearest v (when this happens C is called a "perfect code").[1]

Obviously the last assertion resembles axiom (NP) and suggests some relationship between the two. In fact the suggestion is that C is a gated subset of the enveloping Hamming scheme.

Now when an error-correcting code is nearly perfect, it is a fact that the supports of its minimal weight vectors forms a t-design for fairly high values of t. This thesis was first introduced by Assmus and Mattson [5].

A *t-design* $S(t, k, v)$ is a collection V of v points, and a family $\mathcal{B}$ of k-subsets of v such that, for every t-subset that you name, there is a *unique* block containing that subset.

Applying the Assmus theory to the binary Golay code one obtains a t-design $S(5, 8, 24)$.

I beg the reader to allow me to tax his or her sense of reality with the following example. The president of a midwestern state university finds he can trust only 24 faculty members to support his bizarre policies. In an effort to maintain an appearance of faculty democracy, he is determined to fill 759 commitees only with faculty members. In order not to waste faculty time, and still maintain an appearance of having solicited widely-differing viewpoints, each commitee shall have exactly eight members and any two of these committees will share at most four common members, so that the view point of one committee cannot dominate another. Even under those rules, it is possible for the president to choose these members from just his 24 trusted faculty cohorts.

If this were done one would have 759 commitees of size 8, chosen from only 24 persons, with the property that any 5 of the 24 persons serve together on exactly one commitee. In other words this is a Steiner system $S(5, 8, 24)$.

[1] For a good account of the Golay code, see [39] or [12].

Now we connect this with the error-correcting codes. We simply watch the subspace of the additive subspace of the Boolean algebra generated by the 759 blocks of the Steiner system $S(5, 8, 24)$. This turns out to be a 12-dimensional space C with the property that the input 759 commitees of size 8 are the only non-empty commitees of this size or smaller in C. In other words, we have come full circle in defining this structure: block designs vs error-correcting codes.

Finally we remark that both the Steiner system $S(5, 8, 24)$ and the even-weight minimal-weight-8 code containing 759 weight-8 vectors are unique up to the appropriate isomorphisms. We will use both points of view in describing the following near polygon.

Our points $\mathcal{P}$ will be the 759 "octads" of $S(5, 8, 24)$. Our lines $\mathcal{L}$ will be the *triads*: partitions of the 24 letters into 3 pairwise disjoint octads. Obviously an exercise in counting shows that there exist lots of octads B which intersect an octad A trivially. When this happens, the coding theory tells us that the complement of their union is a weight-eight vector of the code, and hence is a further octad C. Then $\{A, B, C\}$ is a line.

Now, from the coding theory, any two distinct octads intersect at 0, 4, or 2 points. Suppose $L = \{A, B, C\}$ is a line of Γ – that is $A + B + C$ is a partition of the 24 letters Ω into 3 sets of 8. Now if D is a further octad, the partition A,B, C induces a generalized partition[2] $(D \cap A, D \cap B, D \cap C)$. The cardinality of the parts $D \cap A$, etc. is one of the following triplets:

$$(|A \cap D|, |B \cap D|, |C \cap D|) = (4, 4, 0) \text{ or } (2, 2, 4).$$

Since "distance one" means empty intersecton, the first triplet means that octads which intersect at four points are at distance two in the collinearity graph of Γ. Then the second triplet tells us that octads which meet at two points represent points at distance three in Γ. The distribution of numbers within each triplet reveals that there is a unique lowest distance. Thus Γ is a near hexagon.

We shall return to this example to illustrate a number of exceptional principles.

A Near Polygon Based on M_{12}

The ternary Golay code, is a six-dimensional subspace C of $V = \mathbf{G}F(3)^{(12)}$. Its vectors possess only the weights 0, 6, 9, and 12. In C there are exactly 24 vectors of weight 12 forming a set X with $X = -X$. Each vector in X exhibits six -1s and six $+1$s. Of course then, these vectors are the non-zero vectors of exactly 12 1-subspaces $L_1, L_2, \ldots, L_{12}$ which we regard as "lines" on the zero vector.

Now let $\mathcal{P} = V$ and let $\mathcal{L}$ be the vector translates, $L_i + c$, $1 \leq i \leq 12$, $c \in C$, of the 12 weight-12 lines in the code C. Thus one sees that two distinct vectors represent collinear points if and only if their difference is in X. Then the geometry

[2] In a partition, the components are defined to be non-empty. In a "generalized partition" this requirement is relaxed.

$\Gamma = (\mathcal{P}, \mathcal{L})$ is a near hexagon with 729 points; each point is on 12 lines, and each line has size 3. (See [119] for the details proving this assertion.)

8.1.4 Near Polygons Arising from Chamber Systems of Buildings

We have met the notion of a chamber system of a residually connected geometry over a type set I in Chap. 2 (p. 54). In many cases the rank one residues of the chamber system are strongly gated with respect to any other chamber. (Buildings are examples of this condition, but they are not all examples.) Then, taking chambers as points, and rank one residues as lines, one obtains a near polygon.

Consider, for example, the *projective geometry* $\mathbf{P}G(3, D)$, of all proper subspaces of a four-dimensional right vector space over a division ring D. The points of our near polygon will be the collection of all maximal flags – that is, chains of subspaces $(W_1 \leq W_2 \leq W_3)$ where $\dim W_i = i$. The lines are the rank-one residues – that is, residues of flags of corank one. There are thus three kinds of lines, the residues of types 1,2 and 3. For example, if we select a 1-space W_1 contained in a 3-space W_3, then the full collection of flags $\{(W_1, X, W_3) | W_1 \leq X \leq W_2\}$ is a line of our geometry.[3]

For example, when $D = \mathbf{G}F(2)$, we have a near polygon with 315 points, and 315 lines of size 3, 105 of each type. Each point is on nine lines, three of each type.

8.1.5 Near Polygons of Fischer Type

In this construction, the points are the maximal point-cocliques of a reduced Fischer space. Since these words might not be universally meaningful among students, allow me to describe this another way.

Suppose t^G is a conjugacy class of elements of order two (involutions) in a finite group G with these properties:

1. Any product of two members of t^G is an element of order 1, 2, or 3. (In these respective cases the selected involutions generate a group of order 2, the Klein four-group, or the symmetric group on three letters, respectively.)
2. The class t^G generates G.
3. For any two involutions $x, y \in t^G$, $C_G(x) \cap t^G = C_G(y) \cap t^G$ if and only if $x = y$.

Then G is a *reduced* 3*-transposition group*. For our new geometry $\Gamma = (\mathcal{P}, \mathcal{L})$, let $\mathcal{P}$ be the collection of all maximal sets of pairwise commuting involutions in t^G. Next set

$$\mathcal{L} = \{C_G(y) \cap M | M \in \mathcal{P}, y \in t^G - M\}.$$

[3] Projectively, this is a pencil of projective lines on a projective point within a projective plane

Thus if $M \in \mathcal{P}$ is such a maximal set of mutually commuting involutions, and y is a 3-transposition of t^G not in this set, then those elements of M which commute with y, the set $C_G(y) \cap M$, will be considered to be a "line". Amazingly, this set $C_G(y) \cap M$ always has the same cardinality.[4]

It was noticed by J. Hall that $(\mathcal{P}, \mathcal{L})$, as just defined, is a near polygon. See [11] for the proof of this. Some cases are examples of classical dual polar spaces over $\mathbf{G}F(2)$ where the 3-transpositions are symplectic transvections. But there are many other Fischer spaces providing other near polygons. Here is a simple example:

Let $n = 2m$ be an even integer, and let K_n be the complete graph on n vertices. A 1-*factor* is a (necessarily spanning) subgraph with exactly $m = n/2$ edges which are pairwise vertex-disjoint. Of course, for the complete graph, the set of 1-factors is bijective with the set of partitions of the vertex set all of whose components (or blocks) have size 2. Suppose we have two 1-factors which agree in $m - 2$ components, $b_1, \ldots, b_{m-2}$. Then these disjoint pairs b_i cover exactly $n-4$ vertices, leaving just four vertices to be partitioned into two pairs to make a 1-factor. This final act can be performed in exactly three ways, giving the two original 1-factors as well as a new one. We think of these three 1-factors sharing $m - 2$ edges, as being the three "points" of a "line."

So in general, we let $\mathcal{P}$ be the set of all 1-factors of K_n, and let $\mathcal{L}$ be the collection of all 1-factors of any subset of $n - 4$ vertices of K_n. A line L is incident with a point p if and only if L, as a 1-factor on a set of $n - 4$ vertices, can be completed to the 1-factor p.

Note that the union of the edges of two 1-factors is a disjoint union of polygons of even length, where, we understand a "polygon of length two" to be two vertices connected by a repeated edge. For example, if p and q are distinct collinear points of the point-line geometry $(\mathcal{P}, \mathcal{L})$, the two 1-factors make $m - 2$ "polygons" of length 2 and a 4-circuit. In general, this structure reveals the distance between any two points p and q. These points are at distance d if and only if the largest polygon in the subgraph defined by the edges of p and q has girth $2d + 2$. The resulting association scheme on points has many nice combinatorial properties (see the paper of Brouwer and Wilbrink [13]).

Now, the conclusion that $(\mathcal{P}, \mathcal{L})$ is a near polygon follows from the fact that a 1-factor is a maximal set of mutually commuting transpositions of the symmetric group, and the observation of J. Hall mentioned above.

8.1.6 Classical Near Polygons: The Dual Polar Spaces

Let Δ be a non-degenerate polar space of finite rank $r \geq 3$. Then each maximal singular subspace of Δ is isomorphic to the projective space derived from

[4] Any graph for which (1) maximal cliques have a common cardinality, and (2) the intersections $y^{\perp} \cap M$, where M is a maximal clique and y is a vertex not in M, also have a common cardinality, is said to possess the *Zara property*. This fundamental property was brought to prominence in the study of 3-transposition groups by Francois Zara [148].

$\mathbf{P}G(r-1, D)$, D a fixed division ring, by truncation to its points and lines. Let $\mathcal{M} = \mathcal{M}(\Delta)$ denote this collection of maximal singular subspaces of Δ.

Consider the point-line geometry $\Gamma = (\mathcal{M}, \mathcal{H})$, where $\mathcal{H}$ is the collection of all hyperplanes of all projective spaces in $\mathcal{M}(\Delta)$, using containment as incidence. It is obvious that Γ is a partial linear space, whose point-collinearity graph $G = (\mathcal{M}(\Delta), \sim)$ is defined by declaring two vertices A and B of G to be adjacent if and only if they intersect at a hyperplane $H = A \cap B$ of both A and B.

Theorem 8.1.5 (The polygonal structure of maximal elements of a polar space.) *Let Δ be a polar space of finite rank r at least three, relative to the sesquilinear form (V, f). Let $G = (\mathcal{M}(\Delta), \sim)$ be the adjacency graph on the maximal singular subspaces of Δ. Then:*

1. *The graph G is connected.*
2. *For any two vertices A and B be of $\mathcal{M}(\Delta)$, the distance $d_G(A, B)$ in the graph G is the codimension of $A \cap B$ in A.*
3. *Let A and B be elements of $\mathcal{M}(\Delta)$, and let H be a hyperplane of B which contains* $\mathrm{rad}(V)$. *Then, among the elements of $\mathcal{M}(\Delta)$ which contain H, there is a unique element which is nearest A in the graph G.*

Proof The first two parts are proved together.

Since Δ has finite rank, $\mathrm{rad}(\Delta)$ has finite codimension in every element of $\mathcal{M}(\Delta)$. Choosing A and B in $\mathcal{M}(\Delta)$ we see that $d = \mathrm{codim}_A(A \cap B) = \dim(A/(A \cap B))$ is finite. We shall prove part 2 by induction on d.

If $d = 1$, then A and B are adjacent vertices in G. So we assume $d > 1$.

Choose vector b in $B - (A \cap B)$. Then $b^\perp \cap A = H$ is a hyperplane of A containing $A \cap B$. Now form the subspace $A_1 := \langle b, H\rangle$. We claim that A_1 is an element of $\mathcal{M}(\Delta)$. First of all, $A_1 \in \Delta$ since $\langle b\rangle$ and H both belong to Δ and $\langle b\rangle \subseteq H^\perp$. Second, A_1 must be maximal in Δ. For suppose otherwise. Then A_1 is properly contained in a space T belonging to Δ. Thus H has codimension at least two in T, and, as $A = \langle H, a\rangle$, we see that $A^\perp \cap T$ is a hyperplane of T. But since $(A^\perp \cap T) + A \in \Delta$ and A is maximal in Δ, we must have that $A^\perp \cap T \leq A$ so A itself is a hyperplane of T. But in that case A is not maximal in Δ, a contradiction. Thus, we see that $A_1 \in \mathcal{M}(\Delta)$.

Now A_1 is adjacent to A in the graph G. On the other hand since $A_1 \cap B \supseteq \langle b, A \cap B\rangle$ properly contains $A \cap B$, we may apply induction on the codimension of $A_1 \cap B$ in B to see that $d_G < d$. In particular B is in the same connected component of G as is A_1 and A, and the distance $d_G(A, B)$ is at most d.

On the other hand, we claim that $d_G(A, B)$ is at least $d = \dim(A/(A \cap B))$. Suppose, by way of contradiction, that A and B were connected by a walk $(A, A_1, A_2, \ldots, A_k = B)$ of length k less than d. Then $A \cap B$ contains $A \cap A_1 \cap \cdots \cap A_k$. But since $A_{i+1} \cap A_i$ is a hyperplane of A_i, the codimension of $A \cap A_1 \cap \cdots \cap A_{i+1}$ in A can be at most one more than the codimension of $A \cap A_1 \cap \cdots \cap A_i$ in A. Thus $A \cap A_1 \cap \cdots \cap A_k$ has codimension at most k in A. But since $A \cap B$ contains this intersection, we have

$$\mathrm{codim}_A(A \cap B) \leq k < d = \mathrm{codim}_A(A \cap B),$$

an absurdity.

Thus we see that

$$d_G(A, B) = d = \dim(B/(A \cap B)) = \dim(A/(A \cap B)),$$

exactly.

Now we address part 3. Let A and B be maximal elements of $(\Delta, \leq)$, and let H be a hyperplane of B. We have already seen that the finite rank assumption forces $A \cap B$ to have finite codimension in A, and that there is an isomorphism

$$\bar{\alpha} : (B/(A \cap B) \to (A/(A \cap B))^*.$$

It follows that $H \cap A$ has finite codimension e in A, and codimension $e-1$ in H. This difference in codimensions shows that $H^\perp \cap A$ contains $H \cap A$ as a hyperplane. Then $B_0 = \langle H, (H^\perp \cap A)\rangle_\Gamma$ is an element of $\mathcal{M}(\Delta)$ which meets A at $H^\perp \cap A$ with codimension $e-1$ in A. Thus B_0 has distance $e-1$ from A in the graph G.

Conversely, suppose B' were a member of $\mathcal{M}(\Delta)$ containing H and having distance $e-1$ from A. Then by part 2, $B' \cap A$ has codimension $e-1$ in A and contains $H^\perp \cap A$, also at codimension $e-1$ in A. It follows that $B' = \langle H, (H^\perp \cap A)\rangle_\Gamma = B_0$. So there is a unique element of $\mathcal{M}(\Delta)$ containing H and having distance $e-1$ from A. □

There is one last relevant detail.

Theorem 8.1.6 *Suppose Δ is a polar space of finite rank $r > 0$. Then the following statements are true:*

1. *Fix a maximal singular subspace A in $\mathcal{M}(\Delta)$. Then there exists a maximal space B in $\mathcal{M}(\Delta)$ such that $A \cap B = Rad(\Delta)$.* (In that case we say that B is *opposite* A.)
2. *Let H be a hyperplane of some element of $\mathcal{M}(\Delta)$ chosen so that H contains $rad(\Delta)$. Then H lies in at least two members of $\mathcal{M}(\Delta)$.*
3. *Fix integer $k < r$. Then in the graph G, every geodesic path of length k extends to one of length $k+1$.*

Proof Throughout, we let $\mathcal{T}$ be the poset of all singular subspaces of Δ.

1. In proving this part, we may assume that $\mathrm{Rad}(\Delta) = \emptyset$ and that all members of $\mathcal{M}(\Delta)$ have projective dimension $r-1$. Fix A in $\mathcal{M}(\Delta)$ and choose $B \in \mathcal{M}(\Delta)$ so that the dimension of $A \cap B$ is as small as possible. If $A \cap B = \emptyset$ we are done, so assume $A \cap B$ is non-empty Since $\mathrm{Rad}(\Delta) = \emptyset$ there is a maximal element D in $\mathcal{M}(\Delta)$ which does not contain $A \cap B$, and hence an element x in D so that $x^\perp$ does not contain $A \cap B$.

Now $x \in \mathcal{T}$, and x does not lie in B. Thus $C := \langle x, x^\perp \cap B\rangle$ is in $\mathcal{T}$ and because its dimension is $r-1$, C lies in $\mathcal{M}(\Delta)$. We claim that $C \cap A$ is properly contained in $B \cap A$. Suppose there were a 1-space $\langle z\rangle$ in $C \cap A - B$. Then $z^\perp$ would contain $A \cap B$ and the hyperplane $x^\perp \cap B = C \cap B$ of B, and so would contain $B = (C \cap B) + (A \cap B)$. But this would force $z \in B^\perp$ so $\langle z, B\rangle \in \mathcal{T}$. Since

B is maximal in $\mathcal{T}$ we must then have z in B, contrary to our choice of z. Thus $C \cap A \leq B \cap A \cap x^{\perp}$ which lies properly in $A \cap B$. But that contradicts the minimality of the intersection $A \cap B$. Thus the minimal intersection with A must be 0.

2. Suppose H is a hyperplane of $B \in \mathcal{M}(\Delta)$. By part 1, there exists a maximal element A of $\mathcal{T}$ such that $B \cap A = \text{Rad } \Delta$. It follows that $H^{\perp} \cap A$ contains Rad Δ as a hyperplane, and that $C := \langle H^{\perp} \cap A, H \rangle$ is a member of $\mathcal{M}(\Delta)$ containing H which is different from B.

3. Suppose A and B are vertices of $G = (\mathcal{M}(\Delta), \sim)$ at distance $k < r$. Then $A \cap B$ has codimension k in B and so properly contains Rad Δ. Thus we can select a hyperplane H of B containing Rad Δ but not containing $A \cap B$. Now by part 2, H lies in a maximal element $C \in \mathcal{M}(\Delta)$ distinct from B. Since $H \cap A$ has codimension $k+1$ and $H \cap A$ is either equal to or is a hyperplane of $C \cap A$, $C \cap A$ cannot have codimension $k-1$ in A. Applying Theorem 8.1.5, part 2, all the elements of $\mathcal{M}(\Delta)$ which contain H have distance at least k from A. But by part 3 of the same theorem, among the elements of $\mathcal{M}(\Delta)$ which contain H, there is a unique one closest to A. Evidently that is B. Since $C \neq B$ we must have $d_G(A, C) = k+1$, and our theorem is proved. □

Now Theorem 8.1.5, part 2, shows that the lines of the geometry Γ are always gated in the collinearity graph and so Γ is a *near polygon*.

Note that when Δ is *oriflame*, every element of $\mathcal{H}$ is incident with exactly two members of $\mathcal{M}(\Delta)$. In that case Γ has only thin lines, forcing the graph G to be bipartite. That is why there are two classes of maximal singular subspaces in this case – there is no other reason.

8.2 The Existence of Quads in Near Polygons

8.2.1 Internal and External Diameter

Suppose $\Delta = (V, E)$ is a simple connected graph. Consider any induced subgraph on the vertex set X, a subset of V. As a graph, X possesses a diameter, namely

$$\text{diam}(X) := \sup\{d_X(x, y) | (x, y) \in X \times X\}.$$

We call this the *internal diameter* of the induced subgraph X. But there is also another notion which we call the *external diameter* of X, namely

$$\text{ext diam } (X) := \sup\{d(x, y) | (x, y) \in X \times X\},$$

the maximum distance, as distance is measured by paths throughout the entire graph, between two vertices of X.

Consider a simple hexagon $V = \{x_0, \ldots, x_5\}$ as a valence two graph whose edges are the pairs (x_i, x_{i+1}), where the indices i are taken mod 6. Then $X = \{x_1, x_3, x_5\}$ is a coclique of size three. Thus X has infinite internal diameter, while its external diameter is only 2.

We always have

$$\text{ext diam}\,(X) \leq \text{diam}(X).$$

Now let $(\mathcal{E}_k, \subseteq)$ be the set of all subsets Δ of external diameter k, partially ordered by the inclusion relation. Let $\mathcal{E}_k^*$ be the maximal members of this poset. Thus if $U \in \mathcal{E}_k^*$ and y is a vertex outside U, then we must have $d(y, u) > k$ for some vertex $u \in U$.

Lemma 8.2.1 *Any element of $\mathcal{E}_k$ lies in a member of $\mathcal{E}_k^*$.*

Proof This is a straightforward application of Zorn's lemma. □

8.2.2 Bounded Diameter Point-Sets in Near Polygons

Let us recall the following notation used for any integer-valued metric space $\mathcal{P}$. For each point p, the set

$$\Delta_k^*(p) := \{y \in \mathcal{P} | d(p, y) \leq k\},$$

of all points at distant *at most* k from p, is called the *ball of radius k about p*.

Now we make a simple observation.

Lemma 8.2.2 *Suppose $\Gamma = (\mathcal{P}, \mathcal{L})$ is a point-line geometry. Then Γ is a near polygon, if and only if for each point p, and each integer i not exceeding the diameter d:*

1. *Each $\Delta_i^*(p)$ is a subspace of Γ.*
2. *For $i > 0$, $\Delta_i^*(p)$ is either equal to, or is a geometric hyperplane of $\Delta_{i+1}^*(p)$.*

Proof The proof is left as an exercise (Exercise 8.1 in Sect. 8.6). □

Corollary 8.2.3 *Suppose Γ is a near polygon. Then:*

1. *Each member of $\mathcal{E}_k^*$ is a subspace of Γ.*
2. *If X is an element of $\mathcal{E}_2^*$, then either*
   - *(a) X is totally disconnected – that is, it is a set of pairwise non-collinear points, or*
   - *(b) X is a (possibly degenerate) generalized quadrangle.*

Proof From the maximality, it is clear that if $X \in \mathcal{E}_k^*$, then

$$X = \cap_{x \in X} \Delta_k^*(x).$$

But the right side is an intersection of subspaces of Γ, and so is itself a subspace. That proves the first part.

Now suppose X is a maximal member of $\mathcal{E}_2$. Suppose L is a line connecting two distinct points of X. Then L lies in X, by part one. Suppose a and b are distinct points of L and y is any point of X not on L. Then y has distance at most two from

two members of L, and so, by (NP), y is collinear with a point of L. If y is not itself a member of L, this point of $y^\perp \cap L$ is unique. This is enough to prove the dichotomy asserted in the second statement. □

8.2.3 The Existence of Quads

Now we are in a position to prove the following fundamental result.

Theorem 8.2.4 (Yanushka's lemma.) *Suppose a and c are two non-collinear points in a near polygon $\Gamma = (\mathcal{P}, \mathcal{L})$, having the property that there are two distinct points b and d in $a^\perp \cap c^\perp$ (then of course b and d are not collinear).*

Then, if at least one of the four lines ab, bc, cd, da is thick, then the set

$$\Delta_2^*(a) \cap \Delta_2^*(b) \cap \Delta_2^*(c) \cap \Delta_2^*(d)$$

is a generalized quadrangle.

Remark The quadrangles formed in this Lemma are called *quads*.

Proof Suppose (a, b, c, d, a) is a 4-circuit in the point-collinearity graph of the near-polygon Γ and set

$$Y = \Delta_2^*(a) \cap \Delta_2^*(b) \cap \Delta_2^*(c) \cap \Delta_2^*(d).$$

To show that Y is a generalized quadrangle, we need only show that it belongs to $\mathcal{E}^*$ and, noting that Y contains a 4-circuit, apply part 2 of the previous Corollary 8.2.3 to get that Y is a non-degenerate generalized quadrangle.

But actually, it suffices to prove that $Y \in \mathcal{E}_2$. For, by Lemma 8.2.1, Y then lies in a member Y^* of $\mathcal{E}_2^*$. Then, as $\{a, b, c, d\} \subseteq Y$, we have $\{a, b, c, d\} \subseteq Y^*$, so every point of Y^* has distance at most two from each of a, b, c, and d, whence $Y^* \subseteq Y$, by the initial definition of Y. So $Y = Y^* \in \mathcal{E}_2^*$ and the conclusion of the previous paragraph applies.

So let u and v be arbitrary elements of the defined set Q. We need only show that u and v are at distance at most two from *each other*.

First, we can assume $\{u, v\} \cap \{a, b, c, d\} = \emptyset$, since otherwise $d(u, v) \leq 2$ and we are done.

We make a simple observation:

(*) *Let L be one of the four lines ab, bc, cd, or da, and let y be any point of Y. Then y has distance at most one from some point of L.*

For, suppose false. By relabeling points if necessary, we can assume $L = ab$. Then $d(y, a) = 2 = d(y, b)$. By (NP) and the fact that a and b are distinct, y is collinear to a point of ab, contrary to our choice of y. So (*) holds.

Now we can assume that neither u nor v lies on one of the four lines ab, bc, cd, and da. For if u were on ab, say, then as v is collinear or equal to a point of ab, we should have $d(u, v) \leq 2$ and would be done. Thus:

(**) *Neither u nor v is on any of the four lines ab, bc, cd, or da.*

There are two basic cases:

Case 1. *u and v are each not collinear with some "corner" element* $\{a, b, c, d\}$.
Case 2. *Each member of* $\{a, b, c, d\}$ *is collinear with either u or v.*

Case 1. Without loss of generality, assume neither u nor v is collinear with point a. Now (mindful of (*) and (**)) u is collinear with a point u_1 on ab and a point u_2 on ad, where u_1 and u_2 are both distinct from a. Similarly, v is collinear with points v_1 and v_2 on ab and ad, respectively, each of these points being distinct from a. Also, $u_1 \neq v_1$ and $u_2 \neq v_2$ since otherwise $d(u, v) \leq 2$.

Then $d(v_1, u) = 2 = d(v_1, u_2)$. Then (NP) forces v_1 to be collinear with some point z on the line uu_2. Now z cannot be u_2 since a is the unique point on ad collinear with v_1. Also z cannot be u for then $(v, v_1, z = u)$ would be a path of length two. Thus v has distance two from z and has distance two from u_2 (via the path (v, v_2, u_2)), and so by (NP) must be collinear with a point of the line $zu_2 = uu_2$. Thus $d(u, v) \leq 2$, as required.

Case 2. Here each of the corner points is collinear with either u or with v. We can assume that no corner point is collinear with both u and v since otherwise $d(u, v) \leq 2$ and we are done.

So without loss of generality, we may assume u is collinear with a and c while v is collinear with b and d. Now we invoke for the first time the fact that one of the four lines of the 4-circuit is thick, and without loss of generality we take that line to be ab. So choose point e on line ab distinct from both a and b. Then e has distance two from both c and d, and so is collinear with a point f on cd. Now $C = (e, b, c, f, e)$ is a 4-circuit, whose lines are the same as before, except that ef replaces da. From the first sentence of this paragraph it follows that

$$\{u, v\} \subseteq Z := \Delta_2^*(e) \cap \Delta_2^*(b) \cap \Delta_2^*(c) \cap \Delta_2^*(f).$$

But now neither u nor v is collinear with e, and so by Case 1, with C replacing (a, b, c, d, a), we have $d(u, v) \leq 2$.

The proof is complete. □

8.3 The Convexity of Quads

Let us continue the consideration of a 4-circuit (a, b, c, d, a) having one of its four lines thick. We set

$$Q(a, b, c, d) := \Delta_2^*(a) \cap \Delta_2^*(b) \cap \Delta_2^*(c) \cap \Delta_2^*(d),$$

the generalized quadrangle produced in the previous section. We have already made the observation that

$$Q(a, b, c, d) \in \mathcal{E}_2^*.$$

In fact it is the unique member of $\mathcal{E}_2^*$ containing $\{a, b, c, d\}$.

Now let x be a point of $a^\perp \cap c^\perp - \{d\}$ so (a, x, c, d, a) is also a 4-circuit. Now $\{a, b, c, d\} \subseteq Q(a, x, c, d)$, so by the uniqueness, $Q(a, b, c, d) = Q(a, x, c, d)$. Iterating this step, we see that

$$Q(a, b, c, d) = Q(a, x, y, c), \tag{8.1}$$

for every non-adjacent pair (x, y) of points in $a^\perp \cap c^\perp$, and that $Q(a, b, c, d)$ contains every vertex of $a^\perp \cap c^\perp$ on a geodesic path from a to c. We conclude that:

(1) *$Q(a, b, c, d)$ is the unique member of $\mathcal{E}_2^*$ which contains the pair $\{a, c\}$. It contains all geodesics from a to c.*

Since the description depends only on a and c, we may write $Q(a, b, c, d) = Q(a, c)$. Or it could be written $Q(b, d)$ just as well.

Now note that if $Q(a, c)$ contains a thin line, then the quadrangle $Q(a, c)$ is a grid, and every point lies on one thin line and one thick line. Otherwise all its lines are thick. In either case every 4-circuit in $Q(a, c)$ carries a thick line. Thus applying statement (1) above, we conclude that $Q(a, c)$ is the unique member of $\mathcal{E}_2^*$ containing any 2-coclique $\{u, v\}$ of $Q(a, c)$. Thus we have proved the following.

Theorem 8.3.1 *Suppose a and c are two non-collinear points of a near polygon for which:*

1. $|a^\perp \cap c^\perp| > 1$.
2. *Some geodesic path connecting a and c carries a thick line.*

Then there is a unique member $Q(a, c)$ of $\mathcal{E}_2^*$ containing $\{a, c\}$. $Q(a, c)$ is a convex non-degenerate generalized quadrangle.

In general, in may happen that in some near polygons (even those with all lines thick), some pairs of points at distance two do not belong to a 4-circuit. In that case their is a unique geodesic path connecting them. We call such a pair a *special pair*. Other distance-two pairs do belong to a 4-circuit and if even one such circuit contains a thick line, the distance-two pair lies in a convex non-degenerate generalized quadrangle. We call such a pair a *quad pair* (or in some contexts, a *symplectic pair*).

Consider, for example, the near polygon of Fischer type whose points are the 1-factors of the complete graph K_8 on 8 vertices. Then there are 105 points and 210 lines; each line has 3 points, and each point is on 6 lines.

Given two points, the two corresponding 1-factors on the 8 vertices produce a decomposition of the eight into disjoint polygons of even length. (A polygon of length two is regarded as two vertices connected by two edges – a small multigraph.)

Thus if the two points are equal, the decomposition is $2 + 2 + 2 + 2$, that is, four polygons of length two. If two points are distinct but collinear, the decomposition is $2 + 2 + 4$. For a symplectic pair, the decomposition is $2 + 6$, that is a doubled edge together with either a hexagon, or is $4 + 4$. In the former case, the convex closure is the 15-point quadrangle of order $(2, 2)$, while in the latter case, one obtains the 9-point grid of order $(2, 1)$. There are no special pairs. Finally, if the two 1-factors form a single octagon, the two points are at distance three.

The reader may check that a given point p of this near hexagon, is collinear with 12 further points, forms a $(2, 2)$-quad pair with 32 further points, forms a $(2, 1)$-quad pair with 12 points, and is at distance three from 48 points.

If every distance-two pair of points is a quad pair we say that $\Gamma = (\mathcal{P}, \mathcal{L})$ is a *near polygon with quads*. For example, any dual polar space of finite rank at least three is a near polygon with quads.

8.4 The Geometry of Quads

In this section we assume $\Gamma = (\mathcal{P}, \mathcal{L})$ to be a near polygon with quads. We discuss the relations between a point and a quad, a quad and a quad, and the diagram of the geometry of points, lines, and quads.

8.4.1 The Diagram

Let $\mathcal{Q}$ be the collection of all quads of Γ. Obviously, $(\mathcal{P}, \mathcal{L}, \mathcal{Q})$ is a rank three geometry where a point or line is considered to be incident with a quad if they are contained in it. Clearly Γ is the truncation of this rank three geometry to points and lines.

Consider a point p of Γ and let $\mathcal{L}_p$ and $\mathcal{Q}_p$ be the lines and quads which are incident with p. Then $(\mathcal{L}_p, \mathcal{Q}_p)$, with the induced incidence, is a rank two geometry, called the *residue at the point* p and denoted by $\mathrm{Res}_\Gamma(p)$. Now by (NP) lines are maximal singular subspaces of Γ, so two lines, say A and B on point p, lie in a unique quad Q which could be written as $Q(a, b)$ for any pair of points $(a, b) \in A \times B$ which are distinct from p. In other words, two distinct objects of $\mathcal{L}_p$ are always incident with a unique object of $\mathcal{Q}$. Thus $\mathrm{Res}_\Gamma(p)$ is a linear space.

Similarly, if L is a fixed line of Γ, then the residue $\mathrm{Res}_\Gamma(L) = (L, \mathcal{Q}_L)$ of points and quads incident with line L is a *digon* (all objects of one type are incident with those of the other type). This is because incidence is containment, and $\{p\} \subseteq L \subseteq Q \in \mathcal{Q}$ implies $p \in Q$.

Finally, if Q is a quad, then the points and lines in it form a generalized quadrangle as we already know. But that is still saying $\mathrm{Res}_\Gamma(Q)$ is a generalized quadrangle.

This information is assembled into a diagram for the geometry $(\mathcal{P}, \mathcal{L}, \mathcal{Q})$ in Fig. 8.1. Here the symbol "L" indicates that a point-residue is a linear space; the double edge on the first two nodes indicates that the points and lines incident with

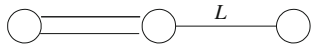

Fig. 8.1 The diagram of the geometry of points ($\mathcal{P}$), lines ($\mathcal{L}$), and quads ($\mathcal{Q}$) of an arbitrary near-polygon with quads

a quad form a generalized quadrangle, and the lack of a bond between points and quads (the extremal nodes) indicates that the point-quad residue of any line is a digon – i.e., quads are subspaces.

8.4.2 The Relations Between Quads

Suppose Q and R are distinct quads of the near polygon Γ. Now if $\{x, y\}$ were a 2-coclique contained in $Q \cap R$ then we would have $Q = Q(x, y) = R$, a contradiction. Thus $Q \cap R$ is a singular subspace. Hence the following applies.

Lemma 8.4.1 *If Q and R are distinct quads of a near polygon Γ, then $Q \cap R$ is either empty, is a single point, or is a line.*

Now suppose $c = (x_0, x_1, x_2, x_3, x_4, x_0)$ is a 5-circuit in the point-collinearity graph of Γ. We say that c is a *pentagon* if and only if x_i is not collinear with x_{i+2} (indices taken mod 5). Another useful feature is the following.

Lemma 8.4.2 *In any near polygon with thick lines for which quads exist, every pentagon lies in a quad.*

Proof This is an easy exercise. □

8.4.3 The Relations Among Points and Quads

Suppose p is a point and Q is a quad. If p is not incident with Q, then from the convexity of Q, $x^\perp \cap Q$ is a singular subspace. However, it cannot be a line since lines are maximal singular subspaces of a near polygon. Thus we see the following.

Lemma 8.4.3 *$x^\perp \cap Q$ is either a single point or is the empty set.*

Given a point p and a quad Q, the *distance from p to Q* is the minimal distance from p to a point of Q, and is denoted by $d(p, Q)$. Thus

$$d(p, Q) = \min\{d(p, q) | q \in Q\}.$$

If $d(p, Q) = 0$, p is one of the points of Q. If $d(p, Q) = 1$, then, as observed in the previous paragraph, there is exactly one point of Q collinear with p. But actually the situation is stronger than that: Q is actually strongly gated with respect to p.

Suppose $d(p, Q) = 1$ and set $\{q\} = p^\perp \cap Q$. Then clearly, $\Delta_1(q) \cap Q \subseteq \Delta_2(p) \cap Q$. Suppose there were a point y in $(Q - q^\perp) \cap \Delta_2(p)$. Then there would exist a point z in $y^\perp \cap p^\perp$. Then for any point $x \in y^\perp \cap q^\perp$ (neccessarily in Q by its

convexity) we would obtain a 5-circuit (p, q, x, y, z, p) and we argue now (in the present tense) as if that were the case.

Now by choice, y is not collinear with q. Since y is collinear with z, it follows that it cannot be that z is on the line pq, for then $z^\perp \cap Q$ would contain the coclique $\{q, y\}$, against Lemma 8.4.3. Thus z is collinear with p, but is not collinear with q. In particular z is not in Q at all, for otherwise it would be a coclique in $p^\perp \cap Q$, also against Lemma 8.4.3. From this it follows that z is not collinear with x, for otherwise z would lie in xy contained in Q. Thus the 5-circuit does not possess any diagonal adjacencies on (z, x) or (z, q). The absence of adjacencies at (p, x), (y, q), and (y, p) are forced by the choices of y and x. Thus c is a pentagon, as defined above. Now by Lemma 8.4.2 there is a quad R containing pentagon c. Then, as $R \cap Q$ contains the coclique $\{q, y\}$, one must have $Q = R$ by Lemma 8.4.1. But that would force $z \in Q$, which we know is impossible.

Thus the hypothesis that there exists a point of $\Delta_2(p) \cap Q$ not collinear to q is untenable, and so

$$\Delta_2(q) \cap Q = \Delta_3(p) \cap Q,$$

and so all distances of points of Q to p are those measured in Q to q plus one. Thus the following applies.

Lemma 8.4.4 *If p is any point at distance one from a quad Q, then Q is stongly gated with respect to p.*

Indeed, the following also applies.

Lemma 8.4.5 *Q is always strongly gated with respect to p whenever there is a unique point of Q nearest p.*

Proof Suppose $d = d(p, Q)$ and $\Delta_d(p) \cap Q = \{q\}$. Then all points of $Q \cap q^\perp$ lie in $\Delta^*_{d+1}(p) \cap Q$. If there were a further point y in this last set (so that $d(q, y) = 2$), then for any point $x \in y^\perp \cap q^\perp \subseteq Q$, by (NP), the line xy (which now contains two distinct points at distance $d + 1$ from p) would contain a point at distance d from p, against the hypothesis that q is the unique such point. Thus all points of $Q - q^\perp$ are at distance $d + 2$ from p, and so Q is strongly gated with respect to p. □

But of course, it could happen that there is no unique point of Q which is nearest p. What happens then?

In order to describe the phenomenon, we must introduce a special sort of geometric hyperplane of a generalized quadrangle. We say that a subset O of the set of points of a generalized quadrangle is an *ovoid* of the generalized quadrangle Q of and only if:

(O) No two points of O are collinear and every line of the quadrangle Q meets O at a single point.

Of course that means O is a geometric hyperplane of Q, but also, there are no "secant" lines of O, or lines contained within O – that is, O is a point-coclique which meets every line. We have already encountered many examples.

Example 1 Utilizing the $\bar{T}_6$-model of the 15-point quadrangle of order (2, 2), where the points are the 2-sets $\{a, b\}$ chosen from a 6-set $\Omega = \{1, 2, \ldots, 6\}$, one sees that the 5-coclique

$$O = \{\{1, y\} | 2 \leq y \leq 6\}$$

forms an ovoid in this quadrangle.

Example 2 Let H be the Hermitian variety determined by a non-degenerate σ-Hermitian form h on a vector space V of dimension four over a field F, $\sigma \neq 1_F$. Then the variety meets any non-degenerate 3-subspace W at a *unital* which is an ovoid of the generalized quadrangle obtained by the lines of the "ruling" on H (in plain language, the 2-spaces which are singular with respect to h).

Example 3 Let Q be the quadrangle of totally singular points and lines of a non-degenerate quadratic form on a 5-space V over a field of characteristic not 2. Let O be the singular 1-spaces belonging to a 4-subspace W, for which the quadratic form induces an *elliptic quadric* of type $O^-(4, F)$ of rank (Witt index) one. Then O is an ovoid of Q.

The next lemma gives a natural setting for ovoids.

Lemma 8.4.6 *Let $\Gamma = (\mathcal{P}, \mathcal{L})$ be a near polygon with thick lines and quads. Fix a quad Q and a point p with $d = d(p, Q) \geq 2$. Suppose there are at least two points in $\Delta_d(p) \cap Q$. Then $\Delta_d(p) \cap Q$ is an ovoid of Q, and consequently, $Q \subseteq \Delta^*_{d+1}(p)$.*

Proof Suppose u and v are distinct points of $D := \Delta_d(p) \cap Q$, where $d = d(p, Q)$. Clearly u and v cannot be collinear, for then (NP) would imply that the line uv must carry a point at a smaller distance $d - 1$ from p, contrary to our definition of d. Thus the set D is a coclique of Q.

We claim that it suffices to show that $Q \subseteq \Delta^*_{d+1}(p)$. Suppose that were the case. Then Q cannot possess any line external to D by the axiom (NP). Thus D is a geometric hyperplane of Q, and we are done.

Now $\Delta^*_{d+1}(p) \cap Q$ is a subspace of Q containing the two distinct geometric hyperplanes $u^\perp$ and $v^\perp$. Since Q has thick lines, each of these cones has a connected complement and so is a maximal subspace of Q. Thus $Q \subseteq \Delta^*_{d+1}(p)$, as desired. □

We may capture the whole section in the following.

Theorem 8.4.7 *Suppose p is a point and Q is a quad in a near polygon Γ all of whose lines are thick. Suppose p is not a point in the quad, so the distance $d = d(p, Q)$ is a positive integer. Then there are just two possible relationships between the point and the quad:*

1. *The quad Q is gated with respect to the point p.*
2. *The points of quad Q at distance d from Q form an ovoid of Q.*

In case $d = 1$, only the first alternative can occur.

8.5 A Characterization of Dual Polar Spaces

8.5.1 The Axioms

Throughout this section $\Gamma = (\mathcal{P}, \mathcal{L})$ is assumed to be a point-line geometry satisfying these axioms:

(DP1) Γ *is a near polygon with quads, all lines being thick.*
(DP2) *Each quad is gated.*

Note that as a subgraph of the point-collinearity graph of Γ, each quad is a convex induced subgraph and so is an isometrically embedded subgraph. This means there is no distinction between a quad being gated, and a quad being strongly gated (see Chap. 1 for the definitions).

Axiom (DP2) strengthens our previous hypotheses. As a result one has the theorem appearing just below. Recall that two lines are said to be *d-opposite* if there is a bijection between the lines matching each point of one line with its nearest point (at distance d) on the other line (see part 3 of Lemma 8.1.3 on p. 252).

Theorem 8.5.1 *(The existence of opposite lines when gated quads exist.) Suppose $\Gamma = (\mathcal{P}, \mathcal{L})$ is a near polygon satisfying the (DP)-axioms. Suppose $p = z_0, z_1, \ldots, z_d = r$ is a geodesic from point p to r of length d. Suppose in addition that L is a line on p carrying no point at distance $d-1$ from r. Then there exists a line L' on point r which is d-opposite L.*

Proof If $d = 0$ we have $r = p$ and the conclusion holds with $L' = L$, for L is 0-opposite itself. We now proceed by induction on d. Thus z_{d-1} lies on a line N which is $d-1$-opposite L. Then the quad $R = \langle N, z_d \rangle$ contains a line L' on z_d distinct from the line $z_d z_{d-1}$. Then, of course L' is 1-opposite N. Since R is gated with respect to p, with the exception of the point $z_d = r$, all other points of L' have distance $d+1$ from $z_0 = p$.

Now choose any point $x \in L - \{p\}$. Since N is $(d-1)$-opposite L, there is a unique point x_N on N at distance $d-1$ from x, and a point x' on L' collinear with x_N. It follows that $d(x, x') \leq d$. But since $d(x', p) = d+1$, we have $d(x, x') = d$ exactly. It now follows from the gatedness of R that the conditions of Lemma 8.1.3 are fulfilled, forcing L' to be d-opposite L. □

8.5.2 The Structure of the Residue of a Point

We begin with a simple result.

Lemma 8.5.2 *Suppose Q_1 and Q_2 are two distinct quads on a point p and that there is a point y not in $Q_1 \cup Q_2$ collinear with some point $g_i \in Q_i - p^{\perp}$, $i = 1, 2$. Then $d(y, p) = 3$ and the two quads intersect at a line $L = Q_1 \cap Q_2$.*

Proof By the last part of Theorem 8.4.7, each Q_i is gated with respect to y since $y^\perp \cap Q_i = \{g_i\}$ is a single point.[5] If $d(p, y)$ were 2, then, by the gatedness, each g_i would lie in $p^\perp$, contrary to assumption. Thus $d(p, y) = 3$.

Suppose, by way of contradiction, that $Q_1 \cap Q_2 = \{p\}$. Now $d(g_1, p) = 2 = d(g_1, g_2)$ (the last equality caused by the presence of y). Now, by (DP2), Q_2 is gated with respect to g_1, and so the equations just listed imply that g_1 is collinear with a unique point h of Q_2 and that g_2 and p both lie in $h^\perp$. Now by convexity of Q_1, h lies in Q_1 since it lies on a geodesic of length two connecting g_1 and p. We have shown that $Q_1 \cap Q_2$ carries the line hp contrary to assumption. The proof is complete. □

Now fix a point p. We have remarked earlier that the *residue of point* p – that is, the rank two geometry $\mathrm{Res}(p) = (\mathcal{L}_p, \mathcal{Q}_p)$ of all lines and quads incident with point p – forms a *linear space*. That was because any two lines on p determine a unique quad on p, namely the convex closure of the two lines.

Suppose z is any point distinct from p. If $d(p, z) = d$, $L(p; z)$ will denote the collection of lines on p which carry a point at distance $d - 1$ from z. Similarly, we let $Q(p; z)$ be the set of those quads on p which carry points at distance $d - 2$ from z, if $d \geq 2$, and if $d(p, z) = 1$, let it be the empty set. Then the incidence system $\mathbf{P}(p; z) := (L(p; z), Q(p; z))$ is defined for every point z distinct from p. Since $L(p; z) \subseteq \mathcal{L}_p$ and $Q(z) \subseteq \mathcal{Q}_p$, it is natural to notice the following.

Lemma 8.5.3 $\mathbf{P}(p; z) = (L(p; z), Q(p; z))$ *is a subspace of the linear space* $Res(p) = (\mathcal{L}_p, \mathcal{Q}_p)$.

Proof Suppose L_1 and L_2 are two distinct lines of $L(p; z)$ and let R be the quad which is their convex closure. Since R now contains p at distance d from z and contains the two distinct points in the L_i at distance $d - 1 = d(p, z) - 1$ from z, axiom (DP2) forces R to contain a unique point at distance $d - 2$ from z. Then clearly, every line of R on p belongs to $L(p; z)$ — that is, all "points" of the "line" R through "points" L_1 and L_2 of $\mathbf{P}(z)$ belong to $L(p; z)$. The proof is complete. □

Lemma 8.5.4 *If* $(p = z_0, z_1, \ldots, z_d)$ *is a geodesic (in the point-collinearity graph), then for* $i = 2, \ldots, d$, $L(p; z_i)$ *properly contains* $L(p; z_{i-1})$. *In fact, the latter is a geometric hyperplane of the former.*

Proof We need only prove the conclusion for $i = d$. First, by Theorem 8.5.1, there is line L on p which is $(d - 1)$-opposite the line $z_d z_{d-1}$. Then L lies in $L(p; z_d)$ but not in $L(p; , z_{d-1})$, so the latter is a proper subspace of the former.

We now show that $L(p; z_{d-1})$ is a geometric hyperplane of $L(p; z_d)$. Suppose $Q \in Q(p; z_d) - Q(p; z_{d-1})$. Then Q carries a point g at distance $d - 2$ from z_d. Now by assumption, z_{d-1} is not at distance $d - 3$ from any point of Q, and $d(z_{d-1}, p) = d - 1$ and $d(z_{d-1}, g)$ while g is not collinear with p. Then axiom

[5] Of course, the gatedness follows from axiom (DP2), but it is nice to observe that the axiom is really not needed at this stage of the proof.

(DP2) forces z to be distance $d-2$ from a point $h \in p^{\perp} \cap g^{\perp}$. Now we see that ph is a line of Q on p lying in $L(p; z_{d-1})$. Thus $L(p; z_{d-1})$ is a geometric hyperplane of $L(p; z_d)$. The proof is complete. □

Lemma 8.5.5 *Now fix a point z at distance three from p. Then $\mathbf{P}(p; z) := (L(p; z), Q(p; z))$ is a generalized projective plane.*

Proof First of all, the previous lemma shows that $\mathbf{P}(p; z)$ is a linear subspace of $\mathrm{Res}(p)$ which is not just a "line" that is, there are at least two distinct quads R and T in $Q(z)$. But then z is collinear with a point of each, while not lying in either. By Lemma 8.5.2, R and T meet at a line on p. It now follows that $\mathbf{P}(p; z)$ is a (generalized) projective plane.[6] □

Let $L(a; b)$ be the collection of lines on point a which carry a point closer to b than a is (this is consistent with the notation $L(p; z_i)$ introduced above). Thus if a has distance $d(a, b) = d$ from b, then each line of $L(a; b)$ carries a point of $\Delta_{d-1}(b)$.

We now wish to prove the following.

Lemma 8.5.6 *Suppose L is a line on p not lying in $L(p; z)$ for some point z at distance d, say, from p. Then there exists a point y collinear with z with $d(p, y) = d + 1$ such that $L(p; y)$ contains both L and $L(p; z)$.*

Proof By hypothesis the point of L nearest z is p itself, the remaining points of L having distance $d + 1$ from z. By Theorem 8.5.1, there is a line N on z which is d-opposite L. Now any point y of $N - \{z\}$ satisfies the conclusion of the theorem. The proof is complete. □

This lemma has an important consequence.

Theorem 8.5.7 *$Res(p) := (\mathcal{L}_p, \mathcal{Q}_p)$ is a projective space of projective dimension $n - 1$, where n is the maximal distance of a point from p in the point-collinearity graph. Every subspace of $Res(p)$ of projective rank $d - 1$ has the form $\mathbf{P}(p; z)$ for some point z for which $d(z, p) = d$.*

Proof Let us assemble a few observations:

1. We know that the lines and quads on point p form a linear space $\mathrm{Res}(p)$.
2. We know that if z is any further point, then the collection $L(p; z)$ of all lines on p carrying a point on a geodesic from p to z is actually a subspace of the linear space $\mathrm{Res}(p)$ (one that we have denoted as $\mathbf{P}(p; z)$) (Lemma 8.5.3, p. 271).
3. For every geodesic $(z = z_0, z_1, \ldots, z_{d-1}, z_d = p\}$ from z to p, the chain of subspaces

$$L(p; z_{d-1}) \subseteq L(p; z_{d-2}) \subseteq \cdots \subseteq L(p; z_0)$$

is a properly ascending unrefinable chain (Lemma 8.5.4).

[6] It might have thin "lines" if some quads are grids.

4. The linear space $\mathrm{Res}(p)$ is a generalized projective space, and so its poset of subspaces form a ranked poset (of flats) of a matroid (Lemma 8.5.5).
5. Let $\mathcal{S}(p)$ be the subposet of all subspaces of $\mathrm{Res}(p)$ of the form $L(p; z)$ for some z. We see that if L is a line on p (i.e., a "point" of $\mathrm{Res}(p)$), and $L(p; z)$ is a member of $\mathcal{S}(p)$ not containing L, then there exists another member of $\mathcal{S}(p)$ of projective rank one higher containing both L and $L(p; z)$ (Lemma 8.5.6).

Thus we see that $\mathcal{S}(p)$ is a subposet of a matroid with these two properties: (1) $\mathcal{S}(p)$ contains all rank one and rank two objects and (2) it is closed under taking joins with atoms. It follows that $\mathcal{S}(p)$ is the full matroid of all subspaces of the generalized projective space $\mathrm{Res}(p)$. Thus:

(MC) (Matroid completeness.) *If p is any point, then $Res(p)$ is a generalized projective space of rank n. Every single one of its subspaces (possessing some projective dimension d–1) has the form $L(p; z)$ where $d(p, z) = d$.*

The theorem is proved. □

Corollary 8.5.8 *Suppose L is a line bearing two distinct points u and v at distance d from point p. Then $L(p; u) = L(p; v)$.*

Proof The hypotheses force L to contain a point w at distance $d - 1$ from p. Then by Theorem 8.5.1 there is a line N on p which is $(d - 1)$-opposite L. Then N does not lie in $L(p; w)$ yet both of these sets of lines lie in both $L(p; u)$ and $L(p; v)$ with $L(p; w)$ being a hyperplane of each. It then follows from the preceding Theorem 8.5.7, that both $L(p; u)$ and $L(p; v)$ are the same unique subspace of $\mathrm{Res}(p)$ generated by $L(p; w)$ and N. □

8.5.3 The (Global) Convex Subspace Structure

Fix point p, and suppose z is a point at distance d from p. We have already seen that $\mathbf{P}(p; z) = (L(p; z), Q(p; z))$, the incidence system of lines and quads on p carrying points at distances $d - 1$ and $d - 2$ (resp.) from z, is a subspace of (projective) dimension $d - 1$ of the projective space $\mathrm{Res}(p)$.

In this section we require the following.

Lemma 8.5.9 *Suppose u and z are distinct collinear points of $\Delta_d(p)$. Which is also at distance d from p. Then $L(p; u) = L(p; z)$. Indeed this is true for all u in the connected component containing z of the induced subgraph $(\Delta_d(p), \sim)$.*

Proof The first statement is immediate from Corollary 8.5.8. The second follows from iteration of the conclusion of the first statement. □

Consider the following set of points:

$$S(p; z) := \{x \in \mathcal{P} | L(p; x) \subseteq L(p; z)\}.$$

Lemma 8.5.10 *The set $S(p; z)$ is a subspace of $\Gamma = (\mathcal{P}, \mathcal{L})$.*

Proof Let L be a line carrying two distinct points u and v of $S(p; z)$. Without loss of generality we may assume $d(p, u) \leq d(p, v) = e$. Then there is a unique point x of L at distance $e - 1$ from p and we have $L(p; x) \subset L(p, v)$. Then by Lemma 8.5.9, $L(p; y) = L(p; v)$ for all remaining points y of $L - \{x\}$. Since $L(p; v) \subseteq L(p; z)$ by hypothesis, we have $L \subseteq S(p; z)$. □

The next lemma is elementary.

Lemma 8.5.11 *Suppose a quad R on point p shares two lines of $L(p; z)$. Then $R \subseteq S(p; z)$.*

Proof The projective line of $L(p; z)$ containing the two given lines lie in a unique quad of $Q(p; z)$. The proof is complete. □

Theorem 8.5.12 *For each point-pair (p, z) for which $d(p, z) = d$, $S(p; z)$ is a convex subspace of Γ of point-diameter d. Each point of S is distance d from at least one further point of S.*

Proof We suppose $d(p, z) = d$. Replacing z by one of its neighbors at distance d from p if necessary, we may assume $d(q, z) = d - 1$ (Lemma 8.5.9). Let W be the set of lines of $S(p; z)$ which are incident with q. We first show that W is a subspace of $\mathrm{Res}(q)$. In fact we show something slightly stronger.

Step 1. *If L_1 and L_2 are two distinct lines of W, then the quad that is their convex closure lies entirely within $S(p; z)$.*

Let L_1 and L_2 be two distinct lines of $S(p; z)$ on point q and let R be the unique quad containing L_1 and L_2.

Observe that if $L_i \neq pq$ (this is true for at least one value of i), then each point $x_i \in L_i - \{q\}$ has distance two from p, and $L(p; x_i) \subseteq L(p; z)$. If $L(p; x_1) = L(p, x_2)$ (or equivalently if $p \in R$) then $L(p; x_1)$ comprises all lines of p in R, and all these belong to $L(p; z)$, so $R \subseteq S(p; z)$.

So assume $L(p; x_1) \neq L(p; x_2)$, which entails $p \notin R$. Then these two sets of lines are projective lines of the projective space $\mathbf{P}(p; z) = (L(p; z), Q(p; z))$ which meet at pq and so generate a plane π of $\mathbf{P}(p; z)$.

Now $d(p, q) = 1$ and $d(p, x_1) = d(p, x_2) = 2$. So, since $d(p, R) > 0$, there is a point y in ${x_1}^{\perp} \cap {x_2}^{\perp} \subseteq R$ at distance 3 from p and one now has

$$\langle L(p; x_1), L(p; x_2) \rangle_{\mathrm{Res}(p)} = \pi = L(p; y).$$

Now every point of $q^{\perp} \cap R$ belongs to $S(p; y)$. Also, since $L(p; y) \subseteq L(p; z)$ and $R - q^{\perp}$ displays a connected point-collinearity graph, we have $R - q^{\perp} \subseteq S(p; z)$ by Lemma 8.5.9. Thus $R \subseteq S(p; z)$. This completes Step 1 and our proof that W is a subspace of $\mathrm{Res}(q)$.

Now, since W is a subspace of $\mathrm{Res}(q)$, by Theorem 8.5.7, $W = L(q; v)$ for some point v. We shall show that v is in the subspace $S(p; z)$. More generally we claim the following statement.

Step 2. *If u is a point such that $L(q;u) \subseteq W$, the set of lines on q in subspace $S(p;z)$, then*

$$L(p;u) \subseteq L(p;z).$$

It suffices to prove this when $d(q,u) = e \geq 3$. Suppose this claim is false. Then there exists a line N on p lying along a geodesic to u which is not in $L(p;z)$. Clearly $N \neq pq$ so one can form the quad $Q = \langle q, N\rangle$. If $d(u,R) = d(q,u) - 2$, then all lines of R on q belong to $L(q;u) = W$ and hence R lies in $S(p;z)$ by Step 1. But this forces $N \in L(p;z)$, a contradiction.

So we must assume that the gate of R with respect to u is either q or a point h collinear with q. In the latter case $ph \in L(p;u) \subseteq W$ and $pq \in W$, so again $R \subseteq S(p;z)$ by Step 1. But then $R \in Q(p;z)$ so $N \in L(p;z)$, a return to the earlier contradiction.

We are left only with the possibility that the gate is q itself. Then $d(p;u) = d(q,u)+1$, and all points of line N distinct from p have points at distance $d(p,u)+1$ – that is, even further from u. This contradicts the assumption that N was along a geodesic to u.

Step 3. *If s is any point of the subspace $S(p;z)$, then $L(q;s) \subseteq W$.*

By way of contradiction assume the assertion of this step to be false. Then there exists a point $s \in S(p;z)$, a distance $e = d(q,s)$, and a line M on q meeting $\Delta_{e-1}(s)$ which does not belong to W. Let R be the quad on M and qp. We shall reach our contradiction by proving that $R \subseteq S(p;z)$, forcing $M \in W$.

Case 1. Assume $d(p,s) = e+1$. Then the unique point w of $\Delta_{e-1}(s) \cap M$ is the gate of R with respect to s. Then all lines of R on p belong to $L(p;s)$. But $L(p;s) \subseteq S(p;z)$ by the assumption on s. Thus all lines of R on p belong to $L(p;z)$. It follows that $R \subseteq S(p;z)$ by Lemma 8.5.11.

Case 2. $d(p,s) = e$ or $e-1$. Then the gate of R with respect to s is a point r at distance $e-2$ from s and pq contains a unique point p' in $r^{\perp} \cap pq$ at distance $e-1$ from s. If $p' = p$, then pq and ph are two lines of $L(p;z)$ in R. If $p' \neq p$ so $d(p,s) = e$, every line of R on p meets $h^{\perp}$, and so lie in $L(p;s) \subseteq L(p;z)$. In either case $R \subseteq S(p;z)$ by Lemma 8.5.11 once more.

Case 3. $d(p,s) < e-1$. This is impossible as $d(q,s) = e$ and $d(p,q) = 1$.

Step 4. *Let W be the collection of all lines of subspace $S(p;z)$ which are incident with point q. If x is any point such that $L(q;x) \subseteq W$, then $S(q;x) \subseteq S(p;z)$.*

If $v \in S(q;x)$ then $L(q;v) \subseteq S(q;x) \subseteq W$, so $v \in S(p;z)$ by Step 2.

Step 5. *Let $pq \in L(p;z)$ and $W = L(q;u)$ be defined as in the previous steps. Then $S(q,u) = S(p;z)$.*

Let V be the lines of $\mathrm{Res}(p)$ which are contained in the subspace $S(q;u)$. We claim $V = L(p;z)$.

First, $V \subseteq L(p;z)$. Clearly pq in V. If $N \in V - \{pq\}$, then the quad $\langle q, N\rangle$ belongs to $Q(q;u)$ (from the definition of V). Then all lines of R on point q belong to W by the definition of u. Then Step 1 forces $R \subseteq S(p;z)$, so $R \in Q(p;z)$ by Lemma 8.5.11. Thus $N \in L(p;z)$ as desired.

In the other direction, choose N in $V - \{pq\}$ and form R. Then R is a quad on q with at least two distinct lines on q belonging to $L(q;u)$. By Lemma 8.5.11 applied to $L(q;u)$, we have $R \subseteq S(q;u)$ and so N belongs to V.

Now the two containments $S(q;u) \subseteq S(p;z)$ and $S(p;z) \subseteq S(q;u)$ follow from Step 3 applied twice (the second time applying transpositions (p,q), (u,z), and (W,V) to the variables).

Step 6. *Set $S := S(p;z)$ and let $(S,\sim)$ be the point-collinearity graph of this subspace. For each point y of S, every geodesic path of $(\mathcal{P},\sim)$ from a point of S to y has all its intermediate vertices in S. Thus $(S,\sim)$ is a convex subgraph of $(\mathcal{P},\sim)$, the point-collinearity graph of Γ. Since S is a subspace of Γ, $(S,\sim)$ is also an induced subgraph. Thus $(S,\sim)$ is a convex isometrically embedded subgraph of $(\mathcal{P},\sim)$.*

First we note that by definition of $S = S(p;z)$, if $u \in S$ then $L(p;u) \subseteq L(p;z)$ and if w is any vertex belonging to a geodesic path from p to u, we have $L(p;w) \subseteq L(p;u)$ so $w \in S$. Thus p has the property desired of the point y of the conclusion. It follows that $(S,\sim)$ is a connected graph, since all its vertices are connected to p by some path in $(S,\sim)$. Let $(p = y_0, y_1, \ldots, y_k = y)$ be a shortest path in $(S,\sim)$ (not known in advance to be a geodesic of $(\mathcal{P},\sim)$). The proof proceeds by induction on k.

If $k = 0$, $y = p$ and the result holds. So we may assume $k > 0$. For $i = 0, 1, \ldots, k$ let W_i be the full collection of lines of S incident with point y_i. Then $W_0 = L(p, z_0)$ where $z_0 = z$. By Lemma 8.5.3, part 4, W_1 has the form $W_1 = L(y_1; z_1)$ for some point z_1. By Step 2, z_1 belongs to $S = S(p, z_0)$ and by Step 5, $S(y_1, z_1) = S(p = y_0, z_0) = S$.

The same sequence of lemmas and steps yields the argument that if, for some index i, $W_i = L(y_i; z_i)$ where $S(y_i, z_i) = S$, then there exists a point z_{i+1} such that $W_{i+1} = L(y_{i+1}; z_{i+1})$ and $S(y_{i+1}, z_{i+1}) = S$.

In the end, this allows us to write S in the form $S = S(y, z_k)$ from which we see that y has the desired property.

Step 7. *The conclusion of the theorem holds: $S := S(p;z)$ is a convex subspace of Γ.*

□

Corollary 8.5.13 *Suppose $d(p;z) = d$. For any point y in $S = S(p;z)$, the subset W_y of all lines of S on point p is a projective subspace of $Res(y)$ of projective rank $d-1$. S contains a point at distance d (as measured globally or in $(S,\sim)$ from y,*

but contains no point at any further distance. In particular the point-diameter of the convex subspace S is d.

Proof Because $(S, \sim)$ is connected, it suffices to prove this when y is collinear with p. We may assume S is larger than a quad. Here we have that W_y has the form $S(y; u)$ for a point u of S with $d(y, u) \geq 2$.

Let us call the subspaces $S(x, z)$ for which $d(x, z) = 3$ *hexagons*. Now one observes that the incidence system $(Q(py, S), H(py, S))$ of all quads and hexagons of S which lie on the line py simultaneously plays two roles: (1) the space of all lines and planes on a point in the projective space $L(y; u)$ and (2) the same thing for $L(p; z)$. It follows that W_y and $W_p = L(p; z)$ possess the same dimension as projective spaces.

The other remarks are just collected from the previous theorem by way of summary. □

We have encountered projective spaces $\mathbf{P}(p; z) = (L(p; z), Q(p; z))$ of dimension $d - 1$ where $d = d(p, z)$. We let $\mathbf{P}((p; z))^*$ be the "dual" projective space $(H(p; z), H_2(p; z))$ of all hyperplanes and co-dimension-2 subspaces of $\mathbf{P}(p; z)$.

Lemma 8.5.14 *Given points p and z at distance d in $(\mathcal{P}, \sim)$, there exists isomorphisms*

$$f_{p,z} : \mathbf{P}(p; z) \to \mathbf{P}((z; p))^*,$$
$$f_{z,p} : \mathbf{P}(z; p) \to \mathbf{P}(p; z)^*.$$

Proof We need only show the first one. Each L of $L(p; z)$ carries a point p_1 at distance $d-1$ from z, and this defines a hyperplane $f_{p,z}(L) := L(z; P_1)$ of $L(z; p)$, a projective space of rank d. Similarly, every quad Q of $Q(p; z)$ carries a point p_2 at distance $d - 2$ from z. Then $f_{p;z}(Q) = L(z; p_2)$ has codimension 2 in $L(z; p)$. (We have freely used Lemma 8.5.4 here.) Thus we have injective mappings

$$L(p; z) \to H(z; p),$$
$$Q(p; z) \to H_2(z; p),$$

which preserve incidence. But $\mathbf{P}(p; z)$ and $\mathbf{P}(z; p)$ are projective spaces of the same rank determined only by $d = d(p, z) = d(z, p)$. It follows that the mappings are surjective and together yield the isomorphism $f_{p,z}$. □

Now let $\mathcal{S}_d$ be the class of all convex subspaces of the form $S(p; z)$ where p and z are arbitrary points such that $d = d(p; z)$. By convention, $\mathcal{S}_0 = \mathcal{P}$, the set of points. $\mathcal{S}_1 = \mathcal{L}$, the set of all lines. Clearly $\mathcal{S}_2 = \mathcal{Q}$, the collection of all quads.

If X is a convex subspace of $\Gamma = (\mathcal{P}, \mathcal{L})$, we write $\mathcal{S}_d(X)$ for those elements of $\mathcal{S}_d$ which are contained in X. Finally, we write

$$\mathcal{S} := \cup_0^\infty \mathcal{S}_d.$$

Before going further, it might be nice to compare these classes of subspaces with the class of all convex subspaces (we don't really need this in what follows, but it is nice to know).

Theorem 8.5.15 *$\mathcal{S}$ is the class of all convex subspaces whose point-collinearity graphs have finite diameter.*

Proof Suppose X is a convex subspace of finite point-diameter d (that is, a subspace whose point-collinearity graph has finite diameter d). Then there exists a pair of points p and z in X with $d(p, z) = d$. Let X_p be the set of lines of X on p. Clearly X_p is a subspace of $\mathrm{Res}(p)$ and so has the form $L(p; u)$, but considerations of rank show that we can take $u = z$. Let x be any point of X. Then any geodesic from x to p finishes with an edge (y, x) with xy in $L(p; x)$. Thus by convexity of X, all the lines of $L(p; x)$ belong to $X_p = L(p, z)$. Thus $x \in S(p, z)$, and one concludes that $X \subseteq S(p; z)$. We leave it as an exercise to show that the convexity of X forces $X = S(p; z)$. □

Lemma 8.5.16 *The following hold:*

1. *Suppose $X \in \mathcal{S}_d$ and $Y \in \mathcal{S}_{d+1}$ and $X \subseteq Y$. Then for each point $z \in Y - X$, there is a unique point $\pi_X(z)$ of X which is collinear with z.*
2. *If X_1 and X_2 are two distinct elements $\mathcal{S}_d(Y)$ where $Y \in \mathcal{S}_{d+1}$, then either $X_1 \cap X_2 = \emptyset$ or else $X_1 \cap X_2 \in \mathcal{S}_{d-1}(Y)$.*
3. *Suppose $X_1, X_2 \in \mathcal{S}_d(Y)$, where $Y \in \mathcal{S}_{d+1}$ as in the previous part, but with the conclusion $X_1 \cap X_2 = \emptyset$. Then the restriction of the projection mapping $\pi_{X_1} : (Y - X_2) \to X_1$ of part 1 to X_2 is an isomorphism $\pi : X_2 \to X_1$ as point-line geometries.*
4. *There exists a poset isomorphism between the elements of $\mathcal{S}$ which contain point p and the subspaces of projective space $Res(p)$, obtained by replacing such a convex subspace on p with its lines on p. (Naturally this induces a bijection from the subspaces of $\mathcal{S}_d$ which contain the point p and the projective subspaces of $Res(p)$ of projective rank $d - 1$.)*

Proof Part 1. One has an element $p \in X$ and element y in Y with $d(p, y) = d$ and $S(p; y) = Y$. Now from the bijectivity of the mapping $f_{y,p}^{-1} : H(p; y) \to L(y; p)$ (a consequence of Lemma 8.5.14) X_p is a hyperplane arising as the lines on p of a subspace $S(p; x)$ where $d(p, x) = d - 1$, and x is collinear with y. Thus y is collinear with a point x of X. It cannot be collinear with further points of X since X is convex and Γ has no singular subspaces properly containing a line.

Part 2. Suppose $X_1 \cap X_2$ contains a point p. Then the set of lines on p in X_i forms a hyperplane X_{ip} of X_p, $i = 1, 2$. Then the intersection $X_1 \cap X_2$ is a convex subspace containing all lines on p in the intersection of the two hyperplanes. It is clear from Theorem 8.5.12 that $S_1 \cap S_2$ belongs to $\mathcal{S}_{d-2}$.

Part 3. From part 1 every element of X_2 is collinear with a unique elements of X_1 and vice versa. So there is a bijection

$$\pi : X_2 \to X_1,$$

carrying each point of X_2 to the unique point of X_1 with which it is collinear. Suppose a and b are distinct collinear points of X_2 and let $c := \pi(a)$, the unique point

of X_1 collinear with a. Then there is a quad $Q = \langle b, a, c\rangle$ which lies in Y. Then in $\mathrm{Res}(c)$ the lines belonging to Y form a projective d-space, and the lines belonging to Q and X_1 on c are respectively a projective line and a projective hyperplane of this subspace, and so share a common line on c, necessarily opposite the line ab in Q. Thus this line of $Q \cap X_1$ also carries $\pi(b)$. Thus collinearity of a and b implies collinearity of $\pi(a)$ and $\pi(b)$. The converse holds by symmetry of X_1 and X_2.

Part 4. If $S \in \mathcal{S}_d$ and p is in S, Corollary 8.5.13 asserts that the lines of S on p is a $(d-1)$-dimensional projective subspace $W(S)$ of $\mathrm{Res}(p)$. Conversely, $S(W) := \{y \in \mathcal{P} | S(p; y) \subseteq W\}$ is an element of $\mathcal{S}_d$. So we obtain a rank-preserving bijection

$$W : \mathcal{S}(p) \to PG(\mathrm{Res}(p)).$$

That the mapping preserves containment is patent from its definiton. □

8.5.4 The Dual Polar Space Structure in Near Polygons of Finite Diameter

Throughout this section, Γ has been a near polygon with quads whose quads are strongly gated. Of course what that abbreviated language means is that the point-collinearity subgraph of any quad is a strongly gated subgraph of the point-collinearity graph of the whole thing – that is of $(\mathcal{P}, \sim)$.

We now make an extra assumption:

(FD) *The point-collinearity graph $(\mathcal{P}, \sim)$ has finite diameter n.*

In this case (with a similar abuse of notation) we say that $\Gamma = (\mathcal{P}, \mathcal{L})$ *has point diameter n* when it actually means that (FD) holds for the point-collinearity graph.

Theorem 8.5.17 (P. Cameron.) *For a near polygon satisfying* (DP1), (DP2)*, and having point-diameter n, the incidence system $DP(\Gamma) := (\mathcal{S}_{n-1}, \mathcal{S}_{n-2})$ (where incidence is containment) is a non-degenerate polar space.*

Proof In order to keep the points and lines of the incidence system $\Gamma = (\mathcal{P}, \mathcal{L})$ separate from the "points" and "lines" of $DP(\Gamma)$ we shall enclose the latter in quotes, just as we have done in this sentence.

Let us select $P \in \mathcal{S}_{n-1}$ and $K \in \mathcal{S}_{n-2}$. If $K \cap P = \emptyset$, then there is a set $N(K) \cap P$ of all points of P which are collinear with a point of K. Since each such point is collinear with just one point of K, and no point of K is collinear with two points of P, one obtains a bijection:

$$\phi : N(K) \cap P \to K,$$

which preserves collinearity (Lemma 8.5.16 parts 1 and 3). Thus, if we select any point r of K, and take the unique line N on r meeting P, then the convex closure

$T := \langle N, K \rangle$ is an element of $\mathcal{S}_{n-1}$ containing K and $N(K) \cap P$. Then T is easily seen to be the unique member of $\mathcal{S}_{n-1}$ which contains K as well as $N(K) \cap P \in \mathcal{S}_{n-2}$ (Lemma 8.5.16 part 4.). That means P is "collinear" with a unique "point" of the "line" K in the point-line geometry $DP(\Gamma)$.

On the other hand, if $P \cap K$ is non-empty, then by Lemma 8.5.16, P meets each element Y of $\mathcal{S}_{n-1}$ containing K at an element of $\mathcal{S}_{n-2}$ – that is P is "collinear" with *all* "points" of the "line" K in $DP(\Gamma)$. It follows that $DP(\Gamma)$ satisfies the basic "one-or-all" axiom of a polar space.

But before we relax, we must show that $DP(\Gamma)$ is *non-degenerate*. That is, for every "point" $P \in \mathcal{S}_{n-1}$ there exists a "point" in $\mathcal{S}_{n-1}$ disjoint from P. It is quite simple. There exists a line L meeting P at a point y. Choose a point x in $L - P$. Now in the projective space $\mathrm{Res}(x)$ of projective dimension $n-1$, there is a hyperplane W not containing L. Then $S(W) := \{y \in \mathcal{P} | S(x; y) \subseteq W\}$ is a member of $\mathcal{S}_{n-1}(x)$. We claim $S(W) \cap P = \emptyset$. If this were not so, we should have $T := S(W) \cap P \in \mathcal{S}_{n-2}(S(W))$, and then x must be collinear to a unique point of T. But it is already collinear with a unique point of P, namely y. Thus $y \in T$. But that is impossible since $S(W) \cap x^{\perp}$ is the union of the lines of W, and by its choice, W does not contain $L = xy$. (We have used Lemma 8.5.16 parts 1, 3, and 4, freely in this argument.) Thus $R \cap P = \emptyset$ as desired. □

Corollary 8.5.18 (Cameron.) *Suppose $\Gamma = (\mathcal{P}, \mathcal{L})$ is a near polygon with thick lines and quads and finite point-diameter. If every quad is gated, then Γ is either a generalized quadrangle, or is the dual polar space of a known polar space.*

Proof Each dual polar point p represents a maximal singular subspace of the polar space $\Delta := DP(\Gamma)$. This is because all members of $\mathcal{S}_{n-1}$ which contain p pairwise intersect at a member of $\mathcal{S}_{n-2}$ according to the dichotomy of part 2 of Lemma 8.5.16. That this collection is even a subspace of Δ is an easy exercise.

The subspaces in $\mathcal{S}_{n-2}$ which contain a line $L \in \mathcal{L}$ also form a singular subspace of $DP(\Gamma)$ of rank one less. These correspondences preserve containment – that is incidence – and so we have an isomorphism of Γ with the dual polar space of the non-degenerate polar space $(\mathcal{S}_{n-1}, \mathcal{S}_{n-2})$. □

Of course, if the polar space $\Delta = (\mathcal{S}_{n-1}, \mathcal{S}_{n-2})$ has polar rank two, then Γ is its dual, also a quadrangle. If Γ is not a generalized quadrangle, Δ has polar rank at least three, and is known by the classification of polar spaces due to J. Tits [138] (whose proof was sketched in Sect. 7.10). □

Remark The special case of Corollary 8.5.18, in which Γ is a near hexagon, was first proved in [119].

8.5.5 Recapitulation

All of the theory of this section (showing that near polygons with finite point-diameter which satisfy (DP1) and (DP2) are generalized quadrangles or classical dual polar spaces) is due to Peter Cameron [24]. We have recast his wonderful paper

a little bit into what I hope are bite-size portions suitable for beginning students. The very last step shows that one has been dealing with a polar space all along, and from there the assertions that these are known objects implicitly invokes the classification of polar spaces partially covered in the previous chapter.

Cameron's theory is more general than what has been presented here in that it accomodates the existence of thin lines. That means we may now have quads whose point-collinearity graph has the form $K_{n,m}$, a complete bipartite graph. The result is still essentially the same: if quads exist and are gated, and the point-diameter is bounded, the near polygon comes from a dual polar space. See Cameron's paper [24] for details. The thin case has important implications. The polar spaces of type $O^+(2n, F)$ give dual polar spaces with thin lines.

The beauty of what has transpired in this chapter is that, up to the very last step, all arguments use only the (NP) and (DP) axioms and nothing else! That is, all arguments are self-contained and proceed by truly elementary steps. No sophistication is required of a student.

Finally, I remark that this chapter is about point-line geometries with all subspaces strongly gated (as graphs). But we only needed lines and quads strongly gated. *Buildings* (which are coming up in Part III) are chamber systems of type M for whom corank-one residues are strongly gated. From this it follows that *all* residues are gated. This is very reminiscent of what we found in dual polar spaces: we could conclude from the strong gatedness of quads and lines that *all* convex closures of two points are gated. Also on another level, the arguments eliciting characterizing properties of buildings from these strong gatedness assumptions on residues have the same elementary flavor as the arguments in near polygons, confined as they are to the interplay of the metric with strongly gated subgraphs. One cannot help wondering whether there is a common background context here. We shall explore this avenue a little further in the exercises.

8.6 Exercises for Chapter 8

8.1 Give a proof of Lemma 8.2.2.

8.2 Give a proof of Lemma 8.4.2.
The six classical finite generalized quadrangles are of these six types (here named after their classical groups):

1. Type $O^+(4, q)$ of order $(q, 1)$, the $(1 + q)$-by-$(1 + q)$ grid.
2. Type $O(5, q)$ of order (q, q).
3. Type $Sp(4, q)$ of order (q, q).
4. Type $U(4, q^2)$ of order (q^2, q).
5. Type $O^-(6, q)$ of order (q, q^2).
6. Type $U(5, q^2)$ of order (q^2, q^3).

The dual of 1 is the complete bipartite graph $K_{q+1,q+1}$. The dual of the quadrangles of type 2 are type 3, but they are isomorphic to their own duals if q is

even (see 1). The dual of 4 is type 5 (via the Klein correspondence). The dual of type 5 are quadrangles of order (q^3, q^2) but each has no embedding as a classical polar space.

8.3 Give a direct proof of the assertion that if q is even, then the quadrangle of type $O(5, q)$ is isomorphic to the quadrangle of type $Sp(4, q)$. [Hint: The quadratic form $Q : W \to \mathbf{GF}(q)$ has an associated symplectic form B with radical $R = \langle n \rangle$ where n is a non-singular vector. Show that each coset $w + R$ has just one singular vector, so there is a bijection between the singular points of $\mathbf{P}(W)$ and the 1-spaces of $V = W/R$. Show that totally singular lines of $\mathbf{P}(W)$ correspond to isotropic 2-spaces of V under the symplectic form induced by B.]

8.4 In this chapter we defined an ovoid of a generalized quadrangle as a collection O of pairwise non-collinear points such that each line meets it in a point. In other words, it is a geometric hyperplane which is a coclique of the point-collinearity graph. Show that if $(\mathcal{P}, \mathcal{L})$ is a generalized quadrangle of order (s, t) then any point-coclique of $1 + st$ points is an ovoid.

8.5 Show that there are "classical" examples of ovoids in many of these cases by showing:

1. An elliptic quadric of type $O^-(4, q)$ is an ovoid of the generalized quadrangle of type $O(5, q)$.
2. A unital of type $U(3, q^2)$ (the "classical unital") is an ovoid of the $U(4, q^2)$.

8.6 Show that in a generalized quadrangle of order (s, t), no ovoid can exist if t exceeds $s^2 - s$. Conclude as follows:

1. No quadrangle of type $O^-(6, q)$ possesses an ovoid.
2. Let Γ be a finite near polygon with quads, and assume that each quad has order (q, q^2) and is proper. Show that Γ is a dual polar space of type $U(2n, q^2)$, $n \geq 3$.

8.7 In a generalized quadrangle, a *hyperbolic line* is a coclique C of points with the property that any external point is collinear with 0, 1 or all of the points of C.

1. In the quadrangle of type $O(5, q)$, q odd, show that every 2-coclique is a maximal hyperbolic line, but is not a maximal coclique.
2. In the quadrangle of type $Sp(4, q)$ show that every 2-coclique lies in a hyperbolic line of size $1 + q$. Such a line is a maximal coclique. (Though not necessarily of maximal size.)
3. In the quadrangles of type $U(4, q^2)$ or $U(5, q^2)$ show that every 2-coclique belongs to a hyperbolic line of size $1 + q$.

8.8 Let U be a classical unital (the points of a non-degenerate 3-space of the Hermitian form (V, h), where V is a 4-space over $\mathbf{GF}(q^2)$. Show that the unital U is a "hyperbolic subspace" in the sense that any hyperbolic line meets U in 0, 1, or all its points. Show that U together with all the hyperbolic lines which it contains forms a 2-design of $q^3 + 1$ points, each point on q^2 lines of size $q + 1$. (Designs with these parameters are defined to be *unitals*.)

8.9 Now let H be a hyperbolic line of the generalized quadrangle G of type $U(4, q^2)$.

1. Show that $H^\perp$ is also a hyperbolic line.
2. If U is a classical unital containing H, show that $(U - H) + H^\perp$ is a non-classical ovoid of G.

8.10 Let G be as in Exercise 8.9. Suppose S is a set of $1 + q^2$ points of G, any three of which span a non-degenerate plane.

1. Show that any point outside this set is collinear with 0 or just 2 of the points of S.
2. The hyperbolic line on any two points of S meets S at just those two points.
3. Fix a point p in S. Let K be the point-set-theoretic union of all the hyperbolic lines on p and one other point of S. Show that K is a non-classical ovoid of G. [Hint: The ovoid part follows from the "0 or 2" condition in the first part. To show that it is non-classical, it is enough to show that it is a spanning set of the Hermitian 4-space. Such sets S really exist!]

8.11 The starting point of this exercise is the assumption that O is an ovoid of the classical quadrangle of type $Sp(4, q)$ where q is odd. The ultimate objective is to show that O does not exist (a result due to J. Thas), so ultimately we wish to reach a contradiction. Of course, one very efficient way to reach a contradiction is to make a mistake. We hope to "ovoid" that route (sorry about that almost pun).

Recall from Exercise 8.5 part 2, that G has a system of hyperbolic lines of size $q + 1$ forming a partial linear space (called a *copolar space* in the literature). Every pair of non-collinear distinct points of G lies in a unique hyperbolic line, and for any such line H, each point outside H is collinear to one or all points of H.

1. Show that each point of G lies on exactly q^2 hyperbolic lines, pairwise meeting at p.
2. Suppose H is a hyperbolic line meeting O at a positive number x of points. Show that $x = 2$, so that we may conclude that any hyperbolic line meets O at zero or two points. [Hint: Count collinear pairs $(u, y) \in (O - H) \times (H - O)$.]
3. Show that for every point p of G, the group $T(p)$ of transvections with center p (1) has order q, (2) acts semiregularly on $\mathcal{P} - p^\perp$, and (3) stabilizes every hyperbolic line on p.
4. Show that if $p \in O$, and $t \in T(p)$ with $t \neq 1$, then $O \cap O^t = \{p\}$.
5. Let $\mathcal{H}_p$ be the orbit $O^{T(p)}$ of all q conjugates of O under the action of the transvection group $T(q)$. These meet pairwise at p and form a so-called "fan" of ovoids. Let H be a hyperbolic line not in $p^\perp$ and not on p. Show that $H - p^\perp$ is even, and reach a contradiction about the parity of q.
6. Show that there exists a hyperbolic line H as in the previous step.

8.12 Suppose Γ is a dual polar space of type $Sp(2n, q)$ where n is at least 3, so Γ is not a quadrangle. Suppose O is now a geometric hyperplane of Γ which is also a point coclique. (The author does not know if such a thing can exist, but as you shall show in these exercises, it would be remarkable if it did.)

1. If there exists such a thing, then there exists one for $n = 3$ (which we assume for the remainder of this exercise).
2. Show that (since $n = 3$) $|O| = (q^2 + 1)(q^3 + 1)$.
3. Every quad Q (which is type $O(5, q)$) meets O at an ovoid of O.
4. Let $\mathcal{B}$ be the collection of point-sets $Q \cap O$ where Q is a quad. Show that $\Gamma := (O, \mathcal{B})$ is a partial linear space. [Hint: Quads are disjoint or meet at a line.]
5. Two distinct points of O are $\mathcal{B}$-collinear if and only if they are at distance two in the dual polar space $\Gamma = \Delta Sp(6, q)$.
6. Γ is a partial geometry with parameters $(s, t, \alpha) = (q^2, q^2 + q, q)$. [Hint: Use the fact that quads are strongly gated in Γ.]
7. There is a family of "planes" Π. Each plane π is a set of $q^2 + q + 1$ points of O which meet each block of $1 + q^2$ points at 0, 1 or $1 + q$ points. The points of π together with the "thick" block intersections forms a projective plane. [Hint: Consider the set $\pi = O \cap y^\perp$ where y is an exterior point. For each point $z \in y^\perp \cap O$, the $q + 1$ quads on the line yz meet O at blocks of $1 + q^2$ points each, of each of which y is collinear with $1 + q$ of them.]

8.13 Let ΔS be a dual polar space of a classical polar space S of finite rank defined by a σ-Hermitian or σ-quadratic form on a vector space V. Let p_1 and p_2 be two non-perpendicular singular or isotropic 1-spaces of V. (In case we are dealing with a (σ, ϵ)-pseudoquadratic form, "perpendicular" derives its meaning from the derived σ, ϵ-hermitian form B.) Thus $A := \langle p_1, p_2 \rangle$ is a non-degenerate 2-space, and we have a perpendicular direct sum decomposition $V = W \perp A$. Let X be the collection of all singular 1-subspaces of A, and let $\mathcal{M}(W)$ be the collection of all maximal singular subspaces of the subspace W. Since W is non-degenerate, $\mathcal{M}(W)$ is the set of points of a dual polar space $\Delta(W)$ of rank one less than ΔS.

1. For each $x \in X$, let $f_x : \Delta(W) \to \Delta S$ be the mapping which sends each maximal singular subspace N of W (point of $\Delta(W)$) to the singular subspace $\langle N, x \rangle$ of ΔS. Show that this mapping is an isomorphism of point-line geometries.
2. For distinct x and y in X, the subspaces $f_x(\Delta(W))$ and $f_y(\Delta(W))$ are disjoint. For each point $\langle N, x \rangle$ in the former, the only point in the latter it is collinear with is $\langle N, y \rangle$ for the same N.
3. For each $N \in \mathcal{M}(W)$, the sets $L_N = \{f_x(N) | x \in X\}$ are a system of transverse lines for the system of pairwise opposite polar spaces.

4. The subspace D of ΔS generated by the subspaces $f_x(\Delta(W)$ is just a product geometry $\Delta(W) \times X$ in the sense of Chap. 2.

8.14 Now let O be the coclique hyperplane of $\Delta Sp(6, q)$ of Exercise 8.12. Let $D = Q_1 + Q_2 + \cdots + Q_{q+1}$ be the subspace D of the previous exercise. That is, the Q_i are now quads of type $O(5, q)$, and there is a unique system of transverse lines $\{L_s | s \in Q_1\}$, such that each point of Q_i lies on a unique tranverse line, and the induced mappings $f_{ij} : Q_i \to Q_j$ produced by these lines are isomorphisms. (This follows from the "product geometry" observation at the end of the previous exercise.)

Set $O'_i = O \cap Q_i$, an ovoid of Q_i. Prove that for each $i > 1$, the isomorphism $f_{i,1} O_i \to Q_1$ defines an ovoid O_i of Q_1 and that the collection $\{O_i | 1 \leq i \leq q+1\}$ is a *spread of ovoids* – that is a family of ovoids partitioning the points of $O(5, q)$.

Remark If q is odd, such a remarkable structure cannot utilize even two of the classical ovoids provided by elliptic quadrics, since these must intersect at 1 or $q + 1$ points.

8.15 *(Identifying the stabilizer of a triangle of type .222 in the Leech lattice.)*

One of the uses of theorems that characterize geometries is to identify groups. Consider the following.

The Leech lattice Λ is an even integral lattice in $\mathbf{Z}^{(24)}$ defined by congruences mod 2 with certain elements of the Golay code over $\mathbf{GF}(2)$. For a simple and full description of this lattice, one is referred to the article by Conway [C] in the book *Simple Groups* by G. Higman and M. Powell (a book every serious student of groups and geometries should own). Other intersesting features – in particular the description by congruences of coordinates – are described in the book by Conway and Sloane [CS]. We do not need to know everything about this lattice in order to describe the problem.

As remarked above, all vectors of the lattice are 24-tuples of integers. First, some notation: The symbol $(a_1, a_2, a_3 | \pm a^\alpha, \pm b^\beta)$, where $\alpha + \beta = 21$, indicates the class of 24-tuples whose first three coordinates are (a_1, a_2, a_3) in that particular order, and among the remaining coordinates, α of them are either a or $-a$, and β of them are b or $-b$, and these may occur in any order. A vector v is *type* k if and only if the dot product $v \cdot v$ is $16k$. The set of such vectors in Λ with this norm is denoted Λ_k.

A triangle of type .222 would be be three vectors u, v, and $u - v$ in Λ_2 with pairwise inner product ± 16. We fix a .222 triangle T with vectors

$$\begin{aligned} u &= (4, 4, 0 | 0^{21}), \\ v &= (4, 0, 4 | 0^{21}), \\ u - v &= (0, 4, -4 | 0^{21}). \end{aligned}$$

The vectors nearest T in Λ_2 form an orbit $\mathcal{O}$ of 891 vectors x for which $(x, u) = (x, y) = 16$ and of course, $(x, u - v) = 0$, for which there is a subgroup $2^9 \cdot PSL(2, 4)$ which acts in four orbits as follows:

$$\begin{aligned} 1 \quad & (0, 4, 4|0^{21}) = \infty, \\ 42 \quad & (4, 0, 0| \pm 4, 0^{20}), \\ 21 \cdot 16 \quad & (2, 2, 2| \pm 2^5, 0^{16}), \text{ non-zero entries at an octad}, \\ 2^9 \quad & (3, 11| \pm 1^{21}). \end{aligned}$$

Now we consider the graph $(\mathcal{O}, \perp)$ where two members of $\mathcal{O}$ are adjacent if and only if they are perpendicular. Then the orbit of length 42 is all the set of all vectors of $\mathcal{O}$ which are perpendicular to ∞.

1. Show that even this set is a union of 21 connected components of size 2. Using the given information that Stab(T) is transitive on $\mathcal{O}$, conclude that triangles are maximal cliques of $(\mathcal{O}, \perp)$ which intersect pairwise in at most one points. Thus letting $\mathcal{T}$ be the set of all these triangles (mutually orthogonal triplets) we see that $S := (\mathcal{O}, \mathcal{T})$ is a partial linear gamma space.
2. Conclude from the action of stab$(T \cup \{\infty\}) \simeq 2^9 \cdot SL(2, 4)$ on these 42 vectors that stab(T) acts transitively on point-line flags of S.
3. Let $d(x, y)$ be the distance from x to y in the grpah $(\mathcal{O}, \perp)$. Show that for any two vectors x and y of $\mathcal{O}$:

   (a) $(x, y) = 0$ if and only x and y are adjacent in $(\mathcal{O}, \perp)$.
   (b) $(x, y) = 16$ if and only if $d(x, y) = 2$.
   (c) $(x, y) = 8$ if and only if $d(x, y) = 3$.

4. Show that S is a near hexagon with quads of type $O^-(6, 2)$ and order $(2, 4)$ (27 points and 45 lines).
5. Conclude from the above data that any pair of quads which share a point in fact share a line. From this $(\mathcal{O}, \mathcal{T})$ is a dual polar space of classical type.
6. Conclude that $(\mathcal{O}, \mathcal{T})$ is exactly the dual polar space of type $U(6, 2^2)$.

Since $\mathcal{O}$ is easily seen to span the vector space $\Lambda \otimes \mathbf{Q}$, the group H of type $2^9 SL(2, 4)$ fixing ∞ acts faithfully on $\mathcal{O}$. Thus the group stab(T) has order $891 \cdot 2^9 |Sl(2, 4)| = |U(6, 2^2)|$. On the other hand, since it is a group of automorphisms of the classical near polygon $\Delta U(6, 2^2)$, it is a subgroup of $U(6, 2^2)$. It follows that the stabilizer in Aut(Λ) of a .222 triangle, is $U(6, 2^2)$.

Remark The situation in part 3 where vertices of an association scheme are embedded in an orthogonal vector space so that the relations between points are signaled by the inner products of the representing vectors (that is, the "angles" between the 1-spaces) is a sort of "representation theory"

of association schemes discussed in the wonderful book *Distance Regular Graphs* by A. Brouwer, A. Cohen, and A. Neumeier [12].[7]

8.16 Complete the proof of Theorem 8.5.15, by showing that if Γ satisfies (DP1) and (DP2), then $S(p, z)$ is the convex subspace hull of the set $\{p, z\}$.

[7] I have placed the word "angles" in quotes, since this could be done in a finite vector space where the word "angle" does not assume its usual meaning.

Part III
Methodology

Chapter 9
Chamber Systems and Buildings

Abstract Chamber systems over a type set I were defined by J. Tits as a family of partitions of a set of vertices called chambers. An equivalent representation as certain classes of graphs with edges labeled by non-empty subsets of I allows one to describe morphisms, truncations, and residues graph-theoretically. Residual connectedness is defined for chamber systems. For a residually connected chamber system, each edge is labeled by exactly one type, while chamber systems in which each chamber lies on infinitely many distinct panels are not residually connected at all. Generalized polygons are presented as both chamber systems and point-line geometries in order to introduce chamber systems of type M. Buildings are plucked out of the sea of all chamber systems of type M by any one of six equivalent conditions involving strong-gatedness of residues, or galleries of reduced type.

9.1 Introduction

In the first part of this book, we undertook to understand the projective spaces and polar spaces as geometries of points and lines. But there are other more complex spaces in which projective and polar spaces are small constituent entities whose interaction with the space as a whole utilizes only modest properties of these subspaces. Because of the natural way they are defined, many of these new spaces may be called "classical."

We now know that the basic classical geometries of this world can be pretty well understood from the unifying point of view of a *building*. Such an assertion is vague, of course. But there are many favorite objects which occur over and over again under the guise of classical mathematics.[1] In a few words, these can be described as groups

[1] For example, Lie groups arise as a certain class of automorphisms of locally Euclidean manifolds. These are associated with Lie Algebras over the real or complex numbers. Then there are algebraic groups which are full automomorphism groups of an affine or projective variety. The definition is very different, but it is true that a non-degenerate algebraic group possesses a group-structure called a (B, N)-pair. Until the concept of building, the similarities between simple Lie groups and indecomposable algebraic groups hung on the frail sign that both were determined by a Dynkin diagram.

E.E. Shult, *Points and Lines*, Universitext, DOI 10.1007/978-3-642-15627-4_9,

© Springer-Verlag Berlin Heidelberg 2011

of Lie type acting on coset spaces of parabolic subgroups. The theory of buildings arose between 1964 and 1967 as an attempt to unite the objects of such a theory without assuming the presence of a group, a topology or a Lie Algebra. This great achievement is entirely due to Professor J. Tits, to whom this book is dedicated.

The projective spaces and polar spaces, together with an appropriate class of subspaces, are just two examples of buildings. But there are many more, associated with the exceptional groups of Lie types G_2, F_4, E_6, E_7, and E_8. All of these are generic schemes which can be defined over any field, finite or infinite. It is also important to note that these are geometries of rank equal to the subscript. Thus G_2 is a rank two geometry – that is, a point-line geometry (which happens to be a generalized hexagon). On the other hand, E_8 is a rank eight geometry – that is, an incidence geometry with 8 sorts of objects. I do not ask the reader to understand this all at once. Just let it be said that there is a very nice set of geometries that seems to come to us from the sky.

So, in order to continue, we must describe buildings. The traditional way of defining buildings involves axioms insuring that a chamber system C over I possesses a rather dense system of isometrically embedded Coxeter chamber systems (each isomorphic to $C(M)$ where M is a fixed Coxeter matrix) called *apartments*. Exacting properties of the building then rests heavily upon special properties of the Coxeter chamber system.

Of course many renditions merely need to *introduce* the topic in just enough detail to harvest the results. This is especially true of books on finite group theory [77, 78, 125, 126] or books and articles on groups of Lie type [27, 44]. Generally, they can live with the "apartment axioms" alone, or, at worst, with a few extra properties of Coxeter groups that they may need by presenting an abbreviated auxilliary section on Coxeter groups.

Then there is another class of books and articles which wish to present a theory of buildings which analytically displays and motivates the interlocking properties from a more or less pedagogical point of view. With only one exception (see last paragraph of this section), these books follow the format of Tits' original book: a development in the language of simplicial complexes leading to Coxeter complexes and finally the apartment axioms [14, 63, 106]. The idea is that the apartment axioms have a natural background in the language of simplicial complexes, foldings, etc., and one's understanding of concepts can be checked there.[2] These pedagogical treatments quoted here are excellent.

The basic paradigm is that algebraic properties concerning the factorizations of words that represent the same element of the Coxeter group can be used to deduce properties of the building.

In this book we take a different approach in which the basic properties of buildings are deduced directly from the strong-gatedness of their residues and nothing

[2] Of course Tits' book was probably intended for experts, but many early sections go to great lengths to give expositions about geometries, uniqueness of mappings, and pseudoquadratic forms that anyone can understand.

more than that. As a result, no special section about properties of Coxeter groups is required. This would then be a treatment independent of the type M hypothesis, and would not even once require mention of the unnecessary simplicial complexes which have nothing to do with any of the key properties.

There are reasons for producing an alternative view:

1. (A pedagogical one.) Some students (such as the author) who don't understand simplicial complexes, or think the whole topic a lot of unecessary baggage, shouldn't be prevented from being introduced to buildings. Everyone at least understands graphs (perhaps before they enter University) whether or not they understand simplicial complexes.
2. (More pedagogy.) The simpler the axioms one begins with, the easier it is to teach the subject and unfold the tree of implications. We don't want to start with something too elaborate.
3. (Pedagogy and boredom.) Why should a new customer always get the same menu?
4. (Theoretical economy.) Virtually every property of Coxeter groups that one assembles to prove properties of buildings can in fact be proved directly for all chamber systems of type M from a simpler hypothesis.
5. (More theoretical economy.) The approach given here entails a number of equivalent conditions never dispayed together in the cited expositions of this subject. Naturally that invites mathematical connections.

On p. 292 I mentioned that there was an expositional textbook which I thought to be an exception to the standard. That was the book by Mark Ronan which introduced buildings as a chamber system with a $W(M)$-valued measure. Looking back, I think it would be fair to say that Ronan's book and Rudolph Scharlau's historic insights more or less determined the path that I have taken here.

9.2 Chamber Systems

9.2.1 The Chamber System of a Geometry

Suppose $\Gamma = (V, \tau, *)$ is a geometry over the type set I. Recall that this means that $\tau : V \to I$ is a mapping assigning a *type* from the type-set I to each *object of* V. Then the symmetric binary relation $*$ on V occurs only between objects of different types. Thus $(V, *)$ is a multipartite graph (the *incidence graph of the geometry*) whose components are the (coclique) fibers of the type function τ.

A *flag* is simply any clique of $(V, *)$. One notes that the type function is injective when restricted to any flag. For any flag F, the subset $\tau(F) := \{\tau(x) | x \in F\}$ is called the *type of the flag* F. A *flag-chamber* of Γ is simply a flag whose type is I. Let $\mathcal{F}$ be the full collection of all flag-chambers of Γ. We are going to form a graph $C = (\mathcal{F}, E, \lambda)$ whose vertex set is $\mathcal{F}$ and whose edge set E is accompanied by an edge-labeling function $\lambda : E \to I$ as follows. Two flag-chambers form an

edge e of C labeled $\lambda(e)$ if and only if they are distinct flags which differ only in their constituent object of type $\lambda(e)$. In this case we say the flag-chambers are $\lambda(e)$-adjacent: that is, flag-chambers F_1 and F_2 are said to be i-adjacent if and only if they differ only in their constituent objects of type i. Thus distance in this chambersystem graph C is nothing more or less than the induced Hamming distance obtained when one regards C to be a subgraph of the Cartesian product over the type-fibers. Note that if F_1, F_2, and F_3 are three chamber flags such that (F_1, F_2) and (F_2, F_3) are edges of C bearing the same label, then (F_1, F_3) is forced to be an edge of C with the same label. Thus for any $X \in \{\mathcal{P}, \mathcal{L}\}$, the relation of being equal or i-adjacent is an equivalence relation on the vertices of C (the set of flags $\mathcal{F}$).

9.2.2 Abstract Chamber Systems

Well, that is the starting point of the notion of an abstract chamber system. An *(abstract) chamber system over I* is a triple (V, E, λ) where (V, E) is a simple graph, and λ is a mapping from the 2-subsets of V into the collection $F(I)$ of finite subsets of I with these properties:

(Support axiom.) $\lambda(e)$ is non-empty if and only if $e \in E$.
(Triangle axiom.) For every two edges (a, b) and (b, c) of E,

$$\lambda(a, b) \cap \lambda(b, c) \subseteq \lambda(a, c).$$

Fix two vertices x and y of V. We say that vertex x is *i-adjacent to y* if and only if $(x, y) \in E$ and i is one of the labels to be found marking this edge – i.e., $i \in \lambda(x, y)$. Then the triangle axiom shows us that i-adjacency together with the identity relation form an equivalence relation on the vertices of V.

The cardinality of the type set I is called the *rank of the chamber system*. Note that $\lambda(E)$ may very well be a proper subset of I.[3]

We have seen in Chap. 2 that, although geometries are just multipartite graphs, it is customary to use a special language for these graphs: vertices are called "objects," edges are "incident pairs of objects," and cliques are "flags." Similarly in chamber systems it is customary to call the vertices *chambers* and to refer to a path $G = (v_0, v_1, \ldots, v_n)$ as a *gallery*. The sequence $(\lambda(v_0, v_1), \lambda(v_1, v_2), \ldots, \lambda(v_{n-1}, v_n))$ is then the *type of the gallery G*.

Let us begin with some simple examples. Our examples are somewhat atypical in that in each underlying graph, each edge bears only one label chosen from $I = \{1, 2, 3\}$, while for a general chamber system, an edge could bear any non-empty set

[3] Again, note the contrast with geometries. For geometries, the rank was the number of types of objects $|\tau(V)|$ in the geometry. But in chamber systems the rank is the number of elements of I whether or not each type appears as an edge label or not. There is a simple reason for this distinction. Geometries and chamber systems will soon be seen to be connected by two functors, and the definitions of rank are geared so that both of these functors are rank-preserving.

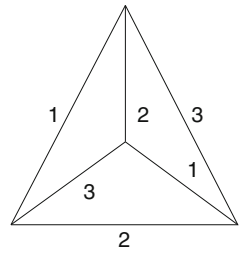

Fig. 9.1 A chamber system of rank three on the graph of the tetrahedron

of labels. Despite this special property, these examples will serve to illustrate many of the concepts which are to follow.

Example 1 (The edge-labeled tetrahedron.) Here the underlying graph is the graph of the vertices and edges of a tetrahedron – that is, the complete graph K_4. Notice that for each edge, there is a unique edge sharing no vertex with the former. In this way, the six edges are partitioned into three pairs, each pair forming a 1-factor of the graph.[4] The two edges of one 1-factor are all labeled "1," those of a second 1-factor are labeled "2," and those of the third are labeled "3." Thus each vertex of the tetrahedron lies on three edges bearing (one each) the three labels 1, 2, and 3. This example is illustrated in Fig. 9.1. We represent the labels attached to each edge by placing the label adjacent to the relevant edge in the figure.

Example 2 (The thin building of trigon type.) Here the underlying graph is the graph of the vertices and edges of the cube. In a normal embedding of the cube in the Euclidean plane, the edges come in three parallel classes, each class forming a 1-factor of the graph. We now label all the edges of one of these 1-factors by "1," the edges of the next 1-factor by "2," and the edges of the third 1-factor by "3." The result is the graph of Fig. 9.2a.

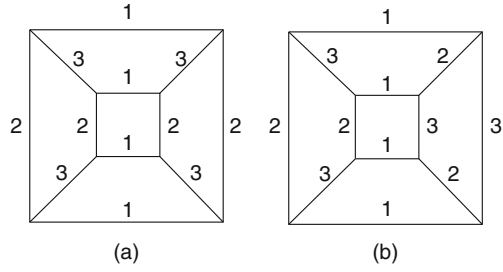

Fig. 9.2 Two chamber systems based on the graph of the cube

[4] This term "1-factor" was defined early in Chap. 8 in discussing an example of a near-polygon.

Example 3 (Another chamber system based on the cube.) Again we have the graph of the cube as our underlying graph, but this time the labels are attached differently as illustrated in Fig. 9.2b.

Example 4 (The shaved cube.) Imagine a perfect cube made of wood. Suppose one cut along each edge of the cube with a knife in a plane at a 135° angle with the planes of the two faces bordering the edge being cut. We do not cut deep enough to meet the plane of the cut in an edge opposite an edge in an original face of the cube. In other words, the shaving of the edges is rather slim. The result is a faceted figure with faces which are either squares or hexagons. In fact each edge of the cube is now a hexagonal face, while each face of the cube is still a square face, but one slightly smaller than the original. Our new graph has valence 3 and has 32 vertices.

Now there is (up to a permutation of I) just one way to label each edge by an element of I so that:

1. The three edges on a vertex bear distinct labels.
2. The edges of each external hexagonal or square face carry just two of the labels, alternating as one tours the face.

This chamber system is presented in Fig. 9.3.

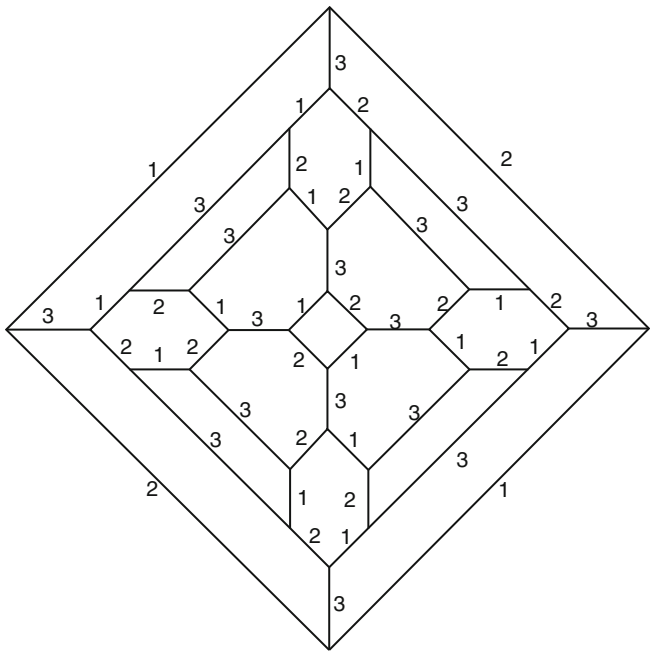

Fig. 9.3 The chamber system of the shaved cube

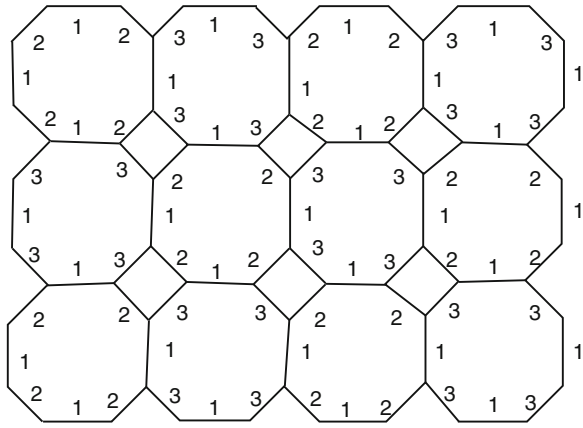

Fig. 9.4 A tiling of the Euclidean plane viewed as a chamber system

Example 5 (Tiling of the plane by squares and octagons: a thin building of type $\tilde{C}_2$.) This tiling of the Euclidean plane can be acheived as follows. At each integral point of the $\mathbf{Z} \times \mathbf{Z}$ lattice, one inserts a small diamond centered on the lattice point, with vertices pointed up and down and left and right. Each diamond should use $(1 - \sqrt{2}/2)$ of the line segment joining two lattice points. Now retain the middle $(\sqrt{2} - 1)$ of all the edges connecting adjacent lattice points. These segments are horizontal or vertical, connect vertices of adjacent diamonds, and are all labeled by "1." The diamonds alternate in their labeling. For adjacent diamonds, one has the label "2" on its northwest and southeast borders and "3" on its northeast and southwest borders, while it is the other way round for its adjacent partner diamond. This causes the octagonal faces to have their edges labeled by "1" and just one of the other labels. A portion of this infinite rank-three chamber system is given in Fig. 9.4.

Chamber Subsystems

Suppose $C := (C, E, \lambda)$ is a chamber system over I. For every subgraph $X := (X, E')$ of C (recall from Chap. 1 that this entails $X \subseteq C$, and $E' \subseteq E_X$, the subset of edges of E having both their vertices in X) we can define an edge-labeling $\lambda' : E' \to I$ which is just the restriction of λ to the edges of E'. Suppose, now that X is *triangle-closed* relative to (C, E), that is:

(ST) *If $\{a, b, c\}$ is a triangle of (C, E) (i.e., a 3-clique) with at least two of the three edges (a, b), (b, c), and (a, c) lying in E', then the third edge also lies in E'.*

Then the subgraph $X = (X, E')$ is converted into a chamber system $X = (X, E', \lambda')$ over J, where $J := \lambda(E')$. We call such a triangle-closed X a *label-induced chamber subsystem* of C. The adjective "label-induced" is there because all labels that are provided by the edge set E' are displayed.

More generally we call a chamber system (X, E', λ') over J a *chamber subsystem* of the chamber system $C = (C, E, \lambda)$ over I if and only if:

1. $X \subseteq C$ and $E' \subseteq E_X$ so (X, E') is a subgraph of (C, E).
2. For each edge $e \in E'$, $\lambda'(e) \subseteq \lambda(e)$; in particular $J = \lambda'(E') \subseteq \lambda(E) \subseteq I$.

Morphisms of Chamber Systems

Suppose $C := (C, E, \lambda)$ and $C' := (C', E', \lambda')$ are two chamber systems over I. A *morphism of chamber systems over* I is a morphism $f : C \to C'$ of the underlying graphs $(C, E) \to (C', E')$ such that for every edge $e = (c_1, c_2)$ of E, for which $f(e)$ is an edge of E' (rather than a single vertex), we have

$$\lambda(e) \subseteq \lambda'(f(e)).$$

Consider a morphism f of chamber systems over I:

$$f : C = (C, E, \lambda) \to C' := (C', E', \lambda').$$

We say that f is *chamber surjective (injective, bijective)* if and only if the underlying graph morphism is vertex surjective (injective, bijective, resp.). The morphism is said to be a *full morphism* if and only if:

1. Every edge $e' \in E'$ is the image of an edge e of E.
2. For any $e \in E$ for which $f(e)$ is an edge e' in E', one has $\lambda(e) = \lambda'(e')$.

A bijective full morphism is called an *isomorphism of chamber systems*. As one might expect, an isomorphism of a chamber system with itself is called an *automorphism of the chamber system*, and such automorphisms of a chamber system C form a group which we denote by $\mathrm{Aut}(C)$.

Let us consider automorphisms of our examples.

In Example 1, the underlying graph is K_4, the graph of the tetrahedron, and the automorphism group of this graph is the symmetric group on four letters, Sym(4). But as a chamber system, its automorphism group is seen to be the Klein four-group acting regularly on the four chambers.

In Example 2, the trigon, the full automorphism group of the graph of the cube is $\mathrm{Sym}(4) \times Z_2$ of order 48. But the automorphism of the chamber system is easily seen to be the largest normal 2-subgroup of this group: $K \times Z_2$ where K is the Klein four subgroup of the Sym(4)-direct factor. Note that K acts in two regular orbits on chambers.

In Example 5, each diamond (involving only the labels 2 and 3) is centered on a lattice point of the integral lattice $L = \mathbf{Z} \times \mathbf{Z}$. Let σ denote the translation of the Euclidean plane which maps each vector (a, b) to $(a, b) + (1, 1) = (a + 1, b + 1)$ – that is translation by the vector $(1, 1)$. Similarly, let τ denote translation by the vector $(2, 0)$. Then the reader can verify that both σ and τ are automorphisms of L which induce automorphisms of the chamber system C of this example.

Now let $A = \langle \sigma, \tau \rangle \leq \mathrm{Aut}(C)$. Then it is easy to see that A acts freely (that is semiregularly) on chambers and that any "fundamental region" contains just eight chambers. For example, if we take two diamonds, one immediately to the right of another, their union is a "fundamental region" – that is, a subset F of C (now regarded as the set of chambers) such that (1) $F \cap F^a = \emptyset$ for each non-identity element a of A and (2) $C = \cup\{F^a | a \in A\}$.

Any time we have a subgroup A of the automorphism group of a chamber system (C, E, λ) over I, we can define a chamber-surjective morphism from C to a chamber system $C/A := (C/A, E/A, \lambda/A)$ as follows. First, regarding C as the set of chambers (the vertex set of (C, E)), let C/A denote the collection of A-orbits on C. An edge of E/A will be a pair of distinct A-orbits (K, L) for which at least one member of K and one member of L form an edge of E. Then one sees that the mapping $C \to C/A$, which maps each chamber to the A-orbit to which it belongs, is a morphism of graphs $(C, E) \to (C/A, E/A)$ given as an example of a graph-morphism in Chap. 1, p. 16. Finally we define the set of labels $(\lambda/A)(K, L)$ to be attached to the edge (K, L) as the set

$$\cup\{\lambda(k, l) | (k, l) \in K \times L, (k, l) \in E\}.$$

(Note that since K and L are A-orbits, the set above could be written $\cup\{\lambda(x, y) | y \in x^{\perp} \cap L\}$ for any fixed element x of K.) It is easy to check that the triangle axiom holds for λ/A and so $C/A := (C/A, E/A, \lambda/A)$ is indeed a chamber system and that the canonical graph morphism is a morphism of chamber systems.

In Example 2, let A be the subgroup of the automorphism group generated by the involution α which takes each vertex to its opposite vertex in the cube. This mapping is easily seen to preserve the labels and so is an automorphism of the chamber system. The canonical projection morphism $C \to C/A$ produces a morphism onto a chamber system C/A isomorphic to that of Example 1.

In the exercises, the student will be asked to show that if A is the group of automorphisms of the chamber system C of Example 5 generated by the automorphisms σ and τ described above, then its canonical image C/A is isomorphic to the chamber system C' of Example 3, displayed in Fig. 9.2b.

A morphism $f : (C, E, \lambda) \to (C', E', \lambda')$ of chamber systems over I is called a *fibering of chamber systems* if and only if f is a full morphism for which the underlying graph morphism is a fibering. Recall from Chap. 1, p. 16, that the latter entails these two properties: (1) f is chamber surjective and (2) if c is a chamber, then f bijectively maps the edges of C on c to the edges of C' on $f(c)$. In particular, all vertices of a fiber $f^{-1}(c')$ are pairwise at distance at least three in (C, E). Notice that in a fibering, given any gallery $g' = (c'_0, \ldots, c'_n)$ of the image, and a chamber c_0 of the fiber $f^{-1}(c'_0)$, there exists a unique lift $g = (c_0, \ldots, c_n)$ which is a gallery in C of the same type.

The canonical projection morphisms $C \to C/A$ from the chamber system of Example 2 onto the chamber system of Example 1 ($A = \langle \alpha \rangle$), and from Example 5 onto the chamber system of Example 3 ($A = \langle \sigma\tau \rangle$), are examples of fiberings.

The Functor C from Geometries to Chamber Systems

It is worth remarking that chamber systems over I (like geometries over I) form a category with respect to the chamber system morphisms (the label-preserving graph homomorphisms) that we have just defined. That means we have identity mappings and we can compose morphisms when the domains and ranges allow it.

Let $\Gamma = (V, E, \tau)$ be a geometry over I, where we suppose $I = \tau(V)$ is the full set of types realized by objects of the geometry. Our construction mentioned ealier produced a chamber system $\mathbf{C}(\Gamma)$ over I whose chambers are the flag-chambers of Γ, the set of flags of type I, two of them being i-adjacent if and only if they differ only by an object of type i. Note that if no chamber flags exist, there may be no chambers at all. Nonetheless, this empty set of chambers is considered to be a chamber system over I. Also, it is clear that if i and j are distinct types, then two flag-chambers cannot be both i- and j-adjacent at the same time. That means that in $\mathbf{C}(\Gamma)$, the labeling function λ assumes only singleton sets of I. Thus $\mathbf{C}(\Gamma)$ is not entirely a typical abstract chamber system.

Now consider a morphism $f : \Gamma \to \Gamma'$ of geometries over I, where I and Γ are as in the preceding paragraph. Then the geometry Γ' also has an object of each type in I. In addition, if F is a flag chamber of Γ, then $f(F)$ is a flag chamber of Γ'. Moreover, if F_1 and F_2 are i-adjacent flag chambers of Γ, then either $f(F_1) = f(F_2)$ (which happens if the objects of type i by which the F_i differ are mapped to a common chamber) or else $f(F_1)$ is i-adjacent to $f(F_2)$. Thus we see that the geometry morphism f induces a chamber system morphism $\mathbf{C}(f) : \mathbf{C}(\Gamma) \to \mathbf{C}(\Gamma')$. One verifies that if $f_1 \circ f_2$ is the composition of two morphisms of geometries over I, then $\mathbf{C}(f_1 \circ f_2) = \mathbf{C}(f_1) \circ \mathbf{C}(f_2)$. Also, if f is an identity mapping, then so is $\mathbf{C}(f)$.

Thus our construction of $\mathbf{C}(\Gamma)$ from Γ in fact provides a functor $\mathbf{C}$ from the category of geometries over I into the category of chamber systems over I.

Residues of a Chamber System

Let $C = (V, E, \lambda)$ be a chamber system over I. For each subset J of I, let E_J be the set of edges e of E for which $\lambda(e) \cap J$ is non-empty – that is, those edges whose non-empty set of labels contains at least one element of J. Each connected component $R := (W, E_W \cap E_J)$ of the graph (V, E_J) is called a *residue of type J of the chamber system C*.[5] The set $I - J$ is called the *cotype* of the residue. The cardinality of J and $I - J$ of a residue of type J is called the *rank* and *corank* of the residue $R := (W, E_W \cap E_J)$, respectively. Note that R may be regarded as a chamber system over J in its own right. Of couse, R is a chamber system over the possibly smaller set $\lambda(E_W \cap E_J)$, but is nevertheless called a residue of type J and taken to be a chamber system over J. A residue of rank one is just an i-adjacency

[5] The notation here is important. Since W is a subset of the vertex set V, the symbol E_W refers to the edges which have both their vertices in W (see Chap. 1). But when J is a subset of the index set I, E_J is the set of edges which bear a label in J.

equivalence class for some type i. Such residues of rank one are often called *panels*. Let us consider residues in our examples. In Example 1, for each type $i \in \{1, 2, 3\}$, there are two residues of type i, namely two edges labeled i forming a 1-factor of the tetrahedron. In fact, something like this is true for residues of type i in all of our examples: the residues of type i form a 1-factor of the graph.

In Example 1, for every 2-subset $\{i, j\}$ of I, the residue of type $\{i, j\}$ is unique and is a spanning subgraph which is a square. Note that these subgraphs are not induced subgraphs.

In Example 2, the subgraph on V obtained by using only edges which bear labels i or j ($i \neq j$) has two connected components each of which is an induced subgraph which is a square; they are opposite faces of the cube.

In Example 3, there are again two residues of type $\{2, 3\}$ representing opposite faces of the cube. But here, the residues of type $\{1, 3\}$ and type $\{1, 2\}$ are spanning subgraphs which are octagons. Clearly neither is an induced subgraph. Moreover, these two residues intersect at the disconnected subgraph consisting of the 1-factor of all edges labeled '1.'

In Example 4, the residues of rank two correspond to the faces of the "shaved cube": 6 are squares (corresponding to the original faces of the cube) and 12 are hexagons (corresponding to the 12 edges of the original cube). The residues of type $\{i, j\}$, a 2-subset of $I = \{1, 2, 3\}$, consist of exactly two squares (corresponding to opposite faces of the original cube) and four are hexagons (corresponding to the parallel class of edges connecting these two faces in the original cube). Thus we see that the rank two residues of a given type do not belong to a fixed isomorphism type in this example. However all rank two residues are induced subgraphs and intersect pairwise at the empty set or a single rank one residue.

In Example 5, all residues of type $\{2, 3\}$ are squares; these are the "diamonds" which were centered over each integral lattice point. Otherwise the residues of type $\{i, j\}$ are octagons. All are induced graphs with any two meeting at the empty set or a rank one residue.

The Functor Γ

Now one can define a geometry $\Gamma(C)$ over I as follows. The objects of type i in $\Gamma(C)$ are exactly the residues of C of cotype i – that is, the connected components of $(V, E_{I-\{i\}})$. Two objects (corank one residues) are incident in this geometry if and only if they contain a common chamber. Note that two distinct objects of the same type cannot be incident since these are chamber-disjoint connected components of the same subgraph.

Remark Then Γ is a functor from the category of chamber systems over I to the category of geometries. It isn't really important, but it does allow us to sling this word "functor" around, and that is a lot of fun.

Let us examine the geometries $\Gamma(C)$ that are obtained as C ranges over our five examples. In Example 1, there is exactly one rank two residue of each type. Since these are the residues of corank one, we obtain a geometry with exactly three

objects, one of each type. The objects are pairwise incident and so together they form the unique flag chamber of the geometry. Thus $\mathbf{\Gamma}(C)$ is the graph K_3 (viewed as a tripartite graph). Then $\mathbf{C}(\mathbf{\Gamma}(C))$ is still a chamber system over $I = \{1, 2, 3\}$ but it has only one chamber.

In Example 2 there are exactly two residues of corank one of each type, and such residues of different types always intersect in a residue of rank one. Thus the geometry $\mathbf{\Gamma}(C)$ is the complete tripartite graph $K_{2,2,2}$ with each component 2-coclique being the objects of a specific type.

In Example 3 the geometry $\mathbf{\Gamma}(C)$ is the complete tripartite graph $K_{2,1,1}$ where the unique 2-coclique represents the set of objects of type 1, and of the other two objects, one is type 2 and the other is type 3.

In Example 4 the geometry $\mathbf{\Gamma}(C)$ contains six objects of each type. The graph is clearly not a complete tripartite graph since there is a residue of type $\{1, 2\}$ meeting a residue of type $\{1, 3\}$ at the empty set. We leave it as an exercise for the student to work out the complete incidence graph.

In Example 5 the geometry $\mathbf{\Gamma}(C)$ has three sorts of objects: the diamonds, and the two types of suboctagons. If D is a diamond (type 1 object) it is incident in this geometry with exactly two objects of type 2 and two objects of type 3 and its residue in this geometry (the subgraph $D^{\perp} - \{D\}$) is the complete subgraph $K_{2,2}$, a subgeometry over $\{2, 3\}$.

Truncations of a Chamber System

Let $C = (V, E, \lambda)$ be a chamber system over I. Fix a subset J of I and set $K = I - J$. Let C/J denote the collection of all residues of C of type J (these are components of a partition of C). For each such pair, (A, B) of residues of type J, let $\lambda^K(A, B)$ denote the set of all types $k \in K$ for which $A \cup B$ lies in a residue of type $J \cup \{k\}$ (k-adjacency). Let E^K be the collection of distinct pairs (A, B) of residues of type J such that $\lambda^K(A, B) \neq \emptyset$. Then $C_K := (C/J, E^K, \lambda^K)$ is a chamber system over K, which we call the *truncation of C of type K*.

Remark Note that if C is connected, so is C_K. There is a simple reason for this. The mapping which takes each chamber c of C to the unique residue of type J which contains it, say $R_J(c)$, is a vertex surjective morphism of the underlying graphs of the two chamber systems.

Lemma 9.2.1 (Functorial properties of truncation.)

1. *If $f : C' \to C$ is a morphism of chamber systems over I and $K \subseteq I$, then there is an induced morphism $f_K : C'_K \to C_K$ as chamber systems over K in which, for each $R \in C'/(I - K)$, $f_K(R)$ is the unique residue of C of type K containing the connected set of chambers $f(R)$.*
2. *Suppose C is a chamber system over I and $K \subseteq I$. Then there is an isomorphism*

$$(\mathbf{\Gamma}(C))_K \to \mathbf{\Gamma}(C_K).$$

Chamber Systems Defined by Cosets

This is a simple construction which will help us describe the spherical buildings that are soon to arrive on the scene. Let G be a group and let B be a distinguished subgroup. Let $\mathcal{H} := \{H_1, H_2, \ldots, H_n\}$ be a collection of subgroups of G indexed by I, where each subgroup of $\mathcal{H}$ contains B. Just from this data, we may define a chamber system $C := C(G, B; \mathcal{H})$ whose chambers (i.e., vertices) are the right cosets G/B of the subgroup B. Coset Bg is i-adjacent to vertex Bh if and only both of these cosets lie in a common coset of H_i (this would be $H_i g = H_i h$). Thus i-adjacency (together with the identity relation) forms an equivalence relation, and so a chamber system over $I = \{1, \ldots, n\}$ is the result. Such a chamber system is called a *coset chamber system*.[6] (Note that, in this example, if the groups H_i do not meet pairwise at B, it is possible for an edge of the chamber system to bear multiple labels.) The notation is important: the first two key groups G and B appear first, followed by a distinguished semi-colon. Then the rest.

9.2.3 Residually Connected Chamber Systems

At the beginning of Sect. 9.2 we saw that if Γ is a geometry over I, then there is a well-defined chamber system of flag-chambers $\mathbf{C}(\Gamma)$ without multiply-labeled edges. Of course, if there are no flag chambers, $\mathbf{C}(\Gamma)$ could very well be a meagre landscape.

Conversely, given a chamber system C over I as given earlier Sect. 9.2, there is a geometry $\mathbf{\Gamma}(C)$ whose objects of type i are the residues of C of cotype i. Two such residues are incident if and only if they share a common chamber.

We have remarked that these constructions produce (1) a very natural functor $\mathbf{\Gamma}$ between the category of chamber systems and the category of geometries and (2) another functor $\mathbf{C}$ from the category of geometries over I to the category of chamber systems over I. Now there is a property called residual connectedness for geometries, and another notion by the same name for chamber systems which was first graph-theoretically formulated by Arjeh Cohen. When either one of these properties is present, the functors are inverses of each other. The theorem is this.

Theorem 9.2.2 (Arjeh Cohen, [10].)

1. *If geometry G is residually connected of finite rank, then so is $\mathbf{C}(G)$, and there is a geometry isomorphism $\mathbf{\Gamma}(\mathbf{C}(G)) \simeq G$.*
2. *If C is a residually connected chamber system, then $\mathbf{\Gamma}(C)$ is residually connected, and $\mathbf{C}(\mathbf{\Gamma}(C)) \simeq C$.*
3. *There exists an isomorphism between the subcategory of residually connected geometries over a finite typeset I, and the subcategory of residually connected chamber systems over the same finite I.*

[6] In analogy with the term "coset geometry."

In other words, the residually connected objects on both sides form subcategories where the restricted functor yields an isomorphism. However, it turns out that the conditions for residual connectedness can never be realized when the typeset is infinite and the edges leaving any chamber realize all types. So this is primarily a concept that will kick in only for chamber systems and geometries of finite rank. It will turn out that buildings of finite rank in fact possess this property – thus allowing us to speak of buildings as either geometries or chamber systems in the case of finite rank.

Residual Connectedness of Geometries, Revisited

Let $\Gamma = (V, E, \tau)$ be a geomety over $I := \tau(V)$. In Chap. 2, we said that Γ was a *residually connected geometry* if and only if:

(RC1) Every flag of corank one has a non-empty residue (that is, it lies in a flag chamber).
(RC2) For every flag F of corank at least 2 (including the empty flag if appropriate), the residue $\mathrm{Res}_\Gamma(F)$ is a non-empty connected geometry.

In Exercise 3 of Chap. 2, we showed that this condition implies:

- Every residue of Γ is residually connected (Lemma 2.5.1 of Chap. 2).
- If the rank $|I|$ is finite, Γ is chamber connected (part 2 of Lemma 2.5.1. Exercise 3 of Chap. 2 is to show that this fails for infinite rank).
- If the rank $|I|$ is finite, then every truncation of rank at least two is residually connected (Corollary 2.5.3 of Chap. 2).

Residual Connectedness for Chamber Systems

Let $C = (C, E; \lambda)$ be a chamber system over I. The chamber system C is said to be *residually connected* if and only if:

(CRC1) For any family $\mathcal{F} = \{R_t\}$ of residues of C which intersect pairwise non-trivially, the global intersection $\cap\{R_t \in \mathcal{F}\}$ is non-empty and connected.
(CRC2) For any chamber c the intersection of all corank one residues of C which contain c is the set $\{c\}$ itself.

An immediate consequence of this property is recorded in the following.

Lemma 9.2.3 *If C is a residually connected chamber system over I, and J is a proper subset of I, then any residue of type J is the intersection of all the corank one residues which contain it.*

The proof of this lemma is left as Exercise 9.8 at the end of this chapter.

Theorem 9.2.4 *Assume $C = (C, E; \lambda)$ is a residually connected chamber system over I. Suppose $e = (x, y)$ is an edge bearing the label i – that is $e \in E$ and*

$i \in \lambda(e)$. Suppose $G = (x = x_0, x_1, \ldots, x_n = y)$ is any gallery connecting to x to y. Then for some integer j in the interval $[1, n]$, we have

$$\lambda(x_{j-1}, x_j) = \{i\}.$$

In particular:

1. *The type function λ never assumes multiple values – that is, for every edge $e \in E$, $\lambda(e)$ is a single-element subset I.*
2. *Each residue of cotype i is an induced subgraph of (C, E).*
3. *All residues are induced subgraphs.*

Proof Take C, e, $i \in \lambda(e)$, and G as in the hypothesis. Suppose each edge of G bears a label in $I - \{i\}$, and so x and y belong in the same residue of cotype i. On the other hand, since (x, y) is itself an edge bearing label i, x, and y lie together in a common residue of cotype t for each $t \in I$ distinct from i. It follows that every residue of corank one which contains x also contains y, and that contradicts (CRC2). So for some edge e of the gallery G, we must have $\lambda(e) = \{i\}$, as required.

The remaining three statements are immediate consequences of the first. If $\lambda(e)$ contained two distinct type-labels, say i and j, then the first statement would be contradicted with $G = (x, y)$. Thus λ cannot assume multiple values.

Suppose, for a residue R of cotype i, $(R, E_R \cap E_{I-\{i\}})$ (which by definition is connected) was not an induced graph. Then there must exist a pair of chambers, x and y in R, such that $e = (x, y)$ is an edge bearing no label of $I - \{i\}$ – forcing $\lambda(x, y) = \{i\}$. But since the graph for R is connected, there is a gallery G connecting x to y, each of whose edges bears a label not equal to i. That contradicts the first statement of the theorem. One can only conclude that the underlying graph of R is an induced graph.

Now statement 3 of the theorem follows from Lemma 9.2.3 upon observing than any arbitrary intersection of induced graphs is an induced graph. □

Residual Connectedness in Infinite Rank

The next theorem reveals that the residual connectedness property does not normally hold for a chamber system of infinite rank.

We say that a chamber system C over I is *firm* if and only if each panel contains at least two chambers. This could be stated another way: *for each chamber c and each type i, there exists a further chamber c' which is i-adjacent to c.*

Theorem 9.2.5 (Kasikova and Shult [84].) *No firm chamber system of infinite rank is residually connected.*

Proof Assume by way of contradiction that C is a residually connected firm chamber system of infinite rank $|I|$. Then each vertex c is i-adjacent to at least one other vertex, for each $i \in I$, and by Theorem 9.2.4, each edge bears a unique type label.

Thus any gallery $G = (c_0, c_1, \ldots)$ has a type $(t_1, t_2, \ldots)$ associated with it, where $t_i := \lambda(c_{i-1}, c_i)$. (The gallery can be finite or countably infinite.) We are particularly interested in those galleries in which *no type is repeated* – i.e., $t_i \neq t_j$ if $0 \leq i < j$. It is an easy consequence of firmness and infinite rank that any such gallery of finite length can be extended. Thus the following result holds:

Step 1. *There exists an infinite gallery* $G = (c_0, c_1, \ldots)$ *of type* $(t_1, t_2, \ldots)$ *where no type is repeated in the sequence* $\{t_i := \lambda(c_{i-1}, c_i) | i = 1, 2, \ldots\}$.

Now, for each chamber c_i in the gallery G, let $R_i := R_{I-\{i\}}(c_i)$ be the unique residue of cotype t_i containing c_i. Since the types t_j, $j > i$ are pairwise distinct, the residue R_i is forced to contain all of the chambers c_k, for all $k \geq i$. But is it possible for R_i to contain any of the earlier chambers in the sequence? The answer appears in the following step.

Step 2. *The residue* R_i *contains no chamber* c_j *of* G, *for* $j < i$.

Suppose by way of contradiction that $c_j \in R_i$ for $j < i$. Then the chambers $c_{j+1}, \ldots c_{i-1}$ are also in R_i. But now (c_{i-1}, c_i) is an edge connecting two vertices of R_i which is not itself an edge of R_i. Thus R_i is not an induced graph, contrary to Theorem 9.2.4. Thus the assertion of Step 2 is established.

[Note that Step 2 implies that the gallery G never crosses itself, that is, the chambers c_k are pairwise distinct.]

Step 3. *The intersection* $\cap_{i \in \mathbf{N}^+} R_i$ *is empty.*

Suppose c were a chamber in each R_i, $i = 1, 2, \ldots$. Clearly c is not one of the chambers c_k of the gallery G for R_{k+1} does not contain c_k by Step 2. Since C is connected (for it is assumed residually connected) there is a gallery of finite length $H = (c = h_0, h_1 \ldots h_d = c_0)$ connecting c to the initial vertex c_0 of the gallery G. Now select any chamber c_k, $k > 1$ of the gallery G. Since c and c_k both belong to the residue R_k, which by definition is connected, there exists a finite gallery H_k connecting c_k to c, having no edge of type k. Thus if H had no edge of type t_k, we should conclude that $c_0 \in R_k$ against Step 2. It follows that some edge of the finite gallery H bears the label t_k. But this statement must be true for *all* k greater than 1 – that is, for infinitely many values of k. On the other hand, since the residual connectedness of C forces each edge of C to bear exactly one type-label (Theorem 9.2.4), H can only accomodate d types. So we have a contradiction. It follows that no such c exists, and so the intersection of all the R_i is empty. That establishes Step 3.

Now we can complete the proof of the theorem. Clearly if $i < j$, $R_i \cap R_j$ contains c_j, so the R_i pairwise intersect non-trivially. On the other hand, their global intersection is empty. This contradicts (CRC1) and so C cannot be residually connected. □

Residual Connectedness and Functors C and Γ

We begin with the following.

Lemma 9.2.6 *Suppose C is a chamber system over I having the property (CRC1). Then $\mathbf{\Gamma}(C)$ is a residually connected geometry over I. If, in addition, C has (CRC2), then there is an isomorphism*

$$C \simeq \mathbf{C}(\mathbf{\Gamma}(C)).$$

Proof Let $F := \{R_\sigma | \sigma \in I - \{i\}\}$ be a flag of cotype i in $\mathbf{\Gamma}(C)$; here the R_σ are residues of C of cotype $\sigma \in I$, $\sigma \neq i$, which intersect pairwise non-trivially. By (CRC1), the intersection $\cap\{R_\sigma \in F\}$ is a (non-empty) panel P of type i. Then for any chamber $c \in P$, the unique residue R_i of cotype i containing c is an object in the residue of F in the geometry $\mathbf{\Gamma}(C)$. That proves (RC1).

Now suppose $H := \{R_\sigma | \sigma \in J\}$ is a flag of $\mathbf{\Gamma}(C)$ of type J, a proper subset of I. That means each R_σ can be assumed to be a residue of cotype σ, and any two members of H possess a non-empty intersection. It follows from (CRC1) that the intersection $S := \cap\{R_\sigma \in H\}$ is a residue of C of type $K := I - J$. Select any chamber $c \in S$, and type $k \in K$. Then the unique residue $R_k(c)$ of cotype k containing c is an object of type k in $\mathbf{\Gamma}(C)$ incident with every member of H. Thus the residue of H in this geometry contains an object of every type. Conversely, any object R of type k in the residue of H in $\mathbf{\Gamma}(C)$ has the form $R_k(c)$ for a chamber $c \in S \cap R$ since the latter intersection is, by (CRC1), a non-empty intersection of the family of corank one residues of $H \cup \{R\}$ which pairwise have a non-empty intersection. Since c lies in $R_k(c)$ as k ranges over K we see that any object $R = R_{k'}(c), c \in S$ of $\mathrm{Res}_{\mathbf{\Gamma}(C)}(H)$ lies in a flag-chamber $F_K(c) = \{R_k(c) | k \in K\}$ of this residue (of type K). Another application of (CRC1) shows that every flag chamber of $\mathrm{Res}_{\mathbf{\Gamma}(C)}(H)$ has the form $F_K(c)$ for some (not necessarily unique) chamber $c \in S$.

It is now easy to check that if c_1 and c_2 are ℓ-adjacent chambers of S ($\ell \in K$) then the corresponding flag-chambers $F_K(c_1)$ and $F_K(c_2)$ are ℓ-adjacent. It follows that the residue $\mathrm{Res}_{\mathbf{\Gamma}(C)}(H)$ is chamber connected, and so, by Exercise 2.5 of Chap. 2, p. 58, $\mathrm{Res}_{\mathbf{\Gamma}(C)}(H)$ is a non-empty connected geometry. Thus (RC2) holds for $\mathbf{\Gamma}(C)$, completing the proof of the first statement of the lemma.

Now suppose C is residually connected. For each type $k \in I$, and chamber c, we write $R_k(c)$ for the unique residue of cotype k containing c. Thus, for each chamber c we obtain a unique flag-chamber $F(c) := \{R_k(c) | k \in I\}$ of $\mathbf{\Gamma}(C)$ – that is, an element of $\mathbf{C}(\mathbf{\Gamma}(C))$. Note that if c and c' are distinct chambers, condition (CRC2) forces $F(c) \neq F(c')$. Conversely, if $H := \{R_k\}_{k \in I}$ is a pairwise non-trivially intersecting family of residues with R_k of cotype k, then by (CRC1) the intersection $\cap_{k \in I} R_k$ is non-empty and by (CRC2) consists of a single chamber c. Thus $H = F(c)$ so $F : C \to \mathbf{C}(\mathbf{\Gamma}(C))$ is surjective. Similarly, if c and c' are distinct chambers, they cannot both live in the same collection of corank one residues, so $F(c) \neq F(c')$. Thus F is a bijection.

Finally c and c' are i-adjacent if and only if $R_k(c) = R_k(c')$ for all $k \neq i$ and $R_i(c) \neq R_i(c')$ (for by (CRC2) the latter condition forces c and c' to belong to a common residue of type i). Thus F is an isomorphism of chamber systems. □

Lemma 9.2.7 *Let $\Gamma = (V, E; \tau)$ be a residually connected geometry over $I = \tau(V)$, where the rank $|I|$ is finite. Then $\mathbf{C}(\Gamma)$ is a residually connected chamber system. Moreover, there is a geometry isomorphism*

$$\Gamma \simeq \mathbf{\Gamma}(\mathbf{C}(\Gamma)).$$

Proof Assume $I = \{1, 2, \ldots, n\}$. By part 2 of Lemma 2.5.1, Chap. 2, p. 55, every flag of Γ lies in a chamber flag. If R is a residue of cotype k in $\mathbf{C}(\Gamma)$ then all flag-chambers of R contain a unique common object $f(R)$ of type k. More than that, since $|I|$ finite implies $\mathrm{Res}_\Gamma(f(R))$ is chamber-connected (Lemma 2.5.1, Chap. 2, once more), the flag chambers in R form the full set of flag-chambers of the form $f(R) + F$ where F ranges over all of the flag-chambers (of type $I - \{k\}$) of $\mathrm{Res}_\Gamma(f(R))$.

Suppose $\{R_k | k \in J\}$ is any family of residues of $\mathbf{C}(\Gamma)$ which pairwise possess a non-trivial intersection and with each R_k of cotype k, k ranging over J, a subset of I. Let X_k be the unique object of type k common to all flag-chambers of R_k. Then for any ℓ in J distinct from k, the existence of a flag-chamber in $R_k \cap R_\ell$ forces X_k to be incident with X_ℓ. Thus $F := \{X_k | k \in J\}$ is a flag of type J in Γ. There exists a flag chamber $\hat{F}$ containing F and clearly $\hat{F}$ is in the intersecion of all the $R_k, k \in J$ since any residue of cotype k containing F contains all flag chambers that contain F (by the argument at the end of the previous paragraph). Thus (CRC1) holds for $\mathbf{C}(\Gamma)$.

Now consider any flag chamber $c = (X_1, X_2, \ldots, X_n)$. Let $R_k(c)$ be the unique cotype k residue of $\mathbf{C}(\Gamma)$ containing c. As observed in the first paragraph above, $R_k(c)$ is the set of all flag chambers which contain X_k. Then the intersection $\cap_{k \in I} R_k$ is the flag chamber containing each X_k and there is only one – c itself. Thus (CRC2) holds for $\mathbf{C}(\Gamma)$.

It remains to exhibit the isomorphism. The function f given at the end of the first paragraph produces a mapping

$$\mathbf{\Gamma}(\mathbf{C}(\Gamma)) \to \Gamma,$$

taking each cotype k residue R of $\mathbf{C}(\Gamma)$ (itself a collection of flag-chambers agreeing in their type k object) to the unique object $f(R)$ of type k which they contain. Since $\mathrm{Res}_\Gamma(f(R))$ is chamber connected, any flag-chamber containing $f(R)$ belongs to R. Thus f is injective.

Also, for any object X_k of type k in Γ, and any flag chamber F' containing it, we have $X_k = f(R_k(F'))$. Thus f is surjective.

If R_k and R_ℓ contain a common flag chamber F', the objects $f(R_k)$ and $f(R_\ell)$ lie in F' and so are incident. Conversely, if X_k and X_ℓ are incident objects of Γ, then any flag-chamber F' containing F lies in both $R_k = f^{-1}(X_k)$ and $R_\ell = f^{-1}(X_\ell)$,

the sets of all flag chambers containing X_k and X_ℓ, respectively. Thus f is an isomorphism of geometries. The proof is complete. □

Remark Now note that Theorem 9.2.2 is just a combined statement of Lemmas 9.2.6 and 9.2.7.

Corollary 9.2.8 (In the presence of residual connectedness, the functors $\mathbf{\Gamma}$ and $\mathbf{C}$ commute with the taking of residues.)

1. *Suppose Γ is a residually connected geometry of finite rank and let $C = \mathbf{C}(\Gamma)$ be its associated chamber system. Let F be a flag of cotype $J \neq \emptyset$. Let $R(F)$ be the set of all chamber flags of Γ which contain F. Then by residual connectedness and the finiteness of rank, $R(F)$ is a residue of $\mathbf{C}(\Gamma)$ of type J. Our assertion is that there is an isomorphism*

$$R(F) \simeq \mathbf{C}(Res_\Gamma(F))$$

   as chamber systems.
2. *Suppose C is a residually connected chamber system over I of arbitrary rank and let $\Gamma = \mathbf{\Gamma}(C)$ be its associated (residually connected) geometry. For every residue R of type J, there corresponds a flag $F(R)$ of cotype J of the associated geometry $\mathbf{\Gamma}(C)$, and there is a geometry isomophism*

$$Res_{\mathbf{\Gamma}(C)}(F(R)) \simeq \mathbf{\Gamma}(R).$$

Proof The isomorphisms in question are just the restrictions to residues of the two isomorphisms of Theorem 9.2.2. The beginning student is encouraged to make a formal proof of this corollary, noting the places at which the hypotheses of residual connectedness are used. □

Example 6 Let Γ be a geometry with exactly one object X_k of each type $k \in I = \{1, \ldots, n\}$. Then we see that Γ is indeed a residually connected geometry of finite rank n with exactly one flag chamber $c = (X_1, \ldots X_n)$. Then $\mathbf{C}(\Gamma)$ has one chamber and all residues of corank one coincide as sets of chambers. Nonetheless each such residue is attached to a unique type and so are distinct as residues. Thus $\mathbf{\Gamma}(\mathbf{C}(\Gamma))$ is indeed isomorphic to Γ. This example shows why one wants to take the rank of the chamber system as the cardinality of the relevent I rather than the set of types i for which i-adjacencies are actually exhibited. This way, the functor $\mathbf{\Gamma}$ preserves rank.

Example 7 Now suppose Γ is a geometry over $I = \{0, 1, \ldots\}$, the set of natural numbers with exactly two objects X_k and Y_k of each type k. We assume that any two objects of distinct types are incident. Then it is easy to see that Γ is connected, but is not chamber-connected since no gallery connects the flag chambers $X := \{X_k | k \in I\}$ and $Y = \{Y_k | k \in I\}$. (See Exercise 2.3 of Chap. 2, p. 58.) The arguments in the preceding Lemma 9.2.7 depended heavily on residues of flags being chamber connected. Thus the hypothesis that Γ has finite rank is crucial for this lemma.

9.3 Chamber Systems with Strongly Gated Residues

9.3.1 Introduction

In this section, C is a chamber system over I. It is not assumed that C is type M, or that it is even connected. (In Sect. 9.6 we shall resume the hypothesis that C is type M.) However, we do assume that C is a chamber system over I satisfying the condition (typ). We shall discuss three distinct hypotheses concerning a chamber system C:

(RG) Every residue is strongly gated in C.
(RG^1) Every residue of corank one is strongly gated.
(RG_2) (Scharlau's condition.) Every residue of C of rank at most 2 is strongly gated in C.

On the face of it, condition (RG) implies the other two. It is not difficult to show that (RG^1) implies (RG) (Corollary 9.3.6 below), and one would hope for an equally "metrical" proof that (RG_1) implies (RG). This is not quite as easy.[7] To keep the argument completely free of the assumption that C is type M, we have to introduce a new type of homotopy, called "Λ-homotopy," which tries to replace the notion of M-homotopy in arbitrary chamber systems which might not be of type M. The beauty of Λ-homotopy is that it always exists for any chamber system C satisfying (typ). It requires no assumptions whatsoever about rank two residues of C.

9.3.2 Basic Properties Concerning Strongly Gated Residues

To start with, the hypothesis that a subgraph $H = (H, E')$ of a graph $\Delta = (V, E)$ is *strongly gated* asserts that for every vertex $c \in V$ there exists a vertex $g = g(c)$ (to indicate that it depends on c) in H such that for *every* $x \in H$,

$$d_\Delta(c, x) = d_\Delta(c, g) + d_H(g, x). \tag{9.1}$$

When $d_\Delta(c, x)$ is finite, both terms on the right side are finite and $g(c)$ is in the same connected component as c and x. It is then easy to prove that the vertex $g(c)$ is uniquely determined by c.

---

[7] The author's original proof in the unpublished Shult *Freiburg Notes* (1989) involved proving that the condition (RG_2) implies condition (P_c) which is an assertion about M-homotopy of galleries, a condition Tits proved was equivalent to the notion of "Building" for chamber systems of type M. The author's proof was hardly original. The arguments had already been outlined by Rudolf Scharlau except that "strongly gated" must replace "gated" in the arguments. But one sees that this proof invokes concepts completely dependent upon the hypothesis that the C is type M. One cannot even say that a gallery is of reduced type without this notion.

The reader might wonder what all this means when the graph Δ is not connected?[8] Well it makes perfect sense.

If two vertices x and y are in separate connected components of a graph G we say that they are "at infinite distance" and write $d_G(x, y) = \infty$.

If $d(c, x)$ is infinite in Eq. (9.1) then clearly at least one of the two terms on the right side of Eq. (9.1) is also infinite.

(1) If the first term is infinite then any vertex of H will serve as the gate and c lies in no common connected component with a single vertex of H.

(2) If the first term $d(c, g)$ is finite, then $d_H(g, x)$ is infinite. That means that x is not connected to g in the subgraph H. But it also means that x is not connected to c.

At first sight this does not seem very significant. But now suppose H is *strongly gated* – that is, it is strongly gated with respect to *any* vertex c. Then we see that two vertices of H are in the same connected component of Δ if and only if they are in the same connected component of H [Hint: Take c to be one of these vertices of H, deduce that it is its own gate $g(c)$, and apply the fundamental (9.1).]

Lemma 9.3.1 *Suppose $H = (H, E')$ is a strongly gated subgraph of a graph $\Delta = (V, E)$ and let $\Delta = \oplus_{\sigma\in\Sigma}(V_\sigma, E_\sigma)$ be a decomposition of Δ into its connected components. Then (as an induced subgraph of H) each non-empty intersection $V_\sigma \cap H$ is connected, and is strongly gated in the connected component $\Delta_\sigma := (V_\sigma, E_\sigma)$.*

Remark So, in a way, the discussion of strongly gated subgraphs of a graph reduces to the relationship of the connected components of both graphs.

Proof The proof is an easy exercise (Exercise 9.10). □

We now recall from Sect. 1.1.4 the following basic facts about subgraphs of a (possibly non-connected) graph.

Lemma 9.3.2 *1. Every strongly gated subgraph of a graph is a convex induced subgraph, and hence is an isometrically-embedded subgraph.*
2. Every convex induced subgraph of a connected graph is connected.
3. An arbitrary intersection of a family $\{R_\sigma | \sigma \in \mathcal{F}\}$ of connected convex induced subgraphs is either empty or is connected.

Proof Part 1 was detailed in Chap. 1. Part 2 is immediate from part 1, noting that the isometric embedding property forces finite distances. Part 3 follows from part 2 and the fact that the family of subgraphs which are convex and induced is closed under taking arbitrary intersections (Chap. 1). □

We wish to transfer these ideas to residues of a chamber system. For that purpose we need an analog of the concept of "induced subgraph" for residues. We say that a residue R of type $J \subseteq I$ in a chamber sytem C over I is an *induced residue of C* if

[8] The author even encountered one referee who insisted that it made no sense at all if Δ were not connected.

and only if, for any edge $e = (x, y)$ of C whose vertices x and y lie in R, we have $\lambda(e) \subseteq J$ – that is, the only adjacencies between vertices of R are the j-adjacencies for $j \in J$.

Note that if condition (*typ*) holds for C, any residue that is embedded as an induced subgraph C becomes an induced residue.

9.3.3 Intersections of Strongly Gated Residues

Corollary 9.3.3 *Let $\{R_\sigma | \sigma \in \mathcal{F}\}$ be a family of convex induced residues of a chamber system C. Let J_σ be the type of the residue R_σ. Then the gobal intersection $\cap\{R_\sigma | \sigma \in \mathcal{F}\}$ is either empty, or is an induced residue of type $\cap\{J_\sigma | \sigma \in \mathcal{F}\}$.*

Proof Set $T := \cap\{R_\sigma | \sigma \in \mathcal{F}\}$, and $J_T := \cap\{J_\sigma | \sigma \in \mathcal{F}\}$. If $T = \emptyset$ there is nothing to prove.

Suppose $e = (t, s)$ is an edge of C with $t \in T$. If $s \in C - T$, then $\lambda(e)$ cannot contain any label of J_T. On the other hand, if $s \in T$, then every type label of $\lambda(e)$ lies in J_σ since R_σ is an induced residue for all σ. Thus T is a union of induced residues of type J_T.

It remains only to show that T is connected. But since our hypothesis provides that each R_σ is connected, this property holds for T by part 3 of the preceding Lemma 9.3.2. □

Corollary 9.3.4 *Suppose $\{R_\sigma | \sigma \in \mathcal{F}\}$ is a family of strongly gated residues of a chamber system C satisfying condition (typ). Then the intersection $\cap\{R_\sigma | \sigma \in \mathcal{F}\}$ is either empty, or is a convex induced residue whose type is the intersection of all the types of the R_σ.*

Proof This is immediate from the preceding Corollary 9.3.3 and the observation that (typ) makes strongly gated residues induced. □

Theorem 9.3.5 (Shult [114].) *Suppose C is a connected chamber system with (typ) so each edge reflects just one type of adjacency. Suppose further that $\{R_\sigma | \sigma \in \mathcal{F}\}$ is a collection of strongly gated residues with a non-empty intersection R. Then R is strongly gated.*

Proof By Corollary 9.3.3, R is a convex induced residue of C of type $J_R := \cap\{J_\sigma | \sigma \in \mathcal{F}\}$ where J_σ is the type of the residue R_σ.

Suppose by way of contradiction that R is not strongly gated. Then there exists a vertex y such that R is not gated with respect to y. This failure cannot occur unless y is in the same connected component of the graph (C, E) as its connected subgraph R. So, among such y such that R is not gated with respect to R, we choose y so that the distance $d := d(y, R)$ – the length of a shortest geodesic connecting y to a vertex of R – is as small as possible. (Since y and R are in the same connected component, this distance d from y to a nearest vertex of R is finite.)

Let us first show that d is not zero – that is, y is not in R. Since R is a convex induced residue, its graph is isometrically embedded in the graph of C. Thus

$d_R(y, r) = d(y, r)$ where d_R is the internal metric of R, and so R is strongly gated with respect to y using y itself as the gate.

So we may assume that y is not in R and $d > 0$. Then there is at least one residue $M := R_\tau$ containing R but not containing y. Since M is strongly gated, there exists a vertex $g \in M$ such that for any vertex $r \in R$,

$$d(y, r) = d(y, g) + d_M(g, r). \tag{9.2}$$

It follows from this statement (universally quantified on r), and the fact that $d(y, g) > 0$, that

$$d(g, R) < d(y, R) = d.$$

By the minimal choice of d, R is strongly gated with respect to g. So, again, there is a gate $h \in R$ such that for any vertex s of R,

$$d(g, s) = d(g, h) + d_R(h, s), \text{ for all } s \in R. \tag{9.3}$$

But as a special case of (9.2):

$$d(y, h) = d(y, g) + d_M(g, h). \tag{9.4}$$

Now, as M is isometrically embedded, $d(g, s) = d_M(g, s)$, for all $s \in R$. By the substitution of bound variables we may replace s by r to get

$$\begin{aligned} d(y, r) &= d(y, g) + d_M(g, r) \\ &= d(y, g) + d(g, r) \\ &= d(y, g) + d(g, h) + d_R(h, r) \\ &= (d(y, g) + d_M(g, h)) + d_R(h, r) \\ &= d(y, h) + d_R(h, r) \end{aligned}$$

for all $r \in R$. Thus R is gated with respect to y after all. This contradiction to the choice of y completes the proof. □

Corollary 9.3.6 *In any chamber system satisfying condition (typ), all residues are strongly gated if and only if all residues of corank one are strongly gated.*

Proof Immediate from Theorem 9.3.5. □

Theorem 9.3.7 *Suppose C satisfies (typ). Suppose $\{R_\sigma | \sigma \in \Sigma\}$ is a finite family of strongly-gated residues of C. Suppose the R_σ pairwise intersect non-trivially – that is, for any $\sigma, \tau \in \Sigma$ we have $R_\sigma \cap R_\tau \neq \emptyset$.*

Then the global intersection is non-empty – that is

$$\cap\{R_\sigma | \sigma \in \Sigma\} \neq \emptyset.$$

Proof Since Σ is finite, we may represent the family as

$$\{R_1, R_2, \ldots, R_n\},$$

where each R_i is strongly gated in C and $R_i \cap R_j$ is non-empty for all $1 \leq i, j \leq n$. We may assume $n \geq 3$ and proceed by induction on n.

If $n \geq 4$ then by induction $R_3' := R_3 \cap \cdots \cap R_n$ is non-empty. By Corollary 9.3.4 R_3' is a strongly gated residue and by a second application of induction on n, $R_1 \cap R_3'$ and $R_2 \cap R_3'$ are non-empty. Then, by a third application of induction on n, we have

$$\emptyset \neq R_1 \cap R_2 \cap R_3' = R_1 \cap \cdots \cap R_n$$

and we are done. Thus we may assume $n = 3$ exactly.

By hypothesis, there is a chamber $y \in R_2 \cap R_3$. Also by hypothesis R_1 is strongly gated with respect to y. So there is a gate $g_1 \in R_1$ and by the convexity and isometric embedding of R_1 in C, for every $x \in R_1$ one can find a geodesic gallery from y to x that passes through g_1. Thus if we take $x \in R_1 \cap R_2$, there is a geodesic gallery from y to x passing through g_1. Then convexity of R_2 forces $g_1 \in R_2$. Similarly, if we take $x \in R_3 \cap R_1$, convexity of R_3 forces $g_1 \in R_3$. Now g_1 is a chamber in the global intersection $R_1 \cap R_2 \cap R_3$. The proof is complete. □

Corollary 9.3.8 *Suppose (RG) and (typ) hold. If C has finite rank, then C is a residually connected chamber system.*

Proof We must show three things. (1) The intersection over any family of pairwise non-trivially intersecting residues is non-empty. (2) Any non-empty intersection of residues is a residue – that is, it is connected. (3) The intersection of all residues containing a chamber c is the set $\{c\}$.

For (1) we can assume that the residues taking part in the intersection are pairwise distinct residues. Since the intersections are non-empty we also see that their types are pairwise distinct subsets of I. Since I is finite, there can only be a finite number of residues participating in the intersection. Then by (RG) and Theorem 9.3.7, the intersection of these residues is non-empty.

Assertion (2) follows from Corollary 9.3.4. Finally, conclusion (3) follows from (2) and the condition (typ) which says that for every pair of adjacent chambers there is a corank one residue containing one but not the other. □

9.3.4 2-Simply Connectedness is a Consequence of Strong Gatedness at Low Rank

In the next few sections, we shall be concerned with this property:

(RG$_2$) *Every residue of rank at most* 2 *is strongly gated.*

With a slight adjustment, this is essentially a notion introduced by R. Scharlau in [105].[9]

Theorem 9.3.9 *Suppose C is a chamber system satisfying (typ) and the hypothesis (RG_2) that all residues of rank at least 2 are strongly gated. Let $\mathcal{C}_2$ be the collection of all pointed circuits each of which is a circuit in some residue of rank two. Then every circuit of C is $\mathcal{C}_2$-contractible. Put another way, every connected component of C is 2-simply connected.*

Remark Of course (bowing to common usage) "2-simply connected" just means "$\mathcal{C}_2$-simply-connected" as defined in Chap. 1.

Proof Suppose, by way of contradiction, that G is a circular gallery which is not $\mathcal{C}_2$-contractible, chosen among such galleries to have minimal length $m = 2d$ or $2d+1$. Then $m > 4$, and any subsegment of G (that is, a subgallery of G) of length at most d is a geodesic path.

Select a chamber c_0 in G whose two edges in G bear labels i and j, and let R be the residue of type $\{i, j\}$ containing c_0. Let k and h be maximal so that the

$$(c_{-k}, c_{-k+1}, \ldots, c_0, c_1, \ldots, c_h)$$

is in a connected component of $R \cap G$ (actually a segment of G of type $\cdots ijijij \cdots$ of maximal length containing c_0), and complete the notation so that

$$G = (c_0, \ldots, c_m = c_0)$$

with the convention that subscripts can be read modulo m. Then by the minimality of G, $d_C(c_0, c_d) = d$ and $d_C(c_{-1}, c_d) = d - 1$ or d according as m is even or odd.

If c_d were in the residue R, then the two geodesic paths

$$(c_0, \ldots, c_d) \text{ and } (c_{-1}, c_{-2}, \ldots, c_d)$$

would both connect vertices of R, and so would lie in R by convexity of the strongly gated R. That would force $G \subseteq R$ against G not being $\mathcal{C}_2$-contractible. Thus we may assume that c_d is not in R.

Now since R is strongly gated by hypothesis, there is a "gate" g in R so that

$$d(c_d, y) = d_C(c_d, g) + d_R(g, y) \tag{9.5}$$

for every chamber y of R.

Let A be a geodesic path from c_d to g and let H^+ and H^- be geodesic paths from g to c_h and c_{-k}, respectively. Also set

[9] In fact Scharlau's work was the starting point that inspired this entire section on strongly gated residues – most of it introduced in the Shult *Freiburg Notes* (1989) as early as 1988.

$$G^+ := (c_h, c_{h+1}, \ldots, c_d) \text{ (length } d-h) \text{ and}$$
$$G^- := (c_{-k}, c_{-k-1}, \ldots, c_d) \text{ (length } m-d-k).$$

We now have the arrangement of galleries given in Fig. 9.5.

Then by (9.5),

$$\ell(G^-) = m - d - k = \ell(A) + \ell(H^-)$$

and

$$\ell(G^+) = d - \ell = \ell(A) + \ell(H^+)$$

where, as usual, $\ell(X)$ denotes the length of the gallery X.

Then $A \circ H^+ \circ G^+$ is a circuit of length $2(d-h) = 2d - 2h$ and $A \circ H^- \circ G^-$ is a circuit of length $2(m-d-k)$.

If $m = 2d$, these lengths are $m - 2h$ and $m - 2k$, respectively. If $m = 2d + 1$, the first length is clearly less than m and the second length is $2(d + 1 - k)$. Since $\min(k, h) \geq 1$, in all cases the circuits have length less than m and so are $\mathcal{C}_2$-contractible by the minimality of G. It follows that G decomposes into three $\mathcal{C}_2$-contractible circuits (the two just mentioned and one circuit entirely in R), all visible in Fig. 9.5, so G is $\mathcal{C}_2$-contractible against the choice of G. The proof is complete. □

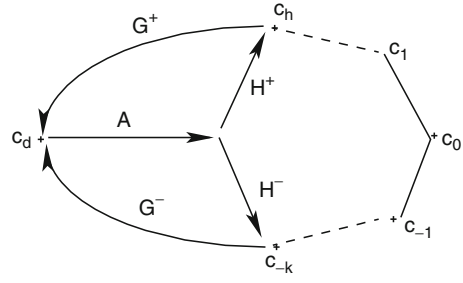

Fig. 9.5 The decomposition of a minimal circuit by condition (RG_2)

9.3.5 Λ-Homotopy

One might well imagine that there should be a corresponding "bottom up" theorem that should read something like this:

(Conjecture.) *Suppose C is a chamber system satisfying the condition (typ) as well as the condition.*

(RG_2) Every residue of rank at most two is strongly gated in C.

Then every residue of C is strongly gated. Stated another way, in the presence of (typ), the condition (RG_2) implies (RG).

Unfortunately the author knows of no way to do this without invoking a more sensitive homotopy – one introduced by Tits in [139] in the context of chamber systems of type M – but one which works for arbitrary chamber systems satisfying the modest hypothesis (typ) (that distinct chambers are i-adjacent for at most one type-label i). We call it "Λ-homotopy," where the "Λ" stands for its key property: it is *length-preserving* on galleries.

Suppose R is a residue of rank two in a chamber system C with the property (typ). A $\Lambda(R)$-*homotopy* is a transformation of minimal galleries of R (viewed as a chamber system) $G \to H$ which connect the same two chambers of R. (Note that it is necessary to keep the residue "R" in view throughout this defintion. A minimal gallery is one of shortest length connecting two chambers of an assumed ambient chamber system. If we say that G is a *minimal gallery of a residue* R, we mean that all vertices and edges of G are vertices and edges of the residue R, and that one can discover no shorter gallery whose vertices and edges belong to the subgraph R connecting the initial and terminal vertices of G. But bear in mind that a minimal gallery of a residue R may not at all be a minimal gallery of C (what we have called a geodesic).

One should note the following simple consequences of the definition of a $\Lambda(R)$-homotopy.

Lemma 9.3.10 *We have:*

1. *Two $\Lambda(R)$-homotopic galleries of R have the same length. (Remember they are minimal galleries of R connecting the same two chambers.)*
2. *If R is type $\{i, j\}$, then two $\Lambda(R)$-homotopic galleries either:*
   - *(a) they have length one – that is, $\ell(G) = \ell(H) = 1$, and $G = H$ is just recording an i-adjacency, or is recording a j-adjacency (condition (typ) prevents both from occurring), or*
   - *(b) $\ell(G) = \ell(H) > 1$ and both i and j-adjacencies occur in each of the galleries.*

So much for R. Now let us say that a transformation $G \to H$ of two galleries of C form an *elementary* Λ-*homotopy* if and only if it has the form

$$A \circ U \circ B \to A \circ W \circ B,$$

where there exists a rank two residue R with respect to which the transformation $U \to W$ is a $\Lambda(R)$-homotopy.

Of course we then say that galleries G and H are Λ-*homotopic* if and only if H is obtained from G by a sequence of elementary Λ-homotopies.

Remark It is time to take note. Two galleries are $\Lambda(R)$-homotopic if and only if they are both minimal galleries in R connecting the same two chambers. But in the wider world of chamber system C, two Λ-homotopic galleries might not be geodesics of C.

Now it is only a matter of sorting through the definitions to deduce the following tautology.

Theorem 9.3.11 (Properties of Λ-homotopy.) *Suppose G and H are Λ-homotopic galleries. Then:*

1. *G and H begin and end at the same vertices.*
2. *G and H have exactly the same length.*
3. *The total collection of edge-labels exhibited in G are exactly the total collection of edge labels exhibited in H: we symbolically render this property as $Typ(G) = Typ(H)$.*

Corollary 9.3.12 *Suppose R is any sort of residue of a chamber system C satisfying (typ). Suppose G and H are two Λ-homotopic galleries connecting two chambers of R. If one of the galleries is in R so is the other.*

Proof If G is a gallery of R, then $Typ(G) = Typ(H)$ which implies the conclusion. □

Remark It is important to note that in defining Λ-homotopy we have not imposed any extra condition upon a chamber system C with (typ) – that is, we have not restricted the range of the definition in any way.

9.3.6 Further Consequences of the Hypothesis (RG_2)

We begin by proving that (RG_2) implies a property that would have been impossible to state, had we not had the diversion of the previous section.

Theorem 9.3.13 *Suppose C is a chamber system with the property (typ). Assume the following:*

(RG_2) *All residues of C of rank at most two are strongly gated.*

Then C has this property:

(Λ-min) *Any two minimal galleries of C which connect the same two chambers are Λ-homotopic.*

Proof Suppose by way of contradiction that C contains pairs (G, H), where G and H are two minimal galleries of C with the same initial and terminal vertices which are not Λ-homotopic. Then of course G and H have the same length, namely the distance in C from their common initial chamber to their common terminal chamber. We imagine the pair (G, H) chosen so that $d = \ell(G) = \ell(H)$ is minimal. We know that $d > 2$, otherwise G and H would belong to the same strongly-gated rank two residue and so would be Λ-homotopic by definition.

We may write $G = (s = g_0, g_1, \ldots, g_d = t)$ and $H = (s = h_0, \ldots, h_d = t)$. If $g_{d-1} = h_{d-1} = s'$, we would obtain factorizations

$$G = G_1 \circ (s', t) \text{ and } H_1 \circ (s', t).$$

Observing that G_1 and H_1 are minimal galleries of length $d-1$, one would conclude that G_1 and H_1 are Λ-homotopic, and then the definition would at once yield the Λ-homotopy of G and H, a contradiction.

Thus we may assume the "next-to-last" chambers of the galleries G and H are distinct. By condition (typ) we may let $i = \lambda(g_{d-1}, t)$ and $j = \lambda(h_{d-1}, t)$, the types attached to the last edge of each gallery.

Next we claim that $i \neq j$. If $i = j$ then g_{d-1}, h_{d-1}, and t would belong to a common panel P containing two distinct vertices at distance $d-1$ from s and one at distance d. But that is impossible since (RG_2) forces panel P to be strongly gated with respect to s.

Let R be the residue of type $\{i, j\}$ containing g_{d-1}, h_{d-1}, and t, and let g be the gate of R with respect to the initial chamber s. As before we write

$$G = G_1 \circ (g_{d-1}, t) \text{ and } H_1 \circ (h_{d-1}, t).$$

Let G_2 and H_2 be minimal galleries in R from g to g_{d-1} and h_{d-1} respectively. Finally, fix a minimal gallery A of C from s to g. Then from the hypothesis of strong gatedness,

$$\begin{aligned} \ell(G_1) &= \ell(A) + \ell(G_2), \\ \ell(H_1) &= \ell(A) + \ell(H_2), \\ d &= \ell(A) + \ell(G_2 \circ (g_{d-1}, t) \\ &= \ell(A) + \ell(H_2 \circ (h_{d-1}, t). \end{aligned}$$

It follows that $A \circ G_2$, $A \circ H_1$, $G_2 \circ (g_{d-1}, t)$, and $H_2 \circ (h_{d-1}, t)$ are all minimal galleries of length less than d. We thus have a series of Λ-homotopies (indicated by "$\sim$"):

$$\begin{aligned} G &= G_1 \circ (g_{d-1}, t) \\ &\sim (A \circ G_2) \circ (g_{d-1}, t) \\ &\sim A \circ (H_2 \circ (h_{d-1}, t)) \\ &\sim H_1 \circ (h_{d-1}, t) \\ &= H. \end{aligned}$$

(We have freely applied the "associativity" of gallery concatenation without explicit intermediate equations.) The proof is complete at this point. □

Theorem 9.3.14 *Assume C is a chamber system with property (typ). If (RG_2) holds, then every residue of C is a convex induced residue. So it is isometrically embedded — that is, any geodesic in R is already a geodesic in C.*

Proof Suppose, among all residues R of C and all geodesics of these residues (their length is the distance between their extremities as measured by the internal metric of R) we chose a pair (R, G) so that G was not a geodesic of C and did this so G has

minimal length d. Since (RG_2) makes all rank two residues isometrically embedded, we must have $d > 2$.

Let $G = (s = g_0, g_1, \ldots, g_{d-1}, g_d)$. By assumption

$$d_(s, g_i) = i, \text{ for } i \leq d-1 \text{ and} \tag{9.6}$$

$$d(s, g_d) \leq d-1. \tag{9.7}$$

Then the panel P (of type i, say) on the edge (g_{d-1}) contains a chamber t at distance $d-2$ from the initial chamber s, and t may or may not be g_d. One thing is certain: (t, g_{d-1}) is labelled i. Let A be a geodesic of C from s to t. Then $A \circ (t, g_{d-1})$ and $(s = g_0, \ldots g_{d-1})$ are geodesics of C from s to $g_d - 1$. By Theorem 9.3.13, these two galleries are Λ-homotopic. By Corollary 9.3.12, A is a gallery in the residue R. Now if $t = g_d$, we have s and g_d connected by A; if $t \neq g_d$, then s and g_d are connected by $A \circ (t, g_d)$. In either case, s and g_d are connected by a gallery of the residue R having length at most $d-1$. This contradicts our choice of G as a minimal gallery of R.

Thus no such pairs (R, G) exist, and the theorem is proved. □

Theorem 9.3.15 *Suppose C is a chamber system with property (typ). If (RG_2) holds, then every residue of C is strongly gated.*

Remark This theorem has been proved for chamber systems of type M using the fact that all residues of a Coxeter chamber system are strongly gated. Since we have not bothered to prove that property of a Coxeter chamber system and since we intend to prove it here without the type M hypothesis, the proof is essentially new. The induction used here is subtle and some care must be taken.

Proof Suppose the theorem is false. Then there exists a chamber system C satisfying (RG_2) and (typ), with a non-empty collection $\mathcal{Y}$ of triples (c, g, R) such that:

1. R is a residue and c is a chamber with $d(c, g) = d(c, R)$.
2. The subset

$$X(c, g, R) := \{x \in R | d(c, x) < d(c, g) + d_R(g, x)\}$$

is non-empty.

Suppose $(c, g, R) \in \mathcal{Y}$, so that the set $X(c, g, R)$ is non-empty. Clearly $g \notin X(c, g, R)$ by definition. We next consider geodesic paths from g to a member of $X(c, g, R)$ which encounters a member of $X(c, r, R)$ only at its terminus. (Clearly in our convex R, such a path is a path of R.) We then sort through such geodesic paths to find one such that the length of the path plus the distance of its terminus from c is as small as possible. Thus if

$$G = (g = g_0, g_1, \ldots, g_{d-1}, g_d = t)$$

is such a path, then for $i < d$, no vertex g_i lies in $X := X(c, g, R)$ and this is done so that $d + d(c, t)$ is as small as possible.

Then we have

$$d(c, g_i) = d(c, g) + i \text{ for } 0 \leq i \leq d-1, \text{ and} \tag{9.8}$$

$$d(c, t) = d(c, g) + (d-1) \text{ or } d-2. \tag{9.9}$$

(Note that, by the preceding Theorem 9.3.14, every residue is isometrically embedded, so we can replace the $d_R(g, y)$s that would normally occur in the equations for the strongly-gated property, by $d(g, y)$s.) Let us write $d_1 := d(c, g)$, since this quantity makes a large number of appearances in this play.

Now suppose

$$d(c, t) = d(c, g_{d-1}) = d_1 + d - 1.$$

Then by (RG_2) the panel P on g_{d-1} and t contains an element t' with $d(c, t') = d_1 + d - 2$, and $d(g, t') \leq d = d(g, t)$. Then as $t \neq t'$, the sum $d(c, t') + d(g(t')$ is smaller than the similar sum $d(c, t) + d(g, t)$ and $d(g, t) = d - 1$ shows that $(g = g_0, \ldots g_{d-2}, t')$ is a geodesic with all members except the last not in the set X. Thus by the minimality of the choice of the gallery G we get a contradiction.

Thus we must suppose that $t = t'$, so we have

$$d(c, t) = d_1 + d - 2 = d(c, g_{d-2}).$$

Moreover, $U := \{t, g_{d-1}, g_{d-2}\}$ cannot lie in a common panel (that would contradict $d(g, t) = d$), and so the edges (g_{d-1}, t) and (g_{d-1}, g_{d-2}) bear distinct labels i and j, respectively. Now let S be the unique residue of type $\{i, j\}$ containing the subgallery (g_{d-2}, g_{d-1}, t). By (RG_2) S is strongly gated with respect to c. Let p be its gate with respect to chamber c. Then (noting the isometric embeddings)

$$d(c, t) = d(c, p) + d(p, t) = d_1 + d - 2. \tag{9.10}$$

Now suppose p is not in X. Then

$$d(c, p) = d(c, g) + d(g, p). \tag{9.11}$$

Substitution of the right side of (9.11) for $d(c, p)$ in Eq. (9.10) yields

$$d_1 + d - 2 = d(c, t) = (d(c, g) + d(g, p)) + d(p, t) \tag{9.12}$$

$$= d_1 + d(g, p) + d(p, t). \tag{9.13}$$

So

$$d(g, p) + d(p, t) = d - 2. \tag{9.14}$$

But the latter is impossible since the "triangle inequality" forces

$$d = d(g, t) \leq d(g, p) + d(p, t). \tag{9.15}$$

So we must assume $p \in X$. Now from the gatedness of S, in addition to Eq. (9.10) we have

$$d(c, g_{d-2}) = d(c, p) + d(p, g_{d-2}) = d_1 + d - 2, \tag{9.16}$$

so

$$d(p, t) = d(p, g_{d-2}). \tag{9.17}$$

Now let

$$H := (d_{d-2} := t_0, t_1, \ldots, t_m := p)$$

be a geodesic in S. Since p is in X but g_{d-2} is not, there is a first index j such that $t_j \in X$. Thus for $0 \leq i < j$,

$$d(c, t_i) = d(c, g) + d(g, t_i) \tag{9.18}$$
$$= d(c, p) + d(p, t_i). \tag{9.19}$$

So as i increases, by Eq. (9.19), $d(c, t_i)$ decreases, and so by (9.18) $d(g, t_i)$ also decreases. But as $t_j \in X$, we have

$$d(g, t_j) = d(g, t_{j-1}) + \epsilon \text{ where } \epsilon = 0 \text{ or } 1. \tag{9.20}$$

Thus

$$d(c, t_j) + d(g, t_j) \leq d(c, t_{j-1}) + d(g, t_{j-1}) \text{ (by (9.20))} \tag{9.21}$$
$$\leq d(c, g_{d-2}) + d(g, g_{d-2}) \tag{9.22}$$
$$< d(c, t) + d(g, t). \tag{9.23}$$

Now since t_{j-1} is not in X, one has

$$d(c, t_{j-1}) = d(c, g) + d(g, t_{j-1}), \tag{9.24}$$

so any geodesic from g to t_{j-1} concatenates with any geodesic from c to g to form a gallery which is a geodesic. This means that all chambers of a geodesic gallery E from g to t_{j-1} cannot belong to X. If $d(g, t_j) = d(c, t_{j-1}) + 1$ (so $\epsilon = 1$ in (9.20)), then $E \circ (t_{j-1}, t_j)$ is a minimal gallery from g to t_j with all chambers but its terminus not in X. On the other hand, if $d(g, t_j) = d(g, t_{j-1})$ (so $\epsilon = 0$), then the geodesic gallery E can be made to factor through a gate g_j of the unique panel

P_j on the edge (t_{j-1}, t_j) – that is $E = E_1 \circ (g_j, t_{j-1})$. But in that case, $E_1 \circ (g_j, t_j)$ is a geodesic from g to t_j with only its terminal chamber in X.

In short, $t_j \in X(c, g, R)$, like t, is approachable from g by a geodesic gallery all of whose other chambers are not in X, except that now (by (9.23)) the sum of the distances of t_j from c and g is now less than the sum of the distances of t from c and g. But t was supposed to be minimal in that respect. Thus the assumption $p \in X(g, c, R)$ also leads to a contradiction.

The proof is complete. □

9.3.7 Equivalence of Various Gatedness Conditions

The following theorem summarizes the results of this section.

Theorem 9.3.16 *We suppose C to be a chamber system over I with condition (typ).*

1. *The three strong-gatedness conditions (RG_2), (RG), and (RG^1) are equivalent to one another. They all imply the condition (Λ-min).*
2. *If C satisfies any one of these three conditions, it is 2-simply connected.*
3. *If C is firm and satisfies one of the equivalent gated conditions, then it is residually connected if and only if its rank $|I|$ is finite.*

Proof Clearly (RG) implies both (RG^1) and (RG_2) since the latter properties are particularizations of the former property. But (RG^1) implies (RG) by Corollary 9.3.6. Also, Theorem 9.3.15 asserts that (RG_2) implies (RG). So the three conditions are now equivalent. Finally, Theorem 9.3.13 shows that (RG_2) implies (Λ-min), so all parts of the first conclusion hold.

The second conclusion follows from the first upon noting that by Theorem 9.3.9, condition (RG_2) implies $\mathcal{C}_2$-simple-connectedness.

If (RG) holds, and C has finite rank, then C is residually connected. But by Theorem 9.2.5, if C has infinite rank and is firm, it cannot be residually connected.

9.4 Generalized Polygons

9.4.1 Panel Homotopy

Introduction

In Chap. 1, we saw that in any graph, a collection $\mathcal{C}$ of circuits defines a $\mathcal{C}$-homotopy. There are various homotopy theories that are useful for chamber systems, depending on the choice of $\mathcal{C}$. We have met one of these: $\mathcal{C}_2$-homotopy where $\mathcal{C}_2$ is the collection of all circuits of the chamber system each of which lies within some residue of rank two. In this section we shall meet another such theory: panel homotopy. The latter is very special, and concerns only the class $\mathcal{C}$ of circuits which are confined to panels (rank-one residues). We must introduce this notion in order to avoid some

awkward terminology that has appeared in some of the literature. There one finds chamber systems that are generalized ∞-gons referred to as "trees." In fact they are not trees in the graph-theoretic sense if any of their panels are thick. The circuits that are important for generalized polygons are actually "panel-reduced" circuits. (One author calls them "proper circuits" without explanation.) The panel homotopy also has some relation to types of galleries that are useful in showing the equivalence of several definitions of "generalized polygon," so it is just something we have to go through. Actually, it is a very simple notion.

Basic Terms

Any *gallery* is a walk $G = (c_0, \ldots c_n)$ where (c_{i-1}, c_i) is an edge. Its *length*, the natural number n, is denoted $\ell(G)$. We usually keep track of the *type of the gallery*, the sequence

$$\lambda(G) := \{\lambda(c_i, c_{i+1}), i = 0, 1, \ldots n-1\}.$$

A *segment of length k of gallery G* is just a subsequence $(c_i, c_{i+1} \ldots, c_{i+k})$, $k \leq n$, itself regarded as a gallery from c_i to c_{i+k} with an inherited type.

A *circuit gallery* (or just plain "circuit") in a chamber system $C = (C, E; \lambda)$ over I is just a circuit in the underlying graph (C, E). The *type of a pointed circuit* $(c_0, c_1, \ldots, c_n)$ is the circular sequence of types

$$\lambda(G) := (\lambda(c_0, c_1), \cdots, \lambda(c_{n-1}, c_n = c_0), \lambda(c_0, c_1)).$$

We say that a gallery $G = (c_0, c_1, \ldots, c_n)$ is *p-reduced* (short for "panel-reduced") if and only if $\lambda(c_{i-1}, c_i) \cap \lambda(c_i, c_{i+1}) = \emptyset$ for $i = 1, \ldots n-1$, and we include the extra requirement that if $c_n = c_0$, so that G is a circuit gallery, then also $\lambda(c_{n-1}, c_n) \cap \lambda(c_0, c_1) = \emptyset$. Thus, in p-reduced galleries and circuits, no two consecutive edges can lie in a common panel.

Panel Homotopy and Reduced Paths

Consider now the collection $\mathcal{C}_1$ of all circuit galleries of a chamber system $C = (C, E; \lambda)$, each of which lies in some residue of rank one (depending on the particular circuit). We call the homotopy with respect to this system of circuits $\mathcal{C}_1$, *panel homotopy* since it is defined by the assertion that all panels are contractible. An example of a panel homotopy of a gallery $G = (x_0, x_1, \ldots, x_{n-1}, x_n)$ is the replacement of G by a gallery

$$G' = (x_0, x_1, \ldots, x_{i-1}, y_i, x_i, \ldots x_{n-1}, x_n),$$

where the two edges (x_{i-1}, y_i) and (y_i, x_i) share a common type. Such a transformation $G \to G'$ or its reverse $G' \to G$ is called an *elementary panel homotopy*.

Like any $\mathcal{C}$-homotopy, it does not change the initial and terminal vertices of the path or gallery.

In the case that C is a chamber system with a *single-valued* type function (that is, exactly one type is assigned to each edge) the discussion is simpler. Every gallery $G = (x_0, x_1, \ldots, x_{n-1}, x_n)$ now has a type consisting of a sequence of types $(\lambda(x_0, x_1), \ldots, \lambda(x_{n-1}, x_n))$ which we render as a word $\lambda(g) := t_1 \cdots t_n$ where $t_i = \lambda(x_{i-1}, x_i)$, with the understanding that $\lambda(G) = \emptyset$ if G has length zero. [We let $\ell(G) = \ell(\lambda(G))$ so that "ℓ" may record the length of a word, as well as a gallery.] Then the type of a concatenation of galleries is just the concatenation of the sequences of types. We now see that if the types of two consecutive edges in a gallery are the same, we can shorten the length-two segment they comprise to one of length one by an elementary panel homotopy. The two consecutive edge types are then coalesced to give the types of the result of the homotopy. This process of coalescing consecutive types in a sequence eventually leads to a shorter sequence in which consecutive edges share no common type. When this is done to the type $\lambda(G)$ of a gallery G, we call the unique result $\hat{\lambda}(G)$ the *p-reduced type* of the gallery G. Thus a gallery of type (11454465776) has unique p-reduced type (14546576). Clearly a gallery G is panel-homotopic to a gallery $\hat{g}$ of p-reduced type (that is, $\hat{g}$ is p-reduced in the basic terminology of the previous section).

When λ is not single-valued, the coalescing process which converts a type to a p-reduced type exists but the p-reduced types obtained may no longer be unique. The lemma below states what we know.

Lemma 9.4.1 *In general, the following hold:*

1. *Every gallery (circuit) is panel-homotopic to a p-reduced gallery (circuit) (as defined in the previous section).*
2. *In the case that the type function λ is single-valued, every gallery is panel homotopic to a gallery of reduced type. Moreover, if gallery G is panel homotopic to gallery H then $\hat{\lambda}(G) = \hat{\lambda}(H)$.*

Proof 1. We prove this part in the case of circuits, but the proof is virtually the same for galleries that are not circuits – just one less edge to consider.

Suppose $G = (x_0, x_1, \ldots, x_{n-1}, x_n = x_0)$ is a circuit with edges $e_1 = (x_0, x_1), \ldots, e_n = (x_{n-1}, x_n)$ of minimal length with respect to not being panel-homotopic to a p-reduced circuit. Then $n > 2$ and G cannot be panel-homotopic to a circuit of shorter length. Then there is a smallest index k such that $\lambda(e_1) \cap \lambda(e_k) = \emptyset$. Then $e_1' := (x_0, x_{k-1})$ is an edge bearing a non-empty set of labels and G is panel-homotopic to the circuit G' defined by the edges $(e_1', e_k, e_{k+1}, \ldots, e_n)$. If $k > 2$, then G' is shorter than G which is impossible. Thus $k = 2$, and so $\lambda(e_1) \cap \lambda(e_2)$ is empty. By a similar argument, $\lambda(e_i) \cap \lambda(e_{i+1}) = \emptyset$, and so G itself is a p-reduced circuit.

2. The statements concerning chamber systems whose type function is single-valued are immediate from the first part and the discussion in the paragraph preceding the statement of the lemma. □

Corollary 9.4.2 *Suppose* $(C, E; \lambda)$ *is a chamber system satisfying these hypotheses:*

(typ) *The type function* λ *assumes only single values on edges.*
(Cir(n)) *Except for the circuits of length zero,* C *has no circuits of p-reduced type of positive length less than* $2n$.

Then every gallery of p-reduced type of length at most n *is a geodesic.*

Proof We first observe that any geodesic path or gallery must have its successive edges labeled by distinct types. Otherwise, it would be panel-homotopic to a shorter gallery connecting the same initial and terminal vertices, and so could not be a geodesic.

Suppose $G = (x_0, \ldots, x_d)$ is a gallery of p-reduced type. Without loss of generality, one may suppose its type to be $\lambda(g) = t_1 t_2 \cdots t_d$, a word of length $\ell(g) = d$ with letters $t_i \in I$ with $t_i \neq t_{i+1}, i = 0, \ldots, d-1$. Now if G is not a geodesic, there is a geodesic gallery $H = (x_0, y_1, \ldots, y_e = x_d)$ of length $\ell(H) = e$ strictly less than d. Then as $d \leq n$, the circuit $F := G \circ H^{-1}$ has length less than $2n$. By Lemma 9.4.1, this circuit is panel-homotopic to a p-reduced circuit $\hat{F}$ of length $\lambda(F) = \ell(\hat{\lambda}(G \circ H^{-1}))$, where

$$\ell(\hat{\lambda}(G \circ H^{-1})) \leq \ell(G \circ H^{-1}) = d + e < n.$$

(Note that the first term is the length of a word, while the second is the length of a gallery.) Then, by (Cir(n)), $\hat{F}$ is length zero, so $G \circ H^{-1}$ must be panel-contractible. It follows from Theorem 1.3.3, p. 1.3.3, Chap. 1, that G is panel-homotopic to H, whence by Lemma 9.4.1, part 2,

$$\hat{\lambda}(G) = \hat{\lambda}(H). \tag{9.25}$$

But the left and right sides are respectively $\lambda(G)$ and $\lambda(H)$ since G and H are both of reduced type. Yet these words have lengths d and $e < d$ respectively, a contradiction. Thus x_d has distance d from x_0 so G is a geodesic gallery. □

Example 8 Consider the following simple chamber system over the typeset $\{1, 2\}$. The graph is a four circuit (a, b, c, d, a) with a diagonal edge (a, c) adjoined. The edges are assigned type sets as follows; edges (a, b) and (b, c) are type $\{1\}$, edges (c, d) and (d, a) are assigned type set $\{2\}$, and edge (a, c) is assigned type set $\{1, 2\}$. Then the panels of type 1 are $\{a, b, c\}$ and $\{d\}$, those of type 2 are $\{a, c, d\}$ and $\{b\}$. Note that there are no non-trivial circuits of p-reduced type of length greater than zero whatsoever.

9.4.2 The Chamber System of a Generalized Polygon

The First Definition of Generalized n-gon

Let $C = (C, E; \lambda)$ be a chamber system over $\{1, 2\}$. Then C is a *generalized n-gon* (n is an integer greater than one or the symbol ∞) if and only if:

(typ) The type function λ assumes only single values on edges.
(Dia(n)) The graph (C, E) has diameter at least n and is firm if $n = \infty$.
(Cir(n)) Except for the circuits of length zero, C has no circuits of p-reduced type of length less than $2n$.
(Ch(n)) For any geodesic gallery G of type $(1, 2, 1, 2, \ldots)$ connecting two chambers x and y at distance n, there is a second geodesic gallery of type $(2, 1, 2, 1, \ldots)$ connecting x to y. A similar statement with the types '1' and '2' transposed is also assumed.

We extend this definition to the case that $n = \infty$ by viewing (1) axiom (Dia(∞)) to mean that C is connected, (2) axiom (Ch(∞)) to be vacuous, and (3) axiom (Cir(∞)) to assert that there are no p-reduced circuits of positive length whatsoever.

If we wish to assert that chamber system C is a generalized n-gon but do not wish to specify n, we say that C is a *generalized polygon*.

Lemma 9.4.3 *Let $C = (C, E; \lambda)$ be a generalized n-gon over type set $\{1, 2\}$. Then the following statements hold:*

1. *C is residually connected.*
2. *Any gallery of length $d \leq n$ is of reduced type if and only if it is a geodesic path of length at most n. When $n = \infty$, all galleries of panel-reduced type are geodesics.*
3. *If $n \neq \infty$, circuits of reduced type and length $2n$ exist. All such circuits are isometrically embedded in the underlying graph (C, E). Every vertex lies in such a circuit. The chamber system is firm.*
4. *Every geodesic path of length d less than n can be extended to one of length $d+1$. As a consequence, every edge is in a circuit of reduced type and length $2n$.*
5. *Suppose n is finite. Then, for any two chambers c and d, there exists a p-reduced circuit of length $2n$ containing them.*
6. *The graph $C = (C, E)$ has diameter n.*
7. *Every panel is strongly gated in (C, E).*

Proof 1. A panel of type 1 cannot intersect a panel of type 2 in two distinct chambers x and y, for otherwise (x, y) would become an edge of type $\{1, 2\}$ against λ being single-valued. Thus every chamber is the intersection of the panels that contain it. Since (Dia(n)) implies C is connected (even when $n = \infty$), we see that C satisfies all the requirements of being residually connected.

2. As commented at the beginning of the proof of Corollary 9.4.2, every geodesic gallery has p-reduced type. Conversely, in view of (typ) and (Cir(n)), Corollary 9.4.2 also tells that any gallery of length at most n is a geodesic if it has p-reduced type. When $n = \infty$, (Cir(n)) holds for every positive integer n, so in that case, the same corollary implies that every path of p-reduced type is a geodesic.

3. Now by (Dia(n)) there exists somewhere a geodesic of length n, and so by (Cir(n)) there is a circuit G of length $2n$ and p-reduced type $1212 \cdots = (12)^n$. Any two antipodal vertices of this circuit – say x_i and x_{i+n} – are connected by a gallery of length n (a segment of G) of p-reduced type. By part 2 just proved, this gallery is a geodesic. It follows that G is isometrically embedded in (C, E).

Let y be any further vertex. We claim that y is distance n from some vertex and so by (Cir(n)) lies in a circuit of p-reduced type $(12)^n$. Now the circuit G of the previous paragraph was itself such a circuit and if y were one of its vertices our claim would be fulfilled. So assume $H = (y = y_0, \ldots, y_d = x_0)$ is a gallery of shortest length connecting y to a vertex x_0 of the circuit G. Then $d \geq 1$ and H has p-reduced type. By transposing the name of the types, if necessary, we assume the edge $(y_{d-1}, y_d = x_0)$ has type "2." Then we name the vertices of G so that $G = (x_0, x_1, \ldots, x_{2n-1}, x_n = x_0)$ in the orientation that produces the p-reduced type $1212\cdots$. Then we may extend H by concatenting it with an initial segment of G to produce a gallery

$$H' = (y_0, y_1, \ldots, y_d = x_0, x_1, \ldots, x_{n-d})$$

of p-reduced type and length n. By part 2, this gallery is a geodesic, and our claim is proved.

Since each vertex lies in a circuit of type $(12)^n$, the panels which contain it have at least two chambers. Thus C is firm.

4. Suppose $G = (x_0, \ldots, x_d)$ were a geodesic path of length d, $0 < d < n$. Then G has p-reduced type. Let $i \in \{1, 2\}$ be the type of its last edge (x_{d-1}, x_d). Since C is firm, x_d lies on an edge $e = (x_d, x_{d+1})$ of type j where $\{i, j\} = \{1, 2\}$. The concatenation $G \circ e$ now has p-reduced type and length at most n and so is a geodesic extending G by part 2.

5. By the diameter assumption, the two chambers c and d are connected by a geodesic path of length $d \leq n$. By iterating part 4, if necessary, this geodesic can be extended so that it is an initial segment P_{cd} of a geodesic path P_{cf} of length n. This extended path P_{cf} has type $12121\cdots$ (n factors), and so by (Ch(n)) there is another path Q_{cf} of type $21212\cdots$ (n factors) connecting c to f, and clearly $P_{cf} \circ (Q_{cf})^{-1}$ is a circuit of panel-reduced type containing the two chambers c and d.

6. Suppose $G = (x_0, \ldots x_n, x_{n+1})$ were a geodesic path of length $n + 1$. Without loss of generality, we may suppose it has type $1212\cdots ij$ where $\{i, j\} = \{1, 2\}$. Now, by (Ch(n)), there is also a gallery $H = (x_0 = y_0, y_1, \ldots y_n = x_n)$ of type $2121\cdots ij$ and length n connecting x_0 and x_n. Since x_n is i-adjacent to both y_{n-1} and x_{n+1}, the latter two are either equal or i-adjacent. But either possibility is prohibited by the fact that x_0 and x_{n+1} are at distance $n + 1$. Thus no such geodesic of length $n + 1$ can exist. So the diameter is at most n which, with (Dia(n)), proves the assertion 6.

7. Suppose P is a panel, say of type 1. Let y be any chamber. We must show that P is strongly gated with repect to y. Since P is a clique it is already isometrically embedded so it is enough to show that P is gated with respect to y when y is not in P. Suppose x_1 and x_2 are two distinct chambers of P achieving the minimal distance d from y to any element of P. Then there is geodesic gallery $G = (y, y_1, \ldots, y_{d-1}, y_d = x_1)$ from y to x_1. Since y_{d-1} is not in P, the last edge (y_{d-1}, x_1) must be type 2. Then the gallery $G \circ (x_1, x_2)$ is of p-reduced type and length $d + 1$. If $d < n$, then by part 2, $G \circ (x_1, x_2)$ is a geodesic, against x_2 being distance d from y. Thus n is finite and by part 5, $d = n$ exactly. In this case (Cir(n))

shows that there is a second gallery $G' = (y, y'_1, \ldots, y'_{n-1}, y_n = x_1)$ of reduced type with the last edge (y'_{n-1}, x_1) being type 1. But in that case, $y'_{n-1} \in P$, against the minimality of d.

Thus P contains a unique chamber nearest y, for any y, and we are done.

All parts of the lemma have been proved. □

When one encounters an omnibus lemma following a definition, one can usually extract other equivalent definitions from various subsets of the listed properties. This case is no exception.

Equivalent Definitions of Generalized Polygon

We first record the following observation.

Lemma 9.4.4 *Let $C = (V, E; \lambda)$ be any chamber system of rank $|I| = 2$. The following statements are equivalent:*

1. *C is residually connected.*
2. *C is connected and each edge of C bears a single label – that is, two adjacent vertices are j-adjacent for a unique $j \in I$.*
3. *C is connected and any two distinct panels intersect in at most one chamber.*

The proof is an easy exercise (part of Exercise 9.10).

In [105] R. Scharlau considers the following axioms for a chamber system C over the type set $I = \{i, j\}$:

(CS0) The chamber system C is residually connected.
(CS1) For any vertex $c \in V$ and $k \in I$, there exists a chamber c' which is k-adjacent to c.
(CS$_n$2) C contains no p-reduced circuit of positive length less than $2n$.
(CS$_n$3) If a gallery from c to d has type type $ijij\cdots$ (n-factors,), then there is also a gallery of type $jiji\cdots$ (n factors) from c to d.

For the purpose of distinguishing axiom systems, we shall refer to the four axioms above as the "Scharlau presentation." Let us compare these axioms with the axioms for a chamber system of the previous section:

- The first axiom (CS0) asserts two things: (1) that C is connected and (2) that any two panels meet in at most one chamber. As a result, an edge connecting two chambers cannot bear both labels i and j. Thus

$$(\text{CS0}) \Rightarrow (\text{typ}).$$

- Axiom (CS1) is just the assertion that C is firm – that is, all panels contain at least two chambers.
- (CS$_n$2) is exactly the assertion (Cir(n)).

- Clearly (CS_n3) implies (Ch(n)). But the latter is conceivably weaker than the former – at least if it is possible to have a gallery of type $ijij\cdots$ of length n which is not a geodesic.

For any chamber system C define the *p-reduced girth* to be the smallest length of a p-reduced circuit of positive length. Then axiom (CS_n2) just says that C either has no p-reduced circuits, or has p-reduced girth at least $2n$.

Now consider the following condition:

(CS_n4) For any two chambers c and d, there exists a p-reduced circuit of length $2n$ containing c and d.

When n is infinite, we interpret axiom ($CS_\infty4$) to assert (CS0) and (CS1) together.

Note that if two chambers x and y of a chamber system over $\{i, j\}$ formed an edge e with $\lambda(e) = \{i, j\}$ — that is, x and y are both i- and j-adjacent – then the edge e could never participate in a p-reduced gallery of length two or more. It follows that (CS_n4) implies (typ) as well as the fact that (C, E) is connected. Thus the following arises.

Lemma 9.4.5 *For a rank two chamber system C, condition (CS_n4) implies condition (CS0) – that C is residually connected and (CS1) that C is firm. (Of course, if n is finite, it also implies that $2n$ is an upper bound for the p-reduced girth.)*

We now have the following.

Lemma 9.4.6 *For a chamber system C of type $\{i, j\}$, the following conditions are equivalent:*

1. *The Scharlau presentation: (CS0), (CS1), (CS_n2) and (CS_n3).*
2. *(CS_n2) and (CS_n4).*
3. *If $n < \infty$, C has p-reduced girth $2n$, and (CS_n4) holds for all n.*

Proof The equivalences are trivial if $n = \infty$. In that case they all assert that C is residually connected and firm and that there are no p-reduced circuits of positive finite length.

So we assume $n < \infty$.

Note that (CS_n2) asserts that the p-reduced girth is at least $2n$, while (CS_n4) implies it is at most $2n$. Thus 2 and 3 are equivalent statements.

First we show that 1 implies 3. Since (CS_n2) implies the p-reduced girth is at least $2n$, and (CS_n4) implies that it is no more than $2n$, it suffices to prove (CS_n4). Now choose two chambers c and d (possibly the same chamber). Then as C is residually connected (condition (CS0)) there is a geodesic G from c to d which is of p-reduced type. Since C is firm it can be extended to a gallery $G \circ H$ of length n of p-reduced type – say $ijij\cdots$ (n factors), transposing i and j, if necessary. Then by (CS_n3) there is a gallery F from c to the terminus of $G \circ H$ of type $jiji\cdots$ (n factors) and $G \circ H \circ F^{-1}$ is the desired circuit of length $2n$ of p-reduced type on c and d. So (CS_n4) holds.

Last, we show that 3 implies 1. Assume 3. By Lemma 9.4.5, (CS_n4) implies (CS0) and (CS1) Also the statement on p-reduced circuits in 3 implies (CS_n2). It remains only to show (CS_n3). Suppose c and d are chambers joined by a gallery of type $ijij\cdots$ (to exactly n factors). Applying (CS_n4) to c and d yields the conclusion of (CS_n3). The proof is complete. □

Now we have the following.

Corollary 9.4.7 *Any chamber system over $I = \{i, j\}$ satisfying any of the three equivalent conditions of Lemma 9.4.6 is a generalized n-gon.*

Proof Assume the Scharlau axioms. If n is infinite, C is connected and firm and (typ) holds. Then C is a generalized polygon. So we assume n is an integer larger than 2.

First, (CS0) implies (typ) and the fact that C is connected. (CS_n2) is the same as $(Cir(n))$ and (CS_n3) implies $(Ch(n))$. It only remains to be shown $Dia(n)$. Since C is connected, it suffices to show that there exists a geodesic of length n and that no geodesic has length $n + 1$.

From firmness, any chamber x is the initial term of a gallery G of length exactly n and p-reduced type $ijij\cdots$. Since (typ) and $(Circ(n))$ hold, G is a geodesic of length n by Corollary 9.4.2.

Now suppose $H = (x_0, x_1, \ldots, x_n, x_{n+1})$ were a geodesic gallery of length $n+1$. Then we can write $H = G \circ (x_n, x_{n+1})$, where G is a geodesic of length n and p-reduced type — say, $ijij\cdots$ (n factors) and ending with type $k = \lambda(x_{n-1}, x_n) \in \{i, j\}$. By (CS_n3) x_0 is also connected to x_n by a gallery $F := (x_0, y_1, \ldots, y_n = x_n)$ of type $jiji\cdots$ (n factors) ending in k', where $\{k, k'\} = \{i, j\}$. Now $\lambda(x_n, x_{n+1})$ is either k or k'. Thus x_{n+1} is distance one from either x_{n-1} or y_{n-1}. In either case, it cannot be distance $n + 1$ from x_0, contradicting the fact that H was a geodesic of length $n + 1$. Thus no such H exists and (C, E) has diameter n exactly. The proof is complete. □

A Metrical Definition of Generalized n-gon

Here is a very important variation due to R. Scharlau.

Lemma 9.4.8 (R. Scharlau [105].) *Suppose $C = (C, E; \lambda)$ is a rank two chamber system over $I = \{i, j\}$. Then C is a generalized n-gon, $n > 1$ if and only if the following hold:*

(typ) *The type function is single-valued.*
(F) *C is firm.*
(PG) *C is connected and each panel is gated.*

Remark When $n = \infty$ we interpret $Diam(n)$ to assert that C is connected and that arbitrarily large distances between points occur.

Proof If C is a generalized n-gon, then (typ) holds *a fortiori* and the remaining properties hold by Lemma 9.4.1.

Suppose C is a rank-two chamber system with all of the properties, (typ), (F), (PG). If there are no p-reduced circuits of positive length, C is a generalized ∞-gon.

So we may assume there exists a p-reduced circuit of minimal possible length — say $E = (c_0, c_1, \ldots c_{n-1}, c'_n, c'_{n-1}, \ldots, c'_1, c_0)$ (the notation is chosen so that c_i and c'_{n-i} occupy antipodal positions in the circuit, $i = 0, \ldots, n-1$. (Note that in rank two, any p-reduced circuit must have even length; so $2n$ is the p-reduced girth.) By the minimality of n, we have the condition $\mathrm{Cir}(n) = \mathrm{CS}_n 2$. Since (typ) is assumed, Lemma 9.4.1 implies

Step 1. *Every p-reduced gallery of length at most n is a geodesic gallery.*

Next we shall show the following.

Step 2. *For any chamber d at distance at most n from c_0, the two chambers c_0 and d lie together in a circuit of p-reduced type and length $2n$.*

This result is trivial if d is on the circuit E. Suppose $G = (c_0, d_1, d_2, \ldots, d_s)$ is a geodesic gallery from c_0 to d where we assume that $s = d_C(c_0, d) \leq n$. Now the type of (c_0, d_1) must agree either with that of (c_0, c_1) (say, type i), or with that of (c_0, c'_1) (say, type j). Without loss of generality we assume d_1 is i-adjacent to both c_0 and c_1, while (c'_1, c_0) and (c_1, c_2) are j-adjacencies. In particular, $d_1 \neq c'_1$. In Fig. 9.6 we have tried to represent the "general position" of this configuration. Note, however, that it is possible that $d_1 = c_2$ and after that $d_2 = c_3$, — and that in fact *the geodesic G from c_0 to d_s may wrap nearly around the bottom half of the circuit E in Fig. 9.6.*

Now, for each t with $1 \leq t \leq s$, the gallery

$$G_t := (d_t, d_{t-1}, \cdots, d_1, c_0, c'_1, c'_2, \cdots c'_{n-t})$$

is p-reduced of length n and hence by Step 1 is a geodesic. We make the following statement:

(Claim) *Given G_t, $1 \leq t \leq s$, there exists a gallery G_t^* from d_t to c'_{n-t} whose type is "opposite" G_t (that is i and j have been transposed).*

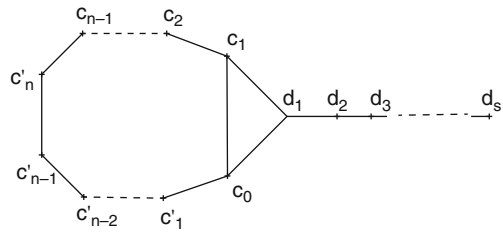

Fig. 9.6 The configuration for Step 2. Here the minimally chosen circuit for this step is $E = (c_0, c_1, \ldots c_{n-1}, c'_n, c'_{n-1}, \ldots, c'_1, c_0)$ (the notation is chosen so that c_i and c'_{n-i} occupy antipodal positions in the circuit, $i = 0, \ldots, n-1$)

We prove the claim by induction on t. To prove it at the induction step t to $t+1$, we replace E by the circuit $E' := G_t \circ (G_t^*)^{-1}$, and apply induction. So it suffices to prove this assertion when $t = 1$.

Now if d_1 is in e there is nothing to prove. So we see that d_1 is distance one from both c_0 and c_1 by the same i-panel. Thus by Step 1, d_1 is distance n from both c'_n and c'_{n-1}, which are distinct elements of some k-panel P. Thus by (GP) there is a chamber p in P at distance $n-1$ from d_1. Let H be a geodesic of this length from p to d_1. Then $G_1 := (d_1, c_0, c'_1, \ldots c'_{n-1})$ and $H^{-1} \circ (p, c_{n-1})$ is the desired G_1^*.

Thus the claim is proved. But now the claim implies the assertion of Step 2.

Step 3. *Any two chambers at distance at most n from one another lie in a common p-reduced circuit of length* $2n$.

Suppose x and y are chambers at distance $s \leq n$. Since C is connected and $n \geq 1$, there is a finite sequence of chambers $(c_0 = x_0, x_1, \ldots, x_{m-1} = x, x_m = y)$ such that x_i has distance at most n from x_{i+1}. Now by Step 1 there is a p-reduced circuit of length $2n$ on both c_0 and x_1, and by iterating Step 1 with x_i in the role of c_0, there is such a circuit on $x_{m-1} = x$ and $x_m = y$.

Step 4. *Axiom* (CS_n4) *holds. That is, any two chambers lie in a p-reduced circuit of length* $2n$.

In view of Step 3, it suffices to show that there are no pairs of chambers at distance greater than n – equivalently (since C is connected), that no geodesic paths of length $n+1$ exist.

Suppose by way of contradiction that there are two chambers x and y at distance $d_C(x, y) = n+1$. By interposing the penultimate chamber z of a geodesic from x to y, we have $d_C(x, z) = n$ and z is i-adjacent (say) to y. Now by Step 3, there is a p-reduced circuit on x and z, say

$$E := (x = x_0, x_1, \ldots, x_n = z, x_{n+1}, \ldots, x_{2n-1}, x_{2n} = x).$$

Now of the two edges of the circuit on z, one is labelled i and carries a chamber $z' = x_{n-1}$ or x_{n+1} at distance $n-1$ from x. But z' is i-adjacent to y by the fundamental property of a chamber system. That conclusion contradicts $d_C(x, y) = n+1$. Thus no geodesics of that length exist, and the proof of Step 4 is complete.

We conclude at this stage that property 2 of Lemma 9.4.6 holds. It follows that C is a generalized polygon by Corollary 9.4.7. □

Remark The minimal p-reduced circuits of length $2n$ which appear throughout the proof of Scharlau's metrical characterization theorem are called *apartments* of the generalized polygon. Basically, that proof establishes that any two chambers lie in a common apartment. Later we shall meet this notion of "apartment" in the general context of buildings. Again we will be interested in the property that any two chambers lie in an apartment.

9.4.3 Generalized n-Gons as Geometries

Having noted that a generalized n-gon is a residually connected chamber system C, one may ask what special properties characterize the geometry $\mathbf{\Gamma}(C)$. Here the definition is especially simple.

In an arbitrary graph, the girth is the minimal length of a circular walk of positive length having no repeated edges.[10]

A geometry over $\{1, 2\}$ is called *a generalized n-gon geometry* if and only if:

(GP1) It is a bipartite graph of girth $2n$.
(GP2) It has diameter n.
(GP3) If n is infinite, then every vertex is on at least two edges.

When n is infinite we interpret the first two axioms in the following way: (1) the girth assumption (GP1) becomes the assertion that there are no circuits (circular tours) of positive length, and (2) the diameter assumption (GP2) is read as the assertion that the graph is connected. Thus if $n = \infty$, Γ is simply a tree with no degree-one vertices ("leaves").

Note that if $\Gamma = (\mathcal{P}, \mathcal{L})$ is a generalized n-gon geometry, then Γ is connected because it has diameter n (as understood allowing n to be infinite). The second observation is that the properties listed make no reference to any particular component of the bipartite graph. That is, if we choose to view this bipartite graph as a rank two incidence geometry of points and lines, then any consequence of the axioms stated in this language would also be true if lines and points were transposed, while preserving incidence. In a few words: *if* $(\mathcal{P}, \mathcal{L})$ *is a generalized polygon, then so is its dual geometry,* $(\mathcal{L}, \mathcal{P})$.

From now on, we view this graph as the incidence graph of a point-line geometry $\Gamma = (\mathcal{P}, \mathcal{L})$. If we wish to say that Γ is a generalized n-gon for some $n \in \mathbf{N} \cup \{\infty\}$, but otherwise do not wish to specify n, we say that Γ is a *generalized polygon geometry*.

We begin with an observation. Suppose Γ is any (simple) bipartite graph viewed as a rank two geometry. Then the chamber system $\mathbf{C}(\Gamma)$ is simply the edge graph of Γ with the edge-labelling recording in which component a vertex of two intersecting edges lies. Then, of course, axiom (typ) holds for $C(\Gamma)$.

Conversely, suppose C is a chamber system of rank two. Then, of course, $\mathbf{\Gamma}(C)$ is the bipartite graph recording when panels of distinct type may intersect non-trivially. We now notice the following very elementary facts.

Lemma 9.4.9 *Suppose* Γ *and* C *are respectively a rank two geometry and a rank two chamber system over the same index set* $I = \{i, j\}$. *Assume either (i)* $C = \mathbf{C}(\Gamma)$, *or (ii)* $\Gamma = \mathbf{\Gamma}(C)$. *Then the following statements hold:*

[10] Some books call such walks a circular tour. Note that the "backtracks" of Chap. 1 are excluded by this requirement. Otherwise all graphs would seem to have girth two!

1. Γ *is connected if and only if* C *is.*
2. *Assume that* C *has (typ). Then* Γ *has no vertices of degree one if and only if* C *is firm.*
3. *Assume* C *is residually connected and firm. Then there is a bijection between the p-reduced circuits of* C *and the circular walks of* Γ *whose succesive edges are distinct.*

Proof The first statement is trivial.

Consider the second statement. Suppose $C = \mathbf{C}(\Gamma)$. Then each panel of C is the collection of all edges on a vertex of Γ so the stated equivalence follows. Now suppose $\Gamma = \mathbf{\Gamma}(C)$. If P were a vertex of Γ on only one edge then it would be a panel of type i, say, intersecting only one panel R of type j. If C were firm each of its chambers would be j-adjacent to another distinct chamber forcing $P \subseteq R$. Since $|P| > 1$, the assumption (typ) (here artificially imposed) must fail. Thus (typ) and firmness of C imply that each vertex of Γ is on at least two edges. Conversely, suppose each vertex of Γ has degree at least two. Then each panel of one type meets at least two other panels of the other type, and so contains at least two chambers (even without assuming (typ)). Thus C is firm.

For the third statement, Γ and C are both derived from each other by the functors (see Lemmas 9.2.6 and 9.2.7). Now any circular walk with successive edges distinct becomes a sequence of edges with any two successive members of the sequence – but never three successive members – sharing a common vertex. This becomes a circuit in the edge graph which is a p-reduced circular gallery. Conversely, any circular p-reduced gallery of C has its edges corresponding to panels of types i and j successively alternating. These panels are thus vertices in Γ forming a circular walk with successive edges distinct. □

Theorem 9.4.10 *The following statements hold:*

1. *A generalized polygon geometry* $\Gamma = (\mathcal{P}, \mathcal{L})$ *is always residually connected, and its associated chamber system* $\mathbf{C}(\Gamma)$ *is a generalized polygon.*
2. *If the chamber system* $C = (C, E; \lambda)$ *is a generalized polygon, then its associated geometry* $\mathbf{\Gamma}(C)$ *is a generalized polygon geometry.*

Proof First we prove the result when the parameter n in each definition is infinite. In this case, a generalized polygon geometry Γ is a tree with no vertices of degree one. So it is connected and each vertex is in a flag-chamber (that is, an edge). Thus it is residually connected. Then by Lemma 9.4.9, $C = \mathbf{C}(\Gamma)$ is residually connected, firm, and has no p-reduced circuits of finite positive length.

Similarly, when $n = \infty$, a generalized polygon $C = (C, E; \lambda)$ is residually connected, firm, and possesses no p-reduced circuits of finite positive length. Then $\Gamma = \mathbf{\Gamma}(C)$ is residually connected and by Lemma 9.4.9 has no circular walks with successive edges distinct except for the trivial walk (and so is a tree) and has no vertices of degree one.

From now on we assume the parameter n is finite in both statements.

The proof of the first statement proceeds through a series of five steps.

Assume Γ is a generalized polygon geometry with parts $\mathcal{P}$ and $\mathcal{L}$. Let $e = (a_1, a_2)$ and $f = (b_1, b_2)$ be two distinct edges of the bipartite graph Γ. Of the four possible distances $d_\Gamma(a_i, b_j)$, choose notation for these vertices so that $d_\Gamma(a_1, b_1)$ is the smallest distance d. Since Γ is bipartite, $d_\Gamma(a_1, b_2) = d_\Gamma(a_2, b_1) = d + 1$. Then there are two cases: (1) $d_\Gamma(a_2, b_2) = d + 2$ and (2) $d_\Gamma(a_2, b_2) = d$. In case (1), for any geodesic path G from a_1 to b_1, the path $(a_2, a_1) \circ G \circ (b_1, b_2)$ is a geodesic path of length $d + 2$, and so $d \leq n - 2$. In case (2), there are geodesic paths G and H of length d connecting a_1 to b_1 and a_2 to b_2, respectively and $W = (a_2, a_1) \circ G \circ (b_1, b_2) \circ H^{-1}$ is a circular walk with at least two distinct edges. It follows that some subset of the edges forms a walk with consecutive edges distinct, which necessarily has length at least $2n$ by the girth assumption. Thus the length of W, which is $2d+2$, is bounded below by $2n$. On the other hand $d \leq n-1$, since the graph diameter n is at least $d_\Gamma(a_1, b_2)$. Thus $d = n - 1$ exactly. Then the girth assumption reveals that the circular walk W is itself a minimal circuit of length $2n$. Thus we conclude the following.

Step 1. *Given any two distinct edges e_1 and e_2 of a generalized polygon geometry Γ of diameter $n < \infty$, either (1) they are the first and last edges of a geodesic path, or (2) they are antipodal edges of a minimal circuit of length $2n$.*

Next, observe the following.

Step 2. *In a generalized n-gon geometry every vertex is on at least two edges.*

We had assumed this when n is infinite, but must now prove it in the finite diameter case. Because we then have finite girth, the collection T of vertices on at least two edges is non-empty. Suppose by way of contradiction that x is a vertex on a unique edge $e = (x, y)$ (there is at least one edge on x since Γ is connected and $n \geq 1$). Choose $t \in T$ so that the distance $d = d_\Gamma(x, t)$ is maximal. Now t is on two distinct edges, $e_i := (t, s_i)$, $i = 1, 2$, and each s_i is at distance $d - 1$ from x and distance $d - 2$ from y from the minimality of d and the uniqueness of y. That means there is a circular walk of length $2(d - 2) + 2 = 2d - 2$ incorporating the distinct edges e_i each exactly once and so there is some circular walk with consecutive edges distinct of length smaller than this, but length at least four. The girth assumption forces $2d - 2 \geq 2n$ so $d > n$, contrary to the assumption on diameter.

Step 3. *In a generalized n-gon geometry any geodesic path of length $d < n$ can be extended to one of length $d + 1$.*

This is an easy step. Consider a geodesic path G from u to v of length $d < n$. By Step 2, there are at least two edges on v. If G cannot be extended, v is adjacent to two vertices, each of distance $d - 1$ from u. Their geodesics to u, together with the path of length 2 from one to another through v, can be assembled into a circular walk of length $2(d - 1) + 2 = 2d$ having two edges used just once. Thus the girth assumption yields $2d \geq 2n$, against $d < n$.

Now let us return to Step 1. Suppose e_1 and e_2 are two edges of Γ which are the first and last edges of a geodesic path G. If this path G has length less than d, then by Step 3, it can be extended to a geodesic path $P := G \circ H$ of length n. Otherwise it already has length n and we write $G = P$. In either case P is a geodesic path with terminal vertex t and initial vertex s incident with edge e_1. By Step 2, t is on an edge $f = (t, r)$ not in path P, and it is easy to see that edges e_1 and f are in Case (2) of Step 1. This means there is a minimal circuit of length $2n$ incorporating the path P. We have just proved the following.

Step 4. *If Γ is a generalized n-gon geometry with n finite, then the rank two chamber system $\mathbf{C}(\Gamma)$ satisfies axiom (CS_n4) – any two chambers are on a p-reduced circuit of length $2n$.*

We can now conclude with the following.

Step 5. *If Γ is a generalized n-gon geometry with n finite, then the rank two chamber system $\mathbf{C}(\Gamma)$ is a generalized n-gon.*

Since Γ has girth $2n$ the bijection in Lemma 9.4.9, part 3, shows that a minimal p-reduced circuit of $\mathbf{C}(\Gamma)$ has length $2n$. Thus $\mathbf{C}(\Gamma)$ now satisfies (CS_n2) and (CS_n4). By Lemma 9.4.6, $\mathbf{C}(\Gamma)$ is a generalized n-gon.

The second statement of the theorem is easier to prove. Assume the chamber system C over $I = \{1, 2\}$ is a generalized n-gon, for finite n. Then C is residually connected and the geometry $\Gamma := \mathbf{\Gamma}(C)$ is also residually connected. Since C contains panel-reduced girth $2n$, the bipartite graph Γ also has girth $2n$ by Lemma 9.4.6. It follows that Γ has diameter at least n, and it remains to be shown that it is no larger. By way of contradiction, suppose $g = (x_0, x_1, \ldots, x_n, x_{n+1})$ were a geodesic path of length $n + 1$ in the graph Γ. Letting e_i denote the edge (x_i, x_{i+1}), we obtain a gallery $G = (e_0, e_1, \ldots, e_n)$ of length n and alternating type – say $1212\cdots i$. Then by (CS_n3) there is a second gallery $G' := (e_0, f_1, f_2, \ldots, f_{n-1}, e_n)$ of type $2121\cdots j$ where $\{i, j\} = \{1, 2\}$. Now the f_i are edges of Γ and, because e_0 and f_1 are 2-adjacent, f_1 contains point x_0 rather than x_1. Similarly, the edge f_{n-1} cannot contain x_n, and so must contain x_{n+1}. Thus putting $f_1 = (x_0, y_1)$, $f_i = (y_{i-1}, y_i)$, $i = 2, \ldots n - 2$ and $f_{n-1} := (y_{n-2}, x_{n+1})$, we have a walk $(x_0, y_1, \ldots y_{n-2}, x_{n+1})$ of length $n - 1$, a contradiction to the assumption that x_0 and x_{n+1} were at distance $n + 1$. Thus Γ has diameter exactly n and Γ is a generalized n-gon geometry.

The proof is complete. □

Generalized Polygons as Point-Line Geometries

As previously noted, a generalized polygon geometry can be regarded as the bipartite incidence graph Γ of a geometry $\Gamma = (\mathcal{P}, \mathcal{L})$ of points and lines. The use of the symbol Γ in both cases is not an abuse of notation, for the expressions attached to the symbol in both cases really represent the same thing. The difference in the two views is just a matter of metaphysical constructions. Normally, we think of G

as a bipartite graph. But if there are no repeated lines, we can also view it as a set of points, together with a distinguished collection of subsets of the point set, namely the lines. This change in viewpoint introduces a change in the way we describe things. For example, one now has notions such as partial linear space, subspace, and more importantly, the metrical features of a new graph – the point-collinearity graph $\Delta = (\mathcal{P}, \sim)$. In this section we wish to reconsider generalized n-gon geometries for $n > 2$ from the point-line point of view.

If the parameter n is infinite, a generalized n-gon geometry Γ is a tree with no leaves. In that case each point is on at least two lines, and each line has at least two points, but otherwise these local incidence numbers may vary wildly from point to point and from line to line.

If $n = 2$, the graph Γ is a complete bipartite graph, so every point of Γ is incident with every line and vice versa. Then we say that Γ is a *generalized digon*.

If $n = 3$, any two distinct lines of Γ are vertices of Γ at distance 2 – the same is true of two points of Γ. Thus any two lines are incident with at least one common point and any two points are incident with at least one line. But since there are no 4-circuits in Γ (that is, no circular walks of length 4 which are not backtracks) the adjective phrases "at least one" of the previous sentence can be replaced by "a unique." Thus Γ is a partial linear space in which any two distinct points are on a unique line (i.e., it is a linear space) and any two distinct lines are together incident with a unique point. By now you will recognize that Γ is a "generalized projective plane" as defined in Chap. 5.

Similarly, if $n = 4$, Γ is a generalized quadrangle as defined in Chap. 7.

For general finite n, a point p and a line L are vertices of the bipartite graph Γ at odd distance from one another. If n is even this distance $d_\Gamma(p, L)$ is less than n so there is a unique geodesic path of Γ connecting them. Thus we have the following.

Theorem 9.4.11 *If n is an even number $2k$, and $\Gamma = (\mathcal{P} \cup \mathcal{L}, *)$ is a generalized polygon geometry, then, as a point-line geometry, $\Gamma = (\mathcal{P}, \mathcal{L})$ has these properties:*

(GPL0) Every point is on at least two lines and each line has at least two points (the latter is required of any point-line geometry).
(GPL1) The point-collinearity graph $\Delta = (\mathcal{P}, \sim)$ of Γ has diameter k.
(GPL2) If p is a point, and L is a line, then there is a unique point q of L nearest p in the metric of Δ and the geodesic path in Δ from p to q is unique.

These properties should remind the reader of the form of the axioms for a (non-degenerate) generalized quadrangle. Similarly we also have the following.

Theorem 9.4.12 *Any point-line geometry $\Gamma = (\mathcal{P}, \mathcal{L})$ satisfying the axioms (GPL0), (GPL1), and (GPL2) is a generalized $2n$-gon geometry.*

The proofs of both of these theorems are left as exercises.

A Characterization of Generalized Polygons

P. Abramenko and H. Van Maldegham [1] proved the following theorem characterizing generalized polygons among the classes of rank-two geometries.

Theorem 9.4.13 *Let $\Gamma = (\mathcal{P} \cup \mathcal{L}, *)$ be (the incidence graph of) a connected firm rank two geometry of finite diameter n, $n \geq 3$. Suppose:*

(AV) *For each pair of vertices x and y at distance $n-1$ in Γ, there is a unique vertex adjacent to y which is at distance $n-2$ from x.*

Then Γ is a generalized n-gon.

Corollary 9.4.14 *Suppose a point-line geometry $(\mathcal{P}, \mathcal{L})$ and its dual $(\mathcal{L}, \mathcal{P})$ are both (connected) near n-gons. Then both are in fact generalized n-gons.*

Remarks (a) At first sight, the condition (AV) seems to do little to prevent the bipartite graph Γ from possessing circuits of length $2d$ smaller than $2n$. But, as we shall see, it is actually a powerful hypothesis.

(b) The proof in the Abramenko-Van Maldeghem paper [1] involves 1-twinning[11] (a special "opposite" relationship among flags) in the "flag graph for Γ" – that is, the chamber system graph $\mathbf{C}(\Gamma)$ with the edge-labels ignored. In order to avoid the introduction of concepts which would be used in this book in only one instance, we venture to present below an alternative proof of the Abramenko–Van Maldeghem theorem, using only the incidence graph $\Gamma = (\mathcal{P} \cup \mathcal{L}, *)$.

Proof of Theorem 9.4.13 Throughout, $\Gamma = (\mathcal{P} \cup \mathcal{L}, *)$ is a connected bipartite graph (with parts $\mathcal{P}$ and $\mathcal{L}$) of diameter n, satisfying the axiom (AV), and having each vertex on at least two edges (firmness). The term "geodesic path" will refer to a path whose length is the graph-theoretic distance between its initial and terminal vertices.

In order to indicate exactly how the axiom (AV) is being used in each instance, for a pair of vertices x and y with $d_\Gamma(x, y) = n-1$, we consider the following assertion:

$[AV](x, y)$: *There is exactly one vertex at distance $n-2$ from x that is adjacent to y.*

The proof proceeds by a series of steps.

Step 1. *Suppose $p = (x_0, x_1, \ldots, x_n)$ is a geodesic path of length n. Suppose $y_1 \in x_0{}^{\perp} - \{x_1\}$. Then $d(y_1, x_{n-1}) = n$.*

[11] The important notion of 1-twinning, due to Mühlherr [90], is studied in a wider context in the beautiful paper of Abremenko and Van Maldeghem. This rich paper contains much more than Theorem 9.4.13 and its Corollary 9.4.14.

Proof Since Γ is bipartite, we know $d(y_1, x_n) = n-1$ so $d(y_1, x_{n-1}) = n$ or $n-2$. If $d(y, x_{n-1}) = n-2$, then, noting that $d(x_0, x_{n-1}) = n-1$ and that $x_1 \neq y_1$, the axiom $[AV](x_{n-1}, x_0)$ is violated. □

Step 2. *Every vertex is distance n from some other vertex.*

Proof Immediate from Step 1 and the connectedness of Γ. □

Step 3. *Let $p = (x_0, x_1, \ldots, x_n)$ and y_1 be as in Step* 1. *Then there exists an isometrically embedded circuit C of length $2n$ containing $(y_1, x_0, \ldots, x_n)$ as a segment.*

Proof In general, a (pointed) circuit $Z = (z_0, \ldots, z_{2n} = z_0)$ of length $2n$ is isometrically embedded in Γ if and only if, for each vertex x_i, $i = 0, \ldots, n-1$, one has $d_\Gamma(x_i, x_{i+n}) = n$.

Now we repeatedly apply Step 1. We presently have, $d(y_1, x_{n-1}) = n$. Next choose $y_2 \in {y_1}^\perp$ so that $d(y_2, x_n) = n-2$ (this is possible because there is a geodesic path of length $n-1$ proceeding from y_1 to x_n). Then $y_2 \neq x_0$ and by Step 1 (with y_2, y_1, and $(y_1, x_0, \ldots, x_{n-1})$ respectively replacing y_1, x_0 and path p) we have $d(y_2, x_{n-2}) = n$.

That was only a model for the following inductive step. Suppose, for some integer d, with $2 < d < n-1$, we have obtained a path $(y_1, y_2, \ldots, y_d)$ with $d(y_i, x_n) = n-i$, and $d(y_i, x_{n-i}) = n$, for $i = 1, 2, \ldots d$. Since $d(y_d, x_n) = n-d$, there exists a vertex y_{d+1} adjacent to y_d and at distance $n-d-1$ from x_n. Now apply Step 1 with y_{d+1}, y_d and $(y_d, \ldots, y_1, x_0, \ldots, x_{n-d})$ in the respective roles of y_1, x_0, and p, to conclude that $d(y_{d+1}, x_{n-d-1}) = n$.

In this way one forms a circuit,

$$C := (x_0, x_1, \ldots, x_n, y_{n-1}, y_{n-2}, \ldots y_2, y_1, y_0 = x_0),$$

in which antipodal pairs of points $(y_i, x_{n-i}), i = 0, 1, \ldots x_{n-i}$ are all at distance n. Thus C is isometrically embedded and contains $(y_1, x_0, \ldots, x_n)$ as a continuous segment. □

Step 4. *Every geodesic path of length k, $1 \leq k < n$, is a segment of an isometrically embedded $2n$-circuit.*

Proof First consider the case $k = 1$. Suppose (x_0, x_1) is an edge. By Step 2, there is a geodesic path of length n, having x_0 as an initial vertex, say $r = (x_0, z_1, z_2, \ldots, z_n)$. If $x_1 = z_2$, then p is an initial segment of r. Otherwise by Step 1, $d(x_1, z_{n-1}) = n$, and $p^{-1} = (x_1, x_0)$ is an initial segment of $r' := (x_1, x_0, z_1, \ldots, z_{n-1})$. Thus either p or p^{-1} is a segment of a geodesic path of length n, which by Step 3 is a segment of an isometrically embedded circuit of length $2n$.

Now suppose, by way of contradiction, that the set $\mathcal{Y}$ of geodesic paths that are not segments of any isometrically embedded $2n$-circuit is not empty. Among such geodesic paths, choose one, say $p = (y, x_0, \ldots, x_{d-1})$, of minimal length d. From the previous paragraph, we know that $d > 1$. By the minimality of d, the segment $s = (x_0, x_2, \ldots x_{d-1})$ lies in an isometrically embedded $2n$-circuit $C = (x_0, \ldots, x_{2n} = x_0)$. If y is a vertex of C, then $y = x_{2n-1}$, and p is a segment of C contrary to our choice. Thus $y \in {x_0}^{\perp} - C$ and we may apply Step 1 (with y, x_0, and $(x_0, \ldots, x_n)$ in the the respective roles of y_1, x_0, and p of Step 1) to conclude that $(y, x_0, \ldots, x_{n-1})$ is a geodesic of length n, which by Step 3, is a segment of an isometrically embedded $2n$-circuit, C'. Thus p is a segment of C', contrary to the choice p as a member of $\mathcal{Y}$. One must conclude that $\mathcal{Y}$ is empty, and Step 4 is proved. □

Step 5. *The graph* Γ *has girth* $2n$.

Proof From Step 3, Γ contains circular tours (that is, pointed circuits with all vertices distinct) of length $2n$. Suppose, by way of contradiction that Γ had girth less than $2n$. Then there exist circular tours with lengths between 2 and $2n$. Among these, select a circular tour $D := (z_0, \ldots, z_2 d)$ of minimal possible length $2d$ (the length is even since Γ is bipartite), where $2 \leq d < n$.

We claim that $d(z_i, z_i + d) = d$ (where the indices are to be read modulo $2d$) so that D is isometrically embedded. Without loss of generality, we may take $i = 0$. Suppose, by way of contradiction, that $d(x_0, x_d) = e < d$. Then there is a geodesic path, $q := (x_0, t_1, \ldots, t_e = x_d)$, as well as the "half-cycle" path $D_1 := (x_0, x_2, \ldots x_d)$. Clearly the paths from x_0 to x_d are different, and so, as we proceed out from x_0, there is a first instance in which the path q departs from D_1 – say $t_i = x_i$ for $i \leq k$, but $t_{k+1} \neq x_{k+1}$. Then, after that divergence of paths, there is a *next* instance at which the two paths join up – say when $t_j = x_m$, $\min(j, m) > k + 1$. Then

$$T = (t_k, t_{k+1}, \ldots, t_j = x_m, x_{m-1}, x_{m-2}, \ldots, x_k = t_k)$$

is a circular tour of length

$$(j - k) + (m - k) < 2d. \tag{9.26}$$

(Note that, since $e < d$, this equation holds even in the extreme case when $k = 0$ and $m = d$ (or equivalently, $j = e$).) But this contradicts the minimality of d. Thus we have $d(x_0, x_d) = d$ as well as a similar result for all antipodal pairs of vertices of D. Thus D is isometrically embedded.

Now, using Step 4, we can extend D_1 to a geodesic path $(x_0, \ldots, x_d, y_{d+1}, \ldots, y_n)$. Then $d(x_0, y_{n-1}) = n - 1$, while, $d(x_{2d-1}, y_{n-1}) = d(x_1, y_{n-1}) = n - 2$, against $[\mathbf{A}V](y_{n-1}, x_0)$. Step 4 is proved. □

Now it follows that Γ satisfies axioms (GP1)–(GP3) of p. 334 with n finite, and so is a generalized n-gon geometry. □

9.4.4 Existence of Generalized Polygons

As we have seen, the class of generalized ∞-gons is bijective with the class of trees with no degree-one vertices. So existence is assured here.

Similarly, when $n = 2$, we are just dealing with the class of complete bipartite graphs with at least two vertices in each component part, so existence is assured here also.

Generalized n-Gons Having Both Thick and Thin Lines

Recall that any vertex of the incidence graph $\Gamma = (\mathcal{P} \cup \mathcal{L}, *)$ of a rank two geometry is said to be *thick* if it is on at least three edges. The object (point or line) is said to be *thin* if it is on exactly two edges. Of course, if Γ is firm, as in the case of generalized n-gons, there are no vertices on just one edge, so all vertices are either thick or thin.

In a beautiful paper by Arthur Yanushka [147], it is shown that all generalized polygons of finite diameter having both thin and thick objects are obtained from a thick generalized polygon by a certain construction. We describe this theory next.

For any finite integer $n > 1$, we have defined a generalized n-gon as a bipartite graph of diameter n and girth $2n$. The *girth*, one may recall, is the minimal length of a *circular tour* – that is, a circular path with no edge repeated. With that understanding there is no need to exclude bipartite graphs with multiple edges, for a graph has multiple edges if and only if it has girth 2.

With this insight, one may define a *generalized* 1-*gon* to be a bipartite graph of diameter 1 and girth 2 – that is, exactly two vertices connected by at least two (possibly infinitely many) edges. (Note the 1-gon is thick, if the two vertices are connected by at least three edges.)

Now we construct a generalized polygon with possible thin objects from a thick polygon. Suppose $\Gamma = (\mathcal{P} \cup \mathcal{L}, *)$ is a thick generalized n-gon. Fix a positive integer k. We construct a generalized kn-gon from the following recipe:

1. Each edge e of the bipartite graph Γ is replaced by a path $p(e)$ of length k whose extremal vertices are those of the original edge e.
2. If e_1 and e_2 are distinct edges of Γ, the two paths $p(e_1)$ and $p(e_2)$ shall share no vertices except possibly the extremal vertices – that is, the original thick vertices of e_1 or e_2.

The result of this replacement is a new bipartite graph $k\Gamma$ whose vertices are the old (thick) vertices of Γ and the new (thin) vertices that are non-extremal vertices of one of the paths $p(e)$, as e ranges over all edges of Γ. Each tour of length ℓ in Γ is then converted to a tour of length $k\ell$ in $k\Gamma$, and any pointed circular tour of Γ becomes a circular tour of length $2k\ell$, pointed at a thick vertex. It follows that the graph $k\Gamma$ is bipartite of girth $2kn$ and diameter kn, and so is a generalized kn-gon.

Of course, if $k = 1$, the edges of Γ remain unadulterated, and the graphs $k\Gamma$ and Γ coincide.

Let us examine 3Γ when Γ is the generalized 1-gon consisting of two vertices connected by three edges. The result is a generalized 3-gon with two thick objects

and six thin ones. Now (unlike the original 1-gon) 3Γ possesses an interpretation as a point-line geometry: namely the (self-dual) projective plane consisting of one thick point incident with three thin points, and one thick point not on that thick line.

This brings us to the remarkable theorem of A. Yanushka [147].[12]

Theorem 9.4.15 (A. Yanushka.) *Suppose* $\Gamma = (\mathcal{P} \cup \mathcal{L}, *)$ *is a generalized n-gon with n finite, having both thin and thick vertices. Then the thick vertices form a generalized m-gon* Γ_0 *where* $n = km$ *and* Γ *is isomorphic to* $k\Gamma$.

The argument is that if two thick vertices of Γ are connected by a path of lenth k, all of whose non-extremal vertices are thin, then all such paths connecting two thick vertices have this length. The geometry Γ_0 is then recovered by replacing each such path (whose extremal vertices are thick, and whose non-extremal vertices are thin) by a simple edge uniquely defined by the path. In this way all thin vertices are erased.

Theorem 9.4.15 reduces the existence question to thick polygons.

Thick Generalized Polygons

For $2 < n < \infty$ we divide the existence question into two parts: (1) the case that $|\mathcal{P}|$ is finite and (2) the case of infinitely many points. We consider the finite case first.

We say that a finite generalized n-gon geometry ($2 < n$) has *order* (s, t) if and only if each point is incident with $t + 1$ lines and each line is incident with $s + 1$ points. The order is forced when all points and lines are thick (that is, they are incident with at least three other objects). When thin objects are allowed, as discussed in the previous section, there may be no order. For example:

1. A generalized 3-gon is a generalized projective plane which may have both thick and thin lines. If both types of lines occur, the 3-gon has no order.
2. We have seen that there are cases where a finite generalized quadrangle does not possess an order. For example, if every point is on just two lines, the quadrangle may be a grid with two distinct line sizes.
3. A similar absence of an order may occur in hexagons and octagons which have thin lines or thin points.

Generalized Polygon Geometries of Finite Order (s, t)

This much is known: if all lines are thick and all points are thick (that is, they lie on at least three distinct lines) then the finite generalized polygon does possess an order (s, t). Of course, the converse fails: there *are* polygons of order (s, t) with $s = 1$ or

[12] Although the context of the paper seems to be finite polygons, the proof of Theorem 9.4.15, which is confined to the structural relation between thick and thin vertices, seems to proceed without using finiteness.

with $t = 1$. If $s = 1 = t$, the generalized polygon is just the vertices and edges of a regular n-gon (or any n-circuit that is a circular path for that matter).

It is easy to show that in a generalized n-gon geometry $(\mathcal{P}, \mathcal{L})$ of order (s, t) $(n > 2)$, the number of points and lines are exactly determined. Either

1. n is odd, one must have $s = t$, and

$$|\mathcal{P}| = |\mathcal{L}| = 1 + s + s^2 + \cdots + s^{n-1},$$

or
2. n is even and

$$|\mathcal{P}| = (1 + s)(1 + (st) + \cdots + (st)^{n-1}),$$
$$|\mathcal{L}| = (1 + t)(1 + (st) + \cdots + (st)^{n-1}).$$

A key result is the Feit-Higman theorem.

Theorem 9.4.16 (Feit and Higman [62].) *Suppose* $\Gamma = (\mathcal{P}, \mathcal{L})$ *is a generalized* n*-gon geometry of finite order* (s, t)*. Then exactly one of the following occurs:*

1. $s = t = 1$ *(the ordinary* n*-gon).*
2. $n = 2$ *(the generalized digon).*
3. $n = 3$ *and* $s = t > 1$ *(the projective plane).*
4. $n = 4$ *(a generalized quadrangle with more than four points).*
5. $n = 6$ *and at most one of* s *and* t *is* 1 *or* $min(s, t) > 1$ *and* st *is a square (generalized hexagons with an order).*
6. $n = 8$ *and at most one of* s *and* t *is* 1 *or* $min(s, t) > 1$ *and* $2st$ *is a square (generalized octagon of order* (s, t)*).*
7. $n = 12$ *and exactly one of* s *and* t *is equal to* 1.

For the cases with $n = 8$ or 12 in which exactly one of the parameters is equal to 1, the generalized n-gon geometry Γ or its dual must arise by the construction of Theorem 9.4.15 given for $2\Gamma^*$.

We have discussed the existence of projective planes of order s at some length in the Appendix to Chap. 5. No planes are known when n is not a prime power, but otherwise they seem to be so plentiful as to discourage any idea of classification.

Finite generalized quadrangles appear in six classical varieties: $Sp(4, q)$, $U(4, q^2)$, $U(5, q^2)$, and their duals $O(5, q)$, $O^-(6, q)$, and the dual of $U(5, q^2)$, which seems not to have a special name. In addition, there are non-classical quadrangles of various sorts:

1. $T_2(\mathcal{O})$, of order (q, q), q even. $\mathcal{O}$ is an oval of $PG(2, q)$ which is not a conic.
2. Generalized quadrangles of order $(s, s + 2)$, with these constructions: (a) $T_2^*(\mathcal{O})$ where $\mathcal{O}$ is a hyperoval of $PG(2, q)$, q even, and (b) the Payne derivative about a regular point of quadrangle of order (s, s).
3. Generalized quadrangles of order (q, q^2): These are (a) the quadrangles arising from Tits' construction $T_3(\mathcal{O})$, where $\mathcal{O}$ is an ovoid of $PG(3, q)$ which is not an elliptic quadric (so q is even) and (b) more than a dozen families of quadrangles

of order $(s, t) = (q, q^2)$ derived from ovoids of $PG(3, q)$ using flocks, Kantor families and q-clans.

These were discussed in the Appendix to Chap. 7. A great many characterizations of the classical quadrangles are surveyed in the excellent article of J. Thas [130] in the *Handbook for Incidence Geometry*. But much more is known and the reader is encouraged peruse the masterful book *Translation Generalized Quadrangles* by J. Thas, K. Thas, and H. Van Maldeghem.

There are only two known infinite families of generalized hexagons with an order (s, t), each parameterized by a finite field. There are the *split Cayley hexagons* of order (q, q) which are related to the group $G_2(q)$ and the *twisted triality hexagons* of order (q, q^3) which are related to the group ${}^3D_4(q)$, and the duals of each type. They are rank two coset geometries defined by the two classes of maximal parabolic subgroups of the indicated group of Lie type.

Finally, the only known octagons of order (s, t), $\min(s, t) > 1$ occur as a coset geometry for the groups ${}^2F_4(q)$ of Ree type. Here q is an odd power of 2. These octagons are called the *Ree–Tits octagons* and have order (q, q^2) (or (q^2, q) in the case of their duals). No other thick octagons are known.

9.5 Diagrams

9.5.1 Introduction

A rank-two diagram is no more or no less than an isomorphism-closed class of rank-two geometries. The word "diagram" comes from our habit of representing nice classes of firm rank two geometries by a labeled edge connecting an ordered pair of vertices, each vertex representing one of the two classes of objects of the geometry. These edges are then incorporated into a larger graph whose vertices (called nodes) are labelled by a set I which we call a "diagram." Geometries over a typeset I are said *to belong to a diagram* D if each of their rank two residues of type $\{i, j\}$ is a member of the class of geometries designated by the label of the edge connecting vertex i and vertex j. Note that a geometry over I cannot belong to a diagram D unless its rank two residues are non-empty and firm.

There is another approach to diagrams. One can also consider diagrams with the labelled edges denoting isomorphism-closed classes of firm rank two chamber systems rather than firm rank two geometries. Then, given a diagram D with vertex set I, we are able to say that a chamber system C over I *belongs to the diagram* D if and only if every rank two residue of type $\{i, j\}$ $(i, j \in I)$ is a chamber system belonging to the class of chamber systems determined by the edge-label of the edge connecting node i to node j.

If $M = (m_{ij})$ is a Coxeter matrix, with rows and columns indexed by I, one can define a diagram $D(M)$ by labeling the edge connecting node i with node j by the class of generalized m_{ij}-gons. Thus in a geometry Γ over I belonging to diagram $D(M)$ (called a *geometry of type* M), the residue $\mathrm{Res}_\Gamma(F)$ of any flag of cotype $\{i, j\}$ is a generalized m_{ij}-gon geometry. Similarly, a *chamber system of*

type M is a chamber system over I belonging to the diagram $D(M)$ – which means that every residue of type $\{i, j\}$ is a generalized m_{ij}-gon (as originally defined as a chamber system). Chamber systems of type M will be the context of our definition of *Building*.

So we have to cover a number of concepts.

9.5.2 Rank Two Diagrams

We begin our glossary of rank two diagrams:

1. (Digons.) These are the rank two geometries whose incidence graph is complete bipartite – that is, *every object of type* 1 *is incident with every object of type* 2.
2. (Generalized trigons = projective planes.) The incidence graph is that of the points and lines of a projective plane.
3. (Generalized quadrangles.) By now the reader should know what these are.
4. (Generalized $2n$-gons.) A class of near polygons $(\mathcal{P}, \mathcal{L})$ with the property that for every line L and point p, there is a unique path from p to its gate (unique nearest point in L).

These cases are indicated by the respective two-vertex diagrams in Fig. 9.7a.

Why are these diagrams left–right symmetric? It is because ultimately a generalized polygon is defined by a hypothesis that reads the same after the words "points" and "lines" are transposed, but "incidence" is left the same. Thus the dual point-line geometry of a digon, projective plane, generalized quadrangle, generalized hexagon, etc., is also one of the same species.

There are other rank two diagrams which are not symmetric. For example, the class of *linear spaces* $(\mathcal{P}, \mathcal{L})$ are those rank two geometries with the property that any two points are incident with a unique line. We denote the class of such geometries by the simple diagram presented in Fig. 9.7b.

As a special case, one might consider the affine planes encountered in Chap. 4. This diagram is depicted in Fig. 9.7 after the label "(c)".

In a dual linear space, any pair of lines meet at a unique point depending on that pair. Since affine planes are missing this property, the class L is not self dual. So it is good that the edge label "L" is not left–right symmetric.

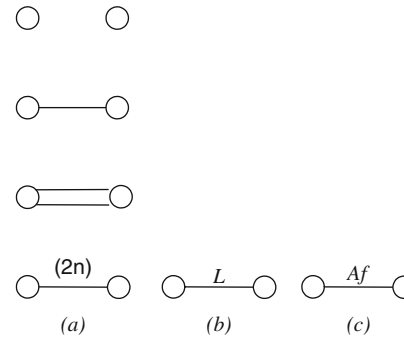

Fig. 9.7 Basic rank two diagrams

9.5.3 Diagram Geometries of Higher Rank

We fix once and for all a type set I, the (global type set). A *diagram (of geometries) over I* is an assignment

$$D : \binom{I}{2} :\to \mathcal{G}_2$$

which assigns to each 2-subset of elements of I an isomorphism-closed class of rank two geometries $D(i, j)$ over $\{i, j\}$. For every subset J of I let $D_{I-J} := D(J)$ denote the restriction of this function to the collection of ordered pairs chosen from J. (We shall always write $D(i, j)$ for $D(\{i, j\})$.) The symbolism above gives us a method of associating with D a graph (also called D) whose nodes are indexed by the elements of I, and for every 2-subset $\{i, j\}$ of I the edge directed from i to j should bear the appropriate rank-two symbol $D(i, j)$ or its graphic representative. The same edge directed from j to i should bear the graphic representative of the class $D(j, i)$. (It is necessary only to assign one of these symbols in the proper orientation since $D(j, i)$ is the class $D(i, j)^*$, the dual point line geometry of $D(i, j)$.) Thus if $D(i, j)$ is the class of linear spaces, any two distinct objects of type i are incident with a unique object of type j and the undirected edge (i, j) is affixed with the symbol "L" written with i to the left and j to the right – or, if it is convenient in drawing the diagram, one can write L^* with j to the left of the symbol.[13]

Next, suppose J is a proper subset of the index set I. Then $D(J)$ is the graph that is obtained when the nodes of $I - J$ and all edges involving at least one of these nodes is removed – that is, it is the labeled graph induced on nodes indexed by J. (Notation is simplified by writing $D(k)$ for $D(\{k\})$. Note that notation has already been arranged so that $D(\{i, j\}) = D(i, j)$.)

In this way we can always describe a graph with $|I|$ nodes associated with the function D : ordered pairs from $I \to$ rank two diagrams.

Conversely, given an edge-labelled graph $D = K_I$, the complete graph over vertex set I, whose edge-labels are prescribed by a function D, directly describes a diagram D.

One says that a geometry Γ over I *belongs to a diagram D* (of geometries) if and only if:

1. For any pair of distinct types $\{i, j\}$, the residue of every flag F of cotype $\{i, j\}$ is a member of the class of rank two geometries $D(i, j)$ over $\{i, j\}$.
2. Every flag F of corank at least three lies in a flag of cotype $\{i, j\}$ for each 2-subset $\{i, j\}$ of the cotype of F.

[13] Of course it is not always possible to represent graphically the diagram so that edges are horizontal, but we still expect the assignment of the asymmetric symbols such as "L" to be oriented with respect to the two vertices so as to reflect the diagram D.

Since the diagram assigns a rank-two geometry for each pair of indices $\{i, j\}$ we see that condition 2 of the definition just given is equivalent to:

2' Every flag of corank at least three lies in a chamber flag.

A *diagram geometry* is simply a geometry which belongs to some diagram.[14]

Are Diagram Geometries Residually Connected?

Of course this question must be answered in the negative if some $D(i, j)$ is a class of disconnected geometries or if Γ itself is not connected.

So let us assume that Γ is a connected geometry belonging to a diagram D with all the $D(i, j)$ being classes of connected geometries. One easily sees that if Γ has rank just three, then Γ is residually connected.

However it is possible to produce a connected geometry of rank four which belongs to a diagram D for which all rank two residues are connected – yet there are rank three residues which are not connected. This is developed in Exercise 9.19 (see p. 397).

So far we have been discussing full geometries belonging to a diagram. In the later chapters we shall discuss how point–line geometries are derived from diagram geometries with special emphasis on classical geometries belonging to diagrams (some of infinite rank).

9.5.4 Chamber Systems Belonging to a Diagram

We may also attach diagrams to chamber systems of higher rank. Here we define a *diagram of chamber systems* as a function

$$D : \binom{I}{2} :\to \mathcal{CH}_2$$

which assigns to each 2-subset $\{i, j\}$ of I a collection $D(i, j)$ of isomorphism classes of chamber systems over $\{i, j\}$.

We say that a chamber system C over I *belongs to the diagram of chamber systems* D if and only if every rank two residue of type $\{i, j\}$, is a rank two chamber system whose isomorphism type is in the class $D(i, j)$.

We allow ourselves to depict the diagram D in the same graphic way as we did for geometries with the understanding that if nodes "i" and "j" are connected by an

[14] Normally, many authors intend only condition 1 for the definition of a geometry Γ belonging to a diagram D. But since a diagram merely expresses the nature of the rank two residues of a geometry or chamber system, this datum may miss anomalies about what is going on at flags of higher corank. Could some of these flags be maximal – that is, have empty residues?

Flags that do not even lie in a flag of corank two are beyond any proscription that a diagram could impose. So it makes sense to propose a more intimate relation between a geometry Γ and a diagram D. That is the reason for condition 2 or 2′.

edge labelled "(n)," then any residue of type $\{i, j\}$ in a chamber system C belonging to diagram D is a generalized n-gon (recall this was defined as a chamber system in the previous section).

Of course edges labelled "L" would now have to denote chamber systems of linear spaces – that is the system of point-line flags with adjacencies determined by their sharing of a common point or line. This may seem a little akward, but we shall not really encounter it. Virtually all chamber systems in this book are "type M" – which means each rank two residue of type $\{i, j\}$, is an m_{ij}-gon, where m_{ij} is an integer greater than 2 or is the symbol ∞. These are studied more deeply in the next major section of this chapter.

9.5.5 Diagrams and the Functors Connecting Chamber Systems and Geometries

Suppose Γ is a residually connected geometry belonging to the diagram D. Then, of course, the associated chamber system $\mathbf{C}(\Gamma)$ is residually connected. But does it belong to the same diagram D? Of course we are now viewing D as a diagram of chamber systems – strictly speaking a $\mathbf{C}(D)$ with the rank two edges being the chamber systems of the geometries $D(i, j)$ – precisely, for each ordered pair (i, j) chosen from I, we have

$$(\mathbf{C}(D))(i, j) = \mathbf{C}(D(i, j)),$$

as classes of chamber systems.

Now we have the following.

Theorem 9.5.1 *Suppose Γ and C are, respectively, a geometry and a chamber system over the same type set I. Assume both are residually connected and either (1) $\Gamma = \mathbf{\Gamma}(C)$ or (2) that I is finite and $C = \mathbf{C}(\Gamma)$.*

1. *If the geometry Γ belongs to the diagram D then so does C.*
2. *If the chamber system C belongs to the diagram D, then so does Γ.*

Proof This is a direct consequence of Corollary 9.2.8. □

9.5.6 Some Examples Concerning Diagram Geometries and Chamber Systems

Example 9 Let us take a simple example. The symbol $D = A_3$ is the name of the following diagram (Fig. 9.8).

Here the type set I is the set $\{1, 2, 3\}$. The diagram tells us that $D_k := D(\{I - k\})$ is a projective plane when $k = 1$ or 3, and that it is a digon D_2 when $k = 2$.

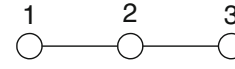

Fig. 9.8 The A_3 diagram

If a geometry Γ belongs to a diagram D, then, according to the definition, for any flag F of cotype $\{i, j\}$ $(i \neq j)$, the residue $\mathrm{Res}_\Gamma(F)$ is a rank two geometry belonging to the class $D(i, j)$. Now consider the geometry $PG(3, K)$ of one-, two-, and three-dimensional vector subspaces of a four-dimensional right vector space V over the division ring K. These are respectively the objects of types 1, 2, and 3 of a geometry in which incidence is vector subspace containment in some order. We have defined and characterized these in Chaps. 3 and 6. The point here is that $PG(3, K)$ *belongs to the diagram* A_3.

Taking the i-dimensional vector subspaces of V to be the objects of type i, saying that $PG(3, K)$ belongs to the diagram A_3, simply means verifying the following three statements:

1. The full collection of objects of types 1 and 2 which are incident with a fixed object of type 3 form a projective plane.
2. Let x, y, and z be objects of types 1, 2, and 3 respectively. If x and z are both incident with y, then they are incident with each other.
3. The collection of all objects of types 2 and 3 which are incident with a given object x of type 1 themselves possess the incidence structure of a projective plane.

This is almost trivial. The last statement holds since the proper subspaces of a three-dimensional vector space over K is a classical projective plane. The first statement it true, since the two- and three-dimensional vector subspaces of V containing a given one-dimensional subspace U are bijective with the one- and two-dimensional subspaces of the three-dimensional factor space V/U (the bijection preserves the incidence relation). The second statement is true simply because incidence is containment.

One might ask whether any geometry belonging to the diagram A_3 is in fact a $PG(3, K)$ for some division ring K? The answer is "yes." One need only show that the truncation to $J = \{1, 2\}$ is a projective space $(\mathcal{P}, \mathcal{L})$ and then use the the Veblen Young theorem. We leave this as an exercise.

9.6 Chamber Systems with a Coxeter Diagram

9.6.1 Coxeter Groups and Coxeter Systems

Coxeter Matrices

Fix an index set I. A *Coxeter matrix over* I is a matrix M whose rows and columns are indexed by I, and whose (i, j)-th entry, m_{ij}, satisfies these properties:

1. m_{ij} is always a positive integer or a formal symbol "∞".
2. $m_{ii} = 1$ for all $i \in I$.
3. $m_{ij} = m_{ji}$ for all $i, j \in I$.

For the moment, the matrix M has no particular linear algebra interpretation. It should simply be regarded as a bank of uncommitted organized data.

The Coxeter Groups $W(M)$

One use of the datum contained in a Coxeter matrix M is that of defining a Coxeter group $W(M)$. Given Coxeter matrix M, the *Coxeter group* $W(M)$ is the factor group

$$F/\langle R^F \rangle,$$

where

1. F is the free group generated by a set $\{x_i\}_{i\in I}$ of generators indexed by I.
2. $\langle R_w^F \rangle$ is the normal closure of the subset of words

$$R_w := \{w_{ij} = (x_i x_j)^{m_{ij}} | i, j \in I, m_{ij} < \infty\}.^{15}$$

Thus, when $m_{ij} = \infty$, $x_i x_j$ is just an element of infinite order in $W(M)$. In effect, $W(M)$ is generated by a collection of elements with trivial square $\bar{R} = \{r_i = x_i \langle R_w^F \rangle\}$, with the property that $\langle r_i, r_j \rangle$ is a dihedral group of order $2m_{ij}$ when m_{ij} is finite, and is the infinite dihedral group when $m_{ij} = \infty$. Moreover, it is the universal group with this property: that is, any other group generated by a class of elements $\{r_i\}_I$ indexed by I with the relations that the product $r_i r_j$ has order m_{ij}, for all $i, j \in I$, is a homomorphic image of the Coxeter group $W(M)$.

Let

$$\rho : F \to F/\langle R_w^F \rangle = W(M)$$

be the natural homomorphism which sends x_i to the element $x_i \langle R_w^G \rangle = x_i^{-1} \langle R_w^G \rangle = r_i$ in $W(M)$. Since each r_i is either an involution or the identity element, and since each element of F is a product of the x_i or their inverses, we can always write an element r of $W(M)$ as a finite product of elements of $\{r_i\}$.

Suppose M is a Coxeter matrix over I and, as above, let $W := W(M)$ be the Coxeter group, and let $R := \{r_i | i \in I\}$, the canonical set of generating elements introduced in the previous two paragraphs.[16] The triple (W, R, M) is called a *Coxeter system*.[17]

---

[15] Recall that the *normal closure* $\langle X^G \rangle$ of a subset X of a group G is the intersection of all normal subgroups of G which contain X.

[16] Not to be confused with R_w, a collection of words in F.

[17] Usually, in the literature the Coxeter matrix M is suppressed, and one simply writes (W, R) for a Coxeter system. I do not understand the reason for this, but the reader should be forewarned of any differences in notation with the standard literature.

Let Z_2 denote the multiplicative group of integers $\{\pm 1\}$. By the fundamental property of a free group, there is a surjection ϕ from the free group on a set of generators X indexed by I onto Z_2 which takes each element of X to the integer -1. If $x, y \in X$, $\phi(xy) = 1$, whether or not x is distinct from y. In particular, if $M = (m_{ij})$ is a Coxeter matrix over I, then

$$1 = 1^{m_{ij}} = \phi(x_i x_j)^{m_{ij}} = (\phi(x_i) \cdot \phi(x_j))^{m_{ij}}.$$

This means the homomorphism ϕ factors through a surjective homomorphism

$$\text{sgn} : W(M) \to Z_2.$$

Clearly if an element r of $W(M)$ is expressible as a product of an odd number of the generators in R, then $\text{sgn}(r) = -1$ and so r cannot be the identity element of $W(M)$. Thus the following applies.

Theorem 9.6.1 *In the Coxeter system (W, R, M), the following assertions are valid:*

1. *There exists a surjective morphism $W(M) \to Z_2$ taking each element of R to -1.*
2. *The elements of the generating set R are involutions.*
3. *Moreover no product of an odd number of elements of R in $W(R)$ can be the identity.*

The Free Monoid Covering the Coxeter Group

We shall need a mechanism to keep track of the ways to express an element of the Coxeter group $W(M)$ as a product of the r_i.

The *free monoid on the alphabet I* is a monoid I^* whose elements are the "words" $w = i_i \cdots i_\ell$ that can be "spelled" with the alphabet I. (Of course each "word" is nothing more than a sequence of elements written as a "string" (the sequence with the intervening commas removed).) The non-negative integer ℓ is call the *length of the word*, and it is intended that ϕ, the empty word (the unique word of length zero), is to be included in I^*. The binary operation on I^* is the concatenation of words. The concatenation of word w_1 with word w_2 is the word $w_1 \circ w_2$ – the word obtained by first writing the word w_1 (from left to right) and then writing w_2 (from left to right) juxtaposed to the right of the first word w_1. Thus if $w_1 = 123$ and $w_2 = 313$ then $w_1 \circ w_2 = 123313$ and $w_2 \circ w_1 = 313123$.[18] One notes that concatenation is an associative binary operation on the set of words, and that the empty word is a two-sided identity with respect to this operation, so the monoid structure is manifest.

[18] The author apologizes for adopting the western-European bias in reading from left to right in defining words and in defining concatenation. Of course there is an opposite monoid more suited to semitic writing (Phonecian, Hebrew, Arabic) and Arabic-script renderings of some non-semitic languages (Parsi and Urdu). How could cultures like the Hittites who wrote *Bostrophedron* ever invent a free monoid? There is an answer.

Now there is a morphism $\rho : I^* \to W(M)$ which substitutes r_i for the letter i in each word of the free monoid to form a finite product of involutions in the group $W(M)$. For example, if $I \supseteq \{1, 2, 3\}$, then ρ takes the word $w = 123123321$ to the product of involutions $r_1r_2r_3r_1r_2r_3r_3r_2r_1$ in $G(R)$. Note that since the r_i are involutions $\rho(w) = \rho(123)$, ρ is far from injective.

Clearly this mapping is a morphism in the categroy of monoids since

$$\rho(w_i \circ w_2) = \rho(w_1)\rho(w_2),$$

where the juxtaposition of elements on the right hand side indicates multiplication in the group $W(M)$. The morphism ρ is surjective by the earlier observation that every element of $W(M)$ is expressible as a product of finitely many r_i.

Now, given an element r of the Coxeter group $W(M)$, the fiber $\rho^{-1}(r)$ lists for us all the possible ways of writing the element r as a product of the r_i. Exactly the library of possibilities we want to keep track of. This way we get to live in two worlds: (1) the world of elements of the group $W(M)$ and (2) the world in the sky of ways to express these elements as a product.

9.6.2 The Cayley Graph of the Coxeter System (W, R, M), and the Coxeter Chamber Systems

We are going to begin this discussion with a simple assumption, which is easily proved in Lemma 9.6.4 at the beginning of the next section.

(*) The elements of R are pairwise distinct, so there is a bijection

$$R = \{r_i\}_{i \in I} \to I.$$

As a consequence, for every element $s \in W(M)$ and generators r_i and r_j in R, the equation

$$sr_i = sr_j$$

implies $i = j$.

In the previous section we produced a surjective monoid morphism $\rho : I^* \to W(M)$, remarking that the set of all preimages of the group element s (that is, the fiber $\rho^{-1}(s)$) is a library of all possible ways to write the element s as a finite product of the generating involutions $\{r_i\}$.

Actually the beginning student has probably met a structure like this in the guise of *Cayley graphs*.[19] The *Caley graph for the Coxeter group $W(M)$ with respect to a*

[19] Indeed the whole idea of representing elements of a group by vertices and words in a set of generators by walks became the basis of both combinatorial topology and combinatorial group theory.

set of generators $R = \{r_i\}_{i \in I}$ is a graph $C(M)$ whose vertex set is the set $W(M)$ of elements of the Coxeter group $W(M)$. We say that (s, t) is an edge *e carrying the label* $j \in I$ if and only if $t = sr_j$ (or equivalently, $s = tr_j$). The result is a graph $C(M)$ (the *Cayley graph of the Coxeter group* $W(M)$), which has edges $e = (s, t)$ which are undirected and which carry with them a set of labels $\lambda(e)$ defined to be the set of indices j in I such that $sr_j = t$.

Since the elements of R are involutions, the Cayley graph $C(M)$ of the Coxeter system (W, R, M) is undirected and has no loops. By the assumption (*) presented at the beginning of this section, the Cayley graph $C(M)$ has no multiple edges. Thus every visible edge defined above carries a unique label, and for every vertex s and label i there is exactly one edge leaving s bearing the label "i."

If 1 denotes the identity element of the Coxeter group $W(M)$, then any walk $p = (1, s_1, s_2, \ldots s_n = t)$ from 1 to t in the Cayley graph $C(M)$ reproduces a unique word $typ(p) = i_1 i_2 \cdots i_n$ $(i = 1, \ldots n)$ of the free monoid I^*, where i_j is an element of $\lambda(s_{j-1}, s_j)$. In that case the element s_n of the Coxeter group is expressible as the product $\prod_{j=1}^{j=n} r_{i_j}$.

We now have a very nice paradigm:

1. *Words in* I^* *correspond one-to-one to walks from* 1 *to any other vertex in the Cayley graph.*
2. *Any relation in* $W(M)$ *among words in its generators corresponds to a circular walk in the Cayley graph* $\mathcal{C}(M)$.
3. *Since* $W(M) = F/\langle R_w^F \rangle$ *is a quotient of a free group by the relation group* $\langle R_w^F \rangle$, *any visible relation (that is, some product over a sequence of the* r_i *is equal to* 1*) must be a logical consequence of basic relations* $(r_i r_j)^{m_{ij}} = 1$ *– that is, every circuit of the Cayley graph is* $\mathcal{C}_2$*-contractible where* $\mathcal{C}_2$ *is the collection of* $2m_{ij}$*-gons defined by taking the orbits of* $\langle r_i, r_j \rangle$ *when* m_{ij} *is finite.*

So it is basically all graphical.

Now it is time to thrust this elementary discussion into a different context. We still retain the assumption (*) maintained at the beginning of this section. Then, as defined, the Cayley graph $C(M)$ of a Coxeter system is a undirected simple graph each of whose edges are labeled by a single element of I.

Lemma 9.6.2 *Under the non-degeneracy condition (*), the Cayley graph* $C(M)$, *with its natural edge-labeling,* λ, *is a connected chamber system over* I. *This chamber system* $C(M)$ *is fully thin (that is, all panels have size exactly two) and satisfies condition (typ) (that the edge labeling assumes a single value on each undirected edge of this simple graph).*

Proof This doesn't really require a bothersome proof. The graph is simple, undirected, and (typ) holds for the edge-labeling. Moreover the "full thinness" condition introduced just above holds. The basic condition defining a chamber system holds automatically – that is, for each $i \in I$, the edges of type i form a "matching" (sometimes called a "1-factor") of the Cayley graph $C(M)$.

We denote the chamber system exhibited by the Cayley graph $C(M)$ by the same symbol $C(M)$.[20]

One observes the following.

Corollary 9.6.3 *1. The Cayley graph $C(M)$ is bipartite. It follows that all its panels (that is its edges) are strongly gated.*

2. The Coxeter group $W(M)$ acts as a group of automorphisms which regularly permutes the chambers.

Proof This is a completely elementary result.

1. According to our paradigm, a circuit in the chamber system $C(M)$ can have type w (a word in the monoid I^*) if and only if $\rho(w)$ is the identity element of $W(M)$. As a consequence, Theorem 9.6.1 implies that $C(M)$ has no circuits of odd length, and so is bipartite.

2. For each element $r \in W(M)$, let r act as a permutation of the chambers by taking chamber c to chamber rc – that is, it acts by *left* multiplication of the chambers. A typical i-adjacency in the chamber system $C(M)$ is an edge $e = (s, sr_i)$ for some $s \in W(M)$. Then r takes e to $re := (rs, rsr_i)$, another i-adjacent pair. This is clearly a group action, and for a chamber c, one has $rc = c$ if and only if $r = 1$, so the action is regular on chambers.

Non-degeneracy of the Coxeter Chamber Systems

The entire subject of groups G *defined by a set X of generators and relations R* (that is, $R \subseteq F(X)$, the free group on generators X, and $G = F(X)/\langle X^{F(X)} \rangle$ is (in the words of J. Humphreys) "notorious" for unexpected results. It could well be that a group G defined in this way by generators and relations is actually trivial. Then of course, the Cayley graph for (G, R) has only one vertex, and all edges would be loops. Even if $G \neq 1$, it is possible that a general Cayley graph has multiple edges.

Because of the *sgn*-epimorphism we know that the first of these pathologies does not occur for the Cayley graph $C(M)$ of a Coxeter system (W, R, M). We know that $C(M)$ has no loops and is bipartite. But what about the second? Can there be multiple edges – or if one prefers, can edges of the chamber system $C(M)$ bear multiple labels? Condition (*) of the previous section asserted that this does not happen. Here we demonstrate (*) by exhibiting a homomorphic image of $W(M)$ in which the elements of R are represented by pairwise distinct elements.

A linear transformation T of a vector space V is *finitary* if and only if the subspace $\mathbf{C}_V(T) : \{v \in V | v^T = v\}$, has finite codimension in V. Working with finitary transformations has the advantage that such transformations are invertible if

[20] This is not an abuse of notation. Rather it is a *use* of definitions. By some extraordinary serendipity which sometimes blesses fumbling authors like myself, the phrases "Cayley" and "Chamber system" begin with the same letter "C." So the mnemonics are preserved when we denote (1) the *Cayley graph* defined by Coxeter matrix M and (2) the *Chamber system* defined by the same Coxeter matrix M, by the very same symbol $C(M)$.

and only if they are injective. We let $GL(V)$ denote the full group of all invertible finitary transformations of V.

Suppose a group G acts as a group of finitary linear transformations on a vector space V. (One normally calls this a *finitary representation of the group G*.) In effect, a finitary representation is simply a group morphism

$$\phi : G \to GL(V).$$

It is said to be a *faithful representation* if and only if $\ker \phi = \{1_G\}$, the identity subgroup of G. Of course that means there is an embedding $G \to GL(V)$.

Let $V = \bigoplus_{i \in I} \mathbf{R} v_i$ be the vector space over the field of real numbers $\mathbf{R}$ whose basis is the set $X := \{v_i\}_{i \in I}$. Let $M = (m_{ij})$ be a Coxeter matrix over I.

Then there exists a symmetric inner product, $B_V : V \times V \to \mathbf{R}$, uniquely defined by the equations

$$B_V(v_i, v_i) = 1, \text{ for all } i \in I, \tag{9.27}$$

$$B_V(v_i, v_j) = B_V(v_j, v_i) = -\cos\left(\frac{\pi}{m_{ij}}\right), \ i, j \in I, i \neq j, \tag{9.28}$$

$$B_V(v_i, v_j) = -1, \text{ when } m_{ij} = \infty. \tag{9.29}$$

This form B_V is called the *Coxeter form* and is completely determined by the Coxeter matrix M.

As an example, for the diagram of type A_3 considered above,

$$M = \begin{pmatrix} 1 & 3 & 2 \\ 3 & 1 & 3 \\ 2 & 3 & 1 \end{pmatrix},$$

while the Gramm matrix is

$$\begin{pmatrix} 1 & -1/2 & 0 \\ -1/2 & 1 & -1/2 \\ 0 & -1/2 & 1 \end{pmatrix}.$$

For each $i \in I$, $\langle v_i \rangle$ is a non-degenerate subspace of V, and so we have a decomposition

$$V = \langle v_i \rangle \oplus v_i^{\perp},$$

where $v_i^{\perp} := \{v \in V | B_V(v, v_i) = 0\}$.

Let σ_i be the reflection on the space V which "inverts" each vector $y \in \langle v_i \rangle$ and which pointwise stabilizes the complementing "perp-space" $v_i^{\perp}$. Since we are over the field of real numbers $\mathbf{R}$ of characteristic not 2, this construction makes sense. Clearly each such reflection σ_i is an involution.

But the form B_V is defined so that the product of the two reflections $\sigma_i \sigma_j$ has order m_{ij} when the latter is a positive integer, or has infinite order when $m_{ij} = \infty$.

Applying the fundamental property of a free group, we see that there is a surjective homomorphism

$$f : W(M) \to S := \langle \sigma_i, i \in I \rangle$$

taking the generator r_i of $W(M)$ to the reflection σ_i of (V, B).

Note that the Coxeter form (V, B_V), the collection $\Sigma := \{\sigma_i | i \in I\}$ of reflections, and the finitary representation $f : W(M) \to S = \langle \Sigma \rangle$ are all uniquely determined (up to isomorphisms) by the Coxeter matrix M alone.

Now we have condition (*) alluded to above.

Lemma 9.6.4 *Let (W, R, M) be a Coxeter system. Then the generators r_i of R are a collection of pairwise distinct involutions.*

Actually, there is quite a bit more.

Theorem 9.6.5 (Bourbaki, Chap V, n^o 4.3) [9].) *The homomorphism $f : W(M) \to S$ defined by the Coxeter form is an isomorphism .*

A very nice proof of this theorem appears in the excellent book of Garret [63, pp. 7–9].

It has a very important consequence.

Corollary 9.6.6 (The Parabolic Subgroup theorem.) *Suppose $(W(M), R, M)$ is a Coxeter system. We suppose I indexes the set of generators R, so that $C(M)$ is a chamber system over I. Let J be any subset of I. Let $R_J = \{r_i \in R | i \in J\}$ and let M_J be the minor of the Coxeter matrix M obtained by restriction to the rows and columns indexed by the elements of J. Let $(W(M_J), S, M_J)$ be the Coxeter system defined by the submatrix M_J, where $S := \{s_j | j \in J\}$ denotes the fundamental set of involutatory generators of the Coxeter group $W(M_J)$.*

Then there is an isomorphism

$$W(M_J) \to \langle R_J \rangle$$

which sends s_j to r_j, for all $j \in J$.

In particular, if $r_{i_1} \cdots r_{i_k} = 1$ is an identical relation in the Coxeter group $W(M)$, and if all indices i_j in this expression belong to J, then this relation is entirely the consequence of the relations

$$(r_I r_j)^{m_{ij}} = 1, \text{ for } i, j \text{ restricted to } J \text{ alone.}$$

Put another way, $(\langle R_J \rangle, R_J, M_J)$ is itself a Coxeter system.

Proof As before, let V be the real vector space with basis $X = \{v_i\}_{i \in I}$, let B be the symmetric bilinear Coxeter form on V defined by Eqs. (9.27), (9.28) and (9.29), and let $\Sigma := \{\sigma_i | i \in I\}$ be the system of reflections which invert the v_i.

For a subset J of I, let

$$V_J := \langle v_j | j \in J \rangle,$$
$$\Sigma_J := \{\sigma_j | j \in J\},$$

let B_J be the restriction of the form B to V_J, and let the glossary for M_J, R_J, and the Coxeter system $(W(M_J), S, M_J)$ be as in the statement of the corollary.

Clearly the reflections in Σ_J induce reflections on the space V_J, each pointwise fixing a subspace of V which complements V_J. For all $j \in J$, let $\bar{\sigma_j} = \sigma_j|_{V_J}$ be the restriction of reflection σ_j to the subspace V_J, and let $\bar{\Sigma}_J$ be the collection of these restricted reflections $\bar{\sigma}_j$, $j \in J$. Then there is an isomorphism

$$\delta : \langle \Sigma_J \rangle \to \langle \bar{\Sigma}_J \rangle,$$

the right side being a subgroup of $GL(V_J)$.

The finitary representation $W(M) \to GL(V)$ restricts to its subgroup $\langle R_J \rangle$ to give a surjection of groups

$$\beta : \langle R_J \rangle \to \langle \Sigma_J \rangle,$$

taking r_j to σ_j, for $j \in J$.

Also, since the R_J and Σ_J satisfy the relations prescribed for R (simply a consequence of the fact that M_J is a submatrix of M), the fundamental property of the free group produces group surjections

$$\alpha : W(M_J) \to \langle R_J \rangle, \text{ and}$$
$$\gamma : W(M_J) \to \langle \Sigma_J \rangle.$$

Now $\beta \circ \alpha = \gamma$ since both sides take the generator s_j to the reflection σ_j, $j \in J$.

Now $W(M_J)$ acts on (V_J, B_J) via

$$\delta \circ \gamma : W(M_J) \to \langle \bar{\Sigma}_J \rangle$$

exactly as in the discussion of the preceding Theorem 9.6.5. An application of that theorem (with M_J and (V_J, B_J) replacing M and (V, B) throughout) implies that $\delta \circ \gamma$ is a group isomorphism. It follows that $\gamma = \beta \circ \alpha$ is a group isomorphism, and so each of the surjective factors α and β are isomorphisms. Now α is an isomorphism sending s_j to r_j, so no relations exist among the elements of R_J other than those dictated by $W(M_J)$. The proof is complete. □

Length Functions in I^* and $W(M)$

We have already discussed the length of a word w in the free monoid I^* over the alphabet I. The length of w is just the number of letters used to "spell" the word w – the empty word being length zero. Given a Coxeter system (W, R, M) defined by a Coxeter matrix M, we have produced a surjective morphism of monoids

$$\rho : I^* \to W(M)$$

in which letter $i \in I$ is mapped to generator $r_i \in R$. We have also seen that ρ induces a bijection $\hat{\rho}$ between I^* and the set of all walks (galleries) in the chamber system $C(M)$ beginning at 1, the identity element of $W(M)$. The inverse of the bijection $\hat{\rho}$ is simply the type function that records the type, typ(G), of a gallery G of $C(M)$ beginning at 1. Note that the terminal chamber of the gallery G is then $\rho(\text{typ}(G))$ and that the length of the gallery $\ell(G)$ is in fact the length $\ell(\text{typ}(G))$ of the word typ(G).

Now there is another notion of length, this time on elements of the Coxeter group $W(M)$. Actually it is a distance function: we define the *length* $\ell(r)$ *of an element* r *of the Coxeter group* $W(M)$ to be the *distance* $d(1, r)$ in the Cayley graph $C(M)$ of the element r from the element 1. That means $\ell(r)$ is the *shortest* length of a word $w \in I^*$ such that $\rho(w) = r$.

We say that a word $w \in I^*$ is a *reduced word* if and only if $\ell(w) = \ell(\rho(w))$. Note that when we use the term "reduced word," there is a Coxeter matrix in the background. It depends on nothing more than that.

Recall from Chap. 1 that a geodesic path in a graph is just a walk of shortest possible length connecting its extremeties. In our case, if w is a reduced word in I^* (with respect to M) there actually is a geodesic from 1 to $\rho(w)$ whose type is w.

One now notes the following.

Lemma 9.6.7 *If r is any element of $W(M)$, and r_i is one of the generators in R, then $\ell(rr_i)$ is either $\ell(r) - 1$ or $\ell(r) + 1$.*

Proof Since $C(M)$ is bipartite, each edge (r, rr_i) of $C(M)$ is gated with respect to the vertex 1. So the two distances $d(1, r)$ and $d(1, rr_i)$ differ by one. The proof is complete. □

Homotopy in the Coxeter Chamber System

Remark We have noted that $\mathcal{C}_2$-homotopy in the Cayley graph $C(M)$ of the Coxeter system (W, R, M) provides a way of asserting (rather than "determining") that, for two words w_1 and w_2 in I^*, one has $\rho(w_1) = \rho(w_2)$. In other words, it provides a "playing board" on which one can graphically restate the fact that two products of elements of R represent the same element of the Coxeter group.[21]

At least $\mathcal{C}_2$-homotopy is the relevent concept.

> *Two walks (or galleries) in $C(M)$ beginning at 1 have the same terminus if and only if they are $\mathcal{C}_2$-homotopic.*

[21] By scrambling the language of the logicians and the topologists, one might say that $\mathcal{C}_2$-homotopy encompasses the "word problem" for the generators and relations. One might say this at a cocktail party, but it would be a bit crude to assert it seriously. Two walks in $C(M)$ might be $\mathcal{C}_2$-homotopic, (that is, they are connected by a finite string of elementary homotopies, each giving a "deduction" of the ρ-equivalence of two words), but it could conceivably be true that there is no *general recipe* to discover this string of elementary homotopies, as would be required by the logician studying word problems in groups.

Referring to our discussion of homotopy in Chap. 1, we may take $\mathcal{C}_2$ to be the collection of all circular walks in the Coxeter chamber system $C(M)$ whose type has the form $(ij)^{m_{ij}}$ for some pair of indices i, j for which $m_{ij} < \infty$.

An *elementary contraction* of a word in I^* is the deletion of a subword of the form ii – thus the concatenation $f = w_1 \circ ii \circ w_2$ is changed to $f' = w_1 w_2$ by an elementary contraction. Similarly, an *elementary expansion* is the insertion of any expression ii (any i in I) at any position in a word w – that is, an elementary contraction performed with the movie film running backwards. These operations *do* change the length of a word by ± 2 – so they do not disturb the parity of the length (its value mod 2).

As with homotopies, we can use the bijection between words of I^* and galleries beginning at chamber c to apply this notion to galleries of C having c as an initial chamber. Thus an elementary contraction of a gallery simply "snips off" a segment a-b-a from the gallery. A sequence of such deletions can be used to trim off all "spurs," as in Fig. 9.9. Given the Coxeter matrix M, an *elementary M-homotopy of words* replaces a subword of type

$$p(i, j; m_{ij}) := ijijij\cdots \text{ of length } m_{ij}$$

by

$$p(j, i; m_{ij}) := jiji\cdots \text{ of length } m_{ij},$$

when m_{ij} is finite.

More generally, an *elementary $\mathcal{C}_2$-homotopy of words* replaces a subword

$$p(i, j; k) := iji\cdots \text{ of length } k$$

by

$$p(j, i; m_{ij} - k) := jiji\cdots \text{ of length } m_{ij} - k.$$

An important observation is as follows.

An elementary M-homotopy does not change the length of the word. Although an elementary $\mathcal{C}_2$-homotopy of words may change the length of a word, it does not change the parity of that length.

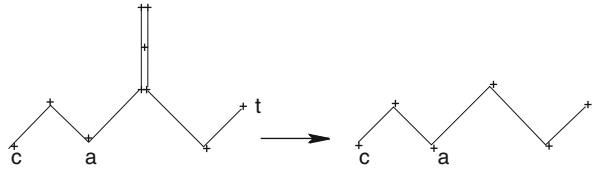

Fig. 9.9 Trimming of spurs

We say two words are *equivalent* if one can be obtained from the other by a finite sequence of expansions, contractions, and elementary $\mathcal{C}_2$-homotopies. The student may easily verify that "equivalent" as just defined is indeed an equivalence relation (the empty sequence of elementary operations is allowed). We write $[w]$ for the equivalence class of all words equivalent to word w.

Clearly if w_i is equivalent to $w_i', i = 1, 2$, then the concatenation $w_1 \circ w_2$ is equivalent to $w_1' \circ w_2'$ (just carry out the elementary operations of contraction, expansion, and $\mathcal{C}_2$-homotopy separately in each factor w_i of $w_1 \circ w_2$). Thus there is a well-defined class containing the concatenations of all words in $[w_1]$ with all words in $[w_2]$ and we may write

$$[w_1] \circ [w_2] = [w_1 \circ w_2]$$

as a well-defined operation on the class of equivalent words. Thus the collection of all word-equivalence classes forms a semigroup $I^*/[\ ]$ and there is an epimorphism $e : I^* \to I^*/[\ ]$.

Lemma 9.6.8 *Two words w_1 and w_2 are equivalent if and only if they have the same image in $W(M)$ under the homomorphism $\rho : I^* \to W(M)$. Thus $I^*/[\ ]$ is isomorphic to the Coxeter group $W(M)$.*

Proof Note that ρ is the restriction to the monoid I^* of the canonical surjection $\hat{\rho} : F(S) \to W(M)$ by which $W(M)$, as a presented group, is represented as a homomorphic image of the free group $F(S)$. But ρ factors as $f \circ [\ ]$ where $f([w])$ is the constant value of ρ on the class $[w]$. We have a surjection $f : I^*/[\ ] \to W(M)$.

It is now clear that $I^*/[\ ]$ is a group. But since all of the relators $R_w = \{(s_i s_j)^{m_{ij}}, (s_i)^2\ i, j \in I\}$ lie in $[\emptyset]$, the group $I^*/[\ ]$ has its subset $S' := \{s_i[\emptyset] \in I\}$ satisfying the same set of relations. Thus $f(s_i[\emptyset]) = s_i$ and the universality of the presented group $W(M)$ shows that f is an isomorphism taking cosets $s_i[\emptyset]$ (which are elements of $I^*/[\ ]$) to $s_i, i \in I$. Thus for any w, w' in $I^*(I)$, $f[w] = e(w')$ if and only if $[w] = [w']$.

9.6.3 Other Properties of Coxeter Chamber Systems

The first property of $C(M)$ that concerns us is the following.

(G_c) If G is a gallery of $C(M)$, then its type $\mathrm{typ}(G)$ is a reduced word of I^* if and only if G is a geodesic path.

To show that $C(M)$ possesses this property, we may apply an automorphism of the chamber system to assure that the initial vertex of G is the identity element 1. In this case G is a gallery of type $w = \mathrm{typ}(G)$ beginning at 1 and ending at $t := \rho(w)$. In fact, by our paradigm, G is the *only* gallery of that type starting at 1. But now, by our definition, w is a reduced word if and only if $\ell(w) = \ell(t)$ — that is, if and only

if $\ell(G) = d(1, t)$. The latter equation holds if and only if G is a geodesic path in $C(M)$.

Remark If we seemed to be "proving" that "every Coxeter chamber system possesses property (G_c)" – one will surely notice that the "proof" is really just a tautology resulting from the way we defined "length" on $W(M)$, and the way we defined "reduced". It should then seem odd if I told you that this property (G_c), applied to arbitrary chamber systems C of type M (of course "C" replaces $C(M)$ in the definition of condition (G_c) given above), becomes one of the equivalent conditions for a building. How can possession of a property be a tautology in the context of Coxeter chamber systems, yet be an important concept in the more general context of chamber systems of type M? Being tautological, it can point to no special property of Coxeter chamber systems to bring to the larger context! But it does, for it suggests that Coxter chamber system-like structures are involved in chamber systems of type M. One can almost "smell" apartments in this property (G_c). But perhaps we are getting ahead of ourselves.

M-Homotopy in a Coxeter Chamber System

Our aim in the next sections is to show that in a Coxeter chamber system $C(M)$ one has condition (RG) – the assertion that all residues are strongly gated. This is proved in Theorem 9.6.14 below.

Remark Most accounts of Coxeter groups in the literature focus on quite different defining properties, such as the "Strong Exchange condition," and "the Deletion condition," which, though they characterize Coxeter groups, do not produce efficient proofs of the (RG)-result just quoted. In fact (excepting Tits' "local approach paper" [139], pp 588–621) it is difficult to find in the literature any mention of Theorem 9.6.14 as a direct consequence that one is dealing with a Coxeter group.

Note that $\mathcal{C}_2$-homotopy of galleries of $C(M)$ can drastically change the length of a gallery in $C(M)$. For example, if $m_{ij} = 3$, a gallery G starting at 1 and type ij can be replaced by one of type $jiji$ – a longer gallery.

So this raises a question.

If G_1 and G_2 are two galleries from 1 to t in $C(M)$ of types w_1 and w_2, respectively, then, of course, the gallery $G_1 \circ G_2^{-1}$ is a circuit pointed at 1, and by the universality of $W(M)$, this circuit is $\mathcal{C}_2$-contractible. By Chap. 1, this means G_1 is $\mathcal{C}_2$-homotopic to G_2.

Now suppose the two words w_1 and w_2 are ($\mathcal{C}_2$-homotopic) *reduced* words – i.e., the two galleries G_1 and G_2 are geodesic paths from 1 to t. Is it really necessary to *lengthen* the galleries – going from a shortest gallery, to longer galleries, and finally back to a final shortest gallery – in order to achieve the $\mathcal{C}_2$-homotopy?

Put another way, is it possible to pass directly from geodesic gallery G_1 to G_2 without ever passing to a longer gallery? One must ask what the elementary homotopies would be in this case? At best, one is replacing some subsegment of type $p(ij) := iji\cdots$ (of length m_{ij}) by one of type $p(ji) := jiji\cdots$ (also of length m_{ij})

and then only in the case that m_{ij} is finite. We call this an *elementary M-homotopy* of $C(M)$. We say two galleries G_1 and G_2 of $C(M)$ are *M-homotopic* if and only if one can be transformed into the other by a chain of elementary M-homotopies. The property we desire is as follows.

(P_c) Suppose G_1 and G_2 are two $\mathcal{C}_2$-homotopic galleries of $C(M)$ which have the same initial and terminal vertices. If both are geodesics, then they are in fact M-homotopic.

Now we assert that this property is true of Coxeter chamber systems $C(M)$.[22] In the next section we shall see that this property (P_c) makes sense in the context of all chamber systems C of type M (of course with C replacing $C(M)$ in the description above) and that in fact this property is equivalent to (G_c), though that is hardly obvious at this stage. We prove that in the next section. It will follow that $C(M)$ possesses the property (P_c), but it is better to wait for a general proof, rather than one geared specifically to Coxeter chamber systems.

9.6.4 Walls, Roots, and Distance in a Coxeter Chamber System

Let (W, R, M) be a Coxeter system. Now the group W, of course, acts from the left on the coset chamber system $C = C(W, 1; R)$ and any element $r \in R^W = \{g^{-1}r_i g | g \in W, r_i \in R\}$ is called a *reflection*.

Given a reflection r, the collection E_r of all edges (rank one residues) of C which are stabilized by r is called a *wall*. Note that because of the free action of W on the chambers of C, any edge e in E_r necessarily has its two vertices transposed by the involution r. (This means, of course, that r transposes the two opposite roots $D^-(e)$ and $D^+(e)$ defined by the edge e. See Exercise 9.22, part 3.)

Suppose, next, that an edge e belonging to the wall E_r, lies in a rank two residue S. Then as S is a connected $\{i, j\}$-component of C, r stabilizes S. If m_{ij} is finite, S is just a $2m_{ij}$-gon (as a graph), r must stabilize the unique edge e' opposite e in this polygon. This means the following.

(9.6.1) *If e is in the wall E_r and e' is the edge opposite e in some finite rank two residue S containing e, then e' is also in the wall E_r.*

We say that a gallery $G = (c_0, c_1, c_2, \ldots, c_m)$ *crosses a wall E_r k times* if exactly k of the edges $e_1 = (c_0, c_1)$, $e_2 = (c_1, c_2)$, …, $e_n = (c_{n-1}, c_n)$ belong to the wall E_r. We say G *crosses the wall E_r* if and only if it crosses E_r at least once – more precisely: it crosses E_r k times, where $k \geq 1$. We observe the following elementary but far-reaching result.

[22] The beginning reader is invited to try it out on the Coxeter chamber system of type A_3.

Lemma 9.6.9 *Let $E = E_r$ be any wall.*

1. *Any minimal gallery crosses the wall E at most once.*
2. *Given any two chambers x and y, either (a) all galleries from x to y cross E an even number of times, or (b) all such galleries cross E an odd number of times.*

Proof 1. Let $G = (c_0, c_1, \ldots, c_m)$ be a minimal gallery. Then any interior interval $(c_i, \ldots, c_j)$ must also be a minimal gallery. Now suppose G crossed the wall E k times, where $k \geq 2$. Then there is at least a first time – i.e., a smallest i with $(c_{i-1}, c_i) = e_i$ in the wall E – and also a second instance $e_j = (c_{j-1}, c_j) \in E$, with j the second smallest index with this property. Then $A := (c_{i-1}, \ldots, c_j)$ is a minimal gallery, as observed, so $d(c_{i-1}, c_j) = j - i + 1$. But the reflection r takes A to a gallery A^r from c_i to c_{j-1}. But the distance of c_i to c_{j-1} is only $j - i - 1$, since $(c_i, c_{i+1}, \ldots, c_{j-1})$ is a minimal gallery of this length. Thus r maps a minimal gallery A to a gallery A^r which is not minimal. But this is impossible as r is an automorphism of the chamber system C and so must preserve distances in its graph.

2. Let A and B be two galleries from chamber x to chamber y. Then, regarding x and y as elements of $W(M)$ we see that since A and B have the same terminal chamber y and initial chamber x, $\rho(typ(A)) = \rho(typ(B)) = x^{-1}y$ where $typ(A)$ and $typ(B)$ are the *types* (as words in I^*) of the galleries A and B respectively. Thus gallery A can be deformed into gallery B by a sequence of cutting off or adding on spurs (segments that are backtracks) and elementary $\mathcal{C}_2$-homotopies. It thus suffices to show that neither of these elementary processes disturbs the parity of the number of times that the gallery crosses the wall E.

Consider first the cutting off of a spur $(b_i, b_{i+1}, \ldots, b_{i+t}, b_{i+t-1}, \ldots, b_{i+1}, b_i)$. For each edge $((b_j, b_{j+1})$ in the first part of the spur which belongs to E there is a second occurrence of it (b_{j+1}, b_j) encountered on the return trip to b_i. Thus snipping off a spur costs one an even number of edges of E.

Since adjoining a spur is the above process with the direction of time reversed, an even number of edges of E would be acquired by this process.

Finally we consider an elementary $\mathcal{C}_2$-homotopy $G_1 \circ A \circ G_2 \to G_1 \circ B \circ G_2$ where A has type $p(i, j; k)$, and B has type $p(j, i; m_{ij} - k)$ when m_{ij} is finite. Clearly the transformation $A \to B$ is occurring inside a residue R of type $\{i, j\}$ which is a $2m_{ij}$-circuit, and A and B comprise complementary segments of that circuit. By our observation (5.1), for each edge e of E in the segment A, there is a corresponging opposite edge e' in the circuit R, which may or may not be in A. If it is in A, both edges e and e' are lost by the exchange of B for A. But if e' is not in A, the homotopy transformation just exchanges e' for e. The same remarks hold with the terms "A" and "B" transposed.

Thus in all cases a $\mathcal{C}_2$-homotopy of galleries only alters the number of edges in E by an even number. The proof is complete.

Part 2 of the above lemma can be used to define, from a wall E_r, a partition of the chambers into two parts $D^+(E_r)$ and $D^-(E_r)$ which we call – for the moment – *half-apartments*. We describe this.

Once given a wall $E_r = E$, galleries are of two types: those crossing E an odd number of times and those crossing E an even number of times. Call them *odd* and *even galleries*, respectively. Then from this definition the following arises.

(9.6.2) *The concatenation $\alpha \circ \beta$ of two galleries is even if and only if α and β are both even or both odd.*

Now Lemma 9.6.9, part 2 says that two chambers are either connected only by odd galleries or only by even galleries – yielding in this way two symmetric relations on the set of chambers. But now (5.2) shows that the relation of being connected only by even galleries is not only reflexive, it is transitive as well, and that there are exactly two equivalence classes $D^+(E)$ and $D^-(E)$ partitioning C. Two chambers x and y are in opposite classes if and only if they are connected only by odd galleries.

We have seen (from part 2 of Exercise 1.1 of Chap. 1 and elsewhere) that as C is a connected bipartite graph, any edge e determines two sets $D^+(e)$ and $D^-(e)$, which we have called "roots" (the sets closest to one of the vertices of e) and that these sets also partition the vertices of C. Now the wall E is a collection of edges. If e is an edge of the wall E, what is the relationship of the roots $D^{\pm}(e)$ and the half-apartments $D^{\pm}(E)$? In the next theorem it will be shown that these two partitions of C are in fact the same. From this it will follow that if $e, f \in E$ then $\{D^{\pm}(e)\} = \{D^{\pm}(f)\}$ – and more.

Theorem 9.6.10 *Let E be a wall and let $C = D^+(E) + D^-(E)$ be the partition of chambers into half-apartments determined by E.*

1. *The sets $D^+(E)$ and $D^-(E)$ are convex.*
2. *For each edge $e = (x, y)$ bridging $D^+(E)$ and $D^-(E)$ (that is, $x \in D^+(E)$, $y \in D^-(E)$), then*

   (a) $D^+(E)$ and $D^-(E)$ are the opposite roots

$$D_x(e) := \{z \in C | d_C(z, x) < d_C(z, y)\} \text{ and}$$
$$D_y(e) := \{z \in C | d_C(z, y) < d_C(z, x)\},$$

   respectively, and

   (b) e belongs to E.

Proof 1. Let (x, y) be an "even" pair of chambers with respect to the wall E. This means $x \neq y$, and $\{x, y\} \subseteq D^+(E)$ or $\{x, y\} \subseteq D^-(E)$. Let G be a minimal gallery from x to y. On the one hand G must cross E an even number of times while on the other hand, by Lemma 9.6.9 it crosses E at most once. Thus each chamber c of the gallery G also forms an even pair with x and lives in the same half-apartment $D^+(E)$ or $D^-(E)$ as x and y. So half-apartments are convex.

2. Let e be a bridging edge as in part 2. Then e itself is a minimal gallery (x, y) connecting x and y. Since, by definition, (x, y) is an odd pair, this gallery crosses the wall $E = E_r$ an odd number of times. Thus the edge e is r-invariant, proving (b).

Now suppose z is any chamber in $D^+(E)$. We must show that $d(z, x) < d(z, y)$. Let us suppose, on the contrary, that z is closer to y than to x, so that by Exercise 9.22, part 3, $d(z, x) = d(z, y) + 1$. Then if G is a minimal gallery from z to y, the augmented gallery $G' := G \circ (x, y)$ is minimal. Then G' crosses E an even number of times. But as it contains the edge $e = (x, y)$ it crosses E at least once, and so must cross E at least twice. But by Lemma 9.6.9, that is impossible for a minimal gallery G'. Thus we must have $d(z, x) < d(z, y)$ after all.

We have shown, then, that

$$D^+(E) \subseteq D_x(e) := \{z | d(z, x) < d(z, y)\}.$$

Similarly, $D^-(E) \subseteq D_y(e)$ and since $D_x(e) \cap D_y(e) = \emptyset$, all containments here are equalities. Thus (a) holds, and the proof is complete. □

Corollary 9.6.11 *For every pair of distinct chambers x and y, there exist exactly $d(x, y)$ roots containing x but not y. In particular if x is adjacent to y, there is exactly one such root.*

Proof First assume x is i-adjacent to y in C. This means C can be coordinatized so that $x = 1$ and $y = s_i$. Then left multiplication by s_i transposes x and y. Since this was just one of many possible coordinatizations, it is correct to say that some conjugate r of s_i transposes x and y. Thus $e := (x, y)$ belongs to the wall E_r, and by the theorem $D_x(e)$ and $D_y(e)$ are the two roots of E_r, with $D_x(e)$ containing x but not y. Suppose now that D was a root containing x but not y. Then D and its opposite root $-D := C - D$ are formed by a wall E_t containing $e = (x, y)$. But then t and r both transpose x and y to tr fixes x. Since G acts regularly on C, $tr = 1$ so $r = t$. Thus each edge (x, y) belongs to a unique wall – and there is a unique root containing x but not y.

Now let $\gamma = (x = c_0, c_1, \ldots, c_m = y)$ be a minimal gallery from x to y. Let D_i be the unique root containing c_{i-1} but not c_i. Then as x is closer to c_{i-1} than c_i, x belongs to D_i since D_i is a root. But similarly D_i contains y but not x, so y is not in D_i. Clearly, then, the $D_i, i = 1, \ldots, m$ comprise a collection of distinct roots which contain x but not y. We claim they comprise *all* such roots. For suppose D' were any root containing x but not y. Then the gallery γ must pass at some point from the set D' to its opposite root $(C - D')$, at some bridging edge – say (c_{j-1}, c_j). Then D' contains c_{j-1} but not c_j. From the uniqueness of roots separating adjacent chambers, we see $D' = D_j$. □

Corollary 9.6.11 also makes it easy to identify those chambers one might encounter along a minimal gallery stretched between two given chambers.

Lemma 9.6.12 *Fix two chambers x and y. The following conditions on chamber z are equivalent:*

(1) z lies on a minimal gallery from x to y.
(2) $d(x, z) + d(z, y) = d(x, y)$.
(3) z lies in every root containing both x and y.

Proof (2) implies (1). If g and h are minimal galleries from x to z, and from z to y, respectively, then (2) shows that the concatenation $g \circ h$ of these galleries is minimal. Since it contains z, (1) holds.

(1) implies (3). If D is a root containing x and y, it must, because of its convexity, contain any minimal gallery between them and in particular must contain z.

(3) implies (2). Suppose z lies in every root containing both x and y. Now if D is a root containing z but not y, it must contain x. For if not, x, like y, must belong to the opposite root $-D$, contrary to our hypothesis on z. Thus we have

$$\begin{aligned} d(z, y) &= \text{no. of roots on } z \text{ but not } y \\ &= \text{no. of roots on } x \text{ but not } y \text{ which contain } z. \end{aligned}$$

Also if D is a root on x but not z, then by hypothesis D does not contain y. Thus

$$\begin{aligned} d(x, z) &= \text{no. of roots on } x \text{ but not } z \\ &= \text{no. of roots on } x \text{ but not } y \text{ which do not contain } z. \end{aligned}$$

But now

$$\begin{aligned} d(x, y) &= \text{no. of roots on } x \text{ but not } y \\ &= \text{no. of roots on } x \text{ but not } y \text{ which contain } z \\ &\quad + \text{no. of roots on } x \text{ but not } y \text{ which do not contain } z \\ &= d(z, y) + d(x, y) \end{aligned}$$

by the preceding two equations. Thus (2) holds and the proof is complete. □

9.6.5 Gatedness and Convexity of Residues

There are nice consequences of Corollary 9.6.11 and Lemma 9.6.12, of the previous section.

Lemma 9.6.13 *All residues of the Coxeter chamber system $C(M)$ are isometrically embedded.*

Proof Let (W, R, M) be the Coxeter system giving rise to the chamber system $C(M)$. Let x and y be two distinct chambers of a residue S of $C(M)$ of type J. Let d_S denote the internal metric of S. It follows from Corollary 9.6.6 that S is isomorphic to $C(M_J)$, the Cayley graph of the Coxeter system (W_J, R_J, M_J). Then

by Corollary 9.6.11, $d_R(x, y)$ is the number of roots of (W_J, R_J, M_J) which contain x but not y. Now suppose $E_J(r) := (D_J^+(r), D_J^-(r))$ is a wall of $S = C(M_J)$ defined by involution $r \in R_J$ (here expressed as a pair of opposite roots). Since $R_J \subseteq R$, we have $D_J^+(r) = D^+(r) \cap S$ and $D_J^-(r) = D^-(r) \cap S$ where $(D^+(r), D^-(r))$ is the partition of $C(M)$ into opposite roots defined by the involution r. The correspondence $(D_J^+(r), D_J^-(r)) \to (D^+(r), D^-(r))$ is therefore an injection. A second application of Corollary 9.6.11 shows that the global distance $d(x, y)$ is *at least as large* as $d_R(x, y)$. But since distance measures the lengths of minimal galleries the inquality also goes the other way. Thus we have $d(x, y) = d_R(x, y)$. Thus S is isometrically embedded. □

Theorem 9.6.14 *In the chamber system C of the Coxeter system (W, R, M), all residues are strongly gated and hence convex.*

Proof Suppose (c, g, x, S) is a quartet where S is a residue of $C(M)$, and c, g, and x are chambers such that

1. $x \in S$, $g \in S$ such that $d(c, S) = d(c, g)$, and
2. $d(c, x) < d(c, g) + d_S(g, x)$.

Clearly c is not in S. Lemma 9.6.13 now tells us that we may drop the subscript S in the right-most term of 2. Then 2 informs us that g lies in no geodesic gallery from c to x. But in that case, Lemma 9.6.12 asserts that there exists a root $D^+(r)$ containing c and x but not g. Thus g is in the opposite root $D^-(r)$.

Now a minimal gallery $G = (c = c_0, \ldots, c_m = g)$ from c to g connects a chamber of $D^+(r)$ to a chamber g of $D^-(r)$, and so there is a first edge $e := (c_{i-1}, c_i)$ which bridges the partition $C(M) = D^+(r) + D^-(r)$. Similarly, in a minimal gallery H of S from g to x, there is a first edge f connecting a chamber of $D^-(r)$ to one in $D^+(r)$ (see Fig. 9.10). By part 2 of Theorem 9.6.10, these bridging

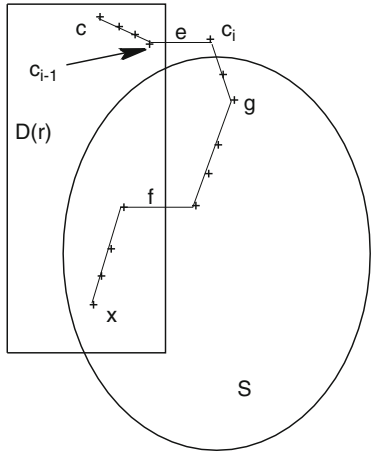

Fig. 9.10 The configuration of a root and a residue in proving strong-gatedness of the residue. The root marked $D(r)$ is $D^+(r)$

edges belong to the wall $E(r)$ and so are r-invariant. Since S is the unique residue of its type containing the edge f, we must have $S^r = S$. In particular $g^r \in S$. But now

$$(c_0, \ldots, c_{i-1}, c_{i+1}^r, c_{i+2}^r, \ldots, c_m^r = g^r)$$

is a gallery from c to g^r whose length is one less than the length of gallery G which is $d(c, S)$ by 1. Thus $g^r \in S$ and yet $d(c, g^r) < d(c, S)$, an absurdity.

Thus no such quartets (c, g, x, S) satisfying 1 and 2 above can exist. This means that for fixed choice of c and residue S, and then for any g for which $d(c, g) = d(c, S)$, we have

$$d(c, x) = d(c, g) + d_R(g, x) \text{ for all } x \in S.$$

Thus S is strongly gated in $C(M)$. □

Corollary 9.6.15 *Let (W, R, M) be a Coxeter system.*

1. (P) *Any two minimal galleries connecting the same two chambers of $C(M)$ are M-homotopic.*
2. *In the monoid I^*, any two reduced words w_1 and w_2 with $\rho(w_1) = \rho(w_2) \in W$ are M-homotopic words.*

Proof 1. We have just proved that every residue of the Coxeter chamber system $C(M)$ is strongly gated – i.e., (RG) holds. Hence (RG_2) holds, and so by Theorem 9.3.13 condition (Λ-min) holds, that is, any two minimal galleries connecting the same two chambers are Λ-homotopic. But in $C(M)$ an elementary Λ-homotopy is an elementary M-homotopy. Thus (Λ-min) becomes condition (P).

2. This M-homotopy among minimal galleries with the same initial and terminal chambers of $C(M)$ induces (via the type-mapping *typ*) the M-homotopy of words described in part 2. □

Theorem 9.6.16 *Suppose (W, R, M) is a Coxeter system over I. Suppose w is a reduced word of I^* and that for some letter $i \in I$,*

$$\ell(wi) < \ell(w).$$

Then w is M-homotopic to a word w' ending in i.

Proof This argument can be played out in the Cayley graph $C(M)$. Suppose G is a minimal gallery of $C(M)$ of type w, beginning at 1 (the identity element of $W(M)$). (It then necessarily terminates at $t := \rho(w)$.) Now on t there is a unique edge (t, t') of type i, and the hypothesis $\ell(wi) < \ell(w)$ tells us that t' is one unit closer to 1 than is t. Thus there is a minimal gallery H of length $\ell(w) - 1 = \ell(wi)$ from 1 to t' and now G and $H \circ (t', t)$ are two minimal galleries running from 1 to t.

By Corollary 9.6.15, part 1, these two galleries are M-homotopic, and so, taking types of the galleries, there is an M-homotopy of words:

$$typ(G) = w \to w' = typ(H \circ (t', t)),$$

where w' ends in i. □

9.6.6 When Is a Coxeter Group Finite?

The title of this section can be recast in this way: *For which matrices M is $W(M)$ a finite group?* This question was completely answered by H.S.M. Coxeter in 1934 [43]. First, it is immediate that if $W(M)$ is to be a finite group, the matrix M can have no entry equal to "∞." Second, M itself can have only finitely many rows and columns, for otherwise, using the Tits' form of Theorem 9.6.5, there would be a chain of subgroups of unbounded orders. Third, if there is a partition $I = A + B$ for which every generating involution in R_A commutes with every generating involution in R_B (so that $m_{ij} = 2$ for all $(i, j) \in A \times B$), then the Coxeter group $W(M)$ is just the direct product

$$W(M) = W_A \times W_B,$$

and so is finite if and only if each direct factor is finite. Thus one need only consider Coxeter systems in which there is no such partition. This is equivalent to saying that in the Dynkin notation, the diagram $D(M)$ is connected. We call such Coxeter groups $W(M)$, *irreducible Coxeter systems*.

Theorem 9.6.17 (Coxeter) *The Coxeter group $W(M)$ is finite if and only if the connected components of its diagram $D(M)$ are among those listed in Fig. 9.11.*

The diagrams listed in Fig. 9.11 are called the *Dynkin diagrams*.

In the next major section of this chapter we are going to define buildings as a class of chamber systems of type M with certain properties which we would like to keep as simple as possible. Our intention is to derive everything from (the seemingly tautological condition) (G_c) without ever once mentioning a simplicial complex. Everything is in terms of graphs.

But our ultimate objective in this book is to characterize point-line geometries. Merely to show that the most interesting geometries of this type are truncations of homomorphic images of geometries whose chamber systems are of type M and satisfy conditions like (G_c) would have no point if one could not classify the latter.

That is why Tits' theorem classifying all buildings of rank at least three and type M, where $W(M)$ happens to be a finite group (the buildings of "spherical type") is such a beacon. Obviously one cannot mention the latter without mentioning Coxeter's classification of finite Coxeter groups mentioned just above. In the eyes of the author, Tits' classification (not to mention his classification of buildings of affine type of rank at least four) is the "theorem of the century" (to borrow an appellation of Shreeram Abhyankar). We can now prove that point-line geometries

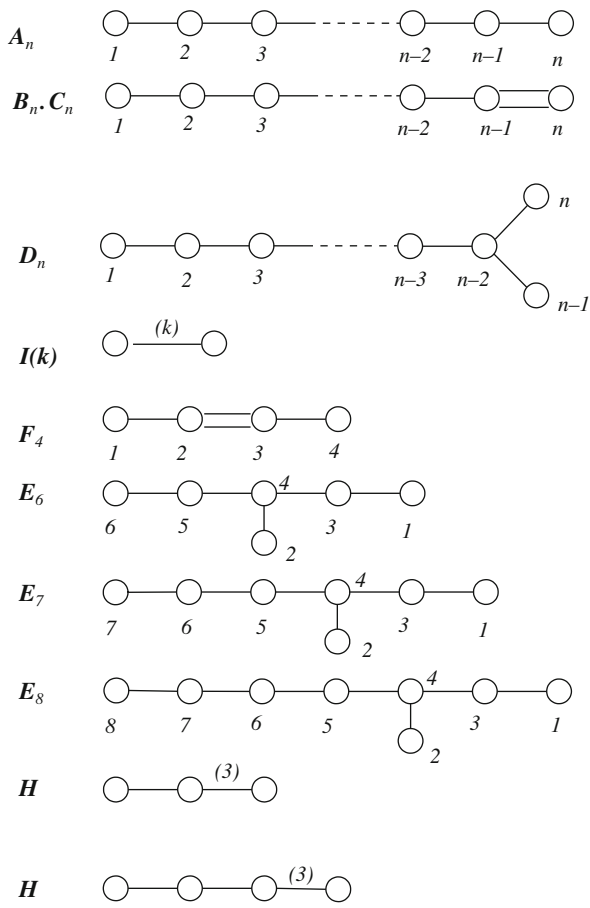

Fig. 9.11 The Dynkin diagrams. These are the diagrams of the irreducible Coxeter systems for which the Coxeter group is finite. (The node labeling is not altogether standardized in the literature. Following Cohen's article in the *Handbook* [35], we have used the numbering adopted by Bourbaki [9]

have real classifiable conclusions because of Tits' theorem. Otherwise one would be condemned to proving that certain axioms imply miscellaneous properties. One could really classify nothing.

9.7 Chamber Systems of Type *M*

9.7.1 Introduction

Let M be a *Coxeter matrix* $M = (m_{ij})$, that is, a symmetric matrix M with all diagonal entries equal to 1, and every other entry equal to either a positive integer greater than one, or the symbol "∞".

Recall that a chamber system $C = (V, E, \lambda)$ over I is said to be a *chamber system of type* M if and only if M is a Coxeter matrix with rows and columns indexed by I, and for every 2-subset $\{i, j\}$ of I, every residue of type $\{i, j\}$ is a generalized m_{ij}-gon.

Note that in a chamber system $C = (V, E, \lambda)$ of type M, each edge of the chamber system must bear a unique label– that is condition (typ) holds. For suppose some edge $e = (c, c')$ carried two or more labels $\lambda(e) = \{i, j. . .\}$. Then setting $J = \{i, j\}$ we see that e and its vertices c and c' belong to a connected component R of the graph (V, E_J), which is a residue of type $\{i, j\}$. But since C is type M, R is the chamber system of a generalized m_{ij}-gon, and so cannot have repeated labels.

The *rank of a chamber system of type* M is the cardinality of the set I indexing the rows and columns of the matrix M.

Along with every chamber system of type M we inherit an entire Coxeter system $(W = W(M), R, M)$ and its associated chamber system $C(M)$ – and in fact, *all* of the paraphenalia of the previous section on Coxeter groups and their Cayley graphs. This means we have the following:

1. A monoid epimorphism $\rho : I^* \to W$ from the free monoid I^* over I onto the Coxeter group $W(M)$ defined by sending letter i to the involution $r_i \in R$.
2. A "coordinatization" μ_c of the chambers of $C(M)$ by the elements of W which assigns the identity element $1 \in W(M)$ to chamber c.
3. A bijection ϕ_c:

$$\text{words in } I^* \longrightarrow \text{walks in } C(M) \text{ beginning at chamber } c.$$

4. The notion of a *reduced* word in W, that is, a word w in the monoid I^* whose length is the distance from 1 to $\rho(w)$ in the Cayley graph $C(M)$.

All these things are in place once a chamber system C of type M is even mentioned. One immediately notices that not one of the concepts listed just above involves the Earthly chamber system C. This is all going on in some sort of Chinese Celestial Heaven above C involving the free monoid and the "Coxeter world"!

However, because of condition (typ), there is associated with each gallery $G = (c_0, c_1, \ldots, c_k)$ (that is, a walk in C) a unique word $\text{typ}(G) = i_0 i_1 \cdots i_{k-1}$ in the monoid I^*, called the *type of the gallery* G, where $i_j := \text{typ}((c_j, c_{j+1}))$ for $j = 0, \ldots, k-1$. (We defined this before for the Coxeter chamber systems; here we are simply observing that the definition makes sense for *any* chamber system of type M.)

Noting that each chamber of C lies on an edge with each possible label, we see that for any chamber c of C, restriction of the type function produces a *surjection*

$$\text{typ}_c : \{\text{walks of } C \text{ beginning at } c\} \to I^*.$$

Unlike the case for $C(M)$, this is normally not an injection. Many galleries of C of the same type can begin at chamber c.

9.7.2 The Three Levels of Homotopy

Homotopy theories for a chamber system of type M simply amount to citing a set of rules (called "elementary homotopies") for replacing a gallery by another without disturbing the initial and terminal chambers of the gallery. The reverse replacement "undoing" an elementary homotopy is also required to be an "elementary homotopy." Then two galleries become "homotopic" (so the theory runs) if and only if they are connected by a chain of elementary homotopies. (The chain may have length zero, so the identity relation is a homotopy.) Then the relation of being "homotopic" is an equivalence relation on galleries and so "homotopy classes of galleries" are naturally defined. (Since we did not even specify the particular definition of homotopy, one can see that this is a rather general concept.)

In Chap. 1, we studied homotopies in arbitrary graphs which were defined by a family of circuits. We proved that for any graph $G = (V, E)$, and any family $\mathcal{C}$ of its circuits, there exists a universal $\mathcal{C}$-cover of G. In this sort of homotopy, backtracks $P \circ P^{-1}$ are always $\mathcal{C}$-contractible so we cannot expect this sort of homotopy to preserve the length of a gallery. We shall be interested in this for $\mathcal{C}_2$-homotopy.

A second homotopy theory, M-homotopy, is a special case of the first – that is M-homotopic strings or galleries will already be $\mathcal{C}_2$-homotopic. So the former is a finer theory. Both theories have already been discussed for the Coxeter chamber systems $C(M)$. We want to establish how they work for general chamber systems of type M, as well as the free monoids.

So our homotopies will be discussed at three levels:

- Homotopy of words in the free monoid I^*.
- Homotopy of walks beginning at the identity element 1 in the Cayley graph of the Coxeter groups $W(M)$ (in other words, the homotopy of galleries beginning at 1 in the Coxeter chamber system $C(M)$).
- Homotopy of galleries of C.

The relation between the first two levels is basically a bijection, as has been discussed in the guise of an earlier paradigm. The relation between the third and the first two is not so tight.

For our two relevent homotopy theories, we shall show two versions of our earlier paradigm:

1. (Going up in levels.) Any homotopy of galleries at the level of the chamber system C of type M determines a corresponding unique homotopy of words in I^* and of galleries beginning at 1 in $C(M)$.
2. (Going down in levels.) A partial converse in the form of an existence theorem will assert that if $w_1 \to w_2$ is a suitable homotopy of words, and G is a gallery of C of type w_1, then there *exists* a corresponding homotopy of galleries of C taking G to a gallery H of type w_2 with the same initial and terminal vertices as G.[23]

[23] Of course now the H is not necessarily unique, as it was in $C(M)$, but this enough to prove equivalence of several properties of C.

Some Notation for Words of I^*

We require a little notation concerning words in the monoid I^*. If $w = a_1 \cdots a_n$ is a word in the free monoid I^* spelled out by the letters $a_i \in I$, then w^* will always denote the *reverse word* $a_n a_{n-1} \cdots a_2 a_1$. This is a useful notation, for if the gallery G of a chamber system C of type M has type $w = \text{typ}(G)$, then the reverse gallery G^{-1} must have type w^*.

If the coefficient m_{ij} of the Coxeter matrix M is finite, and $0 < k < m_{ij}$, define the following words in the monoid I^*:

$$p(i, j) := ijij \cdots (\text{length } m_{ij}), \tag{9.30}$$

$$p(i, j; k) := ijij \cdots (\text{length } k). \tag{9.31}$$

Note that if m_{ij} or k is even, then $p(i, j)$ and $p(i, j; k)$ are respectively words beginning in i and ending in j. Similarly, if m_{ij} or k is odd, these respective words begin and end in i. Thus we have these factorizations in I^*:

$$(ij)^{m_{ij}} = p(ij)^2, \text{ if } m_{ij} \text{ is even}, \tag{9.32}$$

$$(ij)^{m_{ij}} = p(i, j : k)p(i, j; m_{ij} - k) \text{ if } k \text{ is even}, \tag{9.33}$$

$$(ij)^{m_{ij}} = p(i, j)p(j, i) \text{ if } m_{ij} \text{ is odd}, \tag{9.34}$$

$$(ij)^{m_{ij}} = p(i, j; k)p(j, i; m_{ij} - k) \text{ if } k \text{ is odd}. \tag{9.35}$$

In order to avoid constant case divisions according to the parity of m_{ij} or k, it will be convenient to have a common notation for the right-most factors of the equations above. This factor is written $p(i, j)^T$ in the first and third equations, and written $p(i, j : k)^T$ in the second and fourth. Then, whenever m_{ij} is finite, we always have $(ij)^{m_{ij}} = p(ij)p(ij)^T = p(i, j; k)p(i, j; k)^T$.

$\mathcal{C}_2$-Homotopy of C

For any chamber system C of type M, $\mathcal{C}_2$ will denote the collection of all circuits in C that lie within some rank two residue of C. Then along the lines of Chap. 1, $\mathcal{C}_2$-homotopy is defined, and universal $\mathcal{C}_2$-covers always exist.

Now in a chamber system C of type M, any circular gallery G expressible as a concatenation $P_1 \circ P_2$, where P_1 has type $p(i, j : k)$ and P_2^{-1} is a gallery of type $p(i, j; k)^T$ – that is $p(i, j; m_{ij} - k)$ or $p(j, i; m_{ij} - k)$ according to whether k is even or odd – is a circuit of (panel-reduced) type $(ij)^{m_{ij}}$ residing within some rank two residue of C of type $\{i, j\}$ – a residue which is a generalized m_{ij}-gon.[24]

Then an *elementary $\mathcal{C}_2$-homotopy of a gallery of C* is a transformation

$$A \circ P_1 \circ B \to A \circ P_2 \circ B,$$

[24] Note that the assumption that G is circular is necessary. Unlike the Coxeter chamber system $C(M)$, for a general chamber system of type M a gallery of type $(ij)^{m_{ij}}$ need not be circular.

where either:

1. one of $\{P_1, P_2\}$ is a spur (x, y, x) and the other is the gallery of length 0 at x, or,
2. for some $i, j \in I$ and appropriate integer k, P_1 is a gallery of type $p(i, j : k)$ and P_2^{-1} is a gallery of type $p(i, j; k)^T$ — that is $P_1 \circ P_2^{-1}$ is a circuit of type $(ij)^{m_{ij}}$. (Note that P_1 and P_2 have lengths k and $2m_{ij} - k$, respectively.)

But these two transformations determine similar transformations of words of I^*:

$$u \circ ii \circ w \to u \circ w$$

for some i or its reverse, or

$$u \circ p(i, j; k) \circ w \to u \circ (p(i, j; k)^T)^* \circ w$$

either of which we call an *elementary $\mathcal{C}_2$-homotopy of words*.

But now, applying the monoid morphism $\rho : I^* \to C(M)$, this same elementary homotopy of words induces a $\mathcal{C}_2$-homotopy of galleries of the Coxeter chamber system $C(M)$, as explained in an earlier section. So we have $\mathcal{C}_2$-homotopies occurring at three levels: (1) in words of I^*, (2) in galleries of the Coxeter chamber system $C(M)$ which begin at the identity element $1 \in W(M)$, and (3) in galleries of C. The latter determines the two former.

Before leaving $\mathcal{C}_2$-homotopy theory, there is an important fact to observe.

Lemma 9.7.1 (Existence of $\mathcal{C}_2$-homotopies in C.) *Suppose C is a chamber system of type M. Suppose $w_1 \to w_2$ is a $\mathcal{C}_2$-homotopy of words. Then, for any gallery G_1 of C of type w_1, there exists a gallery G_2 of type w_2 $\mathcal{C}_2$-homotopic to G_1.*

Proof One merely verifies the assertion at an elementary homotopy. □

M-Homotopy in C

We have already met *special M-homotopy* in the context of Coxeter chamber systems.

An *elementary M-homotopy of words* of I^* is a replacement of words of the form

$$w \circ p(i, j) \circ v \longrightarrow w \circ (p(i, j)^T)^* \circ v$$

where $p(i, j)$ is the word of type of type $ijij\cdots$ of length m_{ij}, $p(i, j)^T = p(i, j)$ or $p(j, i)$ according to whether m_{ij} is even or odd, and the "star" operator denotes reversal of a word.[25] Then two words of I^* are M-homotopic if and only if they are connected by a series of elementary M-homotopies. Clearly M-homotopic words have the same length.

---

[25] A memnonic device is this: $p(i, j)^T$ must always end in "j." Thus $(p(i, j)^T)^* = p(j, i)$ since it is spelled in $\{i, j\}$ without "double letters," has the right length, and begins with letter j.

Applying the monoid morphism ρ, any M-homotopy of words is converted into an M-homotopy of galleries of the Coxeter chamber system $C(M)$.

Meanwhile back on Earth we have the 'mundane' chamber system C of type M. Even here, Heaven can be "mirrored" by replacing a gallery G_1 of C from chamber s to t expressible as a concatenation $A \circ P_1 \circ B$ by another gallery $G_2 := A \circ P_2 \circ B$ from s to t where $\mathrm{typ}(P_1) = p(i, j)$ and $\mathrm{typ}(P_2) = (p(i, j)^T)^*$). This distortion defines an *elementary M-homotopy of galleries in C* which does not change the length of the gallery. As usual this extends to a definition of *(special) M-homotopy* among the walks of graph C – that is, among the galleries of the chamber system C.

With these notions in place, we may easily deduce the second basic property of chamber systems of type M.

Lemma 9.7.2 *Suppose C is a chamber system of type M.*

(1) (Existence of M-homotopies in C.) *Suppose G is a gallery in C of type $w \in I^*$ connecting chamber x to chamber y. If the word w is M-homotopic to the word v in I^*, then gallery G is M-homotopic in C to a gallery H having type v.*
(2) Conversely, if two galleries of C are M-homotopic, then there is a corresponding M-homotopy of words in $M(I)$ connecting their types.

In other words, any special M-homotopy that can be formed in Heaven can be modeled on Earth (though not necessarily in a unique way). Similarly, any M-homotopy down in C is forever inscribed in Heaven as an M-homotopy of words in I^*.

Proof For any elementary M-homotopy in C, the segment P_1 is a gallery of type $p(i, j) = ijij\cdots$ in a residue which is a generalized m_{ij}-gon. By the standard generalized polygon conditions, there exists a gallery P_2 of type $p(j, i)$ beginning and ending at the same chambers.

The second part is just a consequence of our definitions. □

Geodesics in C have Reduced Type

We are now ready for one more basic property of chamber systems of type M. Recall that a gallery G is a *geodesic* in the chamber system C if and only if it is a shortest possible path connecting its extremeties.

Lemma 9.7.3 *Suppose C is a chamber system of type M. If G is a geodesic in C, then the type of G is a reduced word.*

Proof If the type of G were a word which was *not* reduced then there would be a series of elementary $\mathcal{C}_2$-homotopies of words in I^* which changed w to a shorter word v. Now by the second part of Lemma 9.7.1 there must exist a gallery H of type v which is $\mathcal{C}_2$-homotopic to G. But in that case the length of v, which is the length of H, is shorter than the length of G, but still connects the initial and terminal chambers of G. This contradicts the fact that G was a geodesic. □

9.8 Buildings

9.8.1 Introduction

Buildings are a species of connected chamber systems of type M – possibly of infinite rank. There are many properties which become equivalent in the context of connected chamber systems of type M. We have already shown the equivalence of the three strong-gatedness properties (RG^1), (RG), and (RG_2) for arbitrary chamber systems with condition (typ). In the context of chamber systems of type M any of these conditions are equivalent to the following:

1. Condition (G_c). Any gallery of reduced type beginning at chamber c is a geodesic (or minimal) gallery.
2. Condition (P_c) (Tits). Any two galleries of reduced type connecting chamber c to another chamber are M-homotopic.
3. Condition (G_x). By this we intend to assert condition (G_x) for every chamber x.
4. Condition (P_x). Similarly, the assertion that (P_x) holds for every chamber x.

In the next section we shall demonstrate the equivalence of these conditions for connected chamber systems of type M by means of the following chain of proved results:

1. (Theorem 9.8.1) The conditions (P_c) and (G_c) are equivalent.
2. (Theorem 9.8.2) The conditions (G_c) and (G_x) are equivalent.
3. (Theorem 9.8.4) The condition (G_c) implies (RG).
4. (Theorem 9.8.5) The condition (P_c-min) together with (RG) implies (G_c).

How does this imply the equivalence? We are missing only (RG) implies (G_c). But from the section on chamber systems with strongly gated residues, we proved that (RG) trivially implies (RG_2) which in turn implies condition (Λ-min) (Theorem 9.3.13) . In the context of chamber systems C of type M, the condition (Λ-min) is manifestly as follows.

(P_c-min) Any two minimal galleries connecting the same two chambers of the chamber system C are M-homotopic.

So the assumption (RG) gives us (P_c-min) as well, and then by the fourth result enumerated above, one has (G_c).

We consider a connected chamber system of type M satisfying any one of the properties listed above, as well as any one of the three strongly gated conditions (all seven of which conditions are equivalent) to be a *building*.

In a subsequent section, we show that these conditions are also equivalent to the formulation of Ronan and Tits in which there is a Coxeter-group-valued metric with relatively simple properties[26] and then show (as in Ronan's book) that all of this is

[26] See Ronan [103], Chap. 3.

equivalent to the traditional formulation of the definition of "building" involving a Tits' system of apartments.

Of course this means that now there are very simple ways to define a building in the context of connected chamber systems of type M. What could be simpler than saying that all corank one residues are gated (RG^1)? Or to assert condition (G_x) that all galleries of reduced type are minimal galleries? To this author, simple definitions have an appeal well beyond the fact that such definitions are easier to teach: rather, it is that the more simple a concept, the more one acquires the esthetic feeling that it is forced on us by the nature of the universe – or at least our mental universe. But there is a down side: simple definitions can give a false illusion of understanding, making invisible all the structural complexity underneath. That is in fact the way it is with buildings. We have merely caressed the upper surface.

9.8.2 The Conditions (G_c) and (P_c)

In the previous section we noted that in any chamber system of type M, the type of any geodesic gallery is a word of reduced type (Lemma 9.7.3). The relevant property here is in fact the converse of this assertion. Suppose C is a chamber system of type M and fix a chamber c in C. Then we define this property.

(G_c) *Every gallery of reduced type beginning at chamber c is a geodesic.*

Tits introduced the following (equivalent) condition in [139].

(P_c) *If two galleries G and H have reduced type, and have the same extremities, then these two galleries are M-homotopic.*

The Equivalence of (G_c) and (P_c)

Theorem 9.8.1 *The properties* (P_c) *and* (G_c) *are equivalent.*

Proof (G_c) implies (P_c). Let G and H be galleries of reduced types $w = \mathrm{typ}(G)$ and $u := \mathrm{typ}(H)$ beginning at chamber c and terminating at chamber t. Let s be the next-to-last chamber of gallery H, so s is i-adjacent to t. Now by (G_c), $d = d(c, t)$ is the length of words $\mathrm{typ}(H)$ and $\mathrm{typ}(G)$; but as $d(c, s) = d - 1$, by (G_c) the gallery $G \circ (t, s)$ cannot be of reduced type. Thus the concatenation $\mathrm{typ}(G) \circ i$ is not reduced and so its length is less than $\mathrm{typ}(G)$. Then by Theorem 9.6.16, $\mathrm{typ}(G)$ is M-homotopic in I^* to a word $w' \circ i$ which is reduced. Then by Lemma 9.7.2 above, G is homotopic to a gallery $G' \circ (r, t)$ where r is i-adjacent to t and G' is of reduced type $w' = \mathrm{typ}(G')$.

If $r = s$, then by induction on the length, G' is homotopic to H', the subgallery of H from c to s. Thus

$$G \sim G' \circ (s, t) \sim H' \circ (s, t) = H$$

and we are done.

If $r \neq s$, $G' \circ (r, s)$ is the same reduced type as $G' \circ (r, t)$. But the former is not a geodesic as $d(c, s) = d - 1$, contradicting (G_c).

(P_c) implies (G_c). Now let $G = G' \circ (y, t)$ be a gallery of reduced type $w' \circ i$ (that is, G' has type w' and the final edge (y, t) is labelled i) chosen of minimal length so $w' \circ i$ is reduced, but the distance of its terminus t from c is less than its length. Then, as w' is reduced, G' is a geodesic from c to y of length $d = d(c, y)$, and $d(c, t) = d$ or $d - 1$. Let H be any geodesic from c to t. Then H has reduced type $\mathrm{typ}(H)$ by Lemma 9.7.3. By (P_c), H and G are M-homotopic galleries. But that is impossible as g has length $d + 1$ while H has length d or $d - 1$. Thus no such gallery G can exist – i.e., (G_c) holds. □

Transporting (G_c) to Other Chambers

Theorem 9.8.2 *A chamber system C of type M satisfies condition* (G_c) *if and only if it satisfies condition* (G_x) *for all its chambers x.*

Proof We need only prove the forward implication, and, as the chamber system C is connected, we need only show that (G_c) implies $(\mathrm{G}_{c'})$ where c' is any chamber j-adjacent to c. Let G be a gallery of reduced type w beginning at c'. We must prove that G is a geodesic, and as the result is obvious if G has length 1, we may assume all shorter galleries of reduced type starting at c' are geodesics.

Case 1: $\ell(jw) > \ell(w)$ – i.e., jw is a reduced word. Then the augmented gallery $(c, c') \circ G$ has reduced type jw and so is a geodesic by (G_c). Then its long "tail," G, is also a geodesic.

Case 2: $\ell(jw) < \ell(w)$. Then jw is not a reduced word. It then follows from Theorem 9.6.16 that the word w is homotopic to a word beginning with j, and so by Lemma 9.7.2 (2), G is M-homotopic to a gallery beginning with a j-adjacency. Since the latter gallery is a geodesic if and only if the former gallery is a geodesic, we may, without loss, assume G *is* this latter gallery. Thus we write

$$G = (c' = c_0, c_1, \ldots, c_d)$$

of reduced type jw', so (c_0, c_1) is a j-adjacency.

Subcase 2.1. $c_1 \neq c$. Then letting G' be the subgallery of G running from c_1 to c_d, we see that $(c, c_1) \circ G'$ is also of reduced type jw' and so is a geodesic. Thus $d(c, c_d) = d$. If then $d(c', c_d) = d$ we are done as G is then a geodesic. Thus we must assume $d(c', c_d) = d - 1$. There is thus a geodesic gallery H' of length $d - 1$ running from c' to c_d, say of reduced type w''. Then as $d(c, c_d) = d$, $(c, c') \circ H'$ is a minimal gallery, so its type jw'' is reduced. Now $(c, c') \circ H'$ and $(c, c_1) \circ G'$ are two galleries of reduced types jw'' and jw' terminating at c_d, and so by condition (P_c), the galleries are M-homotopic. By Lemma 9.7.2 (2), their respective types jw'' and jw' are homotopic words. It follows that w'' and w' are also homotopic words. Thus by Lemma 9.7.2 (1), there exists a gallery H from c' to c_d of reduced type w'. We have now the configuration of Fig. 9.12.

Now let H be the gallery $(c' = h_1, \ldots, h_d = c_d)$ of length $d - 1$. Then h_{d-1} and c_{d-1} are both k-adjacent to c_d, since h and g' are *both* type w'.

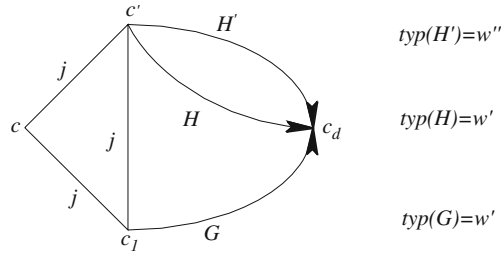

Fig. 9.12 The configuration of Subcase 2.1

Assume now $c_{d-1} \neq h_{d-1}$. Then the gallery

$$(c, c', h_2, \ldots h_{d-1}, c_{d-1})$$

(marked by the dotted path in Fig. 9.13a) is reduced type jw' but travels only a distance $d(c, c_{d-1}) = d - 1$. This contradicts (G_c).

On the other hand, if $h_{d-1} = c_{d-1}$, the gallery

$$(c', c_1, \ldots, c_{d-1})$$

(marked by the dotted path in Fig. 9.13b) has reduced type (a factor of jw') and length $d - 1 > d - 2 = d(c', c_{d-1})$ against the minimality of G.

Subcase 2.2. $c = c_1$. We have $d(c', c_d) \leq d(c, c_d) = d - 1$ despite the fact that $G = (c', c) \circ G'$ has reduced type jw'. We may then supply a minimal gallery K of reduced type v from c' to c_d so $|K| \leq d - 1$. We then have the configuration of Fig. 9.14.

Suppose first that jv is a reduced word. Then by (P_c), jv is a word homotopic to w'. But in that case jw' would be a word M-homotopic to jjv and so could not be a reduced word.

Thus we may assume that jv is not reduced. Then, as v is reduced and jv is not, Theorem 9.6.16 tells us that v is M-homotopic to a word jv', and so by Lemma 9.7.2 (1), K is M-homotopic to another minimal gallery $(c', a) \circ K'$ of type jv'. The picture is now that of Fig. 9.15, since $d(c, c_d) = d - 1 \neq d(a, c_d)$ forced $a \neq c$.

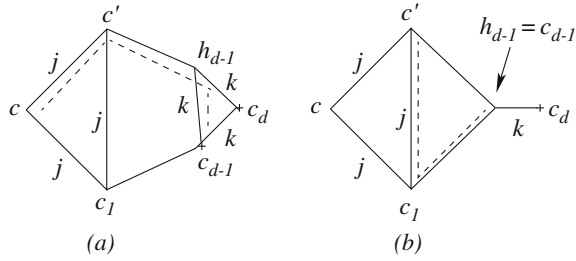

Fig. 9.13 Two impossible galleries

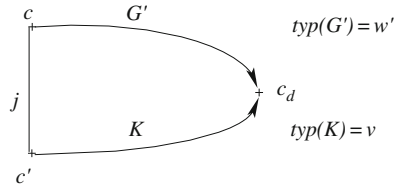

Fig. 9.14 The configuration of Subcase 2.2

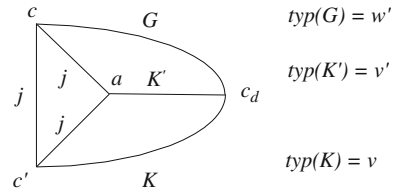

Fig. 9.15 Refinement of Subcase 2.2

Then $(c, a) \circ K'$ is a gallery of reduced type jv' and so by (P_c) is M-homotopic to G', whence by Lemma 9.7.2 (2), w' is homotopic to jv'. Then jw' is homotopic to jjv' and this contradicts the case assumption that jw' is a reduced word. The proof is complete. □

The Condition (G_c) Implies Condition (RG)

We begin with a lemma.

Lemma 9.8.3 *The following statements hold:*

1. *Let C satisfy condition (G_c). Then all residues of C are isometrically embedded induced subgraphs of C. Moreover, they are convex.*
2. *As a chamber system, any residue R also satisfies (G_c).*

Remark One may note that if our definition of building was (G_c) then part 2 asserts that every residue of a building is also a building. This can be seen as an extension of the so-called "parabolic subgroup theorem" to chamber systems of type M. Part 1 of this theorem is the most important. It is in fact the perfect analog of Lemma 9.6.13 for Coxeter chamber systems moved into the more general realm of chamber systems of type M. But now we cannot talk about walls and roots and methods of measuring distance by the number of roots contining one chamber but not the other. We need an entirely new proof.

Proof 1. Suppose x and y are two chambers in a residue R of type J in C. Then at least R is a chamber system of type M_J, the Coxeter matrix restricted to the rows and columns indexed by J. Let G be a shortest gallery of the chamber system R connecting x and y. By Lemma 9.7.3 applied to R, the type w of gallery G is a *reduced* word in $M(J)$.

By a well-known property of Coxeter systems (Lemma 9.6.13, for example), w is a reduced word of I^*. Similarly, if H is any geodesic of C connecting x to y, H also has reduced type. By condition (P_x), H is M-homotopic to G, since these two galleries have the same extremeties and are of reduced type. Thus we see two things:

1. G is a geodesic of C.
2. H is already a geodesic of the residue R.

Statement 1 shows that R is isometrically embedded. Statement 2 shows that R is an induced convex subgraph.

2. From 2 it is clear that (G_c) holds relative to the chamber system R.

The proof is complete. □

Theorem 9.8.4 *For a chamber system C of type M, condition (G_c) implies condition (RG).*

Proof In view of Lemma 9.8.3 all residues are convex induced subgraphs of C, and so ordinary gatedness (in the sense of Scharlau and Dress) implies strong gatedness. Thus it suffices to show that any residue is gated in the ordinary sense.

Choose a residue R of type J and a chamber c and let p_1 be a chamber of R nearest c. If we show $d(c, p_1) + d(p_1, r) = d(c, r)$ for all $r \in R$, the uniqueness of p_1 and the gatedness of R will be shown in one stroke. For this purpose, it is clear that c can be assumed not to lie in R (for then the distance equation follows with $c = p_1$).

We prove the distance equation by induction on $d(p_1, r) = d$, it being true when $d = 0$. So, if the equation is false for some r, it is false for some t of minimal distance from p_1. Thus we have a minimal gallery $G' \circ (s, t)$ of length d and reduced type wj, from p_1 to t and a minimal gallery G_1 of reduced type u from c to p_1. Minimality of $d(c, p_1)$ (which follows from the choice of p_1) implies that the word u *is not M-homotopic to a word ending in a letter of J, the type of the residue R.*

Now uw is reduced as $d(c, s) = d(c, p_1) + d(p_1, s)$ holds by the minimality of d. But uwj is *not* reduced, since otherwise, by (G_c), $G_1 \circ G' \circ (s, t)$ would be a geodesic against the choice of t. Thus in the Coxeter system W, $\ell(uwj) < \ell(uw)$. Thus by Theorem 9.6.16 uw is M-homotopic (in the monoid I^*) to a word vj ending in j.

Let ρ denote the canonical monoid homomorphism $\rho : I^* \to W$ which evaluates words in I^* as products of generating involutions of W. As we have seen in the previous section, the Coxeter chamber system $C(W)$ has its residues gated. This means that if J is the type of the residue R of C and $C(W)$ is coordinatized by W, then $\rho(u)$ is the gate of the residue $\rho(u)W_J$ with respect to the chamber coordinatized by the identity element (that is, since u is not M-homotopic to a word ending in a letter from J, u is the word of shortest length representing an element in the coset $\rho(u)W_J$). But v is a word representing an element in this coset $e(u)W_J$. Thus v is M-homotopic to a word uv'. We now have these M-homotopies:

$$uw \sim vj \sim uv'j \text{ so } w \sim v'j.$$

But this contradicts the fact that as $G' \circ (s, t)$ was a geodesic, the word wj must be reduced. Thus no such chamber t exists. □

Theorem 9.8.5 *If condition* $(P_c$ *min* $)$ *holds for all chambers* c *of* C*, and the gated hypothesis (RG) holds, then* (G_x) *holds for all chambers* $x \in C$.

Proof We assume (G_c) false, and select a chamber c and gallery $G = (c = c_0, \ldots, c_d)$ of reduced type wij of minimal length so that G is not a geodesic. Then $d(c, c_{d-2}) = d - 2 = d(c, c_{d-1}) - 1$ and $d(c, c_d) \leq d(c, c_{d-1}) = d - 1$. Since wij is reduced, $i \neq j$. If $d(c, c_d) = d(c, c_{d-1}) = d - 1$, there exists, via the gatedness of rank one residues, a vertex s which is j-adjacent to c_{d-1} but distance $d - 2$ from c. On the other hand, if $d(c, c_d) = d - 2$, set $s = c_d$, so s has the same property in either case. Then for any minimal gallery H' from c to s, we see that $H' \circ (s, c_{d-1})$ is a minimal gallery, of type ending in j, stretching from c to c_{d-1}. But $G' := (c_0, \ldots, , c_{d-1})$ is also a minimal gallery from c to c_{d-1}. By $(P_c$ min), $H' \circ (s, c_{d-1})$ is M-homotopic with G'. Thus wi, the type of G', is M-homotopic to a word ending in j and this contradicts the hypothesis that wij, the type of G, is reduced.

This completes the proof. □

9.9 Apartments

Our aim is to show that in any building C of type M, there exists a family $\mathcal{A}$ of isometric embeddings $C(M) \to C$ of the Coxeter chamber system defined by M whose images are convex, and which for every pair of chambers, possesses an image covering them.

9.9.1 The Tits Metric of a Building

Let C be a building of type M. Then (P_c) holds for each chamber c. Thus given two chambers x and y, there is just one M-homotopy class of minimal galleries stretched from x to y, and hence just one M-homotopy class $w(x, y)$ of reduced words representing the types of these minimal galleries. Then $w(x, y)$, being an M-homotopy class of reduced words, can be regarded as an element of the Coxeter group, $W = W(M)$. Thus there is a well-defined function

$$w : C \times C \longrightarrow W$$

satisfying

$$(i)\ \ w(x, x) = 1_W \tag{9.36}$$

$$(ii)\ w(x, y) = (w(y, x))^{-1} \tag{9.37}$$

which we call the *Tits metric on* C.

We next observe:

- *The mapping* $\rho_a : x \to w(a, x)$ *is a morphism of* C *onto* W *as chamber systems.*

This is clear, for if x is j-adjacent to y, then $w(a, y) = w(a, x)$ or $w(a, x)\,j$ according to whether $d(a, y) \leq d(a, x)$ or not.

Since the morphism cannot increase distances (as with any morphism of chamber systems) we see

$$(iii)\ \ \ell(w(x, y)) \leq d(x, y) \tag{9.38}$$

Let X be any subset of W. A *strong isometry* $\alpha : X \to C$ is a mapping such that for all $x, y \in X$,

$$w(\alpha(x), \alpha(y)) = x^{-1}y.$$

We have the following.

Theorem 9.9.1 *Any strong isometry $\alpha : X \to C$ can be extended to an isometry of W into C.*

Proof (The proof we give here is essentially that appearing in Ronan's book [103], pp. 31–32.) By Zorn's lemma it is sufficient to show that for any proper subset X of W, a strong isometry can be extended to a larger set. We may assume $X \neq \emptyset$, and recoordinatizing W by pre-left-multiplication if necessary, that $1_W = x_0 \in X$. Since X is proper in W, $X \neq Xr$ for some involution r in S, and the coordinatization can be chosen so $x_0 r = r \notin X$. We therefore need only extend the isometry $\alpha : X \to C$ to $X \cup \{r\}$.

Case 1: Suppose $\ell(rx) > \ell(x)$ for all $x \in X$. Then we may choose $y = \alpha(r)$ to be any chamber r-adjacent to $\alpha(1_W)$ in C. This is a strong isometry on $X \cup \{r\}$ since rx reduced implies $w(\alpha(r), x) = rx$.

Case 2: $\ell(rx_1) < \ell(x_1)$ for some $x_1 \in X$. Then x_1 is homotopic to a word rf in W and we have the configuration in the Coxeter chamber system W given in Fig. 9.16

Then in C there is a unique gallery of reduced type rf from $\alpha(1_W)$ to $\alpha(x_1)$ by the fact that condition $(P_{\alpha(1)})$ holds in C (Theorem 9.8.1). Let y be the second

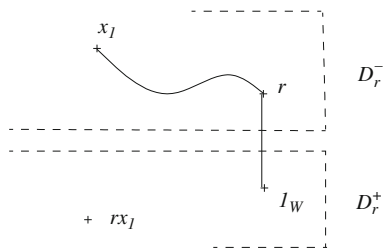

Fig. 9.16 The relation of x_1 to the wall E_r in W

member of this gallery, so $w(\alpha(1_W), y) = r$ and $w(y, \alpha(x_1)) = f$. Now for any $x \in X$, we have

$$\rho_{\alpha(x)}(a) = w(\alpha(x), y) \tag{9.39}$$
$$= w(\alpha(x), \alpha(1_W)) = x \text{ or} \tag{9.40}$$
$$= w(\alpha(x), \alpha(1_W))r = xr \tag{9.41}$$

since y is r-adjacent to $\alpha(1_W)$ and $\rho_{\alpha(x)}$ is a morphism.

Thus, using Eqs. (9.39) and (9.40) or (9.41) above,

$$w(y, \alpha(x)) = x \text{ or } xr \text{ for each } x \in X.$$

Now for each $x \in X$ define

$$\beta(x) := r \circ w(y, \alpha(x)).$$

Then $\beta : X \to W$ is the composition of a strong isometry $\alpha : X \to C$, a morphism $\rho_y : C \to W$, and left multiplication by r on W which is an automorphism of W. Thus β cannot increase distances as none of the factors $\alpha, \rho_y, x \to rx$, do so. Moreover, from the above,

$$\beta(x) = rw(y, \alpha(x)) = r(x \text{ or } rx) = rx \text{ or } x$$

and one calculates that

$$\beta(1_W) = r \circ w(y, \alpha(1_W)) = r \circ r = 1_W \text{ in } D_r^+$$

and

$$\beta(x_1) = r \circ w(y, \alpha(x_1)) = rf = x_1 \text{ in } D_r^-.$$

At this point one can show that β is the inclusion mapping. Suppose otherwise. Then for some $x \in X$, we have $\beta(x) = rx$. Now, if $x \in D_r^+$ (that is, its distance to 1_W in $C(M)$ is shorter than its distance to r) then $\beta(x) = rx$ is further from $\beta(1_W) = 1_W$ than was x – i.e., β has increased distance, a contradiction. On the other hand, if $x \in D_r^-$ (that is, it, like x_1, is closer to r than it is to 1_W) then $\beta(x) = rx$ is further from $\beta(x_1) = x_1$ than is x, and so x has again increased distance. Since β cannot do this, we conclude that no such x exists. Thus for all $x \in X$,

$$\beta(x) = x = r \circ w(y, \alpha(x))$$

so

$$w(y, \alpha(x)) = rx \text{ for each } x.$$

But since $y = \alpha(r)$, this shows that α satisfies

$$w(\alpha(u), \alpha(v)) = u^{-1}v \text{ for } u, v \in X \cup \{r\}$$

– i.e., the "new" α is a strong isometry $X \cup \{r\} \to C$.

This completes the proof. □

9.9.2 Strong Isometries and the Standard Apartment Axioms for a Building

We now let $\mathcal{A}'$ be the class of all strong isometries $W \to C$ obtained by extending possible strong isometries $X \to C$. Next let $\mathcal{A}$ denote the set of *images* of W in C under the various morphisms of $\mathcal{A}'$.[27] We have the following.

(9.9.1) *Any two chambers x and y of C lie in a common member of $\mathcal{A}$.*

Let G be a minimal gallery of type w from x to y. Then let X be the unique gallery of the Coxeter chamber system $W = C(S, M)$ of type w beginning at 1_w. Then there is a clear strong isometry $X \to G$ which extends to a member of $\mathcal{A}'$ whose image in C is an element of $\mathcal{A}$ contained in C . This step is finished. □

(9.9.2) *Any sub-chamber system $A \in \mathcal{A}$ is convex.*

Let x and y be chambers in A and let G be a minimal gallery of reduced type w connecting x and y. By Theorem 9.9.1 there is a morphism $\alpha : W \to C$ such that $A = \alpha(W)$ contains x and y. Since α is a strong isometry,

$$w(x, y) = [w] = (\alpha^{-1}(x))^{-1}(\alpha^{-1}(y)),$$

so $\alpha^{-1}(x)$ is connected to $\alpha^{-1}(y)$ in W by a unique gallery H of reduced type w. Then $\alpha(H)$ is a gallery of type w connecting x and y. By (Q_x), $\alpha(H) = G$ so $G \subseteq A$. This proves that A is convex. □

(9.9.3) *If A_1 and A_2 are two members of $\mathcal{A}$ containing chambers x and y, then there is a chamber system isomorphism $\mu : A_1 \to A_2$ fixing x and y and every vertex on any minimal gallery connecting them.*

First we observe that if H_1 and H_2 are two galleries of reduced type in the Coxeter chamber system W, then there is a unique automorphism of W taking H_1 to H_2.

[27] Note that distinct members of $\mathcal{A}'$ can lead to the same member of $\mathcal{A}$ by composing an isometry with an automorphism of its image.

This is because: (1) (G_c) for all $c \in W$ implies H_1 and H_2 are minimal, (2) on each chamber c there is a unique gallery of type w, and (3) $\mathrm{Aut}(W) \simeq W$ is regular on W.

Now let A_1 and A_2 be two members of $\mathcal{A}$ containing chambers x and y and let G be a minimal gallery of reduced type w connecting x and y and let $\alpha_i : W \to A_i, i = 1, 2$ be the strong isometries. Then by (9.9.3) $G \subseteq A_1 \cap A_2$. Then $H_i := \alpha_i^{-1}(G), i = 1, 2$ are two galleries of reduced type w in W. By the preceding paragraph, there is a unique automorphism $\beta : W \to W$ so that $\beta(H_1) = H_2$ (with the orientation preserved if w is a palindrome). Then the composition of mappings $\alpha_2 \circ \beta \circ \alpha_1^{-1}$ is an isomorphism $A_1 \to A_2$ fixing gallery G chamberwise.

If G' is a second minimal gallery from x to y of type w', then $e(w^{-1}w') = 1_W$ and so, for $i = 1, 2$, the two lifts $H_i' := \alpha_i^{-1}(G')$ are such that $H_i^{-1} \circ H_i'$ are circuits in W. It follows that $\beta(H_1') = H_2'$ and so the the composition $\alpha_2 \circ \beta \circ \alpha_1^{-1} : A_1 \to A_2$ fixes the vertices and edges of G' as well as G. □

In general, if $\mathcal{A}$ is a collection of thin chamber systems of type M isometrically embedded in chamber system C of type M so that (9.9.1), (9.9.2), and (9.9.3) hold, we say that $\mathcal{A}$ *is a system of apartments* for C.

Theorem 9.9.2 *Any building possesses a system of apartments.*

9.10 Appendix to Chapter 9: Spherical Buildings and (B, N)-Pairs

9.10.1 Tits Systems

It had been recognized that many classical groups are generated by two subgroups. The first is the stabilizer of a chamber of C: this is traditionally called the *Borel subgroup* and is denoted B for that reason. The other group is the stabilizer of an apartment A; this group is traditionally denoted N.[28] Of course at the very beginning one didn't have a clear notion of chamber system or what an apartment was. One knew from examples in finite rank that B was a maximal group of upper triangular matrices and that N acted on some sort of spanning frame to be realized as a monimial group. It took a while to recognize the axioms these groups obeyed.[29]

What are they?

---

[28] N seems not to have been named after any person, however obscure. Perhaps it was for "normalizer," for in practice N is the normalizer of $B \cap N$.

[29] Of course Tits did it all. But others were thinking about it. The infrequently-quoted paper "*Geometric ABA groups*" by Jack McCaughlin and Don Higman probably anticipated this development to a certain extent, even if only by a few months [75].

9.10.2 (B, N)-Pairs and Tits Systems

A *Tits system* is a quadruple (G, B, N, S) subject to these axioms:

(T1) G is a group generated by two subgroups B and N. Moreover, $B \cap N$ is a normal subgroup of N.

(T2) The factor group $W = N/(B \cap N)$ has a distinguished system S of generating involutions such that:

(a) (W, S) is a Coxeter system, that is S is a generating set of involutions, subject only to the relations given by the Coxeter matrix $M = (m_{ij})$.

(b) Moreover, for every involution $s \in S$ and element $w \in W$,

$$sBw \subseteq BwB \cup BswB.$$

(c) Finally,

$$sBs \neq B \text{ for any } s \in S.$$

Note that if n and n' are two elements of N which are congruent mod $N \cap B$, so $w = (B \cap N)n = (B \cap N)n' \in W$, then $Bn = Bn'$ which can unambiguously be written as Bw. This slight abuse of notation is convenient for expressing the above axiom (T2)(b).

What is the significance of the axiom (T2)(b)? Suppose W_J is the subgroup of the Coxeter group generated by the involutions $\{r_i | I \in J\}$. Let N_J be the preimage of this group in N — that is $W_J = N_J/B \cap N$. The significance is manifest in the following.

Lemma 9.10.1 *BN_JB is a subgroup of G. In particular $G = BNB$. Conversely, if H is any subgroup containing B, and $B \neq 1$, then $H = BN_JB$ for some subset J of the set I indexing the involutions S of W.*

The first statement is an elementary consequence of the relation given in the second part of axiom (T2). A simple proof of the second statement is given in Ronan's book [103], pp 59–60.

These subgroups containing B are called *parabolic subgroups*. Note that when $|S| = |I| = k$ is finite, there are exactly 2^k parabolic subgroups.

Since each parabolic subgroup is determined by a subset J of I, we write $P_J := BN_JB$ and write P_s for $P_{\{i\}}$, when i indexes involution $s \in S$. The poset of parabolic subgroups is easy to describe, for the following reason.

Lemma 9.10.2 *$P_J \cap P_K = P_{(J \cap K)}$ for any subsets J and K of I.*

Of course B is at the bottom on this poset, and the subgroups $P_s = B \cup BsB$, where $s \in S$, are the *minimal parabolic subgroups* (actually minimal among those properly containing B).

Theorem 9.10.3 *The chamber system $C(G, B; \{P_s | s \in S\})$ is a building.*

Proof This is a direct consequence of the fact that the N-orbit BN is an apartment whose conjugates under G form a system of apartments. (There are other proofs. For example, the W-valued metric ∂ can be directly defined from the (B, N)-pair axioms.) □

Theorem 9.10.4 *Suppose C is a building of type M and let $W := W(M)$. Assume the following:*

(B1) *The automorphism group of C transitively permutes chambers, and the stabilizer B of a chamber c is transitive on all chambers d for which $\delta(c, d)$ has a fixed value $r \in W$.*
(B2) *The stabilizer N of an apartment A induces the full Coxeter group $W(M)$ on that apartment.*

Then $G := Aut(C)$ forms a Tits system (G, B, N, S). (S, of course is the canonical set of generating involutions for $W(M)$.)

9.10.3 Sphericity

Recall that a *spherical building* is simply a chamber system of type M with these properties:

1. All corank one residues are strongly gated (among other equivalent conditions).
2. The matrix M produces a finite Coxeter group $W(M)$.

We have also seen that when $W(M)$ is finite, the connected diagrams are as given in Fig. 9.11. Let us suppose from here on in that we have a building whose type matrix M defines a spherical irreducible Coxeter system.

What happens next is a quite remarkable discovery of J. Tits – a transformation of a purely geometric property into a group-theoretic one.

Let C be the chamber system of a spherical building of type M. Clearly the Weyl group is finitely generated and so apartments exist. So the covering properties of apartments show us that the diameter of this chamber system (as a graph) is bounded by the diameter of an apartment. (Apartments cover all distance-paired vertices, and are isometrically embedded.) That means that for every chamber c, there exists a chamber c' which lives as far as possible from c – that is, as opposite chambers of a finite apartment. The opposite relation is quite strong. If c is opposite c' and d is adjacent to c, then there is a unique vertex d' adjacent to c' which is opposite d. If the rank is at least three, it even works for triangles. So that means that there is a bijection from the neighbors of c to the neighbors of c' and it is not difficult to see that this mapping preserves the edge labels.

So one has an isomorphism $N(c) \to N(c')$ of the neighborhood of c to that of its opposite vertex c'. But of course c' possesses many neighbors c'' which are also opposite c (in fact, except for a small rank pathology, all but one of its neighbors has this property). That means we can locally map a neighborhood $N(c')$ to $N(c'')$

when c' is adjacent to c''. It is easy to see that this consistent neighborhood mapping extends to triangles, and so the thing extends globally to automorphisms of C.[30]

I hope I have led the reader to believe that any spherical building of rank at least three automatically possesses a rich group of automorphisms. In fact it is so rich that the two conditions of Theorem 9.10.4 hold, and a (B, N)-pair ensues. Thus the following arises.

Theorem 9.10.5 *Suppose C is a chamber system of spherical type M with thick panels. Call its automorphism group G.*

1. *The stabilizer B of a chamber c and the stabilizer N of an apartment, form a Tits system (G, B, N, S).*
2. *If $B \neq 1$, the diagram $D(M)$ belongs to one of the diagrams listed in Fig. 9.11, excluding H_3 and H_4 and $I(k)$ for $k = 5, 7$, or $k \geq 8$.*
3. *Let $n = |S|$ be the finite rank of the Coxeter matrix M. Then there are exactly 2^n parabolic subgroups P containing B. In particular there is a system of maximal parabolic subgroups $P_1, \ldots, P_n$ where each $P_i = P_{J-\{i\}}$.*
4. *In the chamber system $C = C(G, B; M_1, \ldots, M_n)$, where the M_i are the minimal parabolic subgroups, the double cosets BgM_i and BgP_i are the residues of type i (and rank one), and cotype i (and corank one), respectively.*
5. *The chamber system C is residually connected and so is derived as $\mathbf{C}(\mathbf{\Gamma}(C))$ from the geometry $\Gamma := \mathbf{\Gamma}(C)$.*
6. *The building geometry Γ is the coset geometry: $\Gamma = \Gamma(G; P_1, \ldots, P_n)$, where the P_i are the maximal parabolic subgroups, as above.*

A Final Comment

It is always good to understand what a theory does. As we shall see in the next chapters of this book, the theory of buildings allows point-line characterizations on a level far beyond projective spaces and polar spaces. But local characterizations suffer the occupational hazard that any global universal object can be folded up a little bit by a homomorphism that does not distort things locally. That means one must be prepared to accept the conclusion "X is a homomorphic image of a universal object." If the universal objects have been classified, at least that is the best one can expect for a classification.

But Tits' theory of buildings actually introduced far more. If we start over again, one sees that Coxeter's classification of the finite "Coxeter groups" was really a

[30] Some authors (for example Ziechang [150]) call this the "reduction theorem" because a geometric problem is reduced to a group-theoretic one. One might see it in the other direction: a rare priveledge to "go upstairs" to group theory. In fact Professor Ziechang has shown that this phenomenon of producing groups from opposite local isomorphisms has a general life in association schemes from which viewpoint the building arguments are "ad hoc" – that is, they depend on enough special properties of buildings that they are not directly transportable to the world of association schemes. More general arguments must be supplied – and these are described in his great book.

way of saying that certain generators and relations actually determinine the group. Another way of viewing Tits' accomplishment is that it provided an extension of the generator-relations of Coxeter to many classical groups — after all, the (B,N)-pair axioms are nothing less than generator-relation specifications. This discovery received its codification for finite group theorists in the famous Curtis–Tits theorem – a key part of the classification of finite simple groups (see the review of Ron Solomon in a recent issue of the *Bulletin of the American Math Soc* [120]). And then, finally, there is the still unexplored logical reason why these axioms work, and here we enter the realm of mathematical logic.

Aside from Tits' classification of the spherical buildings, few theorems of the twentieth century gave so much focus to so many different fields of mathematics.

9.11 Exercises for Chapter 9

9.11.1 Exercises on Chamber Systems

9.1 Prove both parts of Lemma 9.2.1. (The second part utilizes the fact that any corank one residue of C_K, when viewed as collection of chambers of C, is a corank one residue of C with cotype in K.)

9.2 Suppose C is a set. Let $\{\pi_i | i \in I\}$ be a collection of partitions of C indexed by I. The reader is asked to note the slight ambiguity involved in the previous sentence. It could simply mean (1) that there is a mapping $t : I \to \Pi(C)$, where $\Pi(C)$ denotes the lattice of all partitions of C. (2) On the other hand, we may view the collection $P \subseteq \Pi(C)$ of partitions as already existing, so that the indexing is a bijection $\iota : P \to I$. Under this view the mapping $(\iota)^{-1} := t : I \to \Pi(C)$ given in (1) is injective. So, when we say "let $\{\pi_i | i \in I\}$ be a collection of partitions of C indexed by I" this may allow π_i and π_j to be the same partition even if $i \neq j$ (interpretation (1)) or it might means that π_i and π_j are assumed distinct when i and j are distinct (interpretation (2)).

Under either interpretation, given $\{\pi_i | i \in I\}$, say that two chambers c_1 and c_2 of C form an *edge* if and only if there exists at least one partition π_i in the collection for which c_1 and c_2 belong to the same component of π_i – that is π_i does not separate the chambers. In that case let $\lambda(c_1, c_2)$ be the set of all indices i for which π_i does not separate c_1 and c_2. Let E be the set of all edges, so that λ is now a mapping $E \to 2^C - \emptyset$, taking the edges to non-empty subsets of I.

1. Under interpetation (1), show that (C, E, λ) is an abstract chamber system over I as defined on p. 294. [Check the triangle axiom.]
2. Given an abstract chamber system over I as defined on p. 294, the relation of being equal or i-adjacent is an equivalence relation $\sim_i$, as already remarked. Let π_i be the partition of all chambers C into the $\sim_i$-equivalence classes (or panels of type i). Then $\{\pi_i | i \in I\}$ is a collection of partitions of C. Show by example that it is possible for two relations $\sim_i$ and $\sim_j$ to be the

same equivalence relation, so that $\pi_i = \pi_j$. Conclude that *an equivalent definition of an abstract chamber system is the hypothesis of a collection of partitions of a set C indexed by I, under interpretation (1).*

3. Conclude that a collection of partitions indexed by I under interpretation (2) is an abstract chamber system with this property: (sep) *Given any two distinct types i and j there exists a pair of distinct chambers* (x, y) *such that either (a) x is i-adjacent but not j-adjacent to y or (b) x is j-adjacent but not i-adjacent to y.*

9.3 Let $A = \langle \sigma, \tau \rangle$, a subgroup of the automorphism group of the chamber system C of Example 5. Show that the canonical homomorphic image C/A is isomorphic to the chamber system of Example 3.

9.4 Display the tripartite 18-vertex graph of $\mathbf{\Gamma}(C)$ where C is the "shaved cube" of Example 4.

9.5 Let $C = (G, B; \mathcal{H})$ be a coset chamber system. As usual, $\mathcal{H}$ is a collection of subgroups of G containing B whose (not necessarily distinct) members are indexed by I. In practice, one often assumes for coset chamber systems that $\mathcal{H}$ is an antichain in the poset of subgroups of G containing B. This condition is not really necessary. It is only there to avoid having i-adjacency imply j-adjacency, for distinct i and j – that is, having partition π_i refine partition π_j in the formulation of Exercise 1. The condition implies (sep). One still gets an abstract chamber system without this condition on $\mathcal{H}$.

1. Now choose any subset J of the set I indexing the subgroups in $\mathcal{H}$ and let H_J be the subgroup of G generated by the subgroups $\{H_j | j \in J\}$. Show that any residue of C of type J is a coset $H_J g$ (regarded as a collection of right cosets of B so that it is a subset of $C = G/B$, the cosets of B). [One must show that this chamber subsystem is connected.]
2. Show that in the coset chamber system C, it is possible for two chambers to be both i- and j-adjacent for distinct i and j. What does this imply about H_i and H_j?
3. Show right multiplication of all cosets by an element $g \in G$ induces an automorphism of C. Thus we have a morphism $G \to \mathrm{Aut}(C)$. Show that this morphism is injective if and only if B contains no non-identity normal subgroups of G.
4. Let $f : G \to K$ be a surjective homomorphism of groups. Show that f induces a morphism of chamber systems

$$c(f) : (G, B; \{H \in \mathcal{H}\}) \to (K, f(B); \{f(H) | H \in \mathcal{H}\}).$$

[Hint: The mapping $c(f)$ takes coset Bg to $f(Bg) = f(B)f(g)$. Show that if $Bg \cup Bh \subseteq H_i h \in \mathcal{H}$ then $f(gh^{-1}) \in f(B)$ (equivalently,

$f(B)f(g) = f(B)f(h))$ or $f(B)f(g) \cup f(B)f(h) \subseteq f(H_i)f(h)$ – that is, $c(f)$ preserves i-adjacency on edges whose image is an edge.]

5. Suppose M is a subgroup of G. Show that there is a chamber-injective chamber system morphism

$$\sigma_M : (M, M \cap B; \{H_i \cap M | H_i \in \mathcal{H}\}) \to (G, B; \mathcal{H}).$$

[Let σ_M send $(M \cap B)m$ to Bm for each $m \in M$. Show that i-adjacency is preserved by σ_M on edges whose images are edges.]

6. As in part 1 of this exercise, let $H_J = \langle H_j | j \in J \rangle$. Suppose K is a proper subset of I and set $J = I - K$. We can form a chamber system over K whose chambers are the cosets G/H_J. Two distinct cosets $H_J g$ and $H_J h$ are k-adjacent, for $k \in K$, if and only if $H_J g \cup H_J h \subseteq H_{J \cup \{k\}} g$. Show that this chamber system is the coset chamber system $(G, H_J; \{H_{J \cup \{k\}} | k \in K\})$. Show that it is the truncation C_K of type K of the chamber system $C = (G, B; \{H_i | i \in I\})$.

9.11.2 Exercises on Residual Connectedness

9.6 Show that the chamber system C of Example 1 is not residually connected, while its associated geometry $\mathbf{\Gamma}(C)$ is residually connected.

9.7 Show that the chamber systems of Examples 2, 4, and 5 are residually connected, while the chamber system of Example 3 is not residually connected.

9.8 Prove Lemma 9.2.3.

9.9 Formalize the argument for Corollary 9.2.8.

9.11.3 A Few Exercises on Gatedness

9.10 Prove Lemma 9.3.1. Also prove Lemma 9.4.4.

9.11 A subgraph (X, E') of a graph (V, E) is said to be *gated* in (V, E) if, for every vertex $v \in V$, there is a vertex $g(v) \in X$ such that for every vertex $x \in X$,

$$d(v, x) = d(v, g(v)) + d(g(v), x).$$

(Note that, unlike strong gatedness, the second summand on the right involved global distance, not the distance in the subgraph (X, E'). This concept was first studied by A. Dress and R. Scharlau [59].) Show that every residue of

the rank three chamber system of the following example is gated in the full chamber system.

Example 10 Let m be an odd integer greater than two. Consider this vertex set $V := \{x_0, \ldots, x_{2m-1}\}$. Writing $x_{2m} = x_0$ and reading subscripts modulo $2m$ where necessary, we define (x_{2i}, x_{2i+1}) to be an edge labelled "1"; we define (x_{2i-1}, x_{2i}) to be an edge labelled "2"; and we define $(x_i, x_{i+m}$ to be an edge labelled "3" – for all i. Then clearly $(x_0, x_1, \ldots, x_{2m-1}, x_{2m} = x_0)$ is a (pointed) circular gallery of type $1212\cdots2$ of length $2m$. The antipodal pairs of vertices of this circuit form the panels of type "3." Since m is odd, all rank two residues involve all of the chambers.

9.11.4 Exercises on Generalized Polygons

9.12 Prove Theorem 9.4.12.

9.13 Consider the geometry $(\mathcal{P}, \mathcal{L})$ where

$$\mathcal{P} := \{p, q, a_i, b_i | i = 1, \ldots, t+1\},$$

and

$$\mathcal{L} = \{(p, a_i), (q, b_i), (a_i, b_i) | i = 1, \ldots, t+1\},$$

a collection of thin lines. Show that $(\mathcal{P}, \mathcal{L})$ is a generalized hexagon.

9.14 Suppose $\mathcal{P} = \{p_i, a_i, b_i, q_i | i = 1, \ldots, s+1\}$. Set $P := \{p_1, \ldots p_{s+1}\}$ and $Q := \{q_1, \ldots, q_{s+1}\}$. Let L be the two thick lines P and Q together with the collection of thin lines

$$\{(p_i, a_i), (q_i, b_i), (a_i, b_i) | i = 1, \ldots, s+1\}.$$

1. Show that $(\mathcal{P}, \mathcal{L})$ is a generalized octagon with each point on just two lines.
2. Show that $((\mathcal{P}, \mathcal{L}))$ is $2\Gamma^*$ where Γ^* is the 2-by-3 grid.

9.15 The Feit–Higman theorem shows that the only generalized 5-gon of order (s, t) is the ordinary pentagon where $s = t = 1$.

1. If, hypothetically, there were a generalized 5-gon of order (s, s), $s > 1$, show that there are three relations between a point and a line and describe these relations.
2. The student might speculate upon whether there could be a generalized 5-gon with infinitely many points on each line and infinitely many lines on each point.

9.11.5 An Interplay of Examples and Exercises on Diagram Geometries

9.16 Show that any geometry Γ belonging to the finite rank diagram below (Fig. 9.17) is a linear space when truncated to $(\mathcal{P}, \mathcal{L})$ (the first two nodes from the left), and that each remaining node is a subspace with respect this linear space. [Hint: One could first do this for rank three in order to truncate an $L - L$-diagram to a single L and use induction. On the other hand it can also be done directly.]

9.17 Using the previous exercise, show that any geometry belonging to the diagram below (Fig. 9.18) is a projective space when truncated to types $(\mathcal{P}, \mathcal{L})$ represented by the first and second nodes from the left.

This a good time to introduce a few warnings. The main one is not to expect too much from a diagram. The tendency is to use a nice simple diagram, such as A_n (Fig. 9.19) as a means of getting the concept across. But actually this example is quite atypical. For example, if we designate the leftmost two nodes as points and lines, respectively, we obtain the standard projective space $(\mathcal{P}, \mathcal{L})$. It is then a fact that all objects of other types are *subspaces* when realized by their point-line shadows (i.e., their residues truncated to points and lines). The same occurs when we assign some internal node (say the one labeled k) to be the points, and let the flags whose type is represented by adjacent nodes (in this case those of type $(k-1, k+1)$) be the lines. This separation of the point-node from the rest of the diagram by the type of the flags destined to be lines gives us the desirable axiom as follows.

(A) *Any object not a point or line, which is incident with a line L, is in fact incident with every point that is incident with that line L.*

We can hardly hope objects to be subspaces without such an axiom. Indeed, something like that works for diagrams A_n, D_n, E_6, E_7, and E_8. But these are not typical. Consider the following highly pathological example.

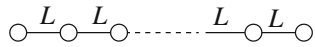

Fig. 9.17 A string diagram with linear-space residues

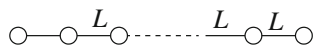

Fig. 9.18 Another diagram with linear-space residues

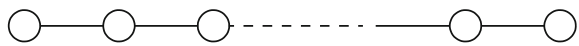

Fig. 9.19 The A_n diagram

Example 11 (The flat Neumaier geometry) Consider first the collection $\mathcal{L}$ of all 35 3-subsets of a 7-set $\mathcal{P} = \{0, \ldots, 6\}$. This is acted on by the group $G = \mathrm{Sym}(7)$. Now consider the Fano plane $F = (\mathcal{P}, \mathcal{L}_F)$ whose seven lines are the set $\mathcal{L}_F$ of additive translates mod 7 of the line $\{1, 2, 4\}$. As we know, its stabilizer in G is the simple group H of order 168 and index 30 in G. G acts imprimitively on the cosets G/H in two blocks which are orbits of length 15 under the alternating group $A := \mathrm{Alt}(7)$. Let π_0 and π_1 be these two orbits. If F and F' are planes in different orbits, then these two planes share three lines or no lines at all. On the other hand, if F and F' are planes within a common A-orbit, they always share exactly one line.

Now form the geometry $N := (\mathcal{P}, \mathcal{L}, \pi_0)$, of 7 points, 35 lines, and 15 planes from one class. Incidence has been defined by the definition of π_0. Let us examine the residue of a point – say, the point $p = 0$. The remaining points form a 6-set $\Omega = \{1, \ldots, 6\}$. The lines incident with p are thus identified with the 15 2-subsets of Ω. Any plane of π_0 produces three lines on p which induce a partition of Ω into three 2-sets. There are 15 such partitions. On, the other hand, any two distinct planes of $\mathcal{L}_0$ share just 1 line, and so these 15 planes each induce a different partition of Ω into 2-subsets. So there is a bijection between $2 - 2 - 2$-partitions of Ω and planes on p.

This is enough information to see that the residue of a point in N is faithfully modelled by the rank-two geometry of 2-subsets and $2 - 2 - 2$-partitions of a 6-set Ω which, as we know, is the generalized quadrangle of order $(2, 2)$. Thus N belongs to the diagram below (Fig. 9.20) which we call the C_3 diagram.

The message here is that the 15 planes of this geometry are *not* subspaces of the truncation to points and lines.

9.18 The reader might consider what happens with the truncation $(\mathcal{P}', \mathcal{L}') = (\pi_0, \mathcal{L})$ to planes and lines. Here we have 15 "points", 35 "lines," and 7 "quadrangles." Are the quadrangles subspaces of $(\mathcal{P}', \mathcal{L}')$?

9.19 Show that the rank three chamber system of Example 10 (p. 394) belongs to a diagram consisting of three nodes with the edge connecting any two of them labelled by "(m)", indicating a generalized m-gon.

Example 12 Let $K_i, i = 1, 2, 3$, be three disjoint copies of the complete tripartite graph $K_{2,2,2}$, each regarded as a geometry over the typeset $\{2, 3, 4\}$. The geometry Γ of this example is over the type-set $\{1, 2, 3, 4\}$. Its truncation to $\{2, 3, 4\}$ is the union $K_1 \cup K_2 \cup K_3$ – a disconnected graph. The set of objects of type 1 is $\{x_1, x_2, x_3\}$. If $\{i, j, k\} = \{1, 2, 3\}$, the object x_i is incident with every object of $K_j \cup K_k$. There are no further incidences.

Fig. 9.20 The C_3 diagram

9.20 In the geometry Γ of Example 12, show the following:

1. Every flag of Γ lies in a chamber flag.
2. Every residue of rank two of Γ is connected.
3. Γ is not residually connected.

9.11.6 Exercises Concerning Chamber Systems of Type M

9.21 Show that in the chamber system $C = C(G, 1 : \langle s_1 \rangle, \ldots, \langle s_n \rangle)$ – which we have been writing as $C(G, 1; S)$ – a gallery is of reduced type if and only if it is a *minimal* gallery – i.e., a shortest gallery (geodesic path) connecting its initial and terminal chambers. [This is not true of general chamber systems, and occurs here because G is a presented group.]

9.22

1. Show that in C, any circuit has even length. [Hint: Any circuit can be deformed to the trivial circuit by some chain of expansions, contractions, and elementary $\mathcal{C}_2$-homotopies, which do not disturb the length parity.]
2. Conclude that C is a bipartite graph.
3. Show that for any bipartite graph and edge $e = (x, y)$, for any vertex v, $d(v, x) = d(v, y) \pm 1$. In particular this must hold for C. So the edge e determines a partition of the vertices of C into two sets: $D^-(e) = \{v | d(v, x) < d(v, y)\}$ and $D^+(e) = \{v | d(v, x) > d(v, y)\}$. These two sets are called *opposite roots* and exist for any bipartite graph.
4. Observe that the partition $C = D^-(e) + D^+(e)$ does not in general match the partition $C = C_1 + C_2$ into two cocliques which defines the bipartness of C.
5. Suppose $e = (x, y)$ is an edge in a bipartite graph C as in the previous two parts of this exercise. Give an example of a bipartite graph C and a"bridging edge" $f = (u, v)$ where $(u, v) \in D^-(e) \times D^+(e)$ for which the partition $D^-(e) + D^+(e)$ does not coincide with the partition $D^-(f) + D^+(f)$. [Compare this with Theorem 9.6.10 which is for bipartite graphs arising from Coxeter chamber systems. In your example, make sure all circuits have even length. This can be done with five vertices.]

Chapter 10
2-Covers of Chamber Systems

Abstract A full account of J. Tits' "local approach theorem" is given, using the homotopy-theoretic concepts of Chaps. 1 and 9. An additional lemma on covers of truncations due to A. Kasikova is needed for the succeeding chapter.

10.1 Introduction

The study of geometries or chamber systems with a specified local structure naturally leads one to investigate covering mappings for which that local structure is preserved. In particular, when the structure of rank two residues is to be preserved, one is interested in 2-coverings. In the case of chamber systems, the 2-coverings are the $\mathcal{C}_2$-coverings (in the sense of Chap. 1 of Part I) of the underlying graph of the chamber system, where $\mathcal{C}_2$ is the collection of all circuits which are manifest in any residue of rank two.

This chapter contains two basic results on 2-covers of chamber systems. Both of them play a crucial role in the theory of locally truncated diagram geometries to be expounded in the succeeding Chap. 11.

The first result is Kasikova's lemma which concerns a condition under which a 2-covering $\hat{C} \to C$ of a chamber system C having a residually connected truncation C_J of type J induces an injective mapping on the truncation to the same type of a residue of the associated geometry $\mathbf{\Gamma}(C)$.

The second main result of this chapter is the celebrated "Local Approach Theorem" of J. Tits. This theorem asserts the following:

> *Suppose C is a chamber system of type M each of whose residues of rank three have a 2-cover which is a building. Then the universal 2-cover of C is a building.*[1]

Of course, a universal 2-cover (being a universal $\mathcal{C}_2$-cover of an edge-labelled graph with the labels lifted along with the edges) always exists, as shown in the last

[1] Note the linguistic difference between a "covering" and a "cover." The convention is this: if $\hat{C} \to C$ is a $\mathcal{C}$-covering, then the domain $\hat{C}$ is a $\mathcal{C}$-cover.

E.E. Shult, *Points and Lines*, Universitext, DOI 10.1007/978-3-642-15627-4_10,
© Springer-Verlag Berlin Heidelberg 2011

part of Chap. 1 of Part I. Buildings were defined in a host of ways in the previous Chap. 9.

The proof of Tits' Theorem weaves together two kinds of homotopy theories and the reader should be prepared at the outset to handle both of them. They are:

1. The general $\mathcal{C}$-homotopy expounded in Chap. 1. It appears here in two roles:
   (a) To define the 2-coverings involved in both the hypothesis and conclusion of Tits' Local Approach Theorem
   (b) To prove that a certain graph, whose vertices are galleries of reduced type with adjacency defined by elementary M-homotopies, is simply $\mathcal{A}$-connected with respect to a certain collection $\mathcal{A}$ of circuits of elementary M-homotopies.
2. M-homotopy, where galleries can only submit to a series of elementary deformations, where a segment that is a geodesic with respect to a rank two residue is replaced by another geodesic of the rank two residue with the same initial and terminal chambers.[2]

Fortunately any reader of the preceding Chap. 9 has already experienced this duplicity in homotopies. Aside from the $\mathcal{C}$-homotopy of Chap. 1 and the C_2-homotopy pervading Chap. 9, we met there something called Λ-homotopy (see p. 317), which is exactly M-homotopy when applied to chamber systems of type M.

In Tits' original paper which introduced them, M-homotopies were called "special homotopies." Of course in many sources an elementary M-homotopy (or elementary "special homotopy") is instead defined to be a replacement of a segment of a gallery – one of reduced type in a rank two residue – by another connecting the same chambers. The definition is exactly the same as that given above except that the property of the gallery-segment being a "geodesic in a rank two residue" is now replaced by the propery of it being a "segment of reduced type " in that residue. But here – because we are dealing with a chamber system of type M whose rank two residues are generalized polygons – a gallery is a geodesic in one of these residues if and only if it has reduced type. Thus for chamber systems of type M, it does not matter whether one writes "reduced type" or "geodesic" in the definition.

The proof proceeds in steps.

1. First we must show that any building possesses a property (Q_c) (Sect. 10.3).
2. Then we must establish a criterion for recognizing that a type M chamber system is 2-covered by a building. The relevant criterion for this property is a condition denoted (R_c). To prove that a 2-covering image of a building has (R_c), we need to know that buildings enjoy (Q_c). On the other hand, in order to prove that the 2-cover of a chamber system of type M having property (R_c) is a building, one requires an entirely new model for the 2-cover of a type M chamber system – one

[2] Recall that the rank two residues are not assumed to be isometrically embedded subgraphs of the graph of the ambient chamber system C. So a geodesic with respect to the rank two residue may very well not be a geodesic of C.

built on M-homotopy classes of galleries of reduced type beginning at chamber c rather than the usual $\mathcal{C}_2$-homotopy classes of arbitrary galleries beginning at c. This is not easy. All of this is proved in Sect. 10.4.

3. Next, Sect. 10.5 presents an excursion showing that a connected component of the graph depicting an M-homotopy class with elementary M-homotopies as the adjacencies is simply $\mathcal{A}$-connected where $\mathcal{A}$ is a family of circuits in this graph representing three distinct ways to produce a sequence of elementary M-homotopies taking a gallery of the Coxeter chamber system of type M to itself. It is proved using Tits' criterion for simple $\mathcal{A}$-connectedness from Chap. 1.
4. Finally, Sect. 10.6 assembles the proof of Tits' theorem. It is enough to show that a chamber system enjoying the hypotheses of Tits' theorem possesses property (R_c). The condition asserts that M-homotopic galleries of the same reduced type are equal. Since we are dealing with M-homotopy of galleries of reduced type, one only needs to check that one recovers the same gallery after travelling through any of the three sorts of circuits of $\mathcal{A}$.

10.2 Kasikova's Lemma on 2-Coverings of Chamber Systems

This section contains a lemma which concludes that, under certain conditions, the morphism of geometries induced by a 2-covering of chamber systems is injective when restricted to an appropriate residue. In the theory of locally truncated geometries, the goal is to show that under suitable conditions, a locally truncated geometry is the truncation of a homomorphic image of a building. At that point Kasikova's lemma has an important role to play in showing that such a morphism induces a $\mathcal{T}$-covering of the point-collinearity graphs of the geometries where $\mathcal{T}$ is the class of 3-cycles. The lemma itself, however, makes no reference to locally truncated geometries. It involves only 2-coverings of chambers systems, residual connectedness of a truncation, and a conclusion about residues of the associated geometries.

Lemma 10.2.1 *(Kasikova's lemma) Suppose $f : \hat{C} \to C$ is a 2-covering of chamber systems over I. Suppose $\bar{X}$ is a residue of cotype K in C, and X is any lift of $\bar{X}$ to a cotype K residue of $\hat{C}$. Let $h = \Gamma f$ be the (functorially induced) morphism of geometries over I:*

$$h : \Delta := \mathbf{\Gamma}(\hat{C}) \to \mathbf{\Gamma}(C) := \bar{\Delta}.$$

Then $\bar{X}$ may be regarded as a flag of type K of the geometry $\bar{\Delta}$ and X may similarly be regarded as a flag X of the geometry Δ, where $h(X) = \bar{X}$.

Now suppose the following:

1. *The residue $\bar{X}$ is a 2-simply connected chamber system over $I - K$.*
2. *For some subset J of I properly containing K, the truncation C_J is a residually connected chamber system over J.*

Let Δ_J and $\bar{\Delta}_J$ be the truncations to type J of the respective geometries Δ and $\bar{\Delta}$. Then the mapping of geometric residues

$$h_X : Res_{\Delta_J}(X) \to Res_{\bar{\Delta}_J}(\bar{X}),$$

induced by the restriction of h is injective.

Proof Suppose, by way of contradiction, that y_1 and y_2 are distinct objects of type $i \in J - K$ in the geometry Δ which are incident with X while $h(y_1) = h(y_2)$. Then, by the definition of Δ, each y_k can be regarded as a residue Y_k of chamber system $\hat{C}$, $k = 1, 2$, each of cotype i such that:

1. $Y_1 \cap Y_2 = \emptyset$
2. $Y_1 \cap X \neq \emptyset \neq Y_2 \cap X$
3. $f(Y_1) = f(Y_2)$ is a residue of cotype i of $f(\hat{C}) = C$

By assumption (1.) $\bar{X}$ is 2-simply connected and so as f is a 2-cover, the restriction of f to X induces an isomorphism $X \to \bar{X}$ as chamber systems over $I - K$. Thus in chamber system C, we have residues $\bar{Y} := f(Y_1) = f(Y_2)$ and $\bar{X}$ of respective cotypes i and K in C, whose intersection $\bar{Y} \cap \bar{X}$ consists of at least two connected components (see Fig. 10.1). Since all of the "components of the Venn diagram," $C - (\bar{X} \cup \bar{Y})$, $\bar{X} - \bar{Y}$, $\bar{Y} - \bar{X}$, and $\bar{Y} \cap \bar{X}$ are each a union of residues of C of type $I - J$, with $\bar{X}$ and $\bar{Y}$ residues of cotypes K and i of the truncation $\mathcal{C}_J$, we must infer that C_J is not residually connected, contrary to hypothesis.

Thus the lemma holds. □

The hypothesis of the lemma which requires that a truncation be residually connected often happens in real life. In particular this happens when the truncation is a truncation of rank at least 2 of a residually connected geometry. In certain cases it is then useful to recall the following from Chap. 9.

Theorem 10.2.2 *A building of finite rank is always residually connected.*

Proof By definition, a building is a chamber system of type M which satisfies a number of equivalent conditions listed on p. 377. One of these is the condition (RG), asserting that all residues are strongly gated. The type M hypothesis also implies condition (typ) that any two chambers are i-adjacent for at most one label "i." Now Corollary 9.3.8 says that any chamber system of finite rank (whether type M or not) that satisfies (RG) and (typ) is residually connected. Thus every building of finite rank is residually connected.

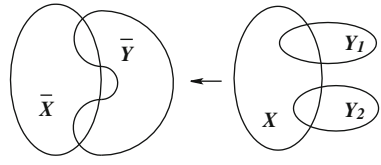

Fig. 10.1 The homomorphism of residues in Kasikova's lemma

10.3 An Important Property of Buildings

For the class of chamber systems of type M, a weaker condition than (P_c) or (G_c) is as follows.

(Q_c) *Between chamber c and any further chamber t there exists at most one gallery of any fixed reduced type w.*

Lemma 10.3.1 *Condition (G_c) implies (Q_c).*

Proof Let G and H be galleries of the same reduced type w beginning at c and ending at t. If $w = i$, $G = H$ as they are the same i-adjacency (c, t). We can prove $G = H$ by induction on the length of gallery G. We set

$$G = (c, c_1, \ldots, c_d = t) \text{ and } H = (c, b_1, \ldots, b_d = t),$$

and since both are type w, both c_{d-1} and b_{d-1} are j-adjacent to t. If $c_{d-1} = b_{d-1}$, we may proceed by induction on the length of G to conclude that the two subgalleries G' of G and H' of H stretching from c to c_{d-1} are identical, whence $G = H$. But if $c_{d-1} \neq b_{d-1}$, the gallery $(c, c_1, \ldots, c_{d-1}, b_{d-1})$ is of reduced type w, but connects two chambers known to be at mutual distance $d - 1$ by (G_c) applied to G. This contradicts (G_c). □

10.4 When Are Universal 2-Covers Buildings?

The answer to the question in the title of this section is very roughly: "most of the time, for chamber systems of type M." For a chamber system of type M the critical property is as follows.

(R_c) *If two minimal galleries with the same reduced type originating at c are M-homotopic, then they are the same gallery.*

Remark Note that the hypothesis of being M-homotopic already entails the statement that the two galleries end at the same chamber.

Theorem 10.4.1 *The universal 2-cover of a chamber system C of type M is a building if and only if C satisfies condition (R_c) for some chamber c.*

Theorem 10.4.1 will follow from Lemma 10.4.2 and Theorem 10.4.7 below. The proof of the latter theorem requires a series of steps. First we record the following.

Lemma 10.4.2 *If the universal 2-cover of chamber system C is a building, then C satisfies condition (R_c).*

Proof Let $\kappa : \hat{C} \to C$ be the universal 2-covering of C. Let G and H be homotopic galleries of the same reduced type beginning at chamber c and ending at chamber t.

Select a preimage $\hat{c} \in \kappa^{-1}(c)$. Since κ is a fibering, the galleries G and H lift to galleries $\hat{G}$ and $\hat{H}$ of the same types in $\hat{C}$ beginning at $\hat{c}$. But since G and H are M-homotopic, and κ is an isomorphism when restricted to any rank two residue, $\hat{G}$ and $\hat{H}$ are also M-homotopic galleries. But by hypothesis, $\hat{C}$ is a building, and so, by Lemma 10.3.1 it satisfies the condition $(Q_{\hat{c}})$. So, as G and H, and hence $\hat{G}$ and $\hat{H}$, have the same type, $(Q_{\hat{c}})$ forces $\hat{G} = \hat{H}$. Applying κ, we have $G = H$. Thus (R_c) holds for every chamber c. □

Our next goal is to prove the converse of Lemma 10.4.2: the assertion that condition (R_c) for C implies that $\hat{C}$ is a building. Let us set the scene.

Suppose C is a chamber system of type M. Let $\hat{C}$ denote the universal 2-cover of C with covering morphism $\kappa : \hat{C} \to C$. Fix a chamber $\hat{c}$ in the fiber $(\kappa)^{-1}(c)$. For each gallery G of C beginning at c, there is a unique lift $\lambda(G)$ of G starting at $\hat{c}$. Since $\kappa : \hat{C} \to C$ is a 2-covering, we see:

1. *If $\hat{G}, \hat{G}_1, \ldots, \hat{G}_m = \hat{H}$ is an M-homotopy of galleries in $\hat{C}$, then so is its image under $\hat{\kappa}$.*
2. *Conversely, if $G, G_1, \ldots, G_m = H$ is an M-homotopy in C, then the sequence of lifts*

$$\lambda(G), \lambda(G_1), \ldots, \lambda(G_m) = \lambda(H)$$

is an M-homotopy of $\hat{C}$.

For any gallery G in C, we let $\tau(G)$ denote the *type* of G, a word in $M(I)$.

Let W and $C(W)$ be the Coxeter group and Coxeter chamber system defined by the Coxeter matrix M. We know from Chap. 9 that each equivalence class of words contains within it a unique M-homotopy class of reduced words. In this way we may regard the chambers of $C(W)$ as M-homotopy classes of reduced words.

In Chap. 1, a model for the universal 2-cover $\hat{C}$ was constructed; its chambers were the $\mathcal{C}_2$-homotopy classes of galleries of C which have c for an initial chamber. The adjacency of one homotopy class to another was determined by adjacency of the terminal chambers for each class.[3]

In the presence of condition (R_c) it is possible to construct a second model of the universal 2-cover of C, one whose set of chambers is the set $\bar{C}$ of all M-homotopy classes of galleries of *reduced type* beginning at c. This new model for the universal 2-cover of C more easily exhibits the fact that it is a building than does the former model.

Now let $\bar{C}$ be the set of M-homotopy classes of galleries of C which begin at chamber c and have reduced type. There are two mappings of interest. The first is

$$\hat{\kappa} : \bar{C} \to \hat{C}$$

[3] This comes with the *proviso* that distinct homotopy classes with the same terminal chambers are not i-adjacent for any type i.

which takes any M-homotopy class $[G]$ to the terminal chamber of $\lambda(G)$. (This is well-defined, for if H is in the class $[G]$, H then by (2) above, $\lambda(H)$ is homotopic to $\lambda(G)$.)[4]

The second mapping is

$$\omega : \bar{C} \to C(W)$$

which takes $[G]$ to the M-homotopy class of reduced words

$$[\tau(G)] = \{\tau(H) | H \in [G]\},$$

regarded as a single element of W.

Set

$$\bar{\kappa} = \kappa \circ \hat{\kappa} : \bar{C} \to C.$$

(Despite the circumlocuted definition, $\bar{\kappa}([G])$ is just the common terminal chamber of all the galleries in the M-homotopy class $[G]$.)

Note that at this point, $\bar{C}$ is just a set. It has not been endowed with the structure of a chamber system. Consequently, the two mappings $\hat{\kappa}$ and $\bar{\kappa}$ are just maps from a set into the sets of chambers of some chamber systems. They are not (so far) chamber system morphisms. To make them so embarks us on the most delicate part of Tits' proof of Theorem 10.4.1.[5]

Lemma 10.4.3 *Now fix a subset J of I and let $\bar{x}$ be a chamber of $\bar{C}$.*

The M-homotopy class $\omega(\bar{x})$ contains a reduced word with the factorization

$$w' \cdot w''$$

where w' is J-reduced – that is w' is a reduced word not M-homotopic to a word ending in a letter from J – and w'' is a reduced word in the free monoid $M(J)$ (reduced with respect to both Coxeter matrices M and M_J).

Proof Since the residues of W are gated, we see that the residue of type J which contains the terminal vertex t of $\omega(\bar{x})$ is gated with respect to $1_W := \omega(\bar{c})$, with gate g. Thus we have a minimal gallery from 1_W to t that passes through g which is M-homotopic to $\omega(\bar{x})$. The type of the initial segment 1_W to g is J-reduced, and the type of the segment from g to t is a reduced word in $M(J)$. □

This factorization is not at all unique, but *for any two such factorizations, the J-reduced prefixes are homotopic words*, since they represent geodesics in W from

[4] Note that in the interpretation of $\hat{C}$ as the set of all $\mathcal{C}_2$-homotopy classes of galleries that begin at c, the mapping $\hat{\kappa}$ takes the M-homotopy class $[G]$ of a gallery of reduced type, to the full $\mathcal{C}_2$-homotopy class of G, regarded as a chamber of $\hat{C}$.

[5] Indeed, the most delicate part of his entire "Local Approach" paper [139].

1_W to the 'gate' g of the residue gW_J. (This is an application of condition (P_c) to the building W, if nothing else.) Similarly, the maximal length suffixes in W_J are also homotopic, since these represent two geodesics from the 'gate' to $[w] = t$.

Now suppose $\bar{x} = [G] = [H]$. Then since C is type M, G is homotopic to a gallery $G' \circ G''$ where G' is type w' and G is type w''. Now suppose w is M-homotopic to a second factorization $w_1' \cdot w_1''$ where w_1' is J-reduced and $w_1'' \in M(J)$, and H is any other gallery in $[G]$. Then, similarly, H is also homotopic to a gallery $H_1' \circ H_1''$ with $\tau(H_1') = w_1'$ and $\tau(H_1'') = w_1''$. Now, as observed at the end of the previous paragraph, w_1' is homotopic to w', and independently, w_1'' is homotopic to w''. Thus by the standard property of chamber systems of type M, H_1' is homotopic to a gallery H' of type w' and H_1'' is homotopic to a gallery H'' of type w''. At this stage, $G' \circ G''$ and $H' \circ H''$ are homotopic galleries of C starting at c of the same reduced type. By condition (R_c), we must have $G' \circ G'' = H' \circ H''$, so

$$G' = H' \sim H_1' \tag{10.1}$$

$$G'' = H'' \sim H_1'' \tag{10.2}$$

For any gallery H of reduced type of the form $H = H' \circ H''$ where $\tau(H')$ is J-reduced and $\tau(H'')$ is in $M(J)$, we say H *possesses a J-factorization* with an *initial J-reduced segment* $H'$and a *terminal J-segment* H''. The upshot of the previous paragraph and the Eq. (10.1) is the following.

(In_J) (A consequence of (R_c).) The initial J-reduced segments of all galleries in $[G]$ (G of reduced type) having a J-factorization form a single non-empty M-homotopy class of galleries.

Thus if $\bar{x} = [G]$, we let $\rho_J(\bar{x})$ denote this unique M-homotopy class of all initial J-reduced segments of galleries in $[G]$. Thus $\rho_J : \bar{C} \to \bar{C}$ is a mapping with these properties:

1. If $\omega(\bar{x})$ is a J-reduced homotopy class (that is, $\bar{x} = [G]$ where G is a gallery of J-reduced type), then $\rho_J(\bar{x}) = \bar{x}$.
2. $\rho_J(\bar{x})$ is a homotopy class of J-reduced type, whence $\rho_J(\rho_J(\bar{x})) = \rho_J(\bar{x})$.

Let $\bar{C}_J$ denote the set of all homotopy classes of galleries of J-reduced type – that is, the set $\rho_J(\bar{C})$. Now if $\bar{x} \in \bar{C}_J$, we set

$$\rho_J'(\bar{x}) := \rho_J^{-1}(\bar{x}) = \{\bar{y} \in \bar{C} | \rho_J(\bar{y}) = \bar{x}\}.$$

We now see the following.

Lemma 10.4.4 *If $|J|$ is at most 2, $\bar{\kappa}$ induces a bijection of $\rho'(\bar{x})$ to $R_J(x)$, the residue of type J in C that contains $x := \bar{\kappa}(\bar{x})$.*

Proof Suppose $\bar{y}_i, i = 1, 2$, are elements of $\rho'(\bar{x})$ with the same image $t = \bar{\kappa}(\bar{y}_i)$. This means that the $\bar{y}_i$ are M-homotopy classes $[X \circ Y_i]$ where X is a gallery from

c to $x = \bar{\kappa}(\bar{x})$ of J-reduced type, and the Y_i are galleries of reduced type in M(J) – that is galleries of reduced type in $R_J(x)$ – from x to t. But R_J is either a rank one residue (a clique) or a generalized polygon (by the type M hypothesis) and so Y_1 and Y_2 are either the same identical edge, or are galleries of reduced type with the same initial and terminal chambers in a generalized polygon. In either case they are homotopic galleries, forcing $\bar{y}_1 = \bar{y}_2$. Thus the induced mapping $\rho'_J(\bar{x}) \to R_J(x)$ is injective.

Also, if y is any chamber in $R_J(x)$, let Y be a gallery of reduced type in $R_J(x)$ from x to y, and form $G := X \circ Y$. If we can show that G has reduced type, we see that $\bar{y} := [G]$ lies in $\rho'_J(\bar{x})$ and that $\bar{\kappa}(\bar{y}) = y$. But $\tau(X) := w'$ is J-reduced, and $\tau(Y)$ is a reduced word in $M(J)$. That $w' \cdot w''$ is a reduced word follows from (P_c) and gatedness of residues of the Coxeter chamber system W. Thus the induced mapping $\rho'_J(\bar{x}) \to R_J(x)$ is onto. □

If i is one of the types, we write ρ_i for $\rho_{\{i\}}$. Note that if $\bar{x} = [G' \circ G'']$ with $[G'] = \rho_J(\bar{x})$, and if $\tau(G'') = i_1 i_2 \cdots i_n$ then

$$\rho_J(\bar{x}) = \rho_{i_1}(\rho_{i_2}(\cdots(\rho_{i_n}(\bar{x}))\cdots)). \tag{10.3}$$

Now as $\bar{x}$ ranges over $C_i := \rho_i(C)$, the sets $\rho'_i(\bar{x})$ form a partition $\bar{\pi}_i$ of $\bar{C}$. We thus have a chamber system defined on $\bar{C}$ by the family of partitions $\{\bar{\pi}_i | i \in I\}$. Applying (10.3) with $J = I$, we see that $\bar{C}$ has become a connected chamber system.

We then see that the induced bijections discussed in Lemma 10.4.4 are always chamber system isomorphisms. Thus $\bar{\kappa}$ is a 2-covering. But the 2-covering $\bar{\kappa}$ factors through the universal 2-covering κ as $\bar{\kappa} = \kappa \circ \hat{\kappa}$. It follows that $\hat{\kappa} : \bar{C} \to \hat{C}$ is an isomorphism of chamber systems of type M.

Thus we have the following.

Lemma 10.4.5 *Let C be a chamber system of type M which satisfies the condition (R_c) for some chamber c. With respect to the family of partitions $\{\pi_i\}_{i\in I}$, the set $\bar{C}$ of all M-homotopy classes of galleries of reduced type beginning at c is a chamber system. (Two distinct chambers $\bar{x_1}$ and $\bar{x_2}$ are i-adjacent if and only if $\rho_i(\bar{x_1}) = \rho_i(\bar{x_2})$.) The mappings $\hat{\kappa} : \bar{C} \to \hat{C}$ and $\bar{\kappa} : \bar{C} \to C$ are both 2-coverings. As a consequence of the former 2-covering, the chamber system $\bar{C}$ is isomorphic to the universal 2-cover, $\hat{C}$.*

For the last result we let $\bar{c} = [(c)]$, denote the homotopy class of the trivial gallery (c).

Lemma 10.4.6 *Suppose $\bar{G}$ is a gallery in $\bar{C}$ of reduced type beginning at $\bar{c}$, and ending at $\bar{x}$. Then its image $\bar{\kappa}(\bar{G})$ is a gallery in C (of the same reduced type) that belongs to the M-homotopy class $\bar{x}$, representing the terminal vertex of $\bar{G}$.*

Proof Let

$$\bar{G} = ([c], [G_1], [G_2], \ldots [G_n])$$

be a gallery of $\bar{C}$ from $\bar{c}$ to $\bar{x}$ of reduced type $w = i_1 \cdots i_n$. This means

$$[G_{n-1}] \in \rho_{i_n}([G_n]) \text{ and } [G_n] \in \rho'_{i_n}([G_{n-1}]) \text{ and} \tag{10.4}$$

$$\text{for } 1 \le j \le n,\ [G_{j-1}] \in \rho_{i_j}([G_j]) \text{ and } [G_j] \in \rho'_{i_j}([G_{j-1}]). \tag{10.5}$$

Now starting with 1 and working upward, we can choose

$$G_{j+1} = G_j \circ (\bar{\kappa}([G_j], \bar{\kappa}([G_{j+1}]).$$

Then

$$\bar{\kappa}(\bar{G}) = (c, \bar{\kappa}(\bar{G}_1), \ldots, \bar{\kappa}(\bar{G}_n))$$

is a representative gallery in the class $\bar{x} = [G_n]$. □

And now we have our desired result.

Theorem 10.4.7 *Let C be a chamber system of type M satisfying condition (R_c), for some chamber c. Then its universal 2-cover $\hat{C}$ is a building.*

Proof It is now easy to show that $\bar{C}$ enjoys property $(P_{\bar{c}})$.

Suppose $\bar{G}$ and $\bar{H}$ are two galleries of reduced type extended from $\bar{c}$ to $\bar{x}$. We have just seen from Lemma 10.4.6 that their image galleries, $\bar{\kappa}([G])$ and $\bar{\kappa}([H])$ are both representative galleries in the homotopy class $\bar{x}$. Thus these galleries have homotopic types. Thus $\bar{H}$ is M-homotopic to a gallery of the same type w as G. So without loss of generality we may assume $\bar{H}$ is also type w.

Now the two images $\bar{\kappa}([\bar{G}])$ and $\bar{\kappa}([\bar{H}])$ are each homotopic (since they are in $\bar{x}$) and of the same reduced type and so they must be equal by condition (R_c) of C. But since $\bar{\kappa}$ is a 2-covering, $\bar{G} = \bar{H}$ by the uniqueness of lifts at $\bar{c}$.

Thus $\bar{C}$ is a building, and so its isomorphic copy $\hat{C}$ must also be a building. This completes the hard part of the proof of Theorem 10.4.1. □

One may observe that in the easy first part of the theorem, we actually proved that if the universal 2-cover $\hat{C}$ of chamber system C is a building then C satisfies condition (R_c) at *at every chamber* c. Thus we see that Theorem 10.4.1 must hold.

10.5 The Structure of an M-Homotopy Class

All along in this chapter we have been dealing with two sorts of homotopy theories: the special (length-preserving) homotopy – which we have called M-homotopy – that has occupied this chapter on chamber systems of type M – and the $\mathcal{C}_2$-homotopy (the crude homotopy) based on a class of pointed circuits as described in Chap. 1. Most of our work in this chapter has utilized the former, the special theory, and we have invoked a version of the latter only because it defines the evaluation epimorphism $e : M(I) \to W$.

In this section, the two theories symphonically combine: we need the crude theory once again, but this time in the ironic role of helping to explicate the special theory.

We have remarked earlier on the correspondence:

$$\mu_c : M(I) \to \text{galleries of } C(W, S) \text{ beginning at chamber } c$$

This mapping preserves elementary homotopies and therefore (special) homotopies – that is elementary homotopies among words in the monoid $M(I)$ map to elementary homotopies as defined among the galleries of C beginning at c.

The purpose of this section is to elucidate the actual structure of the special M-homotopy class $[w]$ by studying the following graph $G[w]$ whose vertices are the walks w' that are M-homotopic to w, two of them being adjacent if and only if they form an elementary M-homotopy. (One supposes this could also be done when w is not reduced, but it is extra work which offers no utility to our ultimate crusade: to characterize buildings.) So we assume w is a reduced word.

The graph $G[w]$ is obviously not a tree. In contains three very important classes of circular walks:

1. (Backtracks.) $(w_1, w_2, \ldots, w_k, w_{k-1}, w_{k-2}, \ldots, w_1)$ where each $(w_i, w_{i\pm1})$ is an *elementary* M-homotopy.
2. (Disjoint Squares.) $(w = w_1 w_2, w_1 w_2{}', w_1{}' w_2{}', w_1{}' w_2, w_1 w_2)$ where $(w_1, w_1{}')$ and $(w_2, w_2{}')$ are elementary M-homotopies.
3. (Rank Three Circuits.) These have the form

$$(uv_0w, uv_1w, uv_2w, \ldots, uv_{n-1}w, uv_nw = uv_0w)$$

where the words v_i, $i = 0, 1, \ldots$ are all in $M(J)$, $|J| = 3$, and (v_i, v_{i+1}) (indices modulo n) indicate elementary M-homotopies with respect to the 3×3 matrix M_J of rows and columns indexed by J.

Now let $\mathcal{A}$ be the closure of the above collection of circular walks under the operations orientation reversal, expansion and contraction.[6] Our theorem is the following.

Theorem 10.5.1 *For any reduced word w (relative to the Coxeter system (W, S)), the graph $G[w]$ is simply $\mathcal{A}$-connected.*

We first establish the following.

Lemma 10.5.2 *If wi and $w'i$ are M-homotopic reduced words, then w and w' are also M-homotopic.*

Remark This lemma fails spectacularly when wi is not reduced.

[6] See Sect. 1.3.

Proof Since wi and $w'i$ are reduced homotopic words relative to matrix M, they are in fact types of minimal galleries from 1_W to t in the thin chamber system C. (Note that $t = e(wi) = e(w'i)$.) Since C is thin, the two corresponding galleries ($\mu_1(wi)$ and $\mu_1(w'i)$) have the same next-to-last vertex s, which is i-adjacent to t. Thus the penultimate prefixes w and w' are the types of geodesic galleries from 1_W to s in C. By Lemma 9.7.2, these galleries are homotopic in C, whence w and w' are homotopic. □

Proof of Theorem 10.5.1: Let K be the set

$$\{j \in I | l(wj) < l(w)\} = \text{suffixes of words in } [w].$$

We have a homomorphism of graphs:

$$\alpha : G[w] \longrightarrow K$$

where K is viewed as the complete graph with vertex set K, and α simply reads off the final suffix of a word in $[w]$. Since any two words of the fiber $\alpha^{-1}(j)$ are reduced M-homotopic words $w_1 j$ and $w_2 j$, by Lemma 10.5.2, any circuit of M-homotopies within $\alpha^{-1}(j)$ is a circuit of elementary M-homotopies of shorter words obtained by removing the final suffix j from each word in the sequence of elementary M-homotopies. By induction on the length of w_1, we must have w_1 $\mathcal{A}$-homotopic to w_2. Thus the collection $\mathcal{A}$ satisfies condition (1) of Tits' Homotopy Theorem (Theorem 1.3.9, Sect. 1.3.3 of Chap. 1).

We next verify condition (2) of that Theorem 1.3.9. Let (i, j) be an edge in the complete image graph K. Two preimage edges in $G[w]$ must be elementary M-homotopies

$$(w_1 i, w_2 j) \text{ and } (u_1 i, u_2 j).$$

This means the elementary M-homotopies are in fact

$$\begin{aligned} w_0 p(j,i)^* &\to w_0 p(i,j)^*, \\ u_0 p(j,i)^* &\to u_0 p(i,j)^*, \end{aligned} \tag{10.6}$$

where $p(j,i)^* = \ldots iji ji$, the word of length m_{ij} alternating in i and j and arranged so that it *ends* in i. Since all factors on either side of (10.6) are reduced, w_0 and u_0 are reduced. But by repeated application of Lemma 10.5.2, w_0 and u_0 are M-homotopic. These homotopies provide paths f and g from $u_0 p(j,i)^*$ to $w_0 p(j,i)^*$ in $\alpha^{-1}(i)$ and from $w_0 p(i,j)^*$ to $u_0 p(i,j)^*$ in $\alpha^{-1}(j)$, respectively. Then the circuit

$$\begin{array}{ccc} w_0 p(j,i)^* & \to & w_o p(i,j)^* \\ f \uparrow & & \downarrow g \\ u_0 p(j,i)^* & \leftarrow & u_0 p(i,j)^* \end{array}$$

becomes a sum of disjoint squares when the columns are broken up into a sequence of elementary homotopies. Thus $\mathcal{A}$ has property (2) of Theorem 1.3.9.

It remains, then, to show property (3) of the hypothesis of Tits' Theorem 1.3.9 for the collection $\mathcal{A}$, namely that $\alpha(\mathcal{A})$ controls the cycle structure of K. Since the triangles of K, at the very least, control the cycle structure of this complete graph, it suffices to show that any triangle $\{i, j, k\}$ of the complete graph is the α-image of a circular walk in $G[w]$.

So far there is nothing special about any 3-subset of K; so, without loss of generality, we may assume $\{i, j, k\} = \{0, 1, 2, \} := J$, a subset of K.

Now for what follows let s be a parameter which varies over $Z/(3)$. From Lemma 10.4.3 we may find in the homotopy class $[w]$ a word with a desirable factorization:

$$w' = w_s \cdot p(s+1, s+2)^*$$

which is connected by a path η_s in $G[w] \cap \alpha^{-1}(s+2)$ to $w_{s+1} \cdot p(s, s+2)^*$. Thus

$$l(w_s \cdot (s+1)) = l(w_s \cdot (s+2)) > \ell(w_s).$$

But there is more. Each word w_s is the type of a minimal gallery running from 1_W to the chamber $proj_{R(s+1,s+2)}(1_W) = p_s$ where $R(s+1, s+2)$ is the residue of C of type $\{s+1, s+2\}$ containing the element $t = [w]$. (See Fig. 10.2.) But if we set $J = \{0, 1, 2\}$, each w_s is M-homotopic to a gallery in C from 1_W to p_s which passes through the gate $p_J = proj_R(w)$, where R is the residue of type $J = \{0, 1, 2\}$ containing chamber $t = [w]$. Without loss we take w_s to be this gallery, and so we have a factorization

$$w_s = w_K \cdot w_s'$$

(see Fig. 10.3).

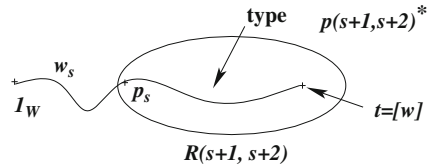

Fig. 10.2 The gallery in C from 1_w to the gated residue of type $(s+1, s+2)$ that contains $[w]$

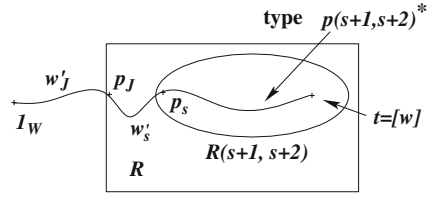

Fig. 10.3 The factorization of w_s using the gated residue of type J containing $[w]$

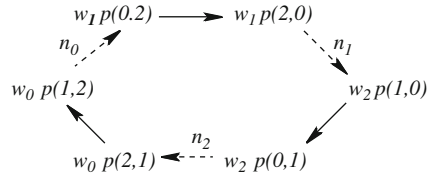

Fig. 10.4 A six-circuit of homotopies. (For technical reasons, η_s is written as n_s, $s = 0, 1, 2$)

In this way w_J is a constant prefix-factor to each w_s. This means that the M-homotopies

$$\eta_s : w_s \cdot p(s+1, s+2)^* = w_K \cdot w'_s \cdot p(s+1, s+2)^* \to w_J \cdot w'_{s+1} \cdot p^*(s, s+1)$$

not only keep the last suffix $s + 2$ fixed, but leave the prefix w_K undisturbed – in other words, η_s is composed of elementary homotopies whose replacements involve substitutions $p(ij)^* \to p(ji)^*$ for j and i in $J = \{0, 1, 2\}$.

It thus follows that the six-circuit presented in Fig. 10.4 is a circuit C' of $G[w]$ whose elementary M-homotopies take place entirely within a rank three residue, and so C' belongs to $\mathcal{A}$. But it is clear that $\alpha(C')$ is the triangle $\{0, 1, 2\}$ in the graph K. So property (3) of Theorem 1.3.9 holds for $\mathcal{A}$. The conclusion of Theorem 1.3.9 now yields the desired conclusion. □

10.6 Tits' "Local Approach" Theorem

Theorem 10.6.1 *A chamber system of type M is 2-covered by a building if and only if each of its rank three residues is 2-covered by a building.*

Proof It suffices to show that a chamber system C of type M satisfies condition (R_x) for all chambers x if and only if this condition holds within every rank three residue. The "if" part of this implication follows trivially.

So suppose condition (R_x) holds within every rank three residue containing chamber c. Let G and H be two M-homotopic galleries of the same reduced type w which begin at chamber c. We are obliged to show that $G = H$, for that will establish (R_c) for C, which is enough. We proceed by induction on the length $\ell(w)$.

Let $(G, G_1, \ldots, G_n = H)$ be an M-homotopy of galleries from G to H. (Successive pairs in this sequence represent elementary M-homotopies.) By recording the successive types of these galleries we obtain a circuit $K = (w, w_1, \ldots, w_n = w)$ in the graph $\mathcal{X} = ([G], \sim)$ of the elementary M-homotopies among the galleries of the Coxeter chamber system W. By Theorem 10.5.1, K is $\mathcal{A}$-homotopic to the trivial circuit (ϕ), where $\mathcal{A}$ are all circuits of these three types: (1) backtracks of elementary homotopies of words, (2) disjoint squares, and (3) circuits of elementary homotopies of words involving only spherical 3-subsets of I (i.e., 3-subsets J of I such that M_J is the Coxeter matrix of a finite Coxeter group). Such an $\mathcal{A}$-homotopy of circuits, $(K = K_0, K_1, \ldots, (\phi))$, forces a sequence of entire chains of galleries, each new gallery existing from the 'type M' hypothesis.

That G =H will then follow (exploiting induction on $\ell(w)$ where necessary) if we show the following.

(10.4.1) *Suppose $K' = (u_0, u_1, \ldots, u_n = u_0)$ is a circuit of $\mathcal{A}$, and G_0 is a gallery on chamber c of type u_n. Then, whenever there exists a sequence of galleries $G_1, \ldots G_n$ of respective types $u_1, \ldots, u_n = u_0$, with each (G_i, G_{i+1}) forming an elementary M-homotopy, then necessarily, $G_n = G_0$.*

So first suppose K is a backtrack $(u_o, \ldots, u_f, u_{f-1}, \ldots, u_n = u_0)$, $(n = 2f+1)$. then $G_{f-1} = G_{f+1}$, since G_{f+1} results from G_{f-1} by two elementary M-homotopies within a generalized m_{ij}-gon. Similarly, $G_{f-2} = G_{f+2}$, etc, follows by iterating these arguments, using at each stage the equality of galleries gained from the previous step. So after finitely many steps, $G = G_n$ in this case.

If K is a disjoint square, the circuit results from the application of elementary M-homotopies within two disjoint regions; specifically the $\mathcal{A}$-circuit is

$$K = (w_0w_1w_2w_3w_4, w_0w_1'w_2w_3w_4, w_0w_1'w_2w_3'w_4, w_0w_1w_3'w_4 = w_0w_1w_2w_3w_4)$$

where $w_1 \sim w_1'$ and $w_3 \sim w_3'$. Then there is a sequence of galleries

$$(G_0 \circ G_1 \circ G_2 \circ G_3 \circ G_4, G_0 \circ G_1' \circ G_2 \circ G_3 \circ G_4, G_0, G_1' \circ G_2 \circ G_3' \circ G_4,$$
$$G_0 \circ G_1'' \circ G_2 \circ G_3'' \circ G_4 \sim G_0 \circ G_2 \circ G_3 \circ G_4)$$

where $G_1 \sim G_1' \sim G_1''$ and $G_3 \sim G_3'' \sim G_3''$ as a homotopy of galleries and G_j has the same type as G_j'' for $j = 1, 3$. Then as neither G_1 nor G_3 is a trivial gallery, $\max(\ell(w_1), \ell(w_2)) < \ell(w)$, so by induction, $G_1 = G_1''$ and $G_3 = G_3''$. So once again $G_0 = G_n$.

Finally, we may assume K is a circuit in $(G, \sim)$ all of whose elementary homotopies involve only elementary homotopies within a segment of G within a rank three residue of type M. Since (R_x) holds for any x within such a gallery, the end result of this homotopy is the same as the initial one.

The proof is complete. □

Chapter 11
Locally Truncated Diagram Geometries

Abstract Roughly speaking, a locally truncated diagram geometry is a geometry over J whose residues behave as if the the geometry were a truncation to type J of a larger ambient geometry with diagram D over I. There are such geometries which in fact are not truncations of the indicated ambient geometry. This chapter gives an account of the Ronan/Bouwer–Cohen theory which gives conditions under which the ambient geometry seeming to give rise to the truncated version actually exists. To get things going, the truncated rank must be at least four. The treatment follows the sheaf-theoretic version of this theory, but with assumptions that residues are truncated buildings replaced by more modest "rigidity" conditions, and more general choices for D. The sheaf produces a chamber system over I which is often covered by a building via the "local approach theorem."

11.1 Locally Truncated Geometries: A Quick Review

Suppose D is a diagram over the typeset I and suppose J is a subset of I. A geometry Γ over J with type function *typ* is said to be *locally truncated of type D (over I)* if and only if:

(LT1) Γ is a connected geometry.[1]

(LT2) *for every flag F of Γ, the residue $Res_\Gamma(F)$ is isomorphic to the truncation to $J - typ(F)$ of a geometry $\Delta(F)$ belonging to the diagram $D_{I-typ(F)}$ (that is, the diagram D with the nodes of $typ(F)$ suppressed).*

The reader may appreciate our resistance to the temptation to recast condition (LT2) in the form that asserts that $\mathrm{Res}_\Gamma(F)$ is isomorphic to the appropriate

[1] The requirement that Γ be connected imposes no real restriction to the theory offered here, since, in the general case, the theory of this chapter is acted out piecemeal on the connected components of Γ. On the other hand the assumption simplifies the definition of a locally connected sheaf since one has the problem that residual connectedness asserts something about the residue of the empty flag, while sheaves only refer to residues of non-empty flags.

E.E. Shult, *Points and Lines*, Universitext, DOI 10.1007/978-3-642-15627-4_11,
© Springer-Verlag Berlin Heidelberg 2011

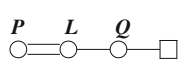

Fig. 11.1 A locally truncated diagram

truncation of a residue of a geometry Ψ belonging to diagram D. No one knows whether such a geometry Ψ exists, so this existence question should not impair the definition above. For further descriptions of locally truncated geometries of type D the reader is referred to Ronan [101] and Brouwer and Cohen [10].

A locally truncated geometry Ξ of type D can always be rendered by presenting the diagram D, and then changing all nodes not in the "real-world" set J to square nodes. For example, the diagram of Fig. 11.1 refers to a rank three geometry of points and lines and quadrangles (denoted $\mathcal{P}$, $\mathcal{L}$, $\mathcal{Q}$, respectively) with the property that for each point p, the rank two incidence geometry $\mathrm{Res}_{\Xi}(p) = (\mathcal{L}_p, \mathcal{Q}_p)$ of all lines and quadrangles on p form a geometry isomorphic to the "points" and "lines" of a $PG(3, D)$.

Like diagram geometries, locally truncated diagrams are a device for axiomatizing a geometry, except that now residuals can be certain proper truncations of diagram geometries.

The basic question is this. When is a locally truncated geometry realizable as a truncation of a geometry belonging to a diagram D? This question was first answered by Ronan [101], whose proof was replaced by a more transparent one due to A. Brouwer and A. Cohen [10]. This chapter follows the sheaf-theoretic approach of Brouwer and Cohen, hopefully fleshing it out in places, and explicating it in others. One embellishment in Sect. 11.8 is a triangle-covering of the point-collinearity graph of a suitable point-line geometry derived from the original locally truncated geometry. When the objects of the locally truncated geometry are subspaces of this point-line geometry, one obtains the main conclusion from simpler point-line data.

11.2 The Idea of Sheaves

11.2.1 Flags and Residues

Let Δ be a geometry over J which is locally truncated of type D (whose type set is I). The definition means that there is an "overall diagram" D over I, such that $J \subseteq I$ and if F is a flag of Γ of type K (necessarily K is a subset of J), then the residue $\mathrm{Res}_{\Gamma}(F)$ is the truncation to $J - K$ of a geometry belonging to the diagram D_{I-K} (the diagram remaining when the nodes of K are suppressed). Note the extra property that if $F \subseteq G$ is a containment of flags in Γ, then $typ(F) = K$ is contained in $typ(G) = L$, and that in that case one has the following.

Lemma 11.2.1 *With F, G, K, L, J, I as in the previous paragraph, the following relationships between residues and flags must hold:*

1. *$G - F$ is a flag of type $L - K$ in $Res_\Gamma(F)$.*
2. *The residue of $G - F$ in $Res_\Gamma(F)$ is naturally isomorphic to $Res_\Gamma(G)$ which "naturally" records the correspondence between any super-flag $H - F$ of $G - F$ and superflag H of G.*

11.2.2 Sheaves

Suppose now that Γ is a geometry over J which is locally truncated of type D (over I). A *sheaf* is a function $\mathcal{F}$ which assigns to each *non-empty* flag F (whose type is rendered by $\text{typ}(F)$), a geometry $\mathcal{F}(F)$ over $I - \text{typ}(F)$, whose truncation to $J - \text{typ}(F)$ is the geometry $\text{Res}_\Gamma(F)$. We must also have "connecting morphisms" in the poset of non-empty flags of Γ. For any containment of flags of Γ, $F_1 \subseteq F_2$, one has an embedding $e(F_2, F_1) : \mathcal{F}(F_2) \to \mathcal{F}(F_1)$ which induces the identity map on the objects in Γ and whose image is the residue of $F_2 - F_1$ in the codomain, thus inducing an isomorphism

$$(*)\quad \mathcal{F}(F_2) \simeq \text{Res}_{\mathcal{F}(F_1)}(F_2 - F_1)$$

as geometries over $I - \text{typ}(F_2)$. These morphisms respect compositions. If $F_1 \subseteq F_2 \subseteq F_3$ is a chain of flags, then

$$e(F_1, F_3) = e(F_2, F_3) \circ e(F_1, F_2).$$

If such a function $\mathcal{F}$ exists, we say that *a sheaf exists*.

By convention, for each object x of Γ, we regard x as also denoting the rank one flag $\{x\}$, so that we may write $\mathcal{F}(x)$ instead of $\mathcal{F}(\{x\})$.

The existence of a sheaf has been worked out for several important cases in the seminal papers of Ronan [101] and Brouwer and Cohen [10]. Fortunately their arguments apply (virtually unchanged) to a more general class of locally truncated geometries, as demonstrated in Sect. 11.3.[2]

11.2.3 The Chamber System Associated with a Sheaf

Suppose that $\mathcal{F}$ is a sheaf for the locally truncated geometry Γ with truncated diagram D. We have $J \subseteq I$ where J is the typeset for Γ and I indexes the nodes of D. Fix a flag F of Γ and select chamber flag c_F of $\mathcal{F}(F)$. Now c_F has type $I - \text{typ}(F)$ and so is a sequence of objects with those objects whose type belongs to $J - \text{typ}(F)$ forming a flag F_1 of Γ of that type. We can always write any desired subsequence

[2] Of course, in the generalized version, these authors' assumption that certain residues of flags are truncated buildings must be replaced by the more general notion of "rigidity." Beyond that their arguments do not require change.

of objects with type in J segregated, and written first, followed by the objects of the other types. This segregation can be indicated by the "$\oplus$" sign. Thus c_F can be written as $F_1 \oplus c'$ where c' is a flag chamber in $\mathrm{Res}_{\mathcal{F}(F)}(F_1) = \mathcal{F}(F \cup F_1)$. In this way we can produce a sequence involving an object of every type, with more than one way to indicate it. Thus

$$F \oplus c_F = (F \cup F_1) \oplus c' \tag{11.1}$$

would be such a sequence. We call such a sequence (where an object of each type in I occurs) an $\mathcal{F}$-*chamber* or simply a *chamber of* $C(\mathcal{F})$. Each such chamber can always be written in the extremal form where the segregation puts *all* objects having type in J first. The right-hand term of the expression in (11.1) above is of that form.

We can now convert $C(\mathcal{F})$ into a chamber system as follows. Suppose $F_1 \oplus c_1$ and $F_2 \oplus c_2$ are two $\mathcal{F}$-chambers in $C(\mathcal{F})$, where $F_i \in \mathbf{C}(\Gamma)$ and $c_i \in \mathbf{C}(\mathcal{F}(F_i))$. If $j \in J$, we say that these two $\mathcal{F}$-chambers are j-*adjacent* if and only if F_1 is j-adjacent to F_2 in Γ and $c_1 = c_2$ (as flags of type $I - J$ in the geometry $\mathcal{F}(F_1 \cap F_2)$ over $(I - J) + \{j\}$). (Note that if $|J| > 1$, then $F_1 \cap F_2$ is non-empty, so this makes sense.) But if, on the other hand, $i \in I - J$, then these two $\mathcal{F}$-chambers are i-*adjacent* if and only if $F_1 = F_2$ and c_1 and c_2 are i-adjacent chamber flags of the common geometry $\mathcal{F}(F_1)$ of type $I - J$.

With these adjacencies, $C(\mathcal{F})$ becomes a chamber system over I.

Of course, it may happen that for some flag F, the geometry $\mathcal{F}(F)$ possesses no flag chambers at all. In that case there are no $\mathcal{F}$-chambers of the form $F \oplus c_F$. In the worse scenario, $C(\mathcal{F})$ itself might be empty. Usually this is prevented by the nature of the diagram D. Such a diagram specifies rank-two residues, and we would like these to be residually connected geometries, so that a diagram $D_{\{i,j\}}$ can be attached to residues of the chamber system of type $\{i, j\}$.

Theorem 11.2.2 *Suppose* Γ *is a locally truncated diagram geometry of type* J *over a diagram* D *of type* I *where* $|J| \geq 3$. *Suppose a sheaf* $\mathcal{F}$ *exists. Then there is a canonically defined chamber system* $C(\mathcal{F})$ *over* I. *Suppose, for each object* $x \in \Gamma$, *that every rank two residue of the geometry* $\mathcal{F}(x)$ *is connected. Then the chamber system* $C(\mathcal{F})$ *also belongs to the diagram* D.

Proof Let $K = \{i, j\} \subseteq I$. Let R be a residue of $C(\mathcal{F})$ of type K. We must show that the rank two residue R belongs to the diagram $D_{\{i,j\}}$ (the relation in D between the nodes i and j). Since R is by definition non-empty, and $|J| > 2$, there exists a non-empty flag F of Γ of type $J - (K \cap J)$, such that

$$R \subseteq R_F := \{F \oplus c_F | c_F \in C(\mathcal{F}(F))\},$$

showing that R is a residue of type K of the chamber system $C(\mathcal{F}(x))$ for some $x \in F$. Since $\mathcal{F}(x)$ belongs to the diagram $D_{I-typ(x)}$ (the diagram D with the node corresponding to the type of the object x removed) and has connected rank two residues, the residue R of its chamber system $C(\mathcal{F}(x))$ belongs to the diagram $D_{\{i,j\}}$. □

11.3 Sheaves for Locally Truncated Geometries

11.3.1 Modelling Geometries on Truncations

Suppose Δ is a geometry over the set I and let $\Gamma = trunc_J(\Delta)$ be the truncation of Δ to type J, a subset of I. Thus Γ is all elements of Δ whose type is in J, with the induced incidence. Then for every element $y \in \Delta$ one may recognize the *J-shadow of y*, which is the set $U(y)$ of all elements of the truncation Γ which are incident with the element y. We say that Γ is a *shadow model* for the geometry Δ if and only if:

(M1) Every element of Δ possesses a J-shadow shared by no other element. (This is just a way of saying that the mapping $y \to U(y)$ is injective.)
(M2) There is a recognizable way to describe incidence in terms of properties of the J-shadows and the types of the objects.

Remarks 1. Sometimes we need only use the j-shadow $\Gamma_j \cap U(y)$ instead of the full J-shadow $U(y)$. Suppose Δ belongs to a diagram D over I and that D contains a node n marked by a type j in J whose removal leaves every node marked by a type in $J - \{j\}$ in a different connected component of D than that for a node marked by a type in $I - J$. From the conventions about diagrams, this means that an object of Γ is incident with an object of $\Delta - \Gamma$ whenever they are both incident with an object of type j. In this case, for each object y of Δ, the set $U(y)$ can be written

$$U(y) = \cup\{U(t) | t \in \Gamma_j \cap U(y)\},$$

and so is completely determined by the j-shadow of y — that is, the set $\Gamma_j \cap U(y)$ of all elements of type j which are incident with y in Δ.

In this case we can model the geometry in the following way. Let the objects of Γ be represented by themselves (as single objects x, not sets $U(x)$). Let the objects of $\Delta - \Gamma$ be represented by their j-shadows. Then we can check incidence among j-shadows X and Y just by checking incidence of the sets

$$X^* = \cup\{U(x) | x \in X\} \text{ and } Y^* = \cup\{U(x) | x \in Y\},$$

as posited in the assumption (M2) above. If one can recognize incidence among X^* and Y^* then that is itself a method of recognizing incidence among the j-shadows.

2. Under a mild assumption we can check incidence of an object in Γ with an object in $\Delta - \Gamma$. If they are incident with a common object of type j, they are incident. The converse is true if we assume that the residue of every flag of type $\{i, j\}$, $(i, j) \in (I - J) \times (J - \{j\})$ contains an element of type j.[3]

[3] This assumption certainly holds if every flag of Δ lies in a chamber flag, or if Δ is residually connected.

11.3.2 Geometric Rigidity

Suppose Γ is a truncation to type J of a geometry Δ over I. We say that the geometry Δ is *J-rigid* if and only if the following hold:

1. Γ is a shadow model of Δ (as in the previous section).
2. Γ can be a shadow model for a geometry isomorphic to Δ in only one way. Let us make this precise. Suppose Γ models two geometries over I which are connected by an isomorphism: $\hat{\sigma} : \Delta' \to \Delta$. Then (1) the two collections of J-shadows (one from the elements of Δ' and one from the elements of Δ) are the same collection. Moreover (2) there is an automorphism σ of Γ which preserves the recognizable incidence among J-shadows and emulates the isomorphism $\hat{\sigma}$ in the models. That is,

$$U(y)^{\sigma} = U(\hat{\sigma}(y)),$$

for all $y \in \Delta'$.

Example 1 Suppose Δ is the classical projective geometry $PG(V)$ of all proper finite-dimensional subspaces of a right vector space V over a division ring K. (This is a geometry belonging to the diagram A_n if $\dim V = n + 1$ or otherwise A_∞ with type set the positive integers.) Let Γ be the associated projective space $\mathbf{P}(V)$ of projective points and lines – that is, the truncation of Δ to the set of types $J = \{1, 2\}$. Then the J-shadow of any element of Δ is simply a proper finite-rank subspace (along with its contained lines) of the projective space $\mathbf{P}(V)$. Conversely, the points and lines of any proper finite-rank subspace of the projective space is the truncation to points and lines, of a proper finite-dimensional subspace of V – that is, an element of Δ. Thus, any projective geometry $PG(V)$ is $\{1, 2\}$-rigid.

In fact, it will be seen later that for any positive integer k, the space $\mathbf{P}(V)$ is also $\{k, k + 1\}$-rigid if $\dim V \geq k + 2$, where here $\dim V$ need not be finite. (This will be a consequence of Theorem 14.2.1 of Chap. 14, p. 506.)

Example 2 Similarly suppose Δ is the geometry of all finite-dimensional singular subspaces of a non-degenerate classical polar space geometry (see Chap. 7 for any definitions that seem unfamiliar). This is a geometry over $I = \{1, \ldots, n\}$ or over the type set $I = \{1, 2, \ldots\}$ of all positive integers, depending on the finiteness of I. In either case the type i is the vector-space dimension of the underlying vector space of the totally singular subspace comprising an object of Δ; in the finite case Δ belongs to the diagram C_n but in the infinite-rank case it belongs to the diagram A_∞, as in the previous example. Its truncation to points and lines (a truncation of type $\{1, 2\}$) yields the point-line geometry that we have called a polar *space* – rather than a polar *geometry*. Once again, the point-line geometry *alone* determines the full collection of all finite rank singular subspaces, which are the elements of Δ, so there is only one way that Δ can be modeled as a geometry of $\{1, 2\}$-shadows. One concludes that polar geometries are also $\{1, 2\}$-rigid.

11.3.3 Constructing Sheaves

Extending Sheaves Defined over Smaller Flag Posets

At the end of the previous section we defined the notion of a sheaf for a locally truncated geometry. We wish to generalize this slightly. As before, Γ is a geometry of type J which is locally truncated with respect to a diagram D over I. Let P be some subposet of the full poset of all flags of Γ. We say that function $\mathcal{F}$ attaching a geometry to the members of P is a *sheaf over flag poset* P if and only if it possesses the two conditions used to define a sheaf for at least those flags in poset P. Thus:

1. For each flag F in P, the geometry $\mathcal{F}(F)$ contains $\mathrm{Res}_\Gamma(F)$ and is isomorphic to $\Delta(F)$.
2. For any chain $F_1 \subseteq F_2$ in poset P,

$$\mathcal{F}(F_2) = \mathrm{Res}_{\mathcal{F}(F_1)}(F_2 - F_1).$$

Of course, if $\mathcal{F}$ is a sheaf, then by restriction it induces a sheaf $\mathcal{F}|_P$ over poset P. The next lemma shows that the process can be reversed when P is the induced poset of flags of rank one and two.

Lemma 11.3.1 *Suppose Γ is a geometry over J which is a locally truncated geometry with respect to diagram D. Suppose $\mathcal{F}$ is a sheaf over P_2, the family of all non-empty flags of Γ of cardinality ≤ 2. Then $\mathcal{F}$ extends to a sheaf $\mathcal{F}'$ defined over all non-empty flags of Γ.*

Proof Set $\mathcal{F}'(x) = \mathcal{F}(x)$ for each object $x \in \Gamma$. By the fact that $\mathcal{F}$ is a sheaf over P_2 we have

$$\mathcal{F}'(\{x, y\}) := \mathrm{Res}_{\mathcal{F}(x)}(y) = \mathrm{Res}_{\mathcal{F}(y)}(x) \tag{11.2}$$

for all incident pairs of distinct objects $\{x, y\}$ of Γ. This defines $\mathcal{F}'(F)$ for all flags $F = \{x, y\}$ of rank two.

Now let F be a flag of rank at least two. We define $\mathcal{F}'(F)$ to be the residue $\mathrm{Res}_{\mathcal{F}(x)}(F - \{x\})$, where x is some object in flag F. To show that $\mathcal{F}'(F)$ is well defined, suppose y is a second object in F distinct from x. Then

$$\begin{aligned}\mathrm{Res}_{\mathcal{F}(x)}(F - \{x\}) &= \mathrm{Res}_{\mathrm{Res}_{\mathcal{F}(x)}(y)}(F - \{x, y\}) \\ &= \mathrm{Res}_{\mathrm{Res}_{\mathcal{F}(Y)}(x)}(F - \{x, y\}) \\ &= \mathrm{Res}_{\mathcal{F}(y)}(\{x\} \cup (F - \{x, y\}) = \mathrm{Res}_{\mathcal{F}(y)}(F - \{y\})\end{aligned}$$

and $\mathcal{F}'(F)$ is defined independantly of the choice of x. Clearly for any such flag F

$$\begin{aligned}\mathrm{trunc}_J(\mathcal{F}'(F)) &= \mathrm{trunc}_J(\mathrm{Res}_{\mathcal{F}'(x)}(F - \{x\})) \\ &= \mathrm{Res}_{\mathrm{trunc}_J(\mathcal{F}'(x))}(F - \{x\}) \\ &= \mathrm{Res}_{\mathrm{Res}_\Gamma(x)}(F - \{x\}) = \mathrm{Res}_\Gamma(F),\end{aligned}$$

since the processes of truncation and taking residuals commute. Thus $\mathcal{F}'$ satisfies the correct identity about truncations.

It remains to be shown that $\mathcal{F}'$ is a sheaf. Let $\emptyset \neq F_1 \subseteq F_2$. We must show that

$$\mathcal{F}'(F_2) = \mathrm{Res}_{\mathcal{F}'(F_1)}(F_2 - F_1).$$

Choose $x \in F_1$. Then $x \in F_2$, so, by the definition of $\mathcal{F}'$ (used twice)

$$\begin{aligned}\mathcal{F}'(F_2) &= \mathrm{Res}_{\mathcal{F}(x)}(F_2 - \{x\}) \\ &= \mathrm{Res}_{\mathrm{Res}_{\mathcal{F}(x)}(F_1 - \{x\})}(F_2 - F_1) \\ &= \mathrm{Res}_{\mathcal{F}'(F_1)}(F_2 - F_1).\end{aligned}$$

The proof is complete. □

Sheaves for Some General Diagrams

The constructions we have for a sheaf rely on two principles: (1) the ability to model the geometries $\Delta(F)$ using J-shadows and (2) the J-rigidity of some of these same geometries.

We begin with a general situation. Γ is a locally truncated geometry over $J = \{1, 2, 3, 4\}$ with respect to the diagram D depicted in Fig. 11.2.

This is actually an entire class of diagrams called a *basic diagram*. A basic diagram indicates which rank two residues must be digons. Otherwise a bond indicates a rank-two geometry which may or may not be a digon. Unlike an ordinary diagram it does not particularly specify projective planes. (This important notion is due to Francis Buekenhout.)

In the diagram presented in Fig. 11.2, D_1 and D_2 are derived diagrams over type sets I_1 and I_2 respectively, while the type set of the geometry Γ is $J = \{1, 2, 3, 4\}$. The 2-subsets of $I = I_1 + J + I_2$ which are declared to be digons are the pairs $\{i, j\}$ in $I_1 \times I_2$, in $I_1 \times (J - \{1\})$, in $I_2 \times (J - \{4\})$ as well as pairs $\{1, 3\}$, $\{3, 4\}$, and $\{1, 4\}$ in $J \times J$. The other edges, those within the D_i, edges within Γ, edges connecting 1 to vertices of D_1 or edges connecting 4 to vertices of D_2, are simply unspecified rank-two geometries – some of them may also be digons.

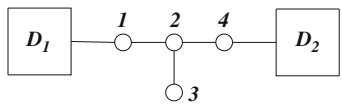

Fig. 11.2 One of the basic diagrams D considered in this section. As a basic diagram it is intended only to point out certain rank two residues which are digons, but does not give further information about rank two residues. Thus the *single bonds* here simply indicate that a digon is *not* insisted on. They do not specifically indicate projective planes as in non-basic diagrams

The construction depends on the following assumptions:

(SH1) Γ *is itself a residually connected geometry.*
(SH2) *Let* F *be a flag of* Γ *of type* $\{1, 2\}$*, and let* $\Delta(F)^R$ *be the truncation of* $\Delta(F)$ *to* $\{3, 4\} + I_2$*. Then* $\Delta^R(F)$ *is* $\{3, 4\}$*-rigid.*
(SH3) *Let* H *be a flag of* Γ *of type* $\{2, 4\}$ *and let* $\Delta^L(H)$ *be the truncation of* $\Delta(H)$ *to* $\{1, 3\} + I_1$*. Then* $\Delta^L(H)$ *is* $\{1, 3\}$*-rigid.*

(Recall that the "ideal" geometries $\Delta(F)$ or $\Delta(H)$ referred to in these hypotheses are provided by statement (*LTZ*) of the definition of locally truncated geometry, on p. 415.)

Assume x_1 is an object of type 1 in Γ. By the (LT) assumption, its residual in Γ, $\mathrm{Res}_\Gamma(x_1)$, is a truncation to $\{2, 3, 4\}$ of a geometry $\Delta(x_1)$ which has the disconnected diagram of Fig. 11.3.

In this geometry, every object whose type is in I_1 is incident with every object with type in $\{2, 3, 4\} \cup I_2$. As a result, we have no way of discerning the supposed structure of its truncation to I_1 – the part with diagram D_1 – by studying its elements in Γ.

However, the situation is different for the elements of

$$\Delta^R(x_1) := \mathrm{trunc}_{\{2,3,4\}+I_2}(\Delta(x_1)).$$

By assumption (LT),

$$\mathrm{Res}_\Gamma(x_1) = \mathrm{trunc}_J(\Delta^R(x_1)),$$

and we at least have the opportunity to model the geometry on subsets of Γ. Indeed, since the node marked "4" in the right-hand part of the diagram separates the remaining nodes with type in J from those with types in I_2, we can utilize the special model in which objects of types 2, 3, or 4 in $\Delta^R(x_1)$ are modeled by themselves and those objects y of $\Delta^R(x_1)$ whose type is in I_2 are represented by their 4-shadows $U(x_1; y)$. (The parameter x_1 is attached to the "U" to remind us that we obtained this subset of $\Gamma_{\{4\}}$ from the geometry $\Delta^R(x_1)$.) Incidence among these elements is that of $\Delta^R(x_1)$. But the disconnected derived diagram $D_{I-\{1\}}$ given just above tells us a little more, as follows.

(1) *A subset* $U(x_1; y)$ *of* $\Gamma_{\{4\}}$ *representing an element* y *of* $\Delta^R(x_1)$ *of type* $j \in I_2$ *is incident with an object* $z \in Res_\Gamma(x_1)$ *if and only if*

$$U(x_1; y) \cap Res_\Gamma(\{x_1, z\}) \text{ is non-empty},$$

– that is, z *is incident with an object of* $U(x_1; y)$.

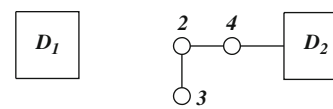

Fig. 11.3 The residue of an object of type 1

In this way we have a complete model of the geometry $\Delta^R(x_1)$ as elements of $\mathrm{Res}_\Gamma(x_1)$ and a family $\mathcal{U}^R(x_1)$ of subsets $U(x_1; y)$ with y ranging over $\Delta(x_1)_{I_2}$. Whereas $\Delta^R(x_1)$, a truncation of $\Delta(x_1)$, is to some degree an abstraction, our model of it is not, and so we give this model a separate name: $\Sigma^R(x_1) := \mathrm{Res}_\Gamma(x_1) + \mathcal{U}^R(x_1)$. Note that we have an isomorphism

$$\Sigma^R(x_1) \simeq \Delta^R(x_1)$$

as geometries over $\{2, 3, 4\} \cup I_2$.

Now let x_3 be an object of type 3 which is incident with x_1. Then we can utilize (1) to describe the residue of x_3 in the model geometry $\Sigma^R(x_1)$ as

$$\Sigma^R(x_1, x_3) := \mathrm{Res}_\Gamma(\{x_1, x_3\}) + \{U \in \mathcal{U}^R(x_1) | \mathrm{Res}_\Gamma(x_3) \cap U \neq \emptyset\}. \tag{11.3}$$

Next consider the object x_3 of type 3 in Γ on its own. Then by (LT) its residual in Γ is a truncation to J of a geometry $\Delta(x_3)$ over $I - \{3\}$ belonging to the diagram of Fig. 11.4.

We can provide a model $\Sigma(x_3)$ for this geometry as well. Let x be an arbitrary object of $\Delta(x_3)$. If the type, $\tau(x)$, of object x is in I_1, we replace x by its 1-shadow $U^L(x_3; x)$. If $\tau(x) \in \{1, 2, 4\}$, x is represented by itself, an element of Γ. If $\tau(x) \in I_2$ then x is represented in the model by its 4-shadow $U^R(x_3; x)$. Thus

$$\begin{aligned}
&\Sigma(x_3) : = \mathcal{U}^L(x_3) + \mathrm{Res}_\Gamma(x_3) + \mathcal{U}^R(x_3), \text{ where} \\
&\mathcal{U}^L(x_3) := \{U^L(x_3; x) | x \in \mathrm{trunc}_{I_1}(\Delta(x_3))\} \text{ and} \\
&\mathcal{U}^R(x_3) := \{U^R(x_3; x) | x \in \mathrm{trunc}_{I_2}(\Delta(x_3))\}.
\end{aligned}$$

Incidence is as determined in $\Delta(x_3)$, but the diagram tells us the following.

(2) *An object z in Γ is incident in $\Sigma(x_3)$ with an element $U^L(x_3; y)$ or $U^R(x_3; y)$ if and only if it is incident in Γ with any member of those sets.*

Since $\{x_1, x_3\}$ is a flag of Γ we can take the residue of x_1 in the geometry $\Sigma(x_3)$. It has two parts: (1) the $U^L(x_3; x)$, $\tau(x) \in I_1$, which contain x_1, and (2) the set

$$\Sigma^R(x_3, x_1) := \mathrm{Res}_\Gamma(\{x_1, x_3\}) \cup \{U \in \mathcal{U}^R(x_3) | \mathrm{Res}_\Gamma(x_1) \cap U \neq \emptyset)\}, \tag{11.4}$$

each part reflecting the two connected components of the derived diagram $D_{I-\{1,3\}}$ given in Fig. 11.5.

Let us denote the right part of this diagram as D^R rather than the more awkward $D_{\{2,4\}+I_2}$. Then our geometry $\Delta^R(\{x_1, x_3\}) := \mathrm{trunc}_{\{2,4\}\cup I_2}(\Delta(\{x_1, x_3\}))$ belongs to diagram D^R. By our hypothesis, $\Delta^R(\{x_1, x_3\})$ is $\{2, 4\}$-rigid. On the

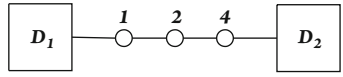

Fig. 11.4 The residue of an object of type 3

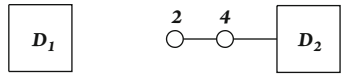

Fig. 11.5 The diagram $D_{I-\{1,3\}}$

other hand, we have inherited two distinct models for this geometry: $\Sigma^R(x_1, x_3)$ and $\Sigma^R(x_3, x_1)$, the first being a residue of $\Sigma^R(x_1)$, the second a residue of $\Sigma^R(x_3) := \text{trunc}_{\{1,2,4\}\cup I_2}(\Sigma(x_3))$. The rigidity assumption forces them to be the same model up to an isomorphism. Identifying objects y with their images under this isomorphism we can now assert the following.

(3) *For each $y \in \Delta^R(\{x_1, x_3\})$ the sets $U(x_1; y)$ and $U^R(x_3; y)$ coincide. For future reference let us write $U(y) := U^R(x_3; y) = U^R(x_1; y)$.*
Thus as geometries,

$$\Sigma^R(x_1, x_3) = \Sigma^R(x_3, x_1).$$

(This is an equal sign, not just an isomorphism.)
Moreover

$$\Sigma^R(x_1, x_3) = \text{Res}_{\Sigma^R(x_1)}(x_3) = \text{Res}_{\Sigma^R(x_3)}(x_1). \tag{11.5}$$

Now suppose $F = \{x_1, x_3, Z\}$ is a flag in Γ, where Z is either (1) an object of type 2, (2) an object of type 4, or (3) is a flag of type $\{2, 4\}$. In each case $Z \subseteq \text{Res}_\Gamma(x_1, x_3)$.

The most important set here is $\mathcal{U}^R(x_1, x_3, Z)$, which will be the set of $U(y)$s, as y ranges over $\Delta^R(\{x_1, x_3\})$, for which either $\text{Res}_\Gamma(Z) \cap U(y) \neq \emptyset$, in case (1) where Z is an object of type 2, or $Z \cap U(y) \neq \emptyset$ in cases (2) and (3) where Z contains an object of type 4. We then declare

$$\Sigma^R(x_1, x_3, Z) := \text{Res}_\Gamma(F) + \mathcal{U}^R(x_1, x_3, Z). \tag{11.6}$$

From this definition $\Sigma^R(x_1, x_3, Z)$ is the residue in $\Sigma^R(x_1, x_3)$ of the flag (or object) Z. Note that every term in Eq. (11.6) can be regarded as a subset of a residue of $\Sigma^R(x_1)$ or a subset of a residue of $\Sigma^R(x_3)$ equally well.

Let x_3' be another object of type 3 in $\text{Res}_\Gamma(x_1, Z)$. Since every element of type four which is equal to or incident with the objects of Z is also incident with x_3', we see that the following two collections of sets coincide:

$$\{U(y) | y \in \Delta^R(\{x_1, x_3\}), (Z \cup \text{Res}_\Gamma(Z)) \cap U(y) \neq \emptyset\} \text{ and}$$

$$\{U(y) | y \in \Delta^R(\{x_1, x_3'\}), (Z \cup \text{Res}_\Gamma(Z)) \cap U(y) \neq \emptyset\},$$

which we render by

$$\mathcal{U}^R(x_1, x_3, z) = \mathcal{U}^R(x_1, x_3', z). \tag{11.7}$$

Similarly we may view the right-hand side Eq. (11.6) as a subset of $\Sigma^R(x_3)$. Thus if x_1' is any element of type 1 in $\mathrm{Res}_\Gamma(x_3, z)$ we have

$$\begin{aligned}\mathcal{U}^R(x_1, x_3, Z) :&= \{U(y) | y \in \Delta^R(\{x_1, x_3\}), (Z \cup \mathrm{Res}_\Gamma(Z)) \cap U(y) \neq \emptyset\}\\ &= \{U(y) | y \in \Delta^R(\{x_1', x_3\}), (Z \cup \mathrm{Res}_\Gamma(Z)) \cap U(y) \neq \emptyset\}\\ &:= \mathcal{U}^R(x_1', x_3, Z). \end{aligned} \tag{11.8}$$

Now, since Γ is residually connected, the incidence graph of $\Gamma_{\{1,3\}} \cap \mathrm{Res}_\Gamma(z)$ is a connected bipartite graph.[4] Thus iterated usage of Eqs. (11.7) and (11.8) shows that for any two flags of type $\{1, 3, \tau(Z)\}$, containing Z — say $\{v_1, v_3, Z\}$ and $\{u_1, u_3, Z\}$ – one has

$$\mathcal{U}^R(v_1, v_3, Z) = \mathcal{U}^R(u_1, u_3, Z). \tag{11.9}$$

We may denote the common collection of sets in Eq. (11.9) by $\mathcal{U}^R(Z)$, since it is independent of the other parameters.

Thus for any flag Z of type 2, 4, or $\{2, 4\}$ we define

$$\Sigma^R(Z) = \mathrm{Res}_\Gamma(Z) + \mathcal{U}^R(Z).$$

Simply by symmetry of the diagram, and the assumption that $\Delta^L(x_3, x_4)$ is $\{1, 3\}$-rigid, one obtains similar definitions of $\mathcal{U}^L(x_4), \mathcal{U}^L(x_3)$ — and hence definitions of the "left sheaves" $\Sigma^L(x_i)$, $i = 1, 2$ as well as the natural one $\Sigma(x_4)$ and the symmetric one $\Sigma(x_3)$ which has a natural truncation $\Sigma(x_3)^L$ of type $I_1 \cup \{1, 2, 3, 4\}$ as well as the the previously defined truncation $\Sigma^R(x_3)$.

Now set

$$\Sigma(F) = \Sigma^R(F) \cup \Sigma^L(F).$$

Considering each of the definitions of Σ^R and Σ^L one may assert that $\Sigma(F)$ is in fact defined for any of the rank-one-or-two flags F of Γ.

Note that if F is a flag of rank two in Γ (and so contains an object of type distinct from 3, $\Sigma^R(F) \cap \Sigma^L = \mathrm{Res}_\Gamma(F)$ — so there is an overlap.

From the definitions, we then see that if $F = \{x, y\}$ is a flag of Γ of rank two then

$$\Sigma(F) = \mathrm{Res}_{\Sigma(x)}(\Sigma(y)). \tag{11.10}$$

So, by Lemma 11.3.1 we have produced a sheaf.

A very similar argument works for the locally truncated diagram given in Fig. 11.6.

So we conclude the following.

[4] See Corollary 2.5.3.

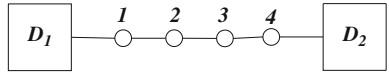

Fig. 11.6 Another locally truncated diagram. As in any basic diagram the single bonds indicate any rank two residue. (The hypotheses that for rank-two flags F_{ij} of type $\{i, j\}$, the geometries $\Delta^R(F_{1,2})$ and $\Delta^L(F_{3,4})$ are rigid with respect to their truncations to $J = \{1, 2, 3, 4\}$ normally prevents the diagram bonds for the pairs $\{1, 2\}$ and $\{3, 4\}$ from being digons)

Theorem 11.3.2 *Suppose* Γ *is a geometry over* $J = \{1, 2, 3, 4\}$ *which is a locally truncated diagram geometry over the type set* $I = I_1 + J + I_2$ *belonging to one of the locally truncated basic diagrams depicted in Figs. 11.2 and 11.6.*

1. *Suppose* Γ *belongs to the locally truncated diagram of Fig. 11.2:*

   (a)For each flag F *of type* $\{1, 3\}$*, assume that the truncation of* $\Delta(F)$ *to type* $J + I_2$ *(we called this* $\Delta^R(F)$*) is* $\{2, 4\}$*-rigid.*
   (b)For each flag F *of type* $\{3, 4\}$ *assume that the truncation of* $\Delta(F)$ *to* $J + I_1$ *(we called this* $\Delta^L(F)$*) is* $\{1, 2\}$*-rigid.*

   Then a sheaf exists.
2. *Suppose* Γ *belongs to the locally truncated diagram of Fig. 11.6:*

   (a)For each flag F *of type* $\{1, 2\}$*, assume that the truncation of* $\Delta(F)$ *to type* $J + I_2$ *is* $\{3, 4\}$*-rigid.*
   (b)For each flag F *of type* $\{3, 4\}$*, assume that the truncation of* $\Delta(F)$ *to type* $J + I_1$ *is* $\{1, 2\}$*-rigid.*

Then a sheaf exists.

Proof This was shown for flags of rank at most two in the case of the locally truncated geometry belonging to Fig. 11.2. Note that the verification of the fact that the value of a sheaf is that of a residue on the sheaf-geometry of any of its proper non-empty subsets can take place in the truncations Σ^R and Σ^L. Although this is apparent for only rank one and rank two flags in these truncations, this is sufficient to produce a sheaf by Lemma 11.3.1. □

It is worth noting that these theorems work when the extreme left or right parts of their diagrams are empty, for the simple reason that, by definition, any rank two geometry is rigid with respect to these types. Thus any locally truncated geometries belonging to the diagrams of Fig. 11.7 possess sheaves.

11.4 Recovering Γ from $C(\mathcal{F})$

The following definition is required: a geometry X over K is *strongly chamber-connected* if and only if:

1. Every flag lies in a flag chamber (a flag of type K).
2. The geometry is chamber connected – that is, the chamber system $\mathbf{C}(X)$ of chamber flags is connected.

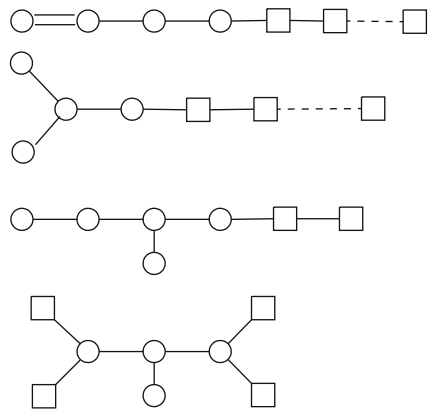

Fig. 11.7 Some locally truncated diagrams whose locally truncated geometries must possess sheaves as consequences of Theorem 11.3.2. (Note that these are not just the "basic" locally truncated diagrams which appear in the constructions preceding Theorem 11.3.2. Here the rank two residues are fully specified.) The fourth example might seem fanciful, but is included to illustrate the sweep of the possibilities

For a geometry of finite rank, this property is implied by residual connectedness; but examples show that it is weaker. (See the example of Exercise 11.2, p. 437.)

On the other hand, a connected geometry X for which the residue of every non-empty flag is strongly chamber connected is in fact residually connected. (See Exercise 11.1(ii), p. 437.)

We now continue with the hypothesis that Γ is a geometry of type J which is a locally truncated geometry with respect to the diagram D over I. In this locally truncated setting, we further assume that a sheaf $\mathcal{F}$ exists. Elements of the associated chamber system $C(\mathcal{F})$ will be called *$\mathcal{F}$-chambers* in order to distinguish them from flag chambers of Γ or of any of the geometries $\mathcal{F}(F)$.

Here is a useful result.

Lemma 11.4.1 *The following statements hold:*

1. *Suppose F is a flag of type K in the geometry Γ (so that $K \subseteq J$). Set*

$$R_F := \{F \oplus c_F | c_F \in \mathbf{C}(\mathcal{F}(F))\}.$$

   If $\mathcal{F}(F)$ is strongly chamber connected, then R_F is a residue of cotype K in the chamber system $C(\mathcal{F})$.
2. *Suppose $\mathcal{F}(F')$ is strongly chamber-connected for any flag F' of Γ of type K. Let R be a residue of $C(\mathcal{F})$ of cotype K. Then there exists a flag F of type K such that $R = R_F$.*

Proof For the first part, the $\mathcal{F}$-chambers in R are connected under the i-adjacency relations as i ranges over $I - J$, for the reason that $\mathcal{F}(F)$ is chamber connected. Suppose now, an $\mathcal{F}$-chamber $r := F_1 \oplus c_1$ in R_f were i-adjacent to some $\mathcal{F}$-chamber $c = F_2 \oplus c_2$ in $C(\mathcal{F})$, for some $i \in I - J$ (we suppose $typ(F_i) = K$).

Then $F = F_1 = F_2$ whence $c \in R$. Thus no chamber of R is i-adjacent to any chamber of $C(\mathcal{F}) - R$, for any $i \in I - K$. It follows that R_F is a residue of type $I - K$.

For the second part, we let R be an arbitrary residue of $C(\mathcal{F})$ of cotype K (i.e., it is a residue of type $I - K$). The definition of adjacency in $C(\mathcal{F})$ shows that there exists a flag F of type K such that $R \subseteq R_F$. By part 1, R_F is already a residue of type $I - K$, so $R = R_F$. □

Theorem 11.4.2 *Suppose, for each object x of the geometry Γ, the geometry $\mathcal{F}(x)$ is strongly chamber-connected. Then there is an isomorphism of geometries:*

$$\phi : \Gamma \to \mathbf{\Gamma}(C(\mathcal{F}))_J,$$

where the right side is the truncation of type J of the geometry functorially defined by the chamber system $C(\mathcal{F})$.

Proof For each object x of the geometry Γ, set $\phi(x) := R_x$ as defined in Lemma 11.4.1 (recall that x is regarded as a flag of rank one, so we don't have to write $\{x\}$ here). By the hypothesis on $\mathcal{F}(x)$ and Lemma 11.4.1, R_x is a residue of cotype $typ(x) \in J$ and hence is an object of type $typ(x)$ in the geometry $(\mathbf{\Gamma}(C))_J$. Conversely, any object of $\mathbf{\Gamma}(C(\mathcal{F}))_J$ is a residue R of $C(\mathcal{F})$ of cotype j, for some $j \in J$. Now by Lemma 11.4.1, part 2, R has the form $R = R_x$ for some object x of type j. Thus ϕ is a surjective mapping.

Now suppose (x, y) is an incident pair of (necessarily distinct) objects of Γ. By strong chamber-connectedness of $\mathcal{F}(x)$, there is a flag-chamber of $\mathcal{F}(x)$ containing $\{x, y\}$, and hence an $\mathcal{F}$-chamber $c \in R_x \cap R_y$. Thus the images of x and y under ϕ are incident objects of $(\mathbf{\Gamma}(C))_J$. Thus ϕ is a morphism of geometries. □

It is also a full morphism, for if $R_x \neq R_y$, then $R_x \cap R_y$ contains an $\mathcal{F}$-chamber of the form $\{x, y\} \oplus c'$, when $\{x, y\}$ is a flag of rank two, whence x is incident with y.

Finally, suppose x and x' are distinct objects of Γ of the same type. Then $R_x \neq R_y$ since they contain no $\mathcal{F}$-chamber in common. Thus $\phi(x) \neq \phi(x')$. So ϕ is injective. Now ϕ is an isomorphism of geometries.

11.5 Residual Connectedness of Γ and $(C(\mathcal{F}))_J$

Suppose Γ is a locally truncated geometry of type J with respect to a diagram D over a type set I containing J. Suppose $\mathcal{F}$ is a sheaf for Γ. We say that *the sheaf $\mathcal{F}$ is residually connected* if and only if, for each object x of Γ, the geometry $\mathcal{F}(x)$ is a residually connected geometry. (Quite obviously, these conditions also force $\mathcal{F}(F)$ to be residually connected for each non-empty flag F (Exercise 11.3).)

Lemma 11.5.1 *Assume the sheaf $\mathcal{F}$ is residually connected. Then the following statements hold:*

1. *The geometry* Γ *is residually connected.*
2. *There is an isomorphism*

$$\mathbf{C}(\Gamma) \simeq C(\mathcal{F})_J,$$

   as chamber systems over J.
3. *The chamber systems of the preceding statement are residually connected chamber systems.*

Proof Let F be any fixed non-empty flag of Γ. From the definition of a sheaf,

$$\mathrm{Res}_\Gamma(F) \simeq (\mathcal{F}(F))_J. \tag{11.11}$$

Moreover, $\mathcal{F}(F)$ is a geometry over $I - t(F)$, and, for any object $x \in F$, the residue

$$\mathrm{Res}_{\mathcal{F}(x)}(F - \{x\})$$

is a residue of the residually connected geometry $\mathcal{F}(x)$, and so is itself residually connected. As a result:

(a) $\mathcal{F}(F)$ is non-empty if $I - t(F)$ is non-empty.
(b) $\mathcal{F}(F)$ is chamber-connected – that is, the chamber system $\mathbf{C}(\mathcal{F}(F))$ is connected, rank one or is empty (only if $t(F) = J = I$).
(c) Any truncation of $\mathcal{F}(F)$ of rank at least two is also residually connected.

Now, if F is a flag of Γ of corank one so $\{J - t(F)\} = \{j\}$ for some $j \in J$, then $\mathcal{F}(F)$ has chambers (by (b)), and so contains an object of type j, so $\mathrm{Res}_\Gamma(F)$ is non-empty. On the other hand, if F is non-empty of corank at least two, then, by Eq. (11.11) and (b) and (c), $\mathrm{Res}_\Gamma(F)$ is connected. Finally, if F is empty, its residue is Γ itself, of rank at least three, and connected by the initial hypothesis (see the definition of "locally truncated," p. 415). Thus Γ is a residually connected geometry.

The chamber system $C(\mathcal{F})$ has as its chambers, elements of the form $c = F \oplus c_F$, where F is a flag of Γ of type J (that is, a flag-chamber of Γ) and c_F is a flag-chamber of $\mathcal{F}(F)$. The residue of $C(\mathcal{F})$ of type $I - J$ which contains chamber c must consist of chambers of the form $F \oplus c'_F$ where c'_F wanders over a connected component of the chamber system $\mathbf{C}(\mathcal{F}(F))$ containing c'_F. Thus the mapping

$$\psi : C(\mathcal{F}) \to \mathbf{C}(\Gamma),$$

which takes each chamber $F \oplus c_F$ to F (that is, it reads off the J-part of each chamber flag) satisfies this important property:

(P) If chamber c is i-adjacent to chamber c' in $C(\mathcal{F})$, then either $\psi(c) = \psi(c')$ and $i \in I - J$, or else $\psi(c)$ and $\psi(c')$ are distinct and i-adjacent for some i in J.

Notice that (c) above implies that every flag-chamber F of Γ is the ψ-image of a chamber $F \oplus c_F$, and so ψ is surjective. But in particular (P) implies ψ is a morphism of chamber systems and that the kernel of ψ is a partition of $C(\mathcal{F})$ into fibers which are unions of $(I - J)$-residues — that is, the map ψ factors through a morphism $\bar{\psi} : C(\mathcal{F})_J \to C(\Gamma)_J$. Now, finally, (b) above, tells us that in fact each of these fibers is a single $I - J$-residue of $C(\mathcal{F})$. Thus the induced mapping $\bar{\psi}$ is an injective morphism of chamber systems over J.

It only remains to be shown that $\bar{\psi}$ is a full morphism. Suppose A and B are distinct elements of $C(\mathcal{F})_J$ (that is, residues of type $I - J$ in $C(\mathcal{F})$), such that $F := \psi(A)$ is j-adjacent to $G := \psi(B)$ in $C(\Gamma)$. Then as $|J| \geq 2$, $H := F \cap G$ is a non-empty flag of cotype j in Γ. Then since $\mathcal{F}(H)$ is residually connected, $F - H$ and $G - H$ lie in respective flag chambers c_F and c_G of $\mathcal{F}(H)$, and $a := H \oplus c_F = F \oplus c'_F$ and $b := H \oplus c_G = G \oplus c'_G$ are connected by a gallery whose type is a word in $(I - J) \cup \{j\}$ (this gallery corresponds to one in $C(\mathcal{F}(H))$). But $R_F = A$ and $R_G = B$ contain a and b, respectively, and so lie in a common residue T of type $(I - J) \cup \{j\}$. Thus $\bar{\psi}$ is a full bijective morphism, and so is an isomorphism of chamber systems. This proves the second statement.

The third conclusion is immediate for if Γ is residually connected, $\mathbf{C}(\Gamma)$ is residually connected as a chamber system. The proof is complete.

11.6 Configurations Produced by a 1-Covering of $C(\mathcal{F})$

This section concerns what happens when we have a 1-covering of the chamber system $C(\mathcal{F})$. We must standardize both the notation and the hypotheses:

1. (The locally truncated geometry.) As has been standard so far, Γ is a geometry over J which is locally truncated with respect to the diagram D over I.
2. (Existence of a sheaf.) We assume there is a sheaf $\mathcal{F}$ defined for this local truncation, and we let $\bar{C}$ denote its associated chamber system over I. If each $\mathcal{F}(x)$ is strongly chamber-connected, we know from Theorem 11.4.2 that there is an isomorphism $\phi : \Gamma \to \mathbf{\Gamma}(C(\mathcal{F}))_J$.
3. (The covering of the chamber system of the sheaf.) We assume there is a 1-covering $\kappa : C \to \bar{C}$.
4. (The morphisms of geometries.) We let $\bar{\Delta} := \mathbf{\Gamma}(C(\mathcal{F}))$ and $\Delta := \mathbf{\Gamma}(C)$ be the geometries over I defined by the chamber systems $\bar{C}$ and C, respectively. (Recall that a flag in $\mathbf{\Gamma}(C)$ is just a collection of pairwise intersecting corank-one residues of C.) Such a flag is called a $\mathbf{\Gamma}$-*image* if all these residues lie on a common chamber. So some flags of $\mathbf{\Gamma}(C)$ are $\mathbf{\Gamma}$-images, and some might not be. (Of course all flags are $\mathbf{\Gamma}$-images if C is residually connected as a chamber system.) A similar distinction applies to the flags of $\bar{\Delta} := \mathbf{\Gamma}(\bar{C})$.) We let $h := \Gamma(\kappa)$ be the functorially defined geometry morphism $\Delta \to \bar{\Delta}$. We let $h_J : \Delta_J \to \bar{\Delta}_J$ be the morphisms induced by h on their truncations of type J.

Lemma 11.6.1 *Assume these hypotheses:*

(i) $\kappa : C \to \bar{C} = C(\mathcal{F})$ *is a* 1*-covering of chamber systems.*
(ii) The sheaf $\mathcal{F}$ *is strongly chamber-connected.*

Then the following statements hold:

1. *The induced* $h_J : \Delta_J \to \bar{\Delta}_J$ *is a full epimorphism. Thus there is a full epimorphism of geometries*

$$f = \phi^{-1} \circ h_J : \Delta_J \to \Gamma.$$

2. *Every flag of* Γ *is the image under* f *of a flag of* Δ_J *which is a* $\mathbf{\Gamma}$*-image with respect to* $\mathbf{\Gamma}(C)$.
3. (The flag-lifting property.) *Suppose* X' *is a flag of* Δ_J *which is a* $\mathbf{\Gamma}$*-image such that* $h_J(X') = X$. *Suppose further that* F *is any flag of* Γ *such that* $X \subseteq F$. *Then there exists a flag* F' *of* Δ_J *(also a* $\mathbf{\Gamma}$*-image) incident with* X' *such that* $h_J(F') = F$.

Proof Consider a containment of non-empty flags, $F_1 \subseteq F_2$ of Γ, and an object $x \in F_1$. Since $\mathcal{F}$ is strongly chamber-connected, by Lemma 11.4.1, $R_{F_2} \subseteq R_{F_1} \subseteq R_x$ is a containment of residues of $\bar{C}$ of cotypes $typ(F_2)$, $typ(F_1)$, and $typ(x)$, respectively. Let c be a chamber in R_{F_2}. Since κ is surjective, there is a preimage c' of c in C. Then let $R_2 \subseteq R_1 \subseteq R(x)$ be the residues of C of respective cotypes $typ(F_2)$, $typ(F_1)$, and $typ(x)$, containing c'. Then, since κ is a 1-cover, R_2, R_1, and $R(x)$ are the lifts of R_{F_2}, R_{F_1}, and R_x at c', respectively.

We apply the development of this paragraph to prove the first two parts of the conclusion of the lemma.

Taking $F_1 = F_2$, we see that $F_1' := \{R(x) | x \in F_1\}$ is a flag of Δ of type $typ(F_1)$ which is a $\mathbf{\Gamma}$-image and which maps onto F_1. This proves part 2. But it also proves part 1, since (a) $x = f(R_x)$ for all $x \in \Gamma$, and (b) any incident pair of objects in Γ lies in a flag F_1 which is the image of the flag F_1' just described.

Part 3. Suppose X' is a $\mathbf{\Gamma}$-image of Δ_J mapping to a flag X of Γ which lies in a larger flag F of Γ. Since X' is a $\mathbf{\Gamma}$-image, it is a collection of corank one residues $R(x)$ of C, whose intersection contains a residue R_1 of cotype $typ(X)$ which maps via κ onto R_X. Choose a chamber $c \in R_F$. It possesses a preimage $c' \in R_1$. Now let R_2 be the residue of cotype τF on c'. Now we have

$$c' \in R_2 \subseteq R_1.$$

Then the full set of corank one residues of C containing R_2 is the desired $\mathbf{\Gamma}$-image flag F' of part 3. □

11.7 Applying the Local Approach Theorem: A Useful Tool for Point-Line Characterizations

Theorem 11.7.1 *Suppose* Γ *is a locally truncated geometry of type* J *with respect to the diagram* D *of type* M *over the type set* I*. We suppose this structure supports a strongly chamber connected sheaf* $\mathcal{F}$*, with the property that every rank three residue of* $C(\mathcal{F})$ *is* 2*-covered by a building.*[5]

Then Γ *is a full homomorphic image of the truncation to type* J *of a building geometry* B *belonging to diagram* D *over* I*.*

Remark Recall that the adjective "full" indicates that every incidence pair of objects in the image geometry is an image of an incident pair in the domain geometry.

Proof Since $\mathcal{F}$ is strongly chamber connected, Theorem 11.2.2, implies that the chamber system $C(\mathcal{F})$ belongs to the diagram D. Since every rank three residue of the chamber system $C(\mathcal{F})$ is covered by a building, Tit's "local approach theorem" (Theorem 10.6.1, p. 412) produces a 2-covering

$$\kappa : C \to C(\mathcal{F}).$$

Then, applying the functor $\mathbf{\Gamma}$ one obtains a new morphism

$$\mathbf{\Gamma}(\kappa) : B := \mathbf{\Gamma}(C) \to \mathbf{\Gamma}(C(\mathcal{F})).$$

A second functorial morphism – that of truncation to type J – induces a second morphism

$$\mathbf{\Gamma}(\kappa)_J : B_J \to \mathbf{\Gamma}(C(\mathcal{F}))_J.$$

Now the fact that $\mathcal{F}$ is strongly chamber connected implies two things: (1) that the morphism $\mathbf{\Gamma}(\kappa)_J$ is a full epimorphism (Lemma 11.6.1, part 1) and (2) that there is an isomorphism

$$h : \Gamma \to \mathbf{\Gamma}(C(\mathcal{F}))_J.$$

Simply composing the appropriate morphisms reveals that Γ is indeed a homomorphic image of B_J, the truncation to the type set J, of a building geometry belonging to diagram D . □

[5] This would happen under several hypotheses; for example (i) the type M diagram D has no rank three residue of type C_3, or (ii) $|J| \geq 4$, and for every object $x \in \Gamma$, every residue of a flag of corank three in $\mathcal{F}(x)$ is 2-covered by a building geometry.

11.8 A Covering of a Point-Collinearity Graph Derived from a Locally Truncated Geometry

In this section we assume the notation of items 1–4 at the beginning of Sect. 11.6. In addition we assume:

1. D is a Coxeter diagram over the finite set I whose rank three residues are covered by buildings.
2. $\kappa : C \to \bar{C}$ is a universal 2-covering.
3. The sheaf $\mathcal{F}$ is residually connected.

Since D is a Coxeter diagram with rank three residues covered by buildings, C is a chamber system over I which is a building and so is residually connected. Accordingly, the associated geometry $\Delta := \mathbf{\Gamma}(C(\mathcal{F}))$ is a building geometry over I.

Now assumption 3 implies $\mathcal{F}(F)$ is chamber-connected for every non-empty flag F of Γ. Thus hypotheses (i) and (ii) of part 1 of Lemma 11.6.1 hold, so f induces the full epimorphism

$$f : \Delta_J \to \Gamma \tag{11.12}$$

of geometries over J.

Again, since the sheaf is residually connected, part 2 of Lemma 11.6.1 forces the morphism (11.12) to possess the flag-lifting property restated as follows.

(FL) *If X' is a flag of Δ_J that is a $\mathbf{\Gamma}$-image such that $f(X') = X$, and F is a flag of Γ incident with X (i.e., $X \cup F$ is a flag), then there exists a flag F' of Δ_J with $f(F') = F$ and $F' \cup X'$ is a flag of Δ_J.*

Finally, our hypothesis that Γ is connected and the assumption 3 that $\mathcal{F}$ is residually connected, imply three statements: (S1) that Γ is a residually connected geometry, (S2) that $\mathbf{C}(\Gamma) \simeq C(\mathcal{F})_J$ as chamber systems over J, and (S3) that $C(\mathcal{F})_J$ is residually connected as a chamber system. All three conclusions are a direct application of Lemma 11.5.1.

We are now ready to introduce the main theorem of this section.

It is sometimes possible to produce rank-two truncations of the geometries Δ_J and Γ of Eq. (11.12) to obtain point-line geometries $(\mathcal{P}_\Delta, \mathcal{L}_\Delta)$ and $(\mathcal{P}_\Gamma, \mathcal{L}_\Gamma)$. Then the morphism of (11.12) induces a morphism $\phi : (\mathcal{P}_\Delta, \sim) \to (\mathcal{P}_\Gamma, \sim)$ of their point-collinearity graphs. The theorem below is sometimes useful: it gives a condition under which the mapping ϕ of point-collinearity graphs is a $\mathcal{T}$-covering, where (in the terminology of Chap. 1) $\mathcal{T}$ is the collection of all pointed 3-circuits in the graph $(\mathcal{P}_\Gamma, \sim)$.

This observation can be useful, for if the latter graph is $\mathcal{T}$-simply connected and the geometry Γ is $\{\mathcal{P}_\Gamma, \mathcal{L}_\Gamma\}$-rigid, then the morphism in (11.12) is an isomorphism.

Theorem 11.8.1 *Suppose $t_{\mathcal{P}}$, $t_{\mathcal{L}}$, and t_{Π} are three pairwise disjoint subsets of J. Let $(\mathcal{P}_\Gamma, \mathcal{L}_\Gamma, \Pi_\Gamma)$ and $(\mathcal{P}_\Delta, \mathcal{L}_\Delta, \Pi_\Delta)$ be the triples of collections of all flags of these respective types in the geometries Γ and Δ_J, respectively. We assume:*

(i) For $E \in \{\Gamma, \Delta\}$, and each object $X \in \Pi_E$, $Res_E(X) \cap \mathcal{P}_E$ is a singular subspace of $P_E := (\mathcal{P}_E, \mathcal{L}_E)$, regarded as a point-line geometry.
(ii) For any triple $(a, b, c) \in \mathcal{P}_\Gamma \times \mathcal{P}_\Gamma \times \mathcal{P}_\Gamma$ of pairwise collinear points of Γ there exists an object $X \in \Pi_\Gamma$ such that $\{a, b, c\} \subseteq Res_\Gamma(X) \cap \mathcal{P}_\Gamma$.
(iii) If $X \in \mathcal{P}_\Gamma \cup \mathcal{L}_\Gamma$, then $\mathcal{F}(X)$ is a 2-simply connected geometry.
(iv) $(\mathcal{P}_\Gamma, \mathcal{L}_\Gamma)$ is a partial linear space.

Then the graph morphism

$$\phi : (\mathcal{P}_\Delta, \sim) \to (\mathcal{P}_\Gamma, \sim)$$

of point-collinearity graphs of $(\mathcal{P}_\Delta, \mathcal{L}_\Delta)$ and $(\mathcal{P}_\Gamma, \mathcal{L}_\Gamma)$ induced by the truncated mapping

$$f : \Delta_J \to \Gamma$$

is a $\mathcal{T}$-covering of graphs, where $\mathcal{T}$ is the collection of all pointed 3-circuits in $(\mathcal{P}_\Gamma, \sim)$.

Remark 1. Conditions (i), (ii), and (iv) hold, for example, if $(\mathcal{P}_\Gamma, \mathcal{L}_\Gamma)$ is a parapolar space and Π_Γ is the full collection of all its projective plane subspaces.
2. Condition (iii) is satisfied if $\mathcal{F}(X)$ is a building geometry over $I - typ(X)$.

Proof of Theorem 11.8.1 For the first two steps below, we need to establish the hypotheses of Lemma 10.2.1. We have on hand the 2-covering of chamber systems: $\kappa : C \to \bar{C} = C(\mathcal{F})$ belonging to diagram D.

Let X be a flag of Γ whose type is an element of $t_{\mathcal{P}} \cup t_{\mathcal{L}}$ — that is, X is either a point or a line of the geometry $(\mathcal{P}_\Gamma, \mathcal{L}_\Gamma)$. By Lemma 11.4.1

$$R_X := X \oplus C(\mathcal{F}(F)) = \{X \oplus c_X | c_X \in \mathbf{C}(\mathcal{F}(F)\}$$

is a residue of cotype $typ(X)$. By hypothesis (iii) of the theorem, $\mathbf{C}(\mathcal{F}(X))$ is 2-simply connected. Since the mapping $X \oplus c_X \to c_X$ induces an isomorphism of $\mathbf{C}(\mathcal{F}(X))$ and R_X, we see that:

(A-1) *R_X is 2-simply connected as a chamber system over $I - typ(X)$.*

But by property (S3) just preceding the statement of this theorem, we conclude:

(A-2) *$C(\mathcal{F})_J$ is residually connected.*

We have now assembled all of the hypotheses of Lemma 10.2.1, where C, $C(\mathcal{F})$, and $typ(X)$ fulfil the roles of C, $\bar{C}$, and K of that lemma. Thus the lemma produces this conclusion:

(LI) *If $\hat{X}$ is a flag of Δ_J with $f(\hat{X}) = X$, then the mapping*

$$f_{\hat{X}} : Res_{\Delta_J}(\hat{X}) \to Res_{\Gamma}(X),$$

induced by the restriction of $f : \Delta_J \to \Gamma$, is injective on objects. It is therefore injective when restricted to the sets of flags of any prescribed type of these residues.

If X is a point, then we have:

Step 1. (Local injectivity of lines.) *If L and N are distinct lines of $\mathcal{L}_\Delta$ incident with a common point p of P_Δ then $\phi(L)$ is not equal to $\phi(N)$ – that is, ϕ is injective when resticted to the lines incident with a common point.*

If X is a line we have:

Step 2. (Local injectivity of points.) *If p and q are distinct points of $\mathcal{P}_\Delta$ incident with a common line L of P_Δ, then $\phi(p)$ is not equal to $\phi(q)$ – that is, ϕ is point-injective when restricted to a line.*

We conclude:

Step 3. *$\phi : (\mathcal{P}_\Delta, \sim) \to (\mathcal{P}_\Gamma, \sim)$ is a fibering of graphs, that is:*

1. *Each fiber $\phi^{-1}(x)$ of a vertex $x \in \mathcal{P}_\Gamma$ is a coclique of $(\mathcal{P}_\Delta, \sim)$.*
2. *Moreover, ϕ induces a bijective mapping*

$$P_\Delta(x) \to P_\Gamma(\phi(x))$$

when restricted to neighborhood subgraph of a point $x \in \mathcal{P}_\Delta$.
3. *ϕ is both point and edge surjective.*

Thus for every path in $(\mathcal{P}_\Gamma, \sim)$ and any specified point x in the fiber above the initial point of this path, there is a unique lift in $(\mathcal{P}_\Delta, \sim)$ of this path beginning at x.

Proof of Step 3 Part 1 of this step follows from Step 2. For part 3 we use the flag-lifting property (FL) to conclude that every point and edge of $(\mathcal{P}_\Gamma, \sim)$ is the image of a point or edge of $(\mathcal{P}_\Delta, \sim)$.

To prove part 2 of Step 3, we invoke for the first time the hypothesis (iv) that $(\mathcal{P}_\Gamma, \mathcal{L}_\Gamma)$ is a partial linear space. Suppose (y, y_1) and (y, y_2) are distinct edges of $(\mathcal{P}_\Delta, \sim)$. Then there are (not necessarily unique) lines N_i of $\mathcal{L}_\Delta$ containing $\{y, y_i\}$, $i = 1, 2$. Suppose by way of contradiction that $\phi(y_1) = \phi(y_2)$. By Step 1, this point is distinct from $\phi(y)$ but collinear with it. Then $f(N_1)$ and $f(N_2)$ are lines of $\mathcal{L}_\Gamma$ on the distinct points y and $\phi(y_1) = \phi(y_2)$. Since $(\mathcal{P}_\Gamma, \mathcal{L}_\Gamma)$ is a partial linear space, we have $f(N_1) = f(N_2)$. But since the N_i lie in $\mathrm{Res}_{\Delta_J}(y)$, Step 2 forces $N_1 = N_2$. But in that case y_1 and y_2 are distinct points of N_1 mapping to a common point of $f(N_1)$, against Step 1. Thus $\phi(y_1) = \phi(y_2)$ is impossible, and so ϕ restricts to a vertex-injective mapping of the set of neighborhood vertices of a given vertex. That ϕ restricted to a neighborhood of y in $(\mathcal{P}_\Delta, \sim)$ maps onto the neighborhood of $\phi(y)$ in $(\mathcal{P}_\Gamma, \sim)$ follows from the fact that any line of $\mathcal{L}_\Gamma$ on $\phi(y)$ lifts to a line on y in $\mathcal{L}_\Delta$ by condition (LF). Thus the induced mapping on vertex neighborhoods is bijective on vertices as required.

Step 4. *Every lift of a* 3*-circuit in* $(\mathcal{P}_\Gamma, \sim)$ *is a* 3*-circuit of* $(\mathcal{P}_\Delta, \sim)$.

Proof of Step 4: Suppose (a, b, c) is a 3-circuit of $(\mathcal{P}_\Gamma, \sim)$. By hypothesis (ii) of the theorem, there exists an object $X \in \Pi_\Gamma$, incident with the three points a, b, and c. Since $f : \Delta_J \to \Gamma$ has the flag-lifting property (FL), for every point a' in the fiber $\phi^{-1}(a)$ there exists a flag X' of Δ_J incident with a' such that $f(X') = X$. Also, by (FL), each point $x \in \{a, b, c\}$ has a preimage $x' \in \mathcal{P}_\Delta$ incident with X'. Now by hypothesis (i) of this theorem, $\mathrm{Res}_{\Delta_J}(X') \cap \mathcal{P}_\Delta$ is a singular subspace of $(\mathcal{P}_\Delta, \mathcal{L}_\Delta)$, and the three points a', b', and c' are pairwise collinear. Thus the 3-circuit (a', b', c') is the unique lift of the 3-circuit (a, b, c) at the point a' of the fiber above a.

The proof that ϕ is a T-covering is complete. □

11.9 Exercises for Chapter 11

11.9.1 Exercises Involving Connectedness Properties

11.1 Prove the following statements:

(i) If X is a residually connected geometry of finite rank, then X is strongly chamber connected. [Hint: Use part 3 of Lemma 2.5.1 of Chap. 2.]

(ii) If X is a connected geometry for which the residue of every non-empty flag is strongly chamber connected, then X is residually connected. [Hint: Observe that from the definitions, any non-empty rank-one geometry (a set of objects of the same type) is automatically chamber connected since its objects are adjacent chambers.]

11.2 Consider the following rank three geometry. It consists of eight objects:

$$\begin{array}{ll} \text{Type 1:} & a_1, a_2, a_3 \\ \text{Type 2:} & b_1, b_2, b_3 \\ \text{Type 3:} & c_1, c_2. \end{array}$$

There are eight chamber flags:

$$(a_1, b_1, c_1), (a_1, b_2, c_1), (a_1, b_2, c_2)$$
$$(a_3, b_3, c_2), (a_3, b_2, c_2), (a_3, b_2, c_1)$$
$$(a_2, b_1, c_1), (a_2, b_3, c_2).$$

All other incidences are those that appear in the chamber flags; that is, each flag lies in at least one chamber flag.

(i) Show that this geometry is strongly chamber connected.

(ii) Show that this geometry is not residually connected.

11.3 Suppose $\mathcal{F}$ is a sheaf for the locally truncated diagram geometry Γ over typset J with respect to diagram D over I. Show that $\mathcal{F}$ is a residually connected sheaf if and only if, for each non-empty flag F of Γ, $\mathcal{F}(F)$ is a residually connected geometry. [Compare with the definition on p. 429.]

11.9.2 Morphisms by Chamber-Semiregular Automorphism Groups Do Not Preserve Residual Connectedness

Suppose B is a building geometry over a finite type set I and suppose A is a group of type-preserving automorphisms of B all of whose non-identity elements act freely on the chambers of B. The resulting morphism

$$\alpha : B \to B/A$$

of geometries over I is exactly the sort of morphism that is produced in the sheaf-theoretic approach to locally truncated geometries presented in this chapter.

If F is a flag of corank at least two in B, then the residue $\mathrm{Res}_B(F)$ is connected since B is a residually connected geometry. Also F^A – the A-orbit of F – is a flag of the image geometry B/A of the same type as F. Since any flag G of B that contains F is mapped to a flag G^A of B/A that contains F^A, there is a natural morphism

$$\alpha_F : \mathrm{Res}_B(F) \to \mathrm{Res}_{B/A}(F^A)$$

as geometries over $I - typ(F)$.

Question: *Is the morphism α_F surjective on objects?*

The following series of exercises will show that in general α_F need not be surjective. In fact, in our example, $\mathrm{Res}_{B/A}(F^A)$ is not even a connected geometry. Thus the following applies.

> *In general, when B is a building of finite rank and $A \subseteq Aut(B)$ acts freely on the chambers of B, the factor geometry B/A may fail to be residually connected.*

11.4 Let $V = \mathbb{C}^{(2n)}$ be the vector space of $2n$-tuples with entries from the field of complex numbers $\mathbb{C}$. Let $Q : V \to \mathbb{C}$ be the "sum of squares" quadratic form:

$$Q : (x_1, \ldots, x_{2n}) \mapsto \sum x_i^2.$$

Show that the set of $2n$-tuples of the form

$$M_1 = \{(a_1, \ldots, a_n; ia_1, \ldots, ia_n) | a_j \in \mathbb{C}\}$$

forms a totally singular subspace of V of vector space dimension n.

11.5 For $1 \leq i \leq n$, let $\mathcal{S}_i$ be the collection of all totally singular subspaces of V of vector-space dimension i. Show that $\Gamma = (\mathcal{S}_1, \mathcal{S}_2)$ under the natural incidence is an oriflame polar space of rank n with two classes of maximal singular subspaces $\mathcal{S}_n = \mathcal{M}_1 + \mathcal{M}_2$.[6] One then forms the geometry

$$B := (\mathcal{S}_1, \mathcal{S}_2, \ldots, \mathcal{S}_{n-1}, \mathcal{M}_1, \mathcal{M}_2)$$

where incidence among all subspaces of differing dimensions is containment and each element of $\mathcal{M}_i$ is incident with those members of $\mathcal{M}_{3-i}$ which meet it in a hyperplane, $i = 1, 2$.
Show that B is a building belonging to the diagram of type D_n with the $\mathcal{M}_i$ in the role of the oriflame nodes.

11.6 Let s be the semilinear transformation which transforms each vector $u = (y_1, \ldots, y_{2n})$ to $s(u) := (\bar{y_1}, \ldots, \bar{y_{2n}})$ where the"bar" denotes the taking of a complex conjugate. Then s preserves the collection of all singular subspaces of V, their dimensions and their containment relations and so induces an automorphism σ of B of order 2 which preserves all $\mathcal{S}_i$, $i \leq n-2$, and the set $\{\mathcal{M}_1, \mathcal{M}_2\}$.

(1) Show that if $M_i \in \mathcal{M}_i$, then $M_i \cap M_i^\sigma$ contains only the zero vector.
(2) Show that σ transposes $\mathcal{M}_1$ and $\mathcal{M}_2$ if and only if n is odd.
(3) If n is even, the type-preserving group $A := \langle \sigma \rangle$ acts freely on the chambers of B.

11.7 Now assume $n \geq 8$ is even and forms the geometry morphism

$$\alpha : B \to B/A.$$

We form the flag $\{M_1, M_2\}$ of type $\{\mathcal{M}_1, \mathcal{M}_2\}$ where (as above)

$$M_1 = \{(a_1, \ldots, a_n; ia_1, \ldots, ia_{2n}) | a_j \in \mathbb{C}\}$$

and

$$M_2 = \{(ia_1, a_2, \ldots, a_n; a_1, ia_2, \ldots, ia_n) | a_j \in \mathbb{C}\}$$

so

$$H = M_1 \cap M_2 = \{0, a_2, \ldots, a_n; 0, ia_2, \ldots, ia_n) | a_j \in \mathbb{C}\}$$

is a hyperplane of both M_1 and M_2.

---

[6] The reader may recall that this occurs because the graph $(\mathcal{S}_n, \sim)$ (where $A \sim B$ if and only if $A \cap B$ is a hyperplane of both A and B) is a bipartite graph.

Then

$$\begin{aligned} M_1^\sigma &= \{(\bar{a_1}, \ldots, \bar{a_n}; -i\bar{a_1}, \ldots, -i\bar{a_n}) | a_j \in \mathbb{C}\} \\ &= \{(a_1 \ldots, a_n; -ia_1, \ldots, -ia_n) | a_j \in \mathbb{C}\} \end{aligned}$$

and similarly

$$M_2^\sigma = \{-ia_1, a_2 \ldots, a_n; a_1, -ia_2, \ldots, -ia_n) | a_j \in \mathbb{C}\}$$

and so

$$M_1 \cap M_2^\sigma = \langle (a_1, 0, \ldots, 0; ia_1, 0, \ldots 0) | a_1 \in \mathbb{C}\},$$

a singular polar point p. Then of course p^σ is incident with M_1^σ and M_2. Thus the orbit $p^A = \{p, p^\sigma\}$ is incident with the orbit $F^A = \{M_1, M_2, M_1^\sigma, M_2^\sigma\}$.

(1) Show that p^A is not incident with any object G^A where G is a flag of B properly containing F.
(2) Conclude that

$$\mathrm{Res}_{B/A}(F^A) = \alpha_F(\mathrm{Res}_B(F)) \cup \{p^A\}$$

which (with respect to polar points) is the disjoint union of a $PG(n-2)$ and an isolated point not incident with any subspace of the $PG(n-2)$. So this geometry is not connected.

Chapter 12
Separated Systems of Singular Spaces

Abstract Sometimes in a geometry, a residual of a flag F may exhibit objects A and B that appear to belong to classes of distinct types in the residual, but may in fact belong to one type in some covering geometry. The method of realizing such a covering geometry, due to A. Cohen, is exposed. If the covering geometry is connected, A and B belong to one class; if it is not connected, A and B are objects of distinct types. There is some wrestling with sufficient conditions for the latter choice.

12.1 Introduction

12.1.1 The Basic Context: Paraprojective Spaces

Throughout this chapter, all point-line geometries considered are gamma spaces all of whose singular subspaces are projective spaces. In his survey article in the *Handbook of Incidence Geometry*, Arjeh Cohen has coined the term *paraprojective spaces* for this class of gamma spaces [35]. For all such spaces there are well-defined rank-two[1] geometries at every point p, namely the *point-residuals* $\mathrm{Res}(p) = (\mathcal{L}_p, \Pi_p)$ — the geometry of lines and projective planes which are incident with point p. As we have remarked, the *point residual* $\mathrm{Res}(p)$ itself may be regarded as a point-line geometry whose "Points" are the lines $\mathcal{L}_p$ and whose "Lines" are the planes Π_p which are incident with p. It should be an easy exercise for the student to verify the following:

> *If Γ is a gamma space with projective singular subspaces, having the property that every point lies in a projective plane, then each point-residual Res(p) is itself a gamma space.*

[1] Or smaller rank under the uninteresting pathologies that p lies in no line or no plane.

E.E. Shult, *Points and Lines*, Universitext, DOI 10.1007/978-3-642-15627-4_12,
© Springer-Verlag Berlin Heidelberg 2011

12.1.2 Local and Global Hypotheses

We think of the geometries $\mathrm{Res}(p)$ as "local geometries." As pointed out in the introduction, this entire book is devoted to one aim: to find relatively simple axioms regarding the behavior of points and lines that would allow us to characterize the classical point-line geometries. Obviously axioms that only concern points and lines in the "neighborhood" of a point should be considered more local than axioms which do not.

For example, the famous "Pasch axiom"[2] concerns the behavior of lines transverse to a pair of intersecting lines. Clearly this is a local axiom since its assertions involve only lines entangled with any fixed configuration of two lines meeting at a point (see Chap. 6). In contrast, the powerful "one-or-all" axiom for polar spaces is actually a "global axiom" since it states something about a non-incident point-line pair *no matter where that point might be located relative to the line* (see Chap. 7). Similarly, the near polygon axioms, as well as Cameron's axiom asserting that when quads exist, they are gated in the point-collinearity graph, are both rather global axioms (see Chap. 8). But of course if these "global characterizations" are used in trying to determine subspaces of point-residuals then they are being put to work in a local context as part of a local hypothesis. I hope the theory of parapolar spaces expounded in the next chapter as well as the applications that follow in Part IV will convince the reader that we need to study geometries with global hypotheses if only to make use of them later in a local context.

There are also intermediate degrees of being "local." In the next chapter we shall study the condition of being a parapolar space. One may regard it as a "rather local" hypothesis since its main assumption concerns the behavior of points and lines around pairs of points known to be at distance two in the point-collinearity graph.

But there are limits. As a rule, full knowledge of the point-residuals $\mathrm{Res}(p)$ does not always convey complete knowledge of the geometry Γ. The reason for this disconnect is that for various subspaces T of the geometry Γ there may be no record of one of them containing the point p that can be ascertained from the point-residual p alone.

However, any *singular* subspace of Γ which contains point p leaves a clear imprint in the point-residual geometry $\mathrm{Res}(p)$. For gamma spaces this much is true:

There exists a bijection

$$f : Sing(\Gamma)_p \to Sing(Res(p))$$

between the collection of all singular subspaces that contain point p and the collection of all singular subspaces of the point-residue $Res(p)$.[3]

---

[2] Following a long mathematical tradition, the namesake of this axiom does not seem to have been the one who invented it (see the article of Pickert, Personal communication (1989)).

[3] In the next chapter on parapolar spaces, we will see this correspondence extend to symplecta.

12.1.3 Separated Systems of Singular Subspaces

Recall that the classical Grassmann space (an example introduced in Chap. 2) is a point-line geometry whose points are the k-dimensional subspaces of some vector space V (k is a finite positive integer) and whose lines are the flags (A, B) of $(k-1)$- and $(k+1)$-dimensional vector subspaces. Incidence is the usual incidence of flags in the finitary projective space $PG(V)$. This point-line geometry is a paraprojective space whose maximal singular subspaces are partitioned into two classes of maximal singular subspaces, say $\mathcal{M}_1$ and $\mathcal{M}_2$, such that every line lies in exactly one space of each class (and is in fact the intersection of these as subspaces).

As a result there is a complete separation of the singular subspaces of projective dimension at least two into two distinct classes: (1) those which lie in a member of $\mathcal{M}_1$ and (2) those which lie in a member of $\mathcal{M}_2$. This is visible in any point-residue where $\mathrm{Res}(p) \simeq M_1 \times M_2$, where $M_i \in \mathcal{M}_i, i = 1, 2$. The "vertical" lines in the product correspond to planes in one class while the "horizontal" lines of the product correspond to planes in the other.

In general this happens for Lie incidence geometries for which the points are maximal parabolics associated with a node of degree two in the Dynkin diagram. (It may happen that the two intersecting subspaces are not singular in general.) For example, one might take as Points the lines of a non-degenerate polar space of rank at least four, and take as Lines the collections of polar lines on a point within a polar plane. This geometry of Points and Lines – called a *Polar Grassmannian of lines* – is an example of a geometry of "long-root type," which we shall characterize later on in this book. Here every line is uniquely realizable as the intersection of a plane and a non-degenerate polar space of rank at least three.

In all such "Grassmann" examples where any line is the intersection of exactly two subspaces from two species of subspaces, say $\mathcal{A}$ and $\mathcal{B}$, in the residue of a point p, one sees the lines on p as residual "points" and the two classes of objects $\mathcal{A}_p$ and $\mathcal{B}_p$ appearing in the residue of p simply as two "colorings" of some sort of line. Here is an example of what such a point-residue might look like. It might be a point-thin generalized hexagon. That is, its "points" might be the point-line flags of a projective plane π, and its "lines" are of two types: the points and lines of the projective plane. For example, the residual might consist of the 21 "points" (flags of the Fano plane of order 2) and 14 "lines" of 2 'colours': 7 for the Fano points and 7 for the Fano lines. Is there a point-line geometry whose point-residuals are like this?

So we see that the *global separation of singular subspaces* of rank at least two into two species $\mathcal{A}$ and $\mathcal{B}$ leads to a similar *local separation of singular subspaces* of rank at least one. Is this process reversible? Must a local separation of singular subspaces of rank at least one lead to a global separation of those of rank at least two? If not, when can it? How can local separation at one point be correlated with the local separation at a point collinear with the first? Does any of this apply to separation of subspaces which properly contain lines but are not singular?

In this brief chapter we wish to assemble a few general principles that bear on the following question:

1. Suppose $\mathcal{X}$ is a class of subspaces of Γ, each member of which properly contains a line. Suppose further that the elements X of the set $\mathcal{X}_p$ of members of $\mathcal{X}$ that contain point p are uniquely recognizable as subspaces $X_p = X \cap \mathrm{Res}(p)$ of $\mathrm{Res}(p)$ consisting of all lines and planes of X that are incident with p. (That means the mapping which takes $X \in \mathcal{X}_p$ to X_p is a bijection of the sort that holds for singular subspaces as noted on p. 442 at the end of the previous section.)
2. Finally, suppose that for every point p, we have a partition
$$\mathcal{X}_p = \mathcal{A}_p + \mathcal{B}_p$$
so that if $A \in \mathcal{A}_p$ and $B \in \mathcal{B}_p$ then $A \cap B$ is either $\{p\}$ or is a line.
3. The question: does it follow that there is a corresponding global partition
$$\mathcal{X} = \mathcal{A} + \mathcal{B}$$
such that, for every point p, the members of $\mathcal{A}$ and $\mathcal{B}$ which are incident with p are the sets $\mathcal{A}_p$ and $\mathcal{B}_p$ respectively?

In general, the answer is "no," but there are many cases where the answer to the question is "yes" and we wish to describe conditions for which this affirmative answer can be realized.

12.2 Geometries with Two Systems of Subspaces Meeting at Lines

12.2.1 Introduction

As we encounter a new type of theorem, it is always good to examine the general context introduced by the novel hypotheses of that theorem. The objective of this section is to show that when the local geometry induced at a point *looks* as though it came from a geometry with two universally distinguished classes of geometries with restricted intersections, it is in fact at worst a two-to-one homomorphic image of such a geometry.

12.2.2 Geometries with Two Systems of Subgeometries

We begin with a hypothesis first introduced by G. Tallini [127] which involves the interaction of two classes of subspaces.[4] Suppose $\Gamma = (\mathcal{P}, \mathcal{L})$ is a geometry with two systems of subgeometries Σ and Π with the properties that:

[4] See hypotheses (T1) and (T2) of Theorem 14.1.1, p. 498, where Σ and Π are maximal singular subspaces of a gamma space.

(S1) Γ *is a connected point-line geometry.*

(S2) (a) *Any two subspaces* $(A, B) \in \Sigma \times \Pi$ *intersect at the empty set or a line.*

(b) *Every line lies properly in exactly one subspace from* Σ *and one subspace from* Π.

We must also consider a "local" version of this. Consider the following hypothesis.

(R1) Γ' *is an incidence system* $(\mathcal{A}, \mathcal{P}', \mathcal{B})$.

(R2) (a) *For any* $(A, B) \in \mathcal{A} \times \mathcal{B}$, $A \cap B$ *is empty or a point.*

(b) *Any "point"* $p \in \mathcal{P}'$ *lies in exactly one member of* $\mathcal{A}$ *and in exactly one member of* $\mathcal{B}$*, and is properly contained in each case. In particular, the subspaces of any one of the families – say* $\mathcal{A}$ *– partitions the points.*

Example 1 Suppose Γ' is a *product geometry* $A \times B$. Then $\mathcal{A}$ can be the collection of subspaces $\{(A, x) | x \in B\}$ while $\mathcal{B}$ can be be the subspaces $\{(y, B) | y \in A\}$. Then $(\Gamma', \mathcal{A}, \mathcal{B})$ satisfies axioms (R1) and (R2).

Example 2 There are other examples which are not products. Consider a projective plane $\pi = (\mathcal{P}, \mathcal{L})$, and let Γ' be the geometry whose point set $\mathcal{P}'$ is the set of flags of π, while $\mathcal{A}$ and $\mathcal{B}$ are the points and lines of Π, respectively, with the usual incidence of flags. Then the axioms (R1) and (R2) hold and $(\mathcal{A}, \mathcal{P}', \mathcal{B})$ is certainly connected.

In fact, one could take any rank two chamber system $C = (C, E, \lambda, I = \{1, 2\})$, set $\mathcal{P}' := C$, the set of chambers, and then set $\mathcal{A}$ and $\mathcal{B}$ to be the rank-one residues of types 1 and 2, respectively. Again the (R)-axioms are in effect.

Now what is the relationship between these two sets of axioms – the (S)-axioms and the (R)-axioms?

Suppose, to start with, $\Gamma = (\mathcal{P}, \mathcal{L})$ is a point-line geometry with two families of subspaces Σ and Π satisfying the axioms (S1) and (S2).

Now, for any point p of Γ, let $\mathcal{L}_p$ be all lines of Γ incident with p and let Σ_p and Π_p be the geometries in the classes Σ and Π which are incident with point p. With a slight abuse of notation (due only to the lumping of classes) we set

$$\mathrm{Res}_\Gamma(p) := (\mathcal{L}_p, \Sigma_p \cup \Pi_p),$$

a rank two (point-line) geometry which we call the *point-residual of* Γ *at point* p. Then $\mathrm{Res}_\Gamma(p)$ satisfies the axioms (R1) and (R2) with $\mathcal{P}' = \mathcal{L}_p$, $\mathcal{A} = \Sigma_p$ and $\mathcal{B} = \Pi_p$.

Thus a point-line geometry with two families of subspaces satisfying the (S)-axioms has all of its point-residuals satisfying the (R)-axioms. But what about the reverse direction?

Suppose:

(L) $\Gamma = (\mathcal{P}, \mathcal{L})$ *is a connected point-line geometry possessing a family* Ξ *of subspaces such that for each point* p*, there is a partition* $\Xi_p = \mathcal{A}_p + \mathcal{B}_p$ *such that* $(\mathcal{A}_p, \mathcal{L}_p, \mathcal{B}_P)$ *satisfies conditions (R1) and (R2).*

Must there exist a *global partition* $\Xi = \mathcal{A} + \mathcal{B}$ inducing the partition on all the point-residuals? The answer is "no," but it is almost true.

The principle result of this section is the following "Covering theorem." The basic idea and the proof is due to A. M. Cohen in [34].

Theorem 12.2.1 (Covering theorem.) *Assume* Ξ *is a class of subspaces of the point-line geometry* $\Gamma = (\mathcal{P}, \mathcal{L})$ *subject to the axiom (L). Then either:*

1. *There exists a global partition* $\Xi = \mathcal{A} + \mathcal{B}$ *of the family of subspaces inducing the local partition* $\Xi_p = \mathcal{A}_p + \mathcal{B}_p$ *so* Γ *itself satisfies axioms (S1) and (S2), or*
2. *There exists a morphism*

$$\phi : \hat{\Gamma} = (\hat{\mathcal{P}}, \hat{\mathcal{L}}) \to (\mathcal{P}, \mathcal{L}) = \Gamma,$$

of point-line geometries, where $\hat{\Gamma}$ *possesses two families of subspaces* $\hat{\Sigma}$ *and* $\hat{\Pi}$ *so that* $(\hat{\Sigma}, \hat{\mathcal{P}}, \hat{\Pi})$ *satisfies axioms (S1) and (S2) and for each point* p *of* Γ*,* $\phi(\hat{\Sigma})_p = \mathcal{A}_p$ *and* $\phi(\hat{\Pi})_p = \mathcal{B}_p$*.*
The morphism is surjective on points and on lines and the fiber above each of these objects is exactly two points or two lines as appropriate.

Proof From axiom (L), for each point p we have a decomposition $\Xi_p = \mathcal{A}_p + \mathcal{B}_p$. Suppose $p \in X \in \Xi$. Then we will let $C_p(X)$ denote the class ($\mathcal{A}_p$ or $\mathcal{B}_p$) which contains X, and let $C'_p(X)$ denote the class of Ξ_p which does *not* contain X.

Now we are going to make a graph $(\Xi, \sim)$ in which two subspaces X and Y of Ξ are adjacent if and only if they intersect exactly at a point p, *and* X and Y belong to the same class of Ξ_p – i.e., $C_p(X) = C_p(Y)$. There are then two cases: (1) $(\Xi_p, \sim)$ is not connected, or (2) it is connected.

In the former, the elements of $\mathcal{A}_p$ and $\mathcal{B}_p$ represent two cliques in different connected components of $(\Xi_p, \sim)$. Writing the decomposition of this graph as $\Xi = \mathcal{A} + \mathcal{B}$, one obtains the conclusion of 1.

Suppose now that case (2) holds, so $(\Xi, \sim)$ is connected. Let $\hat{\mathcal{P}}$ be the collection of all pairs (p, C), where p is a point of $\mathcal{P}$ and C is one of the two classes, $\mathcal{A}_p$ or $\mathcal{B}_p$, of Ξ_p. Let L be a line of $\mathcal{L}$ belonging to a subspace X of Ξ (given L, there are just two choices for X). Then the sets

$$\{(x, C_x(X) | x \in L\}$$

comprise the elements of $\hat{\mathcal{L}}$, and define a point-line geometry $\hat{\Gamma} = (\hat{\mathcal{P}}, \hat{\mathcal{L}})$ with membership as incidence. We define the point mapping $\phi : \hat{\mathcal{P}} \to \mathcal{P}$ by the "first coordinate" equation:

$$\phi((p, C)) = p.$$

Then for any line $A = \{(p, C_p(X))|x \in L\}$ where L is some fixed line of subspace $X \in \Xi$, one sees that $\phi(A) = L$.

For any $X \in \Xi$, let

$$\Sigma_X := \{(x, C_x(X))|x \in X\}$$
$$\Pi_X := \{(x, C'_x(X)|x \in X\}.$$

Then clearly, $\Sigma := \{\Sigma_X|X \in \Xi\}$ and $\Pi := \{\Pi_X|X \in \Xi\}$ are two families of subspaces of $\hat{\Gamma}$, each indexed by Ξ, and $\phi(\Sigma_X) = \phi(\Pi_X) = X$. All parts of the second conclusion are in place. □

12.2.3 A Purely Local View

The comment here is just a footnote to the preceding parts of this section. Even to state axiom (L), one has a global family Ξ of subspaces. Could this family be inferred from local information alone?

To set the scene, we suppose that $\Gamma = (\mathcal{P}, \mathcal{L})$ is a connected point-line geometry and that for every point p we can find a family Ξ_p of subspaces containing p, and a partition $\Xi_p = \mathcal{A}_p + \mathcal{B}_p$, such that with respect to the residual $\mathcal{L}_p$ of all lines on p, $(\mathcal{A}_p, \mathcal{L}_p, \mathcal{B}_P)$ satisfies axioms (R1) and (R2).

Now it is transparent that one can recover Ξ provided the following applies.

(Connecting axiom.) *If X is a subspace in Ξ_p, and q is a point of X which is collinear with p, then $X \in \Xi_q$.*

One way to realize this connecting axiom is by having the elements of Ξ_p being subspaces of Γ which are maximal with respect to having some property. For example, if Γ is a gamma space and if the members of Ξ_p are maximal cliques of Γ, then the elements of Ξ_p are recognizable when passed on to another local environment Ξ_q. All in all, a rather prosaic footnote.

12.3 A Criterion for Global Separation of Singular Subspaces

12.3.1 Examples Illustrating Local Separation Without Global Separation

First let us discuss a few cases where local separation of singular subspaces does not imply a global separation of singular subspaces. Suppose V is a vector space of even dimension $2n$ over a division ring K and $f : V \times V \to K$ is a non-degenerate reflexive sesquilinear form (see Chap. 7 for the definitions of these terms). Then there is a well-defined polarity $\phi : \mathbf{P}(V) \to \mathbf{P}(V^*)$ induced by the mapping which takes each k-dimensional vector subspace U to the $2n - k$-dimensional vector subspace $U^\perp := \{v \in V | f(v, U) = 0\}$.

We next restrict our attention to the Grassmann space $G = (\mathcal{P}', \mathcal{L}')$ of type $A_{2n-1,n}$ whose points are the n-dimensional vector subspaces of V and whose lines are the $(n-1, n+1)$-(vector-space)-dimensional flags. Then by restriction ϕ induces an automorphism α of the Grassmann space G. Since f is non-degenerate, one has $U^{\perp\perp} = U$ for every n-space U, and so α is an involution – that is α^2 is the identity automorphism while α is not.

The distance between a point U and its image $\alpha(U) = U^{\perp}$ in the point-collinearity graph of G is the codimension of $U \cap U^{\perp}$ in U. The maximal vector-space dimension of a totally isotropic subspace of (V, f) is uniform for all maximal isotropic subspaces and is called the *Witt index* of the form (V, f). Since f is non-degenerate this number cannot exceed n. Letting w denote the Witt index of the form f, one sees that the minimal distance between a Grassmann point p and its image $\alpha(p)$ is $n - w$.

Now we can follow the recipe in Chap. 2, p. 52, for forming a homomorphism:

$$\rho : G \to \Gamma = (\mathcal{P}, \mathcal{L})$$

where ρ replaces each object by its $\langle\alpha\rangle$-orbit. In this case the orbits have length at most two since α is an involution. In Γ a point orbit $p = \{U, U^{\perp}\}$ is incident with a line orbit $L = \{(A, B), (B^{\perp}, A^{\perp})\}$ if and only if U is incident with at least one of the flags in L. There is another important property of this morphism: it induces a *full morphism* of the bipartite incidence graph. In this case, it means the following:

1. The mapping $\rho : G \to G/\langle\alpha\rangle = \Gamma$ is surjective on both points and lines.
2. If (p, L) is a flag in the image geomtry Γ, then:
   (a) For any point p' for which $\rho(p') = p$, there exists at least one line L' of G incident with p' such that $\rho(L') = L$.
   (b) For any line L' of G for which $\rho(L') = L$, there exists a point p' in the point-shadow of L' in G such that $\rho(p') = p$ – that is, when ρ is restricted to line L', it surjectively maps the point shadow of L' onto the point shadow of its image line L.

These statements follow directly from the way that the morphism is defined by forming $\langle\alpha\rangle$-orbits. (These facts are mentioned in Chap. 2, in the section on morphisms of geometries on p. 52.)

One can see that the morphism ρ is point-injective when restricted to lines (or any singular subspace of G for that matter) if the Witt index w is at most $n - 2$ and induces a fibering morphism if $w \leq n - 3$ (Exercise 12.3).[5]

So if $w \leq 3$, and point $p \in \mathcal{P}$ is an image $\rho(p')$ of a point $p' \in \mathcal{P}'$, then the restrictions of ρ to $p'^{\perp}$ induces a complete isomorphism $PS(p') \to PS(p)$ of the poset $PS(p')$ of all singular subspaces of G on point p' with the poset $PS(p)$ of all singular subspaces of Γ which are incident with point p. Since the point-residues

[5] In general, the lower the Witt index becomes, the larger the isomorphically preserved neighborhood of a point can become.

$\mathrm{Res}(p'))$ and $\mathrm{Res}(p)$ (taken in their respective geometries) are composed of lines and planes, and these lines and planes are singular subspaces, the isomorphism of posets asserted in the previous sentence implies that the point-residuals are isomorphically preserved. That means that in both geometries we have

$$\mathrm{Res}(p') \simeq \mathrm{Res}(p) \simeq PG(n-1, K) \times PG(n-1, K).$$

That means that in a point-residual of Γ, we have a complete separation of singular subspaces into those with only "horizontal lines" and those with only "vertical lines" – that is, a local separation of singular subspaces.

But Γ does not possess a corresponding global separation of singular subspaces. To see why this is so, let us examine carefully what we mean by a "global class" of subspaces. Fix a point p of Γ. We have just remarked that $\mathrm{Res}(p) \simeq M_1 \times M_2$ where each M_i is a maximal singular subspace of the geometry Γ isomorphic to $PG(n-1, K)$. Each of the spaces (M_1, y) in $M_1 \times M_2$ corresponds to a maximal singular $PG(n, K)$ on the point p, and the "parallel class" of such spaces thus forms a family $\mathcal{A}_p$ of $PG(n, K)$s pairwise intersecting at $\{p\}$. (Of course there is a similar family $\mathcal{B}_p$ of $PG(n, K)$s pairwise intersecting at p which produce the other class of singular subspaces of shape (x, M_2) in the point residual at p. Every member $\mathcal{B}_p$ intersects every member of $\mathcal{A}_p$ at a line on p.)

We shall define a "global class" $\mathcal{A}$ which contains every member of $\mathcal{A}_p$. Now take any maximal singular subspace M in $\mathcal{A}_p$ and choose any second point $q \in M - \{p\}$. Then A is a member of one of the two classes of maximal singular subspaces $\mathcal{A}_q$ and $\mathcal{B}_q$ which contain q. Whichever class it is, all members of that class must be incorporated into the global class $\mathcal{A}$. We keep iterating this process forever.

Of course there is a somewhat simpler way to describe this. Let $\mathcal{M}$ be the full class of maximal singular subspaces of Γ (all are $PG(n, K)$s). Define a graph $\mathcal{Y}$ with vertex set $\mathcal{M}$ where we say that two members A and B of $\mathcal{M}$ are *adjacent* if and only if they intersect at a single point. Then $\mathcal{A}$ can simply be defined as the connected component of $\mathcal{Y}$ which contains the "vertex" A.

The problem with Γ is that $\mathcal{A}$ contains $\mathcal{B}_p$ as well as $\mathcal{A}_p$. That means $\mathcal{A} = \mathcal{M}$ – all maximal singular subspaces. Thus in Γ we have *local* separation of singular subspaces without a corresponding *global* separation of singular subspaces.

Of course one sees this situation is feasible because the maximal singular subspace A inducing a direct factor in the product representation of $\mathrm{Res}(p)$ in Γ can have a member in its class $\mathcal{A}$ which also represents the other member. Occasionally students seeing this example for the first time leap to the conclusion that this could not happen if $\mathrm{Res}(p)$ always had the form $A_p \times B_p$ where B_p is not isomorphic to A_p. That conclusion is entirely unwarranted.

Suppose, for example, that $A_p \simeq PG(n-1, K)$ but $B_p \simeq L \times PG(n-1, K)$. Something like this happens in the Lie incidence geometry $G := E_{6,4}$ in which a point-residual has the form

$$\mathrm{Res}(p) = PG(2, K) \times L \times PG(2, K)$$

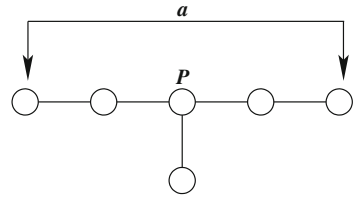

Fig. 12.1 A homomorphic image of the twisted $E_{6,2}$ Lie incidence geometry

where L is a line and K is a field. For suitable fields K, there is an automorphism of the building which is not type-preserving but induces a symmetry of order two of the E_6 diagram as indicated in Fig. 12.1.

One observes that this automorphism induces a polarity on the $A_{5,3}$ diagram corresponding to a sesquilinear form of Witt index zero. Such an automorphism would then induce an involutatory automorphism α of the Lie incidence geometry Γ with the property that, for any point p, the distance between p and $\alpha(p)$ is at least three in the point-collinearity graph. Then, forming the canonical full epimorphism $G \to \Gamma := G/\langle\alpha\rangle$ we see that there is an induced isomorphism on the poset of singular subspaces on a point and the same poset on its image – just as we argued above. Then in Γ there is a local separation $\mathcal{A}_p + \mathcal{B}_p$ of the $PG(3, K)$s which contain point p, but again there is not a global separation.

Now we begin to realize that there may be many geometries Γ for which there is a local separation of singular subspaces without necessarily inheriting a global one. Is there any way for the point-residual itself to tell us that a local separation of singular subspaces signals a global one?

As noted, when the global separation fails, one must find singular subspaces of the global class $\mathcal{A}$ containing all the elements of A_p also entering the class B_p. That means that there is a singular subspace X on a line pq which seems to belong to $\mathcal{B}_p$, while still belonging to $\mathcal{A} \cap \mathrm{Res}(q)$. We are trying to watch the same objects from two perspectives and that certainly raises the question of how much $\mathrm{Res}(p)$ and $\mathrm{Res}(q)$ share? The answer is not much: only what can be perceived in the poset $P(pq)$ of all singular subspaces of Γ which contain the line pq.

This should tell one immediately that one can only hope for sufficient conditions for global separation based on the poset $P(pq)$.

12.3.2 A Sufficient Condition for Local Separation to Imply Global Separation

The best that the author can do is as follows.

Theorem 12.3.1 *Suppose Γ is a gamma space whose singular subspaces are projective spaces (that is, Γ is a paraprojective space). Suppose:*

1. *(The local separation hypothesis.) For each point p, the maximal singular subspaces of Γ which contain p are partitioned into two classes $\mathcal{A}_p$ (each of which is a $PG(n, K)$, where n is a (not necessarily finite) cardinal number not depend-*

ing on p) and $\mathcal{B}_p$ with the property that (a) each line on p lies in exactly one member of class $\mathcal{A}_p$ and at least one member of class $\mathcal{B}_p$ and (b) any member of $\mathcal{B}_p$ intersects a member of $\mathcal{A}_p$ in $\{p\}$ or a line on p.

2. *(The sufficient condition.) For any element B of $\mathcal{B}_p$ which happens to be isomorphic to $PG(n, K)$, and any line L of B on p, there is at least one other maximal singular subspace B' in $\mathcal{B}_p$ which contains L but intersects B at a subspace properly containing L.*

Then global separation of singular subspaces ensues: that is, there is a partition $\mathcal{M} = \mathcal{A} + \mathcal{B}$ of all maximal singular subspaces of Γ so that for every point p, those members of $\mathcal{A}$ which are incident with p are exactly the members of $\mathcal{A}_p$, and those members of $\mathcal{B}$ which are incident with p are exactly the members of $\mathcal{B}_p$.

Proof The proof is very short. For each point-line flag (p, L) of Γ let $A(p, L)$ be the unique member of $\mathcal{A}_p$ on line L. Let q be a point of $L - \{p\}$. We claim the following.

(Step 1) $A(p, L) = A(q, L)$.

Suppose this were false. Then by 1 we see that $A(q, L)$ is a $PG(n, K)$ in $\mathcal{B}_p$. By 2 there exists a maximal singular subspace $B' \in \mathcal{B}_p - \{A(q, L)\}$ such that B' contains line L but intersects $A(q, L)$ at a space properly containing L. But then hypothesis 1 (a) applied to point q (in place of p) forces $B' \in \mathcal{B}_q$ while hypothesis 1 (b) (applied to q) provides the contradiction that $B' \cap A(q, L)$ cannot properly contain L.

(Step 2) *If N is a line on p in the space $A(p, L)$ then $A(p, N) = A(p, L)$.*

This is immediate from 1(a).

(Step 3) *If (q, N) is any flag of the singular subspace $A(p, L)$ then $A(q, N) = A(p, L)$.*

This follows from Steps 1 and 2.

(Step 4) *Setting $\mathcal{A} := \cup\{\mathcal{A}_p | p \in \mathcal{P}\}$ and $\mathcal{B} := \cup\{\mathcal{B}_p | p \in \mathcal{P}\}$, we have the following:*

1. *Any two members of $\mathcal{A}$ intersect at the empty set, or a point.*
2. *Any member of $\mathcal{A}$ intersects any member of $\mathcal{B}$ in the empty set, a point or a line.*
3. *For any point p, the elements of $\mathcal{A}$ which are incident with p are exactly $\mathcal{A}_p$.*

All three parts follow from Step 3.

The proof is complete. □

12.4 Exercises for Chapter 12

12.1 Suppose Γ is paraprojective space (recall the definition on p. 450) whose point-residuals $\mathrm{Res}(p)$ are all of one of these isomorphism types:

1. $PG(n, K) \times PG(m, K), m \neq n$,
2. $PG(a, K) \times PG(b, K) \times PG(b, K)$ where a, b, c are distinct cardinal numbers.
3. $PG(n, K) \times S$ where S is a non-degenerate polar space of rank at least 3. Here, the cardinal number n and the polar rank of S may be infinite.
4. $L \times GQ$ where L is a line and GQ is any generalized quadrangle.

All of these cases exhibit local separation of singular subspaces. Which of them satisfy the sufficient condition of Theorem 12.3.1?

12.2 Let $G = (\mathcal{P}, \mathcal{L})$ be the Grassmann space whose points $\mathcal{P}$ are the n-dimensional vector subspaces of a $2n$-dimensional vector space V, and whose lines are the incident pairs of vector subspaces (A, B) of V where $\dim A = n - 1$ and $\dim B = n + 1$.
Show that two n-dimensional vector subspaces U_1 and U_2 represent points at distance k in the point-collinearity graph of G if and only if $\dim(U_1/(U_1 \cap U_2) = k$.

12.3 Suppose $\alpha : P(V) \to P(V^{\#})$ is a non-degenerate polarity defined by a non-degenerate reflexive sesquilinear form on V (see Chap. 7 for definitions). If A is any vector subspace of V, we write $A^{\perp}$ for A^{α}, and we say that subspace A is a *singular subspace* if and only if $A \subseteq A^{\perp}$. The *Witt index* of α is the maximal possible vector-space dimension of a singular subspace.

1. Show that the Witt index w is the minimal distance between a Grassmann point U and its image U^{α}.
2. Let $\bar{\mathcal{P}} := \{(U, U^{\perp}) | U \in \mathcal{P}\}$. Let $\bar{\mathcal{L}} := \{(L = (A, B), L^{\alpha} = (B^{\perp}, A^{\perp}) | L \in \mathcal{L}\}$. The pair $(U, U^{\perp}) \in \bar{\mathcal{P}}$ is declared to be incident with (L, L^{α}) if and only if Grassmann point U is incident with at least one of the Grassmann lines L or L^{α}. In this way, we obtain a point-line geometry $\Gamma = (\bar{\mathcal{P}}, \bar{\mathcal{L}})$. There is an obvious morphism of point-line geometries:

$$\rho G = (\mathcal{P}, \mathcal{L}) \longrightarrow \Gamma = (\bar{\mathcal{P}}, \bar{\mathcal{L}})$$

taking points and lines of G to their α-orbits.

   (a) Show that if $w \leq n - 2$ then for each point $\bar{p} \in \bar{\mathcal{P}}$, no two Grassmann points of the fiber $\rho^{-1}(\bar{p})$ are collinear.

(b) Under the same conditions show that if $\bar{L} \in \bar{\mathcal{L}}$ then no two Grassmann lines of $\rho^{-1}(\bar{L})$ are incident with a common Grassmann point.

(c) Now suppose $w \leq n - 3$. Show that ρ is injective when restricted to the set of points in $\mathcal{P}$ that are collinear with a given Grassmann point. Similarly show that ρ is then injective on the set of Grassmann lines which intersect a given Grassmann line.

Chapter 13
Cooperstein's Theory of Symplecta and Parapolar Spaces

Abstract B. Cooperstein gave a criterion that symplecta (convex polar subspaces) exist. That theory is developed here in the somewhat wider context of a "polar family." Elementary properties of parapolar spaces are discussed – for example, those of symplectic rank at least three have all their singular subspaces projective. In particular, local connectedness is used to show that symplecta possess deep points only if the whole parapolar space is just one symplecton. Hanssen's principle asserts that a parapolar space with a certain general type of locally-truncated diagram must also possess a class of Grassmannians. In this way, some rank three locally-truncated geometries can be enriched to rank four geometries so that the sheaf-theory of Chap. 11 kicks in.

13.1 Introduction

Suppose $\Gamma = (\mathcal{P}, \mathcal{L})$ is any point-line geometry. A convex subspace S of Γ is called a *symplecton* if and only if it (together with its internal lines) is a non-degenerate polar space of rank at least two. (It may have infinite polar rank.)

In Chap. 8 we met near polygons with quads (for example, the classical dual polar spaces of rank at least two). In that case the quads themselves are symplecta for, by Yanushka's lemma, they are convex subspaces which are generalized quadrangles.

Also, like Yanushka's lemma, there is a theory due to Cooperstein which explains why symplecta must exist in many of the Lie incidence geometries. As will gradually become apparent, the symplecta, more than anything else, seem to play a diagnostic role in the characterizations of the Lie incidence geometries.

13.2 Polar Families

Let $\Gamma = (\mathcal{P}, \mathcal{L})$ be a gamma space. A *polar family* is a collection Σ of non-collinear pairs of distinct points (2-cocliques) such that:

1. If (x, y) is an element of Σ then $x^{\perp} \cap y^{\perp}$ is a non-degenerate polar space of rank at least two.

E.E. Shult, *Points and Lines*, Universitext, DOI 10.1007/978-3-642-15627-4_13,

© Springer-Verlag Berlin Heidelberg 2011

2. If (x, y) is an element of Σ and (u, v) is a pair of distinct non-collinear points of $x^\perp \cap y^\perp$, then (u, v) is also in Σ.

Our goal in this section is to prove, that if Σ is a polar family in a connected gamma space with thick lines, then for any pair $(x, y) \in \Sigma$, the *convex closure* $\langle\langle x, y\rangle\rangle$ (that is, the smallest convex subspace containing $\{x, y\}$) is a symplecton. This important result is essentially a more mobile version of an original argument due to Cooperstein [41].[1]

For this purpose, we assume the following hypothesis for the remainder of this section:

(Coo) 1. $(\mathcal{P}, \mathcal{L})$ is a connected gamma space with thick lines.
2. Σ is a polar family.

Lemma 13.2.1 (The ladder lemma) *Let A and B be two lines for which $A \cap B^\perp$ and $A^\perp \cap B$ are both empty. Assume further that there are at least two distinct points x and y of A, collinear with respective points x' and y' of B. If at least one of the pairs (x, y') or (y, x') belongs to Σ then:*

1. *There is a one-to-one correspondence between the points of A and the points of B with corresponding pairs representing all collinear pairs in $A \times B$.*
2. *Every pair of non-collinear points in $A \cup B$ belongs to Σ.*

Proof From the hypotheses of the lemma, we must have x' and y' distinct. Because of the symmetry of x and y in the hypothesis, we may assume $(x, y') \in \Sigma$. Then by the definition of polar family, $x^\perp \cap (y')^\perp$ is a non-degenerate polar space of rank at least two containing the non-collinear pair (x', y). Thus $x^\perp \cap (y')^\perp \cap (x')^\perp \cap y^\perp$ is a non-degenerate polar space (possibly of rank one) and so contains a non-collinear pair (u, v) which is necessarily a member of Σ. Then $A \cup B$ lies in the polar subspace $u^\perp \cap v^\perp$ and so both conclusions 1 and 2 hold. □

Now we fix a pair (a, b) in Σ. We define the set $S(a, b)$ to be the full set of points p such that for *every* line L of $a^\perp \cap b^\perp$, p is collinear with at least one point of L.

[1] Cooperstein's orginal theorem was presented in the context of strong parapolar spaces, and was later generalized by Buekenhout [21] (with a correction by Surowski David Surowski, Personal Communication (1982)) to the more general context of what he called "polarized spaces." The latter are essentially what we now call "parapolar spaces" – which shall be introduced later in this chapter. They provide a natural platform on which to introduce almost all Lie incidence geometries of spherical type – so it was important to have Cooperstein's theory in this more general context. However, there are Lie incidence geometries of affine buildings and other buildings which are not parapolar spaces, yet symplecta seem to exist for reasons similar to Cooperstein's theory. This observation led the author to present cooperstein's theory in terms of polar families, independent of any particular global context of parapolar spaces or some generalization thereof. This way, it is not necessary to axiomatize fully the ambient space at all; polar families can be identified in a variety of contexts. At the same time, only minor alterations in Cooperstein's original arguments appear here – so this is still his theory. It is just being packaged a little differently.

Remark If x is in $S(a, b)$ then either $x^\perp$ contains $a^\perp \cap b^\perp$ or else $x^\perp \cap (a^\perp \cap b^\perp)$ is a geometric hyperplane of the non-degenerate polar space $a^\perp \cap b^\perp$. In the latter case $x^\perp \cap (a^\perp \cap b^\perp)$ is either (1) a non-degenerate polar space of rank r or $r-1$ where r is the rank of the polar space $a^\perp \cap b^\perp$ or (2) the set $p^\perp \cap a^\perp \cap b^\perp$ for some point p in $a^\perp \cap b^\perp$.

Lemma 13.2.2 *1. Let u be a point of $S(a, b) - (a^\perp \cap b^\perp)$ and v a point of $a^\perp \cap b^\perp$ not collinear with u. Then $(u, v) \in \Sigma$.*
2. If $u \in S(a, b) - a^\perp$, then $(a, u) \in \Sigma$.

Proof 1. Choose two lines C and D on point v and lying in $a^\perp \cap b^\perp$ so that C is not in $D^\perp$ (this is possible). Then setting $\{x\} = u^\perp \cap C$ and $\{y\} = u^\perp \cap D$, one sees that x and y are not collinear and that (x, y) is a member of Σ. Then $x^\perp \cap y^\perp$ contains u and v, whence (u, v) lies in Σ, by one more invocation of the polar family axioms.

2. One can find a non-collinear pair (x, y) of distinct points drawn from $a^\perp \cap b^\perp \cap u^\perp$ belonging to Σ. Since $\{a, u\} \subseteq x^\perp \cap y^\perp$, we must have $(a, u) \in \Sigma$. □

Lemma 13.2.3 *$S(a, b)$ is a subspace of $\Gamma = (\mathcal{P}, \mathcal{L})$ lying in the convex closure of $\{a, b\}$.*

Proof Let X be a line meeting $S(a, b)$ at distinct points x and y, and let z be a third point on X. Suppose, by way of contradiction, that z were a point of X not in $S(a, b)$. Then there must exist a line L of $a^\perp \cap b^\perp$ for which $z^\perp \cap L = \emptyset$.

It follows that $X^\perp \cap L = \emptyset$. But from the definition of $S(a, b)$, $x^\perp \cap L = \{x'\}$, $y^\perp \cap L = \{y'\}$, where x' and y' are distinct. But then by Lemma 13.2.2 $(x, y') \in \Sigma$. Then Lemma 13.2.1 gives a one-to-one correspondence $X \to L$ whose corresponding pairs are the collinear pairs in $X \times L$. Thus z is collinear to a point of L, and that is a contradiction. □

Lemma 13.2.4 *If A is a line on point a meeting $b^\perp$ then $y^\perp \cap A$ is non-empty for any point y in $S(a, b)$.*

Proof Let u be the point of A in $b^\perp$. Then choose C and D on u and lying in $a^\perp \cap b^\perp$ so that C is not in $D^\perp$. As we may assume y is not collinear with u, we obtain a pair $(c, d) \in (C \cap y^\perp) \times (D \cap y^\perp)$, belonging to Σ. Then the polar space $c^\perp \cap d^\perp$ contains y and A, and so we are done. □

Lemma 13.2.5 *For any point y in $S(a, b)$ and any line A of $S(a, b)$ on point a, $y^\perp \cap A$ is non-empty.*

Proof By Lemma 13.2.4 it suffices to show that $b^\perp \cap A$ is non-empty. If $A^\perp \supseteq a^\perp \cap b^\perp$, A and b lie within the polar space $u^\perp \cap v^\perp$ for any non-collinear pair (u, v) in $a^\perp \cap b^\perp$. Thus we may choose u in $a^\perp \cap b^\perp - A^\perp$ and select a line B on u and b. For any point y of $A - \{a\}$, $y^\perp \cap B$ is non-empty by Lemma 13.2.4 with (a, B) in the role of (b, A). Since u is not in $A^\perp$, $y^\perp \cap B := \{y'\} \neq \{u\}$. By Lemma 13.2.2, $(y, u) \in \Sigma$ and so Lemma 13.2.1 shows $b^\perp \cap A$ non-empty as required. □

Lemma 13.2.6 *If A is a line of $S(a, b)$, then a is collinear with a point of A.*

Proof Choose u, v non-collinear in $a^\perp \cap b^\perp$. Let L be a line on a and u. By Lemma 13.2.4, every point of A is collinear with a point of L, and we are done by Lemmas 13.2.1 and 13.2.2 unless we have a point u' in $L \cap A^\perp$. Similarly, if M is a line on a and v, we must suppose there is a point v' in $M \cap A^\perp$. Now, by Lemma 13.2.2, part 2, (u', v') is in Σ, and the polar space $(u')^\perp \cap (v')^\perp$ contains point a and line A and the conclusion follows. □

Lemma 13.2.7 *If A is a line of $S(a, b)$ not lying in $a^\perp$, then for any x in $S(a, b) \cap a^\perp$, x is collinear with a point of A.*

Proof Let N be a line of $S(a, b)$ on a and x. By Lemma 13.2.6, point a is collinear with a point of A. By Lemma 13.2.5, every point of A is collinear with a point of N. Since A is not in $A^\perp$, Lemmas 13.2.4 and 13.2.1 set up a one-to-one correspondence of points of A and N under the collinearity relation. The conclusion follows. □

Lemma 13.2.8 *The condition that A does not lie in $A^\perp$ in Lemma 13.2.7 may be dropped.*

Proof Let x be the given point in $S(a, b) \cap a^\perp$, and as above, let N be a line of $S(a, b)$ on a and x. We wish to show that x is collinear with a point of A.

By Lemma 13.2.7 we may assume that A is contained in $a^\perp$, and, by the definition of $S(a, b)$ that A is not contained in $b^\perp$. By Lemmas 13.2.6 and 13.2.4 there is a point y in $N \cap b^\perp$. Then applying Lemma 13.2.7 with b in the role of a, y is collinear with a point u of A. Then $N \subseteq u^\perp$, and so x is collinear with u on A. □

Lemma 13.2.9 *Suppose $x \in a^\perp \cap S(a, b)$ and $y \in S(a, b) - x^\perp$. Then $(x, y) \in \Sigma$.*

Proof If $x \in a^\perp \cap b^\perp$, the result would follow from Lemma 13.2.2 part 1. Similarly, the conclusion would follow from Lemma 13.2.2 part 2, if $x = a$. So we assume that $x \in a^\perp \cap S(a, b) - (\{a\} \cup b^\perp)$. By Lemma 13.2.8, $b^\perp$ contains a point x_1 on a line N containing $\{a, x\}$, and $x \neq x_1$. Then $x^\perp \cap a^\perp \cap b^\perp = {x_1}^\perp \cap a^\perp \cap b^\perp$, a hyperplane. Obviously we are done if $x^\perp \cap y^\perp \cap a^\perp \cap b^\perp = {x_1}^\perp \cap y^\perp \cap a^\perp \cap b^\perp$ contains a non-collinear pair (u, v) since then $\{x, y\} \subseteq u^\perp \cap v^\perp$ where $(u, v) \in \Sigma$. Thus we see *$x^\perp \cap a^\perp \cap b^\perp$ and $y^\perp \cap a^\perp \cap b^\perp$ are distinct hyperplanes of the polar space $a^\perp \cap b^\perp$ meeting at a singular subspace.*

It follows that there is a point $w \in y^\perp \cap a^\perp \cap b^\perp - x^\perp$. Then by Lemma 13.2.8, if E is a line on y and w, there exists a point z in $x^\perp \cap E$ which is distinct from both w and y.

Now if z is not collinear with a, we see that $a^\perp \cap E = \{w\}$ exactly, and this is distinct from z. Now letting N be a line on a and x, as before, we see that line N is "opposite" E in the sense of Lemma 13.2.1. Since, in addition, Lemma 13.2.2, part 1 gives $(a, z) \in \Sigma$, Lemma 13.2.1 just quoted gives $(x, y) \in \Sigma$ as well. Since that is our desired conclusion, we may suppose that the assumption at the beginning of this paragraph is false.

Thus we may assume that z (as well as w) is collinear with a. Then the gamma space hypothesis gives $E \subseteq a^\perp$, and so we have

$$y \in a^\perp.$$

Now choose a line F on b and x_1. Then y is collinear with a point y_1 on F by Lemma 13.2.5 with b in the role of a. Since y is not collinear with x, but is collinear with a, y is not collinear with x_1 by the gamma space hypothesis. Thus this point y_1 on line F cannot be x_1 (although it *could* be b). Now let M be a line of y and y_1. We claim

$$a^{\perp} \cap M = \{y\}, \tag{13.1}$$

$$x_1{}^{\perp} \cap M = \{y_1\}. \tag{13.2}$$

In both cases, the left side contains the right. So we need only observe that neither a nor x_1 can lie in $M^{\perp}$. It now follows that N and M are in the "opposite" relation of Lemma 13.2.1. But (a, y_1) is a non-collinear pair in $N \times M$, and belongs to Σ by Lemma 13.2.2, part 2. Thus by Lemma 13.2.1, the pair (x, y) belongs to Σ. □

Lemma 13.2.10 *If A is a line in $S(a, b)$ containing no point of $a^{\perp} \cap b^{\perp}$, then any point y of $S(a, b)$ is collinear with a point of A.*

Proof By Lemma 13.2.6, $a^{\perp} \cap A = \{a_1\}$ and $b^{\perp} \cap A = \{b_1\}$, where a_1 and b_1 are distinct by hypothesis. Let A_1 and B_1 be lines containing $\{a, a_1\}$ and $\{b, b_1\}$, respectively. By Lemma 13.2.5, y is collinear with a point u on A_1. By Lemma 13.2.8, we can assume u is not a. Clearly we can also assume u is not a_1. Let U be a line on $\{u, y\}$.

By Lemma 13.2.8, b_1 is collinear with a point w on U. Clearly w is not u for then $b_1 \in A \cap a^{\perp} \cap b^{\perp} = \emptyset$. Also w is not collinear with a_1 for then y would be collinear with a_1 and we would be done again. Thus $u^{\perp} \cap A = \{a_1\}$ and $w^{\perp} \cap A = \{b_1\}$, and since Lemma 13.2.9 reveals that (a_1, w) belongs to Σ, these two lines are "opposite" in the sense of Lemma 13.2.1. Thus y, a point of C is collinear with a point of A, and the proof is complete. □

Lemma 13.2.11 *If A is a line in $S(a, b)$ meeting $a^{\perp} \cap b^{\perp}$ non-trivially, then any point y of $S(a, b)$ is collinear with a point of A.*

Proof We may assume that $A \cap a^{\perp} \cap b^{\perp}$ contains a point u and that u is not in $y^{\perp}$. By Lemma 13.2.8 we may assume that b is not in $y^{\perp}$. Let U be a line on $\{b, u\}$. By Lemma 13.2.5 y is collinear with a point b' on U and b' is distinct from both b and u. Note that a cannot be collinear with b'. Set Y to be a line on y and b'. By Lemma 13.2.6, a is collinear with a point t on Y, and in fact $a^{\perp} \cap Y = \{t\}$, distinct from b'. Now since t is in $a^{\perp} \cap S(a, b)$, by Lemma 13.2.8, t is collinear with a point v of A. Now v is the only point of A collinear with t, otherwise u would be collinear with t as well as b' and hence would be collinear with y, and we would be done. If $v \sim b'$, then $v \sim b$ and $A \subseteq a^{\perp} \cap b^{\perp}$. In that case, y is collinear with a point of A, from the definition of $S(a, b)$, and we would be done. Thus we see that (v, u, b', t) is a 4-circuit with both diagonals (v, b') and (u, t) non-collinear pairs. The hypotheses of Lemma 13.2.1 will thus be met if one of these diagonals is a member of Σ. But $(t, v) \in \Sigma$ by Lemma 13.2.9. Thus there is a bijection, $Y \to A$ taking each point of Y to the unique point of A with which it is collinear. In particular y is collinear to a point of A. □

Theorem 13.2.12 *The subspace $S(a,b)$ is a non-degenerate polar space of rank at least three.*

Proof First, it is a gamma space since it is a subspace of a gamma space (Lemma 13.2.3). It is immediate from this and Lemmas 13.2.10 and 13.2.11 that given any of its points one or all points of any of its lines are collinear with that point. So $S(a,b)$ is a polar space. That it is non-degenerate follows from the fact that $a^\perp \cap b^\perp$ is a non-degenerate polar space of rank at least two. □

It remains to be shown that $S(a,b)$ is a convex subspace.

Lemma 13.2.13 *Suppose $y \in S(a,b) - a^\perp$. Then any line A of $a^\perp \cap b^\perp$ lies in $S(a,y)$. (Note that $(a,y) \in \Sigma$ by Lemma 13.2.2, part 2.)*

Proof We can assume that A is not contained in $a^\perp \cap y^\perp$, or we are done at once. Then $A \cap y^\perp = \{z\}$ as A and y are in the polar space $S(a,b)$. Take $u \in A - \{z\}$, and let au be the unique line of $S(a,b)$ on points a and u. Then there is a unique point $v \in au \cap y^\perp$, again because $S(a,b)$ is a polar space. Then v is distinct from both a and u. Now since a and v are in the subspace $S(a,y)$, so is point u. As u was an arbitrary point of $A - \{z\}$, and $z \in a^\perp \cap y^\perp \subseteq S(a,y)$, we see that $A \subseteq S(a,y)$ as desired. □

Lemma 13.2.14 *For any $y \in S(a,b) - a^\perp$, $(a,y) \in \Sigma$ and $S(a,b) = S(a,y)$.*

Proof Suppose $z \in S(a,y)$ and A is any line of $a^\perp \cap b^\perp$. Then as Lemma 13.2.13 shows, $\{z\} \cup A$ is contained in the polar space $S(a,y)$ and so $z^\perp \cap A \neq \emptyset$. Since A was an arbitrary line of $a^\perp \cap b^\perp$, we now have $z \in S(a,b)$. Since z was arbitrary in $S(a,y)$, we have $S(a,y) \subseteq S(a,b)$. (Note that $(a,y) \in \Sigma$ by Lemma 13.2.2, part 2.)

But now $b \in S(a,y) - a^\perp$, and $(a,b) \in \Sigma$ so the identical argument with b and y transposed yields $S(a,b) \subseteq S(a,y)$. Thus $S(a,y) = S(a,b)$. □

Lemma 13.2.15 *For any 2-coclique (x,y) of $S(a,b)$, $(x,y) \in \Sigma$ and $S(x,y) = S(a,b)$.*

Proof First suppose $x \in S(a,b) - a^\perp$. Then $(a,x) \in \Sigma$ and $S(a,b) = S(x,a)$ by Lemma 13.2.14. So $y \in S(a,b) - x^\perp = S(a,x) - x^\perp$ and another application of the same lemma gives $S(a,x) = S(y,x)$, and we are done.

Similarly, the argument gives $S(a,b) = S(x,y)$ and $(x,y) \in \Sigma$ if $x \in S(a,b) - b^\perp$, $y \in S(a,b) - a^\perp$ or $y \in S(a,b) - b^\perp$. Thus we may assume

$$\{x,y\} \subseteq a^\perp \cap b^\perp.$$

Now, since lines are thick, we can find a point w on the unique line ax of $S(a,b)$ containing a and x, distinct from both a and x. Now three applications of Lemma 13.2.14 give us

$$\begin{aligned} S(a,b) &= S(w,b), (w,b) \in \Sigma, \text{ since } w \in S(a,b) - b^\perp, \\ S(w,b) &= S(y,b), (y,w) \in \Sigma, \text{ since } y \in S(w,b) - w^\perp, \\ S(y,w) &= S(y,x), (x,y) \in \Sigma, \text{ since } x \in S(y,w) - y^\perp. \end{aligned}$$

□

Theorem 13.2.16 *If Σ is a polar family of a point-line geometry $\Gamma = (\mathcal{P}, \mathcal{L})$, then every polar pair $(a, b) \in \Sigma$ lies in a symplecton of Γ.*

Proof It suffices to establish that $S(a, b)$ is a symplecton. Since it is a subspace that is a non-degenerate polar space of rank at least two (actually three) it remains to be shown that it is convex. Since it is a subspace of point-diameter two, it suffices to show that for any non-collinear pair of points (x, y) of $S(a, b)$, that $x^\perp \cap y^\perp \subseteq S(a, b)$. But by Lemma 13.2.15, $(x, y) \in \Sigma$ and $x^\perp \cap y^\perp \subseteq S(x, y) = S(a, b)$. So $S(a, b)$ is a convex subspace. □

13.3 Elementary Properties of Symplecta

Sometimes it is burdensome to broadcast a trivial theorem as an important one. That is because the two concepts (triviality of the proof, and usefulness of the theorem) have nothing to do with each other. If a trivial theorem is extremely useful, it should be so advertised.

The trivial theorem here is the following.

Theorem 13.3.1 *Suppose $\Gamma = (\mathcal{P}, \mathcal{L})$ is point-line gamma space and S is a symplecton of Γ not containing a point x. Then $x^\perp \cap S$ is a singular subspace of S.*

Proof Immediate from Lemma 3.4.4 of Chap. 3. □

The following result is really a basic property of polar spaces.

Lemma 13.3.2 *Suppose C is a convex subspace of a (non-degenerate) polar space of polar rank at least two. Then one of the following holds:*

1. *$C = S$.*
2. *C is a singular subspace. (This includes the case $C = \emptyset$.)*

Proof By way of contradiction, suppose C is a proper convex subspace of S that is not a singular subspace. Then there exists a non-collinear pair of points $\{x, y\}$ in C and so $x^\perp \cap y^\perp \subseteq C$ by convexity.

Let z be any point of $S - C$. Then z cannot be collinear with both x and y. Transposing x and y if necessary, we may assume z is not collinear with y. Since S is non-degenerate of polar rank at least two, $x^\perp \cap y^\perp$ contains a 2-coclique $\{u, v\}$. Then from the polar space property, z must be collinear with a point a on line yu and a point b on line yv, and a and b are both distinct from y. Now a is not collinear with b since otherwise u would be collinear with v. But then $a^\perp \cap b^\perp \subseteq C$ by convexity, forcing $z \in C$, a contradiction.

Thus the conclusion must hold. □

Corollary 13.3.3 *The following statements hold:*

1. *Let S be any symplecton of the point-line geometry Γ. Then S is the convex closure of any of its 2-cocliques.*
2. *If S and T are distinct symplecta of a point-line geometry Γ, then $S \cap T$ is a singular subspace.*

Proof 1. The first assertion is immediate from Lemma 13.3.2 above.

2. Since $S \cap T$ is the intersection of two convex subspaces of Γ, it is also convex. But since it is a proper subspace of one of them, S and T being distinct, it must be singular by the previous Lemma 13.3.2.

The proof is complete. □

The following result will prove to be very useful in Chap. 15.

Theorem 13.3.4 *Suppose A and B are two distinct symplecta of a point-line gamma space $\Gamma = (\mathcal{P}, \mathcal{L})$, and suppose $A \cap B$ is non-empty. If $x \in A - (A \cap B)^{\perp}$, then one of the following alternatives holds:*

1. *$x^{\perp} \cap B$ is a hyperplane of $A \cap B$.*
2. *$x^{\perp} \cap B$ does not lie in $A \cap B$, but contains $x^{\perp} \cap (A \cap B)$ as a hyperplane. In particular, if $A \cap B$ has finite projective rank k, then so does $x^{\perp} \cap B$.*

Remark The words "hyperplane" and "projective rank" appearing in the theorem makes sense since $A \cap B$ and $x^{\perp} \cap B$ are singular subspaces of B (Corollary 13.3.3 and Theorem 13.3.1, respectively) and all singular subspaces of a polar space are projective spaces.

Proof Assume A and B are distinct symplecta with $A \cap B \neq \emptyset$, and choose $x \in A - (A \cap B)^{\perp}$. If $x^{\perp} \cap B \subseteq A$, then the first alternative (1) must hold. So we assume for the remainder of the proof that there is a point z in $x^{\perp} \cap B - A$. Since A is a polar space, $H_0 := x^{\perp} \cap (A \cap B)$ is a hyperplane of $A \cap B$. (Of course if $A \cap B$ is a single point, H_0 is empty.) It is therefore possible to find a point y in $A \cap B - x^{\perp}$. Now $x^{\perp} \cap y^{\perp} \subseteq A$ by the convexity of A. Thus

$$H_0 \subseteq y^{\perp} \cap x^{\perp} \cap B \subseteq x^{\perp} \cap A \cap B = H_0.$$

So

$$y^{\perp} \cap x^{\perp} \cap B = H_0.$$

But since B is a polar space containing y, one must conclude that $H_0 = y^{\perp} \cap x^{\perp} \cap B$ is a hyperplane of $x^{\perp} \cap B$. Thus the first statement of the second alternative (2) holds and the remark on projective ranks follows. □

13.4 Parapolar Spaces: A Stage for Characterizing the Lie Incidence Geometries

13.4.1 The Definition of Parapolar Spaces

We suppose $\Gamma = (\mathcal{P}, \mathcal{L})$ to be an arbitrary point-line geometry. As usual, we let $\Delta = (\mathcal{P} \sim)$ be the point-collinearity graph of Γ, and for each point p we let $p^{\perp}$ denote the set of points collinear with p together with p itself.

The following is equivalent to the definition given in Cohen's article in the *Handbook of Incidence geometry* [35].[2] $\Gamma = (\mathcal{P}, \mathcal{L})$ is called a *parapolar space* if and only the following assertions hold:

(PPS1) $(\mathcal{P}, \mathcal{L})$ is a connected gamma space.
(PPS2) For every pair of distinct non-collinear points x and y in $\mathcal{P}$, either:

(a) $x^\perp \cap y^\perp = \emptyset$,
(b) $x^\perp \cap y^\perp$ is a single point, or
(c) $\{x, y\}$ lies in a symplecton – that is, a convex subspace which is a non-degenerate polar space of rank at least two.

(PPS3) Every line lies in a symplecton.

Let $\{x, y\}$ be a pair of distinct non-collinear points of the parapolar space Γ. If $|x^\perp \cap y^\perp| = 1$, as in the second alternative in (PPS2), then the pair $\{x, y\}$ is called a *special pair*. On the other hand, if $\{x, y\}$ lies in a symplecton, it is called a *polar pair*.

The parapolar space Γ is said to have *sympletic rank r* if every symplecton has polar rank exactly r. Similarly, Γ is said to have *symplectic rank at least r* if every symplecton has polar rank at least r. (In both cases $r \geq 2$.) Finally, we say that Γ is a *strong parapolar space* (or that Γ *is strongly parapolar*) if and only if it is a parapolar space without special pairs. Equivalently, it is strongly parapolar if and only if any pair of points $\{x, y\}$ at distance two in the point-collinearity graph lies in some symplecton.

Immediately one has the following.

Theorem 13.4.1 *Let $\Gamma = (\mathcal{P}, \mathcal{L})$ be a parapolar space.*

1. *Γ is a partial linear space.*
2. *If Γ has symplectic rank at least three, then every singular subspace generated by a line and a plane lies in some symplecton and is therefore a projective plane. In this case, every singular subspace of Γ is a projective space.*
3. *If Γ has symplectic rank at least $k \geq 3$, and p is any point, then each connected component of the geometry $Res_\Gamma(p) := (\mathcal{L}_p, \Pi_p)$ of all lines and projective planes of Γ which are incident with p (called the point-residue at p) is a strong parapolar space of symplectic rank at least $k-1$. In that case there are bijective correspondences:*

   (i) Singular subspaces of Γ on $p \to$ singular subspaces of $Res_\Gamma(p)$
   (ii) Symplecta of Γ on point $p \to$ symplecta of $Res_\Gamma(p)$

   which lower by one, the symplectic or projective rank.

[2] At least this is true if we interpret Cohen's word "quadrangle" (appearing in the definition of the first sentence of Sect. 4.19 of his *Handbook* article) to include 4-circuits with at most one diagonal adjacency.

Proof Part 1. By way of contradiction suppose $(\mathcal{P}, \mathcal{L})$ was a parapolar space that was not a partial linear space. Then there exist distinct lines L_1 and L_2 whose intersection $L_1 \cap L_2$ contains two distinct points. But by (PPS3) there is a symplecton S containing L_1. Since S is a subspace of Γ meeting L_2 in at least two points, L_2 is also a line of S. But that is impossible since S, being a non-degenerate polar space, is itself a partial linear space.

Part 2. Let π be a singular subspace generated by a line L and a point x. If π lies in a symplecton, then π is a projective plane since all singular subspaces of symplecta are projective spaces.

So we suppose, by way of contradiction, that π lies in no symplecton.

Now if there is a point $y \in L^{\perp}$ not collinear with x, then the convex closure $\langle\langle x, y\rangle\rangle$ is a symplecton containing π, contrary to our assumption. Thus x is collinear with all points of $L^{\perp}$. But by the parapolar hypothesis (PPS3), L lies in a symplecton S and we can write $S := \langle\langle y, z\rangle\rangle$, where $\{y, z\}$ is a 2-coclique in $L^{\perp} \cap S$. Then $x^{\perp} \cap S$ contains both y and z, and so $x^{\perp} \cap S$ is not a clique. This forces $x \in S$ as desired.

It follows that every singular space generated by a point and a line lies in a symplecton of rank at least three and so is a projective plane. In turn this implies that every maximal singular subspace of Γ is a projective space.

Part 3. By the assumption of symplectic rank at least three, every line lies in a projective plane. Then any connected component R of the point-residual $\mathrm{Res}_{\Gamma}(p)$ is certainly a connected gamma space. If L_1 and L_2 are two "points" of $\mathcal{L}_p$ at distance two, then ${L_1}^{\perp}$ does not contain L_2 and there is a third "point" L_3 which is coplanar with both L_1 and L_2. Thus if $x_i \in L_i - \{p\}$, $i = 1, 2$ then ${x_1}^{\perp} \cap {x_2}^{\perp}$ contains the line L_3 and so $L_1 \cup L_2 \cup L_3$ lies in a symplecton S of rank at least $k \geq 3$. Clearly the subspace S_p of all lines and planes of S which are incident with p form a symplecton of the point-residual $\mathrm{Res}_{\Gamma}(p)$ of polar rank $k - 1$. We have shown all that is needed to conclude that R is a strong parapolar space of symplectic rank at least $k - 1$. □

Remark All of the conclusions of part 2 of the preceding theorem may fail for parapolar spaces of symplectic rank two. There exist such parapolar spaces of symplectic rank two whose singular subspaces are not projective spaces (see Exercise 13.2). There also exist such parapolar spaces whose singular subspaces are projective spaces of arbitrarily high rank (see Exercise 13.3). Clearly for these, no projective plane can lie in a symplecton.

One can easily recognize parapolar spaces of symplectic rank at least three as follows.

Lemma 13.4.2 *Suppose* $\Gamma = (\mathcal{P}, \mathcal{L})$ *satisfies the following properties:*

(PPS1) $(\mathcal{P}, \mathcal{L})$ *is a connected gamma space.*
(PPS2') *For every pair of distinct non-collinear points* x *and* y *in* $\mathcal{P}$*, either:*

(a) $x^{\perp} \cap y^{\perp} = \emptyset$,
(b) $x^{\perp} \cap y^{\perp}$ *is a single point, or*
(c) $x^{\perp} \cap y^{\perp}$ *is a non-degenerate polar space of rank at least two.*

(PPS3') *For any line L, $L^{\perp}$ is not a clique.*

Then Γ is a parapolar space of symplectic rank at least three.

Proof Because of the rather limited alternatives for $x^{\perp} \cap y^{\perp}$ for two points x and y at distance two, the set $\{(x, y) \in \mathcal{P} \times \mathcal{P} | x \neq y, x^{\perp} \cap y^{\perp} \neq \text{ clique } \}$ forms a polar family. So by Cooperstein's theory, there exists a family $\mathcal{S}$ of symplecta of polar rank at least three which cover every pair (x, y) in the polar family. Condition (PPS3') insures that every line lies in some symplecton. □

Remark The reader should be advised that the three conditions (PPS1), (PPS2'), and (PPS3') of the preceding Lemma 13.4.2 together comprise the definition of parapolar spaces in virtually all of the literature preceding Cohen's article in the *Handbook* [35]. In this older terminology, parapolar spaces always had symplectic rank at least three. Unfortunately, this meant that the point-residuals were not necessarily parapolar spaces in this sense and so had to be given another name, and here the literature shows a tangle of names such as "polarized spaces" without a uniform meaning. It is for this reason that we have adopted Cohen's definition of parapolar spaces. If parapolar spaces in the new sense have symplectic rank at least three, their point-residuals are still parapolar spaces. If their symplectic rank is two, we usually aren't interested in the point-residuals as defined above.

13.4.2 Examples of Parapolar Spaces

Example 1 Of course any polar space of rank at least two is a parapolar space with just one symplecton.

Example 2 The dual polar spaces of Chap. 8 are strong parapolar spaces of symplectic rank two.

Bruce Cooperstein was the first to show that most of the Lie incidence geometries of spherical type are parapolar spaces. [Just to remind the reader of the notation, one indicates the geometry by a symbol $X_{n,k}$ where X_n is one of the spherical diagrams with at least three nodes and k denotes the node representing the class $\mathcal{P}$ of points of the Lie incidence geometry. The lines are the flags whose type is indicated by the nodes adjacent to node k in the diagram. The numbering of nodes used is that of Fig. 9.11 on p. 370.] The following are some of these.

Example 3 The Grassmann spaces are strong parapolar spaces of symplectic rank three.

Example 4 The half-spin geometries are strong parapolar spaces of symplectic rank four.

Example 5 The two exceptional geometries $E_{6,1}$ and $E_{7,7}$ are strong parapolar spaces of symplectic ranks five and six.

Example 6 The following Lie incidence geometries are parapolar spaces which are not strongly parapolar:

1. The metasymplectic spaces (with diagram of type F_4) which have symplectic rank three.
2. The exceptional geometries $E_{6,2}$, $E_{7,1}$, and $E_{8,8}$ which have symplectic ranks four, five, and seven, respectively.
3. Polar Grassmannians. The points are the singular spaces of projective rank k of a non-degenerate polar space of rank r where $0 < k < r - 2$. The lines are the flags (A, B) of singular subspaces of ranks $k - 1$ and $k + 1$, respectively. So this example satisfies the usual recipe for forming a Lie incidence geometry $C_{r,k}$. It possesses symplecta of polar ranks three and $r - k + 1$. These two integers may very well be different.

We shall use the parapolar space paradigm to describe characterizations of homomorphic images of many of these Lie incidence geometries of spherical type in the next five chapters.

13.5 Local Connectedness in Parapolar Spaces

13.5.1 Introduction

Suppose $\Gamma = (\mathcal{P}, \mathcal{L})$ is a gamma space. Recall from Chap. 3, that for any two distinct lines L_1 and L_2 of Γ on a common point p, either $L_1^{\perp} \cap L_2 = L_2$ or $L_1^{\perp} \cap L_2 = \{p\}$. Either equation is equivalent to the same equation with the subscripts 1 and 2 transposed. In the former case it means the two lines generate a singular subspace of the gamma space.

Thus, associated with each point p is a so-called "tangent graph" $G_p = (\mathcal{L}_p, \sim)$ whose vertices are the lines of Γ incident with p, two such lines forming an edge of the graph if and only if they are distinct and generate a singular subspace. We say that the gamma space Γ is *locally connected* if and only if, for each point p, the graph G_p is connected.

If, however, Γ is a gamma space whose singular subspaces are projective spaces (these have been termed *paraprojective spaces* in Cohen's survey article [35]), the tangent graph G_p is in fact the collinearity graph of a "tangent space" $\Gamma_p := (\mathcal{L}_p, \Pi_p)$ of lines and planes incident with p, which we have been calling the *point-residual at* p. In that case the local graph G_p is in fact the collinearity graph of the point-residue geometry

$$\mathrm{Res}(p) := (\mathcal{L}_p, \Pi_p)$$

of all lines and planes of Γ incident with point p. It follows that G_p is connected if and only if the geometry $\mathrm{Res}(p)$ is connected. One may also note that $\mathrm{Res}(p)$ is itself a gamma space.

In the case of parapolar spaces one has an immediate observation whose proof is left to the reader.

Lemma 13.5.1 *Any strong parapolar space of symplectic rank at least three is locally connected.*

By way of contrast, there are classical strong parapolar spaces of symplectic rank two which are not locally connected – for example, dual polar spaces.[3] Also there are many parapolar spaces which are not strong, which are in fact not locally connected. Consider, for example, two polar spaces S_1 and S_2 which intersect at a single point p. We may regard the union of their points and the union of their line sets as a point-line geometry $\Gamma = (\mathcal{P}, \mathcal{L})$. Then Γ is certainly a parapolar space, but it is not locally connected at the point p.

Similarly, we can paste any two locally connected parapolar spaces at a point to get another parapolar space. Indeed one can iterate this process – pasting spaces together at points – to get a collection $\Gamma_i = (\mathcal{P}_i, \mathcal{L}_i)$ of locally connected parapolar spaces where the $\mathcal{P}_i$ pairwise intersect at a coclique of each Γ_i. When is this pasted-together-concoction of locally connected parapolar spaces itself a parapolar space?

We can answer this question for the case that the Γ_i are among the locally connected components. Recall from our study of local connectedness of gamma spaces in Chap. 3 that every gamma space Γ is the point-union of locally connected gamma subspaces $\Gamma_i = (\mathcal{P}_i, \mathcal{L}_i)$ with these properties:

1. Each line L belongs to a unique Γ_i.
2. Each Γ_i is a connected subspace.
3. For each point p and subspace Γ_i containing it, the subgraph $((\mathcal{L}_i)_p, \sim)$ of all lines of Γ_i which are incident with p is a full connected component of the tangent graph $G_p = (\mathcal{L}_p, \sim)$.

Moreover, the subspaces Γ_i with these properties are uniquely determined (Theorem 3.6.1 and Corollary 3.6.2), and we called this unique collection of subspaces the *locally connected components of* Γ (see p. 69).

Suppose now Γ is a gamma space with locally connected components $\Gamma_i = (\mathcal{P}_i, \mathcal{L}_i)$. Now when is Γ a parapolar space? To answer this, let D denote the set of all points of $\mathcal{P}$ which lie in two or more of the sets $\mathcal{P}_i$. We will regard D as an induced subgraph of the point-collinearity graph $\Delta = (\mathcal{P}, \sim)$, the point-collinearity graph of Γ.

Theorem 13.5.2 *Suppose* Γ *is a connected gamma space and*

$$C := \{\Gamma_i = (\mathcal{P}_i, \mathcal{L}_i) | i \in I\}$$

is the full collection of locally connected components of Γ. *Then* Γ *is a parapolar space if and only if each of the following three statements holds:*

[3] Note that singular subspaces and generalized polygons which are not generalized quadrangles are not parapolar spaces since, for these spaces, it is false that every line lies in a symplecton.

1. *Each Γ_i is either a parapolar space or a singular subspace.*
2. *Every 4-circuit that is not a clique of Δ having two of its vertices in D either lies in one of the Γ_i or lies in a convex generalized quadrangle.*
3. *If L is a line belonging to a Γ_i which is not a parapolar space, then L possesses two distinct points x and y which are successive vertices of a 4-circuit that is not a clique.*

There is an immediate corollary:

Corollary 13.5.3 *A connected gamma space is a parapolar space of symplectic rank at least three if and only if each of its locally connected components is a parapolar space of symplectic rank at least three and each 4-circuit of the point-collinearity graph is contained in some locally connected component.*

Before proving this theorem, we require a technical lemma that allows us to recognize a symplecton.

Lemma 13.5.4 *Suppose Γ is a gamma space with the following property:*

(C) Whenever c is a 4-circuit not contained in a clique in the point-collinearity graph Δ of Γ, and two of its vertices lie in the set D, then either c lies in a locally connected component of Γ, or else lies in a convex generalized quadrangle.[4]

Then any symplecton of a locally connected component of Γ is either a symplecton of Γ or lies in some convex generalized quadrangle of Γ. In particular, any symplecton of a locally connected component of Γ which, as a symplecton, has polar rank at least three, is already a symplecton of the entire space Γ.

Proof Suppose Σ is a locally connected component of Γ and suppose S is a symplecton of Σ. If S is a convex subspace of Γ, then it is a symplecton of Γ, and we are done. So we assume that S is not convex.

Since S is not convex in Γ, there exists a point v not in S such that $v^\perp \cap S$ contains two distinct non-collinear points x and y. Then in turn, there exists another pair of distinct non-collinear points $\{a, b\}$ in $x^\perp \cap y^\perp \cap S$. Then $c = (v, x, a, y, v)$ is a 4-circuit that does not lie in a clique and does not lie in Σ or any other locally connected component of Γ. If the line vx were in the same connected component of the tangent graph G_x as the line xa, we would have vx lying in Σ contrary to the choice of v. Thus we must conclude that x lies in at least two locally connected components, and so we conclude that $x \in D$. Similarly we have $y \in D$. By condition (C), the circuit c now lies in a generalized quadrangle Q which is convex in Γ. But then $Q \cap S$ is a convex subspace of the polar space S which is not a singular subspace. By Lemma 13.3.2, $Q \cap S = S$, so $S \subseteq Q$ as desired.

[4] Note that we are not assuming that Γ is a parapolar space in this lemma. In general, the convex closure of a 4-circuit that is not a clique might not be a polar space.

The last statement of the lemma follows on noting that under the conditions of the previous paragraph, S has polar rank two. The proof is complete. □

Proof of Theorem 13.5.2 First we show that 1, 2, and 3 imply that a connected gamma space is a parapolar space. Assuming these three properties, there are just two things to prove: (a) that every line L lies in some symplecton, and (b) that any 4-circuit that is not contained in a 4-clique lies in a symplecton.

The second statement (b) follows easily. Suppose $c := (x_0, x_1, x_2, x_3, x_0)$ is a 4-circuit that is not a clique. If c does not lie in any Γ_i, by hypothesis it must lie in a convex generalized quadrangle, a symplecton of polar rank two. If it lies in some Γ_i, evidently that Γ_i is not a singular subspace, so by 1 it is a parapolar space which necessarily can provide its own symplecton S containing c. Now, noting that 2 is condition (C), by Lemma 13.5.4, S is a symplecton of Γ or is a subpolar space of a convex generalized quadrangle. In either case it lies in a symplecton of Γ as required for (b).

The proof of (a) is even easier. By 3 every line L connects two distinct points that belong to some 4-circuit that does not induce a clique. But we have just proved that every such 4-circuit lies in some symplecton of Γ.

Now we must prove the converse: If Γ is a parapolar space, then 1, 2, and 3 must hold. But 3 and 2 are immediate consequences of the fact that every line and every 4-circuit of Γ not confined to a clique must lie in some symplecton. For in 2, if the 4-circuit is not in a locally connected component, it must be a generalized quadrangle. The problem is 1. We must show that any locally connected component $\Gamma_i = (\mathcal{P}_i, \mathcal{L}_i)$ of Γ runs to one of two extremes: it is either a parapolar space itself, or it is a singular subspace.

We begin with a handy result.

Step 1. *Suppose R is a symplecton of Γ of polar rank at least three which shares a line with Γ_i. Then $R \subseteq \Gamma_i$.*

Proof Let L be a line in $R \cap \Gamma_i$ and suppose x were a point of $R - \Gamma_i$. Then there is a point y in $x^\perp \cap L$ and since S has polar rank at least three, there exists a line N on y which lies in $x^\perp \cap L^\perp$. Then lines xy and L belong to the same connected component of the tangent graph G_y. But that component must consist of all the lines of Γ_i on y forcing xy to lie in Γ_i; against the choice of x. Thus no such x can exist and so $R \subseteq \Gamma_i$. The proof of Step 1 is complete.

Next we partition the lines of Γ_i into two sets

$$\mathcal{L}_i = \mathcal{L}_i^- + \mathcal{L}_i^+,$$

where

$$\mathcal{L}_i^- := \{L \in \mathcal{L}_i | \text{no symplecton on } L \text{ lies in } \Gamma_i\},$$
$$\mathcal{L}_i^+ := \{L \in \mathcal{L}_i | \text{some symplecton on } L \text{ lies in } \Gamma_i\}.$$

Step 2. *If $L \in \mathcal{L}_i^-$ and N is a line of $\mathcal{L}_i$ meeting L at a point, then $N \subseteq L^\perp$.*

Proof Suppose N is not in $L^\perp$. Then $L \cap N = \{p\}$ and by local connectedness of Γ_i, there is a sequence of lines of $\mathcal{L}_i$ on point p — say, $(L = L_0, L_1, \ldots, L_m = N)$ — with $L_{j+1} \subseteq {L_j}^\perp$, for $j = 0, \ldots, m-1$. Then there is a smallest index k such that $L_k \subseteq L^\perp$ while $L_{k+1} \not\subseteq L^\perp$. Since ${L_{k+1}}^\perp \cap L^\perp$ contains the line L_k, the parapolar space Γ contains a symplecton S of polar rank at least three containing the set $L \cup L_{k+1}$. By Step 1, $S \subseteq \Gamma_i$. But that, in turn contradicts the fact that $L \in \mathcal{L}_i^-$.

Thus we must have $N \subset L^\perp$, and the proof of Step 2 is complete.

Step 3. *Suppose $L \in \mathcal{L}_i^-$ and that N is a line of $\mathcal{L}_i$ for which $L \cap N$ is non-empty. Then $N \in \mathcal{L}_i^-$.*

Proof If false the line N belongs to $\mathcal{L}_i^+$, and so lies in a symplecton R of Γ which is contained in Γ_i. Moreover $L \cap N = \{p\}$, a single point. Then in R, there is a second line N' on p which does not lie in $N^\perp$. Now by Step 2, both of these lines N and N' lie in $L^\perp \cap R$. But then convexity of R forces $L \subseteq R$, against $L \in \mathcal{L}_i^-$. So the conclusion of Step 3 is true.

Step 4. *If $\mathcal{L}_i^-$ is non-empty, then $\mathcal{L}_i^+$ is empty.*

Proof If false then there exists both a line $L \in \mathcal{L}_i^-$ and a line $N \in \mathcal{L}_i^+$. Since Γ_i is a connected subspace there is a sequence of lines $(L = L_0, L_1, \ldots, L_m = N)$ with $L_j \cap L_{j+1}$ non-empty for $j = 0, \ldots, m-1$. Since L and N lie in different parts of the line-partition, there exists a smallest integer k for which $L_k \in \mathcal{L}_i^-$ and $L_{k+1} \in \mathcal{L}_i^+$. But this contradicts Step 3.

Step 5. *If $\mathcal{L}_i^-$ is non-empty, then Γ_i is a singular subspace.*

Proof Suppose $\mathcal{L}_i^-$ is non-empty. By Step 4, $\mathcal{L}_i = \mathcal{L}_i^-$. If Γ_i were not a singular space there would exist a 2-coclique $\{x, y\}$ in Γ_i. Since Γ_i is connected, there is a sequence of lines $(L_0, L_1, \ldots, L_m)$ with $L_j \cap L_{j+1}$ non-empty for $j = 0, \ldots, m-1$ with $x \in L_0$ and $y \in L_m$. Now $L_0 \subseteq x^\perp$ while L_m is not in $x^\perp$. Thus there is a smallest integer k such that $L_k \subseteq x^\perp$ while L_{k+1} is not in $x^\perp$. Now, by Step 2, $L_1 \subseteq {L_0}^\perp \subseteq x^\perp$, so we have $k > 1$. Also by Step 2, $L_{k+1} \subseteq {L_k}^\perp$. But in that case, the convex closure of $L_0 \cup L_{k+1}$ is a symplecton S of Γ of polar rank at least three. But again that contradicts $L_0 \in \mathcal{L}_i^-$. So Γ_i must be a singular space and Step 5 is proved.

Now assume the locally connected component Γ_i is not a singular space. Then by Step 5, we see that $\mathcal{L}_i^-$ is empty, and so $\mathcal{L}_i = \mathcal{L}_i^+$. That means that every line of Γ_i lies in at least one symplecton of Γ which is contained in Γ_i. Such a symplecton is clearly a symplecton of the subspace Γ_i. Now consider any 4-circuit $c = (a, b, c, d, a)$ of $\Delta_i := (\mathcal{P}_i, \sim)$, whose vertices do not form a clique. Then, since Γ is a parapolar space, this circuit lies in some symplecton S. If S lies in Γ_i, then S is a symplecton of Γ_i. By Step 1, this must happen if S has polar rank at least three. Suppose then that S does not lie in Γ_i. Then S is a convex generalized quadrangle, and $S \cap \Gamma_i$ is (1) a convex subspace of Γ_i, and (2) a subspace of the generalized quadrangle containing a 4-circuit not contained in a line of S. It follows

that $S \cap \Gamma_i$ is a symplecton of the subspace Γ_i containing c. Thus by definition, Γ_i is a parapolar space.

So (1) has been proved. All parts of Theorem 13.5.2 have been proved. □

13.5.2 Isolating Polar Spaces from Locally Connected Parapolar Spaces

If X is a subspace of some point-line geometry $\Gamma = (\mathcal{P}, \mathcal{L})$, we say that a point p is a *deep point of* X or that X *possesses a deep point* p if and only if $p^\perp \subseteq X$. We can now state the following.

Theorem 13.5.5 *Suppose* Γ *is a locally connected parapolar space of symplectic rank at least* 3. *Suppose* X *is a connected* 2*-convex subspace of* Γ.

If X *possesses a deep point, then* $X = \Gamma$.

In particular, if a symplecton S *of* Γ *contained a deep point, then* Γ *is itself a non-degenerate polar space of rank at least three.*

Proof Assume that Γ is a locally connected parapolar space of symplectic rank at least 3. Let $D(X)$ be the set of all deep points of X, relative to Γ. We assume $D(X) \neq \emptyset$. If $D - X$, so every point of X is a deep point, then X (viewed as an induced subgraph) is a connected component of the point-collinearity subgraph $\Delta = (\mathcal{P}, \sim)$ of Γ. But since parapolar spaces are by definition connected, this would force $X = \Gamma$, and we are done.

So we may suppose that $X - D(X)$ is non-empty. We shall show that $X - D(X)$ is a connected component of $\Delta_X := (X, \sim)$, the point-collinearity graph of the subspace X. Such a conclusion is sufficient to complete the proof since the hypothesized connectedness of Δ_X would then force $D = \emptyset$, contrary to our assumption that X possesses a deep point.

So assume $x \in X - D(X)$ and that y is a point of $X \cap (x^\perp - \{x\})$. Then x and y are connected by a line L lying in X while x also lies on a line N not contained in the subspace X. We wish to show that $y \in X - D(X)$. So by way of contradiction we assume that this is false – that is, y is a deep point of X.

Now by assumption the point-residual $\Gamma_x = (\mathcal{L}_x, \Pi_x)$ is a connected geometry. Thus there is a minimal sequence of lines and planes,

$$(L = L_0, \pi_0, L_1, \pi_1, L_2, \pi_2, \ldots, L_{n-1}, \pi_{n-1}, L_n = N),$$

where, for $i = 0, 1, \ldots, n-1$, each π_i is a plane spanned by the lines L_i and L_{i+1}, both of which are incident with p.

Now $L^\perp \subseteq X$ since y is a deep point of X. Thus $\pi_0 \subseteq X$. On the other hand ${L_{n-1}}^\perp \not\subseteq X$ since the line $L_n = N$ is not in X. So there is a smallest integer $k \geq 2$ such that ${L_{k-1}}^\perp \subseteq X$ while ${L_k}^\perp \not\subseteq X$. Thus we can choose a plane π' on L_k which is not in X. Clearly π' is not in ${\pi_{k-1}}^\perp$. Thus the convex closure $S := \langle\langle \pi', \pi_{k-1}\rangle\rangle$ is a symplecton. Since X is also 2-convex, the intersection $S \cap X$ is a singular subspace. But that is impossible, since the intersection contains ${L_{k-1}}^\perp \cap S$ which, being the

"perp" of a line in a polar space S of polar rank at least three, cannot be a singular subspace.

The proof is complete. □

Corollary 13.5.6 *Suppose Γ is a locally connected parapolar space of symplectic rank at least three. Suppose, for some point p, the point-residual $\Gamma_p := (\mathcal{L}_p, \Pi_p)$ (the incidence geometry of all lines and projective planes incident with point p) is a polar space. Then Γ is itself a polar space of rank at least three.*

Proof Since Γ is a parapolar space, and Γ_p is a (possibly degenerate) polar space of symplectic rank at least two, the convex closure in Γ of $p^\perp$ is some symplecton S. Then Γ_p is the point-residual of S and so is a non-degenerate polar space of polar rank at least two. Moreover $p^\perp \subseteq S$, so p is a deep point of S. Since S is convex and Γ is locally connected and has symplectic rank at least three, Theorem 13.5.5 now forces $S = \mathcal{P}$. The proof is complete. □

13.6 Hanssens' Principle

This section concerns a method of enriching certain parapolar spaces by showing the existence of a class of convex subspaces which are isomorphic to Grassmannians of lines, $A_{n,2}$.

The result is as follows.

Theorem 13.6.1 *Assume n is a (possibly infinite) cardinal number, which is at least* 4.

Suppose Γ is a parapolar space with the property that every point-residual $Res_\Gamma(p)$ is isomorphic to a product geometry $P \times X$, where P is the projective space $PG(n-2, K)$ of projective rank $n - 2 \geq 2$, and X is a given partial linear subspace of the residual such that:

1. *Every point of X lies in at least one line, but*
2. *X does not have $PG(n-2, K)$ as a locally connected component – that is, X does not have a maximal singular subspace M which is isomorphic to $PG(n-2, K)$ and which intersects every other maximal singular subspace of X at the empty set or a point. This cardinal number n appearing in these two hypotheses can be infinite.*

Then the following hold:

1. *Γ possesses a class $\mathcal{M}$ of maximal singular subspaces, each isomorphic to $PG(n-1, K)$ with the property that every line lies in a unique member of $\mathcal{M}$. Any two members of $\mathcal{M}$ are disjoint or intersect at a single point.*
2. *Then Γ possesses two distinct classes of projective planes: (1) those planes which lie in a member of $\mathcal{M}$ (we call these the $\mathcal{M}$-planes), and (2) those projective planes π which appear in a copy of X in the point-residuals of any of its points (the Σ-planes).*

3. *If two members A and B of $\mathcal{M}$ intersect at a point p, then there are exactly two possible relationships between A and B:*
   (a) *For each line L of A which is incident with p, $L^{\perp} \cap B = \{p\}$.* (In this case we write $A \sim_2 B$.)
   (b) *For each line L of A incident with point p, there exists a unique line $f(L)$ whose point-shadow is $L^{\perp} \cap B$. In this case $\langle L, f(L)\rangle$ is a Σ-plane. The mapping $f : \mathcal{L}_p \cap Res(A) \to Res(B) \cap \mathcal{L}_p$ which takes each line L of A on point p to the line $L^{\perp} \cap B$ of B on point p is a bijection. Each plane $\pi := \langle L, f(L)\rangle$ so is a Σ-plane meeting A and B at lines L and $f(L)$, respectively.* (In this case we write $A \sim_1 B$.)

4. *Γ possesses a unique class $\mathcal{D}$ of symplecta of type $D_{3,2}$ – that is they are oriflame of polar rank exactly three. Moreover:*
   (a) *If π_1 and π_2 are respectively an $\mathcal{M}$-plane and a Σ-plane, which intersect at a line, then they lie together in a unique symplecton $D = \langle\langle \pi_1, \pi_2\rangle\rangle$ from $\mathcal{D}$. In fact they are maximal singular subspaces of D belonging to distinct oriflame classes.*
   (b) *If $D \in \mathcal{D}$ and $M \in \mathcal{M}$, then $D \cap M$ is either empty, a single point, or is an $\mathcal{M}$-plane. The intersection is never a line.*
   (c) *Every line lies in a symplecton from $\mathcal{D}$.*
   (d) *If there are any further symplecta of Γ which are not in $\mathcal{D}$, they contain no $\mathcal{M}$-planes.*

5. (Hanssens' principle.)[5] *There exists a class $\mathcal{A}$ of convex subspaces of Γ with these properties:*
   (a) *Each member of $\mathcal{A}$ is isomorphic to a Grassmannian of lines $A_{n,2}$.*
   (b) *If two distinct subspaces A and B of $\mathcal{M}$ intersect at a point p and satisfy the relationship $A \sim_1 B$ defined above, then these two subspaces are contained in a unique member of $\mathcal{A}$.*

Remark The author considers Hanssens' principle to be a theorem that simply cannot be omitted in any exposition that purports to teach students the theory of point-line characterizations via parapolar spaces. As we have observed, in many cases

[5] The author first encountered a special case of this principle (which he has called "Hanssens' principle") in a paper by G. Hanssens [68] where the existence of these convex Grassmannians was invoked in a single sentence – referring to his thesis for a proof. (The author believes Hanssens saw the result as "obvious"). At the time the author, having no opportunity to obtain a copy of this thesis (his university was then in the process of cancelling journals and library services to the Mathematics Department by over 60%), devised his own proof in an unpublished diatribe (essentially a diary entry) called "Hanssens principle." That first proof was rather cumbersome. It recreated the subspaces of $\mathcal{A}$ along the lines of Cameron's proof of the existence of classical near k-gons as outlined in Chap. 8. In recent years, it became clear that the existence of $\mathcal{A}$ could be proved much more simply by arguments that were already at hand in the theory of Grassmannians. Not only that, the result could be proved with the full generality of the original "Hanssens principle" entry in the author's mathematical diary.

where we have a locally truncated geometry, we cannot produce a sheaf unless we have an enriched geometry of rank at least four. The problem with parapolar spaces of symplectic rank at least three is that, although we seem to have four types of objects: points lines, planes and symplecta, in the case that point-residuals are product geometries, we wind up with lines being flags formed from the remaining objects. In fact, sometimes symplecta themselves appear as flags of objects that are really not known to exist. So in this case, when point-residuals are product geometries, we can often only get our hands on just three objects. But with Hanssens' principle operative we can actually enrich our geometry to four objects by incorporating the line-Grassmannians.

If that happens, we can invoke the Brouwer–Cohen/Ronan sheaf theory (as described in Chap. 12) and hopefully Tits' "local approach theorem" (Chap. 10) can then be applied to conclude that the geometry Γ is a symplecton-preserving homomorphic image of a Lie-incidence geometry of a building.

I have insisted on this discussion to show that Hanssens' principle is really quite important. I beleive Hanssens knew this, but supposed that proofs that were easy for him shouldn't require lengthy expositions. He left this for others.

Proof Almost every part of this theorem except that last part entitled "Hanssens' principle" follows directly from the strong hypotheses we have placed upon the point-residues of this parapolar space. So we will touch briefly on the former part, and then go into detail on "Hanssens' principle".

Parts 1 and 2. For any point p, the point-residue $\mathrm{Res}(p) = (\mathcal{L}_p, \Pi_p)$ is assumed to have the form $M' \times X$ where $M' \simeq PG(n-2, K)$, $n-2 > 0$, and X has no locally connected component isomorphic to $PG(n-2, K)$. Thus by Exercise 13.10 there is a class of maximal singular subspaces $\mathcal{M}$, each isomorphic to $PG(n-1, K)$ such that each line L lies in a unique space $M(L) \in \mathcal{M}$. Moreover, if point p is in L, then the lines and planes of $M(L)$ which are incident with p comprise the singular subspaces $(M' \times x)$ partitioning the "points" of $\mathrm{Res}(p) = M' \times X$ (here, the "point" x is the line L).

Part 3 follows directly from the structure of $\mathrm{Res}(p)$ as a product of point-line geometries. If two of the "horizontal spaces", $M' \times x$ and $M' \times y$ are distinct but are connected by a line, then we must have x "collinear" with y in X, and each point $m \times x$ of $M' \times x$ is collinear only with $m \times y$ in $M' \times y$. This accounts for the bijections described in part 3.

Part 4 arises because Γ is a parapolar space. Let A and B be an $\mathcal{M}$-plane and a Σ-plane meeting at a line L. Then for any point p of L, these subspaces determine a convex subspace $A(p) \times B(p) \simeq PG(1, K) \times PG(1, K)$, a subspace "grid" in the point-residue $\mathrm{Res}(p)$, as well as a simplecton $D := \langle\langle A, B\rangle\rangle_\Gamma$ which is necessarily rank three and oriflame. We let $\mathcal{D}$ be the class of all such symplecta. The first two of the three assertions of this part are immediate from the relation of one of these grids to a "point", or to a maximal singular space $M' \times x$, in $\mathrm{Res}(p)$. The last assertion follows from the way we defined $\mathcal{D}$ above.

Part 5. The proof of Hanssens' principle. Let $\Omega = (\mathcal{M}, \sim_1)$, be the graph whose vertices are the elements of $\mathcal{M}$, two being adjacent if they bear the relationship $\sim_1$

of part 3(b). If (A, B) is an edge in Ω we let $\Pi(A, B)$ be the complete set of Σ-planes which meet both A and B at lines. Any such plane must contain the intersection $A \cap B = \{p\}$, and the set $\Pi(A, B)$ is bijective with $\mathcal{L}_p \cap \mathcal{L}(A)$ as well as $\mathcal{L}_p \cap \mathcal{L}(B)$.

Let $\mathcal{M}_1(A, B)$ be the collection of all singular spaces of $\mathcal{M}$ which lie on the point p and intersect some plane $\pi \in \Pi(A, B)$ at a line. Clearly A and B lie in $\mathcal{M}_1(A, B)$. (It should be manifest from its definition that $\mathcal{M}_1(A, B)$ is a clique in the graph Ω.)

Suppose C is an element of $\mathcal{M}_1(A, B)$ distinct from A and B. Then of course $A \cap B = A \cap C = B \cap C = \{p\}$, and there is a plane $\pi \in \Pi(A, B)$ intersecting, A, B, and C at three distinct lines on p – say L_A, L_B, and L_C, respectively. Now let π' be another plane in $\Pi(A, B)$ distinct from π. Then we obtain two new lines $L'_A := A \cap \pi'$ and $L'_B := B \cap \pi'$ distinct from L_A and L_B, respectively. Then there is a symplecton $S \in \mathcal{D}$ containing both planes π and π'. Since $C \cap S$ contains line L_C, by part 4, it is an $\mathcal{M}$-plane. But this intersection is an $\mathcal{M}$-plane that already meets the Σ-plane π' at a point; therefore it meets it at a line since these planes belong to opposite oriflame classes of S. Thus we see the following.

Step 1. *Every member of* $\mathcal{M}_1(A, B)$ *intersects every* Σ*-plane of* $\Pi(A, B)$ *at a line on* p.

Now let $\mathcal{M}_2(A, B)$ be the collection of all members E of $\mathcal{M}$ which do *not* contain $\{p\} := A \cap B$ but nonetheless intersect both A and B at points.

Suppose $E \in \mathcal{M}_2(A, B)$, and set $\{a\} := A \cap E$ and $\{b\} := B \cap E$. Then $\langle a, b, p\rangle_\Gamma$ is a Σ-plane π of $\Pi(A, B)$, and so by Step 1, $E \sim_1 C$ for each $C \in \mathcal{M}_1(A, B)$.

Now let E' be a further member of $\mathcal{M}_2(A, B)$ distinct from E. Then E' also meets A and B at points which together with p generate a Σ-plane π' of $\Pi(A, B)$. If $\pi' = \pi$ then certinly $E \sim_1 E'$. If $\pi' \neq \pi$ then, because $f : \mathcal{L}_p \cap A \to \mathcal{L}_p \cap B$ was an injection, $\pi' \cap \pi = \{p\}$, and so there is a symplecton $S = \langle\langle \pi', \pi\rangle\rangle_\Gamma$ exactly as in the paragraph before last. Then E and E' both meet S at $\mathcal{M}$-planes of the same oriflame class, and so these planes intersect at a point – say q. Also any line on q in one of these planes is coplanar with a line on q in the other. Thus $E \sim_1 E'$. We conclude the following.

Step 2. *The set*

$$\mathcal{M}(A, B) := \mathcal{M}_1(A, B) + \mathcal{M}_2(A, B)$$

is a clique in the graph $\Omega = (\mathcal{M}, \sim_1)$.

Next, suppose $C \in \mathcal{M}_2(A, B)$. Then $A \cap B \cap C = \emptyset$ and so $A \in \mathcal{M}_2(B, C)$ and $B \in \mathcal{M}_2(A, C)$. In this case we say that (A, B, C) is a Σ*-triangle* of the graph Ω for the reason that each such triangle determines a unique Σ-plane $\pi(A, B, C)$ spanned by the three points of their pairwise insections $A \cap B$, $A \cap C$, $B \cap C$.

Suppose (A, B, C) is a Σ-triangle, and set $\{p\} = A \cap B$, $\{a\} = C \cap A$ and $\{b\} = C \cap B$ as before. Then $\langle a, b, p\rangle = \pi(A, B, C)$. We claim that $\mathcal{M}(A, B) \subseteq \mathcal{M}(A, C)$.

If $X \in \mathcal{M}_1(A, B)$, then by Step 1, X intersects the plane $\pi(A, B, C) \in \Pi(A, B)$ at a line. But since $\pi(A, B, C) \in \Pi(A, C)$ as well, we see that $X \in \mathcal{M}_2(A, C)$.

On the other hand, if $X \in \mathcal{M}_2(A, B)$ there are two cases to consider according to whether this intersection contains the unique point a of $A \cap C$. If X contains a then the point in $X \cap B$ lies in $a^\perp \cap B = pb \subseteq \pi(A, B, C)$. Thus X meets $\pi(A, B, C)$ at a line on a and so $X \in \mathcal{M}_1(A, C)$. Suppose now that X does not contain a. Then X and C are two elements of $\mathcal{M}_2(A, B)$ and so X intersects C at a point c by Step 2. This places X in $\mathcal{M}_2(A, C)$.

Having covered all cases of $X \in \mathcal{M}(A, B)$, our claim is proved. By the symmetry of all pairs drawn from a Σ-triangle, the reverse containment holds as well. Thus we have the following.

Step 3. *If* (A, B, C) *is a* Σ*-triangle then*

$$\mathcal{M}(A, B) = \mathcal{M}(A, C) = \mathcal{M}(B, C).$$

Now suppose X were an element of $\mathcal{M}$ adjacent in Ω to all three members of a Σ-triangle (A, B, C). If X does not contain $A \cap B$ it lies in $\mathcal{M}_2(A, B)$. Similarly, if X does not contain $A \cap C$ or does not contain $B \cap C$, it lies in $\mathcal{M}_2(A, C)$ or $\mathcal{M}_2(B, C)$, respectively. But X cannot contain all three intersections $A \cap B$, $A \cap C$, $B \cap C$, for then it would contain the Σ-plane they generate against part 2 above. Thus, using Step 3, we see that $X \in \mathcal{M}(A, B)$. Thus we continue as follows.

Step 4. *If* (A, B, C) *is a* Σ*-triangle in the graph* Ω*, then* $\mathcal{M}(A, B)$ *is the unique maximal clique of* Ω *containing this triangle. It then follows that* $\mathcal{M}(A, B) = \mathcal{M}(X, Y)$ *for any two distinct vertices* X *and* Y *of* $\mathcal{M}(A, B)$.

Remark There are other maximal cliques besides $\mathcal{M}(A, B)$ on the edge (A, B). None of them contains a Σ-triangle. That means that each element of such a clique contains the unique point of $A \cap B$. These cliques are bijective with the collection of all maximial singular subspaces on $A \cap B$ all of whose planes are Σ-planes.

Now suppose (A, B) is an edge in Ω and set $\{p\} := A \cap B$. We define the following set of points:

$$G(A, B) := \cup\{M | M \in \mathcal{M}(A, B)\}.$$

Suppose x and y are two points of G. Then there exists elements X and Y of $\mathcal{M}(A, B)$ containing x and y respectively. Since either $X = Y$ or $X \cap Y = \{z\}$ (since $\mathcal{M}(A, B)$ was a clique of Ω) we see that x and y are points at distance at most two in the subgraph of the point-collinearity graph of Γ induced by the subset $G(A, B)$ of $\mathcal{P}$.

Now suppose x and y are distinct points on some line L. If $X = Y$ then $L \subseteq X \subseteq G(A, B)$. If $X \cap Y = \{z\}$ and z is on L then L is either xz or yz and so L again belongs to $G(A, B)$. Otherwise $\pi = \langle z, x, y \rangle_\Gamma$ is a member of $\Pi(X, Y)$

and every line of π on z lies in a unique member of $\mathcal{M}(A, B)$. So again we obtain $L \subseteq G(A, B)$. Hence we continue as follows.

Step 5. *Suppose* (A, B) *is an edge of* Ω. *Then the point-set* $G(A, B)$ *is a subspace of* $\Gamma = (\mathcal{P}, \mathcal{L})$ *whose point-collinearity graph has diameter two.*

Now suppose $D \in \mathcal{D}$ is a symplecton two of whose $\mathcal{M}$-planes belong to $\mathcal{M}(A, B)$. Since any two $\mathcal{M}$-planes of D meet at a point, it is clear that the collection of all elements of $\mathcal{M}$ which meet D at a plane comprises a clique C in Ω – a clique which contains A and B as well as a Σ-triangle. It follows from Step 4, that $C \subseteq \mathcal{M}(A, B)$, and from this one gets $D \subseteq G(A, B)$.

Now suppose y is a point which is collinear with two non-collinear points u and v of $G(A, B)$. By assumption u lies in $M_u \in \mathcal{M}(A, B)$ and similarly v lies in $M_v \in \mathcal{M}(A, B)$, and since u and v are not collinear, $M_u \cap M_v = \{p\}$ where p is not in $\{u, v\}$. Since $M_u \sim_1 M_v$, we see that $u^\perp \cap M_v$ is a line on p, and so (u, v) cannot be a special pair. So we may form the symplecton $D := \langle\langle u, v\rangle\rangle$ and observe from part 4 (b) that D intersects each of M_u and M_v at an $\mathcal{M}$-plane. From the previous paragraph we conclude that $D \subseteq G(A, B)$. In particular, x is in $G(A, B)$.

With these two paragraphs we have shown the following.

Step 6. $G(A, B)$ *is a convex subspace which is itelf a strong parapolar space.*

We are now ready to prove the following.

Step 7. $G(A, B)$ *is the Grassmannian of lines of a projective space* $\mathbf{P}(V)$. *If the elements of* $\mathcal{M}$ *have finite projective dimension* k, *then* $\mathbf{P}(V)$ *has finite projective dimension* $k + 1$.

It is sufficient to show that $(\mathcal{M}(A, B), G(A, B))$ – viewed as a geometry of points and lines – is a projective space $\mathbf{P}(V)$. For that purpose we only have to show Pasch's axiom holds – that is, in the leftmost configuration of Figure 13.1 (the one at the tail of the arrow), the lines r and s intersect at a point X (not shown). Now we can reinterpret the meanings of all the incidences and the desired conclusions in terms of parapolar points and maximal singular $\mathcal{M}$-subspaces of the strong parapolar space $G(A, B)$. One then obtains the configuration at the head of the arrow in Fig. 13.1. Here it suffices to show that the parapolar points r and s of $G(A, B)$ are collinear. For then the line $L = xy$ is a line of $G(A, B)$ and the unique maximal singular subspace X of $\mathcal{M}$ which contains L, clearly participates in a Σ-triangle with two further members of $\mathcal{M}(A, B)$, and so lies in $\mathcal{M}(A, B)$ by Step 4.

Suppose, by way of contradiction, that the parapolar points r and s of the rightmost configuration of Fig. 13.1 are not collinear. Then by Step 6 they belong to a symplecton $D = \langle\langle r, s\rangle\rangle_\Gamma$ which is itself contained in $G(A, B)$. But that is impossible, for in that case $\langle p, q, r\rangle$ and $\langle p, q, s\rangle$ would be two Σ-planes of D which meet at a line, contradicting the oriflame structure of D.

The proof is complete. □

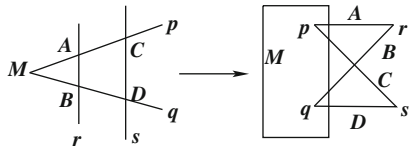

Fig. 13.1 The Pasch configuration for the linear space $(\mathcal{M}(A, B), G(A, B))$ is at left. (M is distinct from A, B, C, and D.) The *arrow* indicates the reinterpetation of these incidences in terms of the strong parapolar space $G(A, B)$. In the former configuration one wishes to show that the "lines" r and s meet at a "point" X; in the latter configuration one wishes to show that r and s lie in a common member of $\mathcal{M}(A, B)$, and for that it is sufficient to show that they are collinear

Remark Note that nowhere in the proof is n assumed to be finite. When n is an infinite cardinal, the expressions "$n - 1$" and "$n - 2$" which appear represent the same cardinal number.

13.7 When Is Either $\Delta_2^*(p)$ or $P_2^*(p)$ a Subspace of a Parapolar Space?

We suppose $\Gamma = (\mathcal{P}, \mathcal{L})$ is a parapolar space, and let $\Delta = (\mathcal{P}, \sim)$ be its point-collinearity graph. For each point p, let $\Delta_2(x)$ be the collection of all points x which are at distance two from vertex p in the graph Δ.[6] In a parapolar space, this set has an important subset: the set $P_2(p)$ of points q such that (p, q) is a polar pair. We are interested in two subsets:

$$\Delta_2^*(p) := p^\perp + \Delta_2(p),$$
$$P_2^*(p) := p^\perp + P_2(p).$$

Thus the first set is the set of all points at distance at most two from p, while the second set is the collection of all points which share a symplecton with p. The title of this section asks when either of these two sets are subspaces.

To discuss the question, one must consider some special properties that a parapolar space Γ may or may not possess.

(NG) (Near-gated) If x is a point and S is a symplecton of Γ such that $|x^\perp \cap S| = 1$, then that S is gated with respect to the point x. (See the definition on p. 12.)
(H2) If L is a line intersecting a plane π at a unique point p, then either none, one, or all of the lines of π on p share a symplecton with line L.
(D2) For each point p, the set $\Delta_2^*(p)$ is a subspace of Γ.
(P2) For each point p, $P_2^*(p)$ is a subspace of Γ.

[6] This is consistent with the notation introduced in Chap. 2.

Many spaces, such as $D_{5,5}$, $E_{6,1}$, and the hexagonic geometries discussed in the next section, have property (NG) for the simple reason that the hypothesis $|x^\perp \cap S| = 1$ never occurs. The property is fairly pervasive among Lie incidence geometries.

Lemma 13.7.1 *Suppose $\Gamma = (\mathcal{P}, \mathcal{L}) \simeq A \times B$ where each of A and B are either a singular space or a parapolar space with property (NG). Then Γ has property (NG).*

We shall see that the property also holds for Grassmann spaces, half-spin geometries, and many other Lie incidence geometries.

Lemma 13.7.2 *For any strong parapolar space, property (NG) implies property (D2).*

Proof Suppose Γ is a strong parapolar space with property (NG). Suppose, for some point p and line L, two distinct points x and y belong to $L \cap \Delta_2^*(p)$. We claim that $L \subseteq \Delta_2^*(x)$. If $L \cap p^\perp \neq \emptyset$ the claim is certainly true; so we assume that $L \cap p^\perp = \emptyset$. Thus $d(x, p) = d(y, p) = 2$. Since Γ is a strong parapolar space, there is a symplecton S containing $\{p, x\}$. If $y \in S$, we have $L = xy \subseteq S \subseteq \Delta_2^*(p)$ and we are done. So we may assume y is not in S. If $y^\perp \cap S$ contains a line N on x, then there is a unique point z in $p^\perp \cap N$, and one now sees that $L \subseteq z^\perp \subseteq \Delta_2^*(p)$ and our claim is realized again. Thus we must assume that $y^\perp \cap S = \{x\}$. Then by property (NG), we must conclude that S is gated with respect to y, with *gate* x. This means

$$\Delta_2^*(y) \cap S = x^\perp \cap S. \tag{13.3}$$

But p is a member of the left side of Eq. (13.3), and so p is collinear with x, contradicting $d(x, p) = 2$ above. □

Theorem 13.7.3 *A parapolar space of symplectic rank at least three has property (H2) if and only if every point-residual $Res_\Gamma(p)$ satisfies property (D2).*

Since point-residuals of parapolar spaces are strong parapolar spaces, we may apply Lemma 13.7.2, and Theorem 13.7.3, to conclude the following.

Corollary 13.7.4 *Any parapolar space of symplectic rank three, all of whose point-residuals satisfy (NG), must possess property (H2).*

Theorem 13.7.5 *Suppose Γ is a parapolar space satisfying properties (NG) and (H2). Then (P2) holds: that is, for each point p, $P_2^*(p)$ is a subspace of Γ.*

Proof This proof is similar to that of Lemma 13.7.2 except now Γ is not a strong parapolar space. Suppose p is a point, and L is a line bearing two distinct points x and y which belong to $P_2^*(p)$. We claim that $L \subseteq P_2^*(p)$. There are a few simple cases that can first be dismissed.

Since $p^\perp \subseteq \Delta_2^*(p)$, we are done if $L \subseteq p^\perp$. So we assume this is not so. Since $p^\perp$ is a subspace, if $L \cap p^\perp$ is not empty, it must be a single point q. Then at least one of x and y (say x) is distinct from q, and the convex closure of x and p is a symplecton S containing q and the line $qx = L$. Since $S \subseteq \Delta_2^*(p)$ we are done.

Thus we can assume from now on that $L \cap p^\perp = \emptyset$ and that $\{x, y\} \subseteq P_2(p)$. Let X and Y be the unique symplecta containing $\{x, p\}$ and $\{y, p\}$, respectively. Of course if $X = Y$, we are also done, so we can assume $X \neq Y$.

Suppose $y^\perp \cap X = \{x\}$. Then by (NG), y is gated with respect to X, with x being the *gate*. But since $x \in P_2(y)$, this would imply $d(y, p) = 3$ against $y \in P_2(y)$.

Thus we must suppose $y^\perp \cap X$ contains a line N incident with x. Then $p^\perp \cap N = \{z\}$, a single point of $y^\perp \cap p^\perp \subseteq Y$. Thus $z \in L^\perp$ and so $\langle z, L\rangle$ is a plane π carrying two distinct lines zx and zy which share a symplecton with the line pz. By (H2), pz shares a symplecton with every line of π on point z. This forces $L \subseteq P_2^*(p)$ as desired. Thus the claim is true and so $P_2^*(p)$ is a subspace of Γ. The proof is complete. □

A *hexagonic geometry* is a parapolar space of symplectic rank at least three satisfying these three axioms:

(H1) If p is a point and S is a symplecton, then $p^\perp \cap S$ is never a set consisting of a single point.
(H2) Suppose π is a plane (that is, a singular subspace of $\Gamma = (\mathcal{P}, \mathcal{L})$ which is a projective plane) and L is a line meeting π at a point p. Then L shares a symplecton with none, one, or all of the lines of π incident with point p.
(H3) Given a point-line flag (p, L) there exists a second line N such that $L \cap N = \{p\}$ and no symplecton contains $L \cup N$ — i.e., (x, y) is a special pair for each $(x, y) \in (L - \{p\}) \times (N - \{p\})$.

The hexagonic geometries with some finite singular rank assumptions are classified in Chap. 16.[7] The geometries that emerged were the classical polar spaces, the polar Grassmannians of lines, and the geometries of exceptional Lie types $F_{4,1}$, $E_{6,2}$, $E_{7,1}$, $E_{8,8}$ (using the Bourbaki numbering of nodes). This class is approximately the class of geometries known to group theorists as "long root" geometries. There are several reasons to insist on a distinction between "hexagonic geometries" and the "long root geometries" as discussed near the beginning of Chap. 16.

Corollary 13.7.6 *Suppose Γ is a hexagonic geometry. Then (P2) holds.*

Proof Property (H1) implies property (NG). Now apply Theorem 13.7.5. □

Theorem 13.7.7 *Suppose Γ is a hexagonic geometry. Then (D2) holds.*

Proof Suppose L is a line carrying two distint point x_2 and x_3 of $\Delta_2^*(p)$ for some point p. We must show that $L \subseteq \Delta_2^*(p)$.

We assume, by way of contradiction, that $L \not\subseteq \Delta_2^*(2)p$, and proceed by a series of steps.

[7] They were first classified in [85]. In their work, Hypothesis 3 of their Theorem 1 can be weakened to (H3) as demonstrated in Chap. 16.

Step 1. $p^\perp \cap L^\perp = \emptyset$.

This is true since otherwise there is a path of length at most two from p to each point of L.

Step 2. $P_2^*(p) \cap L = \emptyset$.

Suppose y were a point in $P_2^*(p) \cap L$. Clearly $y \notin p^\perp$ by Step 1. So there is some symplecton R containing p and y, and if $L \subseteq R$ we have the forbidden conclusion $L \subseteq \Delta_2^*(p)$ at once. So we may assume $L \cap R = \{p\}$. Now (H1) implies the existence of a line N on p lying in $L^\perp \cap R$. Since R is a polar space there is a point u in $p^\perp \cap N$. Then $u \in p^\perp \cap L^\perp$, contradicting Step 1.

Step 3. *No point of $p^\perp \cap {x_2}^\perp$ is in $P_2^*(x_3)$; similarly no point of $p^\perp \cap {x_3}^\perp$ belongs to $P_2^*(x_2)$.*

Because of the symmetry of x_2 and x_3 in their relation to p and L, we need only prove the assertion of the first clause. Suppose there is a symplecton R containing x_3 and a point $y \in p^\perp \cap {x_2}^\perp$. By Step 2, $p \notin R$. Then by (H1) $p^\perp \cap R$ contains a line N on y, and by Step 1, we see that $L^\perp \cap N = \emptyset$ implying that N and L are opposite lines of the symplecton R. Now any point e is collinear with a point of N, and as $N \subseteq p^\perp$, one obtains $L \subseteq \Delta_2^*(p)$.

Step 4. *One of the following two situations occurs:*

1. *There exists a point z in $L^\perp \cap P_2^*(p)$.*
2. *There exists a point $q \in p^\perp$ for which $L \subseteq P_2^*(q)$.*

Choose $x_1 \in p^\perp \cap {x_2}^\perp$ and $x_4 \in p^\perp \cap {x_3}^\perp$. By Step 2, (p, x_2) and (p, x_3) are special pairs, and so the points x_1 and x_4 are uniquely determined. Also, by Step 3, (x_1, x_3) and (x_2, x_4) are special pairs.

Let S be a symplecton containing line L. Then by (H1) there is a line N_2 on point x_2 lying in ${x_1}^\perp \cap S$, and similarly, there is a line N_3 on point x_3 lying in ${x_4}^\perp \cap S$. Since S necessarily has polar rank at least three, there exists a plane π of S meeting ${x_4}^\perp \cap S$ at line N_3.

Suppose π could be chosen so $\pi \cap {x_1}^\perp$ contained a point y. If $y \neq x_2$, ${x_1}^\perp \cap x_3$ contains the distinct points z and x_2 against (x_1, x_3) being special. If $y = x_2$ then ${x_2}^\perp \cap {x_4}^\perp$ contains the line N_3 and so (x_2, x_4) is a polar pair against the observation of the paragraph preceding the previous one. It follows that N_2 and N_3 are opposite lines of the symplecton S. Thus there is a unique point z in ${N_2}^\perp \cap \pi$.

Now (z, x_2) and (z, x_4) are clearly polar pairs (because of the presence of the lines N_2 and N_3, respectively), and so we may form the symplecta

$$R_2 := \langle\langle x_1, z\rangle\rangle \text{ and}$$
$$R_3 := \langle\langle x_4, z\rangle\rangle$$

which contain the respective lines N_2 and N_3. Now $p^\perp \cap R_2$ contains the point x_1 and so, by (H1), contains a line M_2 which bears a unique point q_2 in $z^\perp$; similarly by (H1) $p^\perp \cap R_3$ must contain a line M_3 on x_4 bearing a point q_3 collinear with z.

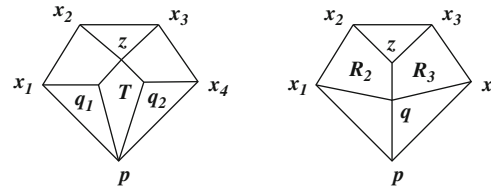

Fig. 13.2 The two configurations which emerge in Step 4

Now we have two cases: (1) if $q_2 \neq q_3$ then these are distinct points of $z^\perp \cap p^\perp$ and so $z \in P_2^*(p) \cap L^\perp$, the first conclusion of Step 4, and (2) if $q_1 = q_2 := q$, then $L_0 := qz$ is a line meeting the plane $\pi' := \langle z, L\rangle_\Gamma$ at the point z. But now L_0 shares a symplecton R_i with the distinct lines zx_i of π on z, for $i = 2, 3$. For the first time we invoke axiom (H2) to conclude that L_0 shares a symplecton with every single line of π' on point z. It follows that

$$L \subseteq P_2^*(q).$$

Thus the second conclusion of Step 3 holds in this case.

The final configurations for both of these cases are presented in Fig. 13.2. (The intermediate constructions S, N_2, N_3 and π are not drawn.)

Step 5. *The first conclusion of Step 4 cannot occur.* Assume there is a point $z \in P_2^*(p) \cap L^\perp$. Then there exists a symplecton T containing p and z. Then for any point $e \in L$, $e^\perp \cap T$ contains a line M_e which meets $p^\perp$, forcing $d(e, p) \leq 2$, for all $e \in L$. Thus $L \subseteq \Delta_2^*(p)$ against our initial hypothesis to the contrary.

Step 6. *The second case of Step 4 cannot occur.* Assume there is a point $q \in p^\perp$ such that $L \subseteq P_2^*(q)$. Then for every point $e \in L$, there is a symplecton R_e containing q and e. Since $q \in p^\perp$, one sees that $p^\perp \cap R_e$ contains a line on q which carries a point of $e^\perp$, forcing $d_\Gamma(p, e) = 2$. Thus $L \subseteq \Delta_2^*(p)$, against our original assumption to the contrary.

Now Steps 4, 5, and 6 are mutually incompatible, and so the assumption that L is not contained in $\Delta_2^*(p)$ is untenable. The theorem is proved. □

13.8 Generalized Parapolar Spaces

13.8.1 Introduction

Parapolar spaces are useful basically because the Lie incidence geometries of the spherical buildings, those whose points are the cosets of a maximal parabolic subgroup, are either a projective space or a parapolar space. After summarizing in *The Handbook for Incidence Geometry* [35] the point-line characterizations that were known in 1994, Arjeh Cohen makes the following observation on p. 721: "Very little has been done with regards to the shadow spaces of affine type" and then he lists the only three such results in this area known at the time.

There is a very good reason for this. With only a few exceptions, Lie incidence geometries of affine or other non-spherical types are not actually parapolar spaces. Yes, they have a rich system of symplecta covering all the lines. The difference is in the possible options that are available to two points at distance two which do *not* lie in a symplecton. In a parapolar space such pairs of points $\{x, y\}$ were called "special pairs" and had the property that there was a unique point collinear with both x and y – that is, $|x^\perp \cap y^\perp| = 1$. In the non-spherical case, there exist "special pairs" – that is, pairs of points $\{x, y\}$ at distance two not lying in a symplecton – for which $x^\perp \cap y^\perp$ is a singular space larger than a single point.

So it is desirable to have a generalization of the notion of parapolar space which accomodates such geometries derived from buildings. But we must do this in a systematic way.

Suppose first that $\Gamma = (\mathcal{P}, \mathcal{L})$ is a gamma space We have noted before[8] that for every point p, there is a "tangent graph," $(\mathcal{L}_p, \sim)$, whose vertices are the lines of Γ which are incident with point p, two such lines on p being adjacent in the graph if and only if together they lie in a common singular subspace of Γ. By the phrase "Γ is locally G" where G is a class of of point-line geometries, we understand that every tangent graph $(\mathcal{L}_p, \sim)$ is isomorphic to the point-collinearity graph of a point-line geometry from G.

Next suppose Γ is "locally a parapolar space." Since a parapolar space is a partial linear space, there is a class of "planes" on p which pairwise intersect at a line, or at the single point p itself. Moreover, every maximal singular subspace M of Γ locally "looks like" a projective space – that is the lines and planes incident with p and lying in M have the structure of a projective space. Is this enough to make M or any singular subspace of Γ a projective space? The answer is "No." The *affine polar spaces* studied by Cohen and Shult [38] are gamma spaces which are locally polar spaces, in the sense just defined. They are obtained by removing a geometric hyperplane from a non-degenerate polar space.[9] All of the singular subspaces of such a space are affine spaces. One can obtain similar locally parapolar gamma spaces whose singular subspaces are "mixed" – either affine or projective – by removing a suitable subspace from any Lie incidence geometry of symplectic rank at least three.

So, in order to achieve a framework encompassing non-spherical Lie incidence geometries, we must assume that Γ is a gamma space for which all singular subspaces are projective spaces – the local information is just not enough to ensure this. In his survey in the *Handbook of Incidence Geometry* [35], Cohen has coined the term *paraprojective space* for a gamma space all singular subspaces of which are projective spaces.[10] Previously there has been little need to invent a special word for this class of gamma spaces, but we seem to have acquired the need at this point.

---

[8] In our discussion of local connectedness of gamma spaces in Chap. 3.

[9] All the terms of this sentence should be well understood from Chaps. 4 and 7.

[10] At first glance the words seem to suggest a type of linear space generalizing a projective space. Nothing could be further from what is intended. On the other hand I can't think of a handy alternative.

Now when we say that a paraprojective space (still a very broad class of gamma spaces) is "locally G" we can say more: in a tangent graph, each clique which represented line of G is now the pencil of lines on p in some projective plane, and so now the "tangent" geometry $\mathrm{Res}_\Gamma(p) := \Gamma$ is the geometry $(\mathcal{L}_p, \Pi_p)$ of lines and projective planes of Γ which are incident with p – what we have been calling a *point-residual*. It is the truncation of the residue $\mathrm{Res}_\Gamma(p)$ of a point p to lines and planes (see Chap. 2, p. 46 for the definition).

Next there is the question of possible symplecta. If the tangent geometry $\mathrm{Res}_\Gamma(p)$ has symplecta, we would like these to be point-residuals of symplecta of Γ. As we shall see below, when Γ is locally parapolar, this problem will take care of itself. The real problem is the existence of "quads" – that is, symplecta of polar rank two. Each such quad shows up in a point-residual $\mathrm{Res}_\Gamma(p)$ as a partial-linear system of cocliques, with each exterior "point" collinear to at most one point of the coclique. As far as I know, such a system does not make an appearence in classical parapolar spaces: usually when there are "hyperbolic lines" like that, there are also exterior points collinear with all points of the hyperbolic line. At least we see that the existence of quads in a paraprojective space that is locally parapolar is not ruled out a priori just by the local tangent parapolar spaces (unless, possibly, when we know what these parapolar spaces actually are). So we should assume that all symplecta (if there are any) have polar rank at least three.

Finally, there is the assumption that any 4-circuit without diagonals lies in some symplecton. We need to have these irreducible circuits controlled by the symplecta in order to be able to enrich the geometry suitably. Again, it has to be assumed, since the locally parapolar hypothesis says nothing here. There is a special way of insuring this. We simply require that Γ is *nexus connected*, which is the condition that if x and y are two non-collinear points, then the subspace $x^\perp \cap y^\perp$ is either empty or is a connected subspace.

We now have the following.

Lemma 13.8.1 *Suppose Γ is a paraprojective space with these properties:*

1. *The geometry Γ is nexus connected.*
2. *Every point-residual of Γ is a parapolar space of symplectic rank at least three.*

Then Γ also has these properties:

1. *Γ is locally connected.*
2. *Let $\{x, y\}$ be a pair of points at distance two in the point-collinearity graph of Γ. Then either:*
   - *(a) $x^\perp \cap y^\perp = \emptyset$,*
   - *(b) $\{x, y\}$ is a "special pair": specifically $x^\perp \cap y^\perp$ is a point or a line, or*
   - *(c) $\{x, y\}$ is a "polar pair": that is, $x^\perp \cap y^\perp$ is a non-degenerate polar space of rank at least three.*
3. *Every polar pair $\{x, y\}$ lies in a symplecton of rank at least four. These comprise all the symplecta on Γ.*
4. *Any $PG(3)$-subspace lies in some symplecton.*

Proof Part 1. This is because by definition, all parapolar spaces are connected.

Part 2. Now let Σ be the collection of all distance-two point-pairs, $\{x, y\}$ for which $x^\perp \cap y^\perp$ is not a singular subspace. By the hypothesis on nexus connectedness, there exists a non-collinear pair of points $\{u, v\}$ in $x^\perp \cap y^\perp$ for which there exists a point w in $x^\perp \cap u^\perp \cap y^\perp \cap v^\perp$.

Thus w is a point of $x^\perp \cap y^\perp$ for which $w^\perp \cap x^\perp \cap y^\perp$ contains a 2-coclique, namely $\{u, v\}$. Then in the point-residue $\mathrm{Res}_\Gamma(w)$ one discerns a 4-circuit of lines (wu, wy, wv, wx, wu) which, by the local hypothesis, lies in a symplecton R_w of rank at least three of $\mathrm{Res}_\Gamma(w)$. Thus the subspace $(wx)^\perp \cap (wy)^\perp$ is a non-degenerate polar space of rank at least 2.

Next let w' be any point of $x^\perp \cap y^\perp$ collinear with w but distinct from it. Then since the lines and planes on w lying in $x^\perp \cap y^\perp$ form a non-degenerate polar space of rank at least two, $(w')^\perp \cap x^\perp \cap y^\perp$ also contains a 2-coclique, and the arguments that were applied to w in the previous paragraph can be recycled for w' to yield the conclusion that the lines and planes of $x^\perp \cap y^\perp$ which are on w' also form a non-degenerate polar space of rank at least two.

Since $x^\perp \cap y^\perp$ is connected, we conclude that it is a paraprojective space which is locally a non-degenerate polar space of rank at least two. By the theorem of Johnson and Shult [81] (see p. 236), the subspace $x^\perp \cap y^\perp$ is a non-degenerate polar space of rank at least three.

Part 3. Now we see that the collection Σ of all distance two point-pairs $\{x, y\}$ for which $x^\perp \cap y^\perp$ is not a singular subspace forms a polar family as defined on p. 455. By Theorem 13.2.16, each such pair of points lies in a symplecton of polar rank at least four.

To prove the last sentence in this part, suppose S is any symplecton of Γ and suppose x and y are two non-collinear points of S. If S had polar rank exactly two, then $x^\perp \cap y^\perp$ would not be connected, contrary to hypothesis. Thus no such rank-two symplecta ("quads") exist. If S has rank at least three, then $x^\perp \cap y^\perp$ is not a singular space and so $\{x, y\} \in \Sigma$, forcing S to be one of the symplecta of polar rank at least four derived from Σ.

Part 4. Now it is clear that if p is a point in a symplecton S then the lines and planes of S on p form a symplecton of the residue $\mathrm{Res}_\Gamma(p)$, and, conversely, every symplecton of the point-residue, $\mathrm{Res}_\Gamma(p)$, is the residue of a point in some symplecton S of Γ. Since these point-residues are parapolar spaces of polar rank at least three, any plane of such a residue lies in some symplecton.[11] It follows that any $PG(3)$ containing p lies in a symplecton S of Γ. Since p is arbitrary, the assertion of part 4 holds. □

[11] Note that this would not be true if the parapolar space $\mathrm{Res}_\Gamma(p)$ possessed symplecta of rank two.

13.8.2 Generalized Parapolar Spaces of Nexus 1

The Definition

The preceding analysis motivates the following definition.

A generalized parapolar space of nexus 1 is a connected paraprojective space $\Gamma = (\mathcal{P}, \mathcal{L})$ with the following properties:

1. Γ is locally connected.
2. If x and y are any two non-collinear distinct points of $\mathcal{P}$, then $x^\perp \cap y^\perp$ is either empty, a point or a line, or $\{x, y\}$ lies in a symplecton of polar rank at least three.
3. Every plane lies in a symplecton.

How to Recognize These Spaces

As in the case of parapolar spaces, we say such a space *has symplectic rank at least r* if and only if every symplecton is a polar space of polar rank at least r. The following two theorems allow us to recognize a generalized parapolar space of nexus 1.

Theorem 13.8.2 *Suppose Γ is a paraprojective space with these properties:*

1. *Γ is locally connected.*
2. *For every plane π, $\pi^\perp$ is not a singular subspace.*
3. *If x and y are any two non-collinear distinct points of $\mathcal{P}$, then $x^\perp \cap y^\perp$ is either empty, a point or a line, or or is a non-degenerate polar space of rank at least two.*

Then Γ is a either a parapolar space of symplectic rank at least three, or is generalized parapolar space of nexus at most one.

Proof One observes that the collection Σ of all pairs of non-collinear points $\{x, y\}$ for which $x^\perp \cap y^\perp$ is a non-degenerate polar space of rank at least two is a polar family, and an application of Theorem 13.2.16 shows that the convex closure of such a pair is a symplecton of polar rank at least three.

It remains only to show that every plane lies in a symplecton. But by assumption, for each plane π, $\pi^\perp$ must contain a pair $\{x, y\}$ of distinct non-collinear points. Then as $x^\perp \cap y^\perp$ properly contains a line, $\{x, y\}$ belongs to the polar family Σ, and so the convex closure of this pair is a symplecton S which, being convex, contains the plane π.

The proof is complete. □

Theorem 13.8.3 *Suppose $\Gamma = (\mathcal{P}, \mathcal{L})$ is a paraprojective space with these properties:*

1. *Γ is locally connected.*
2. *Γ is nexus connected.*

3. *For any plane π, $\pi^\perp$ is not a singular subspace.*
4. *Whenever π_1 and π_2 are two planes not lying in a common singular subspace, but contain a common line L, then $\pi_1{}^\perp \cap \pi_2{}^\perp$ is either equal to L or is equal to a non-degenerate polar space of rank at least two.*

Then Γ is a generalized parapolar space of nexus one and symplectic rank at least four.

Proof It follows from the hypothesis that any point-residual $\mathrm{Res}_\Gamma(p)$ is a parapolar space of symplectic rank at least three. Lemma 13.8.1 now implies all the required conditions for Γ to be a generalized parapolar space of nexus 1 and symplectic rank at least four. □

Some Examples

Let $\Gamma = (\mathcal{P}, \mathcal{L})$ be one of the point-line geometries obtained from a building B with one of the two diagrams in Fig. 13.3. The rank-two truncations to points and lines are indicated by the symbols $\mathcal{P}$ and $\mathcal{L}$. These spaces are generalized parapolar spaces of nexus 1. In the case of Fig. 13.3a, the space has symplectic rank at least four. That means the geometry could conceivably be characterized using Theorem 13.8.3. In the case of Fig. 13.3b, the geometry has symplectic rank three exactly – that is, the symplecta of the point-residuals are quads. Here, Theorem 13.8.3 cannot be used to produce the existence of symplecta from purely local data.

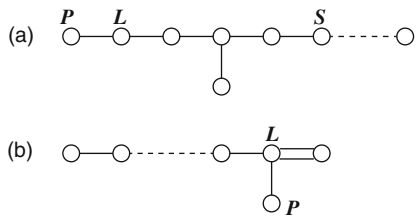

Fig. 13.3 Two Lie incidence geometries which are generalized parapolar spaces of nexus one. The point-line geometry obtained from (**a**) has symplectic rank four, while that obtained from (**b**) has symplectic rank three

13.8.3 Generalized Parapolar Spaces of Higher Nexus

Let Γ be a paraprojective space and let k be an integer greater than -2. We say that Γ is *k-connected* if and only if, for each integer $\ell \leq k$, and each singular subspace A of projective dimension $\ell - 1$, the collection $\mathrm{Res}(A)$ of singular subspaces of projective ranks ℓ and $\ell + 1$ which contain A form a connected geometry. Note that a paraprojective space is connected if and only if it is 0-connected, and is connected and locally connected if and only if it is 1-connected.

We begin our general definition with the following.

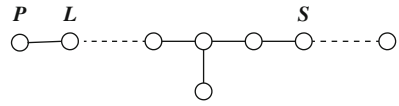

Fig. 13.4 The point-line geometry obtained from the indicated truncation of a building belonging to the depicted diagram is a generalized parapolar space of arbitrarily high nexus

A *generalized parapolar space of nexus k* is a k-connected paraprojective space with these properties:

1. For each pair of points $\{x, y\}$ at distance two, $x^\perp \cap y^\perp$ is either empty, a singular subspace of projective rank at most k, or a symplecton of polar rank at least $k + 2$.
2. Every singular subspace of projective rank $k + 1$ lies in a symplecton.

Buildings belonging to the diagram displayed in Fig. 13.4, when truncated to points $\mathcal{P}$ and lines $\mathcal{L}$, form generalized parapolar spaces of arbitrarily high nexus.

One should note that from the definition, a generalized parapolar space of nexus 0 is just an ordinary parapolar space, and that a generalized parapolar space of nexus -1 is in fact a strong parapolar space.

13.9 Exercises for Chapter 13

13.9.1 Exercises Concerning Parapolar Spaces

13.1 Consider a product geometry $(\mathcal{P}_1, \mathcal{L}_1) = L \times L \times L$, where L is a line. Choose a point a and let G be the subspace $\{(a, x, y) | x \in L, y \in L\}$. Then G is a *grid* – that is, a generalized quadrangle with each point on exactly two lines. We are now going to add a set of lines $\mathcal{L}_2$ to this grid G to make it an affine plane. Then the geometry $\Gamma := (\mathcal{P}_1, \mathcal{L}_1 \cup \mathcal{L}_2)$ contains G as a singular subspace. Show that Γ is not a parapolar space and identify the axiom that is not satisfied.

13.2 Now let A be any linear space that is not a projective space. For each line L of A we adjoin a grid G_L intersecting A at line L and otherwise being independent – that is, for any two lines L and N of A, $G_L \cap G_N = L \cap N$. Then let Γ be the union of A and all the G_Ls, along with the lines which come with these subspaces. Show that Γ is a parapolar space of symplectic rank two.

13.3 Let Δ be a non-degenerate polar space of rank $n \geq 4$. We consider a geometry Γ whose points $\mathcal{P}$ are the singular subspaces of Δ of projective dimension $n - 2$ – that is, they are hyperplanes of maximal singular subspaces of Δ. The lines $\mathcal{L}$ of Γ will be the flags (A, B) where A and B have projective dimensions $n - 3$ and $n - 1$, respectively. Show the following:

1. The geometry Γ is a parapolar space of which all symplecta are generalized quadrangles (these correspond to the singular subspaces of Δ of projective dimension $n-3$).
2. The maximal singular subspaces of Δ (which have projective dimension $n-1$) correspond to a class of maximal singular subspaces $\mathcal{M}$ of Γ also of projective dimension $n-1$.

13.4 Suppose Γ is a parapolar space of symplectic rank three with the property that every point-residual $\mathrm{Res}_\Gamma(p) = (\mathcal{L}_p, \Pi_p)$ is a dual polar space of rank at least three.

1. Show Γ can be enriched to a rank four geometry $\Gamma^* = (\mathcal{P}, \mathcal{L}, \mathcal{A}_2 = \Pi, \mathcal{S})$ of points, lines, planes, and symplecta.
2. Using this enriched geometry Γ^*, show that any two point-residuals are dual polar of the same polar rank.
3. Show that Γ^* is a locally truncated geometry with respect to the multiply extended F_4 diagram (Fig. 13.5)
4. From a theorem in Chap. 11, justify that a sheaf exists. (You must verify the requisite rigidity conditions.)
5. Justify the conclusion that Γ is a homomorphic image of a Lie incidence geometry of a building with the extended F_4-diagram depicted just above. Moreover show that the homomorphism is injective when restricted to any symplecton and any singular space.

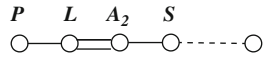

Fig. 13.5 The multiply-extended F_4-diagram

13.5 Suppose Γ is a gamma space of which all singular subspaces are projective spaces.[12] Then for each point p we can form a point-residual $\mathrm{Res}_\Gamma(p) = (\mathcal{L}_p, \Pi_p)$ of all lines and planes on p. Assume:

1. For every pair of non-collinear points $\{x, y\}$, the subspace $x^\perp \cap y^\perp$ is either empty or is a connected subspace. (We met this condition in this chapter. Any gamma space with this condition is said to be *nexus-connected*.)
2. For each point p, $\mathrm{Res}_\Gamma(p)$ is a strong parapolar space of symplectic rank at least three.

Show that Γ is a parapolar space. [Hint: Suppose $\{x, y\}$ is a pair of points at distance two in the point-collinearity graph of Γ. Using the connectedness of $x^\perp \cap y^\perp$ observe that if this space is not a single point, it is a gamma space with

[12] Recall that these are called *paraprojective spaces* in [35].

projective singular spaces whose point-residuals are non-degenerate polar spaces of rank at least two. Use the Buekenhaut–Johnson–Shult theorem [81] (Theorem 7.10.3, p. 236) to conclude that $x^{\perp} \cap y^{\perp}$ is a polar space of rank at least three.]

13.6 Suppose Γ is a connected and nexus-connected gamma space of which all singular subspaces are projective. Suppose each point-residual $\mathrm{Res}_{\Gamma}(p)$ is a half-spin geometry of rank at least five.

1. Show that Γ is a parapolar space.
2. Show that Γ can be enriched to a geometry $\Gamma^* = (\mathcal{P}, \mathcal{L}, \mathcal{A}_2, \mathcal{M}, \mathcal{A}_4, \mathcal{S})$ of point, lines, planes, a class $\mathcal{M}$ of maximal singular spaces of projective rank at least five, a class $\mathcal{A}_4$ of maximal singular subspaces of projective rank four, and a class of oriflame symplecta of polar rank five, respectively.
3. Using this enriched geometry Γ^* show that if p and q are distinct collinear points of Γ then their point-residuals are isomorphic.
4. Show that Γ^* is a locally truncated geometry with respect to the diagram on $n \geq 6$ nodes which appears in Fig. 13.6.
5. Using Theorem 11.3.2 of Chap. 11, conclude that a sheaf exists for this locally truncated geometry Γ^*. (One must check the required rigidity conditions.)
6. Justify the conclusion that Γ is a homomorphic image of the Lie incidence geometry obtained from a building belonging to the diagram presented in the previous part – with points and lines as indicated. Moreover, show that the homomorphism is injective when restricted to any symplecton and any singular space of the Lie incidence geometry.

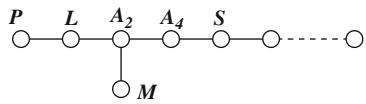

Fig. 13.6 Figure for Exercise 13.6

13.7 Suppose Γ is a parapolar space. Let A be a maximal singular subspace of Γ and suppose, for some point $x \in \mathcal{P} - A$, we have $x^{\perp} \cap A \simeq PG(d, K)$ where d is a positive integer. Show that for any arbitrary point $y \in A$, and subspace $B \simeq PG(d, K)$ contained in A and on point y, there exists a point $z \in \mathcal{P} - A$ such that $z^{\perp} \cap A = B$. [Hint: Use the fact that any pair of points at distance two lies in a symplecton. This takes several stages.]

13.8 Suppose Γ is a parapolar space. Suppose A is a maximal singular subspace of Γ (of possibly infinite projective rank at least 3) and suppose, for some line L in M every $PG(3, K)$ on L lies in A. Then it is true that every $PG(3, K)$ whose intersection with A contains a plane actually lies in A.

13.9 Suppose n is a (possibly infinite) cardinal number at least 2, and that Γ is a parapolar space, for each point p, the point-residue $\mathrm{Res}(p)$ has the form $PG(n-1, K) \times X$, where, for any maximal singular subspace A' of X which is isomorphic to $PG(n-1, K)$, there exists a "point" L of A' with the property that for "line" π on L there exists a "plane" $PG(3; \pi)$ of X which is not in A but contains P . Then show that the same conclusion it true for all "points" L' of A'. [Hint: Use Exercise 13.8.]

13.10 Suppose n is a (possibly infinite) cardinal number greater than 1, and that Γ is a parapolar space each of whose point-residuals $\mathrm{Res}(p)$ has the form $PG(n-1, K) \times X$, where, X has no locally connected component isomorphic to $PG(n-1, K)$. Show that there is a class $\mathcal{M}$ of maximal singular subspaces $PG(n-1, K)$s with the property that for every line $L \in \mathcal{L}$, L lies in a unique member of $\mathcal{M}$. (This means we have complete global separation of singular subspaces in the sense of the previous chapter.) [Hint: Use Exercise 13.9 to verify the sufficient condition for global separation given in Theorem 12.3.1 of Chap. 12.]

13.11 Suppose $\Gamma = (\mathcal{P}, \mathcal{L})$ is a parapolar space of symplectic rank at least three. For each point p, the point-residual $\mathrm{Res}(p) = (\mathcal{L}_p, \Pi_p)$ has the form $L \times Q$ where L is a "line" and Q is a generalized quadrangle not isomorphic to a grid (that is, each of the "points" of Q lies in at least three "lines" of Q). Show that there is a global separation of the planes of Γ – specifically, that there is a partition of the planes into two classes

$$\Pi = \mathcal{M} + \Sigma,$$

such that for each point p, the planes $\mathcal{M}_p$ and Σ_p of each class which are incident with p comprise respectively the "horizontal lines" $\{L \times y | y \in Q\}$ and the "vertical lines" $\{x \times A | A$ a "line" of $Q\}$ in the product $\mathrm{Res}(p) \simeq L \times Q$. [The point of this exercise is that by using the hypothesis that Γ is a parapolar space one is able to obtain global separation of the maximal singular subspaces (in this case the planes) without the aid of either Theorem 12.3.1 or its parapolar-space corollaries, and Exercises 13.9 and 13.10. Hint: There are two classes of symplecta: (1) a class $\mathcal{D}$ of oriflame sympleca of rank three and (2) a class $\mathcal{S}$ of rank three symplecta which are not oriflame. The planes of $\mathcal{M}$ lie only in symplecta belonging to $\mathcal{D}$, while those belonging to Σ lie in one symplecta from each class.]

13.9.2 Exercises Concerning an Application of Hanssens' Principle

In these exercises we suppose Γ is a parapolar space of symplectic rank at least three with this property:

(M-S) *For each point p there exists a cardinal number $n_p \geq 1$ such that the point-residual at p has the form*

$$Res(p) \sim M_p \times S_p,$$

where $M_p \simeq PG(n_p, K)$, K is a division ring independant of the choice of p, and S_p is a non-degenerate polar space whose polar rank r_p is an integer greater than 1 or is the symbol "∞" – indicating that S_p is a polar space of infinite rank – and if $r_p = 2$ then either $n_p > 1$ or else $n_p = 1$ and S_p is not a grid.

13.12 Use Exercise 13.10 (and Exercise 13.11 when $n_p = 1$ and $r_p = 2$) to show that there is a global separation of all singular subspaces of Γ of projective rank at least 2. In particular there is a class $\mathcal{M}$ of maximal singular subspaces with the property that for each point p in a member M of $\mathcal{M}$, the lines and planes of M which are incident with p form one of the subspaces $M_p \times s$ in $\mathrm{Res}(p) \simeq M_p \times S_p$. Every line lies in just one member of $\mathcal{M}$ and so any two members of $\mathcal{M}$ intersect at the empty set or a point. [Hint: S_p can have no locally connected component isomorphic to M_p. Note how $r_p > 2$ is used here.]

13.13 From this, conclude that all elements of $\mathcal{M}$ possess the same projective rank m – that is $m - 1 = m_p$ for all points p.

13.14 Using the global separation established in Exercises 13.10 and 13.11, and the connectedness of Γ, show that all the S_ps possess a common polar rank $r-1$.

13.15 Show that all of the hypotheses of Hanssens' principle are present, and so there is a class $\mathcal{G}$ of line Grassmannians defined by the lines of $PG(m+1, K)$ (remember m can be infinite).

13.16 Show that there are two classes of symplecta, $\mathcal{D}$ and $\mathcal{S}$ in Γ where the symplecta in $\mathcal{S}$ possess only "Σ-planes" (in the sense of Theorem 13.6.1) – that is, a symplecton in $\mathcal{S}$ can share no plane with a maximal singular subspace in $\mathcal{M}$.

13.17 Now assume that the parameters m and r (the projective ranks of elements of $\mathcal{M}$ and the polar ranks of the elements of $\mathcal{S}$, respectively) are both finite. Show that the enriched geometry $\Gamma^* := (\mathcal{P}, \mathcal{M}, \mathcal{S}, \mathcal{G})$ is a locally truncated geometry and draw its associated diagram
(Note that each line lies in a unique element of $\mathcal{M}$ and a unique element of $\mathcal{S}$, and so the lines can be identified with $(\mathcal{M}, \mathcal{S})$-flags.)

13.18 Under the hypothesis of Exercise 13.17, show that a sheaf exists for the locally truncated geometry Γ^*. [Again the two rigidity assumptions must be noted.] Sketch how one then arrives at the conclusion that the original geometry Γ is a homomorphic image of a polar Grassmannian of m-subspaces of a non-degenerate polar space of rank $n := m + r - 1$.

13.9.3 Exercises for Section 13.7

13.19 Prove Theorem 13.7.3.

13.20 Prove that any strong parapolar space of symplectic rank at least three is locally connected.

13.21 Assume Γ is a parapolar space with property (NG). Suppose p is a point and L is a line meeting $P_2^*(p)$ at a point. Show that either $L \cap \Delta_2^*(p) = L \cap P_2^*(p)$ or $L \subseteq \Delta_2^*(p)$.

13.9.4 Exercises for Section 13.8: On Generalized Parapolar Spaces of Higher Nexus

13.22 If Γ is a generalized parapolar space of nexus k where $k \geq 0$, show that any point-residual $\mathrm{Res}_\Gamma(p)$ is a generalized parapolar space of nexus $k-1$.

13.23 Give analogs of Theorems 13.8.2 and 13.8.3 for recognizing the generalized parapolar spaces of nexus k for $k > 1$.

13.24 Is the definition of generalized parapolar space too restrictive? Specifically are there examples of paraprojective spaces for which $x^\perp \cap y^\perp$ can be either singular subspaces of projective rank as high as k, or a polar space of far lower polar rank? (Here the symplecta could be of rank as low as four, while $x^\perp \cap y^\perp$ could be a singular subspace of projective rank four or larger.)

13.9.5 Symplecta Which Cannot Intersect at a Single Point

Suppose Γ is a strong parapolar space with this property:

(*) If S_1 and S_2 are any two symplecta of Γ, it is impossible that $S_1 \cap S_2$ be a single point.

(In Chap. 15 the reader will encounter a Lie incidence geometry called $E_{7,7}$ which is a strong parapolar space with property (*).)

Show the following.

13.25 If p is a point and S is a symplecton such that $p^\perp \cap S = \{x\}$, a single point, then every point of $S - x^\perp$ is distance three from point p. (In terms of the point-collinearity graph Δ and the convention that subspaces are defined as

sets of points with certain properties, this assertion can be stated as follows: *the symplecton S is gated with respect to the point p with gate x*.)

13.26 The point-collinearity graph Δ has diameter at most three. [Hint: Imagine what would happen with a geodesic path of Δ of length four.]

13.27 One has the following condition (D2). *For any point p, the collection $\Delta_2^*(p)$ of all points of distance at most two from p form a subspace of Γ.*

13.28 For every line L, and every point p one has

$$\Delta_2^*(p) \cap L \neq \emptyset.$$

It then follows that $\Delta_2^*(p)$ is always a geometric hyperplane of Γ. [Hint: Suppose x and y are two distinct points at distance three from a point p lying on a common line L. One only needs to show that L carries a point at distance two from p. Suppose (p, a, b, x) and (p, c, d, y) are geodesic paths from p and consider the unique symplecton R on $\{x, b\}$, and a symplecton S on $\{b, d\}$. Their intersection cannot just be $\{b\}$.]

Part IV
Applications to Other Lie Incidence Geometries

Chapter 14
Characterizations of the Classical Grassmann Spaces

Abstract Although Grassmannians are parapolar spaces, here are three characterizations of Grassmann spaces which do not assume that fact. The first starts with a point-line geometry with a partition $\mathcal{M} = \mathcal{A} + \mathcal{B}$ of the maximal singular subspaces, such that each line lies on exactly one space from each component and that a certain intersection property holds. Finally, a member of one of the classes, say $\mathcal{A}$, is assumed to possess a finite saturated chain of subspaces. It is not assumed that the members of $\mathcal{M}$ are projective spaces; this fact emerges by showing the existence of symplecta via Cooperstein's theory of the previous chapter. A second characterization utilizes the first through a "diagonal axiom." The third recasts a theorem of Sprague, replacing his assumption of a line with finitely many points with the existence of a finite saturated chain of subspaces in one of the residues. Sprague's "Linear-Dual Linear" theorem appears in the exercises in this infinite context. A footnote reveals how an implicit assumption of (Int) has caused a misreading of his theorem.

14.1 The First Characterization

14.1.1 Axioms on Two Families of Maximal Singular Subspaces

Theorem 14.1.1 *Assume $(\mathcal{P}, \mathcal{L})$ is a connected point-line geometry with thick lines satisfying these assumptions:*

(T1) *Any three mutually collinear points lie in a singular subspace (equivalently, $(\mathcal{P}, \mathcal{L})$ is a gamma space).*

(T2) *The maximal singular subspaces of $(\mathcal{P}, \mathcal{L})$ fall into two families Σ and Π such that:*

- *(a) Any member of Σ meets any member of Π at a line or at the empty set.*
- *(b) Every line lies properly in exactly one member of Σ and one member of Π.*

E.E. Shult, *Points and Lines*, Universitext, DOI 10.1007/978-3-642-15627-4_14,
© Springer-Verlag Berlin Heidelberg 2011

(T3) (An intersection property.) *If two distinct members P_1 and P_2 of Π each intersect non-trivially two distinct members S_1, S_2 of Σ, then $P_1 \cap P_2$ is non-empty.*

(T4) (Finite rank.) *Some member P of Π has an unrefinable ascending chain of non-empty subspaces of finite length $d+1$ ending at P.*

Then, if $\mathcal{L}_0$ is the set of lines which are actually non-empty intersections of members of Π with members of Σ, then $(\mathcal{P}, \mathcal{L}_0)$ is the classical Grassmann space whose points are the d-subspaces of some vector space.

Remark The vector space of the conclusion may be over any division ring and need not be finite-dimensional. If, however, the members of Σ have finite singular rank, then $(\mathcal{P}, \mathcal{L})$ is a Lie incidence system of type $A_{n,d}(D)$, $d > 1$.

Remark I have taken the liberty of designating the axioms with the letter "T" in honor of G. Tallini who was the first to introduce a hypothesis of this type in order to characterize the Grassmannian of lines of a projective space [127]. His original characterization was extended to a characterization of all Grassmannians of finite singular rank by Shult [108] and by Bichara and Tallini [8] at virtually the same time. Although equivalent, the two characterizations differ in the degree to which provable conclusions are already incorporated in the axioms. For example, in the Bichara–Tallini presentation it is already assumed that $(\mathcal{P}, \mathcal{L})$ is a partial linear space, and that singular spaces are projective spaces. Neither of these assumptions is necessary. For this reason we have followed the proof given in [108].

Proof We begin with several initial observations. First, by axiom (T2) (a) and (b), every line $L \in \mathcal{L}$ lies in a member of Σ and a member of Π and in fact the intersection of these spaces is a possibly larger line belonging to $\mathcal{L}_0$ as well as $\mathcal{L}$. Thus $\mathcal{L}_0$ is precisely the maximal members of $\mathcal{L}$. From here on in we ignore the possibly smaller lines. They play no further role in either the conclusion or the argument that leads to it.

Thus we observe:

1. $\mathcal{L}$-collinearity is coextensive with $\mathcal{L}_0$-collinearity.
2. $(\mathcal{P}, \mathcal{L}_0)$ is a partial linear space which is a gamma space.

The first statement follows from (T2)(b), and the second from (T1) and the first statement.

Let Δ be the point-collinearity graph of $(\mathcal{P}, \mathcal{L})$ (or $(\mathcal{P}, \mathcal{L}_0)$, they are the same graph). We shall let xy denote the unique line of $\mathcal{L}_0$ on points x and y whenever (x, y) is an edge in the graph Δ.

Suppose p and q are two points at distance 2 in the graph Δ. Then there is a point r in $p^\perp \cap q^\perp$ so that one can form the lines pr and rq. Let P_{pr} be the member of Π on pr and let S_{rq} be the unique member of Σ on line rq. Then $P_{pr} \cap S_{rq}$ is a line of $\mathcal{L}_0$ in $p^\perp \cap q^\perp$ on point r. Thus it is fair to say that every point of $p^\perp \cap q^\perp$ lies on some line of that subspace.

Now suppose L is an arbitrary line of $\mathcal{L}_0$ in $p^\perp \cap q^\perp$. By symmetry of p and q, we can assume that the unique maximal singular subspace containing $\langle p, L\rangle$ is an element S_p of Σ. Then $\langle L, q\rangle \subseteq P_q \in \Pi$. Suppose z is any point of $p^\perp \cap q^\perp$ not in line L. If P_p is the unique element of Π containing line pz and S_q is the unique member of Σ containing line qz, then the four intersections:

$$P_p \cap S_q,\ P_p \cap S_p,\ P_q \cap S_q,\ P_q \cap S_p,$$

are all lines of $\mathcal{L}_0$. Then by axiom T(3) $P_p \cap P_q = \{w\}$. Now $\{x\} \cup L \cup \{w\}$ is a clique for $x = p$ and for $x = q$. This forces w to belong to $S_p \cap P_q = L$. Thus $x^\perp \cap L$ is non-empty. Since $p^\perp \cap q^\perp$ is a gamma space, it is a polar space.

Next we show that the polar space $p^\perp \cap q^\perp$ is non-degenerate. By way of contradiction suppose r is a point in the radical Rad $(p^\perp \cap q^\perp)$. Now as before there is a space P_p of Π on line pr and, similarly, there is a space S_q of Σ on line rq. Then $P_p \cap S_q$ is a line L on r. Choose point c on L distinct from r, subspace S_p of Σ on line pc, and subspace P_q of Π on line cq. Then $S_p \cap P_q = N$ is a line on c.

We claim that the radical point r cannot be on line N. If it were, the line rc would belong to both P_p and P_q, forcing $P_p = P_q$ against the fact that p is not collinear with q.

That means that the subspace $\langle r, N\rangle$ is a singular subspace properly containing line N and so lies in one of the two maximal singular subspaces S_p and P_q containing N. But that is impossible. If r were in S_p, then r is a point on $pc = S_p \cap P_p$, and then $p^\perp \cap pc$ would contain distinct points r and c and so, by the gamma space hypothesis, would contain p, a contradiction. The argument against $r \in P_q$ is the same. So no such radical point r can exist, and so $p^\perp \cap q^\perp$ is a non-degenerate polar space.

We have just shown that in $(\mathcal{P}, \mathcal{L}_0)$, the collection $D := \{(p, q) | d_\Delta(p, q) = 2\}$ of all distance-two point-pairs, forms a polar family in the sense of Cooperstein's theory (presented in Chap. 13, defined on p. 455). Moreover, every line lies properly in two distinct maximal singular spaces. Thus the following applies.

Step 1. *There exists a system $\mathcal{S}$ of symplecta of polar rank at least three covering every line and every distance-two pair of points. In fact $(\mathcal{P}, \mathcal{L}_0)$ is a strong parapolar space.*

Now suppose M is a maximal singular subspace in Σ or Π and suppose p is a point not in M for which $p^\perp \cap M$ is non-empty. First we observe that $p^\perp \cap M$ cannot be a set consisting of a single point, say r. For there is always a maximal singular subspace N on the line pr in the class not containing M, and then $M \cap N$ is a line of $\mathcal{L}_0$ in $p^\perp \cap M$. Now, since the subspace $p^\perp \cap M$ is known to contain a line, one may ask whether it could be larger than a line. Whatever happens, $\langle p, p^\perp \cap M\rangle$ is a singular space properly containing a line, and so belongs to a unique maximal singular subspace N in the class opposite that of M (otherwise $p \in M$). It follows that $N \cap M$ contains $p^\perp \cap M$ which is itself a line of $\mathcal{L}_0$.

Suppose (p, L) is an anti-flag of the maximal singular subspace M belonging to Σ (Π). On the line L is exactly one other maximal singular subspace P properly containing L. Then for any $q \in P - L$ there is a symplecton $S = S(q, p)$ containing $L \cup \{p\} \cup \{q\}$. Now such a symplecton S is a polar space of rank at least three, with the unusual property that every line properly lies in exactly two maximal singular subspaces. It follows from our analysis in Sect. 7.4, p. 178, that S is the oriflame polar space of type $A_{3,2}(D_S)$ of all 2-subspaces and $(1, 3)$-subspace flags of a four-dimensional right vector space over a division ring D_S, possibly depending on S. It follows that $\langle p, L\rangle$ is a Desarguesian projective plane coordinatized by the division ring D_S. Since (p, L) was an arbitrary antiflag of M, we see that M is a classical projective space $\mathbf{P}(V)$ where V is a right vector space over the division ring D_S. Note that by symmetry of P and M, P is also a classical projective space over the same division ring D_S. Now since $(\mathcal{P}, \mathcal{L})$ is a connected geometry, and the (possibly four) maximal singular subspaces above any of two intersecting lines are coordinatized by the same division ring, all subspaces in Σ or Π are projective spaces over one common universal division ring D. It then follows that each symplecton has planes $PG(2, D)$ and so are of a fixed isomorphism type.

Before rushing on to other forensics, let us record what we have learned.

Step 2. *There exists a uniform division ring D with these properties:*

1. *Every singular subspace is a classical projective space defined over the division ring D.*
2. *Every symplecton is a convex subspace which is a classical oriflame polar space of rank three over D – that is, the polar space of type $A_{3,2}(D)$.*

Moreover, for any maximal singular subspace $M \in \Sigma \cup \Pi$, and point p not in M, $p^\perp \cap M$ is either empty or is a line $L \in \mathcal{L}_0$.

Now suppose T_1 and T_2 are two members of the same family Ξ of maximal singular subspaces which meet at a point p. Now by Step 1, part (c), for every line L_1 of T_1 on point p there is a unique line $\phi(L_1) := L_2$ of T_2, such that $L_2 := {L_1}^\perp \cap T_2$, a line of T_2 on point p. Also by Step 1, ${L_2}^\perp \cap T_1 = L_1$. This means that there is a bijection

$$\phi : \text{lines of } T_1 \text{ on point } p \to \text{lines of } T_2 \text{ on point } p.$$

We can go beyond this. Suppose π_1 is a projective plane of T_1 incident with point p. Choose any line L_1 of that plane π_1 which lies on p. Select any line N_1 of T_1 on point p so that N_1 is distinct from the line L_1. Set $N_2 = \phi(N_1)$. Then no matter what points $(x, y) \in L_1 - \{p\} \times N_2 - \{p\}$ are chosen, $S := S(x, y)$ is a symplecton S meeting T_2 at a plane $\phi(\pi_1)$. Why is this? Well, for any line N_1 of π_1 that we might choose, ${N_1}^\perp \cap T_2$ lies in $S \cap T_2$. Thus a choice of a projective plane π_1 in T_1 on p, and a member T_2 of the same class meeting T_1 at a point, determines a

projective plane $\phi(\pi_2)$ of T_2 on p such that for each line L of Π on p, $\phi(L)$ is a line of $\pi_2 := \phi(\pi_2)$ on p. Thus the mapping ϕ induces incidence-preserving bijections

$$\begin{aligned} \phi \quad &: \quad \text{lines of } T_1 \text{ on } p \to \text{lines of } T_2 \text{ on } p \\ &: \quad \text{planes of } T_1 \text{ on } p \to \text{planes of } T_2 \text{ on } p. \end{aligned}$$

It follows from connectedness of the geometry that any two maximal singular subspaces of the same class are projective spaces of the same rank. Thus the following applies.

Step 3. *Any two singular spaces of the same class are projective spaces of the same projective dimension. In particular, all maximal singular subspaces belonging to* Π *are classical projective spaces of projective dimension* d.

Let us now form a graph $(\Sigma, \sim)$ whose vertices are the members of Σ, two of which are considered to be adjacent in this graph if and only if they intersect at a point.

Now we first have to treat a special case. Suppose $d = 2$; this means that every member of Π is a projective plane. In this case, we claim that $(\Sigma, \sim)$ is a complete graph – that is, any two members of Σ intersect in a point. Suppose this were false. Then, among disjoint pairs $(T_1, T_2) \in \Sigma \times \Sigma$, we can find one pair in which the distance between points of T_1 and points of T_2 attains a minimal value $s > 0$. In this case, there exist points $t_i \in T_i$, $i = 1, 2$, connected by a geodesic path $(t_1 = x_0, x_1, x_2, \ldots, x_s = t_2)$. If $s = 1$, set $T_0 = T_1$, and if $s > 1$, let T_0 be the unique member of Σ on the line $x_{s-2}x_{s-1}$. (Note that $x_{s-1} = t_1$ if $s = 2$.) Let P be the unique member of Π on the line $x_{s-1}t_2$. Then by (T2), $P \cap T_0 = L_1$ is a line on x_{s-1} while $P \cap T_2$ is a line on $x_s = t_2$. Since P is a projective plane, the two lines L_1 and L_2 meet at a point p. But then $p \in T_0 \cap T_2$ and has distance $s - 1$ from t_1 against the assumption that s was the minimal distance of a point in T_1 from a point in T_2. Thus there are no disjoint pairs in Σ as claimed.

Continuing with the assumption that $d = 2$, we next observe that the point-line geometry $(\mathcal{P}', \mathcal{L}') := (\Sigma, \mathcal{P})$ is a linear space. If two of its "lines," p and q, meet at a "point" (that is they are both incident with one element T of Σ) then there is a unique plane P of Π on the line pq. Then the collection of all elements of Σ meeting P at a line form a "subspace" of the "point-line" geometry $(\mathcal{P}', \mathcal{L}')$ which is a projective plane (the dual of Π). So the following applies.

Step 4. *If* $d = 2$*, then the point-line geometry* $(\Sigma, \mathcal{P})$ *is a projective space* $\mathbf{P}(V)$ *of which* $(\Sigma, \mathcal{P}, \Pi)$ *are its subspaces of projective dimensions* 0, 1, *and* 2, *respectively (that is, the* 1-, 2-, *and* 3-*subspaces of* V*)*

We continue the proof assuming $d > 2$.

We are interested in maximal cliques in the graph $(\Sigma, \sim)$. Obviously Σ_p, the collection of all maximal singular subspaces in Σ which contain the point p, forms a clique. Is it maximal?

Suppose T were an element of $\Sigma - \Sigma_p$ adjacent to two elements T_1 and T_2 of Σ_p. Then the triangle $\{p, T \cap T_1, T \cap T_2\}$ spans a projective plane π lying in a unique maximal singular subspace $P \in \Pi$. Now since $d > 2$, there is a line N of P meeting π at p exactly. Let T_3 be the unique member of Σ on line N. Suppose $q \in T \cap T_3$. Then q is not in π, and $p^\perp \cap T$ would contain the line $T \cap \pi$ as well as the point p, and so $\langle p, T \cap T_1, T \cap T_2, q\rangle$ would be a Π-space meeting T at more than a line, contrary to (T2). Thus $T \cap T_3 = \emptyset$. Thus $\Sigma_p \to T \cap T_3 = \emptyset$, and so Σ_p cannot lie in a larger clique. Thus the following applies.

Step 5. *Since $d > 2$, Σ_p is a maximal clique in the graph* $(\Sigma, \sim)$. (This statement is false if $d = 2$.)

Let us describe another sort of maximal clique. Suppose T_1 and T_2 are two members of Σ meeting at a point p. Let ϕ be the incidence-preserving bijection of lines and planes of T_1 on p to the lines and planes of T_2 on p. We should first describe two subsets of Σ determined by (T_1, T_2).

Let $A(T_1, T_2)$ be the elements A of Σ for which there is a plane π on p meeting both T_1 and T_2 at lines (on p) and A meets π at a line on p. Note that this set $A(T_1, T_2)$ contains the T_1 and T_2. Note also that for any $A \in \Sigma$ as described, there is a symplecton S meeting the T_i at planes, and also meeting A at a plane.

Let $B(T_1, T_2)$ be the collection of all elements T of Σ which meets both T_1 and T_2 at points x_1 and x_2 respectively, which are both distinct from p. In that case $\langle p, x_1, x_2\rangle$ is a Π-plane (that is, a projective plane whose unique maximal singular over-space is in Π). In that case, there is again a symplecton S meeting the T_i and T at planes.

Now, if A_1 and A_2 are two distinct members of $A(T_1, T_2)$, they are adjacent vertices of the graph $(\Sigma, \sim)$ since they both contain point p.

Now suppose $B \in B(T_1, T_2)$. Then B meets T_i at point $b_i \neq p$, and $\langle p, b_1, b_2\rangle$ is a Π-plane π. Next, suppose T is an element of either $A(T_1, T_2)$ or $B(T_1, T_2)$ distinct from B. Then there is again a Π-plane π' meeting T at a line L. Then there are Σ-planes $\sigma_i := \langle \pi \cap T_i, \pi' \cap T_i\rangle$, $i = 1, 2$, and a symplecton S containing all four of the planes $\sigma_1, \pi, \sigma_2, \pi'$, centered on p. Then B and T each meet S at a line, hence meet it at Σ-planes, and these, being maximal singular subspaces of S of the same class, must in turn intersect each other at a point. Thus $B \sim T$. It follows that

$$C(T_1, T_2) := A(T_1, T_2) + B(T_1, T_2)$$

is a clique in $(\Sigma, \sim)$. In fact it is a maximal clique, for if E were an element of Σ adjacent to T_1, T_2, and B as above, but were not in $C(T_1, T_2)$, then E must lie on p but not in the plane π. But then the point $b = E \cap B$ does not lie in π, while $b^\perp$ contains both line $B \cap \pi$ and p, spanning π. Thus $\langle b, \pi\rangle$ is a Π-space meeting B in the plane $\langle B \cap \pi, b\rangle$, which is impossible. Thus the following applies.

Step 6. 1. *The set $C(T_1, T_2)$ is a maximal clique of the graph* $(\Sigma, \sim)$ *which is distinct from the maximal clique Σ_p of all elements of Σ on a point p.*

2. *Let $\mathcal{C}$ be the collection of all maximal cliques of $(\Sigma, \sim)$ of the form $C(T_1, T_2)$. Then any edge (T_1, T_2) of $(\Sigma, \sim)$ lies on just two maximal clicques, the clique $C(T_1, T_2)$ of $\mathcal{C}$, and the clique Σ_p where $\{p\} = T_1 \cap T_2$.*

We wish to produce a linear structure on Σ and for that purpose we need to study the sets $A(T_1, T_2)$ a little more closely.

Suppose T_1 and T_2 are again two members of Σ meeting at a point p. We have seen that any choice of a line L of T_1 on point p determines three unique objects: (1) a line $L^{\perp} \cap T_2 := \phi(L)$, (2) a Π-plane $\pi(L) := \langle L, \phi(L)\rangle$, and (3) the unique element $P(L) \in \Pi$ on $\pi(L)$, which meets T_1 and T_2 at L and $\phi(L)$, respectively. In fact, the choice of any one of the following four objects:

1. L, chosen as a line of T_1 on p
2. $\phi(L)$, chosen as a line of T_2 on p
3. $\pi(L)$ chosen as any Π-plane on p meeting the T_i at lines
4. $P(L)$ chosen as any member of Π on point p

uniquely determines the other three.

Now $A(T_1, T_2)$ was defined to be those elements $T \in \Sigma_p$ for which there is a "bridging" Π-plane π on p meeting T_1, T_2 and T at lines. Thus, for every line N of π on point p, we obtain a unique member T of $A(T_1, T_2)$.

Does this one plane π give us all members of $A(T_1, T_2)$ in this way? Suppose π' were another "bridging" Π-plane meeting T_1 and T_2 at lines L' and $L'_2 := \phi(L')$. Then there is a unique symplecton S containing the two planes π and π' and the four lines $\pi \cap T_i$ and $\pi' \cap T_i$, $i = 1, 2$. Now if $T' \in \Sigma$ is obtained from the line N' of π' on point p, then $T' \cap S$ is a Σ-plane of S whose intersection with plane π of S contains the point p and so is a line N. This means T' is already an element of $A(T_1, T_2)$ which arose from the "bridging" plane π.

Thus we can record the following.

Step 7. *Let π be any plane meeting T_1 and T_2 at lines (so that π contains p). Then there is a one-to-one correspondence between the members of $A(T_1, T_2)$ and the lines of π on p.*

We let

$$\mathcal{A} := \{A(T_1, T_2) | (T_1, T_2) \in \Sigma \times \Sigma, |T_1 \cap T_2| = 1\},$$

and now consider the point-line geometry $\Xi := (\Sigma, \mathcal{A})$. From our previous discussion we see the following.

Step 8. 1. *Adjacency in the graph $(\Sigma, \sim)$ coincides with $\mathcal{A}$-collinearity.*
2. *The subsets Σ_p, $p \in \mathcal{P}$, and $C \in \mathcal{C}$ are maximal singular subspaces of Ξ. Every 3-clique of $(\Sigma, \sim)$ lies in at least one of these spaces, and the two families $\mathcal{C}$ and $\mathcal{P}^* := \{\Sigma_p | p \in \mathcal{P}\}$ exhaust all maximal singular subspaces of Ξ.*

3. Ξ *is a connected partial linear gamma space.*
4. *Every line* $A \in \mathcal{A}$ *is uniquely expressible as an intersection* $A = C \cap \Sigma_p$ *for some* $C \in \mathcal{C}$ *and* $p \in \mathcal{P}$. *It lies properly in each argument of the intersection.*
5. *The point-line geometry* $\Xi = (\Sigma, \mathcal{A})$ *satisfies the axioms (T1) and (T2) with* $(\mathcal{C}, \Sigma, \mathcal{A}, \mathcal{P}^*)$ *in the role of* $(\Sigma, \mathcal{P}, \mathcal{L}_0, \Pi)$.

Now suppose $C_1, C_2 \in \mathcal{C}$ and Σ_p and Σ_q are such that the four intersections $C_i \cap \Sigma_p$ and $C_i \cap \Sigma_q$ contain elements $T_{i,p}$ and $T_{i,q}$ of Σ, respectively, $i = 1, 2$. Since we are trying to show that either $\mathcal{C}_1 \cap C_2$ or $\Sigma_p \cap \Sigma_q$ is non-empty, we may assume that the four $T_{i,x}$ $(x = p, q)$ are distinct, and that p is not $\mathcal{L}_0$-collinear with q. Then for $i = 1, 2$, $T_{i,p} \cap T_{i,q} = \{u_i\}$ for distinct points u_i in $p^\perp \cap q^\perp$. Then the symplecton $S(p, q)$ meets each of the $T_{i,x}$ at a Σ-plane $\sigma_{i,x}$. Since these are all planes of the same type in the oriflame symplecton $S(p, q)$ of rank three, $X := (T_{1,p}, T_{1,q}T_{2,p}, T_{2,q})$ is a 4-clique in the graph $(\Sigma, \sim)$. But how can that be? By Step 8, the 4-clique X either lies in a member of $\mathcal{C}$ or a member of $\mathcal{P}^*$. In the latter case, $X \subseteq \Sigma_z$ forcing $p = z = q$ against $\Sigma_p \neq \Sigma_q$. In the former case, $X \subseteq C_0 \in \mathcal{C}$, forcing $C_1 = C_0 = C_2$ against $C_1 \neq C_2$. Thus we see the following.

Step 9. *The geometry* $\Xi = (\Sigma, \mathcal{A})$ *with its two classes of maximal singular spaces* $\mathcal{C}$ *and* $\mathcal{P}^*$ *satisfies axiom (T3) – that is, if the four intersections* $C_i \cap \Sigma_x$ *(*$i = 1, 2$ *and* $x \in \{p, q\}$*) are non-empty, then at least one of* $C_1 \cap C_2$ *or* $\Sigma_p \cap \Sigma_q$ *is non-empty. (Of course if one of the intersections is non-empty, so is the other.)*

Now it is time to discuss the projective rank of Σ_p as a subspace of Ξ. Let us fix an element P of Π on point p. Then any element $T \in \Sigma_p$ intersects P at a line L on point p. Conversely, whenever L is a line of P on point p, we can look up the unique member T of Σ on line L and observe that $T \in \Sigma_p$. Thus there is a bijection

$$\phi : \Sigma_p \to \text{lines of } P \text{ on point } p.$$

Now any "line" in $\mathcal{A}$ in the subspace Σ_p has the form $A := A(T_1, T_2)$, whose members meet P at lines forming the pencil of all lines of the plane $\pi(A) := \langle T_1 \cap P, T_2 \cap P \rangle$ on point p. The notation makes sense since $\pi(A)$ is generated by any two lines of the pencil. Put another way: for any distinct elements T, T' in A, $T \cap P$ and $T' \cap P$ generate the same plane $\pi(A)$ of P on p as was obtained for T_1 and T_2. Conversely, every plane π on p determines a plane-pencil of lines on p, whose Σ-overspaces comprise a line of $\mathcal{A}$ in Σ (see the discussion just prior to Step 7). Thus we have another bijection

$$\pi : \mathcal{A}\text{-lines in } \Sigma_p \to \text{planes of } P \text{ on point } p.$$

Moreover, the mappings preserve incidence. Precisely, if T is a "point" of Σ_p incident with "line" $A \in \mathcal{A}$ contained in Σ_p, then the ordinary line $\phi(T)$ (of P on p)

is incident with the ordinary projective plane $\pi(A)$ (of P on p). Thus we see the following.

Step 10. *As a subspace of Ξ, Σ_p is a projective space of finite rank $d-1$.*

At this point $\Xi = (\Sigma, \mathcal{A})$ with its two classes of maximal singular subspaces $\mathcal{C}$ and $\mathcal{P}^* := \{\Sigma_p | p \in \mathcal{P}\}$ satisfy all of the axioms (T1)–(T4), except that now the finite rank d asserted in axiom (T4) is relaced by rank $d-1$. So the door is open to apply induction.

Thus the following applies.

Step 11. *There exists a vector space V over a division ring D such that the rank-three geometry $(\mathcal{C}, \Sigma, \mathcal{P}^*)$, with the natural incidences, is isomorphic to the incidence geometry of all $(d-2)$-, $(d-1)$-, d-subspaces of vector space V with their natural incidence.*

Note that each of the proper vector subspaces of V of dimension less that d is represented by a subgraph $\Sigma(U)$ of $(\Sigma, \sim)$ (the vertices of the subgraph $\Sigma(U)$ would be the elements of Σ representing (via Step 11) the $(d-1)$-subspaces of V which contain U). We augment this notation slightly: if $\dim U = d$, we let $\Sigma(U)$ be the point $p \in \mathcal{P}$ which it represents via Step 11. Setting

$$\mathcal{C}_k := \{\Sigma(U) | U \leq V, \dim U = k\},$$

we see that with containment as incidence,

$$(\mathcal{C}_1, \mathcal{C}_2) \simeq (V_1, V_2)$$

is the projective space $\mathbf{P}(V)$. (Recall that V_i denotes the collection of i-dimensional vector subspaces of the vector space V.)

Now choose a particular $P \in \Pi$. Then, as a subspace of $(\mathcal{P}, \mathcal{L})$ we have $P \simeq \mathrm{P}G(d, D)$. We say that an element $\Sigma(U)$ of $\mathcal{C}_k$, $k < d$, is *incident* with P if and only if there is a point $p \in \mathcal{P}$ incident with some member of $\Sigma(U)$. Thus if $k = d$, $\Sigma(U)$ is a point p of P. If $k = d-1$, $\Sigma(U) \in \Sigma$ and meets P at a line. If $k = d-2$, $\Sigma(U) \in \mathcal{C}$ and its Σ-elements all meet P at the lines of a projective subplane of P. In general, if $\dim U = k \leq d-1$, the elements of $\Sigma(U)$ which meet P at a point in fact meet P one-by-one at the collection of all lines of a $\mathrm{P}G(d-k, D)$-subspace of the projective space P, and conversely, the lines of such a subspace define a collection of subspaces in Σ lying in a unique $\Sigma(U) \in \mathcal{C}_k$.

We may form the geometry

$$\Gamma^* := (\mathcal{C}_1, \mathcal{C}_2, \ldots, \mathcal{C}_{d-2} = \mathcal{C}, \Sigma, \mathcal{P}, \Pi)$$

where incidence among all of the $\mathcal{C}_k$ is proper containment, incidence of a $\Sigma(U)$ in $\mathcal{C}_k$ with a point p is that p lies in some subspace $T \in \Sigma(U)$, incidence in $\mathcal{P} \cup \Pi$ is

that of Γ, and incidence of a $\Sigma(U)$ in $\mathcal{C}_k$ with an element $P \in \Pi$ is that described above.

Then from the paragraph preceding the last, the residue of our chosen P in the geometry Γ^* is a projective space which is the dual of the projective geometry $\mathrm{P}G(P) \simeq \mathrm{P}G(d, D)$ of all proper $\mathcal{L}$-subspaces of the projective space P.

Now suppose U_1 and U_1' are two distinct 1-subspaces of V and set $U_2 := U_1 \oplus U_1' \in V_2$.

Step 12. If $\Sigma(U_1)$ and $\Sigma(U_1')$ are both incident with P in the geometry Γ^*, then $\Sigma(U_2)$ is incident with P.

This is because (as $d \geq 3$) there are elements of $\Sigma(U_1) \cap \Sigma(U_1')$ sharing a point with P. It now follows from Step 12 that the residue of P in Γ^* truncated to $(\mathcal{C}_1, \mathcal{C}_2) \simeq \mathbf{P}(V)$ is in fact a subspace of the latter, and so is represented by a vector subspace of V. It is immediate that P is represented by a unique subspace $W(P) \in V_{d+1}$ such that the points and elements of Σ incident with P are elements of V_{d-1} and V_d which are vector subspaces of $W(P)$.

Thus we have shown that the incidence geometry

$$(\Sigma, \mathcal{P}, \Pi) \simeq (V_{d-1}, V_d, V_{d+1}),$$

the classical Grassmannian of d-subspaces of V. The proof is complete. □

14.2 A Variation for Half-Grassmann Spaces

The first theorem characterized classical Grassmann spaces.[1] Much earlier in these notes we gave a particular example of a partial linear space: the incidence system (V_d, V_{d+1}) of d- and $(d+1)$-dimensional subspaces of a vector space V over a division ring k. We call this a *classical half-Grassman space*. Note that it has the same point-set as a classical Grassmann space, but possesses an entirely different set of lines. Here, the lines are one class of maximal singular subspaces of a Grassmann space; the other class is just ignored.

Here is the theorem.

Theorem 14.2.1 *Assume $(\mathcal{P}, \mathcal{L})$ is a connected partial linear space with these properties:*

(G0) For every non-incident point-line pair $(p, L) \in (\mathcal{P}, \mathcal{L})$, intersection $p^\perp \cap L$ is a proper subset of L.[2]

[1] Actually a little more than this, since one of the two classes of maximal singular spaces is allowed to have infinite projective rank.

[2] This chapter basically offers three theorems with similar axioms. Without axiom (G0), Theorem 14.2.1 of this section (just above) would be wildly false. Consider any linear space $(\mathcal{P}, \mathcal{L})$ with thick lines and the ascending chain condition on subspaces. The reader can easily verify that

(G1) *If L is a line, and x is a point not on L, then either*

(a) $x^\perp \cap L = \emptyset$, *or*
(b) $x^\perp \cap L$ *is a subset of L of size at least three.*

(G2) (The diagonal axiom.) *If L is a line and x and y are points not incident with L, then*

$$|x^\perp \cap y^\perp \cap L| \geq 2$$

implies that x is collinear with y.

(G3) *If L is a line, and x and y are distinct collinear points of $\mathcal{P} - L$ with both $x^\perp \cap L$ and $y^\perp \cap L$ non-empty, then $x^\perp \cap y^\perp \cap L \neq \emptyset$.*

(G4) (A finiteness condition.) *For each line L, let $K(L)$ be the collection of non-empty subsets $x^\perp \cap L$ that appear as x ranges over $\mathcal{P} - L$. The assumption is that $(L, K(L))$ has the ascending chain condition on its poset of subspaces.*

Then for some vector space V over a division ring D, and some positive integer d, $(\mathcal{P}, \mathcal{L})$ is the point-line geometry (V_d, V_{d+1}) of d- and $(d+1)$-dimensional subspaces of V with incidence defined by the containment relation.

Remarks 1. In axiom (G1) it is allowed that $x^\perp \cap L = L$.
2. Axiom (G2) is essentially the dual of the Veblen axiom (Pasch axiom in common parlance) for a partial linear space.
3. Of course the finiteness condition $(G4)$ can be dropped if the number of points in a line $L \in \mathcal{L}$ is bounded by a finite number.

14.2.1 Consequences of Axioms (G1) and (G2)

For now, assume $(\mathcal{P}, \mathcal{L})$ is a partial linear space satisfying axioms (G1) and (G2) only.

Fix two distinct collinear points a and b and let $L = ab$ be the unique line incident with these points. Let $C_1(a, b) := \{p \in \mathcal{P} - L | p^\perp \cap L \supseteq \{a, b\}\}$, the points of $\mathcal{P} - L$ collinear with both a and b. By the diagonal axiom $C_1(a, b)$ is a clique.

---

all other axioms (G1), ... (G4) hold, no matter what linear space may may have been chosen by the reader. So why separate (G0) and (G1) as two axioms? Why not include both of them in a modified version of (G1)?

Here is the reason. Sprague's theorem (the third main theorem of this chapter) assumes both (G1) and (G2) but does not require (G0). It uses an induction proof grounded upon what happens at very low dimension. One cannot logically venture into such a program when burdened with an axiom like (G0) which fails for those low-dimensional cases. For Sprague's theorem (Theorem 14.3.1) we must harvest all of the consequences of (G1) and (G2) while shunning the use of (G0).

Lemma 14.2.2 *Assume $(\mathcal{P}, \mathcal{L})$ is a connected partial linear space having at least two lines and which satisfies axioms (G1) and (G2). Let a and b be distinct points of one of its lines L. Then:*

(1) For each $u, v \in C_1(a, b)$, one has $u^\perp \cap L = v^\perp \cap L$. [Since this set depends only on $C_1(a, b)$ and not the particular u or v in $C_1(a, b)$, we denote it by $C_2(a, b)$.]
(2) The set $C_1(a, b)$ is non-empty and $C(a, b) := C_1(a, b) + C_2(a, b)$ is a clique.
(3) $C(a, b) \cap L = \{u^\perp \cap L\}$ for each $u \in C(a, b) - L = C_1(a, b)$.
(4) Either $C(a, b)$ contains L or else $C(a, b)$ and $L = ab$ are the only two maximal cliques of the point-collinearity graph $\Delta = (\mathcal{P}, \sim)$ which contain $\{a, b\}$. Thus if condition (G0) holds, then any line L is a maximal clique.
(5) For any line N for which $|N \cap C(a, b)| > 1$, we have $C(a, b) = C(u, v)$ for any two distinct points u and v in $N \cap C(a, b)$.
(6) For any line N, the intersection $N \cap C(a, b)$ is either empty or has at least two points.
(7) Suppose p is a point not in C. Then $p^\perp \cap C(a, b)$ is either empty or is a non-empty intersection $L \cap C$, where L is a line on p. In particular, if two distinct lines intersect $C = C(a, b)$ non-trivially, and also intersect each other at a point x, then x lies within C.

Proof For the proof of (1), suppose $v \in C_1(a, b)$, and $z \in u^\perp \cap L$. Then, noting that u is collinear with v, we may apply (G2) with (z, v, a, u, ua) in the role of (x, y, a, b, L) to conclude that z is collinear with v. Thus $u^\perp \cap L \subseteq v^\perp \cap L$, and by the symmetry of the hypotheses on u and v, the two sets are equal.

Item (2) is the only part which uses that fact that there are at least two lines.

Suppose $a^\perp \cap b^\perp \subseteq L$. If $a^\perp$ is not contained in L, there is a line N meeting L at point a. Then, as $|b^\perp \cap N| \geq 3$, we have points in $N \cap a^\perp \cap b^\perp - L$, contrary to what we supposed. By a similar argument, $a^\perp \subseteq L$ implies $x^\perp \subseteq L$ for all points x of L. Since $(\mathcal{P}, \mathcal{L})$ is a connected partial linear space, we see that $\mathcal{L} = \{L\}$. But that contradicts $|\mathcal{L}| \geq 2$.

To prove (3) suppose $\{u, v\}$ is a 2-subset of $C_1(a, b)$. Suppose $c \in u^\perp \cap L - \{a, b\}$. Then c and v are both collinear with the two points u and a of the line ua. By (G2) the point c is collinear with v. Since c was any point of $u^\perp \cap L - \{a, b\}$, it follows that $u^\perp \cap L \subseteq v^\perp \cap L$, and equality of the two sets follows from the symmetry of u and v in this context. Now (3) follows from $u^\perp \cap L = v^\perp \cap L$ for all $u, v \in C_1(a, b)$.

Now we prove (4). Suppose for some point $x \in \{a, b\}^\perp$, one has $x^\perp \supseteq L$. Then by step (3), $C(a, b) \supseteq L = u^\perp \cap L$, for all $u \in C(a, b)$, and here $C(a, b)$ coincides with the set $\{p \in \mathcal{P} || p^\perp| \geq 2\}$.

Otherwise, for some point $p \in \{a, b\}^\perp - L$, $p^\perp \cap L$ is a proper subset of L, and $L - p^\perp = L - C(a, b) \neq \emptyset$ and L is a maximal clique.

Suppose C were any maximal clique of Δ distinct from L containing$\{a, b\}$. Then C contains a point $x \in C - L$. Then any such x lies in $C_1(a, b)$. It follows from part 1 that $C \cap L \subseteq C_2(a, b)$, hence $C \subseteq C(a, b)$ and equality holds by the maximality. This distinguishes $C(a, b)$ as well as showing that it itself is a maximal clique. Thus (4) holds.

To prove (5), suppose $\{u, v\}$ is a pair of distinct points on an intersection, $N \cap C(a, b)$, where N is a line. Then $C(u, v)$ and $C(a, b)$ are maximal cliques of the point-collinearity graph that contain $\{u, v\}$ and are distinct from the line on $N = uv$. It follows from part 4 that they are equal.

For part (6), suppose $p \in N \cap C(a, b)$ for some line N. We wish to show that $N \cap C(a, b)$ contains an additional point besides p. So we may assume that $N \neq L = ab$. Thus by partial linearity, N cannot contain both a and b. Transposing a and b, if necessary, we may assume that a is not in N so $p \neq a$. Now by (G1), and the fact that $a \notin N$, we see that $a^\perp \cap N$ contains at least three points. So there is a point $u \in (N - \{p\})$ that is collinear with point a. Now v is collinear with the two distinct points a and p of $C(a, b)$, and does not lie on the line ap. Thus $v \in C(a, p)$. But by part (5), $C(a, p) = C(a, b)$.

Finally, consider part (7). We write C for $C(a, b)$. Suppose $x \in p^\perp \cap C$. If there were a point $y \in p^\perp \cap C - px$, then p cannot be on line xy by partial linearity of $(\mathcal{P}, \mathcal{L})$. But this forces $p \in C(x, y)$, while on the other hand $C(x, y) = C(a, b)$ by (5). Thus there is no such point y, and so $p^\perp \cap C = px \cap C$ as desired.

Now suppose x is a point lying on two distinct lines L_i and L_2, each intersecting C at an $\mathcal{L}_0$-line. If x were not in C, partial linearity would force these intersections to be disjoint. But from the previous paragraph, $x^\perp \cap C$ is a single $\mathcal{L}_0$-line. This contradiction infirms the supposition that x was not in C. □

Now let $\mathcal{C}$ be the collection of all maximal cliques of the form $C(a, b)$ where (a, b) is an edge of Δ for which $a^\perp \cap b^\perp \not\subseteq ab$.

Now let $\mathcal{L}_0$ be the collection of all non-empty intersections $C \cap L$ where $(C, L) \in \mathcal{C} \times \mathcal{L}$. We can then form a new point-line geometry $\Gamma^* := (\mathcal{P}, \mathcal{L}_0)$.

Lemma 14.2.3 *If for some point x not in line L, one has $L \subseteq x^\perp$ (that is, condition (G0) fails), then $\mathcal{L} = \mathcal{L}_0$ and $(\mathcal{P}, \mathcal{L})$ is a linear space.*

Proof Assume $L \subseteq x^\perp$ for some point $x \in \mathcal{P} - L$. Then by Lemma 14.2.2, part (1), L lies in a unique maximal clique C. Let N be any line intersecting L at a point, say, $N \cap L = \{p\}$. We claim that $N \subseteq C$. Suppose otherwise, so that there is a point $y \in N - C$. Now $y^\perp \cap L$ contains p and at least one other point q. But then, by Lemma 14.2.2, part (5), we may write $C = C(p, q)$ forcing $y \in p^\perp \cap q^\perp \subseteq C$, a contradiction.

Next let M be any line with a non-empty intersection with C. If $M \cap L$ is non-empty, one has $M \subseteq C$ by the result of the previous pargraph. So suppose $M \cap L = \emptyset$ and choose a point u in $C \cap M$. Then, noting that $(C, \mathcal{L}_0(C))$ is a linear space, there is a line N on u meeting L at a point. Again, the argument of the previous paragraph gives us $N \subseteq C$. But the same argument then produces $M \subseteq C$ since $M \cap N$ is non-empty. Thus, in any case, $M \subseteq C$.

Thus no point outside C can be collinear with a point of C. Since $(\mathcal{P}, \mathcal{L})$ is connected, one obtains $C = \mathcal{P}$, and so $(\mathcal{P}, \mathcal{L})$ is a linear space. Since all lines are in C, every line of $\mathcal{L}$ is already a line of $\mathcal{L}_0$. □

Now assume $(\mathcal{P}, \mathcal{L})$ has at least two lines. Then, since $(\mathcal{P}, \mathcal{L})$ is connected, we see that for any line L, there is a point y in $\mathcal{P} - L$ with $y^\perp \cap L \neq \emptyset$ – that is, each line in $\mathcal{L}$ meets some clique of $\mathcal{C}$ non-trivially. We have the following result.

Lemma 14.2.4 *1. $\mathcal{L}$-collinearity coincides with $\mathcal{L}_0$-collinearity – that is, Γ and Γ^* have the same point-collinearity graph.*
2. Also from Lemma 14.2.2, $\mathcal{L}_0$ can be described as the collection of non-empty intersections $u^\perp \cap L$ where $L \in \mathcal{L}$ and $u \in \mathcal{P} - L$.
3. We also see that $\Gamma^* = (\mathcal{P}, \mathcal{L}_0)$ is a partial linear gamma space with thick lines.

Proof Only part 3 requires proof. The thickness of the lines of $\mathcal{L}_0$ is immediate from part 2 and (G1). Also, part 2 shows that the partial linearity of Γ forces that of Γ^*. All that remains is the gamma space property. But if a point p not on $L_0 = C \cap L \in \mathcal{L}_0$, $(C, L) \in \mathcal{C} \times \mathcal{L}$, were $\mathcal{L}_0$-collinear with two distinct points a and b of L_0, then p would be $\mathcal{L}$-collinear with them by part 1, and so $p^\perp \cap L = C \cap L$ – that is, p is $\mathcal{L}_0$-collinear with all points of $C \cap L = L_0$. □

We now see the following.

Lemma 14.2.5 *If $\Gamma = (\mathcal{P}, \mathcal{L})$ is connected, has at least two lines, and satisfies (G0), (G1), and (G2), then the partial linear gamma space $\Gamma^* = (\mathcal{P}, \mathcal{L}_0)$ satisfies these axioms:*

(T1) *Any set of three pairwise collinear points lies in a singular subspace of Γ^*.*
(T2) *The maximal singular subspaces of Γ^* fall into two classes $\mathcal{C}$ and $\mathcal{L}$, such that:*
  (a) Any member of $\mathcal{C}$ meets any member of $\mathcal{L}$ at a line of $\mathcal{L}_0$ or at the empty set.
  (b) Every line (of $\mathcal{L}_0$) lies in exactly one member of $\mathcal{C}$ and exactly one member of $\mathcal{L}$.

Proof Axiom (T1) is immediate from the construction of Γ^*. Since maximal cliques of the point-collinearity graph Δ are maximal singular subspaces of Γ^* (a property of gamma spaces), $\mathcal{C}$ and $\mathcal{L}$ are indeed maximal singular subspaces of $(\mathcal{P}, \mathcal{L}_0)$. Finally, properties (a) and (b) follow from the definition of $\mathcal{L}_0$ and from part 4 of Lemma 14.2.2. □

Remark Much of the material in this section specifically exploring axioms (G1) and (G2) in isolation, has been drawn from articles by F. DeClerck and J. Thas [54, 55], where the conclusions are developed for partial geometries, and from I. Debroey [50], where they are examined for semipartial geometries. As one can see from this section, these two axioms say a lot just for partial linear spaces in general. A good reference is the article [56] entitled "Some classes of rank 2 geometries" by F. DeClerck and H. Van Maldegham, appearing in the *Handbook of Incidence Geometry*, Francis Buekenhout, editor, North Holland, Amsterdam, 1994, pp. 433–476.

14.2.2 The Effect of the Other Axioms

So far, we have assumed that $(\mathcal{P}, \mathcal{L})$ was a partial linear space satisfying axioms (G1) and (G2). Now assume, in addition, axioms (G0) and (G3).

Theorem 14.2.6 *Assume* (G1) – (G3) *holds for the connected partial linear space* $\Gamma = (\mathcal{P}, \mathcal{L})$. *Let* $\Gamma^* = (\mathcal{P}, \mathcal{L}_0)$ *and let the two classes of maximal singular subspaces* $\mathcal{C}$ *and* $\mathcal{L}$ *of* Γ^* *be as in the previous section. Then the following holds:*

(T4) *If two distinct elements of* $\mathcal{L}$ *each intersect non-trivially two distinct members of* $\mathcal{C}$, *then the two members of* $\mathcal{L}$ *meet at a point.*

Proof Suppose L_1 and L_2 are distinct elements of $\mathcal{L}$ that have a non-empty intersection with two distinct singular subspaces C_1 and C_2 of $\mathcal{C}$.

By way of contradiction assume $L_1 \cap L_2 = \emptyset$. We claim that this assumption implies that $C_1 \cap C_2 = \emptyset$ as well. For if this claim were false there would be a point z in $C_1 \cap C_2$. Then $z^\perp$ contains the lines $C_1 \cap L_1$ and $C_2 \cap L_1$ of $\mathcal{L}_0$. If z were not in L_1, these two lines would coincide with each other, being as they are, the $\mathcal{L}_0$-line $z^\perp \cap L_1$ (Lemma 14.2.4, part 2). But that would force $C_1 \cap C_2$ to contain this line, against conclusion (T2) (b) of Lemma 14.2.5 and the hypothesis $C_1 \neq C_2$. Thus we must conclude that z is in L_1. By the same argument on $\mathcal{L}_0$-lines $C_1 \cap L_2$ and $C_2 \cap L_2$, we obtain $z \in L_2$ as well as L_1 and this contradicts $L_1 \cap L_2 = \emptyset$.

So we may assume $C_1 \cap C_2 = \emptyset = L_1 \cap L_2$. But now if c and d are points in $C_1 \cap L_2$ and $C_2 \cap L_2$, respectively, then $c^\perp \cap L_1$ and $d^\perp \cap L_1$ are each non-empty. Moreover, by (G3), $c^\perp \cap d^\perp \cap L \neq \emptyset$. But this forces $C_1 \cap C_2 \neq \emptyset$, a contradiction. Thus $L_1 \cap L_2 \neq \emptyset$. □

Proof of Theorem 14.2.1 Now we see that the point-line geometry $\Gamma^* = (\mathcal{P}, \mathcal{L}_0)$ satisfies all of the axioms (T1)–(T4) of the previous section. It follows from the main theorem of that section that there exists a vector space V over a division ring D, such that the incidence system $(\mathcal{C}, \mathcal{P}, \mathcal{L})$ is isomorphic to the truncation of incidence among proper subspaces of V to the $(d-1)$-, d-, and $(d+1)$-subspaces of the vector space V. This implies the conclusion of Theorem 14.2.1. □

14.3 Sprague's Theory

14.3.1 Introduction

When recounting the effect of a series of contributions some 20 years later, one is forced to modify the original presentation for several reasons: (1) there are places in which ad hoc methods of the contributor can be enlarged to a "stand-alone" methodology adaptable to multiple uses, and (2) the work itself is somewhat generalized. We have done this twice before with the work of Teirlinck and the work of Cooperstein. And now we repeat this process with the work of Alan Sprague.[3]

[3] Sprague's case is all the more remarkable: here was some of the most talented work in synthetic geometry in the United States, by someone who did not have the professional opportunities to pursue this line of research for more than a few years. This certainly qualifies as a lost opportunity.

So far, we have encountered two theorems characterizing a Grassmann space, and then a half-Grassmann space by the axioms (G1) and (G2) and some version of the intersection axiom (T3) or (G3). In this section we will consider (G1) and (G2) to be in effect, while (G3) is replaced by an entirely different axiom, the Veblen axiom (H3). Such a theory was first discovered by Alan Sprague, and, were it easily reducible to either of the first two theorems employing (T3) or (G3), it would be a mere footnote to the first two sections. However, despite the common presence of axioms (G1) and (G2), neither of the axioms (G3) and (H3) seem to be deducible from the other. So the Sprague theorem appears to be an entirely new class of theorem. One should note that Sprague's original paper [121] appeared in 1981, only a few months after the Shult–Bichara/Tallini theorems described in Sect. 14.2.

In order to describe the statement of the theorem and related concepts, we need to introduce a minor terminology. A *pencil of lines* is simply some collection of two or more lines which pairwise intersect at a common point – say, p. (If $(\mathcal{P}, \mathcal{L})$ is a partial linear space, as will be assumed in this section, a pencil of lines would just be some subset of $\mathcal{L}_p$ of size at least two, for some point p.) A line N is said to be *transverse to the pencil* $X \subseteq \mathcal{L}_p$ if and only if N is not on p but intersects every line of the pencil X. If $X = \{L_1, L_2\}$ then N is said to be *transverse to* L_1 *and* L_2. Also, we may speak of the set of lines $K_p(N)$ of $\mathcal{L}_p$ which *have been transversed by* N, where N does not contain p but meets at least two lines of $\mathcal{L}_p$.

Here is our version[4] of Sprague's Theorem.

Theorem 14.3.1 (Sprague) *Assume $(\mathcal{P}, \mathcal{L})$ is a connected partial linear space with at least two lines satisfying:*

(G1) If L is a line, and x is a point not on L, then either

(a) $x^\perp \cap L = \emptyset$, or
(b) $x^\perp \cap L$ is a subset of L of size at least three.

(G2) (The diagonal axiom.) *If L is a line and x and y are points not incident with L, then*

$$|x^\perp \cap y^\perp \cap L| \geq 2$$

implies that x is collinear with y.

(H3) (The Pasch-Veblen axiom.) *If L_1 and L_2 are two distinct lines meeting at a point p, then any two transversals of L_1 and L_2 meet at a point.*

(H4) (A finiteness condition.) *Fix a point p, and let K_p be all those non-empty subsets of $\mathcal{L}_p$ which show up as the full set of lines transversed by some line not on p. (By (G1) these sets are empty or entail at least three lines of $\mathcal{L}_p$.)*

[4] Sprague assumed some line had finitely many points. On the other hand, he was able to handle thin cases, where the word "three" in (G1) (b) is replaced by "two" and the geometry of i- and $(i+1)$-subsets of a set X enters the conclusion (see Theorem 2 of [121]).

If L is a fixed line, let $K(L)$ be the full collection of all non-empty subsets $x^{\perp} \cap L$ for $x \in \mathcal{P} - L$. (Again by (G1) these sets are size at least three.) Assume that either:

(a) There exists exists a point p such that the point-line geometry $(\mathcal{L}_p, K_p)$ possesses an unrefinable properly ascending chain of subspaces from $\emptyset$ to $\mathcal{L}_p$ of finite length, or

(b) There exists a line L such that the point-line geometry $(L, K(L))$ possesses an unrefinable properly ascending chain of subspaces from $\emptyset$ to L of finite length.

Then one of the following statements holds:

(1) $(\mathcal{P}, \mathcal{L})$ is a non-classical projective plane.

(2) There exists a vector space V and a positive finite integer d such that either $(\mathcal{P}, \mathcal{L}) \simeq (V_{d+1}, V_d)$ or $(\mathcal{P}, \mathcal{L}) \simeq (V_d, V_{d+1})$. The alternative depends upon which of the respective hypotheses (a) or (b) in (H4) is being assumed.

14.3.2 (G1)-Geometries with the Veblen and Dual-Veblen Axioms, (G2) and (H3)

In this section we assume that $(\mathcal{P}, \mathcal{L})$ is a point-line geometry satisfying the three axioms (G1), (G2), and (H3). First one observes that each of the following properties are self dual (that is, if the property is true of $(\mathcal{P}, \mathcal{L})$, then it is also true of $(\mathcal{L}, \mathcal{P})$):

1. $(\mathcal{P}, \mathcal{L})$ is connected.
2. $(\mathcal{P}, \mathcal{L})$ is a partial linear space.
3. Given $(p, L) \in \mathcal{P} \times \mathcal{L}$ at distance three in the bipartite incidence graph $(\mathcal{P} \cup \mathcal{L}, *)$, then there are at least three intermediate points and three intermediate lines to be found on the geodesic paths in $(\mathcal{P} \cup \mathcal{L}, *)$ connecting p to L. (This is equivalent to (G1).)

Similarly, the Veblen and dual-Veblen axioms (G2) and (H3)[5] are duals of each other. Thus we have a principle of duality operating throughout this section:

(Duality principle.) *Any property of points and lines deducible from axioms (G1), (G2), and (H3) is true of the dual geometry of lines and points.*

Now a perusal of Lemmas 14.2.4 and 14.2.5 of Sect. 14.3 show that (G1) and (G2) alone force the following.

Lemma 14.3.2 (Sect. 14.2 briefly revisited.)

1. For any two distinct points on a line L there exists a clique $C = C(a, b)$ which is a proper maximal clique of the point-collinearity graph $(\mathcal{P}, \sim)$.

[5] Sometimes called the "Pasch" and "dual-Pasch" axioms.

2. *Let* $\mathcal{C} = \{C(a,b) | (a,b)$ *an edge of* $(\mathcal{P}, \sim)\}$. *Let*

$$\mathcal{L}_0 := \{C \cap L | (C, L) \in \mathcal{C} \times \mathcal{L}, |C \cap L| > 1\}.$$

Then $\mathcal{L}_0$*-collinearity is* $\mathcal{L}$*-collineaity on points.* $(\mathcal{P}, \mathcal{L}_0)$ *is a partial linear gamma space.*
3. *Every* $\mathcal{L}_0$*-line is uniquely expressible as the intersection (as sets of points) of two maximal singular* $\mathcal{L}_0$*-subspaces,* $C \cap L$ *for* $(C, L) \in \mathcal{C} \times \mathcal{L}$. *Any* $C \in \mathcal{C}$ *is a maximal singular* $\mathcal{L}_0$*-subspace. If axiom (G0) holds, then* L *is also a maximal singular subspace of* $(\mathcal{P}, \mathcal{L}_0)$.
4. *Given* $L \in \mathcal{L}$ *or* $C \in \mathcal{C}$, *let* $\mathcal{L}_0(L)$ *and* $\mathcal{L}_0(C)$ *be the* $\mathcal{L}_0$*-lines contained in* L *or* C, *respectively. Then both* $(L, \mathcal{L}_0(L))$ *and* $(C, \mathcal{L}_0(C))$ *are linear spaces.*

Now assume (H3). Just by applying duality to the results of the previous lemma we see that there is certain type of clique E in the line-graph $(\mathcal{L}, \sim)$, where two lines are adjacent if and only if they intersect at a point. Suppose L_1 and L_2 are two lines meeting at a point p. Let $c_1(L_1, L_2)$ be the collection of all lines transverse to $\{L_1, L_2\}$. In turn let $c_2(L_1, L_2)$ be all lines on p which meet some transverse line T of $c_1(L_1, L_2)$. Then $e(L_1, L_2) := c_1(L_1, L_2) + c_2(L_1, L_2)$ is a maximal clique of the line graph $(\mathcal{L}, \sim)$ – that is, a bunch of lines which pairwise intersect non-trivially and which lie in no larger bunch of pairwise intersecting lines. This is an easy consequence of the axiom (H3). Now we let $\mathcal{E}$ be the collection of all line-sets of the form $e(L_1, L_2)$ where (L_1, L_2) is an edge in the line-graph $(\mathcal{L}, \sim)$. Now if $\mathcal{E} = \{\mathcal{L}\}$ then by (G2) the incidence system $(\mathcal{E}, \mathcal{L})$ is a projective space, $\mathbf{P}(V)$ and $(\mathcal{P}, \mathcal{L})$ is determined. So we assume $|\mathcal{E}| > 1$.

At this stage we have a rank four incidence system $\Gamma^* = (\mathcal{C}, \mathcal{P}, \mathcal{L}, \mathcal{E})$. Note that any point $p \in L \in E \in \mathcal{E}$ lies on at least two lines of E.

So, blessed by duality, we may record the following.

Lemma 14.3.3 *1. Any two lines* L_1 *and* L_2 *which meet at a point* p *lie together in a unique maximal clique* $E = c(L_1, L_2) \in \mathcal{E}$ *of the line graph* $(\mathcal{L}, \sim)$ *as well as another clique* $\mathcal{L}_p$ *of all lines on point* p. *The intersection of these two cliques is a set of lines* $K_p(N)$ *on* p *transversed by some line* N.
2. *For fixed point* p *let* $K_p = \{K_p(N) | N \in \mathcal{L}, p^\perp \cap N \neq \emptyset\}$ *be the collection of all non-empty subsets of* $\mathcal{L}_p$ *that are the full set of lines of* K_p *transversed by some line not on* p *(see Axiom (H4) on p. 512). Then* $(\mathcal{L}_p, K_p)$ *is a linear space.* [This is the dual of the assertion that $(L, \mathcal{L}_0(L))$ is a linear space.]
3. *For any fixed line-clique* $E \in \mathcal{E}$, *let* $\bar{E}$ *be the union of the point-shadows of the lines of* E. *For any point* $p \in \bar{E}$, *let* E_p *be the lines of* E *which lie on point* p, *and let* $K(E) = \{E_p | p \in \bar{E}\}$. *Then* $(E, K(E))$ *is a linear space.* [This is the dual of the assertion that $(C, \mathcal{L}_0(C))$ is a linear space.]

Now we can do more than dualize: we can integrate the two hypotheses (G2) and (H3). Behold!

Lemma 14.3.4 *For each $C \in \mathcal{C}$, the $\mathcal{L}_0$-subspace $(C, \mathcal{L}_0(C))$ is a projective space.*

Proof Consider two $\mathcal{L}_0$-lines, say $L_1 \cap C$ and $L_2 \cap C$, meeting at a point of p of C. Two $\mathcal{L}_0$-transversals of these lines would be $\mathcal{L}$-lines N and M, not on p meeting L_1 and L_2 at points of C. Then by (H3), N and M meet at a point z. We shall show that z is in fact in C.

The points of intersection produced by the transverse lines can be assumed to be distinct. Then the clique from $\mathcal{C}$ that lies on the triangle $\{z, L_1 \cap N, L_2 \cap N\}$ must be the unique one on $\{L_1 \cap N, L_2 \cap N\}$, namely C itself. Thus $z \in C$.

Since we have established the Veblen axiom for $(C, \mathcal{L}_0(C))$, the latter is a projective space. □

Now there is more. Suppose C_1 and C_2 are two members of $\mathcal{C}$ meeting at a point p. For each line $L \in \mathcal{L}_p$, L must meet both C_1 and C_2 at $\mathcal{L}_0$-lines, $L \cap C_1$ and $L \cap C_2$. In fact $L \cap C_2 = (L \cap C_1)^{\perp} \cap C_2$, and vice versa, so there is a one-to-one correspondence

$$\phi : \text{lines of } C_1 \text{ on point } p \to \text{lines of } C_2 \text{ on point } p.$$

Suppose now L_1 and L_2 are two distinct lines of $\mathcal{L}$ on p. Then $L_{ij} := L_i \cap C_j$ are four distinct $\mathcal{L}_0$-lines, two of them generating a plane π_1 in C_1, the other two generating a plane π_2 in C_2. Now let N be an $\mathcal{L}_0$-line of C_1 on p in plane π_1. Then there is a unique member $L(N) \in \mathcal{L}$ containing N, which meets C_2 at a line $\phi(N)$ on p in C_2. Our objective is to show that line $\phi(N) = L(N) \cap C_2$ lies in plane $\pi_2 = \langle L_1 \cap C_2, L_2 \cap C_2 \rangle$.

Select any transverse T to $L_i \cap C_1$ in π_1 and any $\mathcal{L}_0$-transverse T_2 to $L_i \cap C_2$ in $\mathcal{C}_2$. Let $L(T)$ and $L(T_2)$ be the unique members of $\mathcal{L}$ containing the $\mathcal{L}_0$-transverse lines T and T_2. By (H3) these two lines meet at a point z. Also, if $L(N)$ is the unique member of $\mathcal{L}$ containing N then $L(N) \cap L(T) = N \cap T = \{r\}$, a point of plane π_1. Now by construction $L(T) \cap L_1$ is a point v of plane π_1. Thus $L(T_2)$ and $L(N)$ are both transversals of L_1 and $L(T)$, and so by (H3) meet at a point t. Now $L(N) = tp$ and $t^{\perp} \cap C_2$ contains p and T_2 which is not on p. It follows that $t \in C_1 \cap L(T_2) \subseteq \pi_2$. Moreover, $tp = T(N) \cap C_2 = \phi(N)$, a line of π_2, as promised.

Thus ϕ takes all lines of π_1 that are on p, to all lines of π_2 that are on point p. In this way mapping ϕ induces a mapping

$$\phi : \text{planes of } C_1 \text{ on } p \to \text{planes of } C_2 \text{ on } p.$$

Thus the following lemma:

Lemma 14.3.5 *1. All subspaces $(C, \mathcal{L}_0(C))$, $C \in \mathcal{C}$, are projective subspaces of the same rank.*

2. For $E \in \mathcal{E}$, set $K(E) := \{E_p = E \cap \mathcal{L}_p | p \in E\}$. Then by duality all rank two geometries $(E, K(E))$ are projective spaces of the same rank.

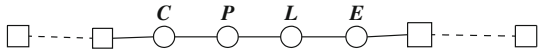

Fig. 14.1 Sprague's locally truncated geometry

At this stage, we have a diagram geometry with this diagram of Fig. 14.1. The important thing is the projective plane residue in the middle: Fix an element E of $\mathcal{E}$ and an element $C \in \mathcal{C}$ that is incident with E. Then the set of points $\bar{E} \cap C$, together with the $\mathcal{L}_0$-lines incident with both C and E, is a projective plane of C defined by E. Thus the residue of E (looking to the left in the diagram of Fig. 14.1) is a linear space whose planes are all projective subspaces. Of course that means the linear space *is* a projective space – which is a way of saying that the residue $(\mathcal{L}_E, \mathcal{P}_E)$ of lines and points of E is also a projective space. This is interesting, for it implies that its "point-residual," the space $(L, K(L))$, is also a projective space – a fact we could have concluded independently by "the magic of duality."

Let us first rid ourselves of a pesky case.

Lemma 14.3.6 *Suppose $E \in \mathcal{E}$, and $E = \{L\}$. Then $\mathcal{L} = \{L\} = E$.*

Proof If the maximal line clique E consists of only one line, L, then no other line can intersect L at a point. For if $N \in \mathcal{L}$ and $N \cap L = p$, the fact that $|x^{\perp} \cap L| \geq 3$ for any $x \in N - \{p\}$, shows that L and N possess a non-empty set of transversals. That would imply that L was a member of a larger line clique, against the maximality of E. Since L can intersect no further line, the connectedness of $(\mathcal{L}, \sim)$ demonstrates that L is the only line there is. □

Lemma 14.3.7 *Assume $|\mathcal{L}| > 1$. For $E \in \mathcal{E}$ and $C \in \mathcal{C}$ let $E(C)$ denote the set of lines in E which intersect the clique C non-trivially, Then either:*

1. *$E(C) = \emptyset$, or*
2. *There exists a projective plane π_E in the projective space $(C, \mathcal{L}_0(C))$ such that the collection $\{L \cap C | L \in E(C)\}$ is precisely the set of $\mathcal{L}_0$-lines of the plane π_E.*

Proof One only need consider the case that $E(C)$ contains a line L. Since $|\mathcal{L}| > 1$, the preceding lemma shows us that $|E| \geq 1$ and so L must intersect a second line N in E, say at a point z. Then for any $x \in L - \{z\}$, the lines $\{xy | y \in x^{\perp} \cap N\}$ comprise at least three lines of E on point x. In general, any point on a line of E lies on at least three lines of E. In particular, any point of C which lies on a line of E, lies on at least three lines of $E(C)$. Even more particularly, $|E(C)| > 2$.

If L_1 and L_2 are distinct lines of $E(C)$, by Lemma 14.2.2, part (7), their point of intersection must lie in C – that is, their corresponding $\mathcal{L}_0$-lines, $L_i \cap C$ intersect in C. If x and y are distinct points of C, each lying on a line of E, then either xy is one of those lines or else xy is a line transverse to those lines. In ether case xy belongs to $E(C)$. It follows that

$$\{\bar{E} \cap C, \{L \cap C | L \in E(C)\})$$

is a projective plane. □

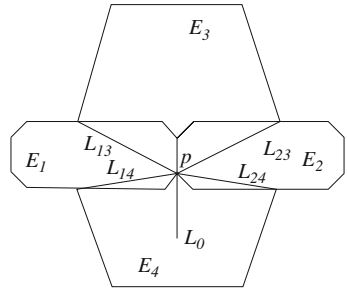

Fig. 14.2 The Veblen configuration for $(\mathcal{L}_p, \mathcal{E}_p)$

Lemma 14.3.8 *Given point p, let $\mathcal{L}_p$ and $\mathcal{E}_p$ be the lines and elements of $\mathcal{E}$ which are incident with p. Then $(\mathcal{L}_p, \mathcal{E}_p)$ is a projective space.*

Proof It suffices to show that $(\mathcal{L}_p, \mathcal{E}_p)$ satisfies the Veblen axiom. We thus consider two distinct elements E_1, E_2, sharing a common line L_0, and two "$\mathcal{E}$-transversals", E_3 and E_4 of E_1 and E_2 so that $E_i \cap E_{j+2}$, $i, j \in \{1, 2\}$, are four lines L_{ij} as in Fig. 14.2. Now all five lines in sight are on point p.

Let C be a singular subspace of $\mathcal{C}$ on point p. Then C intersects each of five lines L_0, L_{ij} at $\mathcal{L}_0$-lines L_{00} and L_{0ij}, respectively, $1 \leq i \leq 2 < 3 \leq j \leq 4$. Then we can form the four projective planes $\pi_j := C \cap E_j$, consisting of the points and lines of E_j incident with C, $j = 0, 1, 2, 3, 4$. (Since the points are those of C these are actually projective planes of the space $(\mathcal{P}, \mathcal{L}_0)$.) Now $\pi_1 \cap \pi_2 = L_{00} \subseteq L_0$, and for $(i, j) \in \{1, 2\} \times \{3, 4\}$, $\pi_i \cap \pi_j = L_{0ij}$, so the planes π_3 and π_4 lie in $\langle \pi_1, \pi_2 \rangle \simeq \mathrm{PG}(3, D)$ and so meet at an $\mathcal{L}_0$-line $L_{0\infty}$ which clearly belongs to both E_3 and E_4. Then $L_{0\infty}$ belongs to a unique line L_∞ of $\mathcal{L}$ belonging to both E_3 and E_4. So the Veblen axiom holds. □

Corollary 14.3.9 *Given line L, let*

$$\mathcal{C}_L := \{C \cap L | C \in \mathcal{C}, C \cap L \neq \emptyset\}.$$

Then $(L, \mathcal{C}_L)$ is a projective space.

The corollary is just the dual of the conclusion of Lemma 14.3.8.

Lemma 14.3.10 *The partial linear space $(\mathcal{P}, \mathcal{C})$ satisfies both the Veblen and diagonal axioms (that is, the Pasch and dual-Pasch axioms, (H3) and (G2)).*

Proof *$(\mathcal{P}, \mathcal{C})$ satisfies the diagonal axiom:*

Suppose $C \in \mathcal{C}$ and x and y are points of $\mathcal{P} - C$, which are $\mathcal{L}$-collinear with two distinct points a and b of C. Then the two triangles $\{x, a, b\}$ and $\{y, a, b\}$ lie in $\mathcal{L}$-lines L and N. But since $L \cap N$ contains the 2-set $\{a, b\}$, we must have $L = N$ since $(\mathcal{P}, \mathcal{L})$ is a partial linear space. However, in that case x and y are collinear.

$(\mathcal{P}, \mathcal{C})$ satisfies the Veblen axiom.

We assume that C_1 and C_2 are distinct elements of $\mathcal{C}$ meeting at a point p, and that C_3 and C_4 are $\mathcal{C}$-transversals of C_1 and C_2, so that $C_i \cap C_{j+2} = \{x_{ij}\}$, $i, j \in \{1, 2\}$

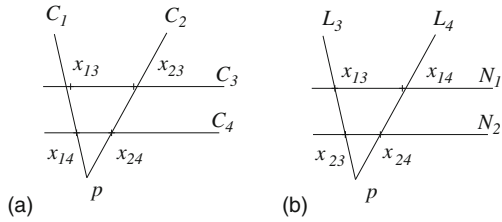

Fig. 14.3 (**a**) The Veblen configuration for $(\mathcal{P}, \mathcal{C})$. (**b**) The derived configuration in $(\mathcal{P}, \mathcal{L})$ when $L_3 \neq L_4$

and p is not in C_{j+2}. We need only show that $C_3 \cap C_4$ is non-empty. Assume, by way of contradiction, that $C_3 \cap C_4 = \emptyset$, so that p and the x_{ij} are five distinct points (see Fig. 14.3a).

Then the two triangles $\{p, x_{13}, x_{23}\}$ and $\{p, x_{14}, x_{24}\}$ lie in respective unique $\mathcal{L}$-lines L_3 and L_4.

Case 1: Suppose $L_3 = L_4$. Then L_3 contains all five points, p and the four x_{ij}, and we can completely duplicate the Veblen configuration with $\mathcal{L}_0$-lines $L_3 \cap C_j$, $j = 1, 2, 3, 4$. But by Corollary 14.3.9 $(L, \mathcal{C}_L)$ is a projective space, so the $\mathcal{L}_0$-lines $L_3 \cap C_3$ and $L_3 \cap C_4$ contain a common point, against $C_3 \cap C_4 = \emptyset$.

Case 2: $L_3 \neq L_4$. Since $L_3 \cap L_4$ contains p, partial linearity forces $x_{23} \neq x_{24}$ and $x_{13} \neq x_{14}$. Then $x_{23}x_{24}$ is an $\mathcal{L}_0$-line of C_2 not on p, and so lies in a unique $\mathcal{L}$-line N_2. Similarly, points x_{13} and x_{14} are incident with an $\mathcal{L}$-line N_1, not on p (see Fig. 14.3b). Since, by hypothesis (H3), $(\mathcal{P}, \mathcal{L})$ satisfies the Veblen axiom, and the lines N_2 and N_1 are both transversals of L_1 and L_2, there exists a point v on both N_1 and N_2. Now triangle $\{x_{13}, x_{24}, v\}$ lies in a unique element of $\mathcal{C}$, and this element has to be C_3. Similarly, triangle $\{x_{14}, x_{24}, v\}$ lies in C_4, so now $C_3 \cap C_4$ contains v, a contradiction.

The proof is complete. □

Corollary 14.3.11 $(\mathcal{L}, \mathcal{E})$ *satisfies the Veblen and dual-Veblen axioms (G2) and (H3).*

Proof This follows at once by duality from Lemma 14.3.10. □

Lemma 14.3.12 *Suppose* $|\mathcal{E}| = 1$. *Then either:*

1. $|\mathcal{L}| = 1$, *that is,* $(\mathcal{P}, \mathcal{L})$ *is just a set of points on one line,*
2. $(\mathcal{P}, \mathcal{L})$ *is a non-Desarguesian plane, or*
3. $(\mathcal{P}, \mathcal{L}) \simeq (V_2, V_1)$, *the incidence system of two- and one-dimensional vector subspaces of some (possibly infinite-dimensional) vector space* V.

Proof Suppose $\mathcal{E} = \{\mathcal{L}\}$ and $|\mathcal{L}| > 1$. Then any two lines intersect at a point. Since there are two lines, such intersecting pairs exists, and so the class of cliques $\mathcal{C}$ is non-empty. Choose an element $C \in \mathcal{C}$. Consider any two lines which intersect C non-trivially. By Lemma 14.2.2, part (7), the point of intersection of these lines must lie in C. Thus any two $\mathcal{L}_0$-lines of C intersect, so $(C, \mathcal{L}_0(C))$ is a projective plane. (One could also conclude this from Lemma 14.3.7.)

Now the line-point geometry $(\mathcal{L}, \mathcal{P})$ is a linear space. Since any three (lines of $\mathcal{L}$ which are not on a common point intersect in a member of C which is a projective plane, one must say that any three "points" of $(\mathcal{L}, \mathcal{P})$, not on a common "line," are incident with a projective plane (the dual plane of $(C, \mathcal{L}_0(C))$). Thus $(\mathcal{L}, \mathcal{P})$ is a projective space.

If $|\mathcal{C}| = 1$, then $C = \mathcal{P}$, and $(\mathcal{P}, \mathcal{L})$ is some sort of projective plane. If $|\mathcal{C}| > 1$, then $(\mathcal{L}, \mathcal{P})$ is a projective space of (projective) dimension at least three, and so, by the famous Veblen–Young theorem, is isomorphic to (V_1, V_2) for some vector space V. The statements of this paragraph are just rephrasings of conclusions (2) and (3). □

We conclude with a technical lemma.

Lemma 14.3.13 *Suppose E_1, E_2, E_3 are three elements of $\mathcal{E}$ which, as line-cliques, pairwise have non-empty intersections, but which together do not contain a common line. Then the three lines $\{L_{ij}\} := E_i \cap E_j$ must contain a common point p.*

Proof The lines L_{ij}, $1 \leq i < j \leq 3$ must be distinct, for if two of them were equal, all three line cliques would contain these equal lines. Then for $\{i, j, k\} = \{1, 2, 3\}$, the lines L_{ij} and L_{ik} must intersect at a single point p_i, since these two lines belong to the common line-clique E_i

Assume for the moment, that the three points p_1, p_2, and p_3 are pairwise distinct, so $p_i p_j = L_{ij}$. Then L_{12} and L_{13} belong to the same set $K_{p_1}(L_{23})$ of lines on p_1 transversed by L_{23}. This forces all three lines L_{ij} to belong to the same member of $\mathcal{E}$, that is $E_1 = E_2 = E_3$, against our hypothesis that the sets together contain no common line.

Thus we see that two of the points P_i are equal. By definition $p_i = L_{ij} \cap L_{ik}$, $\{i, j, k\} = \{1, 2, 3\}$, so if two of these points were equal, then in fact, all three are equal, and that will be the common point p of the conclusion of the lemma. □

14.3.3 Adding the Finiteness Conditions and the Proof of Theorem 14.3.1

We begin with the following.

Lemma 14.3.14 *Suppose $\Gamma = (\mathcal{P}, \mathcal{L})$ is a connected partial linear space satisfying axioms (G1), (G2), and (H3). Suppose, in addition, that the finiteness condition (H4)(a) holds. Then there exists an integer $d \geq 1$ such that for each $C \in \mathcal{C}$ and $p \in \mathcal{P}$:*

1. $(C, \mathcal{L}_0(C)) \simeq PG(d + 1)$.
2. $(\mathcal{L}_p, K_p) \simeq PG(d)$.

Proof Axiom (H3)(a) asserts that, for some point p, the incidence system $(\mathcal{L}_p, K_p)$ possesses a finite unrefinable ascending chain of subspaces. The "points" of this geometry are all lines of $\mathcal{L}$ on point p and K_p is the collection of all subsets $K_p(N)$ of lines of L_p transversed by some line N not on p. (Note that $K_p(N) = \{px | x \in p^\perp \cap N\}$ involves at least three lines.)

Now fix a point-clique $C \in \mathcal{C}$ which contains p. By Lemma 14.2.2 (6), each line of $\mathcal{L}_p$ intersects C at an $\mathcal{L}_0$-line. Suppose L_1 and L_2 are lines on p transvered by a line N. Then there are points u_i in $L_i \cap C - \{p\}$, $i = 1, 2$, and so there is a line $N_0 = u_1u_2$ transversing both L_1 and L_2. Since $(\mathcal{L}_p, K_p)$ is a linear space (Lemma 14.3.3), the lines of $\mathcal{L}_p$ transversed by N are the same as those transversed by N_0. By Lemma 14.3.5 $(C, \mathcal{L}_0(C))$ is a projective space, and the sets $K_p(N_0)$, $N_0 \in \mathcal{L}_0(C) - \mathcal{L}_p$, correspond to projective planes in C lying on point p. Thus $(\mathcal{L}_p, K_p)$ is isomorphic to the geometry of lines and planes of $(C, \mathcal{L}_0)$ that are on p, the point-residual of a projective space. Thus $(\mathcal{L}_p, K_p)$ is also a projective space. Since it possesses a finite unrefinable ascending chain of subspaces, we conclude that for some positive integer d, $(C, \mathcal{L}_0(C)) \simeq PG(d+1)$ and $(\mathcal{L}_p, K_p) \simeq PG(d)$.

Since Lemma 14.3.5 shows that the projective spaces $(C, \mathcal{L}_0(C))$ have the same projective dimension for all the cliques of $\mathcal{C}$, the lemma is proved. □

The main result of this section is the following.

Theorem 14.3.15 *Suppose $\Gamma = (\mathcal{P}, \mathcal{L})$ is a connected point-line geometry with $|\mathcal{L}| \geq 2$ satisfying axioms (G1), (G2), and (H3). Suppose, in addition, that $(C, \mathcal{L}_0(C)) \simeq PG(d+1)$ for some positive integer $d \geq 1$. Then $(\mathcal{P}, \mathcal{L})$ is either a non-classical projective plane, or there is a vector space V such that $(\mathcal{P}, \mathcal{L})$ is isomorphic to the geometry (V_{d+1}, V_d) of all $(d+1)$- and d-dimensional vector subspaces of V under the incidence relation of containment.*

Proof First we gather some preliminaries. Because the first three hypotheses (G1), (G2), and (H3) are in force, all of the results of Sect. 14.2.1 and earlier parts of Sect. 14.3 are at our disposal. Thus $\Gamma = (\mathcal{P}, \mathcal{L})$ enriches to the geometry $\Gamma^* = (\mathcal{C}, \mathcal{P}, \mathcal{L}, \mathcal{E})$ where $\mathcal{C}$ is a class of maximal cliques in the point-collinearity graph $(\mathcal{P}, \sim)$ and $\mathcal{E}$ is a class of maximal cliques in the line-graph $(\mathcal{L}, \sim)$.

Our proof will be by induction on the parameter d. Assume first that $d = 1$. Then $(\mathcal{L}_p, K_p)$ is simply a line. This means $\mathcal{L}_p = K_p$, every line on p lies in a common line clique E. It easy see that any point q collinear with p also cannot lie on a line N not in E. Such a line N would be a transversal of two lines pq and px, $x \in p^\perp \cap N - \{p\}$, which belong to E. By connectedness of the line graph, $(\mathcal{L}, \sim)$, one concludes $\mathcal{L} = E$ so $|\mathcal{E}| = 1$. Then the result follows from Lemma 14.3.12.

So we assume $d > 1$, and $|\mathcal{E}| > 1$, henceforward. (G1) for $(\mathcal{P}, \mathcal{L})$ implied $|K| \geq 3$ for each $K \in K_p$, p any point. Select $E \in \mathcal{E}$ and suppose L is a line not in E, which shares points with two lines N_1 and N_2 of E. The latter intersect at a point p, and as L cannot be a transversal of these lines, it too is on p. Thus L is adjacent to as set $K = K_p \cap E$, which, as remarked involves at least three lines. Thus $(\mathcal{L}, \mathcal{E})$ satisfies (G1).

By Corollary 14.3.11, $(\mathcal{L}, \mathcal{E})$ satisfies axioms (G2) and (H3). Finally, $(\mathcal{L}_p, K_p)$, as remarked earlier, is isomorphic to a point-residual of $(C, \mathcal{L}_0(C)) \simeq PG(d+1)$. Thus:

1. $(\mathcal{P}, \mathcal{L})$ satisfies (G1), (G2), (H3), $|\mathcal{L}| > 1$ and $(C, \mathcal{L}_0(C)) \simeq PG(d+1)$, while
2. $(\mathcal{L}, \mathcal{E})$ satisfies (G1), (G2), (H3), $|\mathcal{E}| > 1$ and $(\mathcal{L}_p, \mathcal{E}_p) = (\mathcal{L}_p, K_p) \simeq PG(d)$.

So we may apply induction to conclude that either $(\mathcal{L}, \mathcal{E})$ is a non-Desarguesian projective plane, or else $(\mathcal{L}, \mathcal{E}) \simeq (V_d, V_{d-1})$, for a vector space V.

In the former case, $(\mathcal{L}, \mathcal{E})$ is a linear space, so any two lines belong to a common E in $\mathcal{E}$, and so must intersect. Thus, the line graph $(\mathcal{L}, \sim)$ is just one vast clique, forcing $|\mathcal{E}| = 1$, a case we have already processed.

Thus we may assume there exists an isomorphism

$$\psi : (\mathcal{L}, \mathcal{E}) \to (V_d, V_{d-1}) \tag{14.1}$$

for some vector space V and integer $d \geq 2$.

For each point p, let $X_p := \{\psi(L) | p \in L\} = \psi(\mathcal{L}_p)$, a subset of V_d. Since $(\mathcal{L}_p, \mathcal{E}_p) = (\mathcal{L}_p, K_p) \simeq PG(d)$, a linear space, X_p is in fact a linear subspace of (V_d, V_{d-1}). Thus any two elements x_i, x_j of X_p contain a unique common element of V_{d-1}.

Now let x_1, x_2, x_3 be three distinct elements of X. Then for $\{i, j\} \subseteq \{1, 2, 3\}$, there is a unique $(d-1)$-space y_{ij} contained in both x_i and x_j, and so there is a unique $(d+1)$-space z_{ij} of V spanned by x_i and x_j. Now the d-space x_3 is spanned by the two $(d-1)$-spaces y_{13} and y_{23}, which in turn are subspaces of the $(d+1)$-space z_{12}. Thus x_3 is contained in z_{12} . This forces $x_{12} = z_{13} = z_{23}$. From the general choice of x_1, x_2, x_3, one realizes the following:

(*) *Any two distinct elements of X_p span the same unique $(d+1)$-space, which we denote by z_p.*

This principal (*) produces a mapping

$$\psi' : \mathcal{P} \longrightarrow V_{d+1}$$

which takes point p to the $(d+1)$-space z_p.

WARNING: If $\psi(L)$ is a subspace of z_p for some line L, one cannot then infer that L is in fact incident with p. Of course, if $p \in L$, then $\psi(L) \leq z_p$, as just described above, but at this stage the converse is not known to hold.[6]

Step 1. *ψ' is injective.*

Proof Let us assume p and q are distinct points for which $z_p = z_q$.

Since $(\mathcal{L}_p, K_p) \simeq PG(d)$, and $d > 1$, $|\mathcal{L}_p|$ and $|\mathcal{L}_q|$ are greater than one. Thus p and q (whether collinear or not) are on respective lines L_1 and L_2 with $L_1 \neq L_2$. Then the $\psi(L_i)$, $i = 1, 2$, are d-subspaces of a common $(d+1)$-subspace $z_p = z_q$, and so must intersect at a $(d-1)$-subspace u. Since $\psi : (\mathcal{L}, \mathcal{E}) \to (V_d, V_{d-1})$ is an isomorphism, $u = \psi(E)$ for some $E \in \mathcal{E}$, and the lines L_1 and L_2 belong

[6] Consider the following two spaces : (1) all subspaces of $V_d \cup V_{d-1}$ contained in z_p, and (2) the collection $\psi(\mathcal{L}_p) \cup \psi(\mathcal{E}_p)$, The fact that both are $PG(d)$s does not make the two spaces equal as sets. A division ring D can be isomorphic to a subring D', and so some $PG(d, D)$ can be isomorphic to a proper subset of itself.

to E. Since E is a line-clique, it follows that $L_1 \cap L_2 = \{r\}$ and that $z_r = z_p$ since z_r is spanned by the $\psi(L_i)$ which are already contained in z_p. It follows that each line of $\mathcal{L}_p$, except possibly pr in the case p and r are collinear, is a transversal to all of the lines of $\mathcal{L}_q$. Thus $\mathcal{L}_p \subseteq E$, and so $(\mathcal{L}_p, K_p)$ is a line, $PG(1)$. Since $(\mathcal{L}_p, K_p) \simeq PG(d)$, one is confronted with $d = 1$, a contradiction. □

Step 2. *ψ' is surjective.*

Proof We must show $\psi'(\mathcal{P}) = V_{d+1}$. The graph $(V_{d+1}, \sim)$, whose edges are pairs of $(d+1)$-subspaces meeting at a d-subspace, is connected. So, if $V_{d+1} - \psi'(\mathcal{P})$ were non-empty, it would be possible to choose an edge (z_0, z_p) in the graph with $z_0 \notin \psi(\mathcal{P})$ and $z_p = \psi'(p)$ for some point p. Now select two d-subspaces of z_0. Since $\psi : \mathcal{L} \to V_d$ is surjective, these two d spaces can be written as $\psi(L)$ and $\psi(N)$ for two distinct lines L, N. But being hyperplanes of a common $(d+1)$-subspace z_0, $\psi(L) \cap \psi(N) \in V_d = \psi(\mathcal{E})$. Thus, as ψ is an isomorphism, lines L and N are incident with a common line clique $E \in \mathcal{E}$. It follows that the lines L and N must intersect at a point q. From the definition of the mapping ψ', we see that $\psi(L)$ and $\psi(N)$ span z_q, and so $z_0 = \psi'(q)$. This last statement contradicts our choice of z_0. Thus the difference set $V_{d+1} - \psi'(\mathcal{P})$ is empty, as required. □

Step 3. *$(\psi' \times \psi) : \mathcal{P} \times (\mathcal{L} \cup \mathcal{E}) \to V_{d+1} \times (V_d \cup V_{d-1})$ preserves incidence.*

Proof Suppose L is a line incident with point p. Since $|\mathcal{L}_p| > 1$ there is a second line N meeting L at p. Then by definition, $\langle \psi(L), \psi(N) \rangle_V = z_p = \psi'(p)$, so $\psi(L) \subseteq \psi'(p)$, which is an incidence in (V_d, V_{d+1}).

Similarly, suppose $p \in \bar{E}$, an incidence in $\mathcal{P} \times \mathcal{E}$. This means there is a line L with $p \in L \in \mathcal{E}$. By the previous paragraph, $\psi(L) \subseteq \psi'(p)$, while $\psi(E) \subseteq \psi(L)$ because ψ is an isomorphism. The incidence $\psi(E) \subseteq \psi'(p)$ is a consequence. □

Step 4. *The inverse mapping*

$$(\psi')^{-1} \times \psi^{-1} : V_{d+1} \times (V_d \cup V_{d-1}) \to \mathcal{P} \times (\mathcal{L} \cup \mathcal{E})$$

preserves incidence.

Proof For some $(p, L) \in \mathcal{P} \times \mathcal{L}$, suppose $\psi(N) \subseteq \psi'(p)$. We must show that p is incident with L. Take a line $N \in \mathcal{L}_p$, so N is distinct from L. Then $\psi(L)$ and $\psi(N)$ are distinct hyperplanes of $\psi'(p)$, so they meet at some $\psi(E)$, and so L and N are incident with E since ψ^{-1} is an isomorphism. Then L and N intersect at a point q, and so $\langle \psi L, \psi N \rangle_V = z_q$. But the two spaces $\psi(L)$ and $\psi(N)$ also lie in z_p so $z_p = z_q$. Since ψ is injective (Step 1), this forces $q = p$, so p is incident with L.

Now suppose $\psi(E) \subset \psi'(p)$ for some $(p, E) \in \mathcal{P} \times \mathcal{E}$. Since $\psi(E)$ is a $(d-1)$-subspace and $\psi'(p)$ is a $(d+1)$-subspace, there exists an intermediate space $\psi(N)$ so that

$$\psi(E) \subseteq \psi(N) \subseteq \psi'(p).$$

Then $L \in E$, since ψ was an isomorphism. But $p \in L$ by the result of the previous paragraph. Hence $p \in \bar{E}$, as desired. □

Steps 1–4 allow us to complete the proof of Theorem 14.3.15, for one now sees that the mapping

$$\sigma := \psi' \times \psi \times \psi : \mathcal{P} \times \mathcal{L} \times \mathcal{E} \to V_{d+1} \times V_d \times V_{d-1}$$

is an isomorphism of rank three geometries. Thus the conclusion of the theorem holds for its truncation $(\mathcal{P}, \mathcal{L})$. □

We are now in a position to prove Sprague's theorem.

Proof of Theorem 14.3.1: We assume $(\mathcal{P}, \mathcal{L})$ satisfies $|\mathcal{L}| > 1$ and the following axioms:(G1),(G2),(G3) and either (H4)(a) or (H4)(b). By (G1), $|\mathcal{P}| > 1$, so the same hypotheses hold for the dual-space $(\mathcal{L}, \mathcal{P})$. The hypotheses imply the existence of a family of point-cliques $\mathcal{C}$ and a family of line cliques $\mathcal{E}$ to form an enriched string geometry $(\mathcal{C}, \mathcal{P}, \mathcal{L}, \mathcal{E})$ subject to an extended duality principle exchanging $\mathcal{C}$ and $\mathcal{E}$ as well as $\mathcal{P}$ and $\mathcal{L}$.[7] Since (H4)(a) becomes (H4)(b), it suffices to assume the former, with the obligation to show that $(\mathcal{P}, \mathcal{L})$ is a non-Desarguesian plane, or is (V_{d+1}, V_d) for some vector space d, and positive integer d.

By Lemma 14.3.14, for any $C \in \mathcal{C}$, one has $(C, \mathcal{L}_0(C)) \simeq PG(d+1)$ for some integer $d \geq 1$. We may now apply Theorem 14.3.15 to obtain the conclusion. □

14.4 Exercises for Chapter 14

14.1 Suppose $\Gamma = (V, E)$ is a graph whose (induced) neighborhood graph $\Gamma(x)$ at every vertex x is a grid. Show that either (1) Γ contains two classes of maximal cliques Σ and Π such that for every vertex v, Σ_v and Π_v meet $\Gamma(x) := x^{\perp} - \{x\}$ at the two clique-partitions of its grid-structure, or (2) there is a twofold covering of graphs $\hat{\Gamma} \to \Gamma$ in which triangles lift to triangles and $\hat{\Gamma}$ possesses two families of maximal cliques $\hat{\Sigma}$ and $\hat{\Pi}$ satifying the hypotheses on Σ and Π in (1). [Hint: Just emulate the proof of the covering theorem.]

14.2 For any gamma-space $\Gamma = (\mathcal{P}, \mathcal{L})$ and point p, there is a graph $\Gamma_p = (\mathcal{L}_p, \sim)$ whose vertices are the lines of $\mathcal{L}$ on point p, two of them being adjacent if and only if they lie together in a singular subspace of Γ. Suppose, for each point p, that Γ_p is a grid. Conclude that the (S) axioms (these are the axioms (S1) and (S2) of Chap. 12, p. 444) hold for Γ with $\mathcal{X}$ the collection of all maximal singular subspaces.

14.3 Assume that the hypotheses of Exercise 14.2 hold for a strong parapolar space Γ. Assume in addition that some maximal singular subspace possesses a finite

[7] The reader should be aware that this is a "duality principle," not a "duality." The geometries $\Gamma^* = (\mathcal{C}, \mathcal{P}, \mathcal{L}, \mathcal{E})$ and $(\Gamma^*)' = (\mathcal{E}, \mathcal{L}, \mathcal{P}, \mathcal{C})$ are *not* isomorphic. It is just that any consequence of the axioms of $(\mathcal{P}, \mathcal{L})$ to Γ^* is also a consequence to $(\Gamma^*)'$ when the type-substitutions are made.

unrefinable chain of subspaces. Conclude that either (1) Γ is the classical Grassmann space whose points are the d-subspaces of some vector space V, or else (2) Γ is the quotient geometry $A_{2d-1,d}(D)/\langle\rho\rangle$ of ρ-orbits of d-subspaces of a right vector space V over a division ring D, of a polarity ρ defined by a non-degenerate σ-Hermitian form (σ is an antiautomorphism of D). [Hint: By the previous exercises, the covering theorem, and the first theorem of this chapter, Γ is either itself a Grassmann space, or has a twofold point-line cover which is. In the latter case, the deck-transformation must be defined by the said polarity.]

14.4 Let V be an infinite-dimensional vector space over a division ring D, and let W be a subspace of both infinite dimension and infinite codimension. Let $S(W)$ be the collection of all subspaces U of V such that $U \cap W$ has the same finite codimension in both U and W.

1. Show that if $U \in S(W)$, then $S(U) = S(W)$.
2. Let $\mathcal{P} = S(W)$, and let $\mathcal{L}$ be the collection of all pairs (A, B) of subspaces of V for which there exists a subspace $U \in S(W)$ with $A < U < B$ (proper maximal inclusions). Show that the geometry $(\mathcal{P}, \mathcal{L})$ satisfies axioms (T1)–(T3).

14.5 Suppose H is the generalized hexagon of order $(2, 1)$ whose points are the 21 flags of the Fano plane, and whose lines are of two varieties corresponding to points and lines of the Fano plane. Then every point is on just two lines, one of each variety, and every line has three points. Now suppose $e : H \to \mathbf{P}(V)$ is a projective embedding of H into the points and lines of the projective space $\mathbf{P}(V)$. (This is possible with V a vector space of dimension 6 over the field of two elements. H is embedded by the embedding of the ambient hexagon of type $G_2(2)$. But one must realize that, even so, H is not a subhexagon of the G_2 hexagon, for any two of its points are at distance at least two in the larger hexagon.) Construct a graph Γ whose system of maximal cliques $\mathcal{M}$ satisfies the (S) axioms (axioms (S1) and (S2) of p. 444) where all lines are thin. [Hint: Let the points be the vectors of V.]

14.6 (Sprague) Suppose $\Gamma = (\mathcal{P}, \mathcal{L}, \mathcal{B})$ is a geometry of points, lines and blocks, subject to these axioms:

1. (SL) *Any two distinct lines on a point p are together incident with a unique common block.*
2. (SL*) *Any two lines incident with a common block are incident with a unique common point.*
3. If line L is incident with a point p and a block B, then p is incident with L.

These three hypotheses state that the rank three geometry of points, lines, and blocks belong to a "dual-linear/linear" diagram (L^*, L). They state nothing more than this.

Assume, in addition, these hypotheses:

(Sp) *Blocks are subspaces of* $(\mathcal{P}, \mathcal{L})$.
(Sp*) (The dual of the preceding axiom.) *Each line* L *is the intersection of any two blocks which contain it.*

Show that both the Pasch axiom, and the dual-Pasch axiom (the "diagonal axiom" expounded in Sect. 14.3) hold for $(\mathcal{P}, \mathcal{L})$. [Hint: By duality of the hypotheses, and axiom (SL^*), it suffices to prove the Pasch/Veblen axiom for $(\mathcal{P}, \mathcal{L})$.][8]

[8] The reader should be warned that this result of Sprague is a highly misunderstood theorem. A straight reading of his paper (Sprague (1985) [123]) would lead a modern reader to believe that geometries with finitely many points satisfying only the "diagram axioms" (SL) and (SL*) are either Grassmann spaces over generalized projective spaces, and in particular, if all lines are thin, the (c^*, c)-geometries of $(i-1)$-, i-, and $(i+1)$-subsets of an n-set. The theorem has been quoted in this guise (that is, as a characterization of (possibly finite) (L^*, L)-geometries) in a number of sources ([25]) and in a footnote regarding theorem in the *Handbook of Incidence Geometry*, clarified later in the article of Buckenhout and Pasini in the same source). This is not really the fault of the editors and writers, for the original paper [123] gives no hint at all in the statement of the main theorem (which allows thin lines) that the two extra axioms (Sp) and (Sp*) are being invoked, as a perusal of his proof (lines 8 through 10 on p. 258 of [123]) shows that they are. The author's guess is that these assumptions were incorporated in the phrase "Buekenhout incidence geometry" or "diagram geometry" as it was understood at that time – some 20 years ago. At times, the intersection hypothesis (Int) was understood to be included by that notion, although that is not the current view.

Let (Γ, τ, I) be a geometry with set I of realized types. Fix $i \in I$. For each flag F let $sh_i(F)$, the *i-shadow of flag F*, be the set of objects of type i incident with flag F. (If $i \in \tau(F)$ this set is a single object of type i). The property being discussed is this:

(1) (Int) *Suppose* A_1 *and* A_2 *are the i-shadows of flags* F_1 *and* F_2*, respectively. Then the intersection* $A_1 \cap A_2$ *must be the i-shadow of some further flag G (despite the relationship of* F_1 *and* F_2*). That is, the i-shadows are closed under intersections.*

This is a very strong axiom. It would explain the two axioms (Sp) and (Sp*), for these are consequences of this general "functorially alien" axiom, whose strength waxes with the number of types involved and wanes under truncation.

What are the counter-examples of the naive understanding of Sprague's theorem? Ironically, they are often presented in the very surveys which were allowed to pass on Sprague's theorem as it seemed to be presented then, without the requisite updating of the notion of "diagram geometries." In the thin case, it has been known for some time that there are many other (c^*, c)-geometries (the most comprehensive survey up to the time of this writing is probably [112] although even that is not at all exhaustive). A simple example is given in Exercise 14.7. Also, it has been well-known for one-and-a-half decades that one can obtain a rank three geometry with diagram (Af^*, Af) simply by removing a geometric hyperplane from a classical Grassmannian. (This is how we get most geometries with Af replacing the projective-plane residues on all edges leaving the node representing points.) The really definitive hard theorem here is the theorem of Cuypers [45] characterizing this Grassmann-minus-hyperplane example. In any event, the (Af^*, Af)-geometries are certainly (L^*, L)-counterexamples to the usual rendering of Sprague's theorem.

14.7 Suppose (V, Q) is a non-degenerate quadratic form $Q : V \to GF(q)$ on a 5-space V over $GF(q)$ where q is odd. There is associated with Q, a bilinear form $B : V \times V \to GF(q)$. We write $X \perp Y$ if, for every pair of vectors $(x, y) \in X \times Y$, one has $B(x, y) = 0$. As we all know, the form (V, Q) forces a partition of the 1-spaces of V into three classes:

1. Singular 1-spaces $\langle v \rangle$, where $v \neq 0$ but $Q(v) = 0$.
2. The set $\mathcal{N}^+$ of 1-subspaces $\langle v \rangle$, where $v^\perp$ is a hyperbolic quadric $Q(3, q)$ (of type $O^+(4, q)$, for the group-theorists).
3. The set $\mathcal{N}^-$ of 1-subspaces $\langle v \rangle$, where $v^\perp$ is an elliptic quadric $Q(3, q)$ (of type $O^-(4, q)$).

Let $\mathcal{P} = \mathcal{N}^-$. Let $\mathcal{L}$ be the collections of subsets of $\mathcal{N}^-$ which lie in a subspace of the form $\langle n \rangle \perp s$, where s is a singular 1-space, $n \in \mathcal{N}^-$. Thus each line contains exactly q points. Finally, let $\mathcal{B}$ be the collection of all non-degenerate 4-spaces which support an elliptic quadric. Show that $(\mathcal{P}, \mathcal{L}, \mathcal{B})$ is a (c^*, c)-geometry.

Chapter 15
Characterizing the Classical Strong Parapolar Spaces: The Cohen–Cooperstein Theory Revisited

Abstract Virtually all parapolar space characterizations of Lie incidence geometries rest on a trail of logical dependencies leading back to Arjeh Cohen's characterization of Grassmannians, presented here. The first tier of consequences of Cohen's theorem is summarized in the two theorems of Cohen and Cooperstein, the first of which falls within this chapter. Further explorations liberate their theory from the assumption of constant symplectic rank.

15.1 Prognosis

One important theorem on the characterization of Grassmannians has thus far not appeared, so in a sense Chap. 14 is incomplete. I refer to the deeper theorem of Arjeh Cohen (Sect. 15.3, Theorem 15.3.7) where the characterization depends on the relation of a point and a symplecton. Cohen's characterization was not presented in the previous chapter for the simple reason that it naturally resides as part of a more general theory introduced by Cohen and Cooperstein which takes place on the explicit stage of strong parapolar spaces. To put it simply, a book like this is about proofs, not conclusions, and so it is the *context of the proofs* which causes one to defer theorems to another chapter, not the taxonomy of their conclusions.

Cohen's theorem requires a lemma which is absolutely crucial to every single characterization of parapolar spaces that is to follow in this and subsequent chapters of this book. Cohen's lemma is necessary for his characterization of Grassmann and twisted-Grassmann spaces and in turn the Cohen–Cooperstein characterizations of classical strong parapolar spaces and their generalizations, which are the subject of this chapter. These theorems, in turn, are used in the succeeding chapters characterizing the long-root geometries and the homomorphic images of the polar Grassmannians. So this one small lemma of Cohen becomes a sort of "logical isthmus" connecting the previously known characterizations, to the world of characterizing parapolar spaces. Sometimes the most uncelebrated results are the most crucial!

It is also interesting to observe that Cohen's crucial lemma (Lemma 15.3.1 of this chapter) depends on Tits' classification of the rank-three polar spaces which was sketched in Chap. 7. So every characterization result in the rest of this book – even

E.E. Shult, *Points and Lines*, Universitext, DOI 10.1007/978-3-642-15627-4_15,

© Springer-Verlag Berlin Heidelberg 2011

those with seemingly transparent proofs – depend on deep theorems of several decades ago. (But see the footnote and reference on p. 530.)

The first theorem of Cohen and Cooperstein concerns strong parapolar spaces, and our version of this theorem together with a number of applications of it is the focus of this chapter. Chapter 17 contains their second theorem on so-called "long root geometries." Both theorems appeared in a remarkable paper [36]. Building on the polar space classification, it became the first paper to characterize – by simple axioms on points and lines of a parapolar space – every single one of the thick spherical buildings of rank at least three which were not already projective or polar spaces.

The first of the two Cohen–Cooperstein theorems had its beginning in Cohen's characterization of the Grassmann spaces ([34]) mentioned at the beginning of the section. At the same time, Cooperstein as well as Cohen envisioned a general characterization of strong parapolar spaces using Cohen's characterization of Grassmannians as an inductive step. So it was natural that both authors should almost immediately collaborate in their famous paper [36], giving their characterizations for the strong parapolar case (their first theorem), and then derive from it their characterizations of the long-root geometries (another class of parapolar spaces) in their second theorem.

In this and the next chapters it is my hope to recast this theory in a theoretically more translucent way by (1) incorporating the methodology concerning covers of geometries in locally truncated geometries and (2) widening the front gate to this theory, so that all sorts of geometries must fall into it.

15.2 Introduction

For any parapolar space, and any non-incident point-symplecton pair (x, S), the set $x^\perp \cap S$ of points of S which are collinear with x must form a singular subspace of S which is a projective space naturally possessing some projective dimension. The first Cohen–Cooperstein theorem concerns strong parapolar spaces of finite uniform symplectic rank r, for which there is a gap in the spectrum of possible projective ranks of $x^\perp \cap S$ among non-incident point-symplecton pairs.

We shall find it necessary to consider such spectrum restrictions outside the context of parapolar spaces of *uniform* symplectic rank. For this reason we introduce the following hypothesis:

(CC)$_k$ *Let k be a fixed non-negative integer. For any non-incident point-symplecton pair (x, S), $x^\perp \cap S$ is either (1) a projective space whose rank is less than $k+1$, or (2) a maximal singular subspace of S.*

Remark The reader is again reminded that in this book the *rank* of a projective space is its rank as a matroid – one more than its projective *dimension* (see p. 141). A projective space has *rank r* if and only if it has *projective dimension* $r-1$.

Two observations are immediate. (1) If $i < j$, then $(\mathrm{CC})_i$ implies $(\mathrm{CC})_j$. (2) If no symplecton has polar rank more than $k+1$ the condition $(\mathrm{CC})_k$ asserts nothing at all. But if the polar rank of S exceeds $k+1$, the hypothesis proposes that there is a "gap" in the spectra of possible projective ranks of $x^\perp \cap S$.

In this case, assuming constant symplectic rank r, Cohen and Cooperstein were able to conclude that Γ is either a highly restricted quotient of a Grassmannian, or a half-spin geometry, or is one of the exceptional Lie incidence geometries $E_{6,1}$ or $E_{7,7}$ (in the Bourbaki scheme).

Because of observation (1), the critical cases occur when $r-k$ and r are both as small as possible. In that case the assumptions are:

(Co1) Γ *is a strong parapolar space of constant symplectic rank three.*
(Co2) *For any non-incident point-symplecton pair* (x, S) *one has* $x^\perp \cap S$ *is either empty, a point, or a maximal singular subspace of* S.

Cohen's theorem appearing in [34], which handles this critical case, asserts that under hypotheses (Co1) and (Co2), Γ is either a polar space, a Grassmann space, or a quotient of a Grassmann space by the orbits of a suitable polarity.

Of course, the case that Γ is itself a polar space is easily isolated. The other two conclusions will hold once it is established that every line lies in exactly two maximal singular subspaces. But establishing this is not an easy endeavor. As indicated in the previous section, Cohen's proof of this fact (described in the next section) uses Tits' classification of all polar spaces of rank three to prove the vital technical Lemma 15.3.1. In contrast, the hypotheses and the stated conclusion of that lemma involve only simple facts about points and lines.

Section 15.4 then gives the full classification of Cohen and Cooperstein for strong parapolar spaces with $(CC)_k$ for appropriate k and constant symplectic rank.

In Sect. 15.5 we ask where these hypotheses come from. There is no apparent reason to believe that God declared this gap in this spectra of dimensions of $x^\perp \cap S$. The idea of this chapter is to reduce this hypothesis to one that sounds and feels "more natural."[1] Our substitution for $(\mathrm{CC})_1$ is this:

$(U)_2$ *If the intersection of two symplecta* A *and* B *contains a projective plane, then each point* x *in* $A - (A \cap B)^\perp$ *is collinear with at least one point of* B *not lying in* A.

This point of view can claim two benefits: (1) it allows us to dispense with the assumption of constant symplectic rank altogether, and (2) provides a way of handling the long root geometries in Chap. 17 in a way that is also free of this assumption.

[1] The quotes are important: whatever "more natural" hypothesis we are talking about will be logically equivalent to the old one, so we are not talking about a logical generalization so much as a psychological one. We shall follow this tradition by showing it is possible to see that the highly specific $(\mathrm{CC})_1$ hypothesis is itself the consequence of one that does not seem particularly drastic.

15.3 Cohen's Characterization of Grassmann Spaces

15.3.1 A Technical Result

In any point-line geometry $\Gamma = (\mathcal{P}, \mathcal{L})$, a *line spread* is a collection $\mathcal{R}$ of lines (i.e., $\mathcal{R} \subseteq \mathcal{L}$) such that every point is incident with a unique member of $\mathcal{R}$.

We begin with an important technical lemma of Arjeh Cohen.

Lemma 15.3.1 (Cohen [34].) *Let $\Gamma = (\mathcal{P}, \mathcal{L})$ be a non-degenerate generalized quadrangle with a line spread $\mathcal{R}$ having these properties:*

1. *The subspace generated by any two distinct members of $\mathcal{R}$ is a grid, one of whose reguli is a subset of $\mathcal{R}$.*
2. *The rank-two incidence geometry $\Pi := (\mathcal{R}, \mathcal{G}, *)$, where $\mathcal{G}$ is the system of generated grids and $*$ is the natural incidence of a line and a grid, is a projective space.*
3. *Γ is the point-residual of a (non-degenerate) polar space of rank three.*

Then Π is a line and Γ is a grid.

Remark Cohen's proof of this technical result employs the full classification of all polar spaces of rank three. For a complete account of his proof the reader is referred to his article [34].[2]

15.3.2 Cohen's Hypotheses and Their Elementary Consequences

For the remainder of this section we assume that Γ satisfies these hypotheses:

(Co1) Γ *is a locally connected parapolar space of constant symplectic rank three.*[3]
(Co2) Γ *satisfies the following condition:*

(CC_1) *For each non-incident point-symplecton pair $(p, S) \in \mathcal{P} \times \mathcal{S}$, one has that $x^\perp \cap S$ is either* (1) *empty,* (2) *consists of a single point, or (3) is a maximal singular subspace of S – in this case, a plane.*[4]

(Co3) Γ *is not itself a polar space.*

Throughout, $\mathcal{S}$ is the full collection of all symplecta of Γ and $\mathcal{M}$ is the full collection of all maximal singular subspaces of Γ.

---

[2] Recently (2007), Professor K. Thas discovered a simpler proof of Lemma 15.3.1 [118]. His proof utilizes two facts: (1) any generalized quadrangle that is a point-residual of a rank three polar space is Moufang, and (2) any subquadrangle of a Moufang quadrangle is also Moufang (Proposition 11.5.2, p. 238 of [132]).

[3] In [34] it was assumed that Γ was a strong parapolar space.

[4] In Cohen's original paper, $(\mathrm{CC})_1$ is replaced with a different but equivalent hypothesis.

Lemma 15.3.2 *Suppose S is a symplecton and x and y are collinear points of $\mathcal{P}-S$ such that $x^\perp \cap S$ is a plane π_x. Then $y^\perp \cap S$ is either contained in π_x or is itself a plane π_y meeting π_x at a single point.*

Proof Suppose $z \in y^\perp \cap S - x^\perp$. Then, since π_x is a maximal singular subspace of the polar space S, $z^\perp \cap \pi_x$ is a line L, and there is a symplecton $R^x := \langle\langle x, z\rangle\rangle_\Gamma$ containing y and the plane $\langle L, z\rangle_S$. Thus $y^\perp \cap S$ contains a line of this plane, and so by hypothesis (Co2), $y^\perp \cap S := \pi_y$ is a plane. Now if $\pi_x \cap \pi_y$ contained a line, this line must be $z^\perp \cap \pi_y$ since $y^\perp \cap S$ is a singular space containing z. But in that case, the symplecton R would contain $\langle y, z^\perp \cap \pi_y, z\rangle \simeq PG(3, D)$, against the assumption of constant symplectic rank three. Thus we must assume that $\pi_x \cap \pi_y$ is a single point. The proof is complete. □

Lemma 15.3.3 *Fix a symplecton S and set*

$$N(S) := \{p \in \mathcal{P} | p^\perp \cap S \text{ contains a plane}\}.$$

Then $N(S)$ is a subspace of Γ.

Remark Note that by definition $S \subseteq N(S)$. We refer to $N(S)$ as the *set of points near S* symplecton, points near.

Proof Suppose x and y are distinct points of $N(S)$ lying on some line L of Γ. Choose any point $z \in L - \{x, y\}$. We must show $z \in N(S)$.

If x and y lie in S, there is nothing to prove as $L \subseteq S \subseteq N(S)$. Suppose instead that $L \cap S = \{u\}$, a single point. Now if z were a point of $(x^\perp \cap S) - (y^\perp \cap S)$, then z and u would be members of a clique $x^\perp \cap S$, and hence would be collinear. By the gamma space property of parapolar spaces, z would belong to $L^\perp \cap S$. It follows that z is collinear to y as well. Thus we see (arguing from the symmetric roles of x and y) that $L^\perp \cap S = x^\perp \cap S = y^\perp \cap S$, so $z \in N(S)$.

So we may assume that $L \cap S = \emptyset$. Now by the previous lemma, either $y^\perp \cap S \subseteq x^\perp \cap S$ or $y^\perp \cap S = \pi_y$ is a plane of S meeting the plane $\pi_x := x^\perp \cap S$ at a point p. In the former case $y \in N(S)$ forces

$$\pi_x = x^\perp \cap S = y^\perp \cap S \subseteq z^\perp \cap S,$$

whence $z \in N(S)$. So we may assume the second alternative holds: $\pi_x \cap \pi_y = \{p\}$ exactly. Choose any point b in $\pi_y - \{p\}$. Then $b^\perp \cap \pi_x$ is a line N on point p. Then $N \subseteq x^\perp \cap b^\perp$ and $R = \langle\langle x, b\rangle\rangle_\Gamma$ is a symplecton containing the point x, the line $L = xy$, and the plane $\pi := \langle N, b\rangle_S$. Then $z^\perp \cap \pi$ is a line B of S, whence $z^\perp \cap S$ is a plane by (Co2), and we have $z \in N(S)$ in this case as well. □

Lemma 15.3.4 *Let M be any maximal singular subspace and let S be any symplecton of Γ. Then $M \cap S$ is either empty, consists of a single point, or is a plane. It is never just a line.*

Proof If $(M, S) \in \mathcal{M} \times \mathcal{S}$ with $M \cap S = L \in \mathcal{L}$, then $L \neq M$ since S has polar rank three, and M is a maximal singular space. Thus there is an element

$x \in M - S = M - L$. Then $x^\perp \cap S$ is a plane π_x by (Co2). Then M is the unique maximal singular subspace containing $\langle x, L\rangle$ and so contains $\langle x, \pi_x\rangle$, whence $\pi_x \subseteq M$, a contradiction. □

Corollary 15.3.5 *If $(x, M) \in \mathcal{P} \times \mathcal{M}$ with x not in M, then $x^\perp \cap M$ is either empty, or is exactly a line.*

Proof We may assume $x^\perp \cap M$ is non-empty. Choose $p \in x^\perp \cap M$ and select $y \in M - x^\perp$. If (x, p) is a polar pair, then $R := \langle\langle x, y\rangle\rangle_\Gamma$ is a symplecton meeting M in a plane π (by Lemma 15.3.4). Then $x^\perp \cap \pi$ is a line, so $x^\perp \cap M$ contains a line. If $x^\perp \cap M$ contained a plane π, then $\langle \pi, x\rangle$ would be a $PG(3)$ in S, against S having polar rank three. Thus $x^\perp \cap M$ is just a line.

There remains the case that (x, p) is a special pair. Since Γ is locally connected, there is a geodesic path in $\mathrm{Res}_\Gamma(p)$ connecting "point" xp to "point" py. In other words, there is a sequence of lines

$$L_0 = px, L_1 = x_1, \ldots L_{n-1} = px_{n-1} = L_n = py$$

such that pairs of successive lines, L_i, L_{i+1}, generate a plane π_i, $i = 0, \ldots, n-1$. Now ${x_{n-2}}^\perp \cap y^\perp$ contains the line L_{n-1}, while, by the minimality of n, x_{n-2} is not collinear with y. It follows that (x_{n-2}, y) is a polar pair, while, by the assumption of this paragraph, (x_0, y) is not. Choose j minimal, so that (x_j, y) is a polar pair. So there is a symplecton T containing $\{x_j, y\}$. Now ${x_{j-1}}^\perp \cap T$ contains the line px_j, so by $(\mathrm{CC})_2$ this intersection contains a plane π_j. Since T is a polar space containing y, $y^\perp \cap \pi_j$ contains a line N which lies in ${x_{j-1}}^\perp \cap y^\perp$. But that forces (x_{j-1}, y) to be a polar pair, against the minimality of j. So the case that (x, p) is a special pair is not viable. This completes the proof. □

Lemma 15.3.6 *For every symplecton S, there exists a maximal singular subspace M which intersects S in exactly one point.*

Proof First we show that there exists a symplecton meeting S at a plane. Since S is not a connected component of Γ (otherwise the connected Γ would be a polar space against (Co3)), there is a line $L_0 = px$ intersecting S at a single point p. If $L_0 \subseteq N(S)$ set $L = L_0$. Otherwise, ${L_0}^\perp \cap S = \{p\}$. In this case, the local connectedness of Γ produces a minimal sequence of lines,

$$L_0, L_1, \ldots L_{n-1}, L_n,$$

each incident with point p, with successive pairs generating a plane, and $L_n \in N(S)$. The minimality of n shows that $L_{n-1} \notin N(S)$. This in turn shows us that $L_n \notin S$. Now set $L = L_n$ in this case. Thus in all cases, we have found a line L on p, not lying in S, but lying in $N(S)$.

Since $L \subseteq N(S)$, there is a line N on p in S such that $\langle\langle L, N\rangle\rangle_\Gamma$ is a symplecton R.[5] Then $R \cap S$ contains the line N. If this intersection contains a plane, we have found our desired symplecton R. Otherwise we have $R \cap S = N$ exactly. Now

[5] Note that we have not invoked any strong parapolar assumption here.

there exist planes π_S and π_R on line N in symplecta S and R respectively so that $\pi_S \not\subseteq \pi_R{}^{\perp}$. Then $T := \langle\langle \pi_S, \pi_R \rangle\rangle$ is a symplecton containing $\pi_S \cup \pi_R$ and meeting S at plane π_S as desired.

Now suppose S_1 and S_2 are two symplecta such that $S_1 \cap S_2 = \pi$ is a plane (guaranteed for any desired S_1 by the previous paragraph). Since S_2 is a polar space of rank three, there is a plane $\pi_2 \in S_2$ with $\pi_2 \cap \pi = \{p\}$, a single point. Now there is a unique maximal singular $M_2 \in \mathcal{M}$ containing plane π_2. Then $M_2 \cap S_1$ is either just the point p, or is a plane π_1 containing the point p (Lemma 15.3.4).

By way of contradiction, assume the latter case. Then we see that for any point $x \in \pi_2 - \{p\}$, the subspace $x^{\perp} \cap S_1$ must contain π_1. But as S_2 is a polar space, $x^{\perp} \cap \pi$ is a line A. In that case $x^{\perp} \cap S_1$ contains $A \cup \pi_2$ which cannot be a clique since S has polar rank three. By convexity of symplecta $x \in S_1 \cap \pi_2 = \{p\}$, contrary to the choice of x.

Faced with this absurdity, one must assume the former case holds, namely that

$$M_2 \cap S_1 = \{p\}.$$

The proof is complete. □

Now we shall study carefully the following situation.

S is an arbitrary symplecton, and M is a maximal singular subspace whose intersection with S is just a single point p.

Clearly the subspace $M \cap N(S)$ contains point p. But actually, this intersection contains a plane.

One may be convinced of the the latter claim by considering any line B of S on point p. Then by Corollary 15.3.5, $B^{\perp} \cap M$ is a line N of M also on point p. Now every point of N belongs to $N(S)$ since its points either belong to S or are collinear with all points of a line B of S. We can repeat this argument with a line B' of S on p with $B' \not\subseteq B^{\perp}$. Then $(B')^{\perp} \cap M = N'$ is a line $N(S) \cap M$ containing p. But we cannot have $N = N'$ since otherwise $N^{\perp} \cap S$ would contain $B \cup B'$ which is not a clique. Thus $M \cap N(S)$ contains at least two lines on point p and, being the intersection of two subspaces (Lemma 15.3.3), it must contain the plane these lines generate.

At this stage, we have an injection:

$$\phi : \text{lines of } N(S) \cap M \text{ on point } p \to \text{planes of } S \text{ on } p,$$

sending line L on point p to $L^{\perp} \cap S$, a plane on p. Let $\mathcal{R}$ be the image of ϕ. Then $\mathcal{R}$ is a collection of planes of S on point p with the following property.

1. *Every line of S on p belongs to a unique member of $\mathcal{R}$.*

The first observation is that the mapping ϕ is injective: for if $\phi(N) = \phi(N') = \pi$ for two distinct lines N and N' of M on p, then $\pi^{\perp} \cap M \supseteq \langle N, N' \rangle$ against Corollary 15.3.5. We have already seen that any line B of S on p is in the plane $\phi(B^{\perp} \cap M)$ and the injectivity of ϕ produces the uniqueness. So the property presented holds.

Now consider an arbitrary pair of distinct planes π_1 and π_2 of $\mathcal{R}$. We let N_1 and N_2 be the unique lines of $M \cap N(S)$ on p for which $\phi(N_i) = \pi_i$, $i = 1, 2$. For each line A_1 of π_1 on point p, we obtain a unique line $A_2 = {A_1}^{\perp} \cap \pi_2$ and a plane $\pi_0 := \langle A_1, A_2 \rangle_S$. Form the symplecton $R := \langle\langle A_1, N_2 \rangle\rangle_\Gamma$. (This exists since ${N_1}^{\perp} \cap {A_1}^{\perp}$ contains the line A_2.) Then R contains N_1 and A_2 and hence contains the two planes $\pi_0 = \langle A_1, A_2 \rangle = R \cap S$ and $\pi = \langle N_1, N_2 \rangle_M = R \cap M$. Now since R is a polar space, for any line C of π_0 on p, one has $C^{\perp} \cap M = C^{\perp} \cap \pi$, and, conversely, for each line N of π on p, one sees that $N^{\perp} \cap S$ is a plane of R meeting π_0 at a line on p. Thus, letting $\mathcal{L}_p \cap \mathrm{Res}_\Gamma(\pi)$ be the pencil of lines of plane π which are on point p, and letting $R(\pi_0)$ be the planes of the symplecton R which meet π_0 at a line on p, π induces a bijection:

$$(*) \quad \mathcal{L}_p \cap \mathrm{Res}_\Gamma(\pi) \to R(\pi_0).$$

This configuration has an especially nice interpretation as a mapping among subspaces of the point-residual

$$\mathrm{Res}(p) := (\mathcal{L}_p, \Pi_p),$$

the incidence geometry of lines and planes of Γ which contain the point p. For any subspace X of Γ containing p, let X_p be the set of lines in X on point p – that is, $X_p := \mathrm{Res}_\Gamma(X) \cap \mathcal{L}_p$. Then X_p is always a subspace of $\mathrm{Res}(p)$.

In particular, the symplecton S and maximal singular subspace M discussed above induce a convex generalized quadrangle S_p and a maximal projective space M_p and its subspace $N_p := (N(S) \cap M)_p$. Then the system of planes $\mathcal{R}$ induces a *line spread* $\mathcal{R}_p$ of the generalized quadrangle S_p, and we have a bijection:

$$\phi_p : N_p \to \mathcal{R}_p.$$

But there is more. For every "line" π_p of N_p, there corresponds the elements $\phi_p(\pi_p)$ which are traversed by unique lines of the quadrangle S_p, so that

$$\phi_p(\pi_p) = \{\phi(L)_p | L \in \mathcal{L}_p \cap \mathrm{Res}_\Gamma(\pi)\}$$

and its traversals form the two parallel classes of lines of a grid. (This is the content of (*).) Moreover, any two lines of $\mathcal{R}_p$ generate such a grid. (This is because our discussion creating $R(\pi_0)$ began with two distinct but otherwise arbitrary planes π_1 and π_2 of R.)

It follows that the collection $\mathcal{G}_p$ of grids generated by members of $\mathcal{R}_p$ are in one-to-one correspondence with the lines of N_p – indeed, that ϕ_p induces an isomorphism

$$(\mathcal{R}_p, \mathcal{G}_p) \to (N_p, \Pi_p \cap \mathrm{Res}_\Gamma(N(S) \cap M)) \tag{15.1}$$

as projective spaces.

Suppose the projective space in (15.1) contains a plane Π whose points and lines are $(\mathcal{R}'_p, \mathcal{G}'_p)$. Then $S'_p := \langle\langle \mathcal{R}'_p, \mathcal{G}'_p \rangle\rangle\rangle$ is a subquadrangle of the classical residual quadrangle S_p.

At this point $(S'_p, \mathcal{R}'_p, \mathcal{G}'_p)$ satisfy the conditions on $(\Gamma, \mathcal{R}, \mathcal{G})$ in Cohen's lemma (Lemma 15.3.1). It follows that the projective space S'_p is not a subquadrangle of a classical quadrangle. But that is a contradiction since S_p is classical. Thus, the projective space $(\mathcal{R}_p, \mathcal{G}_p)$ is a line, S_p is itself a grid, and N_p is a line. We have thus shown the following.

Theorem 15.3.7 (Cohen) *Suppose $\Gamma = (\mathcal{P}, \mathcal{L})$ is a strong parapolar space satisfying the axioms (Co1), (Co2), and (Co3). Then:*

1. *Every symplecton is an oriflame polar space of rank three – that is, a Lie incidence system of type $A_{3,2}(D)$ for some division ring D.*
2. *Every line lies on exactly two maximal singular subspaces.*
3. *If Γ contains a maximal singular subspace of finite projective rank then either:*
   (a) *Γ is a Grassmann space of d-subspaces of a vector space V over a division ring D, where d is a finite integer exceeding 1. (V may be infinite-dimensional.)*
   (b) *$\Gamma \simeq A_{2n-1,n}(D)/\langle\sigma\rangle$, a quotient of a Grassmann space as the orbit geometry of a polarity σ of $PG(2n-1, D)$ of Witt index at most $n-5$.*

Proof Conclusion (1) is a consequence of the fact that S_p is a grid.

For conclusion (2), recall that in a parapolar space every line L lies in some symplecton S. Now Lemma 15.3.4 and the fact that $x^\perp \cap S$ is a clique for each point x outside S together force a bijection between the planes of S on line L and the maximal singular subspaces of Γ on line L. Since the former set has cardinality two, so does the latter.

Now we see that $(\mathcal{L}_p, \mathcal{M}_p)$ is itself a grid. Then the conclusions (a) and (b) of (3) follow from Exercise 14.3, p. 524, of Chap. 14. □

15.4 Strong Parapolar Spaces with Constant Symplectic Rank $r > 3$ and Satisfying $(CC)_{r-2}$

15.4.1 Introduction

As promised, this section gives a slightly updated version of the remarkable theorem of Cohen and Cooperstein for strong parapolar spaces. In this section we assume the following.

(**) *Γ is a strong parapolar space of constant symplectic rank r satisfying the condition $(CC)_{r-2}$. The number r is finite.*[6] *For convenience, we assume that Γ is not itself a polar space.*

[6] Recall that for a parapolar space of symplectic rank exactly r, condition $(CC)_{r-2}$ asserts that if x is a point not incident with a symplecton S, then $x^\perp \cap S$ can never be a hyperplane of a maximal singular subspace of S. Thus there is only one rank, excluded from the possible ranks of $x^\perp \cap S$.

In addition to adding the strong parapolar property, we depart from the previous section by adding two further assumptions: (1) $r \geq 4$, and (2) all singular subspaces of Γ have finite projective rank.

These hypotheses seem to inherit to point-residues extremely well. As usual, for any point p we can form the point-residual geometry $\mathrm{Res}(p) = (\mathcal{L}_p, \Pi_p)$ of lines and planes of Γ which are incident with p, and, in turn, regard it as a geometry of points and lines. Then the following applies.

Lemma 15.4.1 *Each point-residual Res(p) is a strong parapolar space whose point-collinearity graph has diameter two. In addition, since Γ satisfies condition (**) with $r \geq 4$ and statements (1) and (2) of the preceding paragraph, then:*

1. *Res(p) has constant symplectic rank $r - 1 \geq 3$.*
2. *Res(p) has all singular subspaces of finite rank.*
3. *For any non-incident point-symplecton pair (x, S) of the point-residual Res(p), $x^\perp \cap S$ has projective rank at most $r - 3$ (projective dimension at most $r - 4$), or is a maximal singular subspace of S* (the condition $(\mathrm{CC})_{r-3}$ for a parapolar space of symplectic rank $r-1$) *– that is, $x^\perp \cap S$ is not a hyperplane of a maximal singular subspace of S.*
4. *Res(p) is not a polar space.*

Proof The first statement follows from general material about point-residuals of parapolar spaces from Chap. 13. The four conditions in the second statement are direct consequences of the fact that in $\mathrm{Res}(p)$, a symplecton simply consists of the lines and planes of a symplecton of Γ incident with point p, and that a similar statement holds with singular subspaces replacing symplecta (part 3 of Theorem 13.4.1, p. 463).

Now suppose for some point p, $\mathrm{Res}(p)$ were a polar space. Since Γ is a strong parapolar space of symplectic rank at least three, it is locally connected by Lemma 13.5.1, p. 467, of Chap. 13. Thus Corollary 13.5.6 (p. 472 of Chap. 13) can be invoked to conclude that Γ itself is a polar space, contrary to (**). □

15.4.2 The Case $r = 4$

Now by Lemma 15.4.1, $\mathrm{Res}(p)$ satisfies conditions which at least include those of Cohen's Theorem (Theorem 15.3.7). Thus $\mathrm{Res}(p)$ is a Grassmannian or a twisted quotient of one. Since $\mathrm{Res}(p)$ must have point-diameter two, the quotient conclusion is ruled out, and the residual is the Grassmannian of lines of some projective space of finite rank n_p. Four of the objects of this Grassmannian correspond to actual objects of Γ, giving us the diagram of Fig. 15.1.

Here, in addition to the familiar points and lines and symplecta of Γ (denoted $\mathcal{P}$, $\mathcal{L}$, and $\mathcal{S}$, respectively), we have a class of maximal singular subspaces $\mathcal{M}$ of projective dimension at least 4, and a class $\mathcal{A}3$ of maximal singular subspaces which are PG(3)s and play the role of one class of maximal singular subspaces of

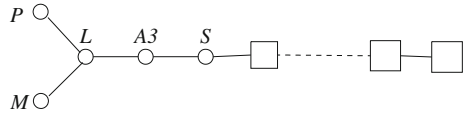

Fig. 15.1 The apparent locally truncated diagram of Γ when $r = 4$

the symplecta which are oriflame of type D_4. This now causes the planes to be coordinatized by a field, rather than an arbitrary division ring.

Of course we cannot assert that Γ is in fact a locally truncated diagram geometry belonging to the diagram of Fig. 15.1 until we establish that all singular subspaces of $\mathcal{M}$ possess the same dimension as projective spaces. But it follows from the local diagram of that figure that any two members of $\mathcal{M}$ lying on a common point p already possess the same rank. Since every line of Γ lies in a member of $\mathcal{M}$, we see that the connectedness of $(\mathcal{P}, \sim)$ now implies that all members of $\mathcal{M}$ have the same projective rank $n - 1$. Thus Γ is a locally truncated diagram geometry belonging to the diagram of Fig. 15.1.

We are now in a position to apply the full theory on locally truncated geometries developed in Chap. 11, where $J = \{\mathcal{P}, \mathcal{L}, \mathcal{M}, \mathcal{A}3, \mathcal{S}\}$ (the round nodes of Fig. 15.1). First of all, by Theorem 11.3.2, p. 427, a sheaf $\mathcal{F}$ exists. Second, the sheaf attains only the values of a polar or projective space (both buildings) at any non-empty flag. That means $\mathcal{F}$ is a strongly chamber-connected sheaf. Moreover, the diagram D is type M and possesses no rank three residue of type C_3. It follows that every rank three residue of $C(\mathcal{F})$ is covered by a building. By Theorem 11.7.1, Γ is the truncation of a homomorphic image of a building. Of course, in order that the image be a parapolar space, this homomorphism is an isomorphism if the point-diameter is less than five. Thus we have the following.

Theorem 15.4.2 $\Gamma = (\mathcal{P}, \mathcal{L})$ *is a homomorphic image of a half-spin geometry of type* $D_{n,n}$, $n > 4$. *This homomorphism is an isomorphism if* $n \leq 9$.

15.4.3 The Case r = 5

Here our strong parapolar space Γ satisfies condition $(CC)_3$, has all singular subspaces of finite rank, and has constant symplectic rank five. Then for any point p, the residual geometry is a half-spin geometry of diameter 2 and so must be $D_{4,4}$ or $D_{5,5}$. The former case is a polar space, which has been excluded in Lemma 15.4.1. In the latter case, every convex subspace of $\mathrm{Res}(p)$ is represented by an actual convex subspace of Γ on p, assembled so that Γ belongs to the diagram E_6 as given in Fig. 15.2. We thus have a rank six geometry Γ^* belonging to the diagram E_6 whose truncation to points and lines is the original parapolar space $\Gamma = (\mathcal{P}, \mathcal{L})$.[7]

Now by the local approach theorem of J. Tits [139] (see Theorem 10.6.1, p. 412), after applying the appropriate functors there is a surjective morphism $f : \Delta \to \Gamma^*$

[7] In the language of Pasini's book [95], Γ^* is called an *enrichment* of Γ.

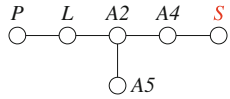

Fig. 15.2 The E_6 diagram of Γ, and the types of subspaces representing its nodes. As usual, $\mathcal{P}$, $\mathcal{L}$, $\mathcal{A}2$, and $\mathcal{S}$ denote the points, lines, planes, and symplecta of Γ. The symbol $\mathcal{A}i$ denotes a class of maximal singular subspaces of Γ which are PG(i,F)s, where F is a field, and $i = 4, 5$. Each symplecton is type D_5 with one of its two classes of maximal singular subspaces belonging to $\mathcal{A}4$; the members of the other class are hyperplanes of members of $\mathcal{A}5$

from the building geometry $\Delta = E_{6,1}(F)$ to Γ (keeping the typset as named in Fig. 15.2). But with respect to the truncation to points and lines, this morphism becomes a morphism of parapolar spaces,

$$f_{\mathcal{PL}} : \Delta_{\{\mathcal{P},\mathcal{L}\}} \to \Gamma,$$

and all of the theory of Sect. 11.8, p. 434 now applies. It follows that, for any symplecton $\hat{S}$ of the building geometry Δ, $S := f(\hat{S}) \in \mathcal{S}$ (that is, its image is a symplecton S of Γ), and

$$f|_{\hat{S}} : \hat{S} \to S$$

is an isomorphism. But the geometry Δ is a strong parapolar space with point-diameter 2, so any two of its points lie together in some symplecton. It follows that the mapping f is bijective when restricted to points. Since all other objects of types $\mathcal{L}$, Π, $\mathcal{A}4$, $\mathcal{A}5$, $\mathcal{S}$ are subspaces of their respective geometries, we see that $f_{\mathcal{PL}}$ is an isomorphism. Hence the following applies.

Theorem 15.4.3 *If Γ is a strong parapolar space all of whose symplecta have polar rank five, and if Γ satisfies the Cohen–Cooperstein condition $(CC)_3$, then Γ is either a polar space of type D_5 or is isomorphic to a Lie incidence geometry of type $E_{6,1}$ for some field F.*

Remark Note that although the hypotheses involve only $(CC)_3$ (that is, $x^\perp \cap S = PG(4)$, if it contains a plane), in fact we have the much stronger condition $(CC)_0$ ($x^\perp \cap S = \emptyset$ or PG(4)) in the conclusion.

15.4.4 The Case r = 6

Once again, Γ is a strong parapolar space. It has constant symplectic rank 6, and for any non-incident point-symplecton pair (x, S), one has $x^\perp \cap S \simeq$ PG(5), if it contains a PG(3). As usual all singular subspaces have some finite rank, and Γ is assumed not to be a polar space.

From the previous Theorem 15.4.3, any point-residual $\mathrm{Res}(p)$ is the $E_{6,1}(F)$-geometry, where F is a field, with all objects of its enrichment induced from

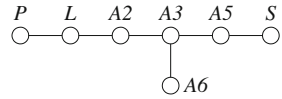

Fig. 15.3 The diagram associated with the diagram geometry formed by subspaces of Γ in the case $r = 6$. Here $\mathcal{P}$, $\mathcal{L}$, $\mathcal{A}2$, and $\mathcal{S}$ are points, lines, planes, and symplecta, while for $i = 3, 5, 6$, $\mathcal{A}i$ are PG(i,F)s, the latter two being maximal singular subspaces. The symplecta are oriflame of type $D_6(F)$. One class of maximal singular subspaces of any symplecton S belong to $\mathcal{A}5$, while each member of the other class is a hyperplane of a unique member of $\mathcal{A}6$. The PG(3)s belonging to $\mathcal{A}6$ lie in members of $\mathcal{A}5$ as well as $\mathcal{A}6$

point-residuals of subspaces of Γ. Thus Γ is enriched to yield a geometry Γ^* belonging to the diagram of Fig. 15.3.

Again, by Tits' "local approach" theorem there is a surjective morphism

$$f : \Delta \to \Gamma^*$$

from the geometry Δ of a building of type $E_7(F)$, which, under truncation to points and lines, induces a morphism of strong parapolar spaces:

$$f_{\mathcal{PL}} : \Delta_{\{\mathcal{P},\mathcal{L}\}} \to \Gamma.$$

Furthermore, this homomorphism induces an isomorphism when restricted to singular subspaces or to symplecta of Γ.

Now the geometry $\Delta_{\mathcal{PL}}$ possesses an interesting property as follows.

(E7) *Given any non-incident point-symplecton pair* (x, S), $x^\perp \cap S$ *is always non-empty.*

Now suppose x and y are two distinct points (objects of type $\mathcal{P}$) of the building geometry Δ for which $f(x) = f(y)$. Clearly x and y cannot be collinear or lie in a common symplecta since f is point-injective when restricted to singular spaces and symplecta. So we may assume $d(x, y) = 3$, where d is the distance metric in the point-collinearity graph $(\mathcal{P}_\Delta, \sim)$ of $\Delta_{\{\mathcal{P},\mathcal{L}\}}$. Now there is a symplecton $\hat{S}$ of Δ containing y and by (E7) above, we have $x^\perp \cap \hat{S} = \{z\}$, a single point. But now the image $S := f(\hat{S})$ is a symplecton of Γ and $f|_{\hat{S}}$ induces an isomorphism (see Sect. 11.8, p. 434). Thus $f(z)$ and $f(y)$ are not collinear in S since z and y are not collinear in $\hat{S}$ (a consequence of $d(x, y) = 3$). But that is impossible since $x \in z^\perp$ forces $f(x) = f(y)$ and $f(z)$ to be collinear.

We conclude that f is bijective on points and, since f is surjective on lines, the mapping f induced on the point-line truncation, $f_{\mathcal{PL}} : \Delta_{\{\mathcal{P},\mathcal{L}\}} \to \Gamma$, is an isomorphism of point-line geometries. Thus the following applies.

Theorem 15.4.4 *If the strong parapolar space* Γ *has constant symplectic rank six, has all singular subspaces of finite rank, and satisfies condition* $(CC)_4$*, then* Γ *is isomorphic to a Lie incidence geometry* $E_{7,7}(F)$ *for some field* F.

Remark Actually $E_{7,7}$ satisfies the stronger $(CC)_1$ as well as (E7). There are geometries with diagram E_7 which are not buildings. For these, the truncation to points and lines at nodes 7 and 6 of the diagram (the positions of the symbols "$\mathcal{P}$" and "$\mathcal{L}$" in Fig. 15.3) cannot yield a strong parapolar space.

15.4.5 The Case $r > 6$

Suppose now that there is a strong parapolar space which is not a polar space with symplectic rank seven satisfying $(CC)_5$ and having all singular subspaces of finite rank. Then by Theorem 15.4.4 a point-residual $\mathrm{Res}(p)$ must be $E_{7,1}(F)$. But the latter geometry has point-diameter three and so cannot be a point-residual of a strong parapolar space.

Of course for any $r > 6$, the case of symplectic rank r and $(CC)_{r-3}$ reduces to the case $r = 7$ of the previous paragraph by considering residuals of larger singular subspaces and repeatedly applying Lemma 15.4.1. Formalizing this argument is left as an exercise for the student.

15.4.6 The First Cohen–Cooperstein Theorem

It is time to sum up what has been proved so far.

Theorem 15.4.5 (Cohen–Cooperstein, updated.) *Let Γ be a locally connected parapolar space all of whose symplecta possess a constant symplectic rank $r \geq 3$. If this rank exceeds three it is also assumed that Γ is a strong parapolar space of which all singular subspaces of Γ possess finite rank.*

Assume these hypotheses:

1. *Γ is not itself a polar space.*
2. (Condition $(CC)_{r-2}$.) *If, for any non-incident point-symplecton pair (x, S), the intersection $x^\perp \cap S$ has projective rank at least $r - 1$ (that is, it contains a $PG(r-2)$), then in fact it has rank r (i.e., it is a maximal singular subspace of S). (In other words, $x^\perp \cap S$ cannot be a hyperplane of a maximal singular subspace of S.)*

Then in fact one of the following conclusions must hold:

1. *If $r = 3$ then Γ is either:*
   (a) *The Grassmannian of k-spaces of a (possibly infinite-dimensional) vector space V over a division ring D, or*
   (b) *The quotient $A_{2n-1,n}(D)/\langle\sigma\rangle$, where σ is a polarity of V of Witt index at most $n - 5$.*
2. *If $r = 4$, $\Gamma = (\mathcal{P}, \mathcal{L})$ is a homomorphic image of a half-spin geometry of type $D_{n,n}$ over a field F. This homomorphism is an isomorphism if $n \leq 9$.*

3. *If $r = 5$, then Γ is the Lie incidence geometry $E_{6,1}(F)$.*
4. *If $r = 6$, then Γ is the Lie incidence geometry $E_{7,7}(F)$ (in the Bourbaki node-numbering scheme).*

(Under no circumstances can r exceed 6.)

Proof The case $r = 3$ is Cohen's theorem (Theorem 15.3.7). The remaining cases review the conclusions of the previous four sections. □

15.5 The Uniformizing Condition

15.5.1 Introduction

This short section concerns a somewhat more natural condition which implies $(\mathrm{CC})_k$.

Throughout this section, Γ is a locally connected strong parapolar space whose symplecta have finite polar rank at least three. (Of course the polar rank may otherwise vary from symplecta to symplecta. We do not assume it is a strong parapolar space except where that condition is specified in a particular result.)

15.5.2 A Basic Lemma

Everything depends upon the following elementary lemma.

Lemma 15.5.1 *Suppose A and B are two distinct symplecta of a parapolar space Γ such that the intersection $A \cap B$ is non-empty and is a proper subspace of a maximal singular subspace of B. In case $A \cap B$ contains a unique point, we also assume that Γ is a strong parapolar space.*

If $x \in A - (A \cap B)^\perp$, and $x^\perp \cap B$ is not contained in $A \cap B$, then there exists a symplecton R containing $\{x\}$ and $x^\perp \cap B$ such that $x \notin (R \cap B)^\perp$ and

$$\dim(A \cap B) = \dim(x^\perp \cap B) = \dim(R \cap B) - 1.$$

(The dimensions are dimensions as projective spaces.[8]*)*

Proof Let (x, A, B) satisfy the hypotheses. We can then choose a point $z \in (x^\perp \cap B) - A$. Clearly, $H_0 := x^\perp \cap (A \cap B)$ is a hyperplane of $A \cap B$. For every point $s \in (A \cap B) - H_0$, we have s not collinear with x and $x^\perp \cap s^\perp \subseteq A$, by convexity of the symplecton A. Since z is not in A, z cannot be collinear with s otherwise it

[8] The singular subspaces of symplecta are always projective spaces. If some symplecta are generalized quadrangles, it is possible that maximal singular subspaces of Γ are not projective spaces, but still, those lying in a symplecton are.

would lie in $s^\perp \cap x^\perp$. Since $z^\perp \cap (A \cap B)$ is a hyperplane of $A \cap B$, containing no element s of $(A \cap B) - H_0$, we have

$$z^\perp \cap (A \cap B) = H_0 = x^\perp \cap (A \cap B). \tag{15.2}$$

Now by hypothesis, $A \cap B$ is a proper subspace of a maximal singular subspace M of B. Since z is not collinear with any $s \in (A \cap B) - H_0$, we see that $H := z^\perp \cap M$ is a hyperplane of M. Now $x^\perp \cap B$ is a singular subspace of B which does not contain any point $s \in (A \cap B) - H_0$. Thus, for such a point s, $s^\perp \cap (x^\perp \cap B)$ is a hyperplane of $x^\perp \cap B$. On the other hand, this hyperplane lies in $x^\perp \cap s^\perp \subseteq A$, so $s^\perp \cap (x^\perp \cap B) = H_0$. Thus $x^\perp \cap B = \langle H_0, z \rangle$, and

$$\dim(A \cap B) = \dim(x^\perp \cap B). \tag{15.3}$$

Set $H := z^\perp \cap M$, a hyperplane of M which intersects $A \cap B$ at H_0. Since M properly contains $A \cap B$, H properly contains H_0. Now if one had $H - H_0 \subseteq (x^\perp \cap B)$, then one would have $H = \langle H - H_0 \rangle_H = H$ contained in $x^\perp$. That is impossible since this would force $H \subseteq x^\perp \cap s^\perp \subseteq (A \cap B) \cap x^\perp = H_0$. Thus, there is a point $t \in (H - H_0)$ which is not collinear with x. Now, $x^\perp \cap t^\perp$ contains $H_0 \cup \{z\}$. Thus, if $H_0 \neq \emptyset$, (x, t) is a polar pair. But if $H_0 = \emptyset$, $A \cap B$ consists of a single point, and by hypothesis (x, t) is again a polar pair. Let R be the unique symplecton containing $\{x, t\}$. Then $R \cap B$ is a singular space containing $\{H_0, t, z\}$. Since x is not collinear with t, we see that $x \notin (R \cap B)^\perp$, as required.

Again, for any $s \in (A \cap B) - H_0$, $s^\perp \cap (R \cap B)$ is a hyperplane of $R \cap B$ containing $\{H_0, t\}$. It, in turn, contains a hyperplane $x^\perp \cap s^\perp \cap (R \cap B)$ which must lie in $s^\perp \cap x^\perp \subseteq x^\perp \cap (A \cap B) = H_0$. Thus we have a chain of singular subspaces:

$$H_0 \subset \langle H_0, t \rangle \subset R \cap B,$$

each subspace being a hyperplane of its successor. The relations among dimensions at the end of the statement of the lemma is now proved. Note that $R \cap B$ contains $\langle H_0, z \rangle = x^\perp \cap B$. □

15.5.3 The Principle

Consider the following condition on a parapolar space Γ:

(U)$_k$ (The Uniformizing Principle.) *Whenever A and B are two symplecta of the parapolar space Γ which intersect in a projective subspace of projective dimension at least k – that is, $A \cap B$ contains a $PG(k)$ – then, for each point x in $A - (A \cap B)^\perp$, the set $x^\perp \cap B$ is not contained in $A \cap B$.*

Obviously if $j \geq k$, then (U)$_k$ implies (U)$_j$.

Theorem 15.5.2 *Let Γ be a parapolar space satisfying condition $(U)_k$, $k \geq 0$. Then the following hold:*

1. *If $k = 0$, assume Γ is a strong parapolar space. Whenever the intersection of two symplecta has projective dimension at least k, the intersection is in fact a maximal singular subspace of each of the intersecting symplecta. In particular, the two symplecta possess the same rank.*
2. *Now assume $k \geq 1$ and that Γ is a strong parapolar space if $k = 1$. Then the condition $(CC)_{k-1}$ holds.*

Proof Part 1. We argue by way of contradiction. Suppose A and B are two distinct symplecta such that $A \cap B$ has projective dimension at least k, but is not a maximal singular subspace of B. Select a point $x \in A - (A \cap B)^\perp$. By condition $(U)_k$, $x^\perp \cap B$ is not contained in $A \cap B$. Then by Lemma 15.5.1, there is a symplecton R containing $\{x, x^\perp \cap B\}$ with $\dim(R \cap B) = \dim(A \cap B) + 1$. Thus we can apply condition $(U)_k$ to the pair of symplecta (R, B). Note that $x \notin (R \cap B)^\perp$. So, by Lemma 15.5.1, we must have $x^\perp \cap B \not\subseteq B$. But that contradicts the fact that $x^\perp \cap B \subseteq R$.

Thus it is false that $A \cap B$ is not a maximal singular subspace of B. By the symmetry of the hypotheses as applied to A or to B, it follows that $A \cap B$ is also a maximal singular subspace of A. Thus A and B are symplecta that possess the same polar rank.

Part 2. Suppose B is a symplecton of Γ and x is a point not in symplecton B such that $x^\perp \cap B$ has projective rank at least k. We must show that $x^\perp \cap B$ is a maximal singular subspace of B.

So, assuming otherwise, there is a maximal singular subspace M of B which properly contains $x^\perp \cap B$. Choose $t \in M - (x^\perp \cap B)$. Either $x^\perp \cap t^\perp$ contains a line, or $x^\perp \cap B$ is a single point, and $k = 1$, forcing Γ to be a strong parapolar space. Thus, in either case, there is a symplecton R containing $\{x, t\}$ and $R \cap B$ contains $\langle t, B \cap x^\perp \rangle$, having projective dimension at least k. Since $x \notin (R \cap B)^\perp$, by condition $(U)_k$, $x^\perp \cap B$ is not contained in R. That is a contradiction since R *does* contain $x^\perp \cap B$. Thus $x^\perp \cap B$ is a maximal singular subspace of B. □

15.5.4 The Uniformizing Principle and Constant Symplectic Rank

Everything is contained in the following theorem.

Theorem 15.5.3 *Assume Γ is a locally connected parapolar space of symplectic rank at least three, satisfying axiom $(U)_2$. Then Γ possesses a constant symplectic rank r.*

Remark So far, in this section, no result requires symplecta to have finite rank. So the constant symplectic rank asserted may be an infinite cardinal.

Proof Let A and B be two arbitrary symplecta. We must show that A and B possess the same rank as polar spaces. The proof proceeds by a series of short steps.

Step 1. *If $A \cap B$ contains a plane, A and B have the same rank.*

Since Γ is a strong parapolar space, satisfying $(U)_2$, part 1 of Theorem 15.5.2 forces $A \cap B$ to be a maximal singular space of both A and B. Thus A and B possess the same polar rank.

Step 2. *If $A \cap B$ contains a line, then A and B have the same rank as polar spaces.*

We may assume that $A \cap B$ is exactly a line L. Now since both A and B have polar rank at least three, L lies in planes π_A and π_B in A and B respectively, chosen so that $\pi_A \cup \pi_B$ does not lie in a common singular subspace. Then the convex closure $\langle\langle \pi_A, \pi_B \rangle\rangle$ is a symplecton R whose intersection with each of A and B contains a plane. By two applications of Step 1, $\text{rank}(A) = \text{rank}(R) = \text{rank}(B)$, which gives the conclusion of this step.

Step 3. *If $A \cap B = \{p\}$, a single point, then $rank(A) = rank(B)$.*

Here one can find lines L and M on p, belonging to A and B respectively, such that $L \not\subseteq M^\perp$. (If this were not true one would have $L^\perp \cap B$, a clique containing all of $p^\perp \cap B$, which is not a clique!)

Suppose, by way of argument, that L and M are contained in a common symplecton T. Then $T \cap S$ and $T \cap B$ each contain lines, so by Step 1,

$$rank(A) = rank(T) = rank(B).$$

So we may assume that L and M do not lie in a common symplecton. Since Γ is locally connected, there exists a minimal sequence of lines on p:

$$L = L_0, L_1, \ldots, L_n = M,$$

such that $\langle L_i, L_{i+1} \rangle$ is a plane π_i, $i = 1, \ldots, n-1$.

By minimality of n, the pair (L_i, L_{i+2}) share a symplecton S_i, $i = 0, \ldots, n-2$. Since $A \cap S_0 \supseteq L$, $B \cap S_{n-2} \supseteq M$, and $S_i \cap S_{i+1}$ share the plane $\pi_{i+1} = \langle L_{i+1}, L_{i+2} \rangle$, $i = 1, \ldots, n-2$, we have

$$rank(A) = rankS_0 = \cdots = rank(S_{n-2}) = rank(B),$$

by Steps 1 and 2. Thus $rank(A) = rank(B)$.

Step 4. *In all cases A and B have the same polar rank.*

The only case left is that $A \cap B = \emptyset$. But as Γ is a connected geometry, there is a path $P := (a = a_0, \ldots, a_m = b)$ in the point-collinearity graph where a is in A and b is in B. The parapolar space hypothesis insures the existence of a symplecton T_i on line $a_i a_{i+1}$, $i = 0, \ldots, m-1$. Thus $A \cap T_0$ contains the point $a_0 = a$, $B \cap T_{m-1}$ contains the point $a_m = b$, and $T_i \cap T_{i+1}$ contains the point a_{i+1} for $i = 0, \ldots, m-2$. By Step 3 applied to consecutive symplecton pairs along this path, one obtains

$$\text{rank}(A) = \text{rank}(T_0) = \cdots = \text{rank}(T_{m-2}) = \text{rank}(B).$$

The conclusion follows. □

We now have the following.

Theorem 15.5.4 *Suppose* Γ *is a locally connected parapolar space each of whose symplecta have a finite polar rank at least three and having at least one maximal singular subspace of finite rank. If* Γ *contains a symplecton of polar rank at least four, we further assume that* Γ *is a strong parapolar space with all singular subspaces of finite rank.*

If Γ *satisfies the hypotheses* $(U)_2$*, then* Γ *is one of the "Cohen-Cooperstein geometries" – that is, a polar space, a Grassmannian, a quotient of a Grassmannian* $A_{2n-1,n}(D)$ *by a polarity of index at most* $n-5$*, an appropriate homomorphic image of a half-spin geometry, or one of the exceptional Lie incidence geometries of types* $E_{6,1}$ *or* $E_{7,7}$*.*

Proof This is simply a matter of assembling the hypotheses of the Cohen–Cooperstein theorem (Theorem 15.4.5). But we have more than these hypotheses since $(U)_2$ implies $(CC)_1$ (Theorem 15.5.2) and by Theorem 15.5.3, Γ has constant symplectic rank r which is finite by hypothesis. □

15.6 Exercises for Chapter 15

15.6.1 General Features

15.1 This exercise concerns the hypotheses of the first Cohen-Cooperstein Theorem (see Theorem 15.4.5).

1. Suppose Γ is a (not necessarily strong) parapolar space satisfying the following Cohen–Cooperstein hypothesis:

   (Co2) If x is a point not incident with a symplecton S, then $x^\perp \cap S$ is either empty, consists of a single point, or is a maximal singular subspace of S.

   Show that if Γ has symplectic rank at least three and is locally connected, then Γ is a strong parapolar space. [Hint: It suffices to show that any point-residual $\mathrm{Res}_\Gamma(p)$ has point-diameter at most two.]
2. Conclude that if the parapolar space Γ satisfies (Co1), has symplectic rank at least three, and has all its singular subspaces of finite projective rank, then every locally connected component is of Cohen–Cooperstein type.
3. Show that the product geometry $D_{7,7} \times E_{7,7}$ satisfies the hypothesis of part 2.

15.6.2 Understanding the Exceptional Strong Parapolar Spaces

Existence Questions

At this point the student has suddenly encountered geometries which are specific point-line truncations of geometries of exceptional Lie types E_6 and E_7 that happen

to be strong parapolar spaces. We have not offered any proof that such geometries even exist. Nor did the original papers of Cohen and Cooperstein, Hanssens, and many others. They did not do so because they were writing to a sophisticated community that knew this answer: thick building geometries of every spherical type or rank at least three exist over any division ring or field that permits it. Of course, any appeal to algebraic groups on this question depends upon the existence of the required algebraic groups. There are two answers to the question – the long older one and the short one, neither one of which is proved in this book.

The long answer says that these geometries exist because Chevalley groups exist. In turn they exist because the simple complex Lie Algebras exist (defined by structure constants on a basis) and Chevalley's Z-form on a suitable basis.[9] Of course more is desired. One would like the matrices $exp(ad(x))$ to have integer entries, so that the so-called "adjoint representation" of the associated Lie group (the linear group generated by the $exp(ad(x))$s, as x ranges over the non-zero roots) becomes an integral representation. Indeed Kostant showed that many complex modules possess Z-forms, so that standard complex representations of these modules can be recast as integral representations.

Why Z? Well any n-by-n matrix over Z can be viewed as a matrix in $GL(n, K)$, for any K and they generate there, what we call a *Chevalley group*. From the construction we get (1) versions of simple complex Lie Algebras over arbitrary fields, (2) very rich groups of automorphisms defined over that same field, as well as (3) modules for these new structures via Kostant's theory.

Finally, we can describe the desired "building geometries" as coset geometries of the Chevalley groups, and on those occasions we wish to think of them as geometries embedded in a projective space, we can pull in Kostant's theory.

There are two things wrong with this picture:

1. From the point of view of this book, the development would be deflected into a two- or three-chapter side-show at a minimum covering the classification of the simple complex Lie Algebras, developing the Chevalley basis, etc.
2. Is it exhaustive? Giving an example of each type is nice. But is that all examples? (It turns out that it does in rank at least four, but to explain that would require still more side-shows.)

The beautiful short answer is this. Given a certain sort of sufficient datum called "blueprints," there is a direct way to show that a building geometry exists. That datum already contains the local information obtained from residues, so it gives an exhaustive existence answer. What I have described is fully outlined in a paper entitled "Building buildings" by Mark Ronan and Jacques Tits [104].

Although I appeal to the beginning student to accept the existence of these geometries, that act of faith (like many acts of faith) does not necessarily increase one's understanding of these supposed objects. For this reason, the following

[9] In simple language this means there is a basis of the algebra such that products of basis elements are integer-linear combinations of the other basis elements.

exercises are intended to give the student a "feel" for how the parapolar spaces of type $E_{6,1}$ and $E_{7,7}$ really work.

Axiomatizing the Exceptional Strong Parapolar Spaces

To begin with, we must describe the geometries E_6, and E_7 in exactly the way they were encountered in the theorems of this chapter – as strong parapolar spaces whose enrichment to singular subspaces and symplecta are higher rank diagram geometries belonging to the diagrams of these types.

Thus we have the following two hypotheses:

(E6) Γ is a strong parapolar space with these properties:

- (E6-1) Any point-residual $\mathrm{Res}_\Gamma(p)$ is a half-spin geometry of type $D_{5,5}$.
- (E6-2) All symplecta are type $D_{5,1}$ – that is, a classical polar space of totally singular 1- and 2-subspaces of a quadratic form (V, Q) of type $O^+(10, F)$ for some field F.

(E7) Γ is a strong parapolar space with these properties:

- (E7-1) Any point-residual $\mathrm{Res}_\Gamma(p)$ is a point-line geometry (of lines and planes incident with p) satisying all of the hypotheses (E6).
- (E7-2) All symplecta are polar spaces of type $D_{6,1}$ – that is, a classical polar space of totally singular 1- and 2-subspaces of a quadratic form (V, Q) of type $O^+(12, F)$ for some field F.

Consequences of the (E6) Axioms

In the following exercises we assume Γ is a strong parapolar space satisfying the axioms of (E6).

15.2 Suppose p is a point and S is a symplecton. Prove that $p^\perp \cap S$ is either the empty set or a maximal singular subspace of S. [Hint: Read this off from $\mathrm{Res}_\Gamma(p)$.]

15.3 Show that if p is a point, then there exists a symplecton S such that $p^\perp \cap S = \emptyset$. [Hint: Suppose false and imagine what consequences that entails for $\mathrm{Res}_\Gamma(p)$.]

15.4 Show that the point-collinearity graph $\Delta = (\mathcal{P}, \sim)$ has diameter at most two. [Hint: Suppose you had a geodesic path (a, b, c, d) of length three, form the unique symplecton S on $\{a, c\}$, and use the previous item to discuss $c^\perp \cap S$.]

15.5 At this point we know that every pair of non-adjacent points lie in a unique symplecton of type $D_{5,1}$, so that there are only two relationships between points. In this exercise we ask the student to show that if S_1 and S_2 are distinct symplecta, then $S_1 \cap S_2$ is either a maximal singular subspace of Γ, or else the intersection is a unique point. [Study $\mathrm{Res}_\Gamma(p)$. Then address the possibility that $S_1 \cap S_2 = \emptyset$ using the previous items.]

Consequences of the (E7) Axioms

In the following exercises we assume Γ is a strong parapolar space satisfying the axioms of (E7).

15.6 Produce an argument showing that (E7-1) alone implies that Γ is a strong parapolar space.

15.7 Show that if S_1 and S_2 are two distinct symplecta, then $S_1 \cap S_2$ is either the empty set, a line, or a maximal singular subspace of each of the S_i.

15.8 Conclude the following:

1. If p is a point and S is a symplecton such that $p^\perp \cap S = \{x\}$, a single point, then every point of $S - x^\perp$ is distance three from point p. (In terms of the point-collinearity graph Δ and the convention that subspaces are defined as sets of points with certain properties, this assertion can be stated as follows: *the symplecton S is gated with respect to the point p with gate x.*)
2. The point-collinearity graph Δ has diameter at most three.
3. One has the following condition (D2). *For any point p, the collection $\Delta_2^*(p)$ of all points of distance at most two from p form a subspace of Γ.*
4. For every line L, and every point p, one has

$$\Delta_2^*(p) \cap L \neq \emptyset.$$

It then follows that $\Delta_2^*(p)$ is always a geometric hyperplane of Γ.

[Hint: Simply exploit Exercises 13.23–13.25 of Chap. 13.]

15.9 Suppose that x and y are two points at distance three in the geometry $E_{7,7}$. Show that there is a one-to-one correspondence

$$\text{lines on } x \rightarrow \text{ symplecta on } y.$$

[Hint: Simply use properties proved in the previous exercise.]

15.6.3 Involutions Acting on Classical Strong Parapolar Spaces

Throughout the next group of exercises Γ is a parapolar space satisfying the axioms (E6) or (E7). With a minor abuse of notation we will also let Γ refer to the enrichment of such a geometry to its singular subspaces and symplecta — geometries that belong to the E_6 or E_7 diagrams. At no point do we assume these are building geometries.[10]

[10] We say this despite the fact that they really are building geometries in the sense that the associated chamber systems $\mathbf{C}(\Gamma)$ are buildings as defined in Chap. 9.

For any diagram geometry Γ, an *involution* is simply a type-preserving automorphism of Γ of order two.

15.10 We begin with a simple case. Prove the following.

Theorem 15.6.1 *Suppose t is an involution of a geometry of type A_n, where either $n \geq 3$ or, when $n = 2$, Γ is a projective plane derived from a vector space.*[11] *If n is even, then t fixes a point. Moreover, if n is odd, then t fixes at least a line.*

[Hint: The projective spaces in question are in fact modeled by the 1- and 2-subspaces of a vector space V. The involution t is derived from exactly two sorts of semilinear transformations $\tau : V \to V$: (1) elements of order two in $\Gamma L(V)$ which are not $-I$, — the scalar transformation which takes each vector v to $-v$, or (2) a semilinear transformation of V whose square is a scalar transformation $\alpha \cdot I$. (This has to be proved.) The rest is just straightforward (semi)-linear algebra.]

15.11 In order to explain what will appear next, we must first define the concept of an *ovoid of the polar space* Γ. An ovoid is a set $\mathcal{O}$ of pairwise non-collinear points with the extremal property that, for every maximal singular subspace M of the polar space Γ, the set $\mathcal{O} \cap M$ is always non-empty (of course by definition the intersection can contain at most one point). In this exercise the student will prove the following.

Theorem 15.6.2 *Suppose Γ is a polar space of type $D_{n,1}$ where n is an odd integer at least three. Suppose that t is an involution of Γ which stabilizes no plane. Then the subset of fixed points, $\mathcal{P}_t$, forms an ovoid of Γ.*

[Hint: (1) First handle the case $n = 3$. Here Γ is the Grassmannian of lines of a $PG(3, D)$, and t is an involution of the $PG(3, D)$ fixing no projective point or projective plane of the $PG(3, D)$. Use Theorem 15.6.1 to show that the stabilized projective lines partition the projective points. (This is called a *line spread* and corresponds to an ovoid of the oriflame polar space Γ of rank three.)

(2) Assume $n > 3$.

A maximal singular subspace M is a $PG(n - 1, F)$. Then $M \cap M^t \simeq PG(k, F)$ for some even $k \leq n$, and is t-invariant, so, by the previous part of this exercise, possesses a fixed point. Next suppose t stabilized a line L of Γ. Then the involution t acts on the geometry $(\Pi_L, \mathcal{A}_L)$ of those planes and $PG(3)$s of the polar space which are incident with line L. The latter is a polar space $res(L)$ of type $D_{n-1,1}$ and must have a fixed residual point as just argued for Γ. That residual point would then correspond to a plane of

[11] In Chaps. 4 and 5 we called such planes "Desarguesian" (not the usual definition).

Γ stabilized by t, which was forbidden by hypothesis. Since t can stabilize no line of the polar space, conclude that for each maximal singular subspace, M, one has $M \cap M^t = \{p\}$, where p is a t-invariant point.]

15.12 In this exercise the student will prove the following theorem.

Theorem 15.6.3 *Assume that $\Gamma = (\mathcal{P}, \mathcal{L})$ is a polar space of type D_6, and that t is an involution of Γ with these properties:*

1. *The set $\mathcal{P}_t$ of t-invariant points is non-nomempty.*
2. *If $p \in \mathcal{P}$ then either $p^t = p$, or p is not collinear with p^t.*
3. *The involution t stabilizes no singular subspace of Γ which is a $PG(3)$.*

Then the collection of t-invariant point and lines forms a thick generalized quadrangle.

Assuming the hypotheses of Theorem 15.6.3, prove the following assertions.

1. If $p \in \mathcal{P}_t$, then the collection $\mathcal{L}_t(p)$ of t-invariant lines which are incident with p forms an ovoid of $\mathrm{Res}_\Gamma(p) = (\mathcal{L}_p, \Pi_p)$ — that is,

   (O) every maximal singular subspace of Γ which is incident with p contains exactly one of the lines of $\mathcal{L}_t(p)$.

   [Hint: The residue space $\mathrm{Res}_\Gamma(p)$ is a polar space of type $D_{5,1}$ admitting the action of the involution t so that none of its planes are stabilized. Apply Theorem 15.6.2 when $n = 5$.]
2. If L is in $\mathcal{L}_t$, the collection of all t-invariant lines, then the point-shadow of L lies in $\mathcal{P}_t$. [Hint: Property 2 of the hypothesis of this exercise.]
3. The geometry $\Gamma_t := (\mathcal{P}_t, \mathcal{L}_t)$, with incidence as in Γ, forms a non-degenerate generalized quadrangle. [Hint: Show the following: (1) that Γ_t is a partial linear space, (2) that if (p, L) is a non-incident pair in $\mathcal{P}_t \times \mathcal{L}_t$, that p cannot lie in $L^\perp$ (use part 1 and (O)), and so $p^\perp \cap L$ is a point of $\mathcal{P}_t$, by the fundamental polar-space axiom, and (3) that it is impossible to have $\mathcal{P}_t \subseteq r^\perp$ for any point $r \in \mathcal{P}_t$.]

15.13 Prove the following.

Theorem 15.6.4 *Suppose Γ is a strong parapolar space satisfying axioms (E6). Suppose t is an involution acting on G, which stabilizes no plane of Γ. Then the subgeometry $(\mathcal{P}_t, \mathcal{S}_t)$ of points and symplecta stabilized by t, is a projective plane under the incidence relation inherited from Γ.*

[Hint: If S is a symplecton, then either $S = S^t$, or $S \cap S^t$ is an invariant $PG(4)$ or an invariant point, and by the previous parts of this exercise, in all cases S contains an invariant point. If t stabilized a line L of Γ, then it would act on

$$\mathrm{Res}_\Gamma(L) \simeq sh_{\mathcal{P}} \times PG(4, F)$$

and so would fix a projective point of the $PG(4, F)$ by Theorem 15.6.1. But that projective point would correspond to a forbidden t-invariant plane of Γ.

So t stabilizes no lines. Conclude that $\mathcal{P}_t$ is a coclique of the collinearity graph of Γ.
If $\mathcal{A}_4$ is the class of maximal singular $PG(4)$-subspaces (each being an intersection of two symplecta), the dual geometry $\Gamma' := (\mathcal{S}, \mathcal{A}_4)$ is also a strong parapolar space satisfying (E6) axioms and admitting t with no planes stabilized. So, dualizing the argument just given, there exist symplecta stabilized by t and any two of them intersect at a fixed point. Conversely, any two fixed points determine a unique t-invariant symplecton.]

15.14 In this exercise the student will prove the following.

Theorem 15.6.5 *Suppose $\Gamma = (\mathcal{P}, \mathcal{L})$ is a strong parapolar space satisfying the axioms (E7). Assume t is an involution of Γ possessing these properties:*

1. *$\mathcal{P}_t \neq \emptyset$ – that is, t-invariant points exist.*
2. *If p is any point of $\mathcal{P}$, then either $p = p^t$ or p^t is not collinear with p.*
3. *The involution t stabilizes no singular $PG(3)$-subspace of the parapolar space Γ.*

Let $\mathcal{S}_t$ be the class of t-invariant symplecta, and, as before, let $\mathcal{L}_t$ be the class of t-invariant lines. Then with respect to the incidence inherited from Γ, the geometry $\Gamma_t := (\mathcal{P}_t, \mathcal{L}_t, \mathcal{S}_t)$ is a dual polar space of rank three.

Assuming the hypotheses of Theorem 15.6.5, prove the following statements:

1. If p is a t-invariant point, then the elements of $\mathcal{L}_t \cup \mathcal{S}_t$ which are incident with p form a projective plane. Moreover, no two lines of $\mathcal{L}_t$ which are incident with p lie in a common singular subspace of Γ. [Hint: Theorem 15.6.4.]
2. It is impossible that any symplecton contains all of the points of $\mathcal{P}_t$. [Hint: Use the preceding part.]
3. Every symplecton stabilized by t contains a point fixed by t. [Hint: Use the preceding part of this exercise and those (E7)-axioms relating to the relation of a point and a non-incident symplecton.]
4. Let S be any t-invariant symplecton. Then its residue in Γ_t is a convex generalized quadrangle. [Hint: Verify hypotheses of Theorem 15.6.3 for S and $t|_S$. Show that the convexity follows from the convexity of the symplecta of Γ.]
5. If (p, L) is a non-incident pair in $\mathcal{P}_t \times \mathcal{L}_t$, then either p is Γ_t-collinear to a single point of L, or else there is a unique point of L at distance two from p in the point-collinearity graph of Γ_t. Thus Γ_t is a near polygon. [Hint: Consider a non-incident point-line pair. By part 1 of this exercise, L is not in $p^\perp$. If $L \cap p^\perp$ or $L \not\subseteq \Delta_2^*(p)$, L possesses a unique point nearest p. This leaves one case: each point of L is distance two from p. Let x_1 and x_2 be two distinct points of L and let S_i be the unique (necessarily t-invariant) symplecta on p and x_i, $i = 1, 2$. Show that $S_1 \cap S_2 = N$ is a t-invariant line on t, and set $\{q_i\} = {x_i}^\perp \cap N$, $i = 1, 2$. Show that

$L \cap p^\perp = \emptyset$ implies $q_1 = q_2$ forcing $\langle q_1, L\rangle_G$ to be a plane containing lines q_1x_1 and q_1x_2 against part 1.]

6. If $S \in \mathcal{S}_t$ and p is a point of $\mathcal{P}_t$ not incident with S, then $p^\perp$ is collinear in Γ_t to a unique point of $\mathcal{P}_t \cap S$.
7. Γ_t is a dual polar space derived from a rank three polar space. [Hint: Use Corollary 8.5.18 of Chap. 8.]

Remark Theorems like those in this section appear in a very fundamental paper of Bernhard Muehlherr [89], but under considerably stronger hypotheses. First, in his paper, these parapolar spaces are assumed to be building geometries (a fact implied by the axioms of (E6) and (E7) via the Cohen–Cooperstein theory, but not actually used in any of the exercises above). Second, Theorem 15.6.3 does not appear, and in his version of Theorem 15.6.5 not only must the generalized quadrangles be hypothesized, but t is assumed not to stabilize any of four types of objects, not just one type.

On the other hand, Muehlherr's paper ranks as a fundamental paper of Lie incidence geometry, for it not only shows that a rank-three polar space emerges from the fixed-point geometry Γ_t, for an appropriate involution t of $E_{7,7}$, he goes on to show that *every non-embeddable rank three polar space arises in this precise way!* This part of his proof is not elementary point-line geometry. One has to know that the polar spaces in question, being Moufang, are characterized by their point-residuals, which are generalized quadrangles constructed from the norm-form of a Cayley–Dixon division algebra. The ability to invent an appropriate involution t that allows the rank-three dual-polar space defined by the Cayley–Dixon–division–algebra–datum to be embedded as Γ_t in some $E_{7,7}$ uses many uniqueness results about extending automorphisms in buildings. So Muehlherr's hypothesis that the geometries in question are building geometries is fully justified by this marvelous result.

Recall from Chap. 8 that the classification of the non-embeddable rank-three polar spaces stood at a level of difficulty beyond what the previous chapters could prepare the student for. In that chapter we could only give a rough sketch of Tits' classification of those objects. So it is rather intriguing to discover that they all arise from an involution acting on a parapolar space that we *can easily* characterize by simple point-line arguments.

Two final remarks:

1. The reader should be aware that involutions with the properties listed in Theorems 15.6.2, 15.6.4, or 15.6.5 do not often exist. They don't exist at all over finite ground fields.
2. The reader is invited to employ these results in a study of interesting fixed-point geometries arising from involutions acting on the long-root geometries of Chap. 17 or even generalized parapolar spaces arising as Lie incidence geometries obtained from affine buildings. (In this connection one notices that, in a more recent paper of B. Muehlherr and H. Van Maldegham [91], a large class of Moufang quadrangles are realized as the fixed-point geometries arising from an appropriate involution acting on a metasymplectic space of type F_4.)

Chapter 16
Characterizing Strong Parapolar Spaces by the Relation Between Points and Certain Maximal Singular Subspaces

Abstract The following hypothesis is investigated. A parapolar space Γ has a class $\mathcal{M}$ of maximal singular subspaces of finite projective rank, covering all lines, and having the property that, for any non-incident pair $(x, M) \in (\mathcal{P}, \mathcal{M})$, $x^\perp \cap M$ is empty or is a $PG(d)$ for a fixed positive integer d. Polar spaces have this property and if $d > 2$, that is the only possibility. Otherwise, if $d = 2$, Γ is a homomorphic image of a half-spin geometry. But when $d = 1$, Γ is a Grassmannian of k-spaces of a possibly infinite-dimensional vector space, or is a space in which each line lies in a unique member of M and has tightly controlled point-residues. No example is known that would realize this last case.

16.1 Introduction

The classic theorems of Cohen and Cooperstein presented in the preceding chapter characterized strong parapolar spaces of Lie type by restricting the possible projective ranks of $x^\perp \cap S$. In this chapter we shall characterize some of the strong parapolar spaces of Lie type by restricting the projective rank of $x^\perp \cap M$, where x is any point and M belongs to a class $\mathcal{M}$ of maximal singular subspaces and M does not contain the point x. Thus there is a superficial similarity with the Cohen–Cooperstein situation with $\mathcal{M}$ replacing the class $\mathcal{S}$ of all symplecta.

But there are some important differences. First, while $\mathcal{S}$ is the class of *all* symplecta, $\mathcal{M}$ need not be the class of *all* maximal singular subspaces. Instead we simply require that every line lies in at least one member of $\mathcal{M}$. Second, in the new theory, one need not make an advance deposit on the strong parapolar polar space hypothesis – rather, in the parapolar space context, that is thrust upon one in the act of replacing $\mathcal{S}$ by $\mathcal{M}$.

Suppose now that we have a class $\mathcal{M}$ of maximal singular subspaces of a parapolar space. Then for each non-incident pair $(x, M) \in \mathcal{P} \times \mathcal{M}$, we suppose the spectrum of possible dimensions of $x^\perp \cap M$ in M to be a collection of finite integers possessing a "gap" just below its maximum value d. That is, the next largest value e of a dimension of $X^\perp \cap M$, where X is a singular subspace not in M, is at most

E.E. Shult, *Points and Lines*, Universitext, DOI 10.1007/978-3-642-15627-4_16,

© Springer-Verlag Berlin Heidelberg 2011

$d-2$. (Recall that if the parapolar space has symplectic rank at least three, M must be a projective space so that we are speaking of projective dimensions here.) Then, when $x^\perp \cap M$ attains the dimension d, there must exist in $x^\perp \cap M$ a subspace U of finite dimension $e+1$ in $x^\perp \cap M$, where $e \leq d-2$. In that case the "residue" of U in the geometry of all singular subspaces containing U becomes a geometry with its own class $\mathcal{M}_U$ of maximal singular subspaces – those members of $\mathcal{M}$ which contain U – except now U has codimension $d-e-1 \geq 1$ in $X^\perp \cap M$ for each singular subspace X in which U is a hyperplane, X is not contained in M, and M is an element of $\mathcal{M}_U$.

In other words, one is permitted to examine a gamma space $\Gamma = (\mathcal{P}', \mathcal{L}')$ (the residue of U), equipped with its own class of maximal singular subspaces $\mathcal{M}'$ with the property that for each non-incident pair $(x, M) \in (\mathcal{P}', \mathcal{M}')$, the singular subspace $x^\perp \cap M$ has a constant finite projective dimension $d \geq 0'$ — a very "clean" and special version of our initial hypothesis about the spectrum of possible dimensions of $x^\perp \cap M$. Assuming the original parapolar space has symplectic rank at least three, we see that the residue Γ enjoying this hypothesis is a strong parapolar space the symplecta of which have rank at least two and, when they are rank two, they are themselves point-residues of polar spaces of rank three – that is, they are classical quadrangles in the sense of Chap. 15.

Hopefully one can classify the parapolar geometries satisfying the "clean" situation, and then use that to approach the case of more general gaps in the spectrum of $\dim(x^\perp \cap M)$.

It is hoped that the remarks of the three previous paragraphs will explain the motivation for the special hypotheses of the following section, and the main characterizing theorem (Theorem 16.2.1) the proof of which is developed in the next four sections. This basic theorem gives rise to further corollaries characterizing two subsequent generations of parapolar spaces whose point-residuals belong to the previous generation (Corollaries 16.6.1 and 16.7.1).

16.2 Parapolar Spaces with the "Clean Hypothesis"

Our "clean" hypothesis reads as follows.

Hypothesis (A)

(A1) $\Gamma = (\mathcal{P}, \mathcal{L})$ is a strong parapolar space. All symplecta of rank two are classical.
(A2) $\mathcal{M}$ is a collection of maximal singular subspaces of Γ, and d is a non-negative integer such that

1. Every point lies in a member of $\mathcal{M}$.
2. For each non-incident pair $(x, M) \in \mathcal{P} \times \mathcal{M}$, we have $x^\perp \cap M \simeq PG(d, D)$ for some division ring D.

3. If $d = 0$, it is assumed that the elements of $\mathcal{M}$ are projective spaces (members of $\mathcal{M}$ are allowed to possess infinite projective rank).

Our goal in the next four sections is to prove the following theorem.

Theorem 16.2.1 *Suppose $\Gamma = (\mathcal{P}, \mathcal{L})$ is a point-line geometry satisfying hypothesis (A) above. Then one of the following alternatives holds:*

1. *Γ is a non-degenerate polar space of finite rank at least two.*
2. *$d = 0$, Γ has symplectic rank two, and one of the following two possibilities holds:*
   (a) *Γ is the product geometry $A \times B$ of two maximal singular subspaces A and B.*
   (b) *Some symplecton is not a grid, and the members of $\mathcal{M}$ partition the points. Any line which is not contained in an element of $\mathcal{M}$ is a maximal singular subspace. No line intersects all members of $\mathcal{M}$.*
3. *The parameter $d = 1$ and Γ is the Grassmannian of lines of some projective space (of possibly infinite projective rank).*

16.2.1 Elementary Consequences of the "Clean Hypothesis"

Lemma 16.2.2 *If S is a symplecton and $M \in \mathcal{M}$, then $S \cap M = \emptyset$ or is a projective space of dimension $d + 1$. In the latter case, the intersection is a maximal singular subspace of S.*

Proof Suppose p is a point in $S \cap M$. Choose a point $x \in S - p^\perp$. Then $x^\perp \cap M$ is a $PG(d)$ in $p^\perp \cap x^\perp \subseteq S$, but does not contain p. Thus $M_0 := \langle p, x^\perp \cap M \rangle$ is a $PG(d+1)$ lying in $M \cap S$. If $M \cap S$ had projective dimension larger than $d + 1$ then for any point $z \in S - M$, $z^\perp \cap M$ would have dimension larger than d against the last part of (A2). Thus $M_0 = M \cap S \simeq PG(d + 1)$.

If M_0 were not a maximal singular subspace of S, then S would contain a point $y \in M_0{}^\perp - M_0$ and $y^\perp \cap M$ would have projective dimension $d + 1$, also against the last part of (A2). Thus M_0 is a maximal singular subspace of S. The result follows. □

Corollary 16.2.3 *Every symplecton of Γ has polar rank $d + 2$.*

Proof Let S be any symplecton. By the first part of (A2), a point of S lies in some member M of $\mathcal{M}$. By Lemma 16.2.2, $M \cap S$ is a $PG(d + 1)$ which is a maximal singular subspace of S. The conclusion follows. □

Lemma 16.2.4 *Fix $M \in \mathcal{M}$. For every subspace U of M which is a $PG(d)$, there exists a point x not in M such that $x^\perp \cap M = U$.*

Proof Let $\mathcal{U}$ denote the collection of all subspaces of M which are of the form $x^\perp \cap M$ for some point $x \in \mathcal{P} - M$. Then $\mathcal{U}$ is a collection of proper subspaces of M, all of projective dimension d. It suffices to show the following.

If $U \in \mathcal{U}$ and W is a subspace of M of projective dimension d meeting U at a hyperplane, then $W \in \mathcal{U}$.

Suppose then that U and W are $PG(d)$-subspaces of M which generate a $PG(d+1)$ and $U \in \mathcal{U}$. Then by the definition of $\mathcal{U}$ there is a point x not in M with $x^\perp \cap M = U$. Choose $w \in W - U$. Then w is not collinear with x, and by the strong parapolar hypothesis, the convex closure $\langle\langle x, w\rangle\rangle$ is a symplecton R with $R \cap M = \langle U, W\rangle \simeq PG(d+1)$ a maximal singular subspace of R. Since W is a hyperplane of $R \cap M$ and R is a polar space, there is a point y in R with $y^\perp \cap (M \cap R) = W$. But since $y^\perp \cap M$ has projective dimension d by hypothesis, we have $y^\perp \cap M = W$ forcing $W \in \mathcal{U}$ as desired. The proof is complete. □

Theorem 16.2.5 (The basic trichotomy.) *One of the following must hold:*

1. $d = 0$.
2. $d = 1$.
3. *Every member of $\mathcal{M}$ has projective dimension $d+1$ and is itself a maximal singular subspace of some symplecton.*

Proof Suppose U_1 and U_2 are two $PG(d)$s of a space $M \in \mathcal{M}$ whose intersection is non-empty and yet has projective dimension at most $d-2$. By Lemma 16.2.4 there exist points u_i, $i = 1, 2$, such that ${u_i}^\perp \cap M = U_i$. Now let U_1' be the singular $PG(d+1)$ spanned by u_1 and U_1. Then for any point v in $U_1' - U_1$, we also have $v^\perp \cap M = U_1$, and this v could without loss of generality replace the original u_1. Now if u_2 were collinear with every such v, then ${u_2}^\perp$ would contain all U_1' since, lines being thick, the latter is generated by $U_1' - U_1$. But that would force U_1 to be contained in ${u_2}^\perp \cap M = U_2$, an absurdity. Thus there is some v in $U_1' - U_1$ which is not collinear with u_2, and we can take that one to be our u_1.

Then the convex closure of $\{u_1, u_2\}$ is a symplecton R which must contain the non-empty $U_1 \cap U_2$. Thus $R \cap M$ is non-empty. Then, by Lemma 16.2.2, $R \cap M$ is a $PG(d+1)$ which is a maximal singular subspace of R. Since the u_i are in R, we see that the ${u_i}^\perp \cap M$ are $PG(d)$s and so must be U_1 and U_2. But then U_1 and U_2 must meet at a $PG(d-1)$ since both lie in the $PG(d+1)$ known as $R \cap M$. That is contrary to the hypothesis.

Thus we cannot encounter a non-empty intersection of $PG(d)s$ in M which is not a hyperplane of each. But if $d \geq 2$, and M had dimension exceeding $d+1$, this would indeed be a possibility. Thus the hypotheses in the previous sentence must be false, hence the conclusion of the lemma holds. □

In the next three sections we consider each of these cases separately. The desire to proceed from the simplest cases to the most difficult will force us to consider them in reverse order.

16.2.2 The Third Case of the Trichotomy

Lemma 16.2.6 *Suppose $d \geq 1$, and that there exists an element of $\mathcal{M}$ of projective dimension $d+1$. Then Γ is itself a polar space.*

Proof There exists an element $M \in \mathcal{M}$ of projective dimension $d+1$. Let S be a symplecton on a point p of M (recall that for any parapolar space, every line lies in a symplecton, so this is true for every point as well). Then $M \subseteq S$ (Lemma 16.2.2). Choose a point $s \in S - p^{\perp}$.

Assume by way of contradiction that $\mathcal{P} \neq S$. At this point we invoke two results from Chap. 13. First, since Γ is a strong parapolar space and $d \geq 1$, Γ must have symplectic rank at least three. Thus by Lemma 13.5.1, it is locally connected. Second, by combining Theorem 13.5.5 (Chap. 13, p. 471) with our assumption that $\mathcal{P} \neq S$, and the two conditions just established (namely, Γ having symplectic rank exactly three, and the local-connectedness of Γ), we may conclude that $s^{\perp} - S$ is non-empty.

Now by (A2), for any $x \in s^{\perp} - S$, $x^{\perp} \cap M$ is a $PG(d)$ contained in the hyperplane $s^{\perp} \cap M$ of M. Thus $x^{\perp} \cap S = \langle s, M \cap s^{\perp} \rangle$. Since x was an arbitrary point of $S^{\perp} - S$, we have

$$s^{\perp} - S \subseteq (M \cap s^{\perp})^{\perp}. \tag{16.1}$$

Let N be a line of S on point s, but not in the clique $x^{\perp} \cap S$, and let R be the symplecton which is the convex closure $\langle\langle x, N \rangle\rangle$. Now by Eq. (16.1) we see that $s^{\perp} \cap R$ is the union of the two cliques $R \cap S$ and $(M \cap s^{\perp})^{\perp} \cap R$. Since R has polar rank at least three, that is impossible.

Thus $\mathcal{P} = S$, and the proof is complete. □

Lemma 16.2.6 yields the following.

Corollary 16.2.7 *In the third case of the trichotomy of Theorem 16.2.5, Γ is a polar space.*

From now on, we may assume in addition to (A1) and (A2), the axiom

(A3) *Γ is not itself a polar space.*

With this assumption, the trichotomy of Theorem 16.2.5 and Lemma 16.2.6 leave just two cases:

Case 2: $d = 1$ and every element of $\mathcal{M}$ has projective dimension larger than two.
Case 1: $d = 0$.

We analyze these cases in the next two sections.

16.3 The Second Case of the Trichotomy

In this section we assume the following.

(H1) $d = 1$ *and every element of* $\mathcal{M}$ *has projective dimension greater than* 2 in order to avoid the case that Γ is a polar space.

Now (A2) (b), Lemma 16.2.2 and (H1) immediately imply the following.

Lemma 16.3.1 (Basic consequences of Case 1.) *The following statements hold:*

1. *For every* $M \in \mathcal{M}$ *and every point* $x \in \mathcal{P} - M$, $x^\perp \cap M$ *is a line.*
2. *Every symplecton has rank three, and intersects any member of* M *at a plane or the empty set. It follows, that all singular subspaces of* Γ *are projective spaces.*

We begin with a study of maximal singular subspaces and symplecta.

16.3.1 The Maximal Singular Subspaces

Lemma 16.3.2 *The intersection of two distinct maximal singular spaces cannot properly contain a line.*

Proof Suppose A and B are distinct maximal singular subspaces for which $A \cap B$ contains a plane. Then choosing $x \in A - B$ and $y \in B - x^\perp$, one sees that the symplecton $\langle x, y\rangle$ has polar rank larger than three. That contradicts Corollary 16.2.3. □

A projective plane of Γ which is contained within some member of $\mathcal{M}$ will be called an $\mathcal{M}$*-plane*.

Lemma 16.3.3 *Suppose* π *is a plane in the symplecton* R*, which is disjoint from an* $\mathcal{M}$*-plane of* R*. Then* π *is a maximal singular subspace of* Γ.

Proof Let $\pi_0 := M \cap R$, $M \in \mathcal{M}$, be an $\mathcal{M}$-plane of R such that $\pi \cap \pi_0 = \emptyset$. Suppose, by way of contradiction, that π was properly contained in a singular subspace A. Then as R has polar rank three, $A \cap R = \pi$ and there is a point x in $A - R$. Since $x^\perp \cap R$ is a singular subspace π of R containing its maximal singular subspace, we must have

$$x^\perp \cap \pi_0 = \emptyset.$$

Now $x^\perp \cap M := L$ is a line, by our hypothesis that $d = 1$. Choose a point y in L and a point p in π_0. Since p is not collinear with x (as shown in the presented equation), there is a unique symplecton S containing p and x. Then S contains the line L as well as the line $N := p^\perp \cap \pi$. Since S is a polar space, there is a point z in $y^\perp \cap N$. Now we see that $z^\perp \cap M$ contains the line $z^\perp \cap \pi_0$ (in $M \cap R$) as well as the point x not in R. This means $z^\perp \cap M$ contains a plane, against $d = 1$.

Thus no such x exists, so no such A exists, and so π is a maximal singular subspace.

Lemma 16.3.4 *Suppose A is a maximal singular space, and that $M \cap A = \emptyset$ for some element $M \in \mathcal{M}$. Then A is a plane.*

Proof First A cannot be a line, since the parapolar space assumption forces every line to lie in some symplecton of rank three.

Next we claim that there exists a point $w \in M$ such that $w^\perp \cap A$ contains a line. Choose distinct points x and y in A. Now $x^\perp \cap M = L$, a line containing some point v. If v is collinear with y, we are done if we set $w = v$. So we may assume v is not collinear with y and form the symplecton $T := \langle\langle v, y\rangle\rangle$. Then $T \cap M = \pi_T$ is a plane. Since $x \in T$, $x^\perp \cap \pi_T$ is a line contained in $x^\perp \cap M = L$, another line, forcing $x^\perp \cap \pi_T = x^\perp \cap M = L$. Since T is a polar space containing y and L, we have $y^\perp \cap L = \{w\}$. Now $w^\perp \cap A$ contains line xy as desired.

To complete the proof, with our $w \in M$ for which $w^\perp \cap A = xy$, we may choose $z \in A - w^\perp$ (this is possible since A is a maximal singular subspace) and form the symplecton $Q := \langle w, z\rangle$. Then $Q \cap M$ and $Q \cap A = \langle x, y, z\rangle$ are opposite planes of the symplecton Q. It follows from Lemma 16.3.3 that the latter is a maximal singular subspace and so coincides with A. □

Corollary 16.3.5 *Any two elements of $\mathcal{M}$ have a non-empty intersection.*

Proof Since all elements of $\mathcal{M}$ have projective rank at least three, the result follows immediately from Lemma 16.3.4.

Lemma 16.3.6 *Any two members of $\mathcal{M}$ intersect in exactly a single point.*

Proof By Corollary 16.3.5, any two members of $\mathcal{M}$ have a non-empty intersection. Suppose some intersection of distinct M_1 and M_2 of $\mathcal{M}$ contained a line L. Since $d = 1$, $M_1 \cap M_2 = L$ exactly and we have the following:

(4.1) *No point of $M_1 - L$ is collinear with a point of $M_2 - L$ and* vice versa.

Now suppose x were a point such that $x^\perp \cap L = \emptyset$. Let $M_x \in \mathcal{M}$ contain x (assumption (A2) (a)). Then by Corollary 16.3.5, we have $M_x \cap M_i$ non-empty and disjoint from L – against (4.1). Thus the following applies.

(4.2) *For every point $x \in \mathcal{P}$, $x^\perp \cap L \neq \emptyset$.*

Now there exists a symplecton S containing L and meeting each M_i at a plane π_i (take $(u_1, u_2) \in (M_1 - L) \times (M_2 - L)$, so $u_1 \not\sim u_2$ by (4.1) and set $S = \langle\langle u_1, u_2\rangle\rangle$, for example). Let N be a line of S opposite L. Choose $s \in N$ and set $\{t\} = s^\perp \cap L$. Suppose $x \in s^\perp - S$. (Such an x exists since (A3) implies $S \neq \mathcal{P}$ so the element of $\mathcal{M}$ containing s cannot be a maximal singular subspace of S.) Now by (4.2), $x^\perp \cap L$ is non-empty. On the other hand this intersection lies in the clique $x^\perp \cap S$ containing s. Thus $x^\perp \cap L = s^\perp \cap L = \{t\}$. Since x was arbitrarily chosen in $s^\perp - S$ we have the following.

(4.3) $s^\perp - S \subseteq t^\perp$.

Since N is not in $t^\perp$ (N was opposite L), the clique $x^\perp \cap S$ does not contain N. Thus there is a unique symplecton $R = \langle\langle N, x\rangle\rangle$. Note that t is not in R since then

$R = \langle\langle t, N\rangle\rangle = S$ against our choice of x. Now by (4.3) $s^\perp \cap R$ is the union of two cliques, $R \cap S$ and $t^\perp \cap R$. That is impossible as $d \geq 1$ forces $s^\perp \cap R$ to contain a non-degenerate generalized quadrangle. □

16.3.2 The Symplecta

The following technical theorem is the crux of our final result.

Theorem 16.3.7 *Let π be a plane meeting M at a point for some element $M \in \mathcal{M}$. Suppose $\pi^\perp \cap M = \{p\}$. Then the following statements hold:*

1. *There is a unique symplecton R containing π.*
2. *Suppose $x \in M - R$. Then $x^\perp \cap \pi = \{p\}$. For any point $a \in \pi - \{p\}$ we have*

$$x^\perp \cap p^\perp \cap a^\perp - M \subseteq \pi^\perp.$$

3. *Any symplecton which intersects π at a line on p is oriflame.*

Proof Part 1. For any line pa of π on point p, $a^\perp \cap M$ is a line pu of M on p. Then, for any $b \in \pi - pa$, the symplecton $R = \langle\langle u, b\rangle\rangle$ meets M at a plane π_R (Lemma 16.3.1, part 2). Since $\pi^\perp \cap M = \{p\}$, we see that $b^\perp \cap M$ is a line pv of M distinct from pu. Thus the plane π_R is spanned by $\{u, v, p\}$.

Now suppose R' is another symplecton containing π. Then by Lemma 16.2.2, $R' \cap M$ is also a plane – but one which contains the lines $a^\perp \cap M$ and $b^\perp \cap M$ which span π. Since R and R' both contain π and π_R they must coincide.

Part 2. Let R be the unique symplecton on π, and set $\pi_R = R \cap M$ as above. Also let a, b, and p span π as in the first part. Then, for any point $c \in \pi - \{p\}$, $c^\perp \cap M = c^\perp \cap \pi_R$. It follows that $x^\perp \cap \pi = \{p\}$ for any point $x \in M - \pi_R$. Choose such an x and form the symplecton $S = \langle\langle x, a\rangle\rangle$. Then clearly $S \cap \pi = pa$. Finally, choose a line pz in $(px)^\perp \cap (pa)^\perp$ which is not in M. This is possible since such a line pz is a line on p in the symplecton S, and all the lines and planes of S on p have the incidence structure of a generalized quadrangle. Now $z^\perp \cap M = px$ not in π_R. But if z were not collinear with b, we could form a symplecton $T := \langle\langle z, b\rangle\rangle$ which would contain all of π. Then we must have $R = T$, by the uniqueness of R in part 1. But in that case, $z^\perp \cap \pi_R$ is a line, against $z^\perp \cap M = px$.

Thus z must be in $b^\perp$ and hence lies in $\pi^\perp$. From the choice of pz, part 2 is proved.

Part 3. As in the proof of part 2, let R be the unique symplecton on π, and let π_R be the plane $R \cap M$. Suppose now S is a symplecton which intersects π at a line on p – say the line pa. Then $S \cap M$ is a plane π_S containing $pu := a^\perp \cap M$. Now S cannot contain π_R since this would force $R = S$, while $S \cap \pi$ is only a line. Thus there is a line px in π_S that is not in π_R. Now the lines on p in $x^\perp \cap a^\perp$ are partitioned into two sets: (1) those on p lying in M (they are in the plane π_S), and (2) those not in M. By Part 2, the latter belong to $b^\perp$ and so are in the clique

$b^\perp \cap S$. Thus $(px)^\perp \cap (pa)^\perp$ is the union of two cliques. This means the generalized quadrangle of all lines and planes of S on p is a grid, and in turn S is an oriflame rank three polar space. □

Corollary 16.3.8 *Every symplecton of* Γ *is oriflame.*

Proof Let S be any symplecton. By Lemma 16.2.2 and (A1), there is an element $M \in \mathcal{M}$ such that $M \cap S = \pi_S$ is a plane. Now select a point $x \in S - \pi_S$. Then $x^\perp \cap M$ is a line, and since S is a polar space, $x^\perp \cap \pi_S$ is also a line contained in the former. Thus $x^\perp \cap M = x^\perp \cap \pi_S := N$, a line of π_S. Now there exists an element M_x of $\mathcal{M}$ which contains x. Then by Corollary 16.3.5, the intersection $M_x \cap M$ is a point p in $x^\perp \cap M = N \subseteq S \cap M = \pi_S$. Thus $\{p\} = M_x \cap M$ is on line N. Now select a line L of π_S which intersects N at exactly the point p. Since M is not a plane, there is a plane π of M which intersects $\pi_S = M \cap S$ at precisely the line L.

We claim that $\pi^\perp \cap M_x = \{p\}$. By way of contradiction, suppose $z \in \pi^\perp \cap M_x - \{p\}$. Then by (A2)(b) and (H1), $z \in M_x \cap M$, forcing $z = p$, a contradiction.

Now (M_x, p, π, S) satisfies the hypotheses of (M, p, π, S) in part 3 of Theorem 16.3.7. By part 3 of Theorem 16.3.7, S is oriflame. The proof is complete.

16.3.3 More Maximal Singular Subspaces

In Lemma 16.3.4 we considered maximal singular subspaces which intersected some member of $\mathcal{M}$ trivially. That still leaves the class of all maximal singular subspaces which intersect non-trivially each member of $\mathcal{M}$. This class, which by Corollary 16.3.5 includes $\mathcal{M}$ itself, is addressed in the next corollary.

Corollary 16.3.9 *Suppose* A *is a maximal singular subspace which is not a plane that intersects some member of* $\mathcal{M}$ *trivially. Then for every point* $x \in \mathcal{P} - A$*, the set* $x^\perp \cap A$ *is a line.*

Proof Fix $x \in \mathcal{P} - A$. Let M_x be a singular subspace of $\mathcal{M}$ which contains x. Lemma 16.3.4 and our hypothesis on A force the existence of a point p in $M_x \cap A$. If $M_x \cap A$ contains a line then $x^\perp \cap A$ contains a line, but in fact it is exactly a line by Lemma 16.3.2. So we have our conclusion.

So we may suppose $M_x \cap A = \{p\} = x^\perp \cap A$ exactly.

We claim that we can find two lines pa and pb in A, such that $a^\perp \cap M_x \neq b^\perp \cap M_x$. If this were not true, we would have a line L in M_x lying on p such that $L \subseteq A^\perp$. Since $x^\perp \cap A = \{p\}$ implies $A \cap M_x = \{p\}$, this would contradict the fact that A is a maximal singular subspace.

With a and b chosen as in the previous paragraph, we see that that $\pi := \langle p, a, b\rangle$ is a plane, $(pa)^\perp \cap M_x$ and $(pb)^\perp \cap M_x$ are distinct lines, and $\pi^\perp \cap M_x = \{p\}$. Since the symplecton that contains $\{x, a\}$ has polar rank three, $x^\perp \cap p^\perp \cap a^\perp$ is not a clique. Thus there is a line pz in $x^\perp \cap a^\perp$ that does not lie in M_x. Since $x^\perp \cap A = \{p\}$, we see that pz is not in A. But by part 2 of the preceding Theorem 16.3.7, one must have $z \in \pi^\perp$. That contradicts Lemma 16.3.2 since z is not in A. The proof is complete. □

Now let $\mathcal{M}'$ be the class of all maximal singular spaces which are either in $\mathcal{M}$ or are not planes intersecting trivially some member of $\mathcal{M}$. Then $\mathcal{M}$' satisfies all the hypotheses on $\mathcal{M}$ appearing in axioms (A1) and (A2). That means that we could have taken $\mathcal{M} = \mathcal{M}'$ all along.

So, from this point onward, we assume the following.

(H2) *If A is a maximal singular subspace not in the class $\mathcal{M}$, then A is a projective plane that is disjoint from at least one member of $\mathcal{M}$.*

Let $\mathcal{A}$ be the class of all maximal singular subspaces which are planes which are disjoint from some member of $\mathcal{M}$. Then (H2) can be taken to be the assertion that the set of all maximal singular subspaces of Γ can be partitioned as $\mathcal{M} + \mathcal{A}$.

Lemma 16.3.10 *If $(M, A) \in \mathcal{M} \times \mathcal{A}$, then $M \cap A$ is the empty set or is a line.*

Proof Because of Lemma 16.3.2, the only way the conclusion could fail would be if $M \cap A = \{p\}$, a single point. Now in that case we have $A^\perp = A$ since A is a maximal singular subspace. In particular, we have $A^\perp \cap M = \{p\}$. Thus (M, p, A) satisfies the hypotheses on (M, p, π) in Theorem 16.3.7. By part 1 of Theorem 16.3.7, A is contained in a unique symplecton R, and by part 2, for any point x of M not in the plane $R \cap M$, and point a of $A - \{p\}$, we have

$$x^\perp \cap p^\perp \cap a^\perp - M \subseteq A^\perp.$$

The points on the left side (which are automatically distinct from p) cannot be in A since $x^\perp \cap A = \{p\}$. Yet the right side is A exactly. It follows that $x^\perp \cap p^\perp \cap a^\perp - M = \emptyset$ so $x^\perp \cap p^\perp \cap a^\perp \subseteq M$. But $\{x, a\}$ lies in a symplecton of polar rank three, so $x^\perp \cap p^\perp \cap a^\perp$ is not a clique and so is not in M. This contradiction completes the proof. □.

Lemma 16.3.11 *Any two distinct members of $\mathcal{A}$ intersect at the empty set or at a single point.*

Proof Suppose by way of contradiction that $A_1 \cap A_2 = L$, where the A_i are planes belonging to $\mathcal{A}$. Choose a point p in L, and an element $M \in \mathcal{M}$ containing point p (M exists by (A1)). Then by Lemma 16.3.10, $M \cap A = L_i$ is a line, $i = 1, 2$. But if $q \in L - \{p\}$, and $L_1 \neq L_2$, we see that q is not in M, yet $q^\perp \cap M$ contains the plane $\langle L_1, L_2 \rangle$ of M, against (H1). Thus we must have $L_1 = L_2 = L$. Now let S be the unique symplecton containing $A_1 \cup A_2$. Then A_1, A_2, and $S \cap M$ are three distinct planes on a common line in the rank three polar space S, and this contradicts the fact that S is oriflame, by Corollary 16.3.8. □

Lemma 16.3.12 *Every line of Γ lies in a unique member of $\mathcal{M}$.*

Proof First, since Γ is a strong parapolar space, every line L lies in some symplecton S of polar rank three. Since S is oriflame by Corollary 16.3.8, L lies in exactly two planes of S. By Lemma 16.3.6, at most one of these planes can lie in a member

of $\mathcal{M}$. By Lemma 16.3.11, at most one of the planes can belong to $\mathcal{A}$. Since $\mathcal{M}$ and $\mathcal{A}$ list all maximal singular subspaces of Γ, it follows, then, that exactly one of the planes of S on line L lies in a member of $\mathcal{M}$. In any event line L lies in a unique member of $\mathcal{M}$. □

16.3.4 The Final Result for Case 2 with (H2)

Theorem 16.3.13 $\Gamma = (\mathcal{P}, \mathcal{L})$ *is the Grassmannian of lines of some (possibly infinite-dimensional) projective space* $P = \mathbf{P}(V)$ *of projective dimension at least four.*

Proof First recall from Lemma 16.3.6 that:

(1) *Any two elements of* $\mathcal{M}$ *are incident with a unique point. That means* $(\mathcal{M}, \mathcal{P})$ *is a linear space.*

Let us say that an element of $\mathcal{M}$ is *incident* with an element of $\mathcal{A}$ if and only if the two of them intersect at a line. Our second statement is as follows.

(2) *Any three members of* $\mathcal{M}$ *which are not incident with a common point are in fact incident with a unique element of* $\mathcal{A}$.

These three members of $\mathcal{M}$ must pairwise meet at three points which span a plane.

(3) *The incidence system of elements of* $\mathcal{M}$ *and all points which are incident with a plane* A *of* $\mathcal{A}$ *is a projective plane.*

This is because Lemmas 16.3.12 and 16.3.6 imply that every line L of A is the intersection $L = M \cap A$ for a unique element M of $\mathcal{M}$ which is by definition incident with the plane A.

As a consequence of these observations, we also have the following:

(4) *If two distinct elements of* $\mathcal{M}$ *are incident with an element* A *of* $\mathcal{A}$*, then the unique point of their intersection is incident with* A*. In other words, the set of elements of* $\mathcal{M}$ *which are incident with* A *is a subspace of the linear space* $(\mathcal{M}, \mathcal{P})$.

The four statements assert that $(\mathcal{M}, \mathcal{P})$ is a linear space in which any three "points" which are not incident with a common "line" are contained in a subspace which is a projective plane. This means $(\mathcal{M}, \mathcal{P})$ is a projective space (of unknown projective dimension at least four) and the theorem is proved. □

16.4 The Final Case 1: $d = 0$

In this section we assume the following hypothesis.

(E) *1.* Γ *is a strong parapolar space.*
2. Γ *contains a class* $\mathcal{M}$ *of maximal singular subspaces with these properties:*

(a) Every member of $\mathcal{M}$ *is a projective space (not necessarily of finite rank).*
(b) Every point lies in a member of $\mathcal{M}$.
(c) If $M \in \mathcal{M}$, *and* x *is a point not in* M, *then* x *is collinear with a unique point* $f_M(x)$. *(Thus* $f_M : (\mathcal{P} - M) \to M$ *is a well-defined mapping, the projection onto* M.*)*

3. Any symplecton is classical.

Remark Part 3 comes from hypothesis (A1).

Lemma 16.4.1 *The following are easy consequences of Hypothesis (E):*

1. *Every symplecton meets a member of* $\mathcal{M}$ *at the empty set, or at a line.*
2. *If* L *is any line disjoint from* $M \in \mathcal{M}$, *then the mapping* $L \to M$ *induced by the projection mapping has as its image either a single point or a line of* M.
3. *Every symplecton is a generalized quadrangle, and consequently* Γ *is a strong parapolar space of symplectic rank exactly two.*

Proof Of course, this is just a re-hash of Lemma 16.3.1, localized at a point, but let's prove it from (E) alone. Suppose the symplecton S intersected the subspace M of $\mathcal{M}$ non-trivially. If $S \cap M$ is not a maximal singular subspace of S, then there exists a point $p \in (S - M) \cap (S \cap M)^{\perp}$, at which point assumption E2(b) forces $S \cap M$ to contain exactly one point, say m. Choose $y \in S - m^{\perp}$. Again by part E2(b), y is collinear with a unique point m_y of M. Then $m_y \in m^{\perp} \cap y^{\perp} \subseteq S$. But that contradicts $S \cap M = \{m\}$, a point.

So we may assume that $S \cap M$ is a maximal singular subspace of S. Then choosing $x \in S - M$, we see that E2(b) implies $x^{\perp} \cap (S \cap M)$ is a point. But the latter is a hyperplane of $S \cap M$. Thus the maximal singular subspace $S \cap M$ of S is a line, and so S is a generalized quadrangle.

Now this argument works for any symplecton S since any of its points lies in an element $\mathcal{M}$. Parts 1 and 3 of the lemma have been proved.

Suppose L is a line disjoint from a member of M and let $f : L \to M$ be the mapping which maps each point of L to the unique point of M with which it is collinear. Choose $x \in L$ and let R be the unique symplecton containing $f(x)$ and L. Then by part 1, R meets M at a line L' opposite L in R. Then $L' = f(L)$, so $f(L)$ is a line as required. □

16.4.1 Fibered Symplecta

We assume hypotheses (E). We suppose $M \in \mathcal{M}$ and that S is a symplecton sharing no point with M. Then the restriction of the projection mapping onto M produces a mapping $f : S \to M$ (namely $f := f_M|_S$) which takes a line of S either to a line of M, or to a point of M. Let $f(S)$ be the collection of all image points – that is, $f(S) = \{m \in M | m^\perp \cap S \neq \emptyset\}$.

Lemma 16.4.2 *Suppose $m \in f(S)$. Then its fiber $f^{-1}(m)$ is either a single point or is a line of S.*

Proof By convexity of S, the fiber $f^{-1}(m)$ is a singular subspace of S. Since S is a generalized quadrangle, this fiber is either a point or a line. □

Lemma 16.4.3 *Let F be the set of points in M whose fibers are lines of S. Then F is a subspace of M. Moreover, if E is a line contained in F, then $f^{-1}(E)$ is a subquadrangle of S which is a grid. The fibers of the points of E form one of the two line-spreads of this grid.*

Proof Suppose a' and b' are two distinct points of F. Let A and B be their respective fibers. Then A and B are opposite lines of S. Then for each point x of A, there is a unique line T_x on x which intersects B at the unique point $b(x)$ of $x^\perp \cap B$. Then this transverse line maps by f to the unique line $E = a'b'$ of M (Lemma 16.4.1, part 2). Thus, for each point $m \in E$, the fiber $D(m) := f^{-1}(m)$ must intersect each transverse line T_x. This gives $E \subseteq F$.

On the other hand, if $z \in D(m)$, then either $D(m) = A$ and $z \in T_z$, or $D(m)$ is a line opposite A, meeting each tranverse line T_x, $x \in A$. It follows that $D(m)$ cannot contain a point not in some T_x.

Thus we see that the fibers of points of E form a line-spread of $G := \cup_{x \in A} T_x$ while the system $\{T_x | x \in A\}$ of transverse lines for A and B form another such system. Since the lines of each system are pairwise opposite, no further collinearities exist among the points. It follows that G is a subspace of S which is a grid with these two systems of line-spreads. □

A Particular Situation

We suppose f is not injective. Then there are distinct points a_1 and a_2 of S collinear with a common point $f(a_1) = f(a_2) = a'$. Since $f^{-1}(a') := a'^\perp \cap S$ is a clique with at least two points, and since lines are maximal singular subspaces of S, we see that the fiber $f^{-1}(a')$ is a line A. Then any other line L on point a_1 is mapped bijectively to a line $f(L) := L'$ of M. Next suppose B is a line meeting L at a point b_1 distinct from a_1, chosen so that its image $f(B) := B'$ is also a line of M. Then B' meets L' at the point $b' = f(b_1)$. Then $\langle L', B' \rangle_M$ is a projective plane π.

Now A and B are opposite lines of the symplecton S. Thus each point a_i of line A is collinear with a unique point b_i of B, and we denote the full collection of lines $\{T_i := a_i b_i\}$ by $\mathcal{T}$ and call them the *tranversal lines for A and B*. It is clear then

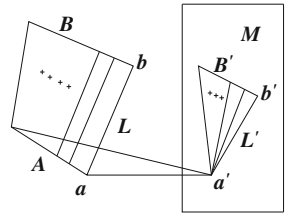

Fig. 16.1 A special situation. In this figure a and b are the respective points a_1 and b_1 of the text

that these transversal lines are mapped onto the full pencil of lines of π on point a', so in fact $\pi \subseteq f(S)$. This situation is illustrated in Fig. 16.1.

Let $R := \cup\{f^{-1}(p) | p \in \pi\}$. Clearly R is a subspace of S since π is a subspace of M. Select any point u on line L so that u is distinct from both a and b. Let T_i be a transversal from A to B which is distinct from (and hence opposite to) L. Then u is collinear with a unique point v of T_i and v is not in B. It follows that line uv is opposite line B in S.

Now $f(u)$ and $f(v)$ are points of $\pi - (B' \cup \{a'\})$, on different members of the line pencil on a'. (In particular, $f(u) \neq f(v)$.) Then f maps the line uv of S to the line $f(u)f(v)$ of π. Since π is a projective plane, $f(u)f(v)$ meets line $B' = f(B)$ at a point $f(b_j)$ for some $b_j \in B$. This means uv contains a point w with $f(w) = f(b_j)$. Since uv and B are opposite lines, $w \neq b_j$ and so are distinct points belonging to the fiber $C := f^{-1}(f(b_j))$. By Lemma 16.4.3 above, $G := f^{-1}(a'f(b_j)) = \langle A, C\rangle_S$ is a grid.

Now R contains $\langle A, B\rangle_S$ and so properly contains the grid G. Then any point of G lies on a line of R which is not in G, and, since such a line lies in R but does not lies in $G = f^{-1}(L')$, it must map onto a line of π. Let us choose distinct points x and y on a line T transversal to two of the line-fibers of G. Let L_x and L_y be lines of R which lie on x and y, respectively, but are not lines of G. One notes that $L_x \cap L_y = \emptyset$ since any point of their intersection would lie outside T while being collinear with distinct points x and y of T. Then $f(x) \neq f(y)$ and $f(L_x)$ and $f(L_y)$ are lines which intersect $L' = f(L')$ at distinct points $f(x)$ and $f(y)$ and so intersect each other at a point $p \in \pi - L'$. But that means $f^{-1}(p)$ contains a point of L_x and a point of L_y and since these are disjoint lines we see that $p \in F \cap \pi$.

It now follows from Lemmas 16.4.2 and 16.4.3 that R is a generalized quadrangle containing a line spread $\mathcal{F}$ (the fibers of $f : R \to \pi$), any two of which generate a grid, and, letting $\mathcal{G}$ be the collection of grids formed in this way, the incidence system $(\mathcal{F}, \mathcal{G})$ is isomorphic to $(\mathcal{P}_\pi, \mathcal{L}_\pi)$, the points and lines of the projective plane π.

By E(4) S is a classical quadrangle. Now Lemma 15.3.1 of Arjeh Cohen shows that this is impossible.

We have proved the following.

Theorem 16.4.4 *Suppose S is a symplecton disjoint from a subspace M of $\mathcal{M}$. Then either:*

1. *The projection into M induces a projective embedding $f : S \to M$, or*
2. *S is a grid and the projection on M induces a mapping $f : S \to M$ onto a line L of M. The fibers of the points of L form one of the line-spreads of the grid S.*

Corollary 16.4.5 *Suppose S is a symplecton which is not a grid. Suppose π is a plane which meets S at a line. Then for any point $x \in \pi - S$, x lies in a unique member M_x of $\mathcal{M}$ and the plane π itself lies in M_x.*

Proof Suppose $\pi \cap S$ is a line L and choose x in $\pi - L$. Suppose x belongs to a singular subspace M' of $\mathcal{M}$ which does not contain π. Then polar rank two forces $\pi \cap M' = \{x\}$. Also, convexity of S forces $M' \cap S = \emptyset$. But now projection on M' forces a mapping $f : S \to M'$ which posseses a non-trivial fiber. By the Theorem 16.4.4, S is a grid, an absurdity. So no such M' exists.

But by hypothesis x lies in some member M_x of $\mathcal{M}$, and so we see that we must have $\pi \subseteq M_x$ for any such singular space. This fact forces the uniqueness of M_x. □

A line is called an *$\mathcal{M}$-line* if it is contained in one of the singular subspaces of $\mathcal{M}$. By axiom E2(a), the $\mathcal{M}$-lines cover all the points.

Corollary 16.4.6 *Suppose S is a symplecton which is not a grid. Then every line of S is either an $\mathcal{M}$-line or is already a maximal singular subspace of Γ.*

Proof If L is a line of S that lies in a larger singular subspace, then there is a plane π containing L. Since S is a generalized quadrangle, $\pi \cap S = L$. Now, by Corollary 16.4.5, L is an $\mathcal{M}$-line. Otherwise L is a maximal singular subspace. □

16.4.2 Unfibered Symplecta: M-Projections which Embed a Symplecton

Let S be a symplecton. We let $N^2(S)$ denote the set of points $x \in \mathcal{P} - S$ for which $x^\perp \cap S$ is a line. Similarly we write $N^1(S)$ for the set of points $x \in \mathcal{P} - S$ for which $x^\perp \cap S$ is a single point, and finally write $N^0(S)$ for the set of points $x \in \mathcal{P} - S$ for which $x^\perp \cap S = \emptyset$. Then we have the following partition of points:

$$\mathcal{P} = S + N^2(S) + N^1(S) + N^0(S).$$

The conclusion of Corollary 16.4.5 motivates another definition. The set U of all points $p \in \mathcal{P}$ which lie in a unique member of $\mathcal{M}$ will be called *the uniqueness set*.

Theorem 16.4.7 *Suppose $M \in \mathcal{M}$ and S is a symplecton which shares no point with M. Suppose further that the projection onto M induces an embedding $f : S \to M$ as in the first case of Theorem 16.4.4. Choose any pair (x, y) of non-collinear points of S, let M_y be a member $\mathcal{M}$ containing y, let L_y be the line $M_y \cap S$, and let t be the unique point of $x^\perp \cap L_y$. Let R be the unique symplecton containing t and $f(x)$. Then R is not a grid.*

The following uniqueness results hold:

1. *The subspace M_y is the unique element of $\mathcal{M}$ containing the point y.*
2. *Since y is arbitrarily chosen in S, we see that every point of S lies in a unique member of $\mathcal{M}$.*
3. *Every point of M lies in a unique member of $\mathcal{M}$.*
4. *Every line of S is either an $\mathcal{M}$-line or is itelf a maximal singular space.*

Proof Let x, y, M_y, L_y, and t be as chosen. Then $f(x)$ is the unique point of M collinear with x. We have $M \cap M_y = \emptyset$ since otherwise $f(L_y)$ is a single point, against f being an embedding. Then $f(x)$ is collinear to a unique vertex z in $M_y - L_y$. Now $(x, f(x), z, t, x)$ is a 4-circuit lying in the symplecton R. Now R meets S at line xt and R meets M_y at line tz. Also, since $R \cap M$ is non-empty, it too is a line – in fact it is the line $f(x)f(t)$. The configuration is illustrated in Fig. 16.2.

Now we see that the point $f(x)$ sits on three distinct lines of R: $f(x)x$, $f(x)z$, and $f(x)f(t)$. Thus R cannot be a grid.

Now we note that $y \in M_y - R$ and M_y meets R at a line. Thus by Corollary 16.4.5 y lies in a unique member of $\mathcal{M}$ confirming part 2.

Similarly every point of $M - (M \cap R)$ lies in a unique member of M. But R is (now, uniquely) determined by the choice of (x, y). But if we choose $u \in S \cap y^{\perp} - t^{\perp}$, then replacement of the pair (x, y) by (u, t) in the construction produces a new symplecton R' which is not a grid, and which meets M at the line $f(u)f(y)$. Note that line uy is opposite line xt in S and so $f(t)f(x)$ and $f(u)f(y)$ are disjoint lines in M since $f : S \to M$ is an embedding. But as before all points of line $f(z)f(t)$ lie in $M - f(u)f(y)$ and so lie in a unique member of $\mathcal{M}$. Thus all points of M have this uniqueness property. This confirms part 3.

There is more. Let x, y, u, t be as in the previous paragraph. Now $M_y \cap R = tz$ (where, as before, z is the unique point of $f(x)^{\perp} \cap M_y$) and $M_y \cap R' = yz'$ where z' is the unique point of M_y collinear with $f(u)$. These lines are disjoint since $f(u)f(y)$ and $f(t)f(x)$ are disjoint lines of M (as observed in the previous paragraph) and the perpendicular relation produces an isomorphism $M_y \to M$. But since the points outside either of these lines are uniqueness points, we have $M_y \subseteq U$.

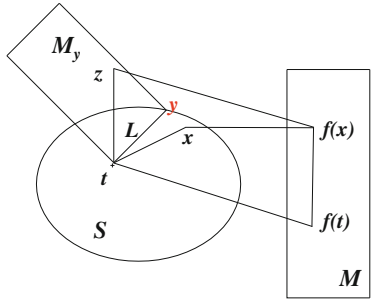

Fig. 16.2 The configuration of Theorem 16.4.7

Now we address part 4. Note that we do not yet know that S is not a grid, otherwise we could apply Corollary 16.4.6 at once. Suppose π is a plane not in an element of $\mathcal{M}$ meeting S at a line N. Then N is not an $\mathcal{M}$-line. Choose distinct points x and t of line N and let M_t be a member $\mathcal{M}$ on point t. Then $M_t \cap S$ is a line L_t on t. Clearly $L_t \neq N$. Now choose y in L_t distinct from t. Then y is not collinear with x. If we rename things, writing M_y for M_t and L_y for L_t, we have exactly the construction of the symplecton R at the beginning of the theorem. Now R is not a grid, but sits on the line $N = xt$ of the plane π. This is contrary to Corollary 16.4.6. Thus no such π exists. It follows that all lines of S which are not $\mathcal{M}$-lines are already maximal singular subspaces. The proof is complete. □

Corollary 16.4.8 *Suppose the symplecton S is not a grid and $M \cap S = \emptyset$ for some $M \in \mathcal{M}$. Then every point of $\mathcal{P}$ lies in a unique member of $\mathcal{M}$.*

Proof Let U be the collection of uniqueness points – those points which lie in a unique member of $\mathcal{M}$. Our objective is to prove that $U = \mathcal{P}$.

By Theorem 16.4.4, since S is not a grid, the mapping $f : S \to M$ induced by the projection into M, is an embedding. By Theorem 16.4.7 part 2, $S \subseteq U$. By Corollary 16.4.6, since S is not a grid, every point of $N^2(S)$ is a uniqueness point. Now suppose $y \in N^1(S) \cup N^0(S)$. Then by our hypothesis, y lies in a singular subspace $M_y \in \mathcal{M}$. Then $M_y \cap S = \emptyset$. Then, as S is not a grid, one obtains an embedding $f_y : S \to M_y$. Upon replacing M by M_y in Theorem 16.4.7 part 3, we see that $M_y \subseteq U$ and in particular $y \in U$.

The discoveries of the previous paragraph can be summarized by asserting

$$\mathcal{P} = S + N^2(S) + N^1(S) + N^0(S) \subseteq U,$$

which we were to prove. □

16.4.3 The Case that Some Member of $\mathcal{M}$ is a Line

Theorem 16.4.9 *If a line L is a member of $\mathcal{M}$, then one of the following holds:*

1. *Γ is itself a generalized quadrangle.*
2. *Γ is a product geometry $L \times P$, where P is a maximal singular subspace. Every symplecton of Γ is a grid.*

Proof Suppose L is a line in M. Since Γ is a strong parapolar space of rank two, L lies in a symplecton S. If $S = \mathcal{P}$, the first conclusion holds and we are done. So we assume $S \neq \mathcal{P}$.

Choose $y \in \mathcal{P} - S$. Then $y^\perp \cap L$ is a single point, say p. We claim that in S there is only one further line on p besides L and that line is $y^\perp \cap S$. If not there would be a line N of S on p not in $y^\perp$. We could then form the symplecton $R := \langle\langle y, N\rangle\rangle$ and choose a point $z \in R - p^\perp$. Then z is collinear with a point q of L distinct from p. Then $q \subseteq p^\perp \cap z^\perp \subseteq R$, and so lies in $R \cap L = R \cap S \cap L = N \cap L = \{p\}$, an absurdity.

Thus, for each $y \in p^\perp - S$ we have $y^\perp \cap S = N$ for any line N in S which contains p and is distinct from L. Since $y^\perp \cap S$ is a clique, the line N is unique. Thus S is a grid. But then we have $p^\perp - S \subseteq N^\perp$, which must be a singular subspace (otherwise Γ would not have symplectic rank two). Thus, for any p of the line L, we have

$$p^\perp = L \cup A_p, \text{ and } L \cap A_p = \{p\},$$

where $A_p = {N_p}^\perp$ and N_p is the unique line of the grid S on p such that N_p is distinct from the line L. Clearly each A_p is a singular subspace which is not a line. Moreover we have a partition into maximal singular subspaces:

$$\mathcal{P} = \uplus\{A_p | p \in L\}. \tag{16.2}$$

Now choose any point x in $\mathcal{P}$. If $x \in L$ set $p(x) = x$ and note that L is the unique line on x not in $A_{p(x)}$. If $x \in \mathcal{P} - L$, let $p(x)$ be the unique point of $x^\perp \cap L$. Then $x \in A_{p(x)}$. We claim that there is a unique line L_x on x which is not in $A_{p(x)}$. First there is at least one such line, since the symplecton $R_x := \langle\langle xp(x), L\rangle\rangle$ contains one. On the other hand, if L' were such a line, then the symplecton $\langle\langle xp(x), L'\rangle\rangle$ would intersect L non-trivially and hence would contain L (Lemma 16.4.1, part 1), forcing it to coincide with the symplecton R_x. But R_x satisfies the hypothesis that we had for S above, and so R_x is a grid. That means there is only one line L_x on x which does not lie in the singular space $A_{p(x)}$, and that line is opposite L.

Thus all lines of Γ which are not in one of the singular subspaces A_p form a spread of lines transverse to the components A_p of the partition in Eq. (16.2). Every symplecton on such a transverse line is a grid with its intersections with the A_p forming a spread. It follows that for any two distinct points x and y of the line L, the point-bijection $A_x \to A_y$ induced by the system of transverse lines takes lines of A_x to lines of A_y. This is the last step needed to conclude that Γ is the product geometry $L \times A_p$. □

16.4.4 Symplecta Disjoint from No Member of $\mathcal{M}$

Theorem 16.4.10 *Suppose S is a symplecton that is not disjoint from any singular subspace of $\mathcal{M}$. Then either (a) Γ is a generalized quadrangle, (b) S is a grid, or (c) all points are uniqueness points.*

Proof We suppose that S is a symplecton disjoint from no member of $\mathcal{M}$. Then S contains the full set of *non-uniqueness* points, $\mathcal{P} - U$. For if point p lies in two distinct elements $i = 1, 2$. Then by hypothesis $S \cap M_i = L_i$ is a line. If S did not contain $\{p\} = M_1 \cap M_2$, then the L_i are opposite lines of the quadrangle S, lying in $p^\perp \cap S$. But that is impossible, since $p \notin S$ implies that $p^\perp \cap S$ is a clique. Thus we must have $p \in S$.

By way of contradiction we assume that S is not a grid and that M_1 and M_2 are two distinct members of $\mathcal{M}$ which meet at a point p in the symplecton S.

If either M_i were a line, we could apply Theorem 16.4.9 to conclude that Γ is a generalized quadrangle, or that $\Gamma = L \times P$ for some singular subspace P. But in the latter case S would be a grid, contrary to assumption. So Γ is a generalized quadrangle, which is one of our conclusions. Thus we may assume that neither M_1 nor M_2 are lines.

For $i = 1, 2$, let $L_i := S \cap M_i$ and suppose N_i is a line of M_i on p distinct from the line L_i (this is possible since M_i is not a line). Let $R := \langle\langle N_1, N_2\rangle\rangle$, the unique symplecton on the N_i. Choose a point $z \in R - p^\perp$ and let M_z be a member of $\mathcal{M}$ on point z. Now by our hypothesis, M_z cannot be disjoint from S, and so $M_z \cap S$ is an $\mathcal{M}$-line A_z of S which is not on point p. Then the unique point a_z of $p^\perp \cap A_z$ is collinear with both p and z and so belongs to R – that is $R \cap S = pa_z$. Also a_z is on neither line L_i since the unique point m_i of $z^\perp \cap M_i$ lies in $N_i - \{p\}$, $i = 1, 2$. Thus the three lines za_z, zm_1, and zm_2 on point z are all distinct and belong to R. It follows that R is not a grid.

Now choose any point $w \in S - p^\perp$ and let M_w be a member of $\mathcal{M}$ containing point w. If $M_w \cap R = \emptyset$, then Corollary 16.4.8 (applied with the non-grid R and singular space M_y replacing S and M respectively) would not allow M_1 and M_2 to intersect at p. Thus we must conclude that $M_w \cap R$ is a line B of R not containing p. Then the point t of $p^\perp \cap B$ lies in R and also lies in $p^\perp \cap w^\perp \subseteq S$, and so is a point of $R \cap S = pa_z$ as well as a point of the $\mathcal{M}$-line $L_w = M_w \cap S$.

But in this configuration we can replace R by a new symplecton $R' := \langle\langle N_1, N_2'\rangle\rangle$ where N_2' is another line of M_2 on p distinct from both N_2 and L_2. Then $R' \cap R = N_1$. But just as we argued for R above, R' is not a grid, and meets S at a line. We note that the two lines $R \cap S$ and $R' \cap S$ are distinct since $R \cap R' = N_1$. But in the last line of the previous paragraph we saw that the $\mathcal{M}$-line $L_w = M_w \cap S$ intersected $R \cap S$ at a point (we called t). So similarly, L_w intersects line $R' \cap S$ at a point s distinct from t. Since $s \in R' \cap S \subseteq p^\perp \cap L_w$ and $t \in R \cap S \subseteq p^\perp \cap L_w$, the observation that $s \neq t$ along with the gamma space property of S then forces $L_w \subseteq p^\perp$, which is absurd since $w \in L_w - p^\perp$. □

Remark Note that the ability to choose line N_2' in the last paragraph of this proof actually utilizes the standing assumption that lines are thick.

16.4.5 When All Symplecta are Grids

In this section we add to (E) the following extra hypothesis.

(G) *Every symplecton of* Γ *is a grid.*

Theorem 16.4.11 *If hypothesis (G) is assumed, then* $\Gamma \simeq M \times A$ *where* $M \in \mathcal{M}$ *and* A *is a singular subspace.*

Proof By way of contradiction suppose the conclusion is false. Then by Theorem 16.4.9 we may assume that no member of $\mathcal{M}$ is a line.

Choose any singular subspace M of $\mathcal{M}$. Suppose, for some point p in M, L_1 and L_2 are two distinct lines on p which are not in M. If L_1 is not in $L_2{}^{\perp}$, then the symplecton R on L_1 and L_2 must meet M at a third line (see Lemma 16.4.1), and so is not a grid, against (G). Thus, always we must have the following.

(Step 1) *For any point p in a singular subspace M of $\mathcal{M}$, there exists another maximal singular subspace $A(p)$ such that*

$$p^{\perp} = M \cup A(p) \text{ where } M \cap A(p) = \{p\}.$$

Suppose, for the moment, that p and M are fixed as in Step 1. Choose any point y not in $p^{\perp}$. Then by hypothesis (E), y is collinear with a unique point y_1 of $M - \{p\}$. Then there is a symplecton G which is a grid on the intersecting lines yy_1 and y_1p. Then p lies on a unique line L of this grid distinct from the line py_1. Then by Step 1, L is a line of the singular space $A(p)$. Moreover y is collinear with a unique point y_2 of line L, since all of this occurs within the grid G.

Now suppose y were collinear with another point y_2' of $A(p)$. Then y_2' would lie in $y^{\perp} \cap p^{\perp} \subseteq G$, and so would lie in $G \cap A(p) = L$. But that forces $y_2' = y_2$. Thus y_2 is the unique point of $A(p)$ which is collinear with such a point y.

Now y lies in some member, say M_y, of $\mathcal{M}$, and so applying Step 1, with M_y replacing M and y replacing p, we see that $y^{\perp} = M_y \cup A$, the union of two maximal singular subspaces which intersect at point y. Since yy_1 and yy_2 are distinct lines of grid G on y, we see that yy_1 lies in one of the maximal singular subspaces (M_y or A) and yy_2 lies in the other.

We have established the following.

(Step 2) *Let p and M be as in Step 1. If $y \in \mathcal{P} - p^{\perp}$, then y is collinear with a unique point y_1 of M and with a unique point y_2 of $A(p)$. Then $y^{\perp}$ is the union of two maximal singular spaces; one is $A(y_1)$ and the other contains yy_2. Thus every line on point y is either in $y_1{}^{\perp}$ or is in $y_2{}^{\perp}$.*

We next show the following.

(Step 3) *Let p, M, y, y_1 and y_2 be as in (Step 2). If a point y' is collinear with both y_1 and y_2, it is either y or p.*
Clearly, any such y' lies in the (convex) grid G which contains them, and the result follows.

Now we can complete the proof of the theorem. First we can uniquely assign coordinates from $M \times A(p)$ to each point at follows. If y is not in $p^{\perp}$, we assign coordinates (y_1, y_2) to y. If m is a point of M, we assign coordinates (m, p) to m. If a is a point of $A(p)$, we assign the coordinates (p, a) to a. Note that p has

coordinates (p, p) and by Step 3 each point of $\mathcal{P}$ receives a unique coordinate by this device. Conversely, if (m, a) is arbitrarily chosen with neither coordinate equal to p, then there is a unique grid $G(m, a)$ on $\{m, a\}$ containing a unique point $y \in G(m, a) - \{p\}$ collinear with both m and a. Thus introducing coordinates produces a complete bijection $\mathcal{P} \to M \times A(p)$.

Now let us consider the collection of coordinates of the points on an arbitrary line of L. If L is in M or in $A(p)$, then one of the coordinates is constantly p while the other coordinates range through a line of M or $A(p)$, respectively. Similarly, if L contains a point $y = (y_1, y_2)$ not in $p^\perp$, then by Step 2, this line is in either ${y_1}^\perp$ or ${y_2}^\perp$ but not both. In the former case the left coordinates of the points of L are constantly y_1 while the right-hand coordinates range over the line py_2 at which the grid $G' = \langle\langle y_1, y_2\rangle\rangle$ intersects $A(p)$. In the latter case the right coordinates of all points of L are constantly y_2 while the right coordinates range over the line py_1.

Thus all lines have the form $m \times L$ where m is a point of M and L is a line of $A(p)$, or else have the form $N \times a$ where N is a line of M and a is a point of $A(p)$.

The geometry on $M \times A(p)$ with this collection of lines is precisely the product geometry, $M \times A(p)$, and the desired isomorphism follows.

This contradicts the assumption the theorem was false, completing the proof. □

16.4.6 The Final Result for $d = 0$

Theorem 16.4.12 *Assume Γ satisfies Hypothesis (E). Then one of the following holds:*

1. *Γ is a generalized quadrangle.*
2. *$\Gamma = A \times B$, the product of two maximal singular subspaces.*
3. *Γ properly contains a symplecton which is not a grid. The members of $\mathcal{M}$ partition the points of Γ. Moreover, if $\mathcal{M}$ is the full collection of all maximal singular subspaces A of Γ with the property that $|p^\perp \cap A| = 1$ for all $p \in \mathcal{P} - A$, then in this case, every line that is not an $\mathcal{M}$-line is disjoint from some member of $\mathcal{M}$ and is itself a maximal singular subspace of Γ.*

Proof As before, let U be the set of points of Γ which lie in a unique member of $\mathcal{M}$.

Step 1. *If $\mathcal{P} \neq U$, then either (1) Γ is a generalized quadrangle, or (2) every symplecton of Γ is a grid.*

Suppose $\mathcal{P} \neq U$. Suppose by way of contradiction that Γ is not a generalized quadrangle and that there exists a symplecton S which is not a grid.

Now by Theorem 16.4.10, if S is disjoint from no member of $\mathcal{M}$, then either (1) Γ is a generalized quadrangle, (2) S is a grid, or (3) $\mathcal{P} = U$. But any of the conclusions (1), (2), or (3) goes against the suppositions of the previous paragraph.

Thus, we must assume that there is a singular subspace $M \in \mathcal{M}$ such that $S \cap M = \emptyset$. Then, since S is not a grid, restriction of the projection mapping into M produces an embedding $f : S \to M$. In that case the hypotheses of Corollary 16.4.8 are in place, forcing us to conclude that $\mathcal{P} = U$, contrary to the hypothesis of the Step 1.

This contradiction completes the proof of this Step 1.

Step 2. *If $\mathcal{P} \neq U$, then either (1) Γ is a generalized quadrangle or (2) Γ is a product geometry $A \times B$ of two singular subspaces A and B.*

By Step 1, either (1) holds or all symplecta are grids. But by Theorem 16.4.11, the latter case forces conclusion (2) above.

Now we can complete the proof of Theorem 16.4.12. Suppose that the first two conclusions, 1 and 2 of this theorem fail – that is, Γ is not a generalized quadrangle nor is it a product of two maximal singular subspaces. Then, from the contrapositive of the statement of Step 2, $\mathcal{P} = U$.

Let us assume now that $\mathcal{M}$ is the full collection of all maximal singular subspaces A with the property that every point outside A is collinear with exactly one point of A. Since $\mathcal{P} = U$ (previous paragraph), the elements of $\mathcal{M}$ partition the points.

Suppose L is a line which is not an $\mathcal{M}$-line, but meets every member of $\mathcal{M}$ non-trivially Then every member of $\mathcal{M}$ meets L in exactly one point. Let B be a maximal singular subspace containing L. Suppose one could find a point x in $B-L$. Then x lies in a unique member M_x of $\mathcal{M}$ and we may set $M_x \cap L = \{y\}$. If $M_x = xy$ then $M_x \subset B$, against M_x being a maximal singular subspace. Now M_x and B are two distinct maximal singular subspaces whose intersection is a line – and that, in turn, would contradict the fact that all symplecta have polar rank two. Thus we see that L is itself a maximal singular subspace.

If $z \in \mathcal{P} - L$, z lies in a unique element $M_z \in \mathcal{M}$, meeting L at a single point. Since L is a maximal singular subspace, $z^\perp \cap L$ contains a single point. But that fact places L into the class of maximal singlar subspaces with the property that any exterior point is collinear with exactly one of its points. Since that is how we defined $\mathcal{M}$, that would make L an $\mathcal{M}$-line, contrary to our assumptions on L.

We have shown, then, that L must be disjoint from some maximal singular subspace of $\mathcal{M}$. (Note that in that case, the possibility that L is not a maximal singular space has now re-emerged.)

Choose $M \in \mathcal{M}$ so that $M \cap L = \emptyset$. Then the set of points of M which are collinear with a point of L themselves form a line L' of M. Choose point y in L, and let M_y be the (now unique) member of $\mathcal{M}$ containing y. Since L is not an $\mathcal{M}$-line, $M_y \cap L = \{y\}$. Since $M_y \cap M = \emptyset$, there is a bijection $f : M_y \to M$ taking each point of M_y to the unique point of M collinear with it. This mapping f is an isomorphism of linear spaces. Note that $f(y) \in L'$. Thus there is a line N in M_y such that $f(N) = L'$. Now let $R = \langle\langle y, L'\rangle\rangle$, the symplecton on y and L'. Each point of $N \cup L - \{y\}$ is collinear with y and a point of $L' - \{f(y)\}$, forcing $N \cup L \subseteq R$. Thus y lies on at least three lines of R, namely L, N, and $yf(y)$, and so R is not a grid.

If L were properly contained in a singular subspace B, then $B \notin \mathcal{M}$ as L is not an $\mathcal{M}$-line. Then, choosing $x \in B - L$, and $M_x \in \mathcal{M}$ containing x we see that $M_x \cap B = \{x\}$, otherwise one has a symplecton of polar rank three or more. Then

$$\begin{aligned} M_x \cap R &\subseteq (x^\perp \cap R) \cap M_x \\ &\subseteq L \cap M_x \\ &\subseteq (B \cap M_x) \cap R \\ &= \{x\} \cap R = \emptyset. \end{aligned}$$

But the mapping $g : R \to M_x$ induced by projection on M_x is not an embedding since g maps the line L to x. By Theorem 16.4.4 R must be a grid. This contradicts the conclusion of the previous paragraph.

Thus L cannot properly lie in another singular subspace – that is, it is itself a maximal singular subspace. □

16.5 Proof of Theorem 16.2.1

We assume $\Gamma = (\mathcal{P}, \mathcal{L})$ satisfies Hypothesis (A) (see p. 554). Then the basic trichotomy of Theorem 16.2.5 produces three cases: (1) $d = 0$, (2) $d = 1$, and (3) every member of $\mathcal{M}$ has projective dimension $d + 1$ and is a maximal singular subspace of some symplecton. In case (3), Corollary 16.2.7 implies that Γ is itself a polar space. So from this point onward, we assume axiom (A3) that Γ is not a polar space, and hence that every element of $\mathcal{M}$ has projective dimension exceeding $d+1$ and that d is either 0 or 1.

In particular, if $d = 1$, hypothesis (H1) of p. 558 holds. Then by Theorem 16.3.13, Γ is the Grassmannian of lines of some projective space $P(V)$ – possibly of infinite dimension.

On the other hand, if $d = 0$, then hypothesis (E) of p. 564 must hold, for this hypothesis merely rehashes hypotheses (A) and (A3) for this value of d. Then the conclusion described in Theorem 16.2.1 for this case is fully implied by Theorem 16.4.12. Thus all parts of Theorem 16.2.1 have been proved. □

16.6 A Corollary of Theorem 16.2.1

This section contains a proof of the following corollary.

Corollary 16.6.1 (First corollary.) *Suppose Γ is a parapolar space with a class of maximal singular subspaces $\mathcal{M}$ with these properties:*

(1) For every line, there exists a member of $\mathcal{M}$ containing that line.
(2) There exists a positive integer $d \geq 1$ such that whenever $(x, M) \in \mathcal{P} \times \mathcal{M}$ is a non-incident pair, then $x^\perp \cap M$ is either empty or is a projective space of projective dimension d.
(3) At least one member of $\mathcal{M}$ possesses a finite unrefinable chain of subspaces.

Then one of the following statements holds:

1. Γ *is a non-degenerate polar space of rank* $d+1$.
2. $d=2$ *and* Γ *is a homomorphic image of a half-spin geometry (necessarily of finite singular rank).*
3. $d=1$ *and one of the following hold:*

   (a) Γ *is the Grassmannian of all* k*-dimensional subspaces of a projective space* $\mathbf{P}(V)$*, possibly of infinite rank, and* k *is a finite integer greater than* 1.
   (b) *All maximal singular subspaces of* Γ *have finite rank and* Γ *is isomorphic to the factor geometry* $A_{2n-1,n}(D)/\langle\sigma\rangle$*, where* σ *is a polarity of Witt index at most* $n-5$.
   (c) Γ *is parapolar space of symplectic rank exactly three for which, every line* L *lies in a unique member* M_L *of* $\mathcal{M}$ *which is of finite projective dimension* $\delta>2$*, and in at least two other maximal singular subspaces – all of the latter maximal singular subspaces being planes. Thus* Γ *has finite singular rank as well. Moreover, each line lies in a proper rank three symplecton which is not oriflame. For each plane* π *that is not contained in a member of* $\mathcal{M}$ *and for each point* p *in* π*, there exists a member of* $\mathcal{M}$ *meeting* π *exactly at point* p.

Remark The last Case 3(c) of the conclusion of the Corollary does not appear to have any known examples living in it. Any geometry in this class must contain a non-oriflame symplecton S of polar rank three containing a class $\mathcal{M}_S$, of maximal singular subspaces with the property that each line of S lies in a unique member of $\mathcal{M}_S$. One will note the $\mathcal{M}_S$ is in fact an ovoidal hyperplane of the dual polar space S^D associated with S. Moreover, S is embedded in Γ so that the members of $\mathcal{M}_S$ lie in larger singular subspaces while all other maximal singular subspaces of S are planes which are already maximal singular subspaces of Γ.

There are many hypotheses that might eliminate the troublesome last Case 3(c) when $d=1$. Any one of the following conditions will accomplish this:

(1) Some line lies in two or more members of $\mathcal{M}$.
(2) All symplecta are oriflame.
(3) Some maximal singular space not in $\mathcal{M}$ properly contains a plane.

Before beginning with the proof of Corollary 16.6.1, a little local preparation is in order.

16.6.1 Uniformity of Point-Residuals

Our purpose here is to identify one of the three conclusions of Theorem 16.4.12 above, by a property of any one of its points. The reason for doing this will emerge in the next section.

Lemma 16.6.2 *We assume that Γ is a gamma space with all singular subspaces projective spaces. Assume that for every point p, the point-residual $Res_\Gamma(p) := (\mathcal{L}_p, \Pi_p)$ satisfies one of the following statements:*

(1) $Res_\Gamma(p)$ is a generalized quadrangle which is not a grid.
(2) $Res_\Gamma(p)$ is the product geometry $A \times B$ of two maximal singular subspaces A and B.
(3) Some symplecton of $Res_\Gamma(p)$ is not a grid, and there is a class $\mathcal{M}$ of maximal singular subspaces which partition the points $Res_\Gamma(p)$. Any line which is not contained in a element of $\mathcal{M}$ is a maximal singular subspace. No line intersects all members of $\mathcal{M}$.

The choice depends on the particular point p.

Let us define three sets of points:

X_1: The set of all points p which lie in some line L with the property that there are at least three maximal singular subspaces containing L, each one of which is a plane.
X_2: The set of all points p which lie in some line L which lies in exactly two maximal singular subspaces.
X_3: The set of points p which lie on some line L such that:

(a) L lies in at least three maximal singular subspaces.
(b) L lies in a unique member M_L of $\mathcal{M}$.
(c) M_L has projective rank at least three.

Then we have a partition $\mathcal{P} = X_1 + X_2 + X_3$, where each set X_i is a union of connected components of the point-collinearity graph $\Delta = (\mathcal{P}, \sim)$.

Proof Suppose $p \in X_1$ and q is another point collinear with p. Now p lies on a line L for which every maximal singular subspace containing L is a plane. Now $\mathrm{Res}_\Gamma(p)$ cannot satisfy hypotheses (2) or (3), and so all singular subspaces of Γ on p are planes. In particular, the line pq has only planes for maximal singular subspaces containing it, so $q \in X_1$.

Suppose $p \in X_2$ and q is another point collinear with p. Then p lies on a line L lying on exactly two maximal singular subspaces. The operative hypothesis is (2), so $\mathrm{Res}_\Gamma(p)$ is a product geometry, all of whose "points" lie in just two maximal singular subspaces. In particular, the line pq is in exactly two maximal singular subspaces, and so $q \in X_2$.

Suppose $p \in X_3$ and q is another point collinear with p. Then p lies on a line L lying in a unique singular subspace M_L of $\mathcal{M}$, with all other singular subspaces on L, of which there are at least two, being planes. By our hypothesis $\mathrm{Res}_\Gamma(p)$ belongs to case (3), and so every line on p has this property – in particular line pq lies in a unique member of $\mathcal{M}$, with all other maximal singular spaces being planes. Thus $q \in X_3$.

Thus for all i, if $p \in X_i$ and q is a point collinear with p, then $q \in X_i$. The proof is complete. □

16.6.2 The Proof of Corollary 16.6.1

Suppose Γ satisfies all of the hypotheses of Corollary 16.6.1. By hypothesis, Γ is a parapolar space. It is easy to see that if (x, p, y) is a geodesic path of length two, then there is a singular space $M \in \mathcal{M}$ containing the line yp, and so, since $d \geq 1$, $x^{\perp} \cap y^{\perp}$ contains a line whenever $d_{\Gamma}(x, y) = 2$. It follows that Γ has symplectic rank at least three. Thus by Theorem 13.4.1 on p. 463 of Chap. 13, we have the following.

Step 1. *Every singular subspace of* Γ *is a projective space.*

Now let p be an arbitrary point of Γ. Then for each line pq and element $M \in \mathcal{M}$ for which q is not in M, we have $q^{\perp} \cap M \simeq PG(d, q)$ for $d \geq 1$. This means that the collection $\mathcal{M}_p$ of all members of $\mathcal{M}$ incident with p produces a class $\bar{\mathcal{M}}_p$ of maximal singular spaces with respect to which hypothesis (A) of Sect. 16.2 (p. 554) holds. Thus we have the following.

Step 2. *For each point* p*, the residue geometry* $Res_{\Gamma}(p) := (\mathcal{L}_p, \Pi_p)$*, satisfies axiom (A2) for parameter* $d' := d - 1 > 0$*, and so is a strong parapolar space of diameter two. Moreover, since the symplecta of* $Res_{\Gamma}(p)$ *are themselves point-residues of polar spaces of rank at least three, they are classical quadrangles or classical polar spaces of rank three or more. Thus at this point* $Res_{\Gamma}(p)$ *satisfies hypotheses (A1) and (A2). Thus by Theorem 16.2.1,* $Res_{\Gamma}(p)$ *is one of the following:*

(1) A polar space ($\bar{\mathcal{M}}_p$ *is any class of maximal singular subspaces that cover all the polar points).*
(2) A Grassmannian of lines ($\bar{\mathcal{M}}_p$ *is the class of singular subspaces* M_x *consisting of all lines on a projective point* x*).*
(3) A product $P \times Q$ *of projective spaces (where* $\bar{\mathcal{M}}_p$ *is a class containing at least one of the two point-partitioning classes of singular subspaces* $\mathcal{P}^* := \{(P, y) | y \in Q\}$ *or* $\mathcal{Q}^* := \{(y, Q) | x \in P\}$*).*
(4) A rank two parapolar space with these properties: Some symplecton of $Res_{\Gamma}(p)$ *is not a grid, and the members of* $\mathcal{M}_p$ *pairwise intersect at* $\{p\}$ *and so the class* $\bar{\mathcal{M}}_p$ *partitions the residual points. Any residual line (that is, a plane of* Γ *on* p*) which is not contained in an element of* $\mathcal{M}_p$ *is a maximal singular subspace of* $Res_{\Gamma}(p)$*. No such residual line intersects all members of* $\mathcal{M}_p$ *at residual points — that is, no plane on* p *which is not in a member of* $\mathcal{M}_p$ *can intersect each member of* $\mathcal{M}_p$ *at a line.*

In cases (1), (2) and (4), for each residue $\mathrm{Res}_{\Gamma}(p)$, the members of $\bar{\mathcal{M}}$ are isomorphic as projective spaces, and in case (4), there exists a subcollection of pairwise isomorphic members of $\bar{\mathcal{M}}$ covering all the residual points.

Now by Step 2, Γ is a strong parapolar space of symplectic rank at least three. Thus by Lemma 13.5.1, p. 467 of Chap. 13 we have the following.

Step 3. Γ *is locally connected.*

Now suppose for some point p of Γ, that the residue $\mathrm{Res}_\Gamma(p)$ is a polar space of rank at least two. Since Γ is locally connected, we can now apply Theorem 13.5.5 of p. 471 to obtain the following.

Step 4. *If some point-residue of* Γ *is a polar space of rank at least two then* Γ *is itself a polar space, one of our desired conclusions. Thus we may assume from this point onward that no residue* $Res_\Gamma(p)$ *is a polar space.*

Now, by hypothesis, at least one member of $\mathcal{M}$ is a projective space of finite rank n. So if we choose a point $p \in M$, we see by the last part of Step 2, that every residual point of $\mathrm{Res}_\Gamma(p)$ lies in a member of $\bar{\mathcal{M}}_p$ of finite projective dimension. In other words, any point q collinear with p lies on a member of $\mathcal{M}$ of finite projective rank. Since the point-collinearity graph of any parapolar space is by definition connected, we conclude the following.

Step 5. *Every point of* Γ *lies in a member of* $\mathcal{M}$ *of finite projective rank* .

At this point, Γ is a gamma space, each of whose point-residues $\mathrm{Res}_\Gamma(p)$ must satisfy one of the cases (2) – (4) of Step 2, depending on the choice of p. In each of these cases, $\mathrm{Res}_\Gamma(p)$ possesses a collection $\bar{\mathcal{M}}_p$ of maximal singular subspaces (induced by the collection $\mathcal{M}_p$) with the property that for some integer d, each residual point not in a member $\bar{M}$ of $\bar{\mathcal{M}}$ is collinear with a $PG(d')$-subspace of M, where $d = d' + 1$, and $d' = 0$ or 1 – that is $d = 1$ or 2. This d does not depend upon the point p.

So, we consider first the case that $d = 2$. In this case $\mathrm{Res}_\Gamma(p)$ is the Grassmannian of lines of some projective space $P(V)$. But in this case, the class $\bar{\mathcal{M}}$ is the class of singular subspaces M_x consisting of all lines on a projective point x of $P(V)$. Since elements of M are assumed to have finite projective rank, we see that $\mathrm{Res}_\Gamma(p) \simeq A_{n,2}(D)$ for a division ring D independent of the choice of point p and some positive integer n. It follows that Γ is a locally truncated geometry belonging to the locally truncated diagram presented in Fig. 16.3.

We thus conclude the following.

Step 6. *For the parapolar space* Γ*, the parameter* d *is either* 2 *or* 1. *If* $d = 2$*, then* Γ *is the homomorphic image* $D_{n,n}/A$ *of a half-spin geometry where* A *is a*

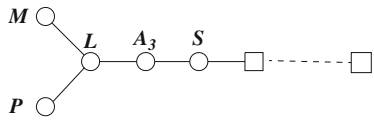

Fig. 16.3 The locally truncated diagram which results when $d = 2$. Aside from points, lines, and symplecta, the symbol $\mathbf{A}_3$ refers to the class of maximal singular $PG(3, D)$s which do not belong to $\mathcal{M}$

chamber-semiregular group of automorphisms of the $D_{n,1}(D)$-polar space such that $\dim N/(N \cap N^a) \geq 5$*, for every maximal singular polar subspace N corresponding to the node labeled $\mathcal{M}$ in Fig. 16.3 and every non-identity element $a \in A$.*

Since this corresponds to one of our conclusions, we may assume henceforward that $d = 1$, and Γ is a parapolar space of symplectic rank three exactly.

At this point, Γ is a gamma space each of whose point-residuals is either in case (3) or case (4) of Step 2. Thus every point belongs to the class X_2 or X_3 described in Lemma 16.6.2. Noting that the parapolar space is connected, we can now invoke that same Lemma 16.6.2 to conclude the following.

Step 7. *The parameter $d = 1$ and one of the following two cases must hold.*

Subcase 1: For every point p, $Res_\Gamma(p)$ is a product space $P \times Q$ at least one of whose partitions P^* or Q^* have components which are isomorphic projective spaces of finite rank r which belong to $\mathcal{M}$ (the parameter r is independent of point p).

Subcase 2: For every point p, $Res_\Gamma(p)$ has the properties described in case (4) of Step 2. All members of $\mathcal{M}$ are pairwise isomorphic as finite-dimensional projective spaces.

Let us consider Subcase 1. For each point p, the symplecta of $\mathrm{Res}_\Gamma(p) \simeq P \times Q$ are just grids $L_P \times L_Q$ where L_P and L_Q are typical lines of P and Q, respectively. Moreover, each residual point not in such a grid G is collinear with either all points of a line of the grid, or no point of G,

The statements of the previous paragraph have the following consequences for Γ. All symplecta of Γ are oriflame of type D_3. If S is any such symplecton, and x is a point not in S, then $x^\perp \cap S$ is either the empty set, a single point, or a plane of S – it is never a line. This last statement is condition (Co2) of Cohen's three conditions listed in Sect. 15.3.2 on p. 530. We already have on hand the other two conditions listed: (Co1) asserting that Γ has symplectic rank exactly three, and (Co3) asserting that Γ itself is not a polar space. Since Γ is also a strong parapolar space we may apply Theorem 15.3.7 of Chap. 15 to conclude that one of the following holds:

1. Γ is a Grassmann space of k-subspaces of a vector space V over a division ring D, where k is a finite integer exceeding 1. (V may be infinite-dimensional.)
2. $\Gamma \simeq A_{2n-1,n}(D)/\langle\sigma\rangle$, a quotient of a Grassmann space as the orbit geometry of a polarity σ of $PG(2n-1, D)$ of Witt index at most $n-5$.

Thus we have shown the following.

Step 8. *In Subcase 1 of Step 7, alternatives (a) and (b) of conclusion 3 of Corollary 16.6.1 hold.*

Thus we are left with Subcase 2. But here, the description of all point-residuals $\mathrm{Res}_\Gamma(p)$ as given in Step 2 (2) also can be translated into assertions about Γ. The result is as follows.

(Step 9) *In Subcase 2 of Step 7, alternative (c) of conclusion 3 of our corollary holds.*

By Steps 7, 8, and 9, the proof of Corollary 16.6.1 is complete.

16.7 A Further Corollary of Theorem 16.2.1

One is now in position to prove the following.

Corollary 16.7.1 *Suppose* $\Gamma = (\mathcal{P}, \mathcal{L})$ *is a locally connected parapolar space of symplectic rank at least four. Assume that there is a class* $\mathcal{M}$ *of maximal singular subspaces of* Γ *with these properties:*

(1) Every projective plane lies in some member of $\mathcal{M}$.
(2) All members of $\mathcal{M}$ *have finite projective rank, and there exists an integer* d, *greater than* 1, *such that if* (x, M) *is any non-incident pair drawn from* $\mathcal{P} \times \mathcal{M}$, *then* $x^{\perp} \cap M$ *is either empty, consists of a single point, or is a* $PG(d)$*-subspace of* M.
(3) In case $d = 2$, *any maximal singular subspace which is not in* $\mathcal{M}$ *is assumed to have projective dimension at least* 4.

Then one of the following statements holds for Γ*:*

(a) Γ *is a polar space of symplectic rank* $d + 1$.
(b) The parameter $d = 3$ *and* Γ *is the truncation to points and lines of a homomorphic image* B/A *of a building geometry* B *belonging to the diagram* $Y_{2,1,m}$ *in Fig. 16.4. (The nodes designating the points and lines of the truncation are so labeled in the figure, as is the node representing the class* $\mathcal{M}(B)$ *of maximal singular subspaces of* $B|_{\mathcal{P},\mathcal{L}}$ *which map onto the class* $\mathcal{M}$ *under the morphism* $\phi : B \to B/A$.*)*
(c) $d = 2$ *and* Γ *is the truncation to points and lines of a homomorphic image* B/A *of a building geometry* B *belonging to the diagram* $Y_{n,1,m}$ *in Fig. 16.5 (nodes representing points and lines being labeled).*
(d) Again, $d = 2$. *There exists a building* B *belonging to the diagram* $Y_{n,1,n}$, *(just as in part (c) except that* $n = m$ *and* $n > 4$*), and an automorphism* σ *of* B *inducing the obvious involution on the* $Y_{n,1,n}$*-diagram so that* Γ *is the truncation to points and lines of a homomorphic image of the factor geometry* $B/\langle\sigma\rangle$.

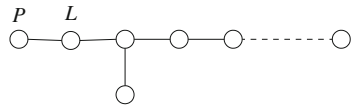

Fig. 16.4 The $Y_{2,1m}$-diagram. The nodes representing the points and lines of the eventual truncation are indicated. The class of maximal singular subspaces $\mathcal{M}(B)$ of $B|_{\mathcal{P},\mathcal{L}}$ which map onto the class $\mathcal{M}$ under the morphism $\phi : B \to B/A$ are represented by the terminal node of the shortest arm

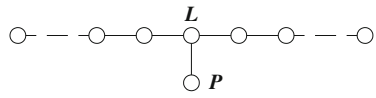

Fig. 16.5 The $Y_{n,1m}$-diagram. The nodes representing the points and lines of the eventual truncation are indicated

For the morphism $\phi : B \to B/A$ in cases (b) and (c), and the induced morphism on the truncation in case (d) – that is the mapping

$$trunc \circ \phi : (B/\langle\sigma\rangle)|_{\mathcal{P},\mathcal{L}} \to \Gamma$$

– the fiber above any point must consist of "points" of the domain geometry at distance at least five from one another.

Remark In one sense Corollary 16.7.1 is a little "mythical" in that many of the buildings from which the conclusion draws its truncation of homomorphic images are not spherical or even affine. The buildings exist, but are they parameterized by the simply-laced diagram and an admissible division ring? One would suspect so. However, one notices that certain Lie incidence geometries of spherical type make their appearance among the conclusions. Using the Bourbaki indexing of nodes these geometries are

$$E_{6,1},\ E_{7,1},\ E_{8,1},\ E_{6,2},\ E_{7,2},\ E_{8,2}.$$

The two examples involving the E_8-building do not seem to be covered by any other theorem of this book.

The proof offered here invokes a number of methodological principles developed in Part III, and for this reason it is hoped that the author will be forgiven for splicing in informal (hopefully instructive) comments emphasizing the following issues:

- Dealing with local partitions of a class of subspaces (Chap. 12).
- Establishing the existence of a sheaf (Chap. 11).
- Dealing with residues that are themselves homomorphic images of buildings, rather than buildings.
- Dealing with residues which themselves possess globally fused systems of singular subspaces ($A_{2n-2,n}/\langle\sigma\rangle$, for example).
- Lifting a non-type-preserving automorphism of a locally truncated geometry to its covering building.

Throughout this informal proof, we write $\mathcal{M}_p$ and $\mathcal{M}_L$ for the subcollection of singular spaces in $\mathcal{M}$ which are incident with a point p or a line L, respectively.

First, one notes that the local connectedness of Γ forces each point-residue $\mathrm{Res}_\Gamma(p)$ to be a strong parapolar space of symplectic rank at least three satisfying the hypotheses of Corollary 16.6.1, with the "$\mathcal{M}$" of that corollary replaced by the

collection of singular subspaces $\mathcal{M}(p)$ of $\mathrm{Res}_\Gamma(p)$ induced by $\mathcal{M}_p$ and the parameter "d" of that corollary replaced by $d-1$. It follows that $\mathrm{Res}_\Gamma(p)$ must satisfy one of the alternatives 1, 2, 3(a), 3(b), or 3(c) of the conclusion of Corollary 16.7.1.

If conclusion 1 holds for some residue $\mathrm{Res}_\Gamma(p)$, then that residue is a polar space of rank at least three, and local connectedness can be invoked to conclude that Γ itself is a polar space via Theorem 13.5.5 of p. 471. We also note that conclusion 3(c) is ruled out for any point-residual by hypothesis (3). The remaining possibilities for $\mathrm{Res}_\Gamma(p)$ are divided into two cases: $d-1=2$ and $d-1=1$. But if a point belongs to one of these cases, it is clear that any point collinear with it belongs to the same case. Thus by the connectedness of Γ only one of these cases prevails for all points. Thus we have the following.

Step 1. *We may assume one of the following:*

1. *Case 1: $d=3$ and for every point p the residue $Res_\Gamma(p)$ is a homomorphic image of a half-spin geometry $D_{n,n}$ where $n-1$ is the projective dimension of each member of $\mathcal{M}$.*
2. *Case 2: $d=2$ and for any given point p, the residue $Res_\Gamma(p)$ is either (1) a Grassmann space of type $A_{m+n-1,n}$, $m>0$ or (2) the geometry $A_{2n-1,n}/\langle\sigma_p\rangle$ where σ_p is a polarity of $PG(2n-1)$ of Witt index at most 5 (so $n\geq 5$). Again, $n-1$ is the common projective dimension of members of $\mathcal{M}$, but the two possibilities (1) and (2) may conceivably vary from point to point.*

Consider first Case 1 of Step 1, in which the residue of any point p is a homomorphic image $D_{n,n}/A_p$ of a half-spin geometry by some chamber-semiregular group of automorphism A_p which may well depend on the point p. There is just one requirement: the fact that $D_{n,n}/A_p$ be a parapolar space forces any A_p-orbit on half-spin points to consist of points which pairwise are at distance at least five. We assert the following.

Step 2. *In Case 1 of Step 1, Γ can be enriched to a rank six locally-truncated geometry Γ^* associated with the diagram of Fig. 16.6, where $\mathcal{P}$, $\mathcal{L}$, and $\mathcal{M}$ have been defined for Γ, $\mathcal{S}$ is the class of symplecta for this parapolar space, $\mathcal{A}_4$ is a class of maximal singular subspaces which lie in some symplecton, and $\mathcal{A}_2$ is the class of planes that are formed from an intersection $M\cap A$ for some incident pair $(A,M)\in\mathcal{A}_4\times\mathcal{M}$ (a flag of Γ^*).*

Comment The student should realize that Step 2 is not just a "slam dunk." Of course there is no problem with the enrichment since the symplecta and prescribed singular spaces derive their existence from the fact that Γ is a parapolar space and that point-residuals $\mathrm{Res}_\Gamma(p)=(\mathcal{L}_p,\Pi_p)$ have the structure prescribed by Step 1, Case 1.

The problem lies in the worry that homomorphic images of half-spin geometries may not be residually connected geometries and that there may be a bizarre behavior among flags which defy the locally truncated geometry

requirements concerning the relation of Δ-residues of flags forming a chain $F_1 \subset F_2$.

For the morphism $f_p : D_{n,n} \to D_{n,n} \simeq A_p$, there is just one bad scenario. If $F := (X, M)$ is a flag consisting of a half-spin point X, and an element M of $\mathcal{M}_p$ (objects of types n and $n-1$ in the building diagram), then the residue in $D_{n,n}/A_p$ of both $f_p(X)$ and $f_p(M)$ are $PG(n-1)$s, while the residue of the flag $f_p(F)$ is the disjoint union of isolated points and $f_p(\mathrm{Res}_{D_{n,n}}((F)) \simeq PG(n-2)$. Thus the morphism does not surjectively map residues in $D_{n,n}$ to their corresponding residues of $D_{n,n}/A_p$.[1] (A specific example of this type was discussed in Sect. 11.9.2, p. 438.)

So, translating this phenomenon in terms of the original objects of Γ, the half-spin point X becomes a line on p and M becomes a member of $\mathcal{M}_p$. We obtain a flag $F_2 := \{p, X, M\}$ of Γ, and setting $F_1 = \{p, X\}$ we seem to fail to have

$$\mathrm{Res}(F_2) = \mathrm{Res}_{\mathrm{Res}(F_1)}(F_2 - F_1)$$

as required for a locally truncated geometry.

However, when we described $\mathrm{Res}(F_2)$ as the disjoint union of several isolated points and a projective space (as in the example), we were referring to isolated *polar points*, not *half-spin points*. In other words, the extra isolated "points" are really a collection of images of half-spin subspaces corresponding to the terminal "ghost node" at the right end of the diagram of Fig. 16.6. What we have is that the $\mathrm{Res}_{\Gamma^*}(F_2)$ is a rank three geometry which is still the truncation to projective, lines, and planes of a $PG(n-2)$. So in this case, we are safe, and have Step 2.]

At this point, we obtain a diagram like that shown in Fig. 16.7, which fits into the mold of Fig. 11.2 and the context of Theorem 11.3.2, p. 427, relative to the type set J which includes $\{\mathcal{L}, \Pi, \mathcal{M}\}$ as well as a class of $PG(4)$s, $\mathcal{A}_4$. In order to apply Theorem 11.3.2 so that we may achieve a sheaf, we must check the two "rigidity" conditions. The right residue of a flag $F = (L, M)$ of type $\{\mathcal{L}, \mathcal{M}\}$ is a $PG(n-2)$ which is $\{\Pi, \mathcal{A}_4\}$-rigid since these types represent the points and lines of this projective space. Similarly, for a flag $G = (A, M)$ of type $\{\mathcal{A}_4, \mathcal{M}\}$, the left residue in Γ is a $PG(3)$ which is $\{\mathcal{L}, \Pi\}$-rigid. Thus by Theorem 11.3.2, a sheaf $\mathcal{F}$ exists.

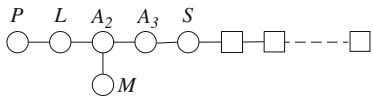

Fig. 16.6 The locally truncated diagram for the rank six geometry Γ^* arising in Case 1 of Step 1. (For technical reasons, the class of planes Π is denoted $\mathcal{A}_2$ in the figure)

[1] Contrary to what seems to be asserted in [10].

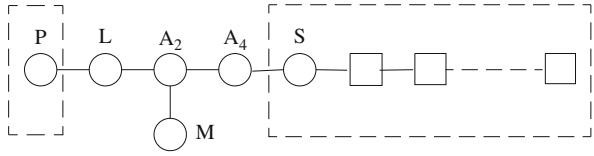

Fig. 16.7 The locally truncated diagram of Step 2. (Again, the class of planes Π is denoted $\mathcal{A}_2$ in the figure)

To see whether this sheaf is residually connected, we need only check its values on the rank-1 flags whose type is in J. But all of these are buildings (mostly with disconnected diagrams). For $(\pi, A, L, M) \in \Pi \times \mathcal{A}_4 \times \mathcal{L} \times \mathcal{M}$, we have:

$$\mathcal{F}(\pi) \simeq PG(n-3) * PG(1) * PG(2), \tag{16.3}$$

$$\mathcal{F}(A) \simeq PG(4) * PG(n-4), \tag{16.4}$$

$$\mathcal{F}(L) \simeq PG(n-1) * PG(1), \tag{16.5}$$

$$\mathcal{F}(M) \simeq PG(n-1), \tag{16.6}$$

where $X * Y$ means every object in geometry X is incident with every object in geometry Y.

So we can employ the machinery of Chap. 11 to conclude with the following.

Step 3. *In Step 2, the locally truncated geometry* Γ^* *of rank six over the locally truncated diagram of Fig. 16.6 possesses a residually connected sheaf. From the theory of Chap. 11,* Γ^* *is the truncation to typeset* $J^* := \{\Pi, \mathcal{A}_4, \mathcal{P}, \mathcal{L}, \mathcal{M}, \mathcal{S}\}$ *of a homomorphic image of a building associated with the diagram of Fig. 16.8.*

Since this step yields part of our conclusion, we may pass on to Case 2 of Step 1.

In this case, for each point p, there are two possibilities for the residue; but for each possibility, each plane on p lies in exactly two maximal singular subspaces (a $PG(n_p)$ and a $PG(m_p)$) and is the only point of their intersection. (Here, parameters n_p and m_p depend only on p, not the choice of the plane.) Thus if point q is collinear with p (but distinct from it), then there is a plane π on the line $L = pq$ with these properties, so we see that $\{n_q, m_q\} = \{n_p, m_p\}$.

Suppose first that $n_p \neq m_p$. Then the theory expounded in Chap. 12 shows that Γ contains two distinct classes of maximal singular subspaces $\mathcal{A}$ and $\mathcal{B}$ (one

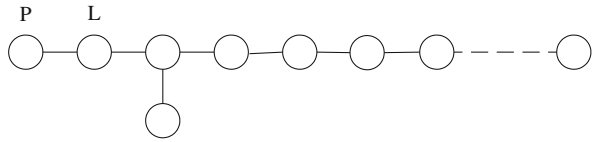

Fig. 16.8 The diagram of the building whose truncation covers the geometry Γ^* of Step 3

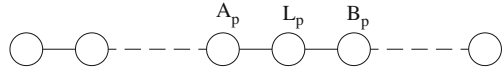

Fig. 16.9 Diagram of $\mathrm{Res}_{\Gamma^*}(p)$

of which – say $\mathcal{B}$ – we can take to be $\mathcal{M}$) such that for each point p, the residue $\mathrm{Res}_\Gamma(p)$ is the Grassmann space $A_{n+m,n}$ with diagram as in Fig. 16.9.

There are no homomorphic images to contend with in these residues and so we obtain a locally truncated rank four geometry Γ^* belonging to the locally truncated diagram of Fig. 16.10 with the extra information that certain "ghost flags" constitute the symplecta of type D_4. The right and left rigidity conditions hold since the relevant truncations of the left and right rank two residues are truncations to points and lines of projective spaces. Thus a residually connected sheaf exists and the theory of Chap. 11 shows us that Γ is the truncation to points and lines of a homomorphic image of a building B belonging to the Coxeter diagram $Y_{n-1,1,m}$. This is contained in conclusion (c) of our corollary. Consequently we have the following.

Step 4. *In Case 2 we may assume that all maximal singular subspaces have the same projective dimension n, and so by assumption (3) every point p lies in a plane which is the intersection of two maximal singular $PG(n)$s belonging to n. Point-residuals are now as described in Case 2 of Step 1 with $m = n$.*

Remark In Γ we have two definite objects, points ($\mathcal{P}$) and lines ($\mathcal{L}$), but we cannot proceed to the construction of a locally truncated diagram which would include the maximal singular subspaces since it is not clear whether there is one or two classes of such objects. What we shall do is change our "point-line" point of view to one in which the lines of Γ are regarded as "points."

Set $\mathcal{P}_0 := \mathcal{L}$, our class of "points." Our "lines," $\mathcal{L}_0$, will be the pencils of Γ-lines incident with some point-plane flag of Γ. Then:

(1) $\Gamma_0 := (\mathcal{P}_0, \mathcal{L}_0)$ is a gamma space.
(2) If M is a maximal singular subspace of Γ, then M becomes a subspace $\mathcal{G}(M)$ of Γ_0 which is the Grassmannian of lines of a $PG(n+1)$.
(3) In Γ_0, the residual of a "point" P_0 is a product geometry

$$PG(n-1) \times PG(n-1) \times L,$$

where L is a residual line.

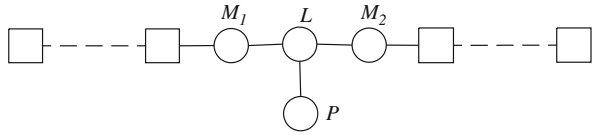

Fig. 16.10 The locally truncated diagram for Γ^* defined for Case 2 of Step 1

(4) If p is a point of Γ, then p corresponds to a subspace $A(p)$ of Γ_0 whose "point-line" shadow in Γ_0 is isomorphic to $\mathrm{Res}_\Gamma(p)$, namely $A_{2n-1,n}$ or $A_{2n-1,n}/\langle\sigma_p\rangle$.

Now, by (3), $\Gamma_0 = (\mathcal{P}_0, \mathcal{L}_0)$ satisfies axiom (L) of p. 446 – that is, the collection Ξ of all Grassmannians of type $A_{n+1,2}$ induce on each point p partitions $\Xi_p = \mathcal{A}_p + \mathcal{B}_p$ satisfying axioms (R1) and (R2) of Chap. 12, p. 445. Now by Theorem 12.2.1 on p. 446 there exist two possibilities which we present as Step 5.

Step 5. *Let Γ_0 and Ξ be defined as in the preceding two paragraphs.*

(1) The family Ξ of Grassmann subspaces of Γ_0 has a partition $\Xi = \mathcal{A} + \mathcal{B}$ such that for any "point" P of Γ_0, $\Xi_P = \mathcal{A}_p + \mathcal{B}_p$.

(2) There exists a morphism of connected point-line geometries

$$\phi : \hat{\Gamma}_0 = (\hat{\mathcal{P}}_0, \hat{\mathcal{L}}_0) \to (\mathcal{P}_0, \mathcal{L}_0) = \Gamma_0$$

which is an isomorphism when restricted to the residue of any $\hat{\Gamma}_0$-point or line. Moreover, $\hat{\Gamma}_0$ possesses two families $\hat{\Sigma}$ and $\hat{\Pi}$ of line-Grassmannian subspaces satisfying axioms (S1) and (S2) of p. 444, such that for each "point" $P \in \mathcal{P}_0$,

$$\phi(\hat{\Sigma})_P = \mathcal{A}_P, \text{ and } \phi(\hat{\Pi})_P = \mathcal{B}_P.$$

The morphism ϕ is surjective on "points" and "lines," and the fiber above each of these contains exactly two points or two lines of $\hat{\Gamma}_0$, respectively.

Now let Γ_{00} be Γ_0 in Case 1, and let it be $\hat{\Gamma}_0$ in Case 2. Then $\Gamma_{00} = (\mathcal{P}_{00}, \mathcal{L}_{00})$ is a gamma space with an enrichment Γ^*_{00} belonging to the locally-truncated diagram of Fig. 16.11. Here $(\mathcal{G}_1, \mathcal{G}_2)$ is $(\mathcal{A}, \mathcal{B})$ or $(\hat{\Sigma}, \hat{\Pi})$ in cases (1) and (2), respectively. The class $\mathcal{A}_{00}$ is a bona fide Grassmann space of type $A_{2n-1,n}$, which is mapped by ϕ to the set of points $\mathcal{P}$ of Γ. The left residue of an element $G_2 \in \mathcal{G}_2$ belongs to the locally truncated diagram of Fig. 16.12 and is the truncation of $A_{n+1,2}$ to types 1, 2, and 3 (in the Bourbaki numbering of the A_{n+1}-diagram), so that lines (that is, (1, 3)-flags) are included. Of course the same statement holds with the indices 1 and 2 transposed. It is then easy to verify the appropriate rigidity conditions and so, by the theory of Chap. 11, to conclude the following.

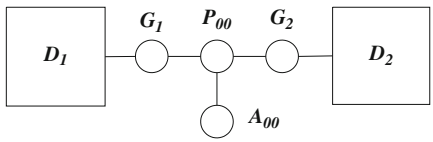

Fig. 16.11 The locally truncated diagram for an enrichment of Γ_{00}

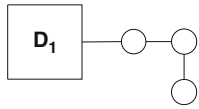

Fig. 16.12 The locally truncated diagram of a left residue of an element G_2 of $\mathcal{G}_2$

Step 6. *The geometry Γ_{00}^* is a homomorphic image of the truncation $B|_T$ of a building B belonging to the diagram $Y_{n-1,1,n-1}$, where the typeset of the truncation, T, corresponds to the branch-point of the "Y" diagram together with its three immediate neighbors. In Case 1 of Step 5, $\Gamma_{00} = \Gamma_0$, and so, re-labeling the types as subspaces of $\Gamma = (\mathcal{P}, \mathcal{L})$, we obtain a special case of conclusion (c) of the corollary.*

Thus we may assume henceforward that Case 2 of Step 5 holds – that $\Gamma_0 = \phi(\hat{\Gamma_0})$. Moreover, an enrichment to subspaces of $\hat{\Gamma_0}$ provides us with a rank four geometry $\hat{\Gamma_0} = (\mathcal{P}_0, \mathcal{A}_{00}, \mathcal{G}_1, \mathcal{G}_2)$ which is a homomorphic image of a truncation $B|_T$ as in the previous paragraph.

But it still remains to be shown that Γ is the truncation to the typeset $T_0 = (\mathcal{P}, \mathcal{L}) = (n+1, n)$ of a homomorphic image of a factor $B/\langle\sigma\rangle$ of a building B by an involutatory (non-type-preserving) automorphism σ inducing the symmetry δ of the Coxeter diagram for B as given in Fig. 16.13.

Thus, δ is a permutation of the typeset $I := \{1, 2, \ldots, 2n\}$, which fixes n and $n+1$ and acts in orbits $(1, 2n), (2, 2n-1), \ldots, (i, 2n-i+1), \ldots$ for $1 \leq i \leq n-2$. Thus $B/\langle\sigma\rangle$ can be regarded as a rank $n+1$ geometry over the typeset $\bar{I}$ of δ-orbits on I, and so the description of Γ as the truncation of a homomorphic image $B/\langle\sigma\rangle$ to the typeset $T_0 = \{n+1, n\}$ makes sense.

Our discussion will involve two further typesets which we define now. These are $T := \{n-1, n, n+1, n+2\}$, a subset of I, and $\bar{T}$ consisting of the set of three δ-orbits: $\{(n-1, n+2), (n), (n+1)\}$, a subset of $\bar{I}$.

Notation About Types and Type-Shifting

For purposes of clarifying the discussion, it will be convenient to adopt the following conventions:

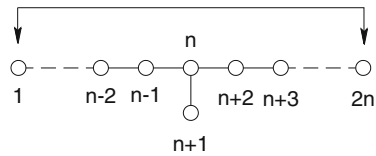

Fig. 16.13 The relevant symmetry of the Coxeter diagram together, displaying the assignment of types from I

(C1) If X is any geometry over I with type function $t : X \to I$, then X^δ is the geometry over I whose type function is $\delta \circ t$. (Since t and $\delta \circ t$ possess the same multipartite graph as domain, X and X^δ share the same objects and incidences; in fact X^δ is simply the geometry X with the types renamed. Moreover, if A is any set of objects of X, then A^δ is the same set viewed as a subset of X^δ. Thus if F is a flag of X, then F^δ is a flag of X^δ, of a different type.

(C2) Similarly, if C is a chamber system over I, then C^δ is a chamber system with the same set of vertices (chambers) but with every i-adjacency replaced by a $\delta(i)$-adjacency.

(C3) A *type-shifting automorphism* of a geometry X (chamber system C) over I is simply a *type-preserving* isomorphism $\alpha : X \to X^\delta$ ($\gamma : C \to C^\delta$). It can also be viewed as an incidence-preserving bijection, $X \to X$, which takes each object of type i to an object of type $\delta(i)$. Similarly, for a chamber system C, it is a bijection $C \to C$ taking each edge $e = (c_1, c_2)$ bearing labels $\lambda(e)$ to an edge $e^\sigma = (c_1^\sigma, c_2^\sigma)$ bearing labels $\lambda(e^\sigma) = (\lambda(e))^\sigma$. In particular, it takes a panel of type i to one of type $\delta(i)$. If we wish to make clear the particular type permutation involved, we refer to these mappings as *δ-shifting automorphisms*.[2]

The Present Situation, and Where it Must Go

Recall that $\Gamma = (\mathcal{P}, \mathcal{L})$ had an enrichment $\Gamma' := (\mathcal{P}, \mathcal{L}, \Pi, \mathcal{S})$, where $\mathcal{S}$ is the collection of all the symplecta of the parapolar space Γ and Π is the class of planes, each plane being the intersection of the unique pair of maximal singular spaces which contain it (see Step 4). We asserted that there was a geometry $\hat{\Gamma}_0$ over T admitting an automorphism $\bar{s}$ such that $\Gamma'' := \Gamma'|_{\{\mathcal{P},\mathcal{L},\Pi\}}$ is isomorphic to $\hat{\Gamma}_0/\langle\bar{s}\rangle$ as geometries over the type set $\bar{T} = \{n+1, n, (n-1, n+2)\}$ whose types are identified with $\{\mathcal{P}, \mathcal{L}, \Pi\}$ in respective order.

But in order to get the geometry $\hat{\Gamma}_0$ and involution $\bar{s}$, we had to employ Theorem 12.2.1 (Chap. 12) which was phrased as a theorem on point-line geometries with different definitions of "points" and "lines" compared to the definitions we were using for Γ. Thus we had to recast Γ as a new point-line geometry $\Gamma_0 = (\mathcal{P}_0, \mathcal{L}_0)$ whose "points" were the lines of Γ and whose "lines" were the line-pencils of point-plane flags of Γ. In this geometry, the type $\mathcal{P} = n + 1$ is a class $\mathcal{A}_0$ of Grassmann spaces of type $A_{2n-1,n}$ so Γ can always be recovered from Γ_0 by a truncation of type $T_0 = \bar{T}_0 = \{n+1, n\}$. Then without further comment, Theorem 12.2.1, produced (1) a "covering geometry" $\hat{\Gamma}_0$ over T, with two distinct classes of line-Grassmannians, $\mathcal{G}_i$ (as well as objects of types n and $n + 1$ covering $\mathcal{P}_0$ and $\mathcal{A}_0$), (2) a type-shifting

[2] The author prefers to retain the word "automorphism" in this definition since the basic incidence and/or adjacency structure is preserved. The adjective preceding the word still serves to distinguish the notion from that of ordinary (type-preserving) automorphisms.

automorphism $\bar{s}$, and (3) the isomorphism of Γ_0 with the truncation of $\hat{\Gamma_0}/\langle\bar{s}\rangle$ to $(n) = \mathcal{P}_0$ and flags of type $\{n+1, (n-1, n+2)\}$ in $\bar{T}$ – that is, the "lines" of $\mathcal{L}_0$.

On the other hand we observed that $\hat{\Gamma_0}$ was a locally truncated geometry over T with locally truncated diagram of type $Y_{n-1,1,n-1}$ over I (as in Fig. 16.10 with both horizontal arms of equal length and the central nodes labeled as in Fig. 16.11), thus allowing us to exploit the sheaf-theoretic construction of Chap. 11 to obtain a morphism $f : B|_T \to \hat{\Gamma_0}$.

So we need to show that the type-shifting automorphism $\bar{s}$ of $\hat{\Gamma_0}$ "lifts" to an involutory type-shifting automorphism σ of B inducing the permutation δ on the type-set I, with a suitable commuting condition $f \circ \sigma|_T = \bar{s} \circ f$ so as to induce a morphism

$$(B/\langle\sigma\rangle)|_{\bar{T}_0} \to (\hat{\Gamma_0}/\langle\bar{s}\rangle)|_{n,n+1} = \Gamma_0 \tag{16.7}$$

as desired. In order to do this we must understand in detail the construction of $\hat{\Gamma_0}$ and $\bar{s}$ which heretofore we simply inherited from Theorem 12.2.1.

Constructing $\bar{s}$

Beginning with Γ, we see that the residual of a point-line flag (p, L) of Γ is a product geometry $A \times B$ of two $PG(n)$s which contains two classes $\mathcal{A}_L$ and $\mathcal{B}_L$ of maximal singular spaces and these sets are the same however point p is chosen in L. So above any Γ-line L there are exactly two classes of maximal singular subspaces; two distinct members of this class belong to the same class if and only if their intersection is L, otherwise their intersection is a plane and they then belong to opposite classes. Given a line L and a maximal singular subspace M containing it, we let $\mathcal{C}(L : M)$ be the class of maximal $PG(n+1)$s of G incident with L which contain element M as a member and let $\mathcal{C}'(L : M)$ be the class which does *not* contain M as a member. We have already mentioned that "points" of Γ_0 are such lines of G, and that the "lines" of Γ_0 are pencils defined by a point-plane flag of Γ.

Now the "points" $\hat{\mathcal{P}_0}$ of $\hat{\Gamma_0}$ are all pairs $(L, \mathcal{C}(M))$ where M is a maximal singular $PG(n+1)$ of Γ containing line L. A "line" of $\hat{\Gamma}_0$ (i.e., an element of $\hat{\mathcal{L}_0}$) is a collection $\{(L, \mathcal{C}(M))\}$ as L ranges over a pencil of Γ-lines defined by a point-plane Γ-flag (p, π) and M is some maximal singular $PG(n+1)$ containing the plane π. Thus a "point" $(L_1, \mathcal{C}(L_1 : M_1))$ of $\hat{\mathcal{P}_0}$ is collinear with a "point" $(L_2, \mathcal{C}(L_2 : M_2))$ if and only if L_1 and L_2 generate a plane $\pi = \langle L_1, L_2\rangle_\Gamma$, and M is a common maximal singular $PG(n+1)$ containing this plane and $\mathcal{C}(L_1 : M_1) = \mathcal{C}(L_2 : M_2) = \mathcal{C}(M)$. By the assumption of the second case in Step 6, this point-collinearity graph is connected.

We now note that if L_1 and L_2 are lines generating a plane in a maximal singular subspace of M, then $\mathcal{C}'(L_1, M) = \mathcal{C}'(L_2, M)$ since there are only two maximal singular subspaces containing this plane, and both classes of the equation contain the one which is not M. Put another way:

In $\hat{\Gamma_0}$*, the "point"* $(L_1, \mathcal{C}(L_1 : M_1))$ *is collinear with* $(L_2, \mathcal{C}(L_2 : M_2))$ *if only if* $(L_1, \mathcal{C}'(L_1 : M_1))$ *is collinear with* $(L_2, \mathcal{C}'(L_2 : M_2))$.

Thus the mapping

$$\bar{s}|_{\hat{\mathcal{P}}_0} : \hat{\mathcal{P}}_0 \to \hat{\mathcal{P}}_0$$

defined by

$$(L, \mathcal{C}(L : M)) \mapsto (L, \mathcal{C}'(L : M))$$

takes "lines" of $\hat{\mathcal{L}}_0$ to other such "lines," giving us an automorphism of order two.

We now note that M defines a subspace $G_1(M)$ of $\hat{\Gamma_0}$ whose "points" are the pairs $(L, \mathcal{C}(L, M))$ as L ranges over the Γ-lines of M, and whose "lines" are collections of such "points" for which L ranges over some point-plane pencil in M. Thus, as a subspace of $\hat{\Gamma_0}$, $G_1(M)$ is a Grassmannian of lines of the projective space M.

It follows that

$$G_2(M) := \{(L, \mathcal{C}'(L, M))|\ L \text{ ranges over the } \Gamma\text{-lines of } M\}$$

is also a line-Grassmannian subspace of $\hat{\Gamma_0}$, since it is $\bar{s}(G_1(M))$. $G_1(M)$ and $G_2(M)$ are recognizably distinct types of subspaces, since in one case M is a member of the second coordinate of its "lines," while it is not in the other case.

We thus obtain the following two disjoint classes of subspaces of $\hat{\Gamma_0}$:

$$\mathcal{G}_i := \{G_i(M)|M \text{ ranging over maximal singular subspaces of } \Gamma\},\ i = 1, 2.$$

As previously noted, $\hat{\Gamma_0} = (\mathcal{G}_1, \hat{\mathcal{P}}_0, \mathcal{A}_0, \mathcal{G}_2)$ is a locally truncated geometry associated with the diagram of Fig. 16.11 on p. 587. What is new is that we have an incidence-preserving automorphism $\bar{s}$ of $\hat{\Gamma_0}$ which interchanges the classes $\mathcal{G}_1$ and $\mathcal{G}_2$, while fixing $\hat{\mathcal{P}}_0$ and $\mathcal{A}_0$. This the first of a series of automorphisms:

$$\bar{s} :\ \hat{\Gamma_0} \to \hat{\Gamma_0}, \tag{16.8}$$

$$s :\ C(\mathcal{F}) \to C(\mathcal{F}), \tag{16.9}$$

$$\hat{s} :\ \hat{C}(\mathcal{F}) \to \hat{C}(\mathcal{F}), \tag{16.10}$$

$$\sigma :\ B = \mathbf{\Gamma}(\hat{s}) \to \mathbf{\Gamma}(\hat{s}) = B, \tag{16.11}$$

$$\sigma_T :\ B|_T \to B|_T. \tag{16.12}$$

Here, $\mathbf{C}(\mathcal{F})$ is the chamber system associated with the sheaf $\mathcal{F}$ defined by the locally truncated geometry Γ_0^*, $\hat{\mathbf{C}}(\mathcal{F})$ is the universal 2-cover of $\mathbf{C}(\mathcal{F})$ and is a building belonging to the $Y_{n,1,n}$-diagram, and $B = \mathbf{\Gamma}(\hat{\mathbf{C}}(\mathcal{F}))$ is the functorially-defined building geometry obtained from the residually connected $\hat{\mathbf{C}}(\mathcal{F})$.

Lifting $\bar{s}$ to a Type-Shifting Automorphism of $C(\mathcal{F})$

We wish to pass from the type-shifting automorphism $\bar{s}$ of the locally truncated geometry $\hat{\Gamma_0}$ with sheaf $\mathcal{F}$ to a δ-shifting automorphism $s : C(\mathcal{F}) \to C(\mathcal{F})$. In order to do this, we must understand how this chamber system is formed from the sheaf, but we need not know it precisely.

The sheaf $\mathcal{F}$ is a function which associates with each non-empty flag F of $\hat{\Gamma_0}$, a geometry $\mathcal{F}(F)$ over $I - t(F)$ ("t" serves here also as the type function $\hat{\Gamma_0} \to T$), which obeys some relations converting intervals in the poset of all non-empty flags of $\hat{\Gamma_0}$ to the taking of residues in the image geometries. The relations (which we shall loosely refer to as "the locally truncated data provided by $\hat{\Gamma_0}$") possess two important features:

A. The relations allow one to describe an object of $\mathcal{F}(F)$ in more than one way (all of them being somehow consistent).
B. The relations admit the complete δ-shifting symmetry induced by $\bar{s}$.

Let $I = I^- + T + I^+$ be the partition of types defined by setting $I^- = \{1, \ldots, n-1\}$ and setting $I^+ = \delta(I^-)$. In this paragraph we wish to exploit the redundancy in item A above to represent the objects of the geometries that are produced by our sheaf $\mathcal{F}$. Let F be any flag of $\hat{\Gamma_0}$.

(a) If $i \in I^-$, then an object of type i in $\mathcal{F}(F)$ is a subgeometry S of $\mathrm{Res}_{\hat{\Gamma_0}}(x_{n+2})$ containing $F - \{x_{n+2}\}$, where x_{n+2} is an object of $\hat{\Gamma_0}$ of type $n+2$ such that $F \cup \{x_{n+2}\}$ is a flag of $\hat{\Gamma_0}$.
(b) If $i \in I^+$, then an object of type i in $\mathcal{F}(F)$ is a subgeometry S of $\mathrm{Res}_{\hat{\Gamma_0}}(x_{n-1})$ containing $F - \{x_{n-1}\}$ where $F \cup \{x_{n-1}\}$ is a flag of $\hat{\Gamma_0}$.
(c) If $i \in T$, then an object of type i in $\mathcal{F}(F)$ is a singleton set $\{y\}$, for $y \in \mathrm{Res}_{\hat{\Gamma_0}}(F)|_i$. (Of course i is in $T - t(F)$ in this case, and no such i appears if F is a chamber flag of $\hat{\Gamma_0}$.)

The point is that objects of $\mathcal{F}(F)$ are always induced subgeometries S from subsets of the set of objects of $\hat{\Gamma_0}$, and so possess images under the action of the type-shifting automorphism $\bar{s}$ of $\hat{\Gamma_0}$. Since $\bar{s}$ preserves all incidences in $\hat{\Gamma_0}$, the following must hold:

(1) X is an object of type i in $\mathcal{F}(F)$ if and only is $\bar{s}(X)$ is an object of type $\delta(i)$ in $\mathcal{F}(\bar{s}(F))$.
(2) If the flag F contains an object whose type is distinct from $n+1$, then for any pair (S_1, S_2) of $\mathcal{F}(F)$ of types i and j, $(i, j) \in I^- \times I^+$, then the set of objects S_1 is considered to be incident with S_2. With this definition, if $X = \{X_{i_1}, X_{i_2}, \ldots, X_{i_k}\}$ is a flag of $\mathcal{F}(F)$ with X_{i_j} of type i_j, then $\bar{s}(X)$ is a flag $\{\bar{s}(X_{i_1}), \bar{s}(X_{i_2}), \ldots, \bar{s}(X_{i_k})\}$ of $\mathcal{F}(\bar{s}(F))$ of type $\delta(\{i_1, \ldots, i_k\})$.

In particular, (2) holds when F is a flag-chamber of Γ and X is a flag-chamber of $\mathcal{F}(F)$.

Now a chamber c of the chamber system $C(\mathcal{F})$ is a triple expressed as

$$c = c^-(F) \oplus F \oplus c^+(F)$$

where F is a flag-chamber of $\hat{\Gamma_0}$, and for $e = +$ or $-$, the symbol $c^e(F)$ denotes a flag-chamber of the truncation of $\mathcal{F}(F)$ to I^e.[3]

It then follows from (2) above that

$$\begin{aligned} s(c) :&= (\bar{s}(c^+(F)) \oplus \bar{s}(F) \oplus (\bar{s}(c^-(F))) \\ &= (c^+(\bar{s}(F)))^\delta \oplus \bar{s}(F) \oplus (c^-(\bar{s}(F)))^\delta \end{aligned}$$

is also a chamber of $C(\mathcal{F})$, and that if chambers c_1 and c_2 are i-adjacent in $C(\mathcal{F})$ then $s(c_1)$ and $s(c_2)$ are $\delta(i)$-adjacent in $C(\mathcal{F})$. Thus

$$s : C(\mathcal{F}) \to C(\mathcal{F})$$

defined by

$$c \mapsto s(c)$$

is a δ-shifting automorphism of the chamber system $C(\mathcal{F})$. This explains Eq. (16.9).

Remark One could summarize this step as follows.

Since the shift-automorphism $\bar{s}$ preserves the "locally truncated data," that is, $\bar{s}$ leaves invariant the poset of subgeometries that define the sheaf values (up to the new assignment of types for these geometries via δ), $\bar{s}$ induces a δ-shifting automorphism of the chamber system $C(\mathcal{F})$.

Lifting s

Now, as noted in (C3), p. 589, the δ-shifting automorphism of Eq. (16.9) can be rephrased as a (type-preserving) isomorphism

$$s^\delta : C(\mathcal{F}) \to C(\mathcal{F})^\delta$$

sending chamber c to $(s(c))^\delta$. Now let

$$\begin{aligned} \kappa &: \hat{C}(\mathcal{F}) \to C(\mathcal{F}), \\ \kappa^\delta &: (\hat{C}(\mathcal{F}))^\delta \to (C(\mathcal{F}))^\delta \end{aligned}$$

[3] The reader should be aware that this assertion uses the fact that the removal of the nodes labeled by types in T from the $Y_{n-1,1,n-1}$-diagram leaves a disconnected diagram. In other words, for $(i, j) \in I^- \times I^+$, each object of type i in $\mathcal{F}(F)$ is incident with each object of type j in $\mathcal{F}(F)$.

be the universal 2-coverings of the chamber systems $C(\mathcal{F})$ and $(C(\mathcal{F}))^\delta$, respectively. Then as s^δ is invertible, it is easy to argue that the composition

$$s^\delta \circ \kappa : \hat{C}(\mathcal{F}) \to (C(\mathcal{F}))^\delta$$

is also a universal 2-covering of $(C(\mathcal{F})^\delta$, and so (by a fundamental property of universal covers) there exists a type-preserving isomorphism

$$(\hat{s})^\delta : \hat{C}(\mathcal{F}) \to (\hat{C}(\mathcal{F}))^\delta$$

with the commuting property

$$\kappa^\delta \circ (\hat{s})^\delta = s^\delta \circ \kappa. \tag{16.13}$$

From this one derives a δ-shifting automorphism

$$\hat{s} : \hat{C}(\mathcal{F}) \to \hat{C}(\mathcal{F})$$

with the property that

$$\kappa \circ \hat{s} = \bar{s} \circ \kappa. \tag{16.14}$$

This explains Eq. (16.10).

The Type-Shifting Automorphism σ

Noting that $s, \kappa, \kappa^\delta, \hat{s}^\delta$ are *type-preserving* morphisms connecting residually-connected chamber systems over I, one can apply the functor $\mathbf{\Gamma}$ to obtain corresponding mappings $\mathbf{\Gamma}(\bar{s}), \mathbf{\Gamma}(\kappa), \mathbf{\Gamma}(\kappa^\delta), \mathbf{\Gamma}(\hat{s}^\delta)$ connecting the geometries $\mathbf{\Gamma}(Z)$, where Z ranges over the chamber systems connected by the listed chamber morphisms.

We also note that by Tits' local approach theorem, $\hat{C}(\mathcal{F})$ and $\hat{C}(\mathcal{F})^\delta$ are buildings associated with diagram $Y_{n-1,1,n-1}$. And so, applying the functor $\mathbf{\Gamma}$, one obtains building geometries $B := \mathbf{\Gamma}(\hat{C}(\mathcal{F}))$ and

$$B^\delta = \mathbf{\Gamma}(\hat{C}(\mathcal{F}))^\delta = \mathbf{\Gamma}((\hat{C}(\mathcal{F}))^\delta), \tag{16.15}$$

connected by an isomorphism

$$\sigma^\delta := \mathbf{\Gamma}(\hat{s}^\delta) : B \to B^\delta.$$

This mapping produces the desired δ-shifting automorphism

$$\sigma := \mathbf{\Gamma}(\hat{s}) : B \to B$$

of Eq. (16.11).

The functor $\mathbf{\Gamma}$ also converts Eq. (16.14) into the following commutative diagram:

$$\begin{array}{ccc} B & \xrightarrow{\sigma^\delta} & B^\delta \\ {\scriptstyle f:=\mathbf{\Gamma}(\kappa)}\downarrow & & \downarrow{\scriptstyle f^\delta:=\mathbf{\Gamma}(\kappa)^\delta} \\ \mathbf{\Gamma}(C(\mathcal{F})) & \xrightarrow{\mathbf{\Gamma}(s^\delta)} & \mathbf{\Gamma}(C(\mathcal{F})^\delta) \end{array}$$

Now the sheaf values $\mathcal{F}(x)$ at objects of $\hat{\Gamma_0}$ are either (1) a disjoint union of two $PG(n-1)$s and a line, subspaces belonging two different components of the union being considered incident, (2) a Grassmannian of lines, or (3) a Grassmannian of type $A_{2n-1,n}$. In all cases these residues are residually connected, and hence strongly chamber connected geometries. We are then permitted to apply Theorem 11.4.2 (Chap. 11, p. 429) to identify the truncation to typeset T of two of these geometries. We essentially have

$$\hat{\Gamma_0} = \mathbf{\Gamma}(C(\mathcal{F}))|_T \text{ and } \hat{\Gamma_0}^{*\delta} = \mathbf{\Gamma}(C(\mathcal{F})^\delta)|_T.$$

Since $T = T^\delta$, and the truncation operator "$|_T$" is a functor, we can apply the functor to the preceding diagram to obtain the diagram below.

$$\begin{array}{ccc} B|_T & \xrightarrow{\sigma^\delta|_T} & B^\delta|_T \\ {\scriptstyle f|_T}\downarrow & & \downarrow{\scriptstyle f^\delta|_T} \\ \hat{\Gamma_0} & \xrightarrow{\mathbf{\Gamma}(s^\delta)|_T} & (\hat{\Gamma_0})^\delta \end{array}$$

We leave it to the reader to verify at a general object x of $\hat{\Gamma_0}$, that

$$\mathbf{\Gamma}(s^\delta)|_T(x) = \bar{s}^\delta(x).$$

Replacing the horizontal arrows by δ-shifting automorphisms one has the commuting diagram

$$\begin{array}{ccc} B|_T & \xrightarrow{\sigma|_T} & B|_T \\ {\scriptstyle f|_T}\downarrow & & \downarrow{\scriptstyle f|_T} \\ \hat{\Gamma_0} & \xrightarrow{\bar{s}} & \hat{\Gamma_0} \end{array}$$

It follows that $f|_T$ converts $\sigma|_T$-orbits on $B|_T$ to $\bar{s}$-orbits on $\hat{\Gamma_0}$. But $\sigma|_T$-orbits are simply σ-orbits on B whose types are in T. From these two statements one deduces the following.

Step 7. *The rank three geometry* $\hat{\Gamma_0}/\langle \bar{s} \rangle$ *is a truncation to* $\bar{T} = \{n+1, n, \{n-1, n-2\}\}$ *of a homomorphic image of the factor geometry* $B/\langle \sigma \rangle$ *where* B *is a building belonging to the* $Y_{n-1,1,n-1}$ *diagram, and* σ *is a* δ*-shifting automorphism of* B*. One more truncation to* $T_0 = \{n+1, n\} = \{\mathcal{A}, \mathcal{P}_0\} = \{\mathcal{P}, \mathcal{L}\}$ *yields the description of* Γ *given in conclusion (d) of the corollary.*

This completes the proof of Corollary 16.7.1. □

16.8 Appendix to Chapter 16

16.8.1 Background

Let $\Gamma = (\mathcal{P}, \mathcal{L})$ be a parapolar space. If Γ has symplectic rank exactly two, assume all maximal singular subspaces are projective spaces. Suppose $\mathcal{M}$ is a class of maximal singular subspaces of finite projective rank with these properties:

1. Every point lies in at least one member of $\mathcal{M}$.
2. For every non-incident pair (p, M) chosen from $\mathcal{P} \times \mathcal{M}$, one has that $x^{\perp} \cap M$ is a projective space of projective dimension d, for a non-negative integer d that is independent of the choice of (p, M).

Then Theorem 16.2.1 forces $d = 0$, $d = 1$, or Γ is itelf a polar space. Moreover, assuming Γ is not a polar space, the following hold:

1. If $d = 1$, Γ is either a polar space of rank three or a Grassmannian of lines of some projective space.
2. If $d = 0$ then either (a) Γ is the product of two projective spaces or (b) the points of Γ are partitioned by $\mathcal{M}$, lines not in an element of $\mathcal{M}$ are maximal singular spaces, and non-oriflame quads exist (conclusion 2(b) of Theorem 16.2.1).

The purpose of this appendix is to establish that examples satisfying the conditions of 2 (b) above do indeed exist.

16.8.2 Admissible Triples

In [51], Bart De Bruyn introduced an important notion, intimately related to quadrangles with a spread of symmetry. Let $U = (P, \Lambda)$ be a 2-design with $s+1$ points per line and t lines on each point. Let G be a group whose binary operation is denoted "*". We suppose

$$\Delta : P \times P \to G$$

to be a mapping with the following property. Whenever p_1, p_2, p_3 are three points lying in a common line $D \in \Lambda$ then

$$\Delta(p_1, p_2) * \Delta(p_2, p_3) = \Delta(p_1, p_3). \tag{16.16}$$

It immediatiately follows that $\Delta(x, x) = e$, the identity element of G, and that $\Delta(x, y) = (\Delta(y, x))^{-*}$, the inverse of the elements $\Delta(x, y)$ in the group $(G, *)$, for any points $x, y \in P$.

Such a triple $T = (U, G, \Delta)$, where Δ is as described above in Eq. (16.16), is called an *admissible triple*.

Starting with an admissible triple T for which $|G| = s + 1$ is also the size of a line of the 2-design U, one can construct a generalized quadrangle $Q(T)$ as follows.

Set $V = G \times P$ and define a graph with vertex set V by asserting pair (g_1, p_1) is adjacent to (g_2, p_2) if and only if either:

$$p_1 = p_2 \text{ and } g_1 \neq g_2, \text{ or} \tag{16.17}$$

$$p_1 \neq p_2 \text{ and } g_2 = g_1 * \Delta(p_1, p_2). \tag{16.18}$$

Then the incidence system $Q(T) := (V, \mathcal{C})$, where $\mathcal{C}$ is the collection of all maximal cliques of the graph $(V, \sim)$, is a generalized quadrangle. (This involves a counting argument - see [52].)

One notes that the collection $\mathcal{S} = \{S_p | p \in U\}$ where $S_p := \{(g, p) | g \in G\}$ forms a spread of $Q(T)$ whose lines are indexed by the points of the design U. If p and q are distinct points of U, there is a unique line D of U on these points, defining a partial spread $\gamma = \{S_p | p \in D\}$. Now by Eq. (16.18) the sets

$$A_g := \{(g * \Delta(p, r), r) | r \in D\}$$

belong to $\mathcal{C}$ and form another system $\gamma^* := \{A_g | g \in G\}$ of transverse lines – so that γ and γ^* together are the lines of a grid. Thus $\mathcal{S}$ is a Hermitian 2-system of the generalized quadrangle $Q(T)$.

Furthermore, the transformation σ_h defined by

$$(g, p) \to (h * g, p),$$

takes the edge $((k, x), (k * \Delta(x, y), y))$ to the edge $(h * k, x), (h * k * \Delta(x, y), y)$, and so G acts as a group stabilizing the spread $\mathcal{S}$ and acting regularly on the points of each component.

There are two primary examples of quadrangles constructed in this way.

Example 1 Let $U = (P, \Lambda)$ be the unital design of points and secants of the Hermitian variety $H(2, q^2)$ defined by a Hermitian form h on the vector space $V(3, q^2)$. Trace- and field-reduction produces an orthogonal form $tr(h)$ of type $O^-(6, q)$ and the points of the unital U induce a Hermitian spread S of the corresponding quadrangle $Q^-(5, q)$. De Bruyn [51] shows that this quadrangle can be reconstructed from an admissible triple (G, U, Δ) where G is the group of $q - 1$-st powers of the elements of $GF(q^2)^*$ (cyclic of order $q + 1$), and Δ is defined by the equation

$$\Delta(\langle x\rangle, \langle y\rangle) = h(a, x)^{q-1} h(x, y)^{q-1} h(y, a)^{q-1}, \tag{16.19}$$

where $(a, x, y) \in (V(3, q)^{\#})^{(3)}$, but a is fixed.

Example 2 $W(3, q)$ is the generalized quadrangle of points and lines that are isotropic with respect to a symplectic form on the underlying vector space $V(4, q)$. Any point x of this quadrangle is regular, which means that if (x, y) is a non-collinear pair, then both $\{x, y\}^{\perp}$ and $(\{x, y\}^{\perp})^{\perp}$ are cocliques of $q + 1$ points. By the construction of Stanley Payne one obtains a second quadrangle $W(3, q, x)$ as follows: its points are the points of $W(3, q) - x^{\perp}$; its lines are first, the lines of $W(4, q)$ not on x, together with the sets $(\{x, y\}^{\perp})^{\perp} - \{x\}$ which form a partition $\mathcal{S}$ of the points of $W(3, q) - x^{\perp}$ far from x. Again, De Bruyn shows that $\mathcal{S}$ forms a Hermitian spread of this quadrangle $W(3, q, x)$ with associated design of lines and grids forming the affine plane $U := AG(2, q^2)$, and that there exists an appropriate definition of Δ such that $T = ((GF(q), +), U, \Delta)$ is an admissible triple from which $W(4, q, x)$ can be recovered by the canonical construction.

Remark For further information on admissible triples, their characterizations, and their relation to near polygons the author recommends [51, 53], and Chap. 12 of Koen Thas' thesis [134].

16.8.3 Constructions of the Relevant Parapolar Spaces

We are now in a position to employ admissible triples to construct some examples of parapolar spaces that should fall under the troublesome last case described at the end of the first section above. (Here $d = 0$, Γ is not a quadrangle and each of its points lies in a unique member of $\mathcal{M}$.)

Suppose $T = (G, U, \Delta)$ is an admissible triple for which $s + 1 = |G|$ is also the size of a line in the design U, and, as above, let $Q(T)$ be the quadrangle obtained from T by the construction of the Sect. 16.8.2 above. Now we suppose s to be a prime power q. Now construct a tower of fields

$$GF(q) = F \subset E \subset K,$$

where $[E : F] = 2$ and $[K : E] = n > 1$. Then $G := E^*/F^* \simeq Z_{q+1}$, the cyclic group of order $q + 1$ and the cyclic group $G' := K^*/F^*$ has order

$$(q^{2n} - 1)/(q - 1) = 1 + q + \cdots + q^{2n-1}.$$

Let $V = K_F$, the $2n$-dimensional vector space obtained by regarding K as an F-module, and let $M = P(V)$ be the associated projective space. G' acts regularly as a Singer cycle on M while G stabilizes and acts regularly on each component of the collection of E-submodules of K_E that form a spread of M. At this point we select a special point p_0 in M so that we can identify the points of M with the elements of

G' so that p_0^g corresponds to element g of G'. Then the lines of the projective space M which are incident with the point p_0 can be thought of as a collection of subsets $\mathcal{D} := \{D_0, D_1, \ldots, D_m\}$ of G', each containing the identity element e (corresponding to point p_0). Here $m + 1 = 1 + q + \cdots q^{2n-2}$, and we can choose the indexing so that D_0 is the subgroup $G = E^*/K^*$. Then, of course, the full set of translates $\mathcal{D}' := \{hD | D \in \mathcal{D}, h \in G'\}$ becomes the line set of the projective space M.

Now set $\mathcal{P} := G' \times P$ where, as above, P is the set of points of the unital $U = (P, \Lambda)$. The set of lines $\mathcal{L}$ that we impose on this set are of two types:

$\mathcal{L}_\lambda$: Sets $L_{h,x,L} := \{(h * \Delta(x, y), y) | y$ ranging over a line L of Λ containing $x\}$. (This notation $L_{h,x,L}$ is not unique, for it also describes $L_{h',x',L}$ if $h' = h * \Delta(x, x')$; but it does suffice to determine our line.)

$\mathcal{L}_\mathcal{M}$: Sets $N_{p,D'} := \{(h, p) | h \in D' \in \mathcal{D}'\}$.

We call lines of the first type "lateral lines," and lines of the second type "$\mathcal{M}$-lines." For each point $p \in P$, the set

$$M_p := \{(h, p) | h \in G'\},$$

together with the $\mathcal{M}$-lines within it, is clearly a maximal singular subspace of $(\mathcal{P}, \mathcal{L})$ isomorphic to the projective space M, and the collection $\mathcal{M} := \{M_p | p \in P\}$ partitions the set $\mathcal{P} = G' \times P$. The "lateral lines" connect points in distinct members of $\mathcal{M}$.

It is also clear that for each coset $\rho := hG$ of the subgroup G in G', the collection of points

$$Q_\rho := \{(h, p) | h \in \rho\}$$

is a subspace which is a generalized quadrangle constructed from the admissible triple (G, U, Δ) as described in the previous section. Then we see that the collection

$$\mathcal{Q} := \{Q_\rho | \rho \in G'/G\}$$

also forms a partition of $\mathcal{P} = G' \times P$.

Each member of $\mathcal{M}$ now intersects each member of $\mathcal{Q}$ at a line, thus inducing a spread of lines in any component of either set.

Now let us consider two arbitrary non-collinear points $a_1 := (h_1, p_1)$ and $a_2 := (h_2, p_2)$. We first note that $p_1 \neq p_2$, since otherwise a_1 and a_2 would lie in a common component of $\mathcal{M}$ and hence would be collinear. On the other hand, a_1 is collinear with $a_1' := (h_1 * \Delta(p_1, p_2), p_2)$ which is collinear with $a_2 = (h_2, p_2)$ since both are members of M_{p_2}. This raises two possibilities: (1) $h_1 * \Delta(p_1, p_2)$ and h_2 belong to the same coset ρ of G in G' or (2) they do not. If (1) holds, the two points in question belong to a common $Q_\rho \in \mathcal{Q}$. If (2) holds, then h_1 and h_2 themselves are not in the same coset of G, and so are connected by a line hD_i, $i \neq 0$. Similarly, a_2 is collinear with

$$a_2' := (h_2 * \Delta(p_2, p_1), p_1) = (h_2 * \Delta(p_1, p_2)^{-1}, p_1)$$

which is collinear with a_1 by a line of M_{p_1}.

We leave it as an exercise that the 4-circuit $(a_1, a_1', a_2, a_2'a_1)$ generates a grid all of whose lines can be be described in the notation given above for lines in $\mathcal{L}_{\mathcal{M}}$ and $\mathcal{L}_\lambda$. One notes that this grid $G_{D',\lambda}$ is complelely determined by the choice of a line $\lambda \in \Lambda$ and a line $D' \in \mathcal{D}' - \mathcal{S}$.

Now fix point $p \in P$ and the corresponding component of $\mathcal{M}$, namely M_p. Suppose $(h, x) \in G' \times P$ is a point not in M_p. Then, of course, $x \neq p$. Now suppose (h, x) is collinear to a point (h', p) of M_p. Then $h = h' * \Delta(x, p)$ and the point (h, p) is uniquely determined. Thus the following applies.

(*) *Given a member M of $\mathcal{M}$, and a point $a \in \mathcal{P}$ not in M, then a is collinear with a unique member of M.*

Finally, as a third exercise, the student, using the fact (*), may prove that the grids $G_{D',\lambda}$ described above are in fact convex subspaces, as are the components of $\mathcal{Q}$.

Thus we see the following.

Lemma 16.8.1 *The point-line geometry $\Gamma = (\mathcal{P}, \mathcal{L})$ constructed as above from a $PG(2n, q)$ and an admissible triple $T = (G, U, \Delta)$, where G is cyclic of order $q + 1$, and each line of the 2-design U possesses $q + 1$ points, is in fact a strong parapolar space with a class of maximal singular spaces $\mathcal{M}$ which are projective spaces with the properties that (1) every point lies in some member of $\mathcal{M}$ and (2) any point not in any designated member M of $\mathcal{M}$ is collinear with exactly one point of M.*

The work of De Bruyn [52] thus gives us these examples.

Example 1 For odd q, the admissible triple for $Q^-(5, q)$ allows a construction of a parapolar space which, following De Bruyn, we may call $PG(n, q) \otimes Q^-(5, q)$.

Example 2 If p is a prime number, the admissible triple (G, U, Δ) that defines the generalized quadrangle $W(3, p, x)$ involves a group G which is cyclic of order p. Thus if $p = q + 1$ where q is a power of two – that is p is a Fermat prime – then we can march through the construction given above to form a parapolar space that would be labeled $PG(n, q = p - 1) \otimes W(3, p, x)$, again following De Bruyn's notation.

Remarks 1. Reviewing the above arguments, there are really only two requirements that make the construction of the parapolar space $M \otimes Q(T)$ possible:

1. The linear space must admit a point-regular group of automorphisms G' which stabilizes a spread $\mathcal{S}$, and contains a normal subgroup G_0' stabilizing and acting regularly on the points each member of $\mathcal{S}$.
2. The group G_0' must be isomorphic to the group G that figures in the admissible triple T.

Thus if one gives up the requirement that members M be projective spaces, one realizes that one could bond affine spaces $AG(n, q)$ with $W(3, q, x)$, where G is the additive group of the underlying field. The result is still a parapolar space of symplectic rank two. Recall that such a parapolar space may not have its singular subspaces projective spaces.

2. Of course the $X \otimes Y$-construction does not depend on finiteness. In the infinite case an admissible triple might conceivably lead to some partial linear space Y which is not neccessarily a generalized quadrangle. (That it was a quadrangle depended on a counting argument.) But otherwise, it is difficult to see why the construction should not proceed provided the two conditions in preceding Remark 1 are satisfied.

3. Finally, why must M even be a linear space? As remarked, all one seems to need is a spread $\mathcal{S}$ of the partial linear space M, a point-regular group H inducing a regular group on $\mathcal{S}$ whose kernel G is isomorphic to the group G of the admissible triple T. But when one discusses two objects X and Y where Y is defined by an admissible triple $T = T_Y$, one does not know in advance that the left bonding object X possesses such a helpful triple $(G', G, \mathcal{S})$. But, as De Bruyn has shown, this can occur when X also arises from an admissible triple Δ_X provided there is an antiautomorphism pasting them together. Perhaps T_X and $(G', G, \mathcal{S})$ amount to the same thing. Perhaps there are admissible triples that define structures of higher point-diameter that are amenable to similar constructions.

Chapter 17
Point-Line Characterizations of the "Long Root Geometries"

Abstract Four axioms characterize the long-root geometries, which include (1) polar Grassmannians of lines, (2) metasymplectic spaces, and (3) certain exceptional geometries of types E_6, E_7, or E_8. ([37, 116] update somewhat the theory presented here).

17.1 Introduction

17.1.1 Long Root Geometries

In work extending over a decade beginning around 1968, finite group theorists developed rather remarkable characterizations of what they call "long root groups."[1] Basically (and without being too precise) these are groups G generated by a conjugacy class A^G of elementary abelian subgroups such that any two distinct subgroups in the class generate one of four possible subgroups. This theory has more recently been redeveloped by Timmesfeld in the infinite case in [135] and in his important book [136], where an axiom on A replaces the assertion that it is elementary abelian. The conclusion is that G belongs to a certain list of known groups possessing a (B, N)-pair, or "Tits system" as it is called. (See p. 388 for the definitions.)

Taking the conjugates of A in G as points, and defining lines by one of the four classes of the generated subgroups, one obtains a geometry which one might justifiably call a "*long root geometry*." For the known G this device yields the following Lie-incidence geometries:

1. The polar Grassmannians of lines $C_{n,2}$, $n \geq 4$ and $D_{n,2}$, $n \geq 5$.
2. The metasymplectic spaces of type F_4 including the Lie incidence geometries $(B/C)_{4,2}$ and $D_{4,2}(F)$.

[1] This area has a long history going back to Thompson's work on quadratic pairs, work of Chat Ho, Stellmacher, Aschbacher-Hall, Cuypers, involving $q = 3^k$, work of Betty Salzmann (Stark), Cooperstein and Kantor – all for the finite case. That involved generational questions, and other questions beyond the mere classification of long-root coset geometries. Later, work of Timmesfield on K-transvections [135] opened up the infinite case and a host of problems (involving strange subrings of division rings) currently being continued by Anja Steinbach and Hans Cuypers [48, 49].

E.E. Shult, *Points and Lines*, Universitext, DOI 10.1007/978-3-642-15627-4_17,

© Springer-Verlag Berlin Heidelberg 2011

3. The exceptional Lie incidence geometries of types $E_{6,2}(F)$, $E_{7,1}(F)$, and $E_{8,8}(F)$ (using the Bourbaki numbering of nodes).
4. The long-root geometries of $A_{\{1,n\}}$ and the Lie incidence geometries $C_{3,2}(D)$.

The two geometries of the last item can be described more explicitly.

1. The first one is a point-line geometry $\Gamma = (\mathcal{P}, \mathcal{L})$ whose points are the point-hyperplane flags (p, H) of a projective space $\mathbf{P}(V)$ where V is a vector space of dimension $n \geq 5$. If one declares the lines to be all 2-space- codimension-2-space flags, you'll get into trouble, for then you won't even have a partial linear space. Instead (following the recipe of Cooperstein's Doctoral thesis) lines are of two types: (1) point-shadow sets $L(p, A) := \{(p, H) \in \mathcal{P} | \text{codim}(A) = 2\}$ where H ranges over all hyperplanes of V which are incident with flag (p, A) and (2) point-shadow sets $L(N, H) := \{(p, H) \in \mathcal{P} \mid N \text{ is a line in hyperplane } H \text{ and } p \in N\}$. If p and q are distinct points of $\mathbf{P}(V)$ and H and H' are two hyperplanes of $\mathbf{P}(V)$ containing the projective line pq, then $y_1 = (p, H)$ and $y_2 = (q, H')$ are two points of $\mathcal{P}$ at distance two in the point-collinearity graph of Γ and the convex closure $\langle\langle y_1, y_2\rangle\rangle$ is a grid associated with the flag $(pq, H \cap H')$ of type $(2, n-2)$. The lines of one parallel class are type (1); the lines of the other class are type (2).

However, if $q \in H \cap H'$ and $p \in H' - H$, then there is a unique point (q, H') collinear with both (q, H) and (p, H').

So this geometry is very much like your typical parapolar space: there are symplecta, polar pairs, and special pairs with respect to these symplecta. The problem is that these symplecta are grids, and this makes classification by point-line axioms follow routes quite distinct from the parapolar space world of this and the previous chapters.

2. In the geometry $C_{n,2}$, the points are the lines of the polar space $C_{n,1}$ and the lines are the point-plane flags of this same polar space Δ. Thus (following the point-shadow description of a line) a *line* of $C_{n,2}$ is the pencil of polar lines on some polar point in a polar plane in Δ. Now there are a number of relations between polar lines A and B of Δ. They could be (1) equal, (2) lie in a common polar plane, (3) be disjoint lines in a singular subspace of Δ, (3') A and B could meet at a polar point, but not lie in a common singular subspace, (4) $A \cap B = \emptyset$ and there is a unique transverse polar line L lying in $A^\perp \cap B^\perp$, or (5) they could be opposite lines of the polar subspace Δ. When A and B are regarded as points of the polar Grassmannian $C_{n,2}$ we see that respectively (1) they are equal, (2) they are distinct collinear points, (3) and (3') they lie in a symplecton, (4) there is a unique point collinear with both A and B, and (5) they are distance three from one another.

Clearly, if $n > 3$, $C_{n,2}$ is a parapolar space with two classes of symplecta of types $A_{3,2}(D)$ and $C_{n-1,1}$.

But if $n = 3$, as in the last item of our list of long root geometries, the point-relation (3) is empty and Γ is a parapolar space whose symplecta are quadrangles.

So we see that the "long root geometries" at the bottom of the list are precisely those which are not parapolar spaces of symplectic rank at least three. The characterizations of this chapter, like the last, take place on the stage of parapolar spaces of symplectic rank at least three, so these last two will not appear among conclusions here.

So these geometries came to be known as "long-root" geometries, because they are the geometries one realizes from long root groups (which play the role of $\mathcal{A}$) of a group of Lie type. But so far, this description is purely group-theoretical. Our goal in this chapter is to characterize them from purely geometric axioms.

A geometric characterization of long root geometries in classes 2 and 3 above was achieved by Cohen and Cooperstein in the second of the two major theorems of their spectacular paper [36]. Their hypotheses require the parapolar space Γ to possess constant symplectic rank, which explains why the long root geometries of class 1, which have two classes of symplecta of possibly different polar ranks, cannot appear in their conclusion. Their hypotheses allow just four relations between a point x and a symplecton S: either x is incident with S, or $x^\perp \cap S$ is empty, a line, or a maximal singular subspace of S.

Following the development in Kasikova and Shult [85],[2] our approach assumes Γ is a parapolar space of symplectic rank at least three satisfying the axioms (H1) and (H2) of a hexagonic geometry (defined below).[3] Our characterization theorem implies the earlier breakthrough theorem of Cohen–Cooperstein.

17.1.2 The Main Results of this Chapter

The goal of the present chapter is to prove the following two theorems.

Theorem 17.1.1 *Suppose Γ is a parapolar space of symplectic rank at least three satisfying these axioms:*

(H1) Given a point x not incident with a symplecton S, the space $x^\perp \cap S$ is never just a point.

(H2) Given a projective plane π and line L meeting π at point p, either (a) every line of π on p lies in a common symplecton with L, or else (b) exactly one such line incident with (p, π) has this property.

(H3) Given a point-line flag (p, L) there exists a second line N such that $L \cap N = \{p\}$ and no symplecton contains $L \cup N$ — i.e., (x, y) is a special pair for each $(x, y) \in (L - \{p\}) \times (N - \{p\})$.

(F) If all symplecta have rank at least four, assume every maximal singular subspace has finite projective rank.

Then Γ is

1. *$E_{6,4}$, $E_{7,7}$, or $E_{8,1}$,*
2. *A metasymplectic space, or*

[2] Professor Antonio Pasini also had a hand in making corrections and filling a gap in that manuscript.

[3] We met these geometries briefly on p. 480 in Sect. 13.7.

3. *A polar Grassmannian of lines of a non-degenerate polar space of (possibly infinite) rank at least four. In the case of finite polar rank, these would be classical Lie incidence geometries of type* $(B/C)_{n,2}$ *or* $D_{n,2}$, $n \geq 4$.

Theorem 17.1.2 *Suppose* Γ *is a strong parapolar space with these three properties:*

(1) For every point-symplecton pair (x, S), $x^{\perp} \cap S \neq \emptyset$.
(2) For each point p, *the ball* $\Delta_2^*(p)$ *of radius* 2 *about any point* p *is a proper subspace of* Γ.[4]
(3) If there is no symplecton of rank two assume every maximal singular subspace has finite projective rank.

Then Γ *is one of the following:*

1. $D_{6,6}$, $A_{5,3}$, *or* $E_{7,1}$.
2. *A dual polar space of rank three.*
3. *A product geometry,* $L \times P$, *where* L *is a line, and* P *is a non-degenerate polar space of rank at least two. (It may have infinite polar rank.)*

Any parapolar space which satisfies axioms (H1) (H2) and (H3) is called a *hexagonic geometry*.[5] In this way there is a clear distinction between the class of geometries defined by an axiomatic hypothesis and the classes of known "long-root geometries." Thus Theorem 17.1.1 can roughly be rendered as follows:

> *Any hexagonic geometry having symplectic rank at least three, and having all of its singular subspaces of finite projective rank is in fact one of the "long root geometries" listed at the begining of this section that have symplectic rank at least three (classes 1, 2, and 3 of that list).*

As usual, our proof concentrates on characterizing the point-residuals of a hexagonic geometry. That, of course, is the purpose of Theorem 2, which we take up next.

17.2 A Special Class of Strong Parapolar Spaces

In this section Γ is a strong parapolar space satisfying the following axioms:

(P1) If x is a point and S is a symplecton, then $x^{\perp} \cap S \neq \emptyset$.
(P2) For every point p, $\Delta_2^*(p)$ is a proper subspace of Γ.
(P3) If every symplecton has rank at least three, all singular subspaces are assumed to have finite projective rank.

---

[4] The reader may recognize this as condition (D2) of Sect. 13.7 with the extra requirement that there is a point at distance three from any point.

[5] This name is borrowed from a scheme devised by Prof. B. Muehlherr.

Remark Note that if p is a point and L is a line not contained in $\Delta_2^*(p)$, then by the parapolar hypothesis, there is a symplecton S containing L. Then by (P1) there is a point $r \in p^\perp \cap S$ which is collinear to a point on L, so L must carry a point of $\Delta_2^*(p)$. Thus we see that axioms (P1)+(P2) are equivalent to (P1)+(P2') where (P2') is the assertion that $\Delta_2^*(p)$ *is always a geometric hyperplane of* Γ.[6]

Remark Note that if all symplecta have rank at least three, then all singular subspaces of Γ are projective spaces. But without this assumption, it is conceivable, for the time being, that singular subspaces are not projective. Indeed, we shall prove that they are projective in a later corollary.

Our objective is to prove the second main theorem (Theorem 17.1.2) by showing the Γ is one of the following:

1. $D_{6,6}$, $A_{5,3}$, or $E_{7,1}$.
2. A classical dual polar space of rank three.
3. A product geometry $L \times P$, where L is a line and P is a polar space of arbitrary rank.

This is accomplished by a series of theorems.

Theorem 17.2.1 *Suppose a point x is distance three from a point p in a symplecton S. Then $x^\perp \cap S = \{r\}$, a single point, and $\Delta_2^*(x) \cap S = r^\perp \cap S$. The symplecton S is gated with respect to x.*

Proof Clearly $x^\perp \cap S$ is a non-empty singular subspace of S. Since S is a polar space and contains a point at distance three from x, $x^\perp \cap S$ is a single point set $\{r\}$. Then $r^\perp \cap S \subseteq \Delta_2^*(x) \cap S$, and equality now follows from the fact that in a polar space of rank at least two having thick lines, all geometric hyperplanes are maximal subspaces. In particular $r^\perp \cap S$ is a maximal singular subspace of S. □

Theorem 17.2.2 *In the point-collinearity graph $\Delta = (\mathcal{P}, \sim)$ of Γ, every geodesic path of length two extends to a geodesic path of length three.*

Proof Let (p, a, b) be a geodesic path of length two in $(\mathcal{P}, \sim)$. We wish to extend it to a geodesic path (p, a, b, c) of length three. Let z be any element in $\Delta_3(p)$ (exists by (P2)) and let S be the unique symplecton containing the geodesic (p, a, b). Then $z^\perp \cap S = \{r\}$ by (P1) and by Theorem 17.2.1 r is the gate of S with respect to z. We need only show that b is collinear to a point of $\Delta_3(p)$. This is implied by the following claim.

(Claim 1) *Every point of $S - p^\perp$ is collinear to a point of $\Delta_3(p)$.*

But since S has thick lines, the induced point-collinearity graph on $S - p^\perp$ is connected. Since r is a point of $S - p^\perp$ which is connected to a point of $\Delta_3(p)$, it suffices to prove the following.

[6] As in Chap. 4, the reader is warned that a few authors do not define geometric hyperplanes to be proper subspaces, as is done throughout this book.

(Claim 2) *If r_1 and x are collinear points of $S - p^\perp$ and r_1 is collinear with a point of $\Delta_3(p)$ then x is also collinear with a point of $\Delta_3(p)$.*

(Proof of Claim 2) By hypothesis r_1 is collinear with a point z_1 of $\Delta_3(p)$. Then x is distance two from z_1 and so there is a symplecton T containing $\{x, z_1\}$. Since S is a polar space, p is collinear with a point g on line xr_1 and g is distinct from both x and r_1 by hypothesis. Now, since T contains a point z_1 at distance three from p, by Theorem 17.2.1 T is gated with respect to p so $\Delta_2^*(p) \cap T = g^\perp \cap T$. But as T is a non-degenerate polar space g is the unique deep point of $g^\perp \cap T$, so there is a point of $x^\perp \cap T$ not contained in $g^\perp \cap T = \Delta_2^*(p) \cap T$. Thus x is collinear with a point of $\Delta_3(p)$ and Claim 2 is proved.

So, as remarked, the entire Theorem 17.2.2 is proved. □

Since we cannot assume that singular subspaces are projective, we must be careful about what we call a "plane." We know that Γ is a partial linear gamma space. We say that a singular subspace is a *plane* if it is generated by a non-incident point-line pair – that is, it has the form $\langle p, L\rangle$. Conceivably, one plane could properly contain another. But because of our strong parapolar hypotheses, we can show that all planes are projective planes.

Theorem 17.2.3 *Every plane $\pi := \langle a, L\rangle$ lies in some symplecton.*

Proof We begin by assuming that π lies in no symplecton.

Choose a point $b \in L$ and select a symplecton R on line ab (one exists by the parapolar hypothesis). Since $a^\perp \cap R$ is not a clique, there is a point x in R with $x^\perp \cap ab = \{a\}$. Now if $u \in x^\perp \cap L$, then the symplecton $\langle\langle x, b\rangle\rangle$ contains b, a, and $u \in L - \{b\}$ and so contains L (Γ is a partial linear space), and so contains π, contrary to assumption. Thus

$$x^\perp \cap L = \emptyset. \tag{17.1}$$

Now, by Theorem 17.2.2, the geodesic (x, a, b) extends to one of length three, say (x, a, b, z). Clearly if $z^\perp$ contained L, then the symplecton $\langle\langle z, a\rangle\rangle$ would contain L and a and hence would contain π, a contradiction. Thus $z^\perp \cap L = \{b\}$.

We observe that any symplecton on L is a rank two symplecton, or "quad." For otherwise, if such a symplecton, Q_0 had polar rank at least three, $L^\perp \cap Q_0$ would not be a clique, and so would contain a point v not in $a^\perp$, and then the symplecton $\langle\langle v, a\rangle\rangle$ would contain π, against our assumption.

Now let $c \in L - \{b\}$ and let $Q = \langle\langle z, c\rangle\rangle$. By the previous paragraph, Q is a quad. Then by Theorem 17.2.1 there is a unique point $w \in x^\perp \cap Q$ and $w^\perp \cap Q = \Delta_2^*(x) \cap Q$ contains L. But as observed, L is a maximal singular subspace of the quad Q, and so $w \in L$. But this contradicts Eq. (17.1), completing the proof. □

Corollary 17.2.4 *All singular subspaces of Γ are projective spaces.*

Proof By Theorem 17.2.3, all planes of such a subspace are projective.

Corollary 17.2.5 *Every pair of distinct intersecting lines lies in a symplecton. If the two lines do not lie in a singular subspace, the symplecton containing them is unique.*

Theorem 17.2.6 *If an intersection of two symplecta contains a point, it also contains a line.*

Proof Suppose, by way of contradiction, that S_1 and S_2 are distinct symplecta whose intersection is a single point p. Choose a geodesic (p, b, x) in S_1 and extend it to a geodesic (p, b, x, q) (permitted by Theorem 17.2.2). Now by hypothesis, $x^\perp \cap S_2$ contains a point y. If $x^\perp \cap S_2$ contained a line, this line would contain a point of $p^\perp - \{p\}$ lying in $S_1 \cap S_2$, contrary to our assumption. Thus $x^\perp \cap S_2 = \{y\}$.

Now either (y, x, q) is a geodesic of length two, or $\{y, x, q\}$ lies in a plane. Using Theorem 17.2.3 in the last case, we see that in all cases, $\{y, x, q\}$ lies in a symplecton R. Now R is gated with respect to p, so $p^\perp \cap R = \{t\}$, and $\Delta_2^*(p) \cap R = t^\perp \cap R$. But the latter set contains both x and y, whence

$$t \in p^\perp \cap x^\perp \cap y^\perp \subseteq S_1 \cap S_2,$$

by the convexity of these symplecta. This forces $t = p$, which is impossible since $d(p, q) = 3$. □

Theorem 17.2.7 *If $x^\perp \cap S$ is a single point for some symplecton S, then S contains a point at distance three from x. In particular, S is strongly gated with respect to point x.*

Proof Suppose, for some point x and symplecton S, that $x^\perp \cap S = \{y\}$. If S were not strongly gated with respect to x, there would be a point z in S at distance two from y which was also distance two from x. In that case there is a symplecton R on $\{x, z\}$. Now by Theorem 17.2.6 the intersection of R and S contains a line L on z. Then as L and x are in symplecton R, $x^\perp$ meets L at a point v. But then $v \in x^\perp \cap S = \{y\}$, so $v = y$. But that contradicts $d(y, z) = 2$. Thus all points of $S - y^\perp$ are at distance three from x and S is gated with respect to x. □

We require a minor lemma.

Lemma 17.2.8 *If all symplecta have rank at least three, then the point-collinearity graph of Γ is simply connected.*

Proof Axiom (P2) of the hypotheses introduced at the beginning of this section shows that the point-collinearity graph has diameter three. Moreover, since symplecta have rank at least three, any circuit within a symplecton is contractible. Thus we need only show that any circuit of length seven or less, decomposes into circuits of length three or four.

Suppose $c = (x_0, x_1, \ldots, x_5)$, $x_5 = x_0$, is a circuit of length five not decomposable into circuits of length three or four. Then x_0 is distance two from both x_2 and x_3, and the intersection of the two symplecta R and S on $\{x_0, x_2\}$ and $\{x_0, x_3\}$ contains a line L on x_0. Now if $R = S$, c is contractible. So we may assume x_2 is not in S. Then $x_2{}^\perp \cap S$ is a clique containing x_3 and a point a on L. Now c decomposes

into the circuits (x_0, x_1, x_2, a, x_0), (x_3, a, x_2, x_3), and (x_0, a, x_3, x_4, x_0) of lengths 4, 3, and 4, a contradiction. Thus all five-circuits are decomposable into circuits of length three or four.

If c is a six-circuit with an antipodal pair of vertices at distance two, it decomposes into two five-circuits. Otherwise, the 6-circuit $(x_0, x_1, x_2, x_3, x_4, x_5, x_0)$ contains an antipodal pair at distance three – say, (x_0, x_3). Then by Theorem 17.2.6, the two symplecta $R = \langle\langle x_0, x_2\rangle\rangle$ and $S = \langle\langle x_0, x_4\rangle\rangle$ must intersect at a line L on x_0. Since R and S are polar spaces, there exist points r and s such that

$$\{r\} = {x_2}^{\perp} \cap L, \{s\} = {x_4}^{\perp} \cap L.$$

Now r and s lie in the subspace $\Delta_2^*(x_3) \cap L$. Since $d(x_0, x_3) = 3$, we see that $\Delta_2^*(x_3) \cap L = \{r\} = \{s\}$. Thus the 6-circuit $(x_0, \ldots, x_5, x_0)$ decomposes into three 4-circuits. Thus all 6-circuits are decomposable.

Suppose now that $c = (x_0, x_1, \ldots, x_7)$, $x_7 = x_0$, is an indecomposable 7-circuit. Then we must have $d_\Gamma(x_0, x_3) = d_\Gamma(x_0, x_4) = 3$. Now the hyperplane $\Delta_2^*(x_0)$ meets the line N on x_3 and x_4 at a point b. Choosing a in ${x_0}^{\perp} \cap b^{\perp}$, we see that c decomposes into the two 6-circuits, $(x_0, x_1, x_2, x_3, b, a, x_0)$ and $(x_0, a, b, x_4, x_5, x_6, x_0)$. This contradicts the indecomposability of c. □

Theorem 17.2.9 *If all the symplecta of* Γ *have rank at least three, then* Γ *is isomorphic to one of the following Lie incidence geometries:*

$$D_{6,6},\ A_{5,3},\ \textit{or}\ E_{7,7}\ \textit{(Bourbaki numbering)}.$$

Proof We shall show that Γ satisfies condition (U_2). Let S_1 and S_2 be two symplecta which intersect at a subspace $A = S_1 \cap S_2$ which contains a plane. Choose an arbitrary point x_1 in $S_1 - A^{\perp}$. We are obliged to show that ${x_1}^{\perp} \cap S_2$ contains a point not in A.

From the way x_1 was chosen, ${x_1}^{\perp} \cap A = H_1$ is a hyperplane of A. Choose any point $z \in A - {x_1}^{\perp}$, so $d(x_1, z) = 2$. By Theorem 17.2.2, any geodesic (z, a, x) connecting z and x can be extended to one of length three — say (z, a, x, y). Since $d(x, y) = 3$, the point y does not lie in $S_1 \cup S_2$. But by (P1), there is a point $x_2 \in y^{\perp} \cap S_2$. Since $z \in A = S_1 \cap S_2$, Theorem 17.2.1, p. 607, tells us that (y, S_i) is a gated pair, so that for $i = 1, 2$,

$$y^{\perp} \cap S = \{x_i\}, \text{ and}$$
$$\Delta_2^*(y) \cap S_i = {x_i}^{\perp} \cap S_i.$$

The subspaces $H_i := {x_i}^{\perp} \cap A$, $i = 1, 2$, are both hyperplanes of A.

Now suppose $H_1 \cap H_2$ were a single point. This could only happen if A were a plane. Then, since (P2) asserts that $\Delta_2^*(y) \cap A$ is a subspace of A containing H_1 and H_2, the existence of $z \in A \cap \Delta_3(y)$ forces $H_1 = H_2$, which contradicts $H_1 \cap H_2$ containing a single point. Thus we see that $H_1 \cap H_2$ contains a line L.

Suppose x_1 were not collinear with x_2. Then there would be a symplecton $T := \langle\langle x_1, x_2\rangle\rangle$, containing

$$x_1{}^{\perp} \cap x_2{}^{\perp} \supseteq \{y\} \cup \{H_1 \cap H_2\} \supseteq \{y\} \cup L.$$

Then, as T is a polar space, there would be a point $w \in y^{\perp} \cap L$ and now $w \in y^{\perp} \cap z^{\perp}$ against $d(y, z) = 3$.

Thus it follows that x_1 is indeed collinear with x_2.

Summarizing what we have proved: if $S_1 \cap S_2$ are two distinct symplecta whose intersection A contains plane, then for any $x_1 \in S_1 - A^{\perp}$, there is a point $x_2 \in x_1{}^{\perp} \cap (S - A)$. Thus condition (U_2) holds for Γ.

Next we may conclude that Γ is locally connected since it is a strong parapolar space of symplectic rank at least three (see Lemma 13.5.1, p. 467).

Finally, by (P3), we are assured that every singular subspace of Γ has finite projective rank whenever Γ has symplectic rank at least four.

At this point, Γ satisfies all of the hypotheses of Theorem 15.5.4, p. 545, and so Γ is one of the following:

1. A polar space.
2. A Grassmann space $A_{n,k}$.
3. $A_{2k-1,k}/\langle\sigma\rangle$, a homomorphic image of a Grassmannian, $A_{2k-1,k}$ where σ is a polarity of Witt index at most $k - 4$.
4. A homomorphic image of a classical halfspin geometry $D_{n,n}$.
5. The Lie incidence geometry $E_{6,1}$ or $E_{7,7}$.

Now we have the additional properties that Γ has point-diameter at least three, and that for any point-symp pair (x, S), $x^{\perp} \cap S$ is non-empty. This eliminates the first case, and in the second case, it eliminates all but the Grassmannian $A_{5,3}$. The third case is eliminated altogether since the geometry there has point-diameter at least four.

In the fourth case above, the parapolar space $\Gamma = (\mathcal{P}, \mathcal{L})$ is enriched to a rank five geometry $\Gamma = (\mathcal{P}, \mathcal{L}, \mathcal{M}, \mathcal{A}3, \mathcal{S})$ over $J = \{n, n-1, \ldots n-4\}$, which is a locally truncated geometry relative to the diagram in Fig. 17.1, where $\mathcal{M}$ and $\mathcal{A}3$ are two classes of maximal singular subspaces, and $\mathcal{S}$ is the class of symplecta. (The planes are $\mathcal{M} \times \mathcal{A}3$-flags. This much is in [36, 60].)

Now by the Remark (1) following Theorem 11.8.1, hypotheses 1, 2, and 4 of that theorem are in place, where Π is the collection of all projective subplanes. Since

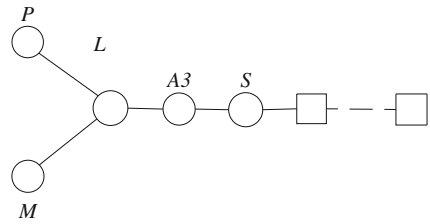

Fig. 17.1 Locally truncated diagram for the fourth case

the residues of any point or line in Γ are truncations of buildings, condition 3 of Theorem 11.8.1 holds (Remark (2) following Theorem 11.8.1).

Thus by Theorem 11.8.1, the geometry epimorphism $f : \Delta_J \to \Gamma$ obtained from Lemma 11.6.1 induces a $\mathcal{T}$-covering

$$\phi : G \to \bar{G}$$

of their point-collinearity graphs. But since the point-collinearity graph $\bar{G}$ is simply connected by Lemma 17.2.8, ϕ is a graph isomorphism. Thus Γ has the point-collinearity graph of the half-spin geometry. But all objects in the rank n building geometry Δ are induced subgraphs of its half-spin collinearity graph, and so the morphism f is also a geometry isomorphism. Then Γ is a bona fide half-spin geometry. Now the two special properties of Γ force $\Gamma \sim D_{6,6}$.

In the last case $E_{6,1}$ is eliminated on the basis of its point-diameter. But for $E_{7,7}$ all the axioms hold, so this case survives. The proof of Theorem 17.2.9 is complete. □

Theorem 17.2.10 *If Γ contains a symplecton which is a grid, and a symplecton which is not, then Γ is the product geometry: $L \times P$, where L is a line and P is a polar space of arbitrary rank*

We begin the proof of Theorem 17.2.10 with a series of lemmas.

Lemma 17.2.11 *Suppose G, S_1, and S_2 are three symplecta whose intersection $G \cap S_1 \cap S_2$ is a point p. We suppose that G is a grid. Then at least one of the S_i is also a grid.*

Proof From Theorem 17.2.6 and the hypotheses, $G \cap S_i = L_i$, $i = 1, 2$ are the two distinct lines of G on p and $S_1 \cap S_2$ is a singular subspace meeting the L_i at p and containing a line N on p. Suppose by way of contradiction that neither S_i was a grid. Then the lines and planes of S_i which are incident with point p is either a rank one polar space (or coclique) with at least three points, or is a polar space of rank at least two. In either case, it is not the union of a clique and the perp of a point. Thus each symplecton S_i contains a line N_i on p which is not in $L_i^{\perp}$ or in the singular space $S_1 \cap S_2$. By Corollary 17.2.5, there is a symplecton R containing $N_1 \cup N_2$. But then by Theorem 17.2.6, $R \cap G$ must be one of the two lines of G on p, L_1, or L_2. But if $R \cap G$ were the line L_1, then R would be the unique symplecton containing $L_1 \cup N_1$, namely S_1, and that would imply that line N_2 was in S_1, contrary to choice. But also $R \cap G$ cannot be the other line L_2 by the complete symmetry of the indices $i = 1, 2$ in the face of the hypotheses and the choice of the lines N_i. This contradiction confirms the assumption that neither S_i was a grid, and the proof of the lemma is complete. □

Lemma 17.2.12 *Suppose G is a symplecton which is a grid, and that G intersects a second symplecton Q, which is not a grid, at a line L_1. Choose a point p on L_1, let L_2 be the other line of G on p distinct from L_1, and choose any point q in $L_2 - \{p\}$. The following statements must hold:*

1. *With the exception of the line L_2, every line on p lies in Q.*
2. *$q^\perp \cap Q = \{p\}$.*
3. *Any symplecton containing p which is distinct from Q is a grid.*
4. *Let $\mathcal{L}_p$ and $\mathcal{L}_q$ be the collections of all lines of Γ on points p and q, respectively. Then there is a bijection*

$$\beta : \mathcal{L}_q - \{L_2\} \to \mathcal{L}_p - \{L_2\},$$

such that for corresponding lines $L \in \mathcal{L}_q - \{L_2\}$ and $\beta(L)$, there is a second bijection $L \to \beta(L)$ taking each point r of L to the unique point of Q to which it is collinear.

Proof 1. Suppose M were a line on p distinct from L_2 and not lying in Q. Then by Corollary 17.2.5, $M \cup L_2$ lies in some symplecton S. By Theorem 17.2.6 the intersection $S \cap Q$ contains a line N on p which is necessarily distinct from L_1. But now Q and S meet G at distinct lines, and so the intersection of all three is just the single point p. Since G is a grid and Q is not, Lemma 17.2.11 implies that S is a grid. But that is impossible since L_2, M, and N are three distinct lines of S on point p. Thus no such line M exists, which proves the result.

2. Suppose by way of contradiction that $q^\perp \cap Q$ contained a line N and let π be the projective plane generated by q and N. Clearly $\pi \cap Q = N$. It follows that there is a line of π on p distinct from N and L_2, against part 1. Thus $q^\perp \cap Q$ can only contain the point p.

3. Suppose S is a symplecton on p distinct from Q. Then $p^\perp \cap S$ consists of the singular subspace $S \cap Q$, which at least contains a line on p, and the unique line L_2 (part 1) on p which is not in Q. Precisely, the collinearity graph induced on $p^\perp \cap S$ is the union of two cliques. It follows that S is a grid.

4. We define the mapping $\beta : \mathcal{L}_q - \{L_2\} \to \mathcal{L}_p - \{L_2\}$ as follows. For each line L on q, there is a symplecton R on $L \cup L_2$ (Corollary 17.2.5). By part 3, R is a grid, so L is not in $L_2^\perp$ and so R is the unique symplecton on L and L_2. Thus, using Theorem 17.2.6, $R \cap Q$ is a line $\beta(L)$ uniquely determined by L. Note that L and $\beta(L)$ are opposite lines of the grid R, and so the desired bijection $L \to \beta(L)$ exists.

Suppose now that $\beta(L) = \beta(L')$ and let R and R' be the unique grids on $L \cup L_2$ and $L' \cup L_2$ as in the previous paragraph. Then R and R' are the unique symplecta on $\beta(L) \cup L_2$, forcing $R = R'$ and $L = L'$. Thus β is injective.

If N is any line of $\mathcal{L}_p - \{L_2\}$, then by part 1, N is not in $L_2^\perp$ and the unique symplecton T on $N \cup L_2$ is a grid by part 3. Then $N = \beta(N')$ where N' is the unique line of T on q which is distinct from L_2. Thus β is onto. The proof Part 4 is complete. □

Lemma 17.2.13 *Again let G, Q, L_1 be as in Lemma 17.2.12. Every line of Q lies in exactly one further symplecton which is a grid. Consequently, every point of Q is incident with exactly one line which is not in Q.*

Proof Choose point p in L_1 and let L_2 be the other line of G on p as in Lemma 17.2.12. Let L be any line of Q. If L is incident with p then the unique

grid containing $L \cup L_2$ is the only symplecton on L besides Q. Suppose then L is not incident with p. Then there is a point r in L collinear with p. Without loss of generality, L_1 can be taken to be the line on p and r and G to be the grid on $L_2 \cup L_1$. Then r enjoys the same hypotheses that p did in Lemma 17.2.12. So there is a unique line L_3 on r which is not in Q. Then the symplecta on L are Q and the unique symplecton on $L \cup L_3$, which, by part 3 of Lemma 17.2.12, is a grid.

The uniqueness of the out-going lines follows. □

Proof of Theorem 17.2.10 Suppose the strong parapolar space Γ contains a symplecton which is a grid and one which is not. Since the point-collinearity graph is connected and every line of Γ lies in at least one symplecton (Corollary 17.2.5 for example), there must be an instance in which a symplecton G which is a grid, intersects a symplecton Q which is not a grid at some point. Then by Theorem 17.2.6, $G \cap Q$ is a line L_1.

Now choose a point p on L_1, let L_2 be the unique second line of G on point p, and choose point q in $L_2 - \{p\}$. We now have the situation of Lemma 17.2.12.

We know that we can choose at least two distinct lines, N_1 and N_2, in $\mathcal{L}_q - \{L_2\}$. Let Q_q be any symplecton containing $N_1 \cup N_2$. We make three claims:

1. $Q_q \cap Q = \emptyset$.
2. The symplecton Q_q is isomorphic to Q, and so is not a grid.
3. Q_q contains all lines on q except L_2.

First suppose Q_q and Q had a non-empty intersection. Then, according to Lemma 17.2.13, Q_q would be a grid intersecting Q in a line L. Moreover, since $q^\perp \cap L$ is non-empty, L must contain p, the unique point of $q^\perp \cap Q$ (Lemma 17.2.12, part 2). It follows that Q_q contains L_2, and so L_2, N_1, and N_2 comprise three distinct lines of Q_q on q. That is impossible since Q_q was a grid in this case division.

We now address the second and third claims. By our basic hypothesis on Γ, each point of Q_q is collinear with at least one point of Q, and, since $Q_q \cap Q = \emptyset$ and each point of Q lives on only one out-going line, this point must be unique. Thus there is an injective mapping $\phi : Q_q \to Q$ taking each point of Q_q to the unique point of Q with which it is collinear. But also by the fundamental hypothesis, each point of Q is collinear with at least one point of Q_q and so ϕ is a bijection. Using the presence of the unique system of interlocking grids, it is easy to see that ϕ and ϕ^{-1} both preserve the collinearity relation on points. Thus ϕ induces a bijection β of the lines of Q_q with those of Q, extending the bijection β of Lemma 17.2.12, part 4. This establishes the second and third claims.

Now set $Q = Q_p$, and for each point x of symplecton Q, let L_x be the unique line on x not in Q. (In this notation, L_2 is now L_p.) From what we have established, each of these "out-going" lines L_x meets each Q_q at a single point, each point in any Q_y lies on a unique one of these L_xs, and has all its remaining lines in Q_y. It follows that the union of the disjoint Q_ys, as y ranges over the points of L_p, is a connected component of the collinearity graph of Γ and hence covers all of the points. Thus every point of $\mathcal{P}$ can be coordinatized as (x, y) where the point x of $Q = Q_p$, indexes the unique line L_x connecting it to Q (or is the point itself, if it

already is in $Q = Q_p$), while the coordinate y is the point of L_p which indexes the unique Q_y in which the point lies. All lines are now either the "horizontal" lines of one of the symplecta Q_y which partition the points, or one of the "vertical" lines L_x, $x \in Q_p$. Thus we have a product geometry $L_p \times Q_p$. The proof is complete. □

Remark Note that in this case it is possible for the vertical lines to possess a different cardinality than that for the horizontal lines.

Theorem 17.2.14 *If all symplecta are generalized quadrangles, Γ is a dual polar space of rank three or the product geometry $L \times Q$ of a line L and a generalized quadrangle Q. (Of course if Q is itself a grid, Γ is just the "Hamming cube" – that is, the product of three lines $L_1 \times L_2 \times L_3$, where three line cardinalities are possible.) In each case, Γ is a near hexagon of classical type.*

Proof By the Theorem of Yanushka–Shult [119] (a special case of Cameron's theorem [24] (Theorem 8.5.17 of Chap. 8)) it suffices to show that Γ is a near hexagon with all quads "classical" – that is each quad is gated with respect to every exterior point.

If there were a plane in Γ, by Theorem 17.2.3, it would lie in some symplecton. But that would be impossible since each symplecton is a generalized quadrangle. So there are no planes. Yet, by hypothesis, for each non-incident point-symplecton pair (p, Q), the intersection $p^\perp \cap Q$ is not empty. Since there are no planes, the intersection $p^\perp \cap Q$ is always a single point. By Theorem 17.2.7, the symplecton Q is strongly gated with respect to x. Thus we see that any symplecton of Γ is a quadrangle with the property that it is strongly gated with respect to every exterior point, and that every such point is collinear with exactly one of its points. This makes Γ a near hexagon of classical type. □

Theorem 17.2.15 *If Γ contains no grids, and at least one symplecton has rank at least three, then all symplecta have rank at least three and the conclusion of Theorem 17.2.9 holds.*

We first prove the following technical lemma.

Lemma 17.2.16 *Suppose Γ contains no grids. Suppose (a_1, x, y) and (a_2, x, y) are two geodesics, $i = 1, 2$. Then there is a point b in $y^\perp$ which is simultaneously distance three from both a_1 and a_2.*

Proof By Theorem 17.2.2, there is a point, b_i such that (a_i, x, y, b_i) is a geodesic of length three for $i = 1, 2$. If $b_1 = b_2$ we are done, so assume the b_i are distinct. By either the strong parapolar hypothesis or Theorem 17.2.3, there is a symplecton R on $\{b_1, y, b_2\}$. Now by Theorem 17.2.1, for $i = 1, 2$, there exist points r_i, such that

$$a_i{}^\perp \cap R = \{r_i\} \text{ and } \Delta_2^*(a_i) \cap R = r_i{}^\perp \cap R.$$

It may happen that $r_1 = r_2$, but in any case, both are distinct from y since $d(a_i, y) = 2$ for both values of i. In any case, the set $\mathcal{L}_i$ of lines of R on point y which lie in $\Delta_2^*(a_i)$ are just those in $(r_i\, y)^\perp$.

If R has rank at least three, the lines and planes of R on y form a polar space with thick lines (with some abuse of notation we may write this as $\mathrm{Res}(y) \cap R$), of which the two sets $\mathcal{L}_i$ form hyperplanes. Since no polar space with thick lines is the union of two hyperplanes, there is a line yb in $\mathrm{Res}(y) \cap R$ which is in neither of these two hyperplanes.

If, on the other hand, R is a generalized quadrangle, one has $(r_i y)^\perp \cap R = r_i y$. In this case, since R is not a grid, there exists a line yb not in either $\Delta_2^*(a_i) \cap R$ for $= 1, 2$.

In all cases $d(b, a_i) = 3$, $i = 1, 2$, as required. □

Proof of Theorem 17.2.15 Under the hypothesis of no grids we shall show that any symplecton which intersects a symplecton of rank at least three non-trivially must itself have rank at least three. It will then follow from the connectedness that all symplecta have rank at least three.

So suppose S_1 is a symplecton of rank at least three and S_2 is a second symplecton intersecting S_1 non-trivially. Our objective is to show that S_2 has rank at least three. By Theorem 17.2.6, $S_1 \cap S_2$ contains a line. If the intersection contains a plane, S_1 has rank at least three and we are done. So we may assume that $S_1 \cap S_2$ is exactly a line L. Choose distinct points x and y on line L, and points a_i in $x^\perp \cap S_i - y^\perp$, $i = 1, 2$. Then (a_i, x, y) are geodesics of length two which, by the previous lemma, can be extended in a common way to geodesics (a_i, x, y, b), $i = 1, 2$. Note that these metric requirements show that b cannot be in either S_1 nor S_2.

Now by hypothesis there is a plane π in S_1 on line L, and so there is a point z in $\pi - L$. Now if b were collinear with z, we would have $b^\perp \cap S_1$ containing the line yz meeting ${a_1}^\perp$, against $d(b, a_1) = 3$. Thus $d(b, z) = 2$, and as Γ is a strong parapolar space, there is a symplecton S on $\{z, y, b\}$. Then $S \cap S_2$ contains a line M.

Obviously, as b is distance three from a_2, $S \cap S_2$ cannot contain a plane. Thus $S \cap S_2 = M$ exactly.

Now the unique point r on ${a_2}^\perp \cap M$ is distinct from y, and by Theorem 17.2.1 is the unique point of ${a_2}^\perp \cap S$ (the gate) and

$$\Delta_2^*(a_2) \cap S = r^\perp \cap S.$$

Why can't $M = L$, so that $r = x$? In that case $S \cap S_1$ contains the plane $\langle L, z \rangle = \pi$ forcing $d(b, a_1) = 2$, a contradiction. Thus M is not in S_1, for otherwise $M \subseteq S_2 \cap S_1 = L$. Thus $M \cap S_1 = \{y\}$, and r is not in S_1.

Since M is a thick line, there is a point s in M distinct from r and y. Moreover, since S is a non-degenerate polar space, there is a point b' in S collinear only with the point s of M. Since b' is not in $r^\perp$, we have that $d(a_2, b') = 3$. Thus we see that

$$S_2 \cap \Delta_2^*(b') = S_2 \cap s^\perp. \tag{17.2}$$

Now $x^\perp \cap S$ contains the line yz and so $d(b', x) = 2$. It follows from Eq. (17.2) that s is collinear with x. We now see that the clique $\{s, x, y\}$ spans a plane in S_2, and so S_2 must have rank at least three. The proof is complete. □

Proof of Theorem 17.1.2: It is given that Γ is a strong parapolar space. By hypotheses (1), (2), and (3), conditions (P1), (P2), and (P3) hold.

We have the following trichotomy:

(A) Γ contains no symplecton that is a grid, and at least one symplecton has polar rank three.
(B) Γ has symplectic rank two.
(C) Γ contains a symplecton that is a grid and another symplecton of polar rank at least three.

In case (A), Theorem 17.2.15 forces symplectic rank at least 3, and (since Γ is then locally connected with all maximal singular subspaces of finite projective rank) one has $\Gamma \simeq D_{6,6}$, $A_{5,3}$ or $E_{7,7}$ by Theorem 17.2.15.

In case (B), Theorem 17.2.14 makes Γ isomorphic to (1) $L \times Q$ where L is a line and Q is a generalized quadrangle, or (2) a dual polar space obtained from a polar space of rank three.

In case (C), Theorem 17.2.10 implies $\Gamma \simeq L \times P$ where L is a line and P is a polar space of arbitrary polar rank.

Thus, in every case some part of our conclusion holds. □

17.3 The Proof of Theorem 17.1.1

In this section, Γ is a parapolar space of polar rank at least three satisfying these hypotheses:

(A1) Given a point x not incident with a symplecton S, the space $x^\perp \cap S$ is never just a point.
(A2) Given a projective plane π and line L meeting π at point p, either (1) every line of π on p shares a symplecton with L, or else (2) exactly one such line incident with (p, π) has this property.
(A3) If L is a line on point p, then there exists at least one further line N on point p such that

$$L^\perp \cap N^\perp = \{p\}.$$

(That means that there is no symplecton which contains $L \cup N$.)
(A4) If all symplecta have rank at least four, then every maximal singular subspace has finite projective rank.

Remark Note that axiom (A3) prevents Γ from being a strong parapolar space.

17.3.1 Simple-connectedness of the Point-Collinearity Graph of Γ

It is useful at the beginning to show that the axioms (A1), (A2), and (A3) alone force the point-collinearity graph $G = (\mathcal{P}, \sim)$ of Γ to be simply $\mathcal{T}$-connected – that is,

every cycle of the graph can be deformed to a single point by some iteration of the processes of either replacing an edge representing one side of a triangle by the other two sides, or the reverse of this process.

We define the *angle* between lines L_1 and L_2 of $\mathcal{L}_p$ as the integer $\alpha(L_1, L_2)$ where

$$\alpha(L_1, L_2) = \begin{cases} 0 & \text{if } L_1 = L_2 \\ 1 & \text{if } L_1 \neq L_2 \text{ but both are in a plane} \\ 2 & \text{if } L_1 \neq L_2 \not\subseteq {L_1}^{\perp} \text{ but both lie in a symplecton} \\ 3 & \text{if } L_1 \text{ and } L_2 \text{ lie together in no symplecton.} \end{cases}$$

(Note that in the last case, if $y_i \in L_i - \{p\}$, then ${y_1}^{\perp} \cap {y_2}^{\perp} = \{p\}$, that is (y_1, y_2) is a special pair.)

Suppose now $w = (x_0, x_1, \ldots, x_n)$ is a path in the point-collinearity graph G. Let L_i be the unique line on $\{x_i, x_{i+1}\}$, and let α_{i+1} be the angle between L_i and L_{i+1}. Then the sequence $(\alpha_1, \alpha_2, \ldots, \alpha_n)$ is the *angle type* of the path w.

Theorem 17.3.1 *The point-collinearity graph $\Gamma = (\mathcal{P}, \sim)$ is simply connected*

Proof We must show that every circuit in G is $\mathcal{T}$-contractible where $\mathcal{T}$ is the collection of all triangles of G. Assume $c = (x_0, x_1, \ldots, x_n = x_0)$ is a circuit of minimal length n subject to being non-contractible. We can assume $n > 3$, and since any 4-circuit lies in a symplecton of polar rank ≥ 3 (whose circuits are easily seen to be contractible), we can assume $n > 4$ as well.

We can also assume that x_i is not collinear with x_j for $n - 1 > |i - j| > 1$. If R were a symplecton on $\{x_1, x_2, x_3\}$, then ${x_0}^{\perp} \cap R$ would contain a line L, and ${x_3}^{\perp} \cap L$ would contain a point z. Then c would decompose as a circuit $c' = (x_0, z, x_3, \ldots, x_n = x_0)$ of length $n - 1$, a triangle, and a circuit of length 4, all of which are contractible by the conditions on n. Since this would make c contractible, no such symplecta can lie on $\{x_1, x_3\}$ – or on $\{x_i, x_{i+2}\}$ for that matter. Thus c has angle type $(3, 3, \ldots, 3)$.

Now choose a plane π on $\{x_1, x_2\}$. Now by axiom (A2) there is a line L of π on x_1 such that L and line x_0x_1 lie in a symplecton R. A second application of (A2) similarly produces a line N of π on x_2, sharing a symplecton S with line x_2x_3. Since both L and N are lines of π they must intersect at a point z. Now by (A1), ${x_4}^{\perp} \cap S$ contains a line B on x_3, which by the polar space property for S, bears a point b of $z^{\perp}$. Similarly ${x_{n-1}}^{\perp} \cap R$ contains a line A on x_0, which in turn also bears a point a of $z^{\perp}$ as $z \in R$. (Note that if $n = 5$ then $x_4 = x_{n-1}$; but this doesn't hurt anything.) Now it is clear that c decomposes into a circuit

$$c' = (a, z, b, x_4, \ldots, x_{n-1}, a)$$

of length $n - 1$, two circuits (x_0, a, z, x_1, x_0) and (x_3, b, z, x_2, x_3) of length at most 4, and three triangles $(x_{n-1}, x_0, a, x_{n-1})$, (x_1, z, x_2, x_1) (in π), and (x_4, b, x_3, x_4), all of which are contractible. Hence c is contractible.

Thus no non-contractible circuits exist, and the theorem is proved. □

17.3.2 The Uniform Structure of the Point-Residuals

As usual, we let $\mathcal{P}$, $\mathcal{L}$, Π, and $\mathcal{S}$ be the set of points, lines, planes, and symplecta of Γ. Then the sets $\mathcal{L}_p$, Π_p, and $\mathcal{S}_p$ are the lines, planes, and symplecta incident with a point p.

Recall that for each point p of Γ, the geometry $\mathrm{Res}(p) := (\mathcal{L}_p, \Pi_p)$ of lines and planes on p is a point-line geometry which is a strong parapolar space with all singular subspaces projective, whose "symplecta" are the lines and planes incident with a flag $(p, S) \in \mathcal{P} \times \mathcal{S}$. By (A1), $\mathrm{Res}_\Gamma(p)$ satisfies the property that each "point" is collinear with at least one point of any "symplecton" which does not contain it. Similarly (A2), (A3), and (A4) force $\mathrm{Res}_\Gamma(p)$ to satisfy the rest of the hypotheses of Theorem 17.1.2. It follows that $\mathrm{Res}_\Gamma(p)$ is isomorphic to one of the geometries appearing in Theorem 17.1.2. But is it the same geometry for each point p?

Suppose $L = pq$ is a line. Then both of the geometries $\mathrm{Res}(p)$ and $\mathrm{Res}(q)$ look the same above a "point" L. Thus when $\mathrm{Res}(p)$ is (a) $D_{6,6}$, (b) $A_{5,3}$, (c) $E_{7,1}$, (d) a dual polar space of rank three, or (e) the product $L \times P$ of a line and a polar space, the subgeometries of symps and singular spaces containing L are respectively, (a') $A_{5,2}$, a Grassmannian, (b') $A_2 \times A_2$, the product geometry of two planes, (c') the exceptional geometry $E_{6,1}$, (d') a projective plane A_2 of lines and symps, or (e') the disconnected union of a point and the point-residual of a symp. So each case is distinctive. That means that $\mathrm{Res}(p)$ is of the same type and defined by the same parameters as $\mathrm{Res}(q)$, except possibly in the last case. If the line pq is the isolated "point" of each point-residue, the symplecta forming a bouquet over p might conceivably be of a different isomorphism type than the symplecta forming a bouquet over q. But even here, the uniformity holds and is discussed fully in Lemma 17.3.4 of the next section.

So we have:

Lemma 17.3.2 *For any two points of* Γ*, the point-residual geometries are uniformly isomorphic.*

From this point onward (with a minor abuse of notation) we regard Γ as a higher rank "enriched" geometry over a typeset J singling out points, lines symplecta, and all singular subspaces. Some of these isomorphism types are sorted into further classes by the nature of the uniform point-residuals.

17.3.3 The Case of Finite Singular Rank

Assume now that every maximal singular subspace of Γ has finite projective rank.

It follows from Lemma 17.3.2 and Theorem 17.1.1 that Γ is a geometry belonging to one of the following locally truncated diagrams over I (the set J which is the recipient of the truncation is indicated by the round nodes in the figure below).

However, in order to define a sheaf, we must be sure that one can define types corresponding to the objects one sees in a point-residual. Any fusion of types would require an automorphism of the locally truncated diagram fixing the node "$\mathcal{P}$",

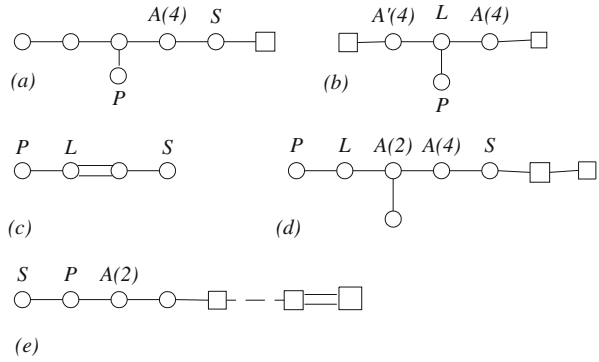

Fig. 17.2 Locally truncated diagrams associated with Γ. In case (b) one must make sure the types are not fused in Γ

and this is only possible in the case that the enriched $\mathrm{Res}(p)$ is a truncation of $A_{5,3}$ as in (b) of Fig. 17.2.[7] Here it is conceivable that the two local classes of $PG(4)$ (denoted $A(4)$ and $A'(4)$) are fused in the global geometry Γ. If so, one can invoke Theorem 12.2.1 (see p. 446) to conclude that there is a geometry morphism $\gamma : \hat{\Gamma} \to \Gamma$ such that, first, the point-residuals of $\hat{\Gamma}$ are mapped isomorphically onto the point-residuals of Γ so that $\hat{\Gamma}$ also belongs to the locally truncated diagram (1) of Fig. 17.2, and, second, in $\hat{\Gamma}$ the two classes $A(4)$ and $A'(4)$ are not fused. But Theorem 11.8.1 also asserts that γ induces a $\mathcal{T}$-covering $\delta = \delta|_Y$ of the point-collinearity graph of the Γ, and by Theorem 17.3.1 that graph covering is an isomorphism. Thus γ induces a bijective mapping on points and since all other objects of $\hat{\Gamma}$ are uniquely determined by their point-shadows, γ is an isomorphism of geometries. Thus from the second property above one can assign a distinct type to the objects of the locally truncated diagram for Γ just as in all the other cases depicted in Fig. 17.2. Thus, because there is an unambiguous assignment of types, Γ is a locally truncated diagram geometry with respect to the diagram D.

Now, by Theorem 11.3.2, there exists a sheaf, and since in each case the ambient diagram is a Dynkin diagram, there exists a building geometry Δ over I and a vertex-surjective morphism $h : \Delta \to \bar{\Delta}$ of geometries such that Γ is isomorphic to $\bar{\Delta}$ truncated to the typeset J: the typeset of Γ which includes $\{\mathcal{P}, \mathcal{L}, \mathcal{S}\}$ and all singular subspaces. Thus h induces a morphism

$$f : \Delta_J \to \Gamma \text{ (enriched to J)}$$

as in Lemma 11.6.1.

Now if we truncate to $\{\mathcal{P}, \mathcal{L}\}$, we recover the original parapolar space Γ (without enrichment) and a truncation of a building geometry $(\mathcal{P}_\Delta, \mathcal{L}_\Delta)$, and both of these point-line geometries are parapolar spaces of polar rank at least three. In each case, let Π and Π_Δ be the full sets of projective planes in these respective geometries. Then Hypothesis 1 of Theorem 11.8.1 holds, just from our choice of Π. Hypotheses

[7] Note that the planes Π are denoted '$A(2)$' in Fig. 17.2.

2 and 4 hold because they are parapolar spaces. We need to check Hypothesis 3 only when X is a line, having already established that the five cases listed above for point-residuals are truncations of rank at least three of buildings. But the same conclusions hold for a line L. This is because the lines are a set of flags which isolate the points $\mathcal{P}$ from all other nodes of the diagram. Thus the sheaf-value $\mathcal{F}(L)$ at a line L is a geometry belonging to a diagram of type

$$A_1 \times Y,$$

where Y is the diagram of a line-residual in the building $\mathcal{F}(p)$ for a point p. Since the latter is a building, so is its residual Y. Thus condition 3 of Theorem 11.8.1 is verified.

Now Theorem 11.8.1 applies to show the following.

(*) *The morphism f induces a graph morphism*

$$(\mathcal{P}_\Delta, \sim) \to (\mathcal{P}_\Gamma, \sim)$$

of the point-collinearity graph of Δ onto the point-collinearity graph of Γ, which is a $\mathcal{T}$-covering of graphs.

Now by Theorem 17.3.1 the latter graph is simply connected, so this graph morphism is actually an isomorphism. This means f induces a bijection on points. Now the fact that all objects of the building geometry Δ_J are distinguished by their point-shadows forces the morphism $f : \Delta_J \to \Gamma$ to be an isomorphism of parapolar spaces, and completes the proof of Theorem 17.1.1 when all maximal singular subspaces possess finite projective rank.

17.3.4 The Case of Infinite Singular Rank

The hypotheses $(A1) - (A4)$ show that if Γ possesses a singular subspace of infinite projective rank, then our point-residual $\mathrm{Res}_\Gamma(p)$ (necessarily a conclusion-geometry of Theorem 17.1.2) is the product geometry $L \times P$ where P is a polar space and L is a line.

Here we shall take this structure of P of $\mathrm{Res}_\Gamma(p)$ as a hypothesis, where P is any polar space that is not a grid. Of course that means we are reproving some of the finite singular rank cases over again, but this time without resorting to the theory of locally truncated geometries. All the better!

We first require a general theorem.

Theorem 17.3.3 *Suppose Γ is a parapolar space satisfying the following hypothesis:*

(A1) If x is a point, and S is a symplecton, then $x^\perp \cap S$ is never a single point.

Then the point-collinearity graph $G := (\mathcal{P}, \sim)$ has diameter at most three.

Proof By definition, parapolar spaces are connected. So it suffices to show that G possesses no geodesic of length four. By way of contradiction assume $g = (x_0, x_1, x_2, x_3, x_4)$ is a geodesic of length four and angle type $(\alpha_1, \alpha_2, \alpha_3)$. Thus each α_i is at least two.

Case 1: *One of the* $\alpha_i = 2$. Suppose first that $\alpha_1 = 2$. Then there is a symplecton R_0 on $\{x_0, x_1, x_2\}$ and ${x_3}^\perp \cap R_0$ contains a line N. Then, as R_0 is a polar space, there is a point $z \in {x_0}^\perp \cap N$, and we have a path (x_0, z, x_3, x_4) of length 3, a contradiction.

Thus $\alpha_1 \neq 2$ and by symmetry, $\alpha_3 \neq 2$.

Next suppose $\alpha_2 = 2$. Then there is a symplecton R_1 containing $\{x_1, x_2, x_3\}$, line $M = {x_4}^\perp \cap R_1$, and point $y \in {x_1}^\perp \cap M$. Then (x_0, x_1, y, x_4) is a path of length three connecting x_0 and x_4, a contradiction.

Case 2: *g has angle type* $(3, 3, 3)$. First choose a symplecton S_1 containing the edge $\{x_1, x_2\}$. Then ${x_3}^\perp \cap S_1$ contains a line M. Now, since S_1 has polar rank at least three, $M^\perp \cap S_1$ is not a clique, while ${x_3}^\perp \cap S_1$ is. Thus there exists a point $u \in S_1 \cap M^\perp - {x_3}^\perp$ and so the polar pair (u, x_3) lies in a symplecton S_2 containing $\{u, x_2, x_3\}$ meeting S_1 in at least the plane $\pi = \langle u, M\rangle$. Now ${x_0}^\perp \cap S_1$ contains a line N_1 on x_1, and in the polar space S_1, ${N_1}^\perp \cap \pi$ must contain a point z. Thus (x_0, z) is a polar pair. Now also (A1) forces ${x_4}^\perp \cap S_2$ to contain a line N_4 which meets $z^\perp$ in at least a point q. Now $w = (x_0, x_1, z, q, x_4)$ is a path connecting x_0 and x_4, of angle type $(\beta_1, \beta_2, \beta_3)$ with $\beta_1 = 2$. But that returns us to Case 1 with w replacing g. We have seen that that case leads to a contradiction.

Thus no such geodesic of length four exists and the proof is complete. □

We are now operating under this hypothesis.

(B1) *Every point-residual geometry* $Res(p) = (\mathcal{L}_p, \Pi_p)$ *is one of the conclusion geometries of Theorem 17.1.2.*

(B2) *There exists a point p for which* $Res(p)$ *is a geometry* $L_p \times P_p$ *where* L_p *is a line and* P_p *is a polar space of rank* $k_p \geq 2$ *that is not a grid.*

In condition (B2) the projective line L_p is thick since we are dealing with the point-residue of a parapolar space with thick lines.

Now suppose $L = pq$ is a line on p. Then L lies in a unique symplecton S of Γ which has rank $k_p + 1$, which is not oriflame of rank three. It also lies in a unique plane lying in no such symplecton. Since q is on such a line, and yields a residue in the conclusion of Theorem 17.1.2, it has a residue of shape $L_q \times S_q$, where $S_q \simeq S_p$ (being point-residue geometries of distinct points of the same symplecton S) and has symplectic rank k_p, and L_q has the same cardinality as L_p. Thus we see that $\mathrm{Res}(q) \simeq \mathrm{Res}(p)$.

Lemma 17.3.4 Γ *is a parapolar space with all point-residuals isomorphic to* $L \times P$ *where the line* L *has a constant cardinality, and the polar space* P *has a uniform polar rank* k *at least two. If* k *is greater than two, then* P *even has a constant isomorphism type.*

Remark The last sentence of the lemma follows from an unpublished theorem of Tits. We don't actually use this fact.

Lemma 17.3.5 *The following statements hold:*

1. *For any symplecton S and point x not in S, $x^\perp \cap S$ is empty, or is a line.*
2. *The symplecta are partitioned naturally into two sets as $\mathcal{S} = \mathcal{D} + \mathcal{S}^+$, where $\mathcal{D}$ is the collection of all oriflame rank three polar spaces, and $\mathcal{S}^+$ are the remaining symplecta (all of polar rank k). We have the following:*
   (a) *Any two distinct members of $\mathcal{S}^+$ intersect at the empty set or at a single point.*
   (b) *If $(D, S) \in \mathcal{D} \times \mathcal{S}^+$, then $D \cap S$ is the empty set or a plane.*
3. *The projective planes Π of Γ are also partitioned into two sets:*
   (a) *The $\mathcal{S}^+$-planes, which are the planes which lie in a (necessarily unique) member of $\mathcal{S}^+$.*
   (b) *The $\mathcal{D}$-planes, which are those planes which lie in no member of $\mathcal{S}^+$ at all. These are maximal singular subspaces of Γ.*
   (c) *Every line lies in a unique member of $\mathcal{S}^+$ and in a unique $\mathcal{D}$-plane.*
4. *If $D \in \mathcal{D}$, then the two oriflame classes of planes of D are the $\mathcal{D}$-planes and the $\mathcal{S}^+$-planes which are contained in D. In particular, any two distinct $\mathcal{S}^+$-planes of D intersect at a single point.*

Proof The symplecta in $\mathcal{D}$ can never be isomorphic to those in $\mathcal{S}^+$, so the two classes of symplecta can never fuse globally. All of the statements follow from the uniform local structure of any point-residual. For example, the $\mathcal{D}$-planes and $\mathcal{S}^+$-planes correspond to the "horizontal" and "vertical" lines respectively in the product geometry $L \times P$ representing a point-residual. □

Corollary 17.3.6 *If two symplecta from $\mathcal{S}^+$ both intersect non-trivially a common symplecton $D \in \mathcal{D}$, then the two symplecta either coincide or intersect at a point.*

Proof Suppose S_1 and S_2 are distinct members of $\mathcal{S}^+$, which intersect non-trivially a symplecton $D \in \mathcal{D}$. Then by Lemma 17.3.5, part 2(b), the intersections $S_i \cap D$ are planes which belong to the same oriflame class of D by part 4 of Lemma 17.3.5. Since the planes meet at a point, so do S_1 and S_2. □

Lemma 17.3.7 *Suppose S_1, S_2, and S_3 are pairwise distinct members of $\mathcal{S}^+$ on a common point p. Suppose R is a member of $\mathcal{S}^+$ which does not contain p. If R intersects S_1 and S_2 non-trivially, then it intersects S_3 non-trivially.*

Proof Let $\{x_i\} := S_i \cap R$, $i = 1, 2$. Let's get rid of an easy case first. Suppose p were collinear with one of the x_i, say x_1. Then by (A1) $p^\perp \cap R$ is a line L on x_1. Since L is not in S_1 (for $S_1 \cap R = \{x_1\}$), the plane $\langle p, L\rangle$ has to be the unique $\mathcal{D}$-plane on px_1. Then $\langle p, L\rangle$ intersects every symplecton of $\mathcal{S}^+$ on p at a line. Therefore L intersects every symplecton of $\mathcal{S}^+$ on p at a point. Thus there is a point in $L \cap S_3 \subseteq R \cap S_3$. Thus $R \cap S_3$ is non-empty.

So we may assume that $d(p, x_1) = d(p, x_2) = 2$ in the point-collinearity graph. If x_1 were collinear with x_2 then ${x_1}^\perp \cap S_2$ would be a line on x_2 carrying a point u

of $p^\perp \cap S_2$. Then $u \in p^\perp \cap {x_2}^\perp \subseteq S_2$, while $u \in p^\perp \cap {x_1}^\perp \subseteq S_1$. This is impossible as $S_1 \cap S_2 = \{p\}$.

Thus x_1 is not collinear with x_2. Select a point $t \in {x_1}^\perp \cap {x_2}^\perp \subseteq R$. Then $t^\perp \cap S_i$ carries a point s_i of $p^\perp$, and $s_1 \neq s_2$. Then there is a symplecton D on p and t, and $D \in \mathcal{D}$ since it meets S_1 and S_2 at lines, at least. Once again, $D \cap R$ and $D \cap S_3$ are non-empty, so Corollary 17.3.6 can be invoked to yield $R \cap S_3 \neq \emptyset$ in this case as well. □

Lemma 17.3.8 *If p is a point and $R \in \mathcal{S}^+$ is a symplecton not on p, then there is a symplecton $S \in \mathcal{S}^+$ containing p and meeting R non-trivially.*

Proof First suppose p is collinear with a point x of R. By Lemma 17.3.5, part 3(c), there is an element $S \in \mathcal{S}^+$ on the line px. Then $S \cap R \neq \emptyset$ and we are done.

Next suppose $d(p, r) = 2$ for some point $r \in R$. Choose $z \in p^\perp \cap r^\perp$. Then $z^\perp \cap R$ is a line L of R. Choose a point $y \in R \cap L^\perp - L$ (this is possible since R has polar rank at least three). Then there is a symplecton $D := \langle\langle z, y \rangle\rangle$ in $\mathcal{D}$. Now by Lemma 17.3.5, part 3(c), the line pz lies in a symplecton S of $\mathcal{S}^+$. But now $S \cap D$ and $R \cap D$ are both non-trivial, so $S \cap R \neq \emptyset$ by Corollary 17.3.6.

So we must assume $d(p, r) \geq 3$ for all points r in R. But by Theorem 17.3.3, we must assume

$$R \subseteq \Delta_3(p). \tag{17.3}$$

Let (p, u, v, r) be a geodesic from p to a point r of R. Then there is a symplecton D of $\mathcal{D}$ on line vr and meeting R at a plane (Lemma 17.3.5, parts 3(c) and 2(b)). Since Eq. (17.3) implies $u^\perp \cap R = \emptyset$, one has $u \notin D$. Then there is an element S of $\mathcal{S}^+$ on the line uv (Lemma 17.3.5, part 3(c) again). Now as S and R are elements of $\mathcal{S}^+$ meeting D non-trivially, there exists a point $s \in S \cap R$, by Corollary 17.3.6. But $p^\perp \cap S$ is a line on u carrying a point of $s^\perp$. Thus $d(p, s) = 2$ against $R \subseteq \Delta_3(p)$. The proof is complete. □

Lemma 17.3.9 *There is no symplecton in $\mathcal{S}^+$ which intersects non-trivially all other symplecta from $\mathcal{S}^+$.*

Proof Suppose $R \in \mathcal{S}^+$ has the "radical" property – that $R \cap S \neq \emptyset$ for all $S \in \mathcal{S}^+$.

Fix a point p in R, and a point $r \in R - p^\perp$. Let T be any symplecton of $\mathcal{S}^+$ on r with $T \neq R$. Then $T \cap R = \{r\}$. Choose $z \in T - r^\perp$. We claim that z is not collinear with any point of R. For if there were such a point, then $z^\perp \cap R \cap r^\perp$ would contain a point in $R \cap T - \{r\}$, an absurdity. Thus

$$z^\perp \cap R = \emptyset.$$

By assumption all symplecta of $\mathcal{S}^+$ which lie on z must meet R. Let S_1 and S_2 be two of these – that is $S_1 \neq S_2$ not in name only. Set $\{x_i\} := S_i \cap R$. Since $z^\perp \cap R = \emptyset$, z is distinct from both x_1 and x_2. If x_1 were collinear with x_2, then, by Lemma 17.3.5, part 1, $x_1 \cap S_2$ contains a line carrying a point w of $z^\perp$. Then by

convexity of S_1, the line zw is in S_1 as well as S_2, against Lemma 17.3.5, part 2(a). Thus we may assume $d(x_1, x_2) = 2$.

Now select a point $t \in x_1{}^{\perp} \cap x_2{}^{\perp} \subseteq R$. Set $t^{\perp} \cap S_i := N_i$, and set $\{y_i\} := z^{\perp} \cap N_i$, $i = 1, 2$. Then $t^{\perp} \cap z^{\perp}$ contains two distinct points y_1 and y_2, so the convex closure $D := \langle\langle t, z\rangle\rangle$ is a symplecton whose intersection with each S_i contains a line zy_i. Thus $D \in \mathcal{D}$. Then $D \cap R$ is a plane π (Lemma 17.3.5, part 2(b)), and so $z^{\perp} \cap \pi$ is a line in $z^{\perp} \cap R$. But by construction the latter set is empty. Thus no such R exists and the proof is complete. □

Theorem 17.3.10 *Under the hypotheses of Lemma 17.3.4, the rank two geometry $P^* := (\mathcal{S}^+, \mathcal{P})$ is a non-degenerate polar space.*

Proof By Lemmas 17.3.7 and 17.3.8 P^* is a polar space. It is non-degenerate by Lemma 17.3.9. □

Corollary 17.3.11 *Under the hypotheses (B-1) and (B-2) which headed this subsection, Γ is the polar Grassmannian of lines of a non-degenerate polar space P^* of rank at least four. Otherwise the nature of P^* is arbitrary.*

Theorem 17.1.1 has now been proved in the two cases in which all singular subspaces have finite projective rank and otherwise.

Remark By now the reader will have noticed how the phantom hypotheses of finite singular rank weave in and out of the two theorems. How they weave in is as dramatic as how they weave out. The original authors (Kasikova and Shult) do not have a good explanation for this. To a large extent the finite rank assumption is in when the proof wishes to invoke one of the following: (1) the Cohen–Cooperstein theorem in some form (for example, Theorem 17.1.2 though free of the assumption of constant symplectic rank, uses the Cohen–Cooperstein theorem) or (2) the theory of locally truncated geometries which requires Tits' "Local Approach Theorem" in order to say anything useful. Indeed, it would not be unfair to say that this theory is the *major application* of that beautiful theorem. But its application requires the covering chamber system C of $C(\mathcal{F})$ to be residually connected so that the functor $\mathbf{\Gamma}$ can be applied. That in turn requires a finite Coxeter matrix, and that means finite rank for all singular subspaces involved.[8]

17.4 Exercises for Chapter 17

17.1 Show that if axiom (H3) is omitted, that the half-spin geometry $D_{5,5}$ satisfies all of the other hypotheses of Theorem 17.1.1.

[8] Theorem 17.1.2 uses the Cohen–Cooperstein theory only under a very special low-diameter circumstance in which $x^{\perp} \cap S$ is always non-empty. If it were possible to prove finite singular rank in such a special case, one could dispense with the finite singular rank assumptions altogether in both Theorems 17.1.1 and 17.1.2. It is worth mentioning.

17.2 Show that if axiom (2). of the hypotheses of Theorem 17.1.2 is altered so that the word "proper" is omitted, the half-spin geometry $D_{5,5}$ satisfies all of the other hypotheses but is not one of the conclusions of that theorem.

17.3 The Lie incidence geometries $(B/C)_{n-1,n}$ are parapolar spaces of symplectic rank two. For $n > 2$ decide whether such geometries satisfy the axioms (H1), (H2), and (H3).

17.4 Make an attempt at the following open question: "Can one characterize parapolar spaces which are *not* of symplectic rank at least three which satisfy the axioms (H1)–(H3)?" This question can be divided into two. Suppose Γ is a parapolar space satisfying axioms (H1)–(H3).

(1) (Symplectic rank two.) Suppose every symplecton is a generalized quadrangle. Can one show that Γ is one of the two long root geometries in class 4, listed at the beginning of this chapter or one of the geometries of the preceding Exercise 17.3?
(2) (The mixed case.) Suppose at least one symplecton is a quadrangle, and at least one has has higher polar rank. Can one show that this is impossible for our prescribed Γ?

17.5 Appendix to Chapter 17: Long Root Geometries Updated

Excercise 17.4 raised the possibility of achieving a characterization of "long root geometries" satisfying axioms (H1)–(H3) when there are symplecta of rank two present.

This problem was approached through a completely different paradigm by Arjeh Cohen and Gabor Ivanyos [37] using the notion of a *root filtration space*

If some line L of the underlying partial linear space of a root filtration space lies in two distinct maximal singular subspaces, the entire space is a parapolar space whose poset of singular subspaces have the ascending chain condition and which satisfies all the hypotheses of Theorem 17.1.1 except for the assumption that this parapolar space has symplectic rank at least three. Remarkably they are able to classify the cases outside the scope of Theorem 17.1.1. The geometry $A_{\{1,n\}}$ (whose "points" are the point-hyperplane flags of a finite-dimensional projective space) is the only candidate when one has a parapolar space, and the generalized hexagon $G_{2,2}$ emerges when the root filtration space is not a parapolar space.

In another development [116], it was discovered that one could dispense with the hypotheses (H2) and (H3) of Theorem 17.1.1, provided one enlarged the conclusion set to include strong parapolar spaces of point-diameter 2 which still enjoy the remaining hypotheses (H1) and (F). The geometries $D_{5,5}$ and $E_{6,1}$ would be included in the latter collection, but it is not known at the present time whether there are more possibilities.

Chapter 18
The Peculiar Pentagon Property

Abstract This is a study of parapolar spaces whose point-residues satisfy the "pentagon property" of A. Cohen. Upon adding two hypotheses on points and symplecta, the only geometries of finite symplectic rank that emerge are metasymplectic spaces and polar Grassmannians of d-dimensional singular subspaces, $d > 0$, for some polar space. The proof requires Hanssens' principle and the machinery of Chap. 11. Some open questions arise when either of the two hypotheses are relaxed.

18.1 Introduction

From the *synthetic* point of view, one could insist that no account of the literature of characterizations of Lie incidence geometries would be complete without touching upon the remarkable consequences of the "Pentagon Property." That will be done in this chapter.

From the *descriptive* point of view, one could say that no book aimed at describing point-line characterizations of Lie incidence geometries would be complete without mentioning the one principal class barely mentioned so far in this book. I refer to the *polar Grassmannians*. These objects will be included in the class of geometries characterized in this chapter.

The reason only one chapter can service both objectives is that the only known characterizations of these geometries which do not begin with an outright declaration of the isomorphism type of every point-residual as a product-geometry of shape "projective space $\times$ polar space"[1] all involve the *Pentagon Axiom*.

[1] The polar Grassmannians are seen to be locally-determined in this way in certain papers of Hanssens, and in an important unpublished preprint of Lehmans (a student of F. Buekenhout). Hanssens wrote many papers in which local properties implied the isomorphism type of a point-residual – many of which were touched on in this book [68–71]. But somehow this doesn't happen with the polar Grassmannians as far as I am aware. The work of Lehmans characterized the polar Grassmannians $C_{n,n-1}$ with the hypothesis (basically) that a point-residual was a product space consisting of a generalized quadrangle and a finite-rank projective space. As one will see, that is the one polar Grassmannian not covered by the theorems of this chapter, and for that reason his contribution is important.

E.E. Shult, *Points and Lines*, Universitext, DOI 10.1007/978-3-642-15627-4_18,

© Springer-Verlag Berlin Heidelberg 2011

18.1.1 The Pentagon Axiom; Its Beginnings

What is this "Pentagon Axiom"? First let us state it. The context is that Γ is a parapolar space. A *pentagon* is described as a 5-cycle $C = (x_0, x_1, x_2, x_3, x_4, x_0)$ in which x_i is not collinear with x_{i+2} (indices taken mod 5). It means that if you drew such a subgraph of the collinearity graph, you would not encounter any "diagonal adjancies" between x_i and x_{i+2}, for all i.

Then the Pentagon Property asserts:

(P) *For some index i, x_i is collinear with at least one point y of the line $x_{i+2}x_{i+3}$ (indices taken modulo 5).*

Immediatelely one sees the following.

Lemma 18.1.1 *A parapolar space has property (P) if and only if every pentagon of the point-collinearity graph lies within some symplecton. Of course a symplecton with such a pentagon cannot be a grid.*

Arjeh Cohen's theorem was first stated as follows.

Theorem 18.1.2 (A. Cohen [33].) *Suppose Γ is a parapolar space with thick lines and constant symplectic rank $k \geq 3$. Suppose further that Γ possesses these two properties:*

(1) Γ possesses property (P).
(2) Given a symplecton S, and a point x not in the symplecton S, then $x^{\perp} \cap S$ is either the empty set or is a line.

Then Γ is a metasymplectic space.

18.1.2 Where Are We Going with this Axiom?

Our goal is to weaken both hypotheses of Cohen's pioneering theorem at the expense of enlarging the class of conclusion geometries so that the polar Grassmannians can finally make their debut on stage. First the Pentagon Property will be weakened by supposing it to hold in the point-residuals rather than holding in the geometry globally. Second we wish to depart from assumption (2) in Cohen's theorem. (It can be shown that keeping this assumption forces one to a subset of the hexagonic (or "long root") geometries of Chap. 16, depriving us of the long awaited polar Grassmannians.)

The development followed in this chapter is based on two sources. The first is the paper in [110] that contains a few typographical errors that I hope are now corrected. The second source is an unpublished preprint "*A characterization of Polar Grassmann Spaces*" by Cecil Ellard and the author.[2]

[2] I regret that I thoughlessly never submitted this paper for publication and that it joined the pile of unpublished papers shared with a limited number of colleagues (such as Professor A. Cohen who

18.1.3 Introduction

Our goal is to prove the following theorem.

Theorem 18.1.3 (Main Theorem, Tits [138], Cohen [33], Shult [110], Ellard and Shult [61].) *Let $\Gamma = (\mathcal{P}, \mathcal{L})$ be a parapolar space of symplectic rank at least three. (It is not assumed in advance that Γ is locally connected.) Assume the following hypotheses:*

(i) Every singular subspace of Γ has finite projective rank.
(ii) The Pentagon Property holds locally (i.e., for each point p the point-residual $\Gamma_p = (\mathcal{L}_p, \Pi_p)$ satisfies the Pentagon Property).
(iii) If S is a symplecton, and $x \in \mathcal{P} - S$ is such that $x^\perp \cap S = \{p\}$, a single point, then there exists a point $y \in p^\perp - S$, such that $y^\perp \cap S$ contains a plane.
(iv) There exists in Γ at least one point-symplecton flag (p, S) such that for every line L on p which is not in S, $L^\perp \cap S$ is just the point shadow of a line L' (possibly depending on the choice of L).

Then Γ is one of the following:

(1) A non-degenerate polar space of finite polar rank at least three.
(2) One of the following three types of metasymplectic spaces classified by Tits:

(a) The Lie incidence geometry of a (non-weak) building of type $F_{4,1}$.
(b) The polar Grassmannian of lines of a non-degenerate non-oriflame polar space of polar rank four – a Lie incidence geometry of type $C_{4,2}$ or $B_{4,2}$.
(c) The polar Grassmannian of lines of a non-degenerate oriflame polar space of polar rank four – type $D_{4,2}$.

(3) A polar Grassmannian of singular $PG(k)$s, $k > 1$, in a non-degenerate polar space of finite polar rank at least $k + 2 > 4$.

Corollary 18.1.4 ([61]) *Assume Γ is a parapolar space satisfying conditions (i), (ii), and (iii) of the preceding Theorem 18.1.3. Suppose condition (iv) is replaced by*

(iv') Γ contains at least one symplecton of polar rank at least four.

Then conclusion (1) or (3) of the preceding theorem holds. (In case (1) the polar space has rank at least four.)

Remark One will note that the parapolar spaces in the theorem and the corollary are not assumed to be locally connected. After some work it will emerge that local connectedness is restored by the assumptions (iv) or (iv') which refer to only one symplecton.

---

included this paper in his article in the *Handbook*). Professor Ellard played an equal part in this paper, and it is my hope this chapter will give his work its overdue recognition.

As usual, we approach this theorem and its corollaries in a series of steps comprising a sort of logical banquet: the hypotheses are served up sequentially in courses, and one tries to prove as much as one can at the end of each course. In this way, the lemmas are presented in the most general setting available, awaiting some future use in a related theorem. This has to be done twice: on the level of point-residuals, and then at the global level. At the global level, the last hypothesis to come in will be the assumption of finite singular rank.[3]

18.2 Strong Parapolar Spaces with the Pentagon Property

The Pentagon Property is really a very strong hypothesis. It is stronger than the "near-gated" property (NG) of parapolar spaces. Recall from Chap. 13, p. 478, that for general parapolar spaces property (NG) was this:

If S is a symplecton and $x \in \mathcal{P} - S$ with $x^\perp \cap S = \{p\}$, a single point, then S is gated with respect to x – that is, $d(x, y) = 3$ for all $y \in S - p^\perp$.

The reader may easily prove the following.

Theorem 18.2.1 *Suppose Γ is a parapolar space. Then the Pentagon Property implies Property (NG).*

18.2.1 Examples

Of course, polar spaces of rank at least two are themselves strong parapolar spaces with the Pentagon Property. But it is easy to manufacture many other examples.

Lemma 18.2.2 *Suppose X and Y are connected point-lines geometries with disjoint but non-empty point sets. Suppose each of X or Y is independantly either a singular space or a strong parapolar space with the Pentagon Property. Then the product geometry $X \times Y$ is also a strong parapolar space with the Pentagon Property.*

Similarly, suppose X and Y are both parapolar spaces with the Pentagon Property. If we identify a point of X with a point of Y, so that $X \cap Y = \{p\}$, then the geometry $X \cup Y$ is also a strong parapolar space with the Pentagon Property.

Also see Exercise 18.3.

18.2.2 The Role of Grids

For the rest of this section we assume the following.

$\Gamma = (\mathcal{P}, \mathcal{L})$ is a strong parapolar space with the Pentagon Property. We let $\mathcal{S}$ denote the set of all symplecta of Γ.

[3] Not to be confused with the stronger notion of bounded singular rank.

Recall that a *grid* is a generalized quadrangle in which every point lies on exactly two lines. Grids are always product geometries $L_1 \times L_2$ for two lines L_1 and L_2.

Lemma 18.2.3 *Suppose $S \in \mathcal{S}$ and $z \in \mathcal{P} - S$ such that $z^\perp \cap S \neq \emptyset$. Then one of the following statements holds:*

1. *$z^\perp \cap S$ is a single point p, and for every $x \in S - p^\perp$, $d(z, x) = 3$.*
2. *S is a grid and $z^\perp \cap S$ is a line. Moreover, for every $x \in S - z^\perp$, $d(x, z) = 2$ and the unique symplecton containing x and z is also a grid.*

Proof Let us suppose S is a symplecton and z is a point not in S which is collinear with at least one point of S. Choose a point $y \in z^\perp \cap S$ and a point $x \in S - y^\perp$. Clearly as $z^\perp \cap S$ is a clique we cannot have x collinear with z.

Suppose for one choice of $x \in S - y^\perp$ we have $d(x, z) = 2$. Since x and z now lie on a symplecton R (the strong parapolar hypothesis), we see that $z^\perp \cap x^\perp$ is not a clique. Since R is distinct from S, $R \cap S$ is a clique. So from the last two statements, we can find a point $u \in z^\perp \cap x^\perp - S$.

Now choose an arbitrary point w in $y^\perp \cap x^\perp \subseteq S$. We have seen that (z, x), (y, x) are non-collinear pairs, and as $u \notin S$, we cannot have (u, y) a collinear pair, either. Choose any point $w \in y^\perp \cap x^\perp \subseteq S$. If there were no diagonal adjacencies at all in the 5-circuit $c = (z, y, w, x, u, z)$, then by the Pentagon Property, the entire 5-circuit would lie in a symplecton which is the convex closure of any two of its 2-cocliques, forcing $R = S$, a contradiction. This means that whatever w we choose in $y^\perp \cap x^\perp$, we must have either $w \in S \cap z^\perp$ or $w \in S \cap u^\perp$. That means $y^\perp \cap x^\perp$ is the union of the two cliques $z^\perp \cap y^\perp \cap x^\perp$ and $u^\perp \cap y^\perp \cap x^\perp$, forcing S to be a grid and $z^\perp \cap S$ to be one of the two lines of S on y.

Finally, $y \notin R$, as this would force $R \cap S$ to contain the non-clique $\{x, y\}$. Yet y is collinear with z while being distance two from $x \in R - z^\perp$. Thus (y, R, z, x) have exactly the same hypotheses as (z, S, y, x) in the previous paragraph. So it similarly follows that R is also a grid. Thus all parts of the second alternative in the conclusion are established.

So we may assume that the hypothesis $d(z, x) = 2$ of the last two paragraphs is not realized for any $x \in S - y^\perp$. Then $z^\perp \cap S$ can only contain the single point y. Thus the first alternative of the conclusion must hold. The proof is complete. □

18.2.3 Why We Assume that Γ *Is Not a Polar Space*

As the examples of the previous section reveal, strong parapolar spaces with the Pentagon Property probably cannot be classified. All sorts of building geometries with bizarre multi-armed diagrams possess truncations to parapolar point-line geometries whose point-residuals have the Pentagon Property.

The important Lemma 18.2.3 of the previous section involves a symplecton and a point not in it. What if there is no such external point? Then (as the parapolar space is by definition connected) we see that $\Gamma = (\mathcal{P}, \mathcal{L})$ is itself a symplecton. In that case Γ could be some unknown quadrangle and we still shall not have pinned things down. However, the strong parapolar Γ that concerns us in this section, will

eventually play the role of a point-residual in a parapolar space Γ' of symplectic rank at least three. If this point-residual $\Gamma = \Gamma'_p$ is a polar space, we know that some symplecton S of Γ' has p for a deep point. If Γ' is locally connected, one may apply Theorem 13.5.5 of Chap. 13, to conclude that Γ' is itself a polar space of polar rank at least three. On the other hand, there are boundless examples of parapolar spaces Γ' of symplectic rank at least three which are not locally connected, but have the Pentagon Property visible in each point-residual. One would not like to assume the burden of classifying these anyway.

So in seeking out a reasonable result, it seems natural to assume that no point-residual of our Γ' will be a polar space. For that reason, in this section, we assume that our strong parapolar space Γ is not a polar space — hoping to handle that alternative at the global stage. (That is exactly what we do in Lemma 18.4.4, p. 649.)

So from now on in this section, we will assume the following:

(PL) Γ *is a strong parapolar space with the Pentagon Property that is not itself a polar space.*

18.2.4 What Extra Axioms Are Needed?

Of course, the geometries described as "bizarre" at the beginning of the previous section probably cannot be classified. Rather than merely guessing at a suitable extra hypothesis, let us ask where our current knowledge leads us. Now that we are allowed to assume that Γ is not a polar space we see that every single symplecton of Γ possesses exterior points, and so, by Lemma 18.2.3 falls into just one of two categories:

$$\mathcal{D} := \{D \in \mathcal{S} |\ max\{|z^\perp \cap D| | z \in \mathcal{P} - D\} > 1\},$$
$$\mathcal{Q} := \{Q \in \mathcal{S} |\ |z^\perp \cap Q| \leq 1 \text{ for all } z \in \mathcal{P} - Q\}.$$

Thus $\mathcal{S}$ partitions as

$$\mathcal{S} = \mathcal{D} + \mathcal{Q}.$$

Thus $\mathcal{D}$ is the class of symplecta for which there exists an external point, collinear with at least two of its internal points. Note that by Lemma 18.2.3 every member of $\mathcal{D}$ is already a grid.

Any axiom we hope to add should respect this partition. Our plan here is to refine the class $\mathcal{Q}$. So we introduce the following axiom:

(PL*) *If $Q \in \mathcal{Q}$, then for any external point $z \in \mathcal{P} - Q$, $z^\perp \cap Q$ is non-empty. That means $z^\perp \cap Q$ is always a single point for all $z \in \mathcal{P} - Q$.*

There are two reasons for adopting this axiom. Many of our favorite classical geometries satisfy (PL*), while the bizarre ones do not. Also, after 10 years we have been unable to characterize strong parapolar spaces with the Pentagon Property satisfying a more general hypothesis than (PL*).

18.3 Classifying the Parapolar Spaces Satisfying (PL) and (PL*)

18.3.1 An Elementary Consequence of (PL)*

So just to touch base, we are assuming for the rest of this section that Γ satisfies both (PL) and (PL*).

Lemma 18.3.1 *Suppose $Q \in \mathcal{Q}$ and $S \in \mathcal{S}$ with S distinct from Q. Then the following hold:*

1. *$Q \cap S$ can never be a single point.*
2. *Suppose $Q \cap S$ is non-empty.*
   - *(a) If Q has polar rank at least three, then S is a grid in $\mathcal{D}$.*
   - *(b) If $S \in \mathcal{Q}$ then both S and Q are generalized quadrangles.*
   - *(c) In any case S is a generalized quadrangle and $Q \cap S$ is a line.*

Proof 1. Suppose $S \cap Q = \{p\}$. Choose $z \in S - p^\perp$. Since $Q \neq S$, z is not in Q. But since $Q \in \mathcal{Q}$, $z^\perp \cap Q$ contains a point y, by (PL*). If y were in $p^\perp$, we would have

$$y \in (p^\perp \cap z^\perp) \cap Q \subseteq P \cap Q = \{p\},$$

forcing $y = p$. That is absurd, since y is collinear with z while p is not. Thus we see that $z^\perp \cap S$ can contain no point of $p^\perp$, and in particular cannot contain a line. Since it is a clique, we have $z^\perp \cap Q = \{y\}$. Now if we choose $s \in p^\perp \cap z^\perp \subseteq S$, we see that s is not in Q, and so $s^\perp \cap Q$ is a singular subspace of Q containing p. Since $p^\perp \cap y^\perp \subseteq S$ is not a clique, we can choose a point $q \in p^\perp \cap y^\perp - s^\perp$. Similarly, $S \cap Q = \{p\}$ implies $s^\perp \cap Q = \{p\}$, and $y^\perp \cap S = \{z\}$, so that (s, y) and (s, q) are non-collinear pairs. Then $P := (p, s, z, y, q, p)$ is a 5-circuit of the point-collinearity graph having no diagonal adjacencies. By the Pentagon Property, P lies in a single symplecton T which is the 2-convex closure of any non-collinear pair of points in it. Thus $S = T = Q$, contrary to the hypothesis that S and Q were distinct symplecta. Thus it is impossible for $S \cap Q$ to be a single point.

2. Now suppose $(S, Q) \in \mathcal{S} \times \mathcal{Q}$ is a pair of distinct symplecta having a non-empty intersection. Then by part 1, $Q \cap S$ contains a line L.

First suppose Q has polar rank at least three. Then, since $Q \cap S$ is a singular space while $L^\perp \cap Q$ is not, we can choose a point $q \in L^\perp \cap Q - S$. Then q is a point external to S for which $q^\perp \cap S$ contains a line. By Lemma 18.2.3, S is a grid in class $\mathcal{D}$. So 2(a) holds.

If both S and Q belong to $\mathcal{Q}$, S does not belong to $\mathcal{D}$ and so by 2(a) neither can have polar rank exceeding two. Thus they are both generalized quadrangles, and 2(b) is proved.

Finally, S is either in $\mathcal{D}$ or in Q, and in either case S is a generalized quadrangle. This forces $L = S \cap Q$. So 2(c) holds and the proof is complete. □

18.3.2 The Case that $\mathcal{D}$ is Empty and a Revisitation of Cohen's Theorem

Just for this section, in addition to (PL) and (PL*), we are assuming that $\mathcal{D}$ is empty. Thus $\mathcal{S} = \mathcal{Q}$.

We observe that Lemmas 18.2.3 and 18.3.1 now imply the following.

Lemma 18.3.2 *1. All symplecta are generalized quadrangles.*
2. *Any two distinct symplecta intersect at the empty set or a line.*
3. *For any point x not incident with symplecton Q, Q is gated*[4] *with respect to x with the gate being the unique point of $x^\perp \cap Q$.*
4. *Since Γ is not a polar space but is a connected strong parapolar space, for each point-symplecton flag $(p, S) \in \mathcal{P} \times \mathcal{S}$, p lies on a line not in S.*

We also note the following.

Lemma 18.3.3 *Every line is a maximal singular subspace.*

Proof If not, there exists a plane π. Now any line L in π lies in a symplecton R (required by the parapolar hypothesis), and since R is a quadrangle, one has $R \cap \pi = L$. But by Lemma 18.2.3 this forces $R \in \mathcal{D}$, contrary to the assumption of this section. □

Lemma 18.3.4 *Suppose Q_1, Q_2, Q_3 are three symplecta which pairwise intersect non-trivially. Then there is a common point to all three symplecta.*

Proof By Lemma 18.3.2, part 2, for any 2-subset $\{i, j\} \subseteq \{1, 2, 3\}$, we may set $Q_i \cap Q_j = L_{ij}$, a line. We are done if any two of the L_{ij} possess a non-empty intersection, so we may presume these intersections to be empty. Suppose x is a point in L_{12}. Then x is collinear with a unique point of Q_3, namely the unique point y of $x^\perp \cap L_{23}$. But $x^\perp \cap L_{13}$ contains a point y', since Q_1 is a polar space. But now $y' \in x^\perp \cap Q_3 = \{y\}$. Thus $y' = y$ lies in L_{13} as well as L_{23}. So y lies in all three of the Q_i. □

Lemma 18.3.5 *The rank three geometry $P := (\mathcal{Q}, \mathcal{L}, \mathcal{P})$ (where containment of subspaces is incidence) is a classical non-degenerate polar space of rank three or is one with thin lines. In other words, $\Gamma = (\mathcal{P}, \mathcal{L})$ is one of the following:*

1. *A classical dual polar space of rank three.*
2. *The product geometry $L \times Q$ where L is a line and Q is a generalized quadrangle.*
3. *The product geometry $L_1 \times L_2 \times L_3$ of three lines L_i of arbitrary independant cardinalities.*

Proof It is necessary only to verify that $(\mathcal{Q}, \mathcal{L})$ is a non-degenerate polar space of rank three. So we are to imagine that the symplecta of $\mathcal{Q}$ are the "Points" and the

[4] Note that since Q is convex of point-diameter two, there is no distinction between (x, Q) being gated and (x, Q) being strongly gated.

collection $\mathcal{Q}(L)$ of all symplecta which contain a fixed line L, form the "Point-shadow" of a "Line." Thus two distinct "Points", A and B are "collinear" if and only if A and B are two symplecta of Γ whose intersection is a line.

Now, since distinct symplecta intersect at a line or the empty set (Lemma 18.3.2, part 2), we see that the collection of all symplecta which are incident with a common point p form a "singular subspace" $\Pi(p)$ of $P = (\mathcal{Q}, \mathcal{L})$ which by Lemma 18.3.2, part 4 properly contains a "Line".

Now, by Lemma 18.3.4, every triangle of the "Point collinearity graph" of $P = (\mathcal{Q}, \mathcal{L})$ lies in a "Plane" $\Pi(p)$, from which we deduce that P is a gamma space.

Now we prove the "one-or-all axiom" of polar spaces. Since P is already a gamma space we need only prove that if $\mathcal{Q}(L)$ is the "Point-shadow" of a "Line", L, and R is a symplecton of Γ not containing L (i.e., a "Point", not on this "Line"), then R is "collinear" with at least one "Point" of $\mathcal{Q}(L)$. Of course we are done if $R \cap L \neq \emptyset$. So suppose $R \cap L = \emptyset$. Now each point ℓ of L is collinear with a unique point $\rho(\ell)$ of R. If $\rho(\ell) = \rho(\ell') = r$ for distinct points ℓ and ℓ' of L, then $r \in L^{\perp}$ against Lemma 18.3.3. So we see that the mapping $\rho : L \to R$ is injective. But the strong parapolar property forces the intersecting lines L and $\ell \cdot \rho(\ell)$ to lie in a sympleton S which meets R at a line N. Since L and N are opposite lines of the symplecton S, we see that $N = \{\rho(\ell) | \ell \in L\}$, i.e., we can write $N = \rho(L)$ with some insult to the notation. Now R is a symplecton of $\mathcal{Q}(L)$ which is "collinear" with the symplecton S.

Next we prove that the polar space is non-degenerate. Suppose S and R are symplecta of Γ sharing a line L. Then R contains a line N which is opposite L in R. Now, by part 4 of Lemma 18.3.2, R is not the only symplecton on the line N, so we can assume that there is symplecton T with $T \cap R = N$. Now if T were "collinear" with S, by Lemma 18.3.4 one should be able to find a point $x \in S \cap R \cap T$. But that is impossible since $L \cap N = \emptyset$. Thus the arbitrary symplecton S can never represent a "Point" in the radical of the polar space P.

Thus $P = (\mathcal{Q}, \mathcal{L})$ is a non-degenerate polar space of rank at least three.

Finally, the last step: we need to know that $P = (\mathcal{Q}, \mathcal{L})$ possesses polar rank three. It is sufficient to show that the "Planes" $\Pi(p)$ are maximal singular subspaces of P. But this is easy. Suppose R is a symplecton not in $\Pi(p)$, which nonetheless shares a line with each member of $\Pi(p)$. Choose any line L on the point p. Since R shares a line with any two members of $\mathcal{Q}(L)$, then by Lemma 18.3.4, R must contain a point ℓ of L. Now choose a second line N on p distinct from L. Then by the same argument R must contain a point n of N. By our choice of R, neither n nor ℓ can be the point p. Thus ℓ and n are distinct points of $p^{\perp} \cap R$. This contradicts $R \in \mathcal{Q}$.

Thus, $P = (\mathcal{Q}, \mathcal{L})$ is a non-degenerate polar space of rank exactly three. The proof is complete. □

Remark Of course, without the previous lemma, one can see from the gatedness of each symplecton that Γ is in fact a classical[5] near polygon of point-diameter three,

[5] The adjective "classical" merely refers to the fact that the rank three polar spaces with thick lines were classified by Tits. As the reader will recall from Chap. 7, not all such spaces are 'classical'

and read off the classification from Cameron's theorem. But Cameron's beautiful theorem is quite non-trivial – as I hope Chap. 8 shows. To employ it would invoke a lot of superb work for a general result in a very special situation – from the point of view of "logical distance," it would not exactly be a geodesic under this metric.

Now this lemma provides an immediate direct route to a generalization of Cohen's remarkable awakening call.

Theorem 18.3.6 (Cohen [33] and [110].) *Suppose* Γ *is a parapolar space of symplectic rank at least three with these three properties:*

(1) Γ *is not itself a polar space.*
(2) The Pentagon Property holds locally – that is, each point-residual $\Gamma_p = (\mathcal{L}_p, \Pi_p)$ *satisfies the Pentagon Property.*

1. *For each symplecton* S *and point* $x \in \mathcal{P} - S$, $x^\perp \cap S$ *is either empty or is a line.*

Then Γ *is a metasymplectic space.*

Proof From the hypotheses, every connected component of the point-residual $\Gamma_p = (\mathcal{L}_p, \Pi_p)$ is a strong parapolar space which is not a polar space, possessing the Pentagon Property. Moreover, the third hypothesis forces each symplecton of Γ_p to be gated with respect to every exterior point. In particular Γ_p is connected, and we conclude that Γ is locally connected.

Thus we may apply Lemma 18.3.5 to conclude that Γ_p is either a non-degenerate polar space of rank at least 2 or a classical dual polar space of polar rank three.

Recall from our study of parapolar spaces in Chap. 13 that since Γ has symplectic rank at least 3, each symplecton of Γ_p is the set of lines and planes on a point in some symplecton of Γ. Thus, if Γ_p should be a polar space for some point p, then p is a deep point of some symplecton S of Γ – i.e., $p^\perp \subset S$. But in that case, since Γ is locally connected, we may apply Theorem 13.5.5 on p. 471 to conclude that Γ is itself is a polar space, against hypothesis (PL).

Thus we see that every point-residual Γ_p is a dual polar space of a polar space (possibly with some thin lines) of rank three. Now the enriched geometry $\Gamma^* := (\mathcal{P}, \mathcal{L}, \Pi, \mathcal{S})$ of points, lines, planes, and symplecta of Γ forms a diagram geometry of type F_4 satisfying all of Tits' axioms for a metasymplectic space. □

Remark In the case that one point-residual $\Gamma_p \simeq L_1 \times L_2 \times L_3$, the product of three lines, then uniformly all point-residuals are isomorphic to each other, the cardinality of the L_i are now equal, and Γ is a Grassmannian of lines of a polar space of type D_4. Similarly, in the case that one point-residual Γ_p has the form $L \times Q$ where L is a line and Q is a generalized quadrangle that is not a grid, then all point-residuals

---

in the sense that this term is used in the expression "classical groups." Not all such polar spaces are defined by quadratic or reflexive sesquilinear forms. For certain fields, they arise from compact forms of Lie Algebras over certain fields in characteristic zero. For some of these polar spaces, the planes are Moufang but not Desarguesian. The famous Bruck–Kleinfeld classification produces a Cayley–Dixon algebra which together with its "norm" function was used by Tits to complete the classification. Of course these polar spaces also deserve to be called "classical."

have this form and Γ is a polar Grassmannian of lines of a non-oriflame polar space of rank four. But these "thin building cases" are included in Tits' classification of the metasymplectic spaces just as they were for the hexagonic classification of the previous chapter.

Remark I have included the proof of this version of Cohen's theorem despite the fact that it could be proved in just a few lines by the theory of hexagonic geometries of the preceding chapter. But the latter course would actually yield a proof of much greater "logical length" since the theorem used was earned by a painfully lengthy proof that was not permitted to exploit the Pentagon Property.

18.3.3 The Case that Some Symplecton has Polar Rank at Least Three

We continue with our assumption that Γ satisfies (PL) and (PL*).

Lemma 18.3.7 *Let Q be an arbitrary symplecton of $\mathcal{Q} \cup \mathcal{D}$.*

1. *If $x \in Q$ then $T(x) := \{x\} \cup (x^{\perp} - Q)$ lies in a maximal singular subspace M of Γ.*
2. *If, in addition, $Q \in \mathcal{Q}$, then $M \cap Q = \{x\}$.*

Proof Part 1. Suppose u and v are distinct non-collinear points of $T(x) - \{x\}$. Then the lines ux and vx are distinct, do not lie in a common singular subspace and so, by the strong parapolar space assumption, must lie together in a unique symplecton R. By Lemma 18.3.1, the non-empty intersection $R \cap S$ is a line L. But also from Lemma 18.3.1, since S has polar rank at least 3, R is a grid belonging to $\mathcal{D}$. But that is impossible since the point x lies on the three distinct lines, xu, xv, and L, all belonging to R.

So such a 2-coclique $\{u, v\}$ cannot exist in $T(x) - \{x\}$ and it follows that the points of $T(x)$ form a clique in the point-collinearity graph of Γ. Since Γ is a gamma space, $T(x)$ generates a singular subspace.

Part 2. Suppose M were a singular subspace of Γ containing $T(x)$. If M contained a line on x not contained in $T(x)$, then from the definition of $T(x)$, that line must be a line $N \subseteq M \cap Q$. But then for any $u \in T(x) - Q$, we should have $u^{\perp} \cap Q$ containing line L and that would force $Q \in \mathcal{D}$ against its being rank at least three. Thus we see $T(x) = M$ is itself a maximal singular subspace of Γ. □

For the rest of Sect. 18.3.3, we assume that Q is a symplecton of polar rank at least three.

Lemma 18.3.8 *Suppose $x \in Q$ and $z \in T(x) - \{x\}$. Then z lies in a unique symplecton $Q_z \in \mathcal{Q}$ and there is an isomorphism $Q_z \to Q$. All lines on z not in Q_z belong to $T(x)$ – so with a slight abuse of notation we have $T(z) = T(x)$. Similarly, all lines on z which are not in $T(z)$ are in Q_z. So, with the usual abuse of notation, we have*

$$\mathcal{L}_z = (\mathcal{L}_z \cap T(z)) + (\mathcal{L}_z \cap Q_z).$$

Proof Let $\mathcal{L}_x \cap Q$ denote the lines on point x which lie in Q and let $\mathcal{L}(z, Q)$ be the lines on z which do not lie in $x^\perp$. Then there is a bijection

$$f : \mathcal{L}_x \cap Q \to \mathcal{L}(z, Q)$$

given as follows. Let L be a line on x in Q and choose a point $p \in L - \{x\}$. Then $d(p, z) = 2$ and so, since Q has polar rank at least three, the unique symplecton $S(L)$ containing $\{z, p\}$ is an element of $\mathcal{D}$, by Lemma 18.3.1, and hence is a grid. Thus there is a *unique* line $f(L)$ on z lying in $S(L)$ that is distinct from the line zx.

Clearly f is injective, for if $f(xp) = f(xq) = N$, then the intersection of the grids $\langle\langle z, p\rangle\rangle$ and $\langle\langle z, q\rangle\rangle$ contains the two lines N and xz, and so are the same convex grid. Thus they intersect S at the same line, forcing $xp = xq \in \mathcal{L}_x \cap Q$.

Now we show that f is surjective. Suppose N is a line on z not in $x^\perp$. We claim that $N^\perp \cap T(x) = \{z\}$. Suppose otherwise. Then there is a line K on x in $T(x) \cap N^\perp$. Clearly K cannot contain x from the definition of N, and so we have two planes $\pi_1 := \langle N, K\rangle$ and $\pi_2 := \langle K, x\rangle$ intersecting in line K and together not lying in a singular space. It follows that the two planes lie in a unique symplecton S of polar rank at least three. But that is impossible since $x \in S \cap Q$ and so by Lemma 18.3.1, S must be a grid.

Now that $N^\perp \cap T(x) = \{z\}$, we see that N and xz lie in a unique symplecton D which by Lemma 18.3.1 meets Q at a line L on x. Then $D \in \mathcal{D}$ as before and $N = f(L)$.

But we have also proved the following:

(*) *If N is a line of Γ meeting Q at the empty set and not lying in $x^\perp$ for any point $x \in Q$, then the mapping $\phi : N \to S$ defined by $\{\phi(z)\} := z^\perp \cap Q$ is injective, and its image $\phi(L)$ is a line of Q (namely the intersection of the unique symplecton on $N \cup z\phi(z)$, for any $z \in N$, with Q).*

Suppose L_1 and L_2 are two distinct lines of $\mathcal{L}_x \cap Q$ which lie in a plane π. Let N be a line of π meeting L_i at a point x_i distinct from x, so we can write $L_i = xx_i$. Then $f(L_i) = zz_i$ where $\phi(z_i) = x_i$ for $i = 1, 2$. Now $(z, z_1, x_1, x_2, z_2, z)$ is a 5-cirucit with no diagonals on z, x_1, or x_2, and does not lie in either of the convex grids $\langle\langle z, x_i\rangle\rangle$. The Pentagon Property then forces z_1 to be collinear to z_2. The line $N' := z_1z_2$ cannot intersect Q and so by (*) we have $\phi(N) = N'$. Thus $L_1 \perp L_2$ implies $f(L_1) \perp f(L_2)$ where "$\perp$" records the coplanar relation on lines. Also, by (*), the converse is true: $f(L_1) \perp f(L_2)$ implies $L_1 \perp L_2$.

Now in $\mathcal{L}_x \cap Q$ there exists a line L_3 such that $L_1 \perp L_2 \perp L_3$ while L_1 and L_3 are not coplanar. Then $f(L_1)$ is not coplanar with $f(L_3)$ and so their convex closure is a symplecton Q_z on z carrying the plane $\langle f(L_1), f(L_2)\rangle$. Thus Q_z has rank at least three and so belongs to $\mathcal{Q}$. Then $Q_z \cap Q$ is empty (otherwise Q_z would be a grid) and it follows that the restriction of ϕ to the symplecton Q_z is a bijection. But by (*) (applied to both ordered pairs (Q_z, Q) and (Q, Q_z)), this restriction $\phi : Q_z \to Q$ is an isomorphism of polar spaces.

In particular, we see that Q_z contains every line on z that is not in $T(z) = T(x)$ since it is true that $\mathcal{L}_x \cap Q$ is precisely all lines of x that are not in $T(x)$. The proof is complete. □

Theorem 18.3.9 *Assume* Γ *satisfies (PL) and (PL*), and contains a symplecton* Q *of polar rank three or more. Then*

$$\Gamma = M \times Q,$$

where M *is a singular subspace that at least contains a line.*

Proof By the previous Lemma 18.3.8 there is now a partition of the points of Γ:

$$\mathcal{P} = \oplus\{T(v) | v \in Q\}. \tag{18.1}$$

Suppose v and w are distinct points of Q. Then as noted $T(v) \neq T(w)$. Let x be any point of $T(v) - Q$. If v is collinear with w in Q, the $\{x, v, u\}$ lie in a unique symplecton $D \in \mathcal{D}$, and the unique line of D on w that is distinct from vw belongs to $T(w)$ and carries on it a point z collinear with x.

Now (whether or not v and w are collinear in Q) suppose $x \in T(v) - \{v\}$ is collinear with a point z' in $T(w) - \{w\}$. Then $\{z', x, v\}$ lies in a unique symplecton D' which must intersect Q at a line on v carrying a point of $z'^{\perp} \cap Q$. But there is only one point in that intersection and that is w. Thus v is collinear with w, and we see (following the notations of the previous paragraph) that $D = D'$ and $z' = z$. Thus we have:

(T) *Assume* v *and* w *are distinct points of* Q. *A point of* $T(v)$ *can be collinear with a point of* $T(w)$ *if and only if* v *is collinear with* w. *Moreover, if that happens, then every point* x *of* $T(v)$ *is collinear with a unique point* $\phi_{vw}(x)$ *of* $T(w)$ *defining an isomorphism*

$$\phi_{vw} : T(v) \to T(w)$$

as linear spaces.

Fix a point $x \in Q$. By the previous Lemma 18.3.8, we know that for each element of $z \in T(x) - \{x\}$, there is a unique symplecton $Q_z \in \mathcal{Q}$ which is the convex closure of all lines of $\mathcal{L}_z$ not in $x^{\perp}$ and that $Q_z \cap T(x) = \{z\}$, and $Q_z \cap Q = \emptyset$. (For notational convenience, we also write $Q_x := Q$.) It now follows that the points are partitioned by the elements of $\mathcal{Q}$.

$$\mathcal{P} = \oplus\{Q | Q \in \mathcal{Q}\}.$$

Let us summarize:

1. Each point x lies in a unique maximal singular subspace of the form $T(v)$, $v \in Q$, having a fixed isomorphism type (where either $x = v$, or $x \notin Q$, and $x^{\perp} \cap Q = \{v\}$).
2. Also x lies in a unique member Q_x of $\mathcal{Q}$.
3. Moreover, the two spaces $T(v)$ and Q_x partition the lines on x so that all lines of $\mathcal{L}$ either lie in some $T(v)$ or some member of $\mathcal{Q}$, but not both.

It follows from these statements that $\Gamma = M \times Q$, where M is the uniform isomorphism type of the $T(v)$s, $v \in Q$.

18.3.4 The Case that $\mathcal{D} \neq \emptyset$ and All Members of $\mathcal{Q}$ Have Polar Rank Two

In this section, we continue with assumptions (PL) and (PL*) which gave us the division of the symplecta into classes $\mathcal{Q}$ and $\mathcal{D}$ with the former non-empty. For this section we assume that $\mathcal{D}$ is non-empty, and that $\mathcal{Q}$ contains no symplecta of polar rank three or more – that is, the symplecta of $\mathcal{Q}$ are all generalized quadrangles.

Since Γ is now a strong parapolar space with symplectic rank exactly two, we no longer know that singular spaces are projective spaces (although that will happen if we know that Γ is itself a point-residue of another parapolar space of symplectic rank at least three). So the reader should be aware that in the results for this section, the word "*plane*" will simply mean *a singular subspace generated by a line and a point not on it* – not necessarily a projective plane. At least every line contained in a plane is a proper subspace of it.

We begin with a useful, if slightly technical, lemma.

Lemma 18.3.10 (The Tent lemma.) *Suppose D is a symplecton and p is a point of $\mathcal{P} - D$ for which $p^\perp \cap D = L$ is a line. (Then, of course, $D \in \mathcal{D}$ and so is a grid.)*

Let c and d be distinct points of a line N of D which is opposite line L and let a and b be the unique points of L which are collinear with c and d respectively. Then, as $E := \langle\langle p, d\rangle\rangle$ is an grid in $\mathcal{D}$, there is a unique point, $q \in p^\perp \cap d^\perp$ which is distinct from b (see Fig. 18.1 (i)). Then c and q are distinct collinear points – that is, Fig. 18.1 (ii) is obtained.

Proof If p were collinear with c, one would have $p \in \langle\langle c, b\rangle\rangle = D$, against our assumption. If a were collinear with q then q then would lie in $\langle\langle a, d\rangle\rangle = D$, and, since $q \neq c$ (for c is not collinear with p and while q is), we should have three distinct lines dc, dq, and db on point d in grid D, an absurdity. Finally, we see that (c, b) and (a, d) are non-collinear pairs from the way a and b were chosen.

Thus if c were not collinear with q the circular path (p, q, d, c, a, p) would be a pentagon. If that were true, its five points would lie in a common symplecton S, by

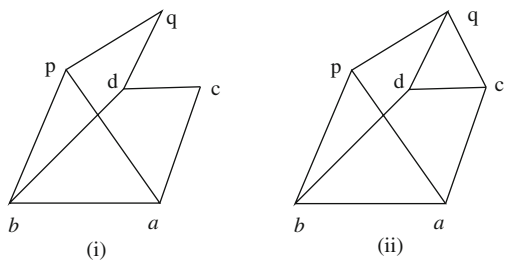

Fig. 18.1 The completion of a configuration in the Tent lemma

the Pentagon Property. But then $E = \langle\langle p, e\rangle\rangle = S$ so $b \in S$. This forces the plane $\pi = \langle a, p, b\rangle$ to lie in S, against the hypothesis that there are no symplecta of rank three or more.

Thus we must have c collinear with q. □

Lemma 18.3.11 *If $Q \in \mathcal{Q}$ and S is any symplecton, then $S \cap Q$ is either empty or is a line.*

Proof This just recalls the statement of Lemma 18.3.1, part 2(c).

Lemma 18.3.12 *The following statements hold:*

1. *Planes exist in* Γ.
2. *If L is a line then $L^\perp$ is a maximal singular subspace. As a result, if π is a plane, then $\pi^\perp$ is a maximal singular subspace.*
3. *If π is a plane, then for each quadrangle $Q \in \mathcal{Q}$, $\pi^\perp \cap Q$ is a single point.*

Proof 1. Since $\mathcal{D} \neq \emptyset$, there exists a point y outside a symp D, with $y^\perp \cap D$ being a line, and clearly y and this line span a plane – as we have defined the term above.

2. Now let L be a line. If (x, y) were a pair of non-collinear points in $L^\perp$ the strong parapolar hypothesis would give us a symplecton $\langle\langle x, y\rangle\rangle$ – the convex closure of $\{x, y\}$ – which contained planes, against the fact that Γ has symplectic rank exactly two. Thus $L^\perp$ is a singular subspace. It follows that if X is a singular subspace containing L, then $X^\perp = L^\perp$.

3. We may suppose π is a plane and $Q \in \mathcal{Q}$, with $\pi \cap Q = \emptyset$. Then there is a mapping

$$\phi_\pi : \pi \to Q,$$

sending each point $x \in \pi$ to the unique point of q with which it is collinear. We claim that ϕ_π maps π to a single point of Q. If this is not true, there exist two distinct points a and b, on some line L of π, which have distinct images $c := \phi_\pi(a)$ and $d := \phi_\pi(b)$, respectively. That means that the two intersecting lines $L = ab$ and $L' := ac$ do not lie in a clique, and so determine a symplecton D. Since L is a proper subspace of π, there is a point $p \in \pi - L$, so $p^\perp \cap D$ contains a line, from which we infer that $D \in \mathcal{D}$. But that means D is a grid. On the other hand, by Lemma 18.3.11 D must must intersect Q at a line, N and now it is easy to see that $N = \phi_\pi(L)$.

Now suppose p is a point of $\pi - ab$, and set $q := \phi_\pi(p)$. Suppose q were a point in line N. Then we can find $u \in \{a, b\}$ such that $\phi_\pi(u) \neq q$. Then $p^\perp \cap D$ contains $\{u, q\}$ which are non-collinear, forcing $p \in D$ by convexity. Then D contains the plane $\langle p, L\rangle_\Gamma$, against the fact that it is a quadrangle.

Thus we see that q is a point of Q not on the line N. Now we can repeat the arguments of the paragraph preceding the last, with grid D replaced by grids $D_a := \langle\langle ap, ac\rangle\rangle$ and $D_b := \langle\langle bp, bd\rangle\rangle$ to infer that $\phi_\pi(ap) = D_a \cap Q$ and $\phi_\pi(bp) = D_b \cap Q$ are the lines of Q containing $\{a, q\}$ and $\{b, q\}$ repectively. But in that case, Q contains the plane $\langle q, N\rangle_\Gamma$, against Q being a quadrangle.

Thus we are forced to conclude that $\phi_\pi(\pi)$ contains a unique point x. Then $x \in \pi^\perp \cap Q$ from the definition of ϕ_π, and this part is proved. □

Let $\mathcal{M}'$ denote the class of maximal singular subspaces which contain planes. Then by Lemma 18.3.12, parts 1 and 3, $\mathcal{M}' \neq \emptyset$ and we have $|M \cap Q| = 1$ for each member M of $\mathcal{M}'$ and each member Q of $\mathcal{Q}$.

Lemma 18.3.13 *The elements of $\mathcal{Q}$ are pairwise disjoint.*

Proof Suppose by way of contradiction that Q_1 and Q_2 are distinct symplecta of $\mathcal{Q}$ for which $Q_1 \cap Q_2 \neq \emptyset$. Then by Lemma 18.3.11, $Q_1 \cap Q_2 := L$ is a line. Now choose $M \in \mathcal{M}'$. Then $M \cap Q_i = \{x_i\}$, a point, $i = 1, 2$.

If $x_1 \in Q_1 - L$ then some point of L is the unique point of Q_2 collinear with x_1. But as x_1 is collinear with x_2 (each being a point of the singular space M), we see that $x_2 \in L \subset Q_1$ forcing $x_2 = x_1$, the unique element of $M \cap Q_1$. Thus we cannot have $x_1 \in Q_1 - L$ and so we see

$$\{x_1\} = \{x_2\} = M \cap (Q_1 \cup Q_2) = M \cap L.$$

Next choose a point $p \in L - M$, and lines L_i in Q_i incident with p but distinct from line L, $i = 1, 2$. Then from the definition of $\mathcal{Q}$, $L_1 \cup L_2$ does not lie in a singular space, and so they lie together in a unique symplecton R.

We claim $R \in \mathcal{Q}$. If not, $R \in \mathcal{D}$ and so there is a plane π' incident with a line N of D. Again, as $Q_i \in \mathcal{Q}$, the line N cannot be either L_1 or L_2. But since D is a grid, it must intersect one of them and be opposite the other. Since there is a complete symmetry between $i = 1, 2$ relative to M, L and R, we may assume without loss of generality that N is opposite L_1. Now $N^\perp = (\pi')^\perp := M'$ is a member of $\mathcal{M}'$ and meets Q_2 at $N \cap L_2$ containing a point $r_2 \in L_2 - L_1$. But by Lemma 18.3.12, part 2(c), $Q_1 \in \mathcal{Q}$ forces $M' \cap Q_1$ to contain a unique point r_1. But then $r_1 \in {r_2}^\perp \cap Q_1 = \{p\}$, forcing $L_1 = pr_2 \subseteq N^\perp$ which is impossible for intersecting lines of a grid.

Now, having established that $R \in \mathcal{Q}$, $M \cap R$ contains a single point r (Lemma 18.3.12 again). But this forces r to be in ${x_1}^\perp \cap R$ which contains p as its unique point since $R \in \mathcal{Q}$. Thus $r = p$. But p was chosen in L distinct from x_1, against $M \cap L = \{x_1\} = \{x_2\}$.

Thus we cannot have two members of $\mathcal{Q}$ intersecting non-trivially, and the lemma is proved. □

Lemma 18.3.14 *Let Q_1 and Q_2 be distinct members of $\mathcal{Q}$ so $Q_1 \cap Q_2 = \emptyset$ (by Lemma 18.3.13). Let $f : Q_1 \to Q_2$ be the mapping which sends each point x of Q_1 to the unique point of Q_2 which is collinear with it. Then f is an isomorphism of polar spaces.*

Proof First f is injective since $Q_1 \in Q$. Each point of Q_2 is likewise collinear with a unique point of Q_1, in this way defining an inverse mapping $f^{-1} : Q_2 \to Q_1$. So f is a bijection of point sets.

Now suppose x_1 and y_1 are distinct points of Q_1 and set $x_2 := f(x_1)$ and $y_2 := f(y_1)$. If x_1 and y_1 are collinear, then there a symplecton $D = \langle\langle x_1y_1, y_1y_2\rangle\rangle$ which

shares point y_2 with Q_2. Since $D \cap Q_i \neq \emptyset, i = 1, 2$, and Lemma 18.3.13 forces $D \in \mathcal{D}$. Thus $D \cap Q_2$ is a line L (Lemma 18.3.11), and x_1 is collinear with a point of L, and this point must be x_2. Thus the collinearity of x_1 and y_1 implies the collinearity of their f-images x_2 and y_2. But the argument also holds with the Q_i transposed – that is, the collinearity of x_2 and y_2 implies the collinearity of their f^{-1}-images. Thus $f : Q_1 \to Q_2$ is an isomorphism of partial linear spaces. □

Lemma 18.3.15 *If $p \in Q \in \mathcal{Q}$, then all lines on p that are not contained in Q lie together in some maximal singular subspace.*

Proof Suppose L_1 and L_2 are distinct lines not contained in Q, but intersecting Q at the common point p. If $L_1 \not\subseteq L_2{}^{\perp}$, then there is a unique symplecton R containing them. But by Lemma 18.3.11, R must intersect Q at a line, and that line on p differs from either L_i. Thus R contains three lines on p and so cannot be a grid. Thus it is not in $\mathcal{D}$ and so it must be in $\mathcal{Q}$. But that contradicts the previous Lemma 18.3.13. Thus $L_1 \subseteq L_2{}^{\perp}$ for any two lines on p not in Q and the conclusion follows. □

Lemma 18.3.16 *Fix a maximal singular subspace M in the class $\mathcal{M}'$, and choose any point $x \in M$. For any two distinct lines N_1 and N_2, each intersecting M exactly at x, we have $N_1 \not\subseteq N_2{}^{\perp}$.*

Proof Suppose $N_1 \cap M = N_2 \cap M = \{x\}$, for distinct lines $N_i, i = 1, 2$. By way of contradiction assume that N_1 and N_2 generate a plane π. Then $\pi \cap M = \{x\}$.

Now choose $Q \in \mathcal{Q}$ (remember, the assumption here is that $\mathcal{Q}$ is not empty). Then $M \cap Q$ contains a unique point p. Now by Lemma 18.3.15 M contains all the lines on p which are not in Q. Whether $p = x$ or not, by Lemma 18.3.12, part 3, $\pi^{\perp}$ is a maximal singular subspace meeting Q at a single point and this point must be p, the unique point of $x^{\perp} \cap Q$.

Suppose $x \neq p$ so xp is a line. Then $\pi^{\perp}$ and M are distinct maximal singular spaces containing $(xp)^{\perp}$, against Lemma 18.3.12, part 2. Thus one must assume $x = p$. Then N_1, N_2, and all lines on x in M belong to the singular space $T(x)$ (see Lemma 18.3.7, p. 637, where Q is *not* assumed to have polar rank at least three).

It follows from the maximality of M that $\pi \subseteq M$. This is impossible because the lines N_i were chosen so that $N_i \cap M = \{x\}$.

The proof is complete.

Lemma 18.3.17 *The elements of $\mathcal{M}'$ are pairwise disjoint.*

Proof This is immediate from the previous Lemma 18.3.16. First, if $M_i \in \mathcal{M}'$, $i = 1, 2$, and $M_1 \neq M_2$, then $M_1 \cap M_2$ cannot contain a line, L for then $L^{\perp}$ is not a maximal singular subspace, as required by Lemma 18.3.12, part 2. But if $M_1 \cap M_2 = \{x\}$, then there are planes in M_2 intersecting M_1 just at x and that contradicts the conclusion of the preceding lemma. □

Lemma 18.3.18 *Fix $x \in M \in \mathcal{M}'$. All of the lines incident with x which do not lie in M belong to a common member of $\mathcal{Q}$.*

Proof First let us satisfy ourselves that there actually are two or more lines on x not in M. Since $\mathcal{Q}$ is non-empty, there is a symplecton Q in $\mathcal{Q}$ which by Lemma 18.3.12,

part 3, intersects M at some point y. If $y = x$, Q contains at least two lines on x. If $y \neq x$, for each line L of Q on y, the symplecton $\langle\langle xy, L\rangle\rangle$ is a grid in $D \in \mathcal{D}$ containing a unique line L' distinct from xy on point x. Since $D = \langle\langle L', xy\rangle\rangle$ as well, the correspondence $L \mapsto L'$ defines a bijection

$$\text{lines of } Q \text{ on } y \to \text{lines of } x \text{ not in } M.$$

Thus there are at least two lines on x not in M.

Choose two distinct lines L_1 and L_2 on point x not in M and form the symplecton $R := \langle\langle L_1, L_2\rangle\rangle$. Suppose $R \in \mathcal{D}$. Then there exists some plane π with $\pi \cap R = N$ being a line. Then N lies in a member M' of $\mathcal{M}'$ distinct from M and so by Lemma 18.3.17 $N \cap M = \emptyset$. Then, reindexing the L_i if necessary, we may assume that $N \cap L_1 = \{a\}$, a single point, and that N is opposite L_2 in R. Now the plane π contains at least another line A on a and the symplecton $R' := \langle\langle A, L_1\rangle\rangle$ is a grid of $\mathcal{D}$. Choose $p \in A - \{a\}$; then R' contains a unique point $q \in p^\perp \cap x^\perp - \{a\}$. Now choosing $b \in N - \{a\}$ and setting $\{d\} := b^\perp \cap L_2$ we see that the points (a, p, b, d, x, q) are precisely in the roles of (a, p, b, d, c, q) of the Tent lemma 18.3.10. It follows that q is in ${L_2}^\perp$. But then L_2 and xq are two distinct lines on x not in M which lie in a plane. But that contradicts Lemma 18.3.16.

Thus we have shown that the symplecton R defined at the beginning of the previous paragraph is a member of $\mathcal{Q}$. Now suppose B were a line on x not in either M or R. Then, by Lemma 18.3.15, B, together with all the lines of x in M form a maximal singular subspace. But that would contradict the fact that M is a maximal singular subspace. Thus no such line B can exist. The proof is complete. □

Lemma 18.3.19 *The points are partitioned in two ways as the point-shadows of the elements of $\mathcal{M}'$ and the point-shadows of elements of $\mathcal{Q}$:*

$$\mathcal{P} = \biguplus\{M | M \in \mathcal{M}'\},$$
$$\mathcal{P} = \biguplus\{Q | Q \in \mathcal{Q}\}.$$

Proof By Lemmas 18.3.17 and 18.3.13 it suffices to show that every point lies in a unique member of each class.

Fix $Q_0 \in \mathcal{Q}$, and $M_0 \in \mathcal{M}'$. This is possible since $\mathcal{Q}$ and $\mathcal{D}$ are both assumed to be non-empty. Moreover, these two objects must intersect, so we set $M_0 \cap Q_0 := \{e\}$. We would like to begin by showing that every point of Q_0 lies in a unique member of $\mathcal{M}'$. By Lemma 18.3.15 the union of all lines on such a point that are not lines of Q_0 form a maximal singular space, and this singular space contains a plane – that is, it is a member of $\mathcal{M}'$ – if and only if there are at least two lines on the point which are not in Q_0. So we begin with the following.

(Step 1) *If x and y are distinct collinear points of Q_0, then there is a bijection*

$$\mu_{xy} : \text{ lines of } x \text{ not in } Q_0 \to \text{ lines of } y \text{ not in } Q_0.$$

Proof This argument is similar to that producing a different bijection of lines in the preceding Lemma 18.3.18, using the "grid trick." If L is any line on x not in Q_0, the symplecton

$$D_{xy} := \langle\langle L, xy\rangle\rangle \in \mathcal{D}$$

is a grid, and so y lies in a unique line $\mu_{xy}(L)$ of D_{xy} distinct from xy. Clearly $D_{xy} = \langle\langle \mu_{xy}(L), xy\rangle\rangle$, so the maps μ_{xy} and μ_{yx} are inverses of each other.

(Step 2) *Each point of Q_0 lies in a unique member of $\mathcal{M}'$ which is the union of all lines on the point not in Q_0.*

Proof Now the point e in Q_0 lies on at least two lines not in Q_0 for the reason that it lies in M_0 which contains a plane. Now we may apply Step 1 and the fact that Q_0 is a connected point-line geometry, to reach the stated conclusion.

(Step 3) *Every point of $\mathcal{P}$ lies in a member of $\mathcal{M}'$.*

Proof The conclusion holds for every point of Q_0 by Step 2, so we suppose p is an arbitrary point of $\mathcal{P} - Q_0$. Then p is collinear with a unique point p' of Q_0. Then pp' is one of the lines of p' not in Q_0, and so by Step 2 lies in a member of $\mathcal{M}'$. So p is in a member of $\mathcal{M}'$.

(Step 4) *Every point lies in a member of $\mathcal{Q}$.*

Proof By Step 3, each point p lies in a member M of $\mathcal{M}'$. Then by Lemma 18.3.18, the lines of p which are not in M are all the lines on p of a unique symplecton $Q_p \in \mathcal{Q}$.

Since we have shown that the union of all the subspaces of each one of the classes $\mathcal{M}'$ or $\mathcal{Q}$ comprise all the points, we have the desired partitions. The proof is complete. □

Corollary 18.3.20 *Assume hypotheses (PL) and (PL*), and assume both classes of symplecta $\mathcal{Q}$ and $\mathcal{D}$ are non-empty, and that no symplecton has polar rank three or more. Then Γ is a product geometry $M \times Q$ where M is some sort of singular space containing a plane (that is, a singular space generated by a non-incident point-line pair), and Q is a generalized quadrangle.*

18.3.5 What We Have Proved

So far we have seen that the hypotheses (PL) and (PL*) force the symplecta of Γ to be of two types: $\mathcal{D}$, the class of symplecta which meet some plane in exactly a line,[6] and those symplecta S such that $|x^{\perp} \cap S| = 1$ for all $x \in \mathcal{P} - S$.

[6] With the meaning of "plane" used before – a singular subspace generated by a line and a point not incident with the line.

We have considered the case that $\mathcal{D} = \emptyset$ in Sect. 18.3.2. There we found that there were three possibilities: (1) $\Gamma = C_{3,3}$, a dual polar space of a non-oriflame polar space with thick lines of polar rank three, (2) $\Gamma \simeq L \times Q$, the product of a line L and a generalized quadrangle Q which is not a grid, or (3) $\Gamma \simeq L_1 \times L_2 \times L_3$, the product of three lines.

We next considered the case that both $\mathcal{D}$ and $\mathcal{Q}$ were non-empty, first in the case that there existed a symplecta of rank at least three (always a member of $\mathcal{Q}$) and in the case that Γ had symplectic rank exactly two. In both cases, $\Gamma \simeq M \times P$ where P is a non-degenerate polar space of (possibly infinite) polar rank at least two, and M is any singular subspace of any sort that properly contains a line.

That leaves the final case $\mathcal{S} = \mathcal{D}$, that is, $\mathcal{Q} = \emptyset$. Here many new bizarre examples are possible and the structure becomes too amorphous to make either a simple, useful, or even pleasing classification.

So we may at least record what we have achieved.

Theorem 18.3.21 (Local Main Theorem) *Suppose* $\Gamma = (\mathcal{P}, \mathcal{L})$ *satisfies the following hypotheses:*

(PL) Γ *is a strong parapolar space satisfying the Pentagon Property;* Γ *is not itself a polar space.*

(PL)* *If* S *is a symplecton for which* $x^\perp \cap S = \emptyset$ *for some point* $x \in \mathcal{P} - S$*, then there exists a point* $y \in \mathcal{P} - S$ *for which* $y^\perp \cap S$ *contains a line. (The contrapositive of this assertion would be this: if a symplecton* S *does not intersect any plane* π *exactly at a line* L *– i.e.,* $S \cap \pi = L$ *– then* $x^\perp \cap S \neq \emptyset$ *for all* $x \in \mathcal{P} - S$*.)*

(Q) *There exists at least one symplecton* S *for which* $|x^\perp \cap S| \leq 1$ *for all* $x \in \mathcal{P} - S$*.*

Then one of the following cases must occur:

1. ($\mathcal{D} = \emptyset$.)
   - *(a)* $\Gamma \simeq C_{3,3}$*, a dual polar space with thick lines.*
   - *(b)* $\Gamma \simeq L \times Q$*,* L *is a line and* Q *is a non-grid generalized quadrangle.*
   - *(c)* $\Gamma \simeq L_1 \times L_2 \times L_3$*, the product of three lines.*
2. ($\mathcal{D} \neq \emptyset$ or a symplecton of rank at least three exists.) $\Gamma \simeq M \times P$ *where* M *is an arbitrary linear space which properly contains some line, and* P *is a non-degenerate polar space of (possibly infinite) polar rank at least three.*

Recall that if (PL) and (PL*) are in place, then each symplecton of polar rank three or more belongs to $\mathcal{Q}$. Then we have the following.

Corollary 18.3.22 *Assume* Γ *satisfies conditions (PL) and (PL*) of the previous theorem. Instead of (Q), assume:*

(Q') There exists a symplecton of polar rank three or more.

Then $\Gamma \simeq M \times P$ where M is an arbitrary linear space which contains some line, and P is a non-degenerate polar space of (possibly infinite) polar rank at least two.

18.4 The Proof of the Main Theorem

18.4.1 The Hypothesis (H)

Throughout this section we suppose Γ to be a parapolar space of symplectic rank at least three.

We shall be concerned with the following hypothesis:

(H) (a) *The Pentagon Property holds in each point-residual $\Gamma_p = (\mathcal{L}_p, \Pi_p)$.*
(b) *Given a symplecton S and a point $z \in \mathcal{P} - S$ for which $z^\perp \cap S$ contains just one point p, there must exist a point $y \in (\mathcal{P} - S) \cap p^\perp$ for which $y^\perp \cap S$ contains a plane.*
(c) *Γ is not itself a polar space.*

Lemma 18.4.1 *A parapolar space Γ of symplectic rank at least three satisfies property (H) if and only if each point-residual $\Gamma_p = (\mathcal{L}_p, \Pi_p)$ satisfies the Pentagon Property and condition (PL*).*

Proof Now suppose every point-residual $\Gamma_p = (\mathcal{L}_p, \Pi_p)$ satisfies the Pentagon Property. Hypothesis (PL*) for Γ_p asserts that if S is a symplecton on p for which there is a line $L := zp \in \mathcal{L}_p$ which is not coplanar with any line of $\mathcal{L}_p \cap S$ (i.e., $z^\perp \cap S = \{p\}$) then there exists a further line $N = py$ which is coplanar with each of the lines of a plane of S lying on point p (in other words, $y^\perp \cap S$ contains a plane on p). So condition (H) holds for Γ.

Clearly if (H) holds for Γ, then any point-residual $\Gamma_p = (\mathcal{L}_p, \Pi_p)$ satisfies the Pentagon Property. Moreover, if S_p represents the collection of points and planes of a symplecton S which are incident with one of its points p, then the existence of a line $L = zx$ which is not coplanar with any line of $\mathcal{L}_p \cap S$, implies the existence of another line $N = yp$ with $N^\perp \cap S$ containing a plane on p – that is, axiom (PL*) holds for $\Gamma_p = (\mathcal{L}_p, \Pi_p)$. □

Remark It is important that we did not assert that if (H) holds for Γ then $\Gamma_p = (\mathcal{L}_p, \Pi_p)$ satisfies (PL) as well as (PL*). The reason is that axiom (PL) asserts (among other things) that Γ_p is a strong parapolar space, and the definition of strong parapolar space (as a species of parapolar space) requires that the geometry be connected. But we have not been handed the fact that Γ_p is connected in hypothesis (H).

Assume Γ satisfies (H). Then for every point-residual Γ_p we have inherited the Pentagon Property and, if Γ_p is not itself a polar space, we also inherit a division

$$\mathcal{S}_p = \mathcal{Q}_p + \mathcal{D}_p$$

of the symplecta on p into two classes: (1) symplecta $D_p \in \mathcal{D}_p$ for which there exists a line $L \in \mathcal{L}_p - \mathcal{L}_p(D_p)$ such that $L^\perp \cap D_p$ contains a plane, and (2) symplecta $Q_p \in \mathcal{Q}_p$ such that *for any line* $L \in \mathcal{L}_p - \mathcal{L}_p \cap Q_p$ we have that $L^\perp \cap Q_p$ is exactly a line.[7]

We say that a point-symplecton flag (p, S) is *type* $\mathcal{D}$ if it has the property (1) and we say that it is *type* $\mathcal{Q}$ if it has property (2).

Remark Note that if Γ_p is itself a polar space, then certainly $\mathcal{D}_p$ is empty since condition (1) asserts the existence of a line on p exterior to a symplecton on p. But on the other hand, condition (2) is vacuously true in this case. In this case we see that p is a deep point of some symplecton S. This is a case we would eventually like to get rid of. We cannot do this until it is established that Γ is locally connected.

But luck is on our side.

Lemma 18.4.2 *If* Γ *satisfies (H), and* (p, S) *is a point-symplecton flag of type* $\mathcal{Q}$*, then the point-residual* $\Gamma_p = (\mathcal{L}_p, \Pi_p)$ *is connected.*

Proof By the definition of type Q, for every line L on p which is not contained in S, L is coplanar with a unique line of S lying on p. But since Γ has symplectic rank at least three, the lines and planes of S which are incident with p form a non-degenerate polar space of rank at least two which is, of course, connected. It follows that Γ_p – the geometry of all lines and planes on p – is connected. □

Lemma 18.4.3 *Suppose* Γ *satisfies (H). Assume, furthermore, that* (p, S) *is a point-symplecton flag of type* $\mathcal{Q}$*.*

1. *If the point-residual* Γ_p *is not itself a polar space then it has one of the following two forms:*
   (a) (p, S) *is type* $\mathcal{Q}$ *for every symplecton* S *on* p *and* Γ_p *is either (i) a classical dual polar space of a non-oriflame polar space of polar rank three having thick lines, (ii) a product geometry* $L \times Q$*, where* L *is a line and* Q *is a generalized quadrangle, or (iii) a product* $L_1 \times L_2 \times L_3$ *of three lines.*
   (b) Γ_p *is a product geometry* $M_p \times Q$ *where* M_p *is a projective space of (possibly infinite) rank at least 1, and* Q *is a non-degenerate polar space of (possibly infinite) polar rank at least two. (In this case the projective space* M *contains a line of the residue geometry* Γ_p*.)*
2. *For each line* L *on point* p*, there exists a symplecton* S' *containing* L*, such that* S' *is isomorphic to* S*, and* (p, S') *is a point-symplecton flag of type* $\mathcal{Q}$*. In more detail, if* Γ_p *is not a polar space and is in case (b) of the preceding conclusion, the choice of* S' *containing* L *so that* (p, S') *is type* $\mathcal{Q}$ *is unique. On the other hand, if* Γ_p *is in case (a), line* L *is in at least two distinct isomorphic symplecta* S_i *for which* (p, S_i) *is type* $\mathcal{Q}$*,* $i = 1, 2$*.*

[7] The reader should appreciate that the division $\mathcal{S}_p = \mathcal{Q}_p + \mathcal{D}_p$ was not called a *partition* since it is possible that one of the two sets $\mathcal{Q}_p$ or $\mathcal{D}_p$ is empty.

Proof 1. By Lemma 18.4.2, Γ_p is connected, and so is a strong parapolar space satisfying axiom (PL) as well as the axiom (PL*), which follows from Lemma 18.4.1. Then if Γ_p is not itself a polar space, Theorem 18.3.21 shows that Γ_p must have one of the indicated structures, (a) or (b).

In case (b) $\Gamma_p \simeq M_p \times Q$. Since Γ is a gamma space, the union M' of the pairwise coplanar lines on p which comprise "Points" of M'_p form a singular subspace of Γ. Since Γ has symplectic rank at least three, M' is a projective space by Theorem 13.4.1 of Chap. 13. Then M_p is also a projective space.

2. If Γ_p is a polar space, then p is a deep point of a symplecton S, and so any line L on p lies in S and the conclusion is true with $S' = S$. If Γ_p is not a polar space, by part 1, there are the two cases, (a) and (b).

In case (a) Γ_p is the dual polar space formed from a non-oriflame polar space of rank three. In this dual polar space each "Point" is in a pair of "Quads" $Q_i, i = 1, 2$ which are isomorphic, are not grids, and belong to the class $\mathcal{Q}$. Thus each Q_i has the form $(S_i)_p$ for non-oriflame symplecta S_i of rank three. The isomorphism of S and S' follows from the theorem of Tits – Theorem 7.5.13 of Chap. 7.

In case (b) Γ_p has the form $M_p \times Q$ where $Q = (S \cap \mathcal{L}_p, S \cap \Pi_p)$, and M_q is a singular subspace. That means every "Point" L of the point-residual Γ_p lies in a "Symplecton" S'_p isomorphic to Q and in a singular space M'_p isomorphic to M_p. Then S_p is just the lines and planes incident with p which lie in a symplecton S' which contains line L. At this stage, S and S' are two symplecta of rank at least three, with $S \cap S' = \{p\}$, and with isomorphic point-residues at point p. The isomorphism of S and S' now follows from the theorem of Tits – Theorem 7.5.13 of Chap. 7.

The proof is complete. □

Now we can prove the following.

Lemma 18.4.4 *Suppose Γ is a parapolar space of symplectic rank at least three which satisfies hypothesis (H). If one of its point-residuals is a polar space, then Γ is itself a polar space.*

Proof Assume the hypothesis and suppose, for some point p, the point-residual Γ_p is a polar space. Then of course this polar space has polar rank at least 2 since it is itself the set of lines and planes incident with a point-symplecton flag (p, S) where p is a deep point of the symplecton S of polar rank at least three.

If we knew in advance that Γ was locally connected, we could simply invoke Theorem 13.5.5 of Chap. 13 to conclude that $\mathcal{P} = S$. Unfortunately, as noted above, the hypothesis (H) does not give us local connectedness. But there is still strength in our hypotheses.

Suppose by way of contradiction that $\mathcal{P} \neq S$. Since $\Gamma = (\mathcal{P}, \mathcal{L})$ is connected, S must contain non-deep points as well as deep points. Since S is connected, there is an instance where a deep point p of S is collinear with a point q of S that is not deep in S.

Now that means Γ_q is not a polar space. Thus S either belongs to $\mathcal{D}_q$ or $\mathcal{Q}_q$.

First suppose the flag (q, S) is type $\mathcal{D}$. Then there exists a point y in $(\mathcal{P} - S) \cap q^\perp$ such that $y^\perp \cap S$ contains a plane π on point q. Then $p^\perp \cap \pi$ contains a line,

so (p, y) is not a special pair of the parapolar space Γ. Thus there is a symplecton R containing p and y, and $R \cap S$ is a singular subspace. It follows that since $p^\perp \cap R$ cannot be a singular space, there is a line of R incident with p which does not lie in S. But that contradicts the assumption that p was a deep point of S.

Thus we are forced to assume that the flag (q, S) is type $\mathcal{Q}$. By Lemma 18.4.2 the point-residual Γ_q is a connected geometry. Then by Lemma 18.4.3 there are two possiblilites for the point-residual at q: (1) Γ_q is a classical dual polar space of polar rank three or (2) $\Gamma_q \simeq M_q \times S_q$ where S_q denotes the lines and planes incident with the flag (q, S) and M_q is the geometry of lines and planes of a maximal singular subspace M which is not itself a plane. Since p is a deep point of S, every plane on the line N must lie in S. That means the line $N := pq$ is a deep "Point" in S_q. But neither alternative (1) nor (2) possess any symplecta of Γ_q with deep "Points." Thus p cannot be a deep point of S contrary to assumption.

Thus $\mathcal{P} = S$ is forced, and the proof is complete. □

18.4.2 The Axiom (Q)*

Where is this going? The theorem that we shall eventually prove will require just one extra axiom:

(Q*) *Somewhere in* Γ *there exists a point-symplecton flag* (p, S) *with this property: for every line* L *on* p *which is not incident with* S, $L^\perp \cap S$ *is simply a line on* p.

Using the definitions above we may rewrite the assertion of axiom (Q*) in the following two equivalent ways:

1. For some point p the set $\mathcal{Q}_p$ is non-empty.
2. There exists a point-symplecton flag of type $\mathcal{Q}$.

We first observe the following.

Lemma 18.4.5 *Suppose* Γ *is a parapolar space satisfying hypothesis (H). If* Γ *is not a polar space, and contains a symplecton of rank at least four, then* Γ *satisfies condition (Q*).*

Proof Let p be any point in a symplecton S of Γ of polar rank at least four. Since Γ is not a polar space, Lemma 18.4.4 tells us that p is not a deep point of S. Now by Lemma 18.3.7 on p. 637, we see that $S \in \mathcal{Q}_p$ – that is, it follows that the flag (p, S) is type $\mathcal{Q}$. Thus (Q*) holds. □

18.4.3 Local Uniformity

In this section we wish to prove certain "uniformity" results for a parapolar space satifying hypotheses (H) and (Q*) – that is, if one point-residual Γ_p of Γ has a

certain shape, then every point-residual of Γ has the same shape. Naturally the argument will rest on the connectedness of Γ.

We already know from Lemma 18.4.4 that if even one point-residual of Γ is a polar space, then Γ is itself is a polar space. Our starting point is thus the existence of a point-symplecton flag (p, S) of type $\mathcal{Q}$. By Lemma 18.4.3 there are two possiblilites for the point-residual at p: (1) Γ_p is a classical dual polar space of polar rank three or (2) $\Gamma_p \simeq M_p \times S_p$ where S_p denotes the lines and planes incident with the flag (p, S) and M_p is the geometry of lines and planes of a maximal singular subspace M which is not itself a line. In either case, each line L on p lies in at least one symplecton S_L, and the pair (L, S_L) satisfies this property:

(Q-local) *Every plane π on line L which is not already contained in the symplecton S_L, satisfies $\pi^\perp \cap S_L = L$.*

That means that if q is a point that is collinear with p via line L, then (q, S_L) is also a flag of type $\mathcal{Q}$. Therefore, in turn, the point-residual $\Gamma_q = (\mathcal{L}_q, \Pi_q)$ is also a connected geometry and has one of the two forms (1) or (2), listed in the previous paragraph.

But the two cases can be distinguished relative to the line L. In case (1) for Γ_p, every line L on p lies in at least two distinct symplecta S and S' such that (L, S) and (L, S') both satisfy the hypothesis (Q-local), whereas in case (2) each line L on p satisfies the property (L, S) and is (Q-local) for a unique symplecton S on L.

It follows that for any point q collinear with point p via a line L, the point-residual Γ_q is connected and belongs to case (1) or (2) accordingly as Γ_p belongs to it.

We can even go beyond that – the two point-residuals Γ_p and Γ_q are even isomorphic.

In case (1) we have two things to help us. (a) Γ_p is a dual polar space and in it a "Point-residual" is a projective plane – that is, the collection of planes and symplecta over a line L has the strucure of a projective plane. The same plane structure exists relative to any point $q \in L - \{p\}$. (b) For each symplecton S on line L, the point-residual of S at p is isomorphic to the point-residual at q – that is, $\Gamma_p \cap S \simeq \Gamma_q \cap S$. It follows from (a) and (b) that the dual polar spaces Γ_p and Γ_q are isomorphic.

In case (2) $\Gamma_p \simeq M_p \times Q$, so that the line $L = pq$ lies in a unique projective space M and a unique symplecton S with the property that (L, S) satisfies condition (Q-local), and that

$$\text{projective rank}(M) + \text{polar rank}(S) \geq 6.$$

The isomorphism type of Γ_p is completely determined by the rank of M_p and the isomorphism type of Q. But M and S sit on line L, and so the point-residual $\Gamma_q \simeq M_q \times S_q$ has M_q being the same projective rank as M_q and the two point-residuals of S at points p and at q are isomorphic – that is, $Q \simeq S_q$. One concludes that Γ_p is isomorphic to Γ_q in this case.

From the connectedness of the parapolar space we may conclude the following.

Theorem 18.4.6 (Uniformity Theorem.) *Assume* Γ *is a parapolar space of symplectic rank at least three satisfying hypotheses (H) and (Q*). Then* Γ *is locally connected. More specifically, four cases hold:*

1. *For each point* p *the point-residual* Γ_p *is a dual polar space of a classical non-oriflame polar space with polar rank three, having a fixed isomophism type.*
2. *For each point* p*, the point-residual* $\Gamma_p \simeq L \times Q$*, where* L *is a line and* Q *is a generalized quadrangle of a fixed isomorphism type that is not a grid.*
3. *For every point* p*,* Γ_p *is a product of three lines* $L_1 \times L_2 \times L_3$*.*
4. *For each point* p*, the point-residual* Γ_p *is a product geometry of the form* $M \times P$ *where* P *is a fixed non-degenerate polar space (a classical quadrangle in the case where* P *has polar rank two) whose isomorphism type does not depend on the particular point* p*, and* M *is a classical projective space defined over the same division ring as* P *— that is* $M \times P$ *has a unique isomorphism type which is independant of the point* p*. Morover, the projective rank of* M *and the polar rank of* P *are each at least two.*

18.4.4 The Metasymplectic Case

As before, we assume $\Gamma = (\mathcal{P}, \mathcal{L})$ is parapolar space satisfying (H) and (Q*). Then one of the four cases of Theorem 18.4.6 holds. In this section we assume it is one of the first three cases (for which $\mathcal{D}_p$ is always empty). Then Γ *is a parapolar space is which each point-residual* Γ_p *is a near hexagon in which each quad is gated.* This case was fully discussed in Theorem 18.3.6, p. 636, which records that Γ is a metasymplectic space as defined and classified by Tits in [138]. Thus, for the three options of a point-residual, we may conclude that (1) Γ is the Lie incidence geometry $F_{4,1}$ of a thick building of type F_4, (2) a polar Grassmannian of lines of a non-oriflame non-degenerate polar space of polar rank four (for example, B_4, C_4, polar spaces defined by σ-Hermitian form, and real forms of certain exceptional Lie Algebras, similar forms over algebraic number fields etc. – see Tits [138]), and (3) the Grassmannian of lines of an oriflame polar space of rank four – the Lie incidence geometries of type $D_{4,2}$.

18.4.5 Beyond the Metasymplectic Case: A Related Geometry

As before, we assume $\Gamma = (\mathcal{P}, \mathcal{L})$ is parapolar space satisfying (H) and (Q*). But now we assume that only Case 4 of Theorem 18.4.6 holds. Specifically:

> *For each point* p *of* Γ*, the point-residual* Γ_p *is a product geometry of the form* $M \times P$ *where* P *is a fixed non-degenerate polar space (a classical quadrangle in the case* P *has polar rank two) and* M *is a classical projective space defined over the same division ring as* P*. Morover the projective dimension of* M *and the polar rank of* P *are at least two.*

This uniform class of isomorphic point-residuals produces certain classes of subspaces of the global point-line geometry Γ.

The singular subspaces M, in $\Gamma_p \simeq M \times P$, correspond to a class $\mathcal{M}^*$ of maximal singular subspaces of Γ of projective dimension at least three with this property:

No member of $\mathcal{M}^*$ is contained in a symplecton.

The residual symplecta P that appear in $\Gamma_p \simeq M \times P$ come from the flags (p, S) of type $\mathcal{Q}$ and so correspond to the class of all symplecta $\mathcal{Q}^*$ of polar rank at least three, with this property:

If $S \in \mathcal{Q}^*$, then for each point x in S, $x^\perp \cap S$ is either empty, or is a line.

Thus, from the point-line geometry Γ, one derives a rank three incidence geometry $\Gamma^* = (\mathcal{Q}^*, \mathcal{P}, \mathcal{M}^*)$ with the following incidences. A point is incident with any subspace of $\mathcal{M}^* \cup \mathcal{Q}^*$ that contains it. A member of $\mathcal{Q}^*$ is incident with a member of $\mathcal{M}^*$ if and only if they intersect at a line.

Remark Of course certain other subspaces of Γ are not singled out in defining the new geometry Γ^*. There are still further symplecta – the oriflame ones of polar rank three which correspond to the grids of the point-residuals. Also, every line of Γ lies in exactly one member of $\mathcal{Q}^*$ and exactly one member of $\mathcal{M}^*$ — thus these lines are just rank-2 flags of Γ^* of type $\{\mathcal{Q}^*, \mathcal{M}^*\}$. In this way, each projective space $M \in \mathcal{M}'$ becomes its rank two residue in the geometry Γ^*. Similarly, each symplecton Q in $\mathcal{Q}$ can be regarded as its residue in Γ^* – that is, the geometry $\mathrm{Res}_{\Gamma^*}(Q)$, a geometry over $(\mathcal{P}, \mathcal{M}^*)$.

18.4.6 Regarding Γ^* as a Locally Truncated Geometry

At this point we serve one final course to our banquet: we adjoin one further assumption to the hypotheses (H) and (Q*) already introduced:

(F) *The parapolar space Γ has finite singular rank – that is, every singular subspace of Γ has finite rank as a projective space.*

So far, the uniform shape of the point-residuals of Γ force Γ^* to be a locally truncated diagram geometry. Because of the finiteness condition (F), Γ^* belongs to the locally truncated diagram given in Fig. 18.2, where (since elements of $\mathcal{Q}^*$ have polar rank at least three) there is at least one box to the right of the node labelled $\mathcal{M}^*$.

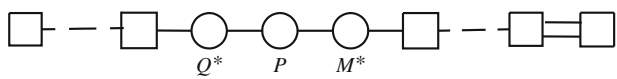

Fig. 18.2 The locally truncated diagram for the rank-three geometry Γ^*

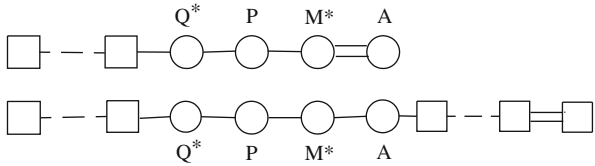

Fig. 18.3 The locally truncated diagrams of Γ^{**}

18.4.7 Enriching Γ^* to a Rank Four Geometry via Hanssen's Principle

So far, Γ^* does not quite convey all of the information that we have on hand. There is the extra fact that $\Gamma = (\mathcal{P}, \mathcal{L})$ is a parapolar space, and that produces a second class $\mathcal{D}$ of oriflame symplecta of polar rank three – that is, symplecta of type D_3.

We can now invoke "Hanssens principle," Theorem 13.6.1 (Chap. 13, p. 472), to infer that in Γ^* there exists a class $\mathcal{A}$ of Grassmannians of lines of type $A_{m,2}$ where $m-1$ is the uniform projective rank of a member of $\mathcal{M}'$. One is now presented with a rank four geometry

$$\Gamma^{**} = (\mathcal{Q}^*, \mathcal{P}, \mathcal{M}^*, \mathcal{A}),$$

belonging to one of the two locally truncated diagrams of Fig. 18.3, depending on whether elements of $\mathcal{Q}^*$ have polar rank three or polar rank greater than three. (The details are spelled out fully in this particular case in Exercises 13–18 of Chap. 13, p. 491.)

18.4.8 The Conclusion that Γ Is a Homomorphic Image of a Polar Grassmannian

In either case, the locally truncated diagram has the form given in the second part of Theorem 11.3.2 of p. 427. Thus by Theorem 11.3.2 there exists a sheaf $\mathcal{F}$ for this locally truncated diagram geometry. Since the C_3-residues of the chamber system $\mathbf{C}(\mathcal{F})$ are derived from classical polar spaces, they are covered by buildings. We thus conclude that $\Gamma^* = (\mathcal{Q}^*, \mathcal{P}, \mathcal{M}^*)$ is the homomorphic image of a truncation to types $(k-1, k, k+1)$ of a building Δ of type $(C/B)_n$ with the diagram in Fig. 18.4, where $k \leq n-2$ and if $k = n-2$, then $k \geq 2$.

Stated more exactly, there exists a geometry epimorphism:

$$f : \mathrm{trunc}_{k,\{k-1,k+1\}}(\Delta) \to \Gamma = (\mathcal{P}, \mathcal{L}). \tag{18.2}$$

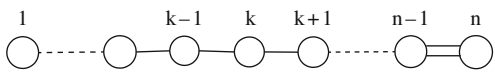

Fig. 18.4 The diagram of the building geometry Δ

The reader will recall from Part I, Sect. 2.4.4 (p. 53) that in the category of geometries, truncation is a functor – that is, truncations commute with morphisms. That means that $f(\Delta) = \Gamma$ is obtained by truncating a homomorphic image Δ/A of Δ, where A is a group of automophisms which acts freely on chamber-flags of Δ. But of course the truncation of that image has to be a parapolar space. Thus we can say the following.

Lemma 18.4.7 *The parapolar space Γ is isomorphic to the truncation to types $\{k, \{k-1, k+1\}\}$ of a homomorphic image Δ/A of a building Δ of type C_n, where $n \geq k+2$ and $n \geq 5$. For each non-identity element a of the group A, and each singular k-space U of the polar space $\Delta_{1,2}$, one sees that U^a has distance at least five from U in the classical Lie incidence geometry $C_{n,k} = trunc_{\{k,\{k-1,k+1\}\}}(\Delta)$.*[8]

18.4.9 The Simple Connectedness of the Geometry Γ in the Case that $k = 2$

Suppose the parameter k is the number 2. That means that every member of $\mathcal{M}'$ is a plane and that the symplecta of $\mathcal{Q}$ all have rank at least four (recall we are in the case that $\mathcal{D}$ is non-empty). Of course we possess other oriflame symplecta $\mathcal{D}$ of rank three exactly. But at least we are in the case that every point-residual has the form $L \times Q$ where L is a line and Q is a symplecton of rank at least three. Clearly each point-residual now possesses two important properties:

1. *In the residual Γ_p, the collection of "Points" at distance at most two from a given "Point" form a subspace.*
2. *For each symplecton S_p of Γ_p (whether in $\mathcal{Q}_p$ or $\mathcal{D}_p$) and every "Point" ℓ of $\Gamma_p - S_p$, ℓ is "collinear" with at least one "Point" of S_p.*

In the language of Γ this means that Γ is a locally connected parapolar space of symplectic rank at least three with these two properties:

1. *If line L intersects a plane π at a unique point p, and if L shares a symplecton with at least two lines of π on point p, then L shares a symplecton with every line of π on p. (See Sect. 13.7, p. 478, where the relationship of these properties are examined in full.)*
2. *For any point $p \in \mathcal{P}$ not incident with a symplecton S of Γ, the set $p^\perp \cap S$ can never be the single point-set $\{p\}$.*

Of course the reader will recognize these two assertions as the defining axioms (H1) and (H2) that make a parapolar space of symplectic rank at least three, a hexagonic geometry as presented Sect. 13.7 and in the preceding Chap. 17. In the latter chapter, we gave an elementary argument from first principles that the point-collinearity graph of such a geometry was simply connected (Theorem 17.3.1).

---

[8] Note that this distance is not the codimension of $U^a \cap U$ in U (as it would be in ordinary Grassmannians). The discrepancy has to do with the size of the radical of $\langle U, U^a \rangle / (U \cap U^a)$.

Now the morphism f of Eq. (18.2) induces a fibering of the point-collinearity graphs. Therefore that graph morphism is an isomorphism. Since every line is defined by its point-shadow (after all f is an epimorphism between two partial linear spaces) it follows that f is an isomorphism, and we are done.

Summing up, Γ is either a polar space, one of the three classes of metasymplectic spaces, or a particular homomorphic image of a classical polar Grassmannian described in Lemma 18.4.7. The main Theorem (Theorem 18.1.3) is proved.

18.5 Final Comments on the Main Theorem

There are certain insights to be gained in introducing the axioms one by one. It isn't only that lemmas appear in their optimal setting. One begins to wonder why one feels forced to introduce a new axiom and begins to wish there were another way around it.

The Main Theorem of this chapter is no exception to this rule. Surely every reader noticed that we had reached exact information about the nature of point-residuals without any need to invoke finite singular rank. And in fact we could proceed with two of the conclusions without using it. We did not need finite singular rank in the case that Γ was a polar space. That was because of Peter Johnson's theory of polar polar spaces of infinite rank, which was painfully incorporated into our development in Chap. 7.[9] We did not need finite singular rank in the metasymplectic case because we inherited this property by case-division. The only place that we needed to invoke the axiom (F) was when all point-residuals had the shape $M \times P$ with the lower bounds on the sum of the ranks of M and P.

Why was this? We had full knowledge of the point-residuals (at least as truncated geometries). We simply do not possess a method of passing from this local information to the the global without invoking (1) the sheaf-theory introduced by Ronan and Brauer–Cohen, coupled with (2) Tits' "Local Approach Theorem" (Theorem 10.6.1, Chap. 10, p. 412). So one must ask whether it was (1) or (2) that required the assumption of finite singular rank. Well it was only partly (1). As developed here, finite rank was never used in defining a sheaf. It all rested on the *rigidity* of certain rank two residues. Whether these rank two truncations were truncations of finite or infinite geometries did not matter. The problem comes in the formation of a chamber system from the sheaf. That is not automatic in infinite rank, though it can occur.

The real problem here is Tits' local approach theorem. We know from [84] that buildings of infinite rank (in the guise of chamber systems) – even if they could play a role as universal covers of a chamber system $\mathbf{C}(\mathcal{F})$ – are never residually

[9] Of course Tits could have attended to these details, but as Oparin said, truly great innovators are condemned to see several generations of people refining the work they had not time for en route to their discoveries. (I confess this is a generous paraphrase of what Oparin (in his book *The Origin of Life*) actually said about Pasteur.)

connected, and so one loses the connection between chamber systems and geometries, so vital in completing the Main Theorem.

One last comment. One notices that the special simply-connected result about polar Grassmannians of lines came from recognizing that these were hexagonic geometries. Yet the analysis of the latter in Chap. 16 never used the Pentagon Property. This suggests that the Pentagon Property was only needed to handle polar Grassmannians of isotropic k-spaces for $k > 2$. Now we know why the Main Theorem of this chapter and the previous one marched their separate ways. The bifurcation of the two theories would not be so obvious if we tried to gorge ourselves with these axioms all at once. Civilized logical meals in which the axioms are introduced to the gourmet one by one in a certain order can be more satisfying.[10]

18.6 Exercises for Chapter 18

The following exercises invite the student to probe the limits of the Main Theorems of this chapter.

18.6.1 Relaxing Hypothesis (PL)*

We have seen that if Γ is a strong parapolar space with the Pentagon Property, there is a natural division of the symplecta into two classes: $\mathcal{D}$, a class of grids one of whose lines properly lies in a larger singular subspace, and $\mathcal{Q}$, a class of symplecta, all of whose lines are maximal singular subspaces.

In proving the "Local Main Theorem" (Theorem 18.3.21, p. 646) we shouldered the burden of two further hypotheses:

(PL*) Every point is collinear with at least one point of each symplecton in class $\mathcal{Q}$.

(Q) The class $\mathcal{Q}$ is non-empty.

18.1 Suppose P is a classical polar space of finite polar rank at least three. Then the associated dual polar space DP is the point-line geometry whose points are the maximal singular subspaces of P and whose lines are are the hyperplanes of those subspaces, with the usual incidence. If P is not an oriflame polar space – that is, one of the polar spaces associated with the classical groups of $O + (2n, F)$ (acting on a quadric $Q^+(2n-1, F)$) – then the lines of DP are thick, and this is a class of strong parapolar spaces which was characterized by Cameron's theorem in Chap. 8.

1. Prove that DP possesses the Pentagon Property. [Hint: Use one of the hypotheses of Cameron's theorem, the gatedness of symplecta.]

[10] At least to the cook.

2. Suppose A is a group of automorphisms of P such that for every non-identity element a of the group A, and maximal singular subspace M of the polar space P, $M \cap M^a$ has co-dimension at least five in either M or M^a (they are necessarily the same). Then, of course, each element of A permutes singular subspaces of projective rank one less than the polar rank among themselves, while preserving incidence with maximal singular subspaces. Thus each element $a \in A$ induces an automorphism $\bar{a}$ of the dual polar space DP. We let $\bar{A}$ denote the group of automorphisms of DP induced in this way. We can now form the epimorphism

$$DP \to DP/\bar{A},$$

whose points and lines are the $\bar{A}$-orbits of points and lines of DP, two such $\bar{A}$-orbits being incident if and only if a flag of DP appears in the Cartesian product of the two orbits.

Show that DP/A is a strong parapolar space which possesses the Pentagon Property, with the class $\mathcal{D}$ of symplecta being empty. Show that (PL^*) must fail for $DP/\bar{A}$ unless P has polar rank at most 3.

18.2 Can there exist strong parapolar spaces with the Pentagon Property in which there are symplecta of rank at least three?

Specifically, which of the following classical strong parapolar spaces possess the Pentagon Property?

1. The Grassmann spaces $A_{n,k}(D)$.
2. The half-spin geometries, $D_{n.n}(K)$.
3. $E_{6,1}(K) \simeq E_{6,6}(K)$.
4. $E_{7,7}(K)$.

Using characterizations and examples from previous chapters, justify your answer.

18.3 Give an explicit proof that if X is a connected convex subspace of a product $\Pi_{\sigma \in I} Y_\sigma$ a collection of point-line geometries $Y_\sigma = (\mathcal{P}_\sigma, \mathcal{L}_\sigma)$ each of which is a strong parapolar space with the Pentagon Property, then X is also a strong parapolar space with the Pentagon Property.

18.4 Give an example of a strong parapolar space with the Pentagon Property in which axiom (Q) holds, while (PL*) fails. Show that there exist disconnected spaces whose connected components satisfy these hypotheses. [This shows that when (PL*) is relaxed, connectedness is also lost unless we artificially impose it with the "strong parapolar space hypothesis" (itself introduced only as a matter of convenience in stating theorems).]

18.5 Employ the preceding exercises to prove that the following geometry is a strong parapolar space with the Pentagon Property in which axiom (Q) holds:

$$\Gamma = DP/\bar{A} \times PG(n, K) \times E_{7,7}.$$

The problems raise an open question — which could reasonably be called the "$(\mathcal{D}, \mathcal{Q})$-factorization." Assuming the hypotheses requested in Exercise 18.4, is it always true that Γ is a product geometry $A \times B$ where A is either a singular subspace or is a strong parapolar space with the Pentagon Property in which the class $\mathcal{Q}_A$ of symplecta whose lines are maximal singular subspaces is empty, while B is either a polar space, or a strong parapolar space all of whose symplecta belong to the class $\mathcal{Q}$, that is, no point exterior to a member of this class is collinear with a line of the symplecton? In such a geometry, the class of grids $\mathcal{D}$ consists of symplecta of A, or products $L_A \times L_B$ where L_A is a line of A and L_B is a line of B.

Consider this example:

$$\Gamma = M \times L_1 \times \cdots \times L_n \times (DP/\bar{A}),$$

where M is a singular space properly containing a line, the L_i are lines, and the $DP/\bar{A}$ is one of the spaces emblazoned with that notation in Exercise 18.1. Only the first factor comprises A, and the factor B is the product $L_1 \times \cdots L_n \times DP/\bar{A}$. Such a factorization theorem would go a long way toward a classification theory. It is not silly for the only interconnecting symplecta, employing a line from A and a line from B, must be grids. Grids are what organize product geometries. But they need help.

18.6.2 Relaxing (Q): What if the Class $\mathcal{Q}$ is Empty?

In the next few exercises we consider the following hypothesis for a point-line geometry $\Gamma = (\mathcal{P}, \mathcal{L})$:

(D) (a) $\Gamma = (\mathcal{P}, \mathcal{L})$ is a strong parapolar space.
(b) Γ possesses the Pentagon Property (P).
(c) Every symplecton S is a grid.

One reason for considering this hypothesis is that hypothesis (D) must hold if we relax hypothesis (Q) in the Main Local Theorem on p. 646. In that context we inherit just one further property. That property is this:

(D4) If D is a symplecton, there exists some external point x such that $x^{\perp} \cap D$ is a line.

Now suppose I is some index set, and that for each index value σ, $P_\sigma = (\mathcal{P}_\sigma, \mathcal{L}_\sigma)$ is a linear space which properly contains a line. Then the product geometry

$$\Pi_{\sigma\in I} P_\sigma$$

can be described in this way. Its points are functions

$$\alpha : I \to \uplus_{\sigma\in I} P_\sigma,$$

from I to the disjoint union of the P_σ, for which $f(\sigma)$ is a point of the linear space P_σ. The lines are partitioned into classes $\mathcal{L}_\sigma$. The point-shadow of a line L in $\mathcal{L}_\sigma$ consists of those functions f possessing a constant specified value at all arguments $\tau \neq \sigma$, and ranging over the values $f(\sigma)$ constituting a line of the linear space P_σ.

The distance metric in the collinearity graph is essentially the Hamming distance between points – that is, the distance between point f and point g (each being functions with domain I) which is the number of values of σ at which the two functions differ. Now let

$$\phi : I \to \Pi_{\sigma\in I} \mathrm{Aut}(P_\sigma)$$

be a function which maps each index σ to an automorphism α_σ of the linear space P_σ.

18.6 Let ϕ be defined as in the previous paragraph. We define an action of ϕ on the point-set of $\Pi_{\sigma\in I} P_\sigma$ in this way. Given $\alpha : I \to \uplus_{\sigma\in I} P_\sigma$ we let α^ϕ be the function that maps $\sigma \in I$ to $\alpha(\sigma)^{\phi(\sigma)}$ – that is, every image value $\alpha(\sigma)$ has been distorted by the automorphism $\phi(\sigma) \in \mathrm{Aut}(P_\sigma)$.

1. Prove that ϕ is an automorphism of the geometry

$$X := \Pi_{\sigma\in I} P_\sigma.$$

   [Hint: One only has to show that ϕ takes lines to lines.]
2. Prove that the subgroup M generated by the action of the functions ϕ is isomorphic to the direct product of the automorphism groups $\mathrm{Aut}(M_\sigma)$.
3. Prove that M is the full group of automorphisms of the product geometry $\Pi_{\sigma\in I} P_\sigma$ which preserves each maximal singular subspace of the space.
4. Let FM be the set of mappings ϕ as above for which $\phi(\sigma)$ is a non-identity automorphism for only finitely many arguments σ. Show that FM is a subgroup of M and that M is isomorphic to the direct sum of the automorphism groups $\mathrm{Aut}(M_\sigma)$.

18.7 Prove that two points f and g of $\Pi_{\sigma\in I} P_\sigma$ have distance d in the point-collinearity graph if and only if the functions f and g differ at exactly d arguments in I.

Conclude that if I is infinite, then $\Pi_{\sigma\in I} P_\sigma$ is not a connected geometry.

18.8 Fix a point (or function) f of $\Pi_{\sigma\in I} P_\sigma$. Let Π_f be the collection of all functions g which differ from f at only finitely many arguments $\sigma \in I$.

1. Prove that Π_f is a connected component of $\Pi_{\sigma\in I} P_\sigma$.
2. In addition demonstrate the following: as a subspace, Π_f is a parapolar space with no 5-circuits – i.e., it is a strong parapolar space in which the Pentagon Property holds. In fact it satisfies hypothesis (D).
3. Prove that the group FM leaves the subspace Π_f invariant.
4. Fix a subset J of the index set I. Let Π_J, be the collection of all functions $g \in \Pi_{\sigma\in I} P_\sigma$ which agree with f outside the index set J. Then clearly Π_J is a subspace of $\Pi_{\sigma\in I} P_\sigma$. Similarly, let M_J be the set of automorphisms in M which induce the identity automorphism on each P_σ for which $\sigma \notin J$ – that is,

$$M_J := \{\phi \in M | \phi(\sigma) = \mathrm{id}(P_\sigma) \text{ for all } \sigma \in I - J\}.$$

   If J is a finite set, note that M_J is a subgroup of FM and so leaves each connected component Π_f of $\Pi_{\sigma\in I} P_\sigma$ invariant. Now suppose J is finite of cardinality at least five and that $B \subseteq \Pi_J$ is a subgroup of Π_J, all of whose non-identity elements induce a fixed-point-free action on each maximal singular subspace M_σ for $\sigma \in J$. Show that the factor space Π_f/B is also a strong parapolar space with the Pentagon Property – that is it satisfies hypothesis (D). [Hint: Note that any two distinct points belonging to a common B-orbit have distance at least five from one another.]
5. Is this fantasy? No it is not. Suppose each M_σ is a copy of the same $PG(n, K)$ where K is a field possessing an irreducible polynomial $p(x)$ of degree d greater than 1 which divides n.

   1. Show that $PG(n, K)$ admits a non-identity automorphism s which acts semiregularly on the points of $PG(n, K)$. [Hint: The underlying vector space for $PG(n, K)$ admits a linear transformation t preserving a direct decomposition of V into n/d direct summands and acting with minimal polynomial $p(x)$ on each summand.]
   2. Show that this means that for each choice of σ there exists an automorphism α_σ of P_σ which acts semiregularly on P_σ derived from a completely reducible transformation with minimal polynomial $p(x)$.
   3. Now, keeping J finite, let B be the cyclic subgroup of M_J generated by the element $\phi : I \to \sum_\sigma \mathrm{Aut}(P_\sigma)$ for which

$$\phi(\sigma) = \begin{cases} \alpha_\sigma & \text{if } \sigma \in J \\ \mathrm{id}(P_\sigma) & \text{if } \sigma \in I - J. \end{cases}$$

      Show that if $|J| > 5$, then Π_f/B has property (D) and if $n \geq 3$ it has property (D4).

18.9 Now we come to a connection with coding theory. Unfortunately the connection is descriptive, rather than analytic, to invoke an old distinction. In plain language, the coding ideas help us describe examples, but otherwise contribute nothing whatsoever toward the classification of these spaces.

Suppose $V = \sum_{\sigma \in I} V_\sigma$ is a direct sum of vector spaces over I. That means each vector $f \in V$ if a function

$$f : I \to \uplus_{\sigma \in I} V_\sigma$$

for which $f(\sigma)$ is a vector in V_σ which differs from the zero vector 0_σ of V_σ only finitely many times. The vectors of V are the points of our geometry. The lines (as usual for any product of point line geometries) are defined by those collections of functions which are constant at all arguments except some $\sigma \in I$, at which argument $f(\sigma)$ is allowed to range over all vectors of a coset of a 1-subspace of the vector space V_σ – that is, $f(\sigma)$ ranges over the points belonging to a line of the affine space A_σ obtained from vector space V_σ.

The *weight* of a vector f is defined to be the number of values $\sigma \in I$ for $f(\sigma) \neq 0_\sigma$, the zero-vector of the vector space V_σ. As in the ordinary coding theory, the function $d(u, v) = wt(u-v)$ defines the point-collinearity-metric in the point-collinearity graph. The only way this presentation differs from the classical presentation of an error-correcting code, is simply that in the latter, the V_σ are one-dimensional.[11]

Now we suppose each V_σ is a right vector space over the same division ring K. We define a *linear code* to be a subspace C – that is a right K-submodule module of V. We are interested in a possible subspace C in which each non-identity element of C has weight at least five.

Prove that the mapping $V \to V/C$ is an epimorphism of geometries preserving the strong parapolar space structure as well as the Pentagon Property. (Note that in this example, the maximal singular subspaces are affine spaces.)

18.10 This is really an open question rather than an exercise. In the preceding exercises of this section, the reader has been introduced to a number of strong parapolar spaces of symplectic rank two satisfying the hypotheses (D) and (D4) above. All of them are homomorphic images of a connected component of a product $\Pi_\sigma(P_\sigma)$ of linear spaces. This raises a question. *Must any strong parapolar space satisfying (D) and (D4) be a homomorphic image of such a space?*.

Suppose Γ is a strong parapolar space satisfying the hypotheses (D) and (D4) above. Show that Γ is already a connected component of a product of

[11] The distinction is not at all important at an elementary level. Products as above occur when one considers a product of error-correcting codes or a field-reduction of such codes. There is nothing novel in these definitions.

linear spaces if either (1) all maximal singular spubspaces of Γ are gated, or if (2) all symplecta (which are necessarily grids) are gated in Γ. [Hint: Use either of the properties and induction on distance to define the appropriate equivalence classes on maximal singular spaces needed to define the product.]

18.11 Suppose $\Gamma = (\mathcal{P}, \mathcal{L})$ is such a strong parapolar space and let $\Delta = (\mathcal{P}, \sim)$ be its point-collinearity graph. Now let $\mathcal{C}$ denote the class of pointed circuits of Δ of lengths three and four. Then we can form a universal $\mathcal{C}$-covering

$$\kappa : \hat{\Delta} \to \Delta,$$

as described fully in the very first chapter of this book. Let $\hat{\mathcal{P}}$ denote the vertex set of $\hat{\Delta}$.

In Δ, the point-shadows of the lines in $\mathcal{L}$ form a special collection of cliques which, with a slight abuse of notation, we also denote by $\mathcal{L}$. Then the cliques of $\mathcal{L}$ lift to a collection $\hat{\mathcal{L}}$ of cliques in the covering graph $\hat{\Delta}$. Show that with respect to the natural incidence, $\hat{\Gamma} = (\hat{\mathcal{P}}, \hat{\mathcal{L}})$ is also a strong parapolar space with properties (D) and (D4).

References

1. Peter Abramenko and Hendrik Van Maldeghem. Characterizations of generalized polygons and opposition in rank 2 buildings. *J. Geom.*, 74:7–28, 2002.
2. R. W. Ahrens and G. Szekeres. On a combinatorial generalization of 27 lines associated with a cubic surface. *J. Austral. Math. Soc.*, 10:485–492, 1969.
3. Martin Aigner. *Combinatorial Theory*. Springer, Berlin, 1997.
4. Michael Aschbacher and Yoav Segev. Extending morphisms of groups and graphs. *Ann. Math.*, 10:297–323, 1992.
5. E. F. Assmus and H. F. Mattson. New 5-designs. *J. Combin. Theory, Ser. A*, 6:122–151, 1969.
6. A. Barlotti. Un' estensione del teorema di Segre-Kustaanheimo. *Boll. Un. Mat. Ital.*, 10: 96–98, 1955.
7. Albrecht Beutelspacher. Projective planes. In F. Buekenhout, editor, *Handbook of Incidence Geometry*, Chapter 4, pages 107–136. North Holland, Amsterdam, 1995.
8. A. Bichara and G. Tallini. On a characterization of Grassmann space representing the h-dimensional subspaces in a projective space. *Ann. Discrete Math.*, 18:113–131, 1983.
9. N. Bourbaki. *Groupes et Algebras de Lie, ch III–IV*. Hermann, Paris, 1968.
10. A. E. Brouwer and A. M. Cohen. Local recognition of Tits' geometries of classical type. *Geom. Dedicata*, 20:181–199, 1986.
11. A. E. Brouwer, A. M. Cohen, J. I. Hall, and H. A. Wilbrink. Near polygons and Fischer spaces. *Geom. Dedicata*, 49:349–368, 1994.
12. A. E. Brouwer, A. M. Cohen, and A. Neumaier. *Distance regular graphs; Ergebnisse der Mathematik*, vol. 18. Springer, Berlin, 1989.
13. A. E. Brouwer and H. A. Wilbrink. The structure of near polygons with quads. *Geom. Dedicata*, 14:145–146, 1983.
14. Kenneth S. Brown. *Buildings*. Springer, New York, 1989.
15. M. R. Brown. Ovoids of $PG(3, q)$, q even, with a conic section. *J. Lond. Math. Soc.*, 62:569–582, 2000.
16. R. Bruck and Erwin Kleinfeld. The structure of alternative division rings. *Proc. Nat. Acad. Sci. USA.*, 37:89–90, 1951.
17. R. H. Bruck and H. J. Ryser. The non-existence of certain finite projective planes. *Can. J. Math.*, 1:88–93, 1949.
18. Francis Buekenhout. Une caractérisation des espaces affines basée sur las notion de droite. *Math. Z.*, 11:367–371, 1969.
19. Francis Buekenhout. A characterization of polar spaces. *Simon Stevin*, 53:3–7, 1979.
20. Francis Buekenhout. The basic diagram of a geometry. In Martin Aigner and Dieter Jungnickel, editors, *Geometries and Groups*, pages 1–29. Springer, Berlin, 1981.
21. Francis Buekenhout. Cooperstein's theory. *Simon Stevin*, 57:125–140, 1982.
22. Francis Buekenhout and Arjeh Cohen. *Diagram Geometries*. Springer, Berlin, To appear.
23. Francis Buekenhout and Ernest Shult. On the foundations of polar geometry. *Geom. Dedicata*, 3:155–170, 1974.

E.E. Shult, *Points and Lines*, Universitext, DOI 10.1007/978-3-642-15627-4,

© Springer-Verlag Berlin Heidelberg 2011

24. Peter Cameron. Dual polar spaces. *Geom. Dedicata*, 12:75–85, 1982.
25. Peter Cameron. *Projective and Polar Spaces, School of Mathematical Sciences, Queen Mary and Westfield College (University of London)*. QMW Maths Notes, London, 1993.
26. Peter Cameron. *Combinatorics: Topics, Techniques, Algorithms*. Cambridge University Press, Cambridge, 1994.
27. Roger Carter. *Simple Groups of Lie Type*. Wiley, New York, 1972.
28. Rey Casse, *Projective Geometry: An Introduction*. Oxford University Press, Oxford, 2006.
29. Chris Charnes and Ulrich Dempwolf. Spreads, ovoids and S_5. *Geom Dedicata*, 56:129–143, 1995.
30. Chris Charnes and Ulrich Dempwolf. The eight-dimensional ovoids over GF(5). *Math. Comp.*, 70:853–861, 2001.
31. W. Cherwitzo. Hyperovals in Desarguesian planes of even order. *Ann. Discrete Math.*, 37:87–94, 1986.
32. W. Cherwitzo, T. Pentilla, I. Pinneri, and G. Royle. Flocks and ovals. *Geom. Dedicata*, 60:193–194, 1996.
33. Arjeh Cohen. An axiom system for metasymplectic spaces. *Geom. Dedicata*, 12:417–433, 1982.
34. Arjeh Cohen. On a theorem of Cooperstein. *Euro. J. Combin.*, 4:107–126, 1983.
35. Arjeh Cohen. Point-line spaces related to buildings. In F. Buekenhout, editor, *Handbook of Incidence Geometry*, Chapter 12, pages 647–737. North-Holland, Amsterdam, 1995.
36. Arjeh Cohen and Bruce Cooperstein. A characterization of some geometries of Lie type. *Geom. Dedicata*, 15:73–105, 1983.
37. Arjeh M. Cohen and Gabor Ivanyos. Root shadow spaces. *Euro. J. Combin.*, 28:1419–1441, 2007.
38. Arheh Cohen and Ernest Shult. Affine polar spaces. *Geom. Dedicata*, 35:43–76, 1990.
39. J. H. Conway. Three lectures on exceptional groups. In M. B. Powell and G. Higman, editors, *Finite Simple Groups*, pages 215–247. Academic Press Inc., New York, 1971.
40. John Conway, Peter Kleidman, and Rob Wilson. New families of ovoids in $O^+(8, q)$. *Geom. Dedicata*, 26:157–170, 1988.
41. Bruce Cooperstein. A characterization of some Lie incidence structures. *Geom. Dedicata*, 6:205–258, 1977.
42. Bruce Cooperstein. A sporadic ovoid in $\Omega^+(8, 5)$ and some non-Desarguesian translation planes of order 5. *J. Combin. Theory, Ser. A*, 54:135–140, 1990.
43. H. M. S. Coxeter. Discrete groups generated by reflections. *Ann. Math.*, 35:588–621, 1934.
44. C. W. Curtis. Chevalley groups and related topics. In M. B. Powell and G. Higman, editors, *Finite Simple Groups*, pages 135–189. Academic Press Inc., New York, 1971.
45. Hans Cuypers. Affine Grassmannians. *J. Combin Theory, Ser. A*, 70:289–304, 1995.
46. Hans Cuypers, Peter Johnson, and Antonio Pasini. On the embeddability of polar spaces. *Geom. Dedicata*, 44:349–358, 1992.
47. Hans Cuypers, Peter Johnson, and Antonoi Pasini. On the classification of polar spaces. *J. Geom.*, 48:56–62, 1993.
48. Hans Cuypers and Anja Steinbach. Linear transvection groups and embedded polar spaces. *Invent. Math.*, 137:169–198, 1999.
49. Hans Cuypers and Anja Steinbach. Special linear groups generated by transvections and embedded projective spaces. *J. Lond. Math. Soc.*, 64:576–594, 2001.
50. I. DeBroey. Semipartial geometries satisfying the diagonal axiom. *J. Geom.*, 13:171–190, 1979.
51. B. DeBruyn. Generalized quadrangles with a spread of symmetry. *Euro. J. Combin.*, 20: 759–771, 1999.
52. B. DeBruyn. *Constructions and Characterizations of Near Polygons*. Ph.D. Thesis, Universiteit Ghent, Ghent, 2000.
53. B. DeBruyn and K. Thas. Generalized quadrangles with a spread of symmetry and near polygons. *Ill. J. Math.*, 46:797–818, 2002.

54. Frank DeClerck and Jef Thas. Partial geometries in finite projective spaces. *Arch. Math. (Basel)*, 30:537–540, 1978.
55. Frank DeClerck and Jef Thas. The embedding of $(0, \alpha)$-geometries in $PG(n, q)$, Part I. *Ann. Discrete Math.*, 18:229–240, 1983.
56. Frank DeClerck and Hendrick Van Maldeghem. Some classes of rank 2 geometries. In F. Buekenhout, editor, *Handbook of Incidence Geometry*, Chapter 10, pages 647–737. North Holland, Amsterdam, 1995.
57. Peter Dembowski. *Finite Geometries*. Springer, Berlin, 1968.
58. Keith Devlin. *The Joy of Sets*. Springer, Berlin, 1979.
59. A. W. M. Dress and R. Scharlau. Gated sets in metric spaces. *Aequationes Math.*, 34: 112–120, 1987.
60. M. El-Atrash and E. Shult. A characterization of certain families of strong parapolar spaces. *Geom. Dedicata*, 71:221–235, 1998.
61. Cecil A. Ellard and Ernest Shult. A characterization of polar Grassmann spaces. Preprint, Kansas State University, 1988.
62. Walter Feit and Graham Higman. The nonexistence of certain generalized polygons. *J. Algebra*, 1:114–131, 1964.
63. Paul Garrett. *Buildings and Classical Groups*. Chapman and Hall, London, 1997.
64. Jon Hall. Locally Petersen graphs. *J. Graph Theory*, 4:173–187, 1980.
65. Jon Hall and Ernest Shult. Locally cotriangular graphs. *Geom. Dedicata*, 18:113–187, 1980.
66. Marshall Hall Jr. *The Theory of Groups*. Macmillan, New York, 1959.
67. Marshall Hall Jr. Affine generalized quadrilaterals. In L. Mirsky, editor, *Studies in Pure Mathematics*, pages 113–116. Academic Press, New York, 1971.
68. Guy Hanssens. A characterization of buildings of spherical type. *Euro. J. Combin.*, 7: 333–347, 1986.
69. Guy Hanssens. A technique in the classification theory of point-line geometries. *Geom. Dedicata*, 24:85–101, 1987.
70. Guy Hanssens. A characterization of point-line geometries for finite buildings. *Geom. Dedicata*, 25:297–315, 1988.
71. Guy Hanssens. Polarized spaces of polar rank three which are locally a direct product. *Euro. J. Combin.*, 12:23–31, 1991.
72. Guy Hanssens and J. A. Thas. Pseudopolar spaces of polar rank three. *Geom. Dedicata*, 22:117–135, 1987.
73. D. G. Higman. Invariant relations, coherent configurations and generalized polygons. *Combinatorics, Part 3, Math. Centre Tract.*, 57:27–43, 1974.
74. D. G. Higman. Coherent configurations, I. *Geom. Dedicata*, 4:1–32, 1975.
75. Don Higman and Jack McCaughlin. Geometric ABA-groups. *Ill. J. Math.*, 5:382–397, 1961.
76. Dan Hughes and Fred Piper. *Projective Planes*. Springer, Berlin, 1973.
77. Bertram Huppert. *Endlichen Gruppen*. Springer, New York, 1967.
78. Bertram Huppert and Norman Blackburn. *Finite Groups III*. Springer, New York, 1982.
79. Nathan Jacobson. *Basic Algebra II*. Freeman and Co., San Francisco, 1980.
80. Peter Johnson. Polar spaces of arbitrary rank. *Geom. Dedicata*, 35:229–250, 1990.
81. Peter Johnson and Ernest Shult. Local characterizations of polar spaces. *Geom. Dedicata*, 28:127–151, 1988.
82. William Kantor. Generalized quadrangles associated with $G_2(q)$. *J. Combin. Theory, Ser. A*, 29:212–219, 1980.
83. William Kantor. Ovoids and translation planes. *Can. J. Math.*, 34:1195–1207, 1982.
84. Anna Kasikova and Ernest Shult. Chamber systems which are not geometric. *Comm. Algebra*, 24:3471–3481, 1996.
85. Anna Kasikova and Ernest Shult. Point-line characterizations of Lie geometries. *Adv. Geom.*, 2:147–188, 2002.
86. Geoff Mason and Ernest Shult. The Klein correspondence and the ubiquity of certain translation planes. *Geom. Dedicata*, 31:29–50, 1986.

87. Eric Moorhouse. Ovoids from the E_8 root lattice. *Geom. Dedicata*, 46:287–297, 1993.
88. Eric Moorhouse. Ovoids and translation planes from lattices. In N. Johnson, editor, *Mostly Finite Geometries*, pages 123–134. Dekker, New York, 1997.
89. B. Muehlherr. A geometric approach to non-embeddable polar spaces of rank 3. *Bull. Soc. Math. Belg. Ser. A*, 42:577–594, 1990.
90. B. Muehlherr. A rank-2 characterization of twinnings. *Euro. J. Combin.*, 19:603–612, 1998.
91. Bernhard Muehlherr and Hendrik Van Maldegham. Exceptional Moufang quadrangles of type $F4$. *Can. J. Math.*, 5:347–371, 1999.
92. C. M. O'Keefe and T. Pentilla. Ovoids with a pencil of translation ovals. *Geom. Dedicata*, 62:19–34, 1996.
93. C. M. O'Keefe and T. Pentilla. Ovals in translation hyperovals and ovoids. *Euro. J. Combin.*, 18:667–683, 1997.
94. J. G. Oxley. *Matroid Theory*. Oxford University Press, New York, 1976.
95. Antonio Pasini. *Diagram Geometries*. Oxford Science Publications, Clarendon Press, Oxford, 1994.
96. Stanley Payne. Nonisomorphic generalized quadrangles. *J. Algebra*, 18:201–212, 1971.
97. Stanley Payne. A garden of generalized quadrangles. *Algebras, Groups, Geom.*, 3:323–354, 1985.
98. Stanley Payne and Jef Thas. *Finite Generalized Quadrangles*. Pitman, Boston, 1984.
99. Nicholas Percsy. On the Buekenhout–Shult theorem on polar spaces. *Bull. Belg. Math. Soc. Ser. B*, 41:283–294, 1989.
100. G. Pickert. *Projectiven Ebenen*. Springer, Berlin, 1975.
101. Mark Ronan. Extending locally truncated buildings and chamber systems. *Proc. London Math. Soc.*, 53:385–406, 1986.
102. Mark Ronan. Embeddings and hyperplanes of discrete geometries. *Euro. J. Combin.*, 8: 179–185, 1987.
103. Mark Ronan. *Lectures on Buildings*. Perspectives in Mathematics, vol. 7. Academic Press, New York, 1989.
104. M. A. Ronan and J. Tits. Building buildings. *Math. Ann.*, 278:291–306, 1987.
105. R. Scharlau. A characterization of Tits buildings by metrical properties. *J. London Math. Soc.*, 32:317–327, 1985.
106. R. Scharlau. Buildings. In F. Buekenhout, editor, *Handbook for Incidence Geometry*, Chapter 11, pages 477–645. North-Holland (Elsevier), Amsterdam, 1995.
107. J.-P. Serre. *A course in Arithmetic*. Graduate Texts in Mathematics no. 7. Springer, Berlin, 1973.
108. Ernest Shult. Characterizations of the Lie incidence geometries. In E. Keith Lloyd, editor, *Surveys in Combinatorics*, London Mathematical Society Lecture Series, vol. 82, pages 157–186. Cambridge University Press, Cambridge, 1983.
109. Ernest Shult. A sporadic ovoid in $\Omega^+(8, 7)$. *Algebras, Groups Geom.*, 2:495–513, 1985.
110. Ernest Shult. Characterizations of spaces related to metasymplectic spaces. *Geom. Dedicata*, 30:325–371, 1991.
111. Ernest Shult. A remark on Grassmann and half-spin geometries. *Euro. J. Combin.*, 15:47–52. 1994.
112. Ernest Shult. Rank three geometries with simplicial residues. In S. Magliveras, editor, *Geometries, Combinatorial Designs and Related Structures (Spetses, 1996)*, London Mathematical Society Lecture Note Series, No. 245, pages 217–239. Cambridge University Press, Cambridge, 1997.
113. Ernest Shult. On Veldkamp lines. *Bull. Belg. Math. Soc. Simon Stevin*, 4:299–316, 1997.
114. E. E. Shult. Aspects of buildings. In Lino di Martino et al., editors, *Groups and Geometries (Siena 1996)*, pages 177–188. Birkhauser, Basel, 1998.
115. Ernest Shult. Characterizing the half-spin geometries by a class of singular subspaces. *Simon Stevin*, 12:883–894, 2005.
116. Ernest Shult. On characterizing the long-root geometries. *Adv. Geom.*, 10:353–370, 2010.

117. Ernest Shult and Jef Thas. Constructions of polygons from buildings. *Proc. London Math. Soc.*, 71:397–440, 1995.
118. E. E. Shult and K. Thas. A question of Cohen on parapolar spaces. *Combinatorica*, to appear.
119. Ernest Shult and Arthur Yanushka. Near n-gons and line systems. *Geom. Dedicata*, 9:1–72, 1980.
120. Ron Solomon. On finite simple groups and their classification. *Notices Amer. Math. Soc.*, 42:231–239, 1995.
121. Alan Sprague. Pasch's axiom and projective spaces. *Discrete Math.*, 33:79–87, 1981.
122. Alan Sprague. Dual linear, linear diagrams. In N. Johnson and M. Kallaher, editors, *A Conference in Honor of T. Ostrom*, pages 413–418. Dekker, New York, 1983.
123. Alan Sprague. Rank 3 incidence structures admitting a dual-linear, linear diagram. *J. Combin. Theory, Ser. A.*, 38:254–259, 1985.
124. Michio Suzuki. Contributions to the theory of finite groups. In Marshall Hall Jr., editor, *1960 Institute on Finite Groups; Proceedings of Symposia in Pure Mathematics*, vol. IV, pages 107–109. American Mathematical Society, Providence, 1962.
125. Michio Suzuki. *Group Theory I*. Springer, New York, 1982.
126. Michio Suzuki. *Group Theory II*. Springer, New York, 1986.
127. G. Tallini. On a characterization of the Grassmann manifold representing the lines of a projective space. In P. J. Cameron and J. W. P. Hirschfeld, editors, *Finite Geometries and Designs: Proceedings of the Second Isle of Thorns Conference, 1980*. London Mathematical Society Lecture Notes 49, pages 354–358. Cambridge University Press, Cambridge, 1981.
128. Luc Teirlinck. Planes and hyperplanes of 2-coverings. *Bull. Math. Soc. Belg.*, 29:73–81, 1977.
129. K. Tent and H. Van Maldeghem. Split BN-pairs and Moufang polygons, I. *Adv. Math.*, 174:254–265, 2003.
130. Jef Thas. Generalized polygons. In F. Buekenhout, editor, *Handbook for Incidence Geometry*, Chapter 9, pages 383–431. North Holland, Amsterdam, 1995.
131. Jef Thas, Stanley Payne, and Hendrik Van Maldeghem. Half Moufang implies Moufang for generalized quadrangles. *Invent. Math.*, 105:153–156, 1991.
132. J. A. Thas, F. Thas, and H. Van Maldeghem. *Translation Generalized Quadrangles*. World Scientific, London, Singapore, 2006.
133. J. A. Thas and H. Van Maldeghem. Flat lax and weak lax embeddings of finite generalized hexagons. *Euro. J. Combin.*,19:733–751, 1998.
134. K. Thas. *Automorphisms and Combinatorics of Finite Generalized Quadrangles*. Ph. D. Thesis, University of Ghent, 2002.
135. F. G. Timmesfeld. Groups generated by k-transvections. *Invent. Math.*, 100:167–206, 1990.
136. F. G. Timmesfeld. *Abstract Root Subgroups and Simple Groups of Lie Type*. Birkhauser, Basel, 2001.
137. J. Tits. Ovoides et groups de Suzuki. *Arch. Math.*, 13:187–198, 1962.
138. J. Tits. *Buildings of Spherical Type and Finite BN-Pairs*, Lecture Notes in Mathematics, vol. 386. Springer, Berlin, 1974.
139. J. Tits. A local approach to buildings. In C. Davis, B. Grunbaum, and F. A. Sherk, editors, *The Geometric Vein (the Coxeter Festschrift)*, pages 519–547. Springer, Berlin, 1981.
140. J. Tits and R. M. Weiss. *Moufang Polygons*. Springer, Berlin, 2002.
141. Hendrik Van Maldeghem. *Generalized Polygons*. Birkhauser, Basel, 1998.
142. O. Veblen and J. Young. *Projective Geometry*. Ginn, Boston, 1916.
143. F. D. Veldkamp, Polar geometry, I–IV. *Indag. Math.*, 21:512–551, 1959.
144. G. M. Weetman. A construction of locally homogeneous graphs. *J. London Math. Soc.*, 50:68–86, 1994.
145. Richard M. Weiss. *The Structure of Spherical Buildings*. Princeton University Press, Princeton, 2003.
146. J. A. Welsh. *Matroid Theory*. Academic Press, London, 1976.
147. A. Yanushka. On order in generalized polygons. *Geom. Dedicata*, 10:451–458, 1981.

148. F. Zara. Une caractërisation des graphes associés aux groupes de Fischer. *Euro. J. Combin.*, 7:187–190, 1986.
149. Hans Zassenhaus. Und endlicher Fastkörper. *Abh. Math. Sem. Hamburg*, 11:187–220. 1935.
150. Paul-Hermann Ziechang. *An Algebraic Approach to Association Schemes*. Springer, Berlin, 1996.
151. M. Zorn. Theorie der alternativen Ringe. *Abh. Math. Sem. Hamburg*, 8:123–147, 1931.

Index

A
Admissable triple, 597
Affine plane, 109
Affine plane, Desarguesian, 110
Affine space, $AG(n, D)$, 145
Affine space, $A_H(P)$, 144
Affine space, $A(V)$, 144
Affine space, definition, 143
Affine translation plane, 114
Alternating forms, 203
Angle, for parapolar space lines, 618
Angle type, 618
Antiautomorphism, of a division ring, 194
Apartments, system of, 387
Automorphism
  type-shifting, 589
Axiom, diagonal, 507, 512
Axiom, dual Pasch, 513
Axiom, dual Veblen, 513
Axiom, hexagonic, (H1), 480
Axiom, hexagonic, (H2), 480
Axiom, hexagonic, (H3), 480
Axiom, Pasch, 442, 512, 513
Axiom, pentagon, 628
Axiom, Veblen, 513

B
Backtrack, 4
(B, N)-pair, 603
Borel subgroup, 387
Building, condition (Q_c), 403
Building, definition, 377
Building, homomorphic image of, characterization, 403
Building, residual connectedness of, 402

C
Chamber-connected geometry, strong, 428
Chamber connectedness
  of a geometry, 54
Chamber, flag chamber of a geometry, 45
Chamber system, automorphism group, 298
Chamber system, coset, 303
Chamber system, definition, 294
Chamber system, induced subsystem, 298
Chamber system, morphism, 298
Chamber system, morphism, full, 298
Chamber system, panels of, 301
Chamber system, residual connectedness, 304
Chamber system, residue of type J, 300
Chamber system, residue, corank of, 300
Chamber system, residue, cotype of, 300
Chamber system, residue, rank of, 300
Chamber systems, as a category, 300
Chamber system, truncation of, 302
Chevalley group, 546
Circuit, 17
Circuit, $\mathcal{C}$-contractible, 33
Circuit, association class of, 17
Circuit, backtrack, 17
Circuits, $\mathcal{C}$-contractible, 22
Circuit, oriented association class of, 17
Circuit, pointed, 17
Classical polar space, of a reflexive form, 74
Clique, 6
Closure, induced-convex, 10
Closure operator, 139
Coclique, 6
Cohen-Coopersein geometries, 545
Collinearity, 61
Collineation, of a plane, perspective, 121
Condition, $(CC)_k$, 543
Condition, (Coo), 456
Condition, (G_c), 377
Condition, (G_x), 377
Condition, (H), 648
Condition, (H1), 480
Condition, (H3), 480

Condition, (MC), 273
Condition, (NG), 630
Condition, (NP), 252
Condition, (P_c), 377
Condition, (P_c-min), 377
Condition, (P_x), 377
Condition, (PL*), 646
Condition, (PL), 645, 646
Condition, (R_c), 403
Condition, $(U)_k$, 542
Condition (RG^1), 310
Condition (RG_2), 310
Condition (RG), 310
Condition on a parapolar space, $(CC)_k$, 528
Condition on a parapolar space, (D2), 478, 479
Condition on a parapolar space, (H), 647, 649
Condition on a parapolar space, (H2), 478
Condition on a parapolar space, (NG), 478, 479, 630
Condition on a parapolar space, (P), 628
Condition on a parapolar space, (P2), 478–480
Condition on a parapolar space, pentagon property, 628, 630, 647
Condition on a parapolar space, (PL*), 632, 645, 647
Condition on a parapolar space, (PL), 632
Condition on a parapolar space, (Q*), 650
Condition on a parapolar space, $(U)_2$, 529
Conformal space, 234
Convex closure, 456
Copolar space, 283
Covering morphism, $\mathcal{C}$-cover, 18, 663
Covering morphism, of graphs, 18
Covering morphism, universal $\mathcal{C}$-cover, 18
Covering, of a point-line geometry, 62
Covering, of a point-line geometry, universal, 62
Covering, of graphs, 18
Coxeter group, 351
Coxeter group, length of a word, 359
Coxeter group, reduced word, 359
Coxeter group, reflection, 363
Coxeter matrix, 350
Coxeter system, 351

D

Deck transformation, 37
Deep points, 650
Deep point of a subspace, 471
Degree, of a vertex, 8
Dependance relation, 138
Diagonal axiom, 507, 512
Diagrams, for geometries, 347
Difference set, 109
Dimension, of a projective space, 141
Division algebra, Cayley-Dickson, 232
Division ring, opposite, 193
Dual, of a projective space, 193

E

Elation, 121
External lines, 212

F

Fano plane, 51
Fibering, 16
Fibering, lift of a walk, 17
Fibre, of a graph morphism, 15
Flag, of a geometry, 45
Flat, in a matroid, 139
Form, symmetric, 203
Form, symplectic, 203

G

Gallery, 294, 324
Gallery, J-factorization of, 406
Gallery, J-reduced segment, M-homotopy class of, 406
Gallery, initial J-reduced segment, 406
Gallery, panel-reduced, 324
Gallery, terminal J-segment , 406
Gallery, type of, 294, 324
Gamma space, locally connected, 68
Gamma space, nexus-connected, 489
Gamma space, paraprojective, k-connected, 487
Gamma space, tangent graph of, 68
Gatedness, between vertex and subgraph, 12
Generalized n-gon, as a chamber system, 326
Generalized n-gon, as a geometry, definition, 334
Generalized n-gon, n even, 253
Generalized parapolar space, of nexus 1, 486
Generalized parapolar space, of nexus k, 488
Generalized polygon, Abramenko-Van Maldeghem characterization, 339
Generalized polygon, as a chamber system, 327
Generalized polygon, generalized 1-gon, 342
Generalized polygon, order of, 343
Generalized polygon, Scharlau's characterization, 331
Generalized projective plane, 87
Generalized projective space, 88, 137
Generalized quadrangle, 171
Generalized quadrangle, non-degenerate, 171
Generalized quadrangle, opposite lines in, 171

Generalized quadrangle, of order (s, t), 172
Geometric hyperplane, 80
Geometric rigidity, 420
Geometry, Cohen-Cooperstein, 545
Geometry, connected, 44
Geometry, definition, 43
Geometry, firm, 60
Geometry, flag of, 45
Geometry, flag chamber, 45
Geometry, J-rigid, 420
Geometry, incidence relation in, 44
Geometry, induced subgeometry, 44
Geometry, locally truncated, 415
Geometry, long root, 605
Geometry, point-line, 60
Geometry, pointline, universal cover of, 62
Geometry, product of, 71, 660
Geometry, rank of, 44
Geometry, residue of a flag, 46
Geometry, shadow model of, 419
Geometry, strongly chamber connected, 428
Geometry, truncation of, 44
Geometry, type function, 43
Geometry, typeset, 43
Graph, automorphism, 15
Graph, automorphism group of, 8
Graph, ball of radius i, 5, 262
Graph, bipartite, 7
Graph, complement of, 7
Graph, complete bipartite, 8
Graph, complete multipartite, 8
Graph, connected, 4
Graph, connected components of, 4
Graph, $\mathcal{C}$-simply connected, 34
Graph, degree of a vertex, 8
Graph, diameter of, 5
Graph, distance in, 5
Graph, distance regular, 9
Graph, external diameter of a set of vertices, 261
Graph, homomorphism of, 15
Graph, image of morphism of, 15
Graph, of an incidence relation, 61
Graph, internal diameter of a set of vertices, 261
Graph, isomorphism of two, 15
Graph, locally X, 8
Graph, morphism, 15
Graph, morphism, $\mathcal{C}$-cover, 18
Graph, morphism, covering, 18
Graph morphism, deck transformation, 37
Graph, morphism, fiber of, 15
Graph, morphism, fibering, 16
Graph, morphism, full, 15
Graph, morphism, locally bijective, 16
Graph morphism, universal $\mathcal{C}$-cover, 18
Graph, multipartite, 7
Graph, point-collinearity, 61
Graph, pointed circuit of, 17
Graph, radius $\Delta_i(x)$, 5
Graph, regular, 8
Graphs, fibering, lifts of walks, 17
Graphs, hypercube, 11
Graph, simple, 3
Graph, simple, finite, 8
Graph, simply $\mathcal{C}$-connected, 19
Graph, strongly regular, 9
Graph, subgraph, 6
Graph, subgraph, clique, 6
Graph, subgraph, coclique, 6
Graph, subgraph, convex, 9
Graph, subgraph, convex closure of, 10
Graph, subgraph, gate, 12
Graph, subgraph, gated, 12
Graph, subgraph, induced, 6
Graph, subgraph, induced closure of, 6
Graph, subgraph, induced-convex closure of, 10
Graph, subgraph, intersection of, 6
Graph, subgraph, isometrically embedded, 11
Graph, subgraph, strongly gated, 12, 13
Graph, subgraph, vertex gated to, 12
Graph, subgraph, vertex strongly gated to, 12
Graph, walk, 4, 5
Graph, walk
initial vertex of, 4
length of, 4
Graph, walk, circular, 5
Graph, walk, concatenation of two, 4
Graph, walk, geodesic, 5
Graph, walk, inverse of, 4
Graph, walk, path, 5
Graph, walk, simple backtrack, 5
Graph, walk, subsegment of, 4
Graph, walk, terminal vertex of, 4
Graph, Zara, 258
Graphs, distance-transitive, 9
Grassman space, characterization, 498
Grid, 631

H

Half-Grassmann space, 506
Hanssens' Principal, 473
Hanssen's principle, 654
(σ, ϵ)-Hermitian form, 202
(σ, ϵ)-Hermitian form, alternating or symplectic, 203

Hermitian form, σ-Hermitan form, 203
(σ, ϵ)-Hermitian form, symmetric, 203
Hexagonic geometry, 606
Homotopy, Λ-, 317
Homotopy, $\mathcal{C}$-, 20
Homotopy, Λ-, elementary, 317
Homotopy, elementary $\mathcal{C}$-, 20
Homotopy, $Lambda(R)$-, local, 317
Homotopy, M-, 361
Homotopy, panel, 325

I
Image graph, 15
Incidence, in a geometry, 44
Incidence graph, 61

K
Klein correspondence, 119
Klein quadric, 119

L
Ladder lemma, 456
Lax embeddings, 66
Lifts, under a fibering, 17
Lines, pencil of, 512
Line spread, 530
Lines, transverse line, 512
Lines, transverse to a pencil, 512
Locally connected, 466
Locally connected components, 69, 467
Locally connected, gamma space, 68
Locally truncated geometry, sheaf over, 417
Locally truncated geometry, type D over I, 415
Long root geometry, 603, 605
Long root groups, 603

M
Matroid, basis of, 139
Matroid, definition, 138
Matroid, flats of, 139
Matroid, rank of, 140
Metasympletic space, 603
M-homotopy, 361
Morphism, fibering, of a point-line geometry, 62
Morphism, of graphs, 15
Morphism, of a point-line geometry, 61
Morphism, of projective space, index of, 160
Morphism, of projective space, kernel of, 160
Moufang condition, for a projective plane, 232

N
Near n-gon, 253
Near polygon, definition of, 252
Near polygons, d-opposite lines, 253
Near polygons, quads in, 263
Near polygons, with quads, 266

O
Opposite division ring, 193
Opposite lines, in a polar space, 173
Opposite roots, 397
Orthogonal group, 52
Ovoid, of a polar space, 549
Ovoid, of a polar space, of type $O^+(6, F)$, 120

P
Panels, 301
Parabolic subgroup, 388
Parallel classes, in affine planes, 110
Parapolar space, 463
Parapolar space, strong, 463
Parapolar space, symplectic rank of, 463
Paraprojective space, 441, 466
Paraprojective space, k-connected, 487
Pencil of lines, 512
Perspective collineation, 121
Perspective collineation, axis of, 121
Perspective collineation, center of, 121
Perspective collineation, elation, 121
Perspective collineation, homology, 121
Plane, affine translation, 114
Plane, general definition, 608
Planes, projective, 608
Point-collinearity graph, 61
Point-line geometry, 60
Point-line geometry, point residue, 466
Point-residual, of a parapolar space, 463
Point-shadow, 60
Point-shadow, of a line, 111
Pointed circuit, 17
Pointed circuit, in a graph, 17
Polar family, 456
Polar geometry, of a relfexive form, 74
Polar Grassmannian, 603, 627
Polar Grassmannian, of lines, 443
Polar pair, 463
Polar space, definition of, 173
Polar space, embedded, defective point of, 213
Polar space, oriflame, 179
Polar space, radical of, 173
Polar spaces, finite polar rank, 177
Polar spaces, infinite polar rank, 177
Product geometry, 71, 630, 660
Projective dimension, 141

Projective plane, 106
Projective plane, finite, 107
Projective plane, finite, order of, 107
Projective plane, generalized, 87
Projective planes, 87
Projective planes, (p, L)-transitive, 124
Projective space, 136
  genereralized
    definition, 88
Projective space, classical or Desarguesian, 73
Projective space, definition, 88
Projective space, dimension of, 141
Projective space, dual of, 193
Projective space, rank of, 141
Pseudoquadratic form, definition of, 208
Pseudoquadratic form, σ-quadratic, 208
Pseudoquadratic form, quadratic, 208
Pure subspace, 162
Pure vectors, in a wedge product, 118

Q
σ-quadratic form, 208
(σ, ϵ)-quadratic form, 208
Quadratic form, 208
Quads, 263
Quasi-polarity, 195
Quasi-polarity, absolute subspaces of, 196
Quasi-polarity, represented by a σ-sesquilinear form, 198
Quasi-polarity, trace-valued, 197

R
Rank, of a chamber system, 294
Rank, of a geometry, 44
Rank, of a matroid, 140
Rank, of a projective space, 141
Reflection, 363
Residual connectedness
  for a geometry, 54
  of a sheaf, 429
Residual connectedness, of a chamber system, 304
Residue, of a chamber system, 300
Root filtration space, 626

S
Secant line
  of a set of points, 63
Secant lines, 212
Secant lines, geometry of, 66
Second Velkamp space, 96
σ-sesquilinear form, 194
Sesquilinear form, associated with a pseudoquadratic form, 208
σ-sesquilinear form, left radical of, 194
σ-sesquilinear form, nondegenerate, 198
σ-sesquilinear form, proportional forms, 194
σ-sesquilinear form, reflexive, 198
σ-sesquilinear form, totally isotropic subspaces of, 198
Shadow model, of a geometry, 419
Sheaf, 417
Sheaf, residually connected, 429
Simple $\mathcal{C}$-connectedness, of graphs, 19
Singular space, 64
Singular subspace, 442
Space, singular, 64
Special pair, 463
Spread, of d-subspaces of a vector space, 111
Spread, of half-subspaces of a vector spaces, 111
Spread, of lines, 549
Strong gatedness, between vertex and subgraph, 12
Strong isometry, 384
Strongly gated, 12
  with respect to a vertex, 310, 311
Subspace, 63
Subspace, convex, 64, 455
Subspace, exterior line of, 63
Subspace, generated by a point-set, 63
Subspace, internal line of, 63
Subspace, proper, 63
Subspace, singular, 64
Subspace, tangent line of, 63
Symmetric forms, 203
Symplectic forms, 203
Symplectic rank, 463
Symplectic rank at least r, 463
Symplecton, definition, 455
Symplecton, points near, 531

T
Tangent graph, 68
Tangent lines, of embedded polar spaces, 212
Tits system, 388, 603
Tits' metric, 383
Totally isotropic subspace, 198
3-transposition group, reduced, 257
Triangle, in a graph, 25
Truncation, of a chamber system, 302
Truncation, of a geometry, 44
Truncation, of type J, 44

U
Uniformizing principle, 542
Universal cover, of a point-line geometry, 62

V
Veblen Axiom, 89
Veldkamp space, 81
  relative to $\mathcal{H}$, 82
  the second one, 96

W
Walk, 4
Walk, backtrack, 4
Walk, circular, 5
Wall, 363
Wedge product, pure vectors, 118
Witt index, 448
Word, J-reduced, 405

Z
Zara graphs, 258