

Buildings

Springer

New York

Berlin

Heidelberg

Barcelona

Budapest

Hong Kong

London

Milan

Paris

Singapore

Tokyo

Kenneth S. Brown

Buildings

With 22 Illustrations



Springer

Kenneth S. Brown
Department of Mathematics
Cornell University
Ithaca, NY 14853
USA

Figures on pages 42 and 43 are reprinted from *Vorlesungen über die Theorie der elliptischen Modulfunktionen Bd. I: Grundlegung der Theorie - 1890* by Klein and Fricke, B.G. Teubner. Reprinted with permission.

Figure on page 134 reprinted from *Trees* by J.-P. Serre, 1980, Springer-Verlag Heidelberg. Reprinted with permission.

Mathematics Subject Classification (1991): 20-02, 20E42, 51E24

Library of Congress Cataloging-in-Publication Data

Brown, Kenneth S.

Buildings / Kenneth S. Brown

p. cm.

Bibliography: p.

Includes index

ISBN 0-387-98624-3

1. Buildings (Group Theory) I. Title.

QA171.B875 1988

512'.22--dc19

Printed on acid-free paper.

© 1989 by Springer-Verlag New York Inc.

All rights reserved. This work may not be translated or copied in whole or in part without the written permission of the publisher (Springer-Verlag, 175 Fifth Avenue, New York, NY 10010, USA), except for brief excerpts in connection with reviews or scholarly analysis. Use in connection with any form of information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed is forbidden.

The use of general descriptive names, trade names, trademarks, etc. in this publication, even if the former are not especially identified, is not to be taken as a sign that such names, as understood by the Trade Marks and Merchandise Marks Act, may accordingly be used freely by anyone.

Camera-ready copy provided by the author.

Printed and bound by Braun-Brumfield, Inc., Ann Arbor, MI.

Printed in the United States of America.

9 8 7 6 5 4 3 2 1

ISBN 0-387-98624-3 Springer-Verlag New York Berlin Heidelberg SPIN 10691382

Preface

For years I have heard about buildings and their applications to group theory. I finally decided to try to learn something about the subject by teaching a graduate course on it at Cornell University in Spring 1987. This book is based on the notes from that course.

The course started from scratch and proceeded at a leisurely pace. The book therefore does not get very far. Indeed, the definition of the term “building” doesn’t even appear until Chapter IV. My hope, however, is that the book gets far enough to enable the reader to tackle the literature on buildings, some of which can seem very forbidding.

Most of the results in this book are due to J. Tits, who originated the theory of buildings. The main exceptions are Chapter I (which presents some classical material), Chapter VI (which presents joint work of F. Bruhat and Tits), and Chapter VII (which surveys some applications, due to various people). It has been a pleasure studying Tits’s work; I only hope my exposition does it justice.

A number of people read parts of a preliminary version of this book and made helpful comments. In particular, I would like to thank H. Abels, R. Alperin, A. Borel, F. Buekenhout, G. Mess, M. Ronan, J-P. Serre, C. Squier, M. Stein, J. Thévenaz, J. Tits, and W. Waterhouse. I would also like to acknowledge the partial support of the National Science Foundation.

Contents

Chapter I. Finite Reflection Groups	1
1. Definitions	1
2. Examples	2
3. Classification	5
4. Cell Decomposition	6
5. The Associated Simplicial Complex	15
Appendix. Abstract Simplicial Complexes	27
Chapter II. Abstract Reflection Groups	33
1. In Search of Axioms	33
2. Examples	37
3. Consequences of the Deletion Condition	46
4. Coxeter Groups	52
5. Loose Ends	53
Chapter III. Coxeter Complexes	58
1. The Coxeter Complex is Simplicial	58
2. Local Properties of Coxeter Complexes	59
3. Construction of Chamber Maps	63
4. Half-spaces	66
Chapter IV. Buildings	76
1. Definition and First Properties	76
2. Examples	79
3. Retractions	85
4. The Complete System of Apartments	87
5. The Spherical Case	92
6. The Homotopy Type of a Building	94
7. The Axioms for a Thick Building	97
Chapter V. Buildings and Groups	99
1. Strongly Transitive Automorphism Groups	99
2. BN-Pairs	107
3. The Building Associated to a BN-Pair	112
4. Historical Remarks	115
5. Example: The General Linear Group	118
6. Example: The Symplectic Group	120
7. Example: The Orthogonal Group	123
8. Example: The Special Linear Group Over a Field With Discrete Valuation	127

Chapter VI. Euclidean Buildings	139
1. Euclidean Reflection Groups	139
2. Euclidean Coxeter Complexes	149
3. Euclidean Buildings as Metric Spaces	151
4. The Bruhat–Tits Fixed-Point Theorem	157
5. Application: Bounded Subgroups	159
6. Bounded Subsets of Apartments	163
7. A Metric Characterization of the Apartments	165
8. Construction of Apartments	169
9. The Spherical Building at Infinity	174
Chapter VII. Applications to Group Cohomology	183
1. Arithmetic Groups Over the Rationals	183
2. S-Arithmetic Groups	189
3. Cohomological Dimension of Linear Groups	194
4. S-Arithmetic Groups Over Function Fields	195
Appendix. Linear Algebraic Groups	198
Suggestions for Further Reading	206
References	207
Notation Index	211
Subject Index	213

I

Finite Reflection Groups

This book is about connections between groups and geometry. We begin by considering groups of isometries of Euclidean space generated by hyperplane reflections. In order to avoid technicalities in this introductory chapter, we confine our attention to *finite* groups and we require our reflections to be with respect to *linear* hyperplanes (i.e., hyperplanes passing through the origin). We will generalize this in Chapter VI, replacing “finite” by “discrete” and “linear” by “affine”.

1 Definitions

Let V be a Euclidean space, i.e., a finite-dimensional real vector space with an inner product. By a *hyperplane* in V we mean a subspace $H \subset V$ of codimension 1. The *reflection* with respect to H is the linear transformation $s_H : V \rightarrow V$ which is the identity on H and is multiplication by -1 on the (one-dimensional) orthogonal complement H^\perp of H . A *finite reflection group* is a finite group W of linear transformations of V generated by reflections s_H , where H ranges over a set \mathcal{H} of hyperplanes. We will also sometimes refer to the *pair* (W, V) as a finite reflection group when it is necessary to emphasize the vector space V on which W is acting.

The requirement that W be finite is a very strong one. Suppose, for instance, that $\dim V = 2$ and that W is generated by two reflections $s = s_H$ and $s' = s_{H'}$. Then the rotation $ss' \in W$ has infinite order (and hence W is infinite) unless the angle between the lines H and H' is a rational multiple of π .

The following criterion is often useful for verifying that a given group generated by reflections is finite:

Proposition. *Let R be a finite set of non-zero vectors in V , and let \mathcal{H} be the set of hyperplanes of the form α^\perp for some $\alpha \in R$. Let W be the group generated by the reflections s_H ($H \in \mathcal{H}$). If R is invariant under the action of W , then W is finite.*

PROOF: We will show that W is isomorphic to a group of permutations of the finite set R . Let V_1 be the subspace of V spanned by R , and let V_0 be

its orthogonal complement. Then

$$V_0 = \bigcap_{\alpha \in R} \alpha^\perp = \bigcap_{H \in \mathcal{H}} H,$$

so V_0 is the fixed-point set $V^W = \{v \in V : wv = v \text{ for all } w \in W\}$. In view of the orthogonal decomposition $V = V_0 \oplus V_1$, it follows that an element of W is completely determined by its action on V_1 and hence by its action on R . \square

It will be convenient to have some terminology for describing the sort of decomposition of V that arose in the proof. Let W be a group generated by reflections s_H ($H \in \mathcal{H}$), where \mathcal{H} is an arbitrary set of hyperplanes. Let V_0 be the fixed-point set

$$V^W = \bigcap_{H \in \mathcal{H}} H.$$

We will call V_0 the *inessential* part of V , and we will call its orthogonal complement V_1 the *essential* part of V . The pair (W, V) will be called *essential* if $V_1 = V$, or, equivalently, if $V_0 = 0$.

The study of a general (W, V) is easily reduced to the essential case. Indeed, V_1 is W -invariant since V_0 is, and clearly $(V_1)^W = 0$; so we have an orthogonal decomposition $V = V_0 \oplus V_1$, where the action of W is trivial on the first summand and essential on the second. We may therefore identify W with a group acting on V_1 , and, as such, W is essential (and still generated by reflections).

2 Examples

There are two classical sources of examples of finite reflection groups:

(i) The theory of root systems, which arose historically from the study of Lie groups and Lie algebras. The precise definition of “root system” can be found in Bourbaki [16], but we will not need to know it. For our purposes, it suffices to know that a root system in a Euclidean space V is a finite subset $R \subset V - \{0\}$ and that it satisfies the hypothesis of the proposition above. Consequently, there is an associated finite reflection group W . It is called the *Weyl group* of R . [This explains why it has become customary to use the letter “ W ” for a reflection group.]

(ii) The theory of regular solids (also called regular convex polytopes), in any dimension ≥ 1 . Every regular solid X has an associated finite reflection group W , which is the group of symmetries of X , i.e., the group of isometries of the ambient Euclidean space which leave X invariant. [We should assume here that X is centered at the origin, so that the symmetries will be *linear* transformations.]

We will not assume that the reader knows anything about either of these two subjects. But it will be convenient to use the language of (i) and (ii) informally as we discuss examples. It is a fact that *all* examples of finite reflection groups can be explained in terms of (i) and/or (ii); we will return to this in the next section.

We proceed now with examples.

1. The group W of order 2 generated by a single reflection is a finite reflection group. Regardless of the dimension of V , this example is “essentially” 1-dimensional, in the sense that $\dim V_1 = 1$, where V_1 is the essential part of V as in §1. Thus we can (and should) think of W as the group $\{\pm 1\}$ acting on \mathbf{R} by multiplication. Note that, in this 1-dimensional setting, W is the group of symmetries of the regular solid $[-1, 1]$ in \mathbf{R} . It also happens to be the Weyl group of the root system $R = \{\pm 1\} \subset \mathbf{R}$, which is called the root system of type A_1 .

2. Let V be 2-dimensional, and choose two hyperplanes (lines) which intersect at an angle of π/m for some integer $m \geq 2$. Let s and t be the corresponding reflections and let W be the group $\langle s, t \rangle$ they generate. [Here and throughout this book we use angle brackets to denote the group generated by a given set.] Then the product $\rho = st$ is a rotation through an angle of $2\pi/m$ and hence is of order m . Moreover, s conjugates ρ to $s(st)s = ts = \rho^{-1}$ and similarly for t , so the cyclic subgroup $C = \langle \rho \rangle$ of order m is normal in W . Finally, the quotient W/C is easily seen to be of order 2; hence W is indeed a finite reflection group, of order $2m$.

This group W is called the *dihedral group* of order $2m$, and we will denote it by D_{2m} . [Warning: Some mathematicians, following the standard notation of crystallography, write D_m instead of D_{2m} .] If $m \geq 3$, one can check that W is the group of symmetries of a 2-dimensional solid, namely, a regular m -gon together with its interior. If $m = 3, 4$, or 6 , then W can also be described as the Weyl group of a root system. For example, the dihedral group of order 12 is the Weyl group of a root system $R \subset \mathbf{R}^2$, called the root system of type G_2 , which is defined as follows: Identify \mathbf{R}^2 with \mathbf{C} , and let R consist of the six 6th roots of unity ζ^j ($\zeta = e^{2\pi i/6}$, $j = 0, \dots, 5$) together with the six vectors $\zeta^j + \zeta^{j+1}$.

3. Let $V = \mathbf{R}^{n+1}$ with its standard inner product, where $n \geq 1$, and let W be the symmetric group on $n+1$ letters, acting on V by permuting the coordinates. Then W is generated by the transpositions s_{ij} ($1 \leq i, j \leq n+1$, $i \neq j$), where s_{ij} interchanges i and j . Now s_{ij} acts on V as the reflection with respect to the hyperplane H_{ij} defined by $x_i = x_j$, so W is indeed a finite reflection group, of order $(n+1)!$. But it is not essential. In fact, $V^W = \bigcap_{i,j} H_{ij}$ is the line $x_1 = x_2 = \dots = x_{n+1}$, which is spanned by the vector $e = (1, 1, \dots, 1)$. Thus the subspace $V_1 \subset V$ on which W is essential is the n -dimensional subspace e^\perp defined by $\sum_{i=1}^{n+1} x_i = 0$.

The interested reader can verify that W is the group of symmetries of a regular n -simplex in V_1 . [HINT: The convex hull σ of e_1, \dots, e_{n+1} is a regular

n -simplex in the affine hyperplane $\sum x_i = 1$, which is parallel to V_1 . The desired regular simplex in V_1 is now obtained from σ via the translation $x \mapsto x - b$, where b is the barycenter of σ .] W is also the Weyl group of a root system in V_1 , called the root system of type A_n . It consists of the $n(n+1)$ vectors $e_i - e_j$ ($i \neq j$), where $\{e_1, \dots, e_{n+1}\}$ is the standard basis of \mathbf{R}^{n+1} .

When $n = 1$, this example reduces to Example 1; when $n = 2$, it reduces to Example 2 with $m = 3$, i.e., W is dihedral of order 6.

4. Now let $V = \mathbf{R}^n$ ($n \geq 1$), again with its standard inner product, and let W be the group of “signed permutations” of n letters, i.e., the group of linear transformations of V leaving invariant the set $\{\pm e_i\}$ of standard basis vectors and their negatives. In terms of matrices, W can be viewed as the group of $n \times n$ monomial matrices whose non-zero entries are ± 1 . [Recall that a *monomial matrix* is one with exactly one non-zero element in every row and every column.] The group W is generated by transpositions s_{ij} as above, together with reflections t_1, \dots, t_n , where t_i changes the sign of the i th coordinate (i.e., t_i is the reflection in the hyperplane $x_i = 0$). Hence W is a finite reflection group of order $2^n n!$, and this time it is essential.

Once again, the interested reader is invited to verify that W is the group of symmetries of a regular solid in V , which one can take to be the n -cube $[-1, 1]^n$. Or, if you prefer, take the solid to be the convex hull of the $2n$ vectors $\{\pm e_i\}$; this is a “hyperoctahedron”. [The hyperoctahedron is the dual of the cube, which means that it is the convex hull of the barycenters of the faces of the cube. Since a solid and its dual have the same symmetry group, it makes no difference which one we choose. We did not mention this in our previous examples because the dual of a regular m -gon is again a regular m -gon, and the dual of a regular simplex is again a regular simplex.]

And, once again, W is the Weyl group of a root system, called the root system of type B_n . Alternatively, W can be described as the Weyl group of the root system of type C_n . [Every root system R has a “dual”, whose Weyl group is isomorphic to that of R . The dual of the root system of type B_n is called the root system of type C_n . The root systems mentioned in Examples 1–3, like the regular solids, are self-dual.]

When $n = 1$, this example reduces to Example 1; when $n = 2$ it reduces to Example 2 with $m = 4$, i.e., W is dihedral of order 8.

We close this section by mentioning an uninteresting way of constructing new examples of finite reflection groups from given ones:

EXERCISE

Given finite reflection groups (W', V') and (W'', V'') , show that the direct product $W = W' \times W''$, can be realized as a finite reflection group acting on the orthogonal direct sum $V = V' \oplus V''$.

A finite reflection group (W, V) is called *reducible* if it decomposes as in the exercise, with V' and V'' non-trivial, and it is called *irreducible* oth-

erwise. We will see later that an essential finite reflection group always admits a canonical decomposition into “irreducible components” (cf. exercise in §5E below).

3 Classification

Finite reflection groups (W, V) have been completely classified up to isomorphism. In this section we list them briefly, with little explanation; see Bourbaki [16] or Grove–Benson [31] for more details. (The notation below is that of [16], which is slightly different from that of [31].) We will confine ourselves to the reflection groups which are *essential* and *irreducible*; all others are obtained from these by taking direct sums and, possibly, adding an extra summand on which the group acts trivially.

First, we list three infinite families of reflection groups:

- Type A_n ($n \geq 1$): Here W is the symmetric group on $n + 1$ letters, acting as in Example 3 above on a certain n -dimensional subspace of \mathbf{R}^{n+1} . This group is the group of symmetries of a regular n -simplex, and it can also be described as the Weyl group of the root system of type A_n .
- Type B_n ($n \geq 2$), also known as type C_n : This is the group W of signed permutations acting on \mathbf{R}^n as in Example 4 above. (We require $n \geq 2$ because Example 4 with $n = 1$ gives the group of type A_1 again.) The group W is the group of symmetries of the n -cube (or n -dimensional hyperoctahedron); it is also the Weyl group of the root system of type B_n (or type C_n).
- Type D_n ($n \geq 4$): This is not one that we saw in §2. It is the Weyl group of a root system, called the root system of type D_n , but it does *not* correspond to any regular solid. It also happens to be a subgroup of index 2 of the reflection group of type B_n (or C_n).

Next, there are seven exceptional groups:

- Type E_n ($n = 6, 7, 8$): This is the Weyl group of the root system of the same name. It does not correspond to any regular solid.
- Type F_4 : This is the Weyl group of the root system of the same name; it is also the group of symmetries of a certain self-dual 24-sided regular solid in \mathbf{R}^4 whose (3-dimensional) faces are solid octahedra.
- Type G_2 : This is the Weyl group of the root system of the same name. As we saw in Example 2 above, W is dihedral of order 12, so we can also describe it as the group of symmetries of a hexagon.
- Type H_n ($n = 3, 4$): This does not correspond to any root system, but it is the symmetry group of a regular solid X . When $n = 3$, X is the dodecahedron (which has 12 pentagonal faces) or, dually, the icosahedron

(which has 20 triangular faces). When $n = 4$, X is a 120-sided solid in \mathbf{R}^4 (with dodecahedral faces) or, dually, a 600-sided solid (with tetrahedral faces).

Finally, we have the dihedral groups D_{2m} (not to be confused with the groups of type D_n listed above!). If $m = 2$, the group is reducible (it is $\{\pm 1\} \times \{\pm 1\}$ acting on $\mathbf{R} \oplus \mathbf{R}$). The cases $m = 3, 4$ correspond, respectively, to the groups of type A_2 and B_2 . And the case $m = 6$ corresponds to the group of type G_2 . This leaves:

- Type $I_2(m)$ ($m = 5$ or $m \geq 7$): The group W is the dihedral group of order $2m$. It is the symmetry group of a regular m -gon, but it does not correspond to any root system.

Remarks

1. The subscript in the notation for each type indicates the dimension of the vector space on which the group acts.
2. If you want to know more about the regular solids mentioned above, there are many books which discuss them; see, for instance, Coxeter [26] or Lyndon [36] or further references cited therein.

4 Cell Decomposition

Let (W, V) be an essential finite reflection group. The hyperplanes H with $s_H \in W$ cut V into polyhedral pieces, which turn out to be cones over simplices. One obtains in this way a simplicial complex Σ which triangulates the unit sphere in V . These assertions will be proved in §5. The purpose of the present section is to lay the groundwork for §5 by studying the polyhedral decomposition of Euclidean space induced by an arbitrary finite set \mathcal{H} of hyperplanes.

This section is somewhat long because it develops from scratch some basic facts about polyhedral geometry. If you are already familiar with these facts, or if you are willing to accept them as “intuitively obvious”, then you can read much of the section quickly, just to get the notation and terminology.

Throughout §4, V will denote a real vector space of finite dimension n , and $\mathcal{H} = \{H_1, \dots, H_k\}$ will denote an arbitrary finite set of linear hyperplanes in V .


4A Definition and properties of the cells

For each $i = 1, \dots, k$, let $f_i : V \rightarrow \mathbf{R}$ be a non-zero linear function such that H_i is described by the homogeneous linear equation $f_i = 0$. The function f_i is uniquely determined by H_i , up to multiplication by a non-zero scalar. A *cell* in V with respect to \mathcal{H} is a non-empty set A obtained by choosing for

each i a sign $\sigma_i \in \{+, -, 0\}$ and specifying $f_i = \sigma_i$. [Here “ $f_i = +$ ” means “ $f_i > 0$ ”, and similarly for “ $f_i = -$ ”.] Thus A is defined by k homogeneous linear equalities or strict inequalities, one for each hyperplane. In more geometric language, we have $A = \bigcap_{i=1}^k U_i$, where U_i is either H_i or one of the open half-spaces of V determined by H_i .

Note that we can recover the sign choices σ_i by looking at $f_i(x)$ for any $x \in A$. So we have a right to talk about *the* description of A by conditions $f_i = \sigma_i$.

The cells A form a partition of V into disjoint convex cones, where a *cone* is a subset closed under multiplication by positive scalars. The picture below shows a simple example, with $\dim V = 2$ and $k = 3$. The three lines partition the plane into 13 cells (6 open sectors, 6 open rays, and the cell consisting of the origin), even though there are actually 27 possible sign choices $(\sigma_i)_{1 \leq i \leq 3}$. This is typical—it simply means that many of the sign choices are inconsistent (i.e., they define the empty set). Note, for instance, that the inequalities $f_1 > 0$, $f_2 > 0$, and $f_3 < 0$ are inconsistent; for the first two of them define one of the sectors shown in the picture, and f_3 happens to be positive on that sector.



The *support* of a cell A is the linear subspace L defined by the equalities $f_i = 0$ that occur in the description of A . [If there are no such equalities, then $L = V$.] Note that A , as a subset of L , is defined by strict inequalities, hence it is an open subset of L . In particular, it follows that L is the linear span of A . For example, the support of each of the rays in the picture above is the line spanned by the ray, and the support of each of the open sectors is the whole space V .

A cell B is called a *face* of A if its description is obtained from that of A by replacing zero or more inequalities by equalities. We write $B \leq A$ in this case. For example, each sector in the picture above has four faces: the sector itself, two rays, and the origin.

Let \bar{A} be the set obtained by replacing the strict inequalities $f_i > 0$ or $f_i < 0$ in the description of A by the corresponding weak inequalities $f_i \geq 0$ or $f_i \leq 0$. We call \bar{A} the *closed cell* associated to A . [The cell A itself, by contrast, will often be called an *open cell*, even though it is not in general an open subset of V . It is, of course, open in its support L and hence open in \bar{A} .] For example, the closed cell corresponding to each of the open sectors above is the corresponding closed sector.

It is immediate from the definitions that

$$\bar{A} = \bigcup_{B \leq A} B.$$

Since the open cells are disjoint, it follows that the face relation can be characterized in terms of the closed cells:

$$B \leq A \iff \bar{B} \subseteq \bar{A}.$$

This shows, in particular, that $B = A$ if and only if $\bar{B} = \bar{A}$. Hence:

Proposition 1. *The function $A \mapsto \bar{A}$ is a bijection from the open cells to the closed cells.* \square

We will find it helpful to have a geometric description of the correspondence between open cells and closed cells, which does not refer to \mathcal{H} :

Proposition 2. *Let A be an open cell.*

- (1) *\bar{A} is the closure of A in V (in the sense of point-set topology).*
- (2) *Let L be the linear span of \bar{A} . Then A is the interior of \bar{A} in L , i.e., the largest open subset of L contained in \bar{A} .*

PROOF: (1) Clearly \bar{A} is closed in V , so it contains the closure of A . Conversely, given $y \in \bar{A}$, choose $x \in A$ and consider the closed line segment from x to y , denoted $[x, y]$. Each equality in the description of A holds on the whole line segment; and each strict inequality holds on the half-open segment $[x, y)$. So $[x, y) \subseteq A$ and hence y is in the closure of A .

(2) Note first that L is simply the support of A . For the support of A contains \bar{A} and is spanned by A , so it is also spanned by \bar{A} . We therefore have $A \subseteq \text{int}_L(\bar{A})$ (= the interior of \bar{A} in L) since A is open in its support. Conversely, suppose $y \in \bar{A} - A$ and consider the segment $[x, y]$ again. Since $y \notin A$, there must be an inequality in the description of A , say $f_i > 0$, such that $f_i(y) = 0$. So if the line segment is continued past y , we immediately have $f_i < 0$, which means we have left \bar{A} (but stayed in L). Hence $y \notin \text{int}_L(\bar{A})$. \square

Our next observation is that we can give a direct definition of what it means to be a closed cell, independent of the notion of “open cell”. Recall that a closed cell is defined by k equalities or weak inequalities, one for each i . Conversely, suppose X is an arbitrary set defined by specifying for each i the equality $f_i = 0$ or one of the weak inequalities $f_i \geq 0$ or $f_i \leq 0$; we will show that X is a closed cell:

Proposition 3. *Let X be a set defined by k equalities or weak inequalities as above. Then X is a closed cell with respect to \mathcal{H} .*

PROOF: Let σ_i be 0 if $f_i = 0$ on X . Otherwise either $f_i \geq 0$ on X or $f_i \leq 0$ on X , and we take σ_i to be $+$ or $-$, accordingly. [Caution: It is possible that our original description of X involved an inequality, say $f_i \geq 0$, but that nevertheless $f_i = 0$ on X ; so σ_i is 0 in this case.] Let A be the set defined by the signs σ_i . If A is non-empty, then it is a cell and $X = \bar{A}$. To prove $A \neq \emptyset$, choose for each i with $\sigma_i \in \{+, -\}$ a vector $x_i \in X$ with $f_i(x_i) \neq 0$. Let x be the sum of these vectors (or 0 if there are none). Then $x \in A$. \square

Corollary. *An intersection of closed cells is a closed cell.* \square

We turn, finally, to the geometric meaning of the face relation. If you visualize a cell in dimension 2 or 3, you certainly have no trouble seeing what its faces are, without knowing the particular system of equalities and inequalities by which A was defined. Roughly speaking, the faces are the flat pieces into which the boundary of A decomposes. The following proposition states this precisely:

Proposition 4. *Let A be a cell. Then two distinct points $y, z \in \bar{A}$ lie in the same face of A if and only if there is an open line segment containing both y and z and lying entirely in \bar{A} . Consequently, the partition of \bar{A} into faces depends only on A as a subset of V , and not on the set \mathcal{H} of hyperplanes.*

PROOF: Suppose y and z are in the same face $B \leq A$. For each condition $f_i = \sigma_i$ in the description of B , we can extend the segment $[y, z]$ slightly in both directions without violating the condition. Since there are only finitely many such conditions, it follows that B contains an open segment containing y and z , hence so does \bar{A} . Conversely, suppose y and z are in different faces of A . Then there is some i such that y and z behave differently with respect to f_i , say $f_i(y) > 0$ and $f_i(z) = 0$. If we now continue the segment $[y, z]$ past z , we immediately have $f_i < 0$, so we leave \bar{A} ; hence \bar{A} does not contain an open segment containing both y and z . \square

The significance of this for us is that if we want to understand the polyhedral structure of a particular cell A , then we can replace \mathcal{H} by any other finite set of hyperplanes for which A is still a cell. In practice, we will want to take a minimal set of hyperplanes for a given A .

4B Chambers and walls

The cells defined by taking all $\sigma_i \in \{+, -\}$ are called *chambers*. They are non-empty open convex sets which partition the complement $V - \bigcup_{i=1}^k H_i$, so they are the connected components of that complement. They can also be described as the cells whose support is V and hence as the cells of maximal

dimension (where the dimension of a cell is defined to be the dimension of its support).

Fix a chamber C . We say that a subset $\mathcal{H}' \subseteq \mathcal{H}$ *defines* C if C is defined by the conditions $f_i = \sigma_i$, where i ranges over the indices such that $H_i \in \mathcal{H}'$.

Proposition 1. *There is a unique minimal subset $\mathcal{H}' \subseteq \mathcal{H}$ which defines C . The elements of \mathcal{H}' are precisely the supports of the codimension 1 faces of C .*

PROOF: Choose an arbitrary minimal $\mathcal{H}' \subseteq \mathcal{H}$ defining C . This is possible because \mathcal{H} is finite. Renumber the H 's so that $\mathcal{H}' = \{H_1, \dots, H_r\}$ for some $r \leq k$. And choose the f_i so that $f_i > 0$ on C for $i = 1, \dots, r$. As we noted above, we can use \mathcal{H}' instead of \mathcal{H} to determine the faces of C . In particular, any codimension 1 face of C has support H_i for some $i \leq r$. [An intersection of more than one H_i has codimension at least 2.] It remains to show that each H_i for $i \leq r$ supports a codimension 1 face, i.e., that we get a non-empty set A by specifying $f_i = 0$ and $f_j > 0$ for $j \neq i$ (and $j \leq r$). By the minimality of \mathcal{H}' , the inequalities $f_j > 0$ for $j \neq i$ ($j \leq r$) define a subset C' strictly bigger than C . Take $y \in C' - C$ and $x \in C$. Since $f_i(x) > 0$ and $f_i(y) \leq 0$, there is a point $z \in (x, y]$ such that $f_i(z) = 0$. Then $z \in A$. \square

The elements of \mathcal{H}' are called the *walls* of C . The second sentence of the proposition makes it clear that they are *intrinsically* associated to C and do not depend on the original set \mathcal{H} of hyperplanes. Here is a useful characterization of them:

Proposition 2. *Let H be a linear hyperplane in V . Then H is a wall of C if and only if C lies on one side of H and $\overline{C} \cap H$ has non-empty interior in H .*

PROOF: If H is the support of a codimension 1 face A of C , then certainly C lies on one side of H and $\overline{C} \cap H$ contains A , which is a non-empty open subset of H . Conversely, suppose H is a hyperplane such that C lies on one side of H and $\overline{C} \cap H$ has non-empty interior in H . Then C is still a cell with respect to $\mathcal{H}^+ = \mathcal{H} \cup \{H\}$, so we can use \mathcal{H}^+ to determine the faces of C . By Proposition 3 of §4A, the set $\overline{C} \cap H$ is a closed cell \bar{A} with respect to \mathcal{H}^+ , and clearly the corresponding open cell A is a face of C (because $\bar{A} \subseteq \overline{C}$). Since \bar{A} is contained in H and has non-empty interior in H , the support of A must be H . Thus A has codimension 1 and its support H is therefore a wall of C . \square

4C The structure of a chamber

We continue to denote by C a fixed but arbitrary chamber and by \mathcal{H}' the set of walls of C . For simplicity of notation we will assume, as in the proof of Proposition 1 above, that the elements of \mathcal{H}' are the hyperplanes $f_i = 0$ for $1 \leq i \leq r$ and that $f_i > 0$ on C for $1 \leq i \leq k$.

Let $V_0 = \bigcap_{i=1}^k H_i$. We will call \mathcal{H} *essential* if $V_0 = 0$. There is no loss of generality in restricting attention to the essential case. For if $V_1 = V/V_0$, then the linear functions f_i pass to the quotient V_1 and define an essential set of hyperplanes there. And the cells determined by these hyperplanes in V_1 are in 1-1 correspondence with the cells in V . More precisely, the cells in V are the inverse images in V of the cells in V_1 . [Geometrically, then, the cells in V are simply the cells in V_1 , “fattened up” by a factor \mathbf{R}^d , where $d = \dim V_0$.]

Assume now that \mathcal{H} is essential. Then we have $\bigcap_{i=1}^r H_i = 0$. For V_0 is the “smallest” cell (i.e., it is a face of every cell), hence it is certainly the smallest face of C ; but the smallest face of C can be determined by using \mathcal{H}' instead of \mathcal{H} , so it is also equal to $\bigcap_{i=1}^r H_i$. In particular, the latter is 0 if $V_0 = 0$.

It follows that $r \geq n = \dim V$. It is easy to visualize examples where inequality holds (e.g., C could be the cone over the interior of a square, in which case $r = 4 > 3 = \dim V$). We will now prove that equality holds if and only if the cone C is *simplicial*, by which we mean that, for some basis e_1, \dots, e_n of V , C consists of the linear combinations $\sum_{i=1}^n \lambda_i e_i$ with all $\lambda_i > 0$. [In other words, C is the interior of the cone over the simplex with vertices e_1, \dots, e_n .]

Proposition. *Assume that \mathcal{H} is essential. Then the following conditions on C are equivalent:*

- (1) C is a simplicial cone.
- (2) C has exactly n codimension 1 faces, i.e., $r = n$.
- (3) f_1, \dots, f_r are linearly independent.
- (4) f_1, \dots, f_r form a basis for the dual space V^* of V .

PROOF: As we noted above, the assumption that \mathcal{H} is essential implies that $\bigcap_{i=1}^r H_i = 0$, i.e., that the equations $f_1 = 0, \dots, f_r = 0$ have only the trivial solution. The equivalence of (2), (3), and (4) follows easily from this by elementary linear algebra.

Suppose now that (2)–(4) hold, and let $(e_i)_{1 \leq i \leq n}$ be the basis of V dual to (f_i) . Then the description “ $f_i > 0$ for $1 \leq i \leq n$ ” of C implies that C consists of the positive linear combinations of the e_i , whence (1).

Conversely, (1) implies that C is defined by $x_i > 0$ for $1 \leq i \leq n$, where x_i is the i th coordinate function with respect to some basis for V . We can use this description of C to determine its walls, which are easily seen to be the coordinate hyperplanes $x_i = 0$; this proves (2)–(4). \square

4D A sufficient condition for C to be simplicial

The result of this subsection will be used later to show that the chambers associated to an essential finite reflection group are always simplicial cones.

Assume that V has an inner product $\langle -, - \rangle$. Then the linear function f_i is given by $\langle e_i, - \rangle$ for a unique vector e_i . Replacing f_i by a scalar multiple, we

may assume $\|e_i\| = 1$; thus e_i is one of the two unit vectors perpendicular to H_i . If there is a fixed chamber C under discussion, as there is at the moment, then we can remove this ambiguity by requiring that e_i point toward the half-space bounded by H_i containing C . This is equivalent to requiring, as above, that $f_i > 0$ on C .

In summary, then, we are now assuming that the chamber C is defined by $\langle e_i, - \rangle > 0$ for $1 \leq i \leq k$, where the e_i are unit vectors, and that the first r of these inequalities in fact suffice to define C . Moreover, no smaller set of inequalities defines C .

Proposition. *Assume \mathcal{H} is essential. If $\langle e_i, e_j \rangle \leq 0$ for each $i \neq j$ ($i, j \leq r$), i.e., if the angle between e_i and e_j is not acute, then C is a simplicial cone.*

PROOF: We must show that e_1, \dots, e_r are linearly independent. If not, then I claim that there is a non-trivial linear relation among them with non-negative coefficients. For let $\sum_{i=1}^r \lambda_i e_i = 0$ be an arbitrary linear relation (with the λ_i not all zero). If the non-zero λ_i all have the same sign, the claim follows at once. Otherwise we can rewrite the relation in the form

$$\sum_{i \in I} \mu_i e_i = \sum_{j \in J} \mu_j e_j,$$

with I and J disjoint non-empty subsets of $\{1, \dots, r\}$ and all coefficients $\mu_i, \mu_j > 0$. Then the inner product of the left-hand side of this equation with the right-hand side is ≤ 0 . But this is the inner product of a vector with itself, so that vector must be 0. Thus both sides of the equation are 0, and the claim is proved.

Note that what we have done so far applies to *any* set of vectors with pairwise non-positive inner products. But now let's add the additional information that the inequalities $\langle e_i, - \rangle > 0$ ($i \leq r$) define the (non-empty) chamber C . This is clearly inconsistent with the existence of a non-trivial non-negative linear relation among the e_i , so we have reached a contradiction. Thus e_1, \dots, e_r are indeed linearly independent. \square

4E Formal properties of the poset of cells

We return to the generality of §4A above, i.e., \mathcal{H} is not necessarily essential (although it might as well be) and V is not assumed to be equipped with an inner product. Let Σ be the partially ordered set (or *poset*) consisting of the open cells, ordered by the face relation. Recall from §4A that Σ is isomorphic to the set of *closed* cells, ordered by inclusion. Recall, also, that any intersection of closed cells is a closed cell; consequently:

Proposition 1. *Any two elements of Σ have a greatest lower bound.* \square

We will denote by $A \cap B$ the greatest lower bound of two open cells A and B . It is, of course, *not* the set-theoretic intersection of A and B , this intersection being empty unless $A = B$; it is, rather, the open cell whose closure is the intersection of the closure of A and the closure of B .

Proposition 2. *Any cell $A \in \Sigma$ is a face of a chamber. If A has codimension 1, then it is a face of exactly two chambers.*

PROOF: Let A be defined by the sign choices $(\tau_i)_{1 \leq i \leq k}$. Choose $x \in A$ and $y \in V - \bigcup_{i=1}^k H_i$. [Such a y certainly exists: V is not the union of finitely many hyperplanes.] For any i with $\tau_i \in \{+, -\}$, set $\sigma_i = \tau_i$. For all other i , let σ_i be the sign of $f_i(y)$. Then all points of the half-open segment $(x, y]$ sufficiently close to x satisfy the conditions $f_i = \sigma_i$ for all i , so these conditions define a chamber C and $A \leq C$.

Suppose now that A has codimension 1. Then there is exactly one index i , say $i = 1$, such that τ_i is 0. Hence there are at most two chambers C with $A < C$, corresponding to the two possible signs σ_1 for C . The proof in the previous paragraph shows that we do get such a C by taking σ_1 to be the sign of $f_1(y)$. But the other choice of σ_1 works just as well, as one sees by continuing the line from y to x a little past x and noting that f_1 changes sign at x . \square

We will call two chambers C and C' *adjacent* if they have a common codimension 1 face. Thus either $C = C'$ or else there is a codimension 1 cell A such that C and C' are the two chambers having A as a face. In the second case, it follows from the proof above that there is a unique i [which we took to be 1 in the proof] such that C and C' are on opposite sides of H_i . This hyperplane $H = H_i$ is the support of A , and A is necessarily $C \cap C'$. [You can prove this last assertion by a dimension argument or simply by checking the definition of $C \cap C'$.] We will often say, in this situation, that “ C and C' are adjacent along the wall H ”.

It is admittedly counter-intuitive to call a chamber C adjacent to itself, as we have agreed to do, but this terminology will prove convenient later when we study “chamber maps” and their effect on “galleries”.

A *gallery* is a sequence of chambers $\Gamma = (C_0, \dots, C_d)$ such that consecutive chambers C_{i-1} and C_i ($i = 1, \dots, d$) are adjacent. We will write

$$\Gamma : C_0, \dots, C_d$$

and say that Γ is a gallery from C_0 to C_d , or that Γ *connects* C_0 and C_d . The integer d is called the *length* of Γ . Note that we allow the possibility that $C_{i-1} = C_i$ for some i ; we say that Γ *stutters* if this happens.

For any two chambers C, D , let $l(C, D)$ be the number of $H \in \mathcal{H}$ which separate C from D , i.e., the number of indices i such that the sign choices defining C and D differ.

Proposition 3. *Any two chambers $C, D \in \Sigma$ can be connected by a gallery of length $l(C, D)$.*

PROOF: We argue by induction on $l = l(C, D)$. If $C = D$ there is nothing to prove, so assume $C \neq D$. Then there must be a wall of C separating C from D . For C is defined by inequalities corresponding to its walls;

and if D also satisfied all these inequalities, then we would have $D \subseteq C$, contradicting the fact that C and D are disjoint.

So choose a wall H of C separating C from D , let A be the face of C whose support is H , and let C' be the chamber different from C having A as a face. Then, as we noted above, H is the only element of \mathcal{H} separating C from C' , so $l(C', D) = l - 1$. By induction we can find a gallery Γ' of length $l - 1$ connecting C' to D ; the gallery (C, Γ') consisting of C followed by Γ' therefore has length l and connects C to D . \square

The *combinatorial distance* between C and D , denoted $d(C, D)$, is defined to be the minimal length of a gallery connecting C to D . And any gallery $\Gamma : C = C_0, \dots, C_d = D$ which achieves this minimum will be called a *minimal gallery* from C to D . Such a Γ is of course non-stuttering, so we have a well-defined sequence $H_1, \dots, H_d \in \mathcal{H}$ such that C_{i-1} and C_i are adjacent along H_i . [Warning: This notation has nothing to do with our original numbering of the elements of \mathcal{H} as H_1, \dots, H_k ; we will have no further need for the older notation.] We will refer to the H_i as the “walls crossed” by Γ . Since H_i is the only element of \mathcal{H} which separates C_{i-1} from C_i , it is clear that H_1, \dots, H_d are the only elements of \mathcal{H} which can possibly separate C from D . Consequently:

$$l(C, D) \leq \text{card}\{H_1, \dots, H_d\} \leq d = d(C, D),$$

where “card” denotes the cardinality of a set. On the other hand, we have $d(C, D) \leq l(C, D)$ by Proposition 3. Hence all of the inequalities above are actually equalities. This proves the first two parts of the following proposition:

Proposition 4.

- (1) $d(C, D) = l(C, D)$.
- (2) If Γ is any minimal gallery connecting C to D , then Γ crosses exactly once each element of \mathcal{H} which separates C from D .
- (3) If C and C' are distinct adjacent chambers and D is an arbitrary chamber, then $d(C, D) = d(C', D) \pm 1$. The sign is “+” if C' and D are on the same side of the wall which separates C from C' .

PROOF OF (3): If $d(-, -)$ is replaced by $l(-, -)$, then the assertion is easy and was essentially proved in the proof of Proposition 3; statement (3) therefore follows from (1). \square

Remarks

1. A minimal gallery should be thought of as the combinatorial analogue of a geodesic. The following restatement of (2) is consistent with this intuition: *Any minimal gallery from C to D crosses precisely the same elements of \mathcal{H} as a straight line segment joining a point of C to a point of D .* [Of course, the minimal gallery need not cross these walls in the same order as the line segment.]

2. At first glance, it may seem that the equality in (3) is immediate from the definition of $d(-, -)$. But this is not true. All that is obvious from the definition is that $d(C, D)$ differs from $d(C', D)$ by at most 1. Some further argument is needed to show that $d(C, D) \neq d(C', D)$.

By the *diameter* of Σ we mean the maximum distance $d(C, D)$ between two chambers C, D . In view of (1) of Proposition 4, it is clear that this diameter is at most $k = \text{card}(\mathcal{H})$. In fact, we have $d(C, D) \leq k$, with equality if and only if C and D differ with respect to all k sign choices σ_i . This proves:

Corollary. *The diameter of Σ is $k = \text{card}(\mathcal{H})$. For any chamber C , there is a unique chamber D with $d(C, D) = k$, namely, $D = -C$. \square*

EXERCISES

1. Given $A, C \in \Sigma$ with C a chamber, consider galleries $C_0, \dots, C_d = C$ with $A \leq C_0$. Such a gallery will be said to connect A to C . Prove analogues of (1) and (2) of Proposition 4 for the minimal length $d(A, C)$ of such a gallery.

2. By a *subcomplex* of Σ we mean a subset $\Delta \subseteq \Sigma$ which contains, for each $A \in \Delta$, all the faces of A . Let Δ be a subcomplex containing at least one chamber, and let X be the corresponding subset $\bigcup_{A \in \Delta} A \subseteq V$. Prove that the following conditions on Δ are equivalent:

- (1) X is convex.
- (2) Given $A, C \in \Delta$ with C a chamber, Δ contains every minimal gallery from A to C .
- (3) Every $A \in \Delta$ is a face of a chamber in Δ , and Δ contains every minimal gallery joining two of its chambers.
- (4) X is an intersection of closed half-spaces bounded by elements of \mathcal{H} .

[This exercise takes some work, but it is well worth the effort. I will therefore not spoil the fun by giving a hint, except to suggest that you prove the equivalence in the usual circular fashion: (1) \implies (2) \implies (3) \implies (4) \implies (1).]

3. If Δ does not contain a chamber, show that (1) and (4) are still equivalent. [HINT: Suppose (1) holds. Take a maximal $A \in \Delta$, let $L = \text{support}(A)$, and show that $X \subseteq L$. Then A is a chamber in L with respect to a suitable set of hyperplanes, and you can now apply Exercise 2.]

5 The Associated Simplicial Complex

We return, finally, to the setup at the beginning of the chapter, where V is assumed to have an inner product, W is a finite reflection group acting on V , and \mathcal{H} is a set of hyperplanes such that the reflections s_H ($H \in \mathcal{H}$) generate W . We assume further that \mathcal{H} is W -invariant. [Such an \mathcal{H} certainly exists. For example, we could take \mathcal{H} to consist of all hyperplanes H with $s_H \in W$; the W -invariance of this set follows from the easily verified identity $s_{wH} = ws_Hw^{-1}$.] Note that \mathcal{H} is necessarily finite, so §4 is applicable. Let Σ be the poset of cells associated to \mathcal{H} .

5A The action of W on Σ

Since Σ is defined in terms of \mathcal{H} and the linear structure on V , it is clear that W permutes the cells and preserves the face relation. In other words, W acts on Σ as a group of poset automorphisms. Note that Σ appears to depend on the choice of \mathcal{H} ; but part (3) of the following theorem shows that there is only one possible \mathcal{H} , so Σ in fact depends only on (W, V) .

Theorem.

- (1) *The action of W is simply-transitive on the set of chambers. In particular, the number of chambers is equal to $|W|$ (the order of W).*
- (2) *If C is any fixed chamber, then W is generated by the reflections s_H such that H is a wall of C .*
- (3) *\mathcal{H} necessarily consists of all hyperplanes H in V with $s_H \in W$.*

The proof will consist of a sequence of observations which will lead to the statements (1)–(3), among other things.

- (a) *If C is a chamber, H is a wall of C , and s is the reflection s_H , then sC is adjacent to C along H .*

PROOF: If A is the face of C supported by H , then sA is the face of sC with support sH . But $sA = A$ and $sH = H$, so C and sC (which are clearly distinct since they are separated by H) have a common codimension 1 face supported by H . \square

Now fix a chamber C and let S be the set of reflections with respect to the walls of C .

- (b) *For any $w \in W$ and $s \in S$, wsC is adjacent to wC along the wall wH , where $s = s_H$.*

PROOF: This follows from (a) by applying the action of w . \square

The following schematic diagram should help you remember the statement of (b):

$$\begin{array}{ccc} C \mid sC & \xrightarrow{w} & wC \mid wsC \\ H & & wH \end{array}$$

- (c) *Given $s_1, \dots, s_d \in S$, there is a non-stuttering gallery*

$$C, s_1C, s_1s_2C, \dots, s_1s_2 \cdots s_dC.$$

Conversely, any non-stuttering gallery starting at C has this form.

PROOF: The first assertion follows immediately from (b). Conversely, suppose C_0, \dots, C_d is an arbitrary non-stuttering gallery with $C_0 = C$. Assume inductively that s_1, \dots, s_i have been constructed and that $C_i = w_iC$, where $w_i = s_1 \cdots s_i$. Then the wall along which C_i and C_{i+1} are adjacent must have the form w_iH for some wall H of C . Letting $s = s_H$, we now have the following situation (cf. (b)):

$$\begin{array}{ccc} C \mid sC & \xrightarrow{w_i} & C_i \mid w_i sC \\ H & & w_i H \end{array}$$

It follows that $C_{i+1} = w_i s C$, so we can complete the induction by setting $s_{i+1} = s$. \square

Let W' be the subgroup $\langle S \rangle$ of W generated by S . (c) shows that W' is transitive on the chambers, since any chamber can be connected to C by a gallery and hence has the form $s_1 \cdots s_d C$. In particular, W is transitive on the chambers, which proves part of (1).

(d) $W = W'$; in other words, (2) holds.

PROOF: It suffices to show that W' contains the generators s_H of W ($H \in \mathcal{H}$). I claim, first, that any $H \in \mathcal{H}$ is a wall of some chamber D . For H cannot be the union of its proper subspaces $H' \cap H$ ($H' \in \mathcal{H}$, $H' \neq H$). So if we pick any $x \in H$ not in this union, then the cell A containing x has support equal to H . Hence H is a wall of either of the chambers D having A as a face, as claimed.

Now we know that $D = wC$ for some $w \in W'$, so that $H = wH'$ for some wall H' of C . Letting $s = s_{H'}$, we obtain $s_H = w s w^{-1} \in W'$. \square

To complete the proof of (1), we must show that the stabilizer of C in W is trivial:

(e) If w is a non-trivial element of W , then $wC \neq C$.

PROOF: In view of (d), we can write $w = s_1 \cdots s_d$, with $s_i \in S$. Choose a minimal such expression (i.e., one with minimal d) and consider the corresponding gallery $\Gamma : (w_i C)_{0 \leq i \leq d}$ from C to wC , where $w_i = s_1 \cdots s_i$. If H_i is the wall of C fixed by s_i , then the walls crossed by Γ are the transforms $w_{i-1} H_i$ ($i = 1, \dots, d$):

$$\begin{array}{ccc} C \big|_{H_i} s_i C & \xrightarrow{w_{i-1}} & C_{i-1} \big|_{w_{i-1} H_i} C_i \end{array}$$

We will show that the first of these walls, H_1 , separates C from wC and hence that $wC \neq C$. [In fact, we could just as easily prove that all of the walls crossed by Γ separate C from wC , but we won't bother.]

To prove this, recall that the wall along which two consecutive chambers of Γ are adjacent is the unique element of \mathcal{H} separating these chambers. Thinking about the sign choices which define the chambers, one concludes easily that H_1 must separate C from wC unless it is crossed more than once. So suppose that $w_{i-1} H_i = H_1$ for some $i > 1$. Passing to the associated reflections, we obtain

$$w_{i-1} s_i w_{i-1}^{-1} = s_1,$$

or

$$w_{i-1} s_i = s_1 w_{i-1}.$$

Now expand w_{i-1} in terms of the s 's and cancel s_1^2 to get

$$s_1 \cdots s_i = s_2 \cdots s_{i-1}.$$

Thus we can delete s_1 and s_i from the expression $w = s_1 \cdots s_d$. This contradicts the minimality of d and completes the proof of (e). \square

The final step is to prove (3):

(f) *If H is a hyperplane in V with $s_H \in W$, then $H \in \mathcal{H}$.*

PROOF: Suppose, to the contrary, that $H \notin \mathcal{H}$. Then $H \not\subseteq \bigcup_{H' \in \mathcal{H}} H'$, since otherwise H would be a finite union of proper subspaces $H \cap H'$. So H must meet a chamber D . Since the element $w = s_H$ of W fixes H , it follows that wD meets D and hence that $wD = D$, contradicting (e). \square

Remarks

1. The form of the contradiction which we obtained in the proof of (e) is rather remarkable, in that we were able to shorten the word $s_1 \cdots s_d$ simply by deleting two of its letters. We will explore the consequences of this in Chapter II, in a more general setting.

2. It follows from (c) and (e) that minimal expressions for w as a word $s_1 \cdots s_d$ are in 1-1 correspondence with minimal galleries from C to wC . In particular, we now know that the gallery Γ considered in the proof of (e) is minimal, and hence that all the walls crossed by Γ do in fact separate C from wC by Proposition 4 in §4E.

5B Examples

1. Suppose that (W, V) is essential and that $\dim V = 2$. One could simply give a direct analysis of this situation, but it will be instructive to see what the theorem says about it. Let $m = \text{card}(\mathcal{H})$. Then $m \geq 2$, and the m lines in \mathcal{H} divide the plane V into $2m$ chambers, each of which is a sector determined by two rays. The transitivity of W on the set of sectors implies that they are all congruent, so each sector must have angle $2\pi/2m = \pi/m$. In view of (2), then, W is generated by two reflections in lines L_1 and L_2 which intersect at an angle of π/m . In other words, W is *dihedral of order $2m$ and (W, V) looks exactly like Example 2 of §2*.

Let us also record, for future reference, the following fact about this example: Let L_1 and L_2 be the walls of one of the chambers C , and let e_i be the unit normal to L_i ($i = 1, 2$) pointing to the side of L_i containing C . Then the inner product of e_1 and e_2 is given by

$$\langle e_1, e_2 \rangle = -\cos \frac{\pi}{m}.$$

[If the minus sign surprises you, draw a picture; you will see that the angle between e_1 and $-e_2$ is π/m .]

2. This is a trivial generalization of Example 1, but it will be useful to have it on record. Instead of assuming that $\dim V = 2$, we will only assume that (W, V) is *essentially* 2-dimensional. In other words, if we write $V = V_0 \oplus V_1$ as in §1, then $\dim V_1 = 2$. By Example 1 applied to (W, V_1) , we have $W \approx D_{2m}$ for some $m \geq 2$. Moreover, if C_1 is a chamber in V_1 with walls L_i and normals e_i as above, then $V_0 \times C_1$ is a chamber in V

with walls $V_0 \oplus L_i$ and the *same* normals e_i . In particular, it is still true that a chamber C has two walls and that the corresponding unit normals (pointing toward C) satisfy

$$\langle e_1, e_2 \rangle = -\cos \frac{\pi}{m}.$$

3. Let W be the symmetric group on $n + 1$ letters, acting on \mathbf{R}^{n+1} as in Example 3 of §2. Then we can take \mathcal{H} to consist of the $\binom{n+1}{2}$ hyperplanes H_{ij} ($i < j$), where H_{ij} is defined by $x_i = x_j$; this set is clearly W -invariant. One chamber C is the set defined by the $\binom{n+1}{2}$ inequalities $x_i < x_j$ for $i < j$. But these inequalities are redundant; C is actually defined by the n inequalities

$$x_1 < x_2 < \cdots < x_{n+1}.$$

Thus C has n walls, the i th of which is given by the equation $x_i = x_{i+1}$ ($i = 1, \dots, n$). [Note: n is the “right” number of walls, since this example is essentially n -dimensional.] The reflection with respect to the i th wall is the transposition $s_i = s_{i,i+1}$ which interchanges i and $i + 1$, so part (2) of the theorem reduces to the well-known fact that these n “basic” transpositions generate the symmetric group. And part (1) is also easy to verify directly: There are exactly $(n + 1)!$ chambers, one for each possible way of imposing a linear ordering on the $n + 1$ coordinates.

Let’s compute, now, the canonical unit vectors e_1, \dots, e_n associated to C . Let v_1, \dots, v_{n+1} be the standard orthonormal basis for $V = \mathbf{R}^{n+1}$. Then the i th inequality defining C can be written $\langle v_{i+1} - v_i, x \rangle > 0$, so the unit vector e_i perpendicular to the i th wall and pointing toward C is given by

$$e_i = \frac{v_{i+1} - v_i}{\sqrt{2}}.$$

In particular, we can calculate the inner product

$$\langle e_i, e_j \rangle = \begin{cases} 1 & \text{for } j = i \\ -1/2 & \text{for } j = i + 1 \\ 0 & \text{for } j > i + 1. \end{cases}$$

Note that $1 = -\cos(\pi/1)$, $-1/2 = -\cos(\pi/3)$, and $0 = -\cos(\pi/2)$. Hence the inner product calculation can be written in the more concise form

$$\langle e_i, e_j \rangle = -\cos \frac{\pi}{m_{ij}},$$

where m_{ij} is the order of $s_i s_j$ (or, equivalently, $2m_{ij}$ is the order of the dihedral subgroup generated by s_i and s_j). This formula should not be shocking, in view of Examples 1 and 2 above.

You might want to similarly analyze Example 4 of §2, where W is the signed symmetric group.

4. This final example is presented to help you develop some geometric intuition about the subject. A number of statements will be made without

proof, and you are advised not to worry too much about this—just convince yourself that the statements are intuitively plausible.

Let W be the reflection group of type H_3 , i.e., the group of symmetries of a regular dodecahedron in $V = \mathbf{R}^3$. It is convenient to restrict the action of W to the unit sphere S^2 and to think of W as a group of isometries of this sphere. As such, it is the group of symmetries of the regular tessellation of the sphere obtained by radially projecting the faces of the dodecahedron onto the sphere. Let P be one of the 12 spherical pentagons which occur in this tessellation. It has interior spherical angles $2\pi/3$ since there are 3 pentagons at each vertex.

The planes of symmetry of the dodecahedron barycentrically subdivide P , thereby cutting it into 10 spherical triangles. A typical such triangle T has angles $\pi/2$, $\pi/3$, and $\pi/5$. [The angle $\pi/5 = 2\pi/10$ occurs at the center of P ; the angle $\pi/3$, which is half of the interior angle $2\pi/3$ of P , occurs at a vertex of P ; and the angle $\pi/2$ occurs at the midpoint of an edge of P , where the line from the center of P perpendicularly bisects that edge.]

Finally, a typical chamber C in V is simply the cone over such a triangle T . There are $12 \cdot 10 = 120$ such chambers, so part (1) of the theorem implies that $|W| = 120$. Thus the dodecahedral group W is a group of order 120 generated by 3 reflections. The calculation of the angles of T above makes it easy to compute the orders of the pairwise products of the generating reflections. One has, for a suitable numbering s_1, s_2, s_3 of these reflections,

$$(s_1 s_2)^3 = (s_2 s_3)^5 = (s_1 s_3)^2 = 1.$$

5C The structure of a chamber

Let (W, V) be a finite reflection group. Fix a chamber C and let its walls be H_1, \dots, H_r , with corresponding unit normals e_1, \dots, e_r , where e_i points to the side of H_i containing C . Let s_i be the reflection s_{H_i} with respect to the i th wall. Thus $S = \{s_1, \dots, s_r\}$ is the set of generators of W described in the last theorem.

The reader who has worked through the examples above will not be surprised by the following result:

Theorem. *With the notation above, we have*

$$\langle e_i, e_j \rangle = -\cos \frac{\pi}{m_{ij}}$$

for $1 \leq i, j \leq r$, where m_{ij} is the order of $s_i s_j$. In particular, $\langle e_i, e_j \rangle \leq 0$ for $i \neq j$. Consequently, C is a simplicial cone if (W, V) is essential.

PROOF: The last assertion follows from the criterion for a chamber to be simplicial given in §4D above, so we need only prove the first assertion. We may assume $i \neq j$ and, to simplify notation, $i = 1$ and $j = 2$. Let W' be the subgroup of W generated by s_1 and s_2 . Then (W', V) is a finite reflection group with $V^{W'} = H_1 \cap H_2 = (\mathbf{R}e_1 \oplus \mathbf{R}e_2)^\perp$. Hence (W', V) is

essentially 2-dimensional. In view of Example 2 above, the desired inner product formula will follow if we show that e_1 and e_2 are the canonical unit vectors associated to some W' -chamber C' in V .

Let $\mathcal{H}' \subseteq \mathcal{H}$ be the set of hyperplanes of the form $w'H_i$, with $i = 1$ or 2 and $w' \in W'$. Then \mathcal{H}' is W' -invariant, and the reflections with respect to the elements of \mathcal{H}' are in W' and generate it. Hence \mathcal{H}' is the set of W' -walls, i.e., the set of hyperplanes which define the W' -cells. Since C is convex and is disjoint from all the elements of \mathcal{H}' , it is contained in a W' -chamber C' . Moreover, H_1 and H_2 are walls of C' by Proposition 2 of §4B. [Note that $\overline{C'} \cap H_i \supseteq \overline{C} \cap H_i$, which has non-empty interior in H_i .] Finally, since $\langle e_i, - \rangle > 0$ on $C \subseteq C'$ for $i = 1, 2$, e_i is indeed the unit normal to H_i pointing toward C' . \square

5D The Coxeter matrix

We continue with the notation above, but we assume, in addition, that (W, V) is essential. Then $r = n = \dim V$ by the theorem [since the simplicial cone C has exactly n walls], and e_1, \dots, e_n form a basis for V . This fact, together with the inner product formula of the theorem, has the following important consequence:

Corollary. *Assume that (W, V) is essential. Then (W, V) is completely determined, up to isomorphism, by the $n \times n$ matrix $M = (m_{ij})$.*

PROOF: Given M , we can recover (W, V) as follows: V is isomorphic to \mathbf{R}^n , endowed with the inner product whose matrix is $A = (a_{ij})$, where $a_{ij} = -\cos(\pi/m_{ij})$. [This means that the inner products of the standard basis vectors e_i of \mathbf{R}^n are given by $\langle e_i, e_j \rangle = a_{ij}$.] And W can be identified with the group of automorphisms of \mathbf{R}^n generated by s_1, \dots, s_n , where s_i is the orthogonal reflection with respect to the hyperplane e_i^\perp . \square

The matrix M is called the *Coxeter matrix* associated to W . More precisely, it is associated to W together with a choice of “fundamental chamber” C and a numbering of the n walls of C .

EXERCISE

What happens if you change the choice of C ? [HINT: W acts simply-transitively on the chambers.]

For future reference, let's record the following explicit formula for s_i in terms of the inner product (and hence in terms of the Coxeter matrix):

$$s_i x = x - 2\langle e_i, x \rangle e_i.$$

If you haven't seen this before, it simply says that you reflect x with respect to e_i^\perp by subtracting off twice the component of x in the direction of e_i , thereby changing the sign of that component.

Remark. The Coxeter matrix has the following formal properties: It is a symmetric matrix of integers m_{ij} , with $m_{ii} = 1$ and $m_{ij} \geq 2$ for $i \neq j$. But not every such matrix can be the Coxeter matrix of a finite reflection group. A further necessary (and, as it turns out, sufficient) condition is that the associated matrix A defined in the proof above must be positive definite; for it is the matrix of a (positive definite) inner product. This fact, together with the corollary above, is the basis for the classification result stated in §3. If you look at the proof of that result in Bourbaki [16] or Grove–Benson [31], you’ll see that it consists of analyzing the possibilities for M , given that A is positive definite.

From the examples given in §5B above, we can easily write down some Coxeter matrices. For example, if W is of type A_n (symmetric group on $n + 1$ letters), then $m_{i,i+1} = 3$ and $m_{ij} = 2$ if $j > i + 1$. [Everything else is determined by the formal properties mentioned above.] Or if W is of type H_3 (dodecahedral group), then $m_{12} = 3$, $m_{23} = 5$, and $m_{13} = 2$.

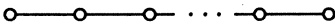
EXERCISE


Work out the Coxeter matrix of the reflection group of type B_n (signed symmetric group).

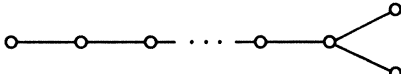
5E The Coxeter diagram

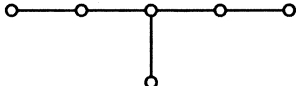
Instead of working directly with the matrix M , one usually works with a picture called the *Coxeter diagram*, which encodes all the information in M . This diagram is a graph, with vertices and edges, defined as follows: There are n vertices, one for each index $i = 1, \dots, n$, and the vertices corresponding to i and j are connected by an edge if and only if $m_{ij} \geq 3$. If $m_{ij} \geq 4$, then there is more than one convention in the literature as to how to indicate this in the diagram; the one we will follow is simply to label the edge with the number m_{ij} in this case. In summary, a labelled edge (with label necessarily at least 4) indicates the value of the corresponding m_{ij} ; an unlabelled edge indicates that $m_{ij} = 3$; and the lack of an edge joining i and j indicates that $m_{ij} = 2$.


The diagrams for all of the irreducible finite reflection groups are shown below. Based on the examples we have given, you should be able to check that the diagrams are correct for the cases A_n , B_n , G_2 , H_3 , and $I_2(m)$.

A_n  ($n \geq 1$ vertices)

B_n/C_n  ($n \geq 2$ vertices)

D_n  ($n \geq 4$ vertices)

E_6 



Remarks

1. Note that the diagrams which occur in this list are very special. For example, the graphs are all trees; there is very little branching in these trees; and the edge labels are rarely necessary (i.e., the numbers m_{ij} are rarely bigger than 3). One does not need the full force of the classification theorem in order to know these properties; in fact, these properties are the first few observations which occur in the *proof* of the classification theorem, as given in [16] or [31].

2. If you are familiar with the “Dynkin diagrams” which occur in Lie theory (cf. [16]), you will note that they are similar to Coxeter diagrams. But a Dynkin diagram contains slightly more information; in particular, it contains enough information to distinguish the root system of type B_n from that of type C_n , even though these root systems have the same Weyl group.

EXERCISE

Show that the essential finite reflection group (W, V) is irreducible if and only if the graph underlying its Coxeter diagram is connected. Deduce, in the reducible case, a canonical decomposition

$$(W, V) \approx (W_1 \times \cdots \times W_t, V_1 \oplus \cdots \oplus V_t)$$

into “irreducible components”, one for each connected component of the Coxeter diagram.

5F Fundamental domain and stabilizers

When studying the action of a group on a set, one wants to know how many orbits there are and what the stabilizers are at typical points of these orbits. Both of these questions have extremely simple answers in the case of W acting on V :

Theorem. *Let (W, V) be a finite reflection group, C a chamber, and S the set of reflections with respect to the walls of C . Then \overline{C} is a set of representatives for the W -orbits in V . Moreover, the stabilizer W_x of a point $x \in \overline{C}$ is the subgroup $\langle S_x \rangle$ generated by $S_x = \{s \in S : sx = x\}$. In particular, W_x fixes every point of \overline{A} , where A is the cell containing x .*

PROOF: Since W is transitive on the chambers, it is clear that every point of V is W -equivalent to a point of \overline{C} . Everything else in the theorem will follow if we prove: For $x, y \in \overline{C}$ and $w \in W$, if $wx = y$ then $x = y$ and $w \in \langle S_x \rangle$. We argue by induction on the smallest integer d such that w can be expressed as a word $s_1 \cdots s_d$ of length d in the generating set S of W .

If $d = 0$ there is nothing to prove, so assume $d > 0$ and choose an expression $w = s_1 \cdots s_d$ of minimal length. As we showed in the proof of the theorem in §5A above, C and wC are separated by the wall H_1 fixed by s_1 . We therefore have

$$wx = y \in \overline{C} \cap w\overline{C} \subseteq H_1.$$

So if we apply s_1 to both sides of the equation $wx = y$, we obtain

$$w'x = s_1y = y,$$

where $w' = s_1w = s_2 \cdots s_d$. By the induction hypothesis, it follows that $x = y$ [whence $s_1 \in S_x$] and that $w' \in \langle S_x \rangle$. So $w = s_1w'$ is also in $\langle S_x \rangle$, and the proof is complete. \square

5G The poset Σ as a simplicial complex

Assume that (W, V) is essential. Since the chambers are simplicial cones, one expects (i) that Σ is the poset of simplices of a simplicial complex; and (ii) that this simplicial complex triangulates the unit sphere in V . [We expect (ii) because our conical “cells” are in 1-1 correspondence with their intersections with the sphere, and a simplicial cone intersects the sphere in a subset homeomorphic to a simplex.] The assertions (i) and (ii) are Propositions 1 and 2 below. Before proceeding to these, however, the reader might find it helpful to look at the first few paragraphs of the appendix, where we explain our conventions regarding simplicial complexes.

Proposition 1. *The poset Σ is a simplicial complex.*

PROOF: According to the appendix, we must check two conditions, (a) and (b). Condition (a), that Σ have greatest lower bounds, has already been proved in §4E. As to Condition (b), concerning the poset $\Sigma_{\leq A}$ of faces of a cell $A \in \Sigma$, we know that A is a face of a chamber, so it suffices to consider the case where A is a chamber. But it is a trivial matter to compute the poset of faces of a simplicial cone, and this poset is indeed isomorphic to the set of subsets of $\{1, \dots, n\}$. \square

I chose to give this somewhat abstract proof of the proposition because I wanted to introduce the unorthodox terminology that I will be using regarding simplicial complexes. But it is easy to chase through the discussion in the appendix in order to describe in more conventional terms how Σ can be identified with an abstract simplicial complex (in which the simplices are certain finite subsets of a set of “vertices”):

Every 1-dimensional cell $A \in \Sigma$ is a ray $\mathbf{R}_+^* v$, where \mathbf{R}_+^* is the set of positive reals and v is a unit vector; the unit vectors v which arise in this way are the *vertices* of our simplicial complex. For each $(q+1)$ -dimensional cell $A \in \Sigma$ ($q \geq -1$), there is a q -simplex $\{v_0, \dots, v_q\}$ in our complex, where the v_i are the unit vectors on the 1-dimensional faces of A . It should be clear that we do indeed obtain a simplicial complex in this way and that Σ can be identified with the poset of simplices of this complex. Notice that we have allowed $q = -1$ above. The cell A is $\{0\}$ in this case, and it corresponds to the empty set of vertices. [Recall from the appendix our convention that the empty set is always included as a simplex of an abstract simplicial complex.]

Proposition 2. *The geometric realization $|\Sigma|$ is canonically homeomorphic to the unit sphere $S^{n-1} \subset V$.*

PROOF: Recall [cf. Appendix] that $|\Sigma|$ consists of certain convex combinations $\sum_v \lambda_v v$, where v ranges over the vertices of Σ , viewed as basis vectors of an abstract vector space. Now the vertices v_0, \dots, v_q of any $A \in \Sigma$ can also be viewed as unit vectors in V , and, as such, they are linearly independent. Hence we have a map $|\Sigma| \rightarrow V - \{0\}$, given by $\sum \lambda_v v \mapsto \sum \lambda_v v$. Composing this with radial projection, we obtain a continuous map $\phi : |\Sigma| \rightarrow S^{n-1}$. Since ϕ takes $|A| \subset |\Sigma|$ bijectively to $A \cap S^{n-1} \subset V$, it is bijective and therefore a homeomorphism (by compactness of $|\Sigma|$). \square

EXERCISE

Suppose W is the group of symmetries of a regular solid X . Make an intelligent guess [and prove it, if you can] as to how to describe Σ directly in terms of X . [HINT: The boundary of X is a topological sphere, decomposed into cells which are not necessarily simplices.]

5H A group-theoretic description of Σ

We close this chapter with one last observation, which will have far-reaching consequences. For simplicity, (W, V) will still be assumed essential.

We started the chapter with a “concrete” group W , given to us as a group of linear transformations (or, if you prefer to think more geometrically, as a group of isometries of Euclidean space, or, even better, as a group of isometries of a sphere). The geometry gave us (after we chose a chamber C), a set S of generators of W . The geometry also gave us a simplicial complex Σ , constructed by means of hyperplanes and half-spaces. We will prove

below, however, that if we forget the geometry and just view W as an abstract group (with a given set S of generators), then we can reconstruct Σ by pure group theory.

Let's look first at the subcomplex $\Sigma_{\leq C}$ consisting of the faces of C . To every face $A \leq C$, we associate its stabilizer $W_A = \{w \in W : wA = A\}$. In view of the theorem in §5F above, W_A can also be described as the stabilizer of any point $x \in A$, and, moreover, it fixes \bar{A} pointwise. That theorem also says that W_A is generated by a subset of our given generating set S . We will call a subgroup of W *special* if it is generated by a subset of S .

Thus we have a function ϕ from $\Sigma_{\leq C}$ to the set of special subgroups of W , and we will show that ϕ is a bijection. In fact, we can construct the inverse ψ of ϕ by taking fixed-point sets: Let W' be a special subgroup of W , generated by a set $S' \subseteq S$; then the fixed-point set of W' in \bar{C} is obtained by intersecting \bar{C} with the walls of C corresponding to the reflections in S' . So this fixed-point set is equal to \bar{A} for some $A \leq C$, and we can set $\psi(W') = A$. Using the stabilizer calculation in §5F, one can easily check that ψ is inverse to ϕ .

Note next that ϕ and its inverse ψ are order-reversing. For ψ , this is immediate from the definition. In the case of ϕ , the assertion follows from the fact that W_A fixes \bar{A} pointwise and hence stabilizes every face of A .

Thus we have a poset isomorphism

$$\Sigma_{\leq C} \approx (\text{special subgroups})^{\text{op}},$$


where “op” indicates that we are using the opposite of the usual order on the set of special subgroups. The picture below illustrates this isomorphism when $n = 3$. Here C is the cone over a triangle, and we have drawn a slice T of C (or, if you prefer, the intersection of C with the unit sphere). Almost every face of T has been labelled with its stabilizer, the one exception being the empty face (which is hard to see in the picture). The empty face corresponds to the cell $A = \{0\}$, which would appear in the picture if we drew the whole chamber C instead of just T . It is the smallest face of T , and its stabilizer is the largest special subgroup of W , namely, W itself. Similarly, the largest face is T itself, whose stabilizer is the smallest special subgroup (generated by $\emptyset \subset S$).

Returning now to the general case, we can use the W -action to extend our isomorphism to one from the whole poset Σ to the set of *special cosets* in W , i.e., the cosets wW' of special subgroups. Indeed, we can send a typical element $wA \in \Sigma$ ($w \in W$, $A \leq C$) to the coset wW_A . It is a routine matter to deduce the following result from what we did above for $\Sigma_{\leq C}$:

Theorem. *There is a poset isomorphism*

$$\Sigma \approx (\text{special cosets})^{\text{op}}$$

which is compatible with the W -action, where W acts on the special cosets by left-translation. \square



EXERCISE

Let W be the symmetric group on $n + 1$ letters and let S be the set $\{s_1, \dots, s_n\}$ of basic transpositions, where s_i interchanges i and $i + 1$. If you have done the exercise at the end of §5G, then you know (or at least suspect) that the complex Σ associated to W is the barycentric subdivision of the boundary of an n -simplex. Deduce this fact from the theorem above. [HINT: First figure out what the special subgroups are. Then look at the barycentric subdivision Δ of the boundary of an n -simplex. Its simplices are the “flags” $I_0 \subset \dots \subset I_q$ of non-empty proper subsets of $\{1, \dots, n + 1\}$, cf. Appendix, §B. Let σ be the simplex

$$\{1\} \subset \{1, 2\} \subset \dots \subset \{1, 2, \dots, n\}$$

of Δ . There is an obvious action of W on Δ , and the stabilizers of the faces of σ are precisely the special subgroups of W .]

Appendix. Abstract Simplicial Complexes

A Definitions

Recall that a *simplicial complex* with vertex set \mathcal{V} is a collection Δ of finite subsets of \mathcal{V} (called *simplices*) such that every singleton $\{v\}$ is a simplex and every subset of a simplex A is a simplex (called a *face* of A). The cardinality r of A is called the *rank* of A , and $r - 1$ is called the *dimension* of A . We include the empty set as a simplex; it has rank 0 and dimension -1 . A *subcomplex* of Δ is a subset Δ' which contains, for each of its elements A , all the faces of A ; thus Δ' is a simplicial complex in its own right, with vertex set equal to some subset of \mathcal{V} .

Note that Δ is a *poset*, ordered by the face relation. As a poset, it has the following formal properties:

- (a) Any two elements $A, B \in \Delta$ have a greatest lower bound $A \cap B$.
- (b) For any $A \in \Delta$, the poset $\Delta_{\leq A}$ of faces of A is isomorphic to the poset of subsets of $\{1, \dots, r\}$ for some $r \geq 0$.

Conversely, any poset Δ satisfying (a) and (b) can be identified with the poset of simplices of a simplicial complex. Namely, take the vertex set \mathcal{V} to

be the set of rank 1 elements of Δ [where the rank of A is defined to be the unique integer r such that (b) holds]; then we can associate to each $A \in \Delta$ the set $A' = \{v \in \mathcal{V} : v \leq A\}$. It is easy to check that $A \mapsto A'$ defines a poset isomorphism of Δ onto a simplicial complex with vertex set \mathcal{V} .

We will therefore extend the previous terminology and call any poset Δ satisfying (a) and (b) a *simplicial complex*. The elements of Δ will be called *simplices*, and those of rank 1 will be called *vertices*.

We visualize a simplex A of rank r as a geometric $(r - 1)$ -simplex, the convex hull of its r vertices. One makes this precise by forming the *geometric realization* $|\Delta|$ of Δ , which is a topological space partitioned into (open) simplices $|A|$, one for each non-empty $A \in \Delta$. To construct this topological space, start with an abstract real vector space with \mathcal{V} as a basis. Let $|A|$ be the interior of the simplex in this vector space spanned by the vertices of A , i.e., $|A|$ consists of the linear combinations $\sum_{v \in A} \lambda_v v$ with $\lambda_v > 0$ for all v and $\sum_{v \in A} \lambda_v = 1$. We then set

$$|\Delta| = \bigcup_{A \in \Delta} |A|.$$

If Δ is finite, then all of this is going on in \mathbf{R}^N , where N is the number of vertices of Δ , and we simply topologize $|\Delta|$ as a subspace of \mathbf{R}^N . The question of how to topologize $|\Delta|$ in the general case is more subtle, and we will not need to deal with it.

The purpose of the remainder of this appendix is to call attention to three special properties of the simplicial complex Σ associated to a reflection group: (i) Σ is a flag complex; (ii) Σ can be labelled; and (iii) Σ is determined by its associated chamber system. You can read it now (at least to find out what the terminology means), or you can wait and refer back to it when you need to.

B Flag complexes

Let P be a set with a binary relation called “incidence”, which is reflexive and symmetric. For example, P might consist of the points, lines, planes, etc., of a geometry (affine, projective, Euclidean, ...), with the usual notion of incidence. Or P might be a poset, with x and y incident if they are comparable (i.e., if $x \leq y$ or $y \leq x$). An important special case of this is the poset of cells of a cell complex (possibly simplicial), ordered by the face relation.

A *flag* in P is a set of pairwise incident elements of P ; if P is a poset, this is the same as a chain, i.e., a linearly ordered subset. The *flag complex* associated to P is the simplicial complex $\Delta(P)$ with P as vertex set and the finite flags as simplices. One example of this construction appears naturally in the foundations of the theory of simplicial complexes: If P is the poset of simplices of a simplicial complex, then $\Delta(P)$ is the *barycentric subdivision* of P . (We won’t make any use of this fact, except for motivational purposes. If you haven’t seen it before, you should draw some low dimensional

pictures to convince yourself that it's plausible.) The requirement that P be simplicial is not necessary here; P could just as easily be the poset of cells of a cube, or dodecahedron, or

Not all simplicial complexes are flag complexes. For example, the boundary of a triangle is not a flag complex. [If it were, then the three vertices would be pairwise incident, hence would form a flag, hence would be the vertices of a simplex of the complex.] The following proposition characterizes the flag complexes. A family of simplices in a simplicial complex Δ is called *joinable* if it has an upper bound in Δ ; in this case it has a least upper bound, which is just the set-theoretic union when Δ is viewed as a set of subsets of its vertex set.

Proposition. *The following conditions on a simplicial complex Δ are equivalent:*

- (1) Δ is a flag complex.
- (2) Every finite set of pairwise joinable simplices is joinable.
- (3) Every set of 3 pairwise joinable simplices is joinable.
- (4) Every finite set of pairwise joinable vertices is joinable.

PROOF: It is immediate that (1) \implies (2) \implies (4) \implies (1) and that (2) \implies (3). The proof that (3) \implies (2) is a straightforward induction and is left as an exercise. \square

We can now prove that the complex Σ associated to a reflection group is a flag complex. [This shouldn't be shocking, since every specific example we've worked out has in fact been a barycentric subdivision.] More generally, we will prove:

Proposition. *Let Σ be the poset of cells associated to a finite set \mathcal{H} of linear hyperplanes. Then Σ satisfies condition (2) above, i.e., every set of pairwise joinable elements of Σ is joinable.*

PROOF: No $H \in \mathcal{H}$ can strictly separate two joinable cells. So if we are given a family of pairwise joinable cells, then we can find for each $H \in \mathcal{H}$ a closed half-space U_H bounded by H which contains all the cells of the given family. Then $\bigcap_{H \in \mathcal{H}} U_H$ is a closed cell which contains all the given cells (cf. §4A, Proposition 3). Hence the corresponding open cell is an upper bound for the family. \square

C Labelled chamber complexes

Let Δ be a finite-dimensional simplicial complex. We will say that Δ is a *chamber complex* if all maximal simplices have the same dimension and any two can be connected by a gallery. [As before, a *gallery* is a sequence of maximal simplices in which any two consecutive ones are *adjacent*, i.e., have a common codimension 1 face.] The maximal simplices will then be called *chambers*.

A *labelling* of the chamber complex Δ by a set I is a function which assigns to each vertex an element of I , in such a way that the vertices of every chamber are mapped bijectively onto I . If Δ can be labelled, then the labelling is essentially unique: Any two labellings (say by sets I and I') differ by a bijection $I \approx I'$. (To see this, just note that if the labelling is known on a chamber C , then it is determined on any chamber adjacent to C .) You might find it helpful to think of a labelling as a “coloring” of the vertices. The number of colors used is required to be the rank of Δ (i.e., the number of vertices of any chamber), and joinable vertices are required to have different colors.

If Δ is a chamber complex which is a barycentric subdivision, then it can be labelled: Every vertex v of Δ corresponds to a cell of the original complex, and we may label v by the dimension of that cell. Similarly, if the chamber complex Δ is the flag complex of a geometry, then we can label its vertices by their type (point, line, ...). But not every chamber complex can be labelled. For example, the boundary of a triangle is a rank 2 chamber complex which cannot be labelled. More generally, the boundary of an m -gon can be labelled if and only if m is even.

It is useful to characterize labellability in terms of chamber maps. Recall first that a *simplicial map* from a simplicial complex Δ to a simplicial complex Δ' is a function ϕ from the vertices of Δ to those of Δ' which takes simplices to simplices. If the image $\phi(A)$ of a simplex A always has the same dimension as A , then ϕ is called *non-degenerate*. A non-degenerate simplicial map is the same as a poset map $\phi : \Delta \rightarrow \Delta'$ such that ϕ maps $\Delta_{\leq A}$ isomorphically to $\Delta'_{\leq \phi(A)}$ for every $A \in \Delta$. Finally, if Δ and Δ' are chamber complexes of the same dimension, then a simplicial map ϕ is non-degenerate if and only if it takes chambers to chambers; in this case ϕ will be called a *chamber map*.

[Note that a chamber map takes adjacent chambers to adjacent chambers and hence galleries to galleries. This would not be true if, in our definition of “adjacent”, we had required adjacent chambers to be distinct.]

An important special case which will arise fairly often in this book is the case where Δ' is a subcomplex of Δ and ϕ is the identity on Δ' ; a chamber map of this type is called a *retraction* of Δ onto Δ' .

Returning now to the question of labellability, let I be a set of cardinality n , where n is the rank of Δ , and let $\Delta(I)$ be the “simplex with vertex set I ”, i.e., the complex consisting of all subsets of I . Then a labelling of Δ by I is exactly the same as a chamber map $\lambda : \Delta \rightarrow \Delta(I)$. Thus Δ is labellable if and only if it admits a chamber map to $\Delta(I)$.

Given a labelling $\lambda : \Delta \rightarrow \Delta(I)$, we will often call $\lambda(A)$ for $A \in \Delta$ the *type* of A ; it is a subset of I . Note that the type of A is simply the set of labels of the vertices of A .

For any chamber C of Δ , our labelling λ maps the subcomplex $\Delta' = \Delta_{\leq C}$ generated by C isomorphically onto $\Delta(I)$; hence we may compose λ with the inverse isomorphism to get a retraction $\phi : \Delta \rightarrow \Delta'$. In concrete terms,

$\phi(A)$ is simply the unique face of C having the same type as A . Conversely, a retraction onto Δ' can be viewed as a labelling of Δ , with the set I of labels being the set of vertices of C . Thus we have another characterization of labellability: Δ is labellable if and only if it admits a retraction onto $\Delta_{\leq C}$.

Proposition. *The chamber complex Σ associated to a finite reflection group is labellable.*

PROOF: Choose a chamber C . Then we can define $\phi : \Sigma \rightarrow \Sigma_{\leq C}$ by letting $\phi(A)$ be the unique face of C which is W -equivalent to A (cf. §5F). It is easy to check that ϕ is a well-defined chamber map and a retraction. \square

D Chamber systems

We wish to show that if Δ is a sufficiently nice chamber complex, then Δ is completely determined by the system consisting of its chambers together with a suitable refinement of the adjacency relation. Thus we can forget about the vertices and, indeed, all the non-maximal simplices, when it is convenient to do so.

To refine the adjacency relation, we assume that Δ is labelled by a set I . Then any codimension 1 simplex of Δ has type $I - \{i\}$ for some $i \in I$. Given $i \in I$, two chambers of Δ will be called *i-adjacent* if they have the same face of type $I - \{i\}$. Note that this is an equivalence relation, unlike the ordinary adjacency relation. The *chamber system* associated to Δ is the set of chambers together with the relations of *i-adjacency*, one for each i .

In order to state conditions under which we can recover Δ from its chamber system, we need to recall some more terminology. The *link* of a simplex A , denoted $\text{lk } A$ or $\text{lk}_{\Delta} A$, is the subcomplex of Δ consisting of the simplices B which are disjoint from A [i.e., $A \cap B$ is the empty simplex] and joinable to A . As a poset, $\text{lk } A$ is isomorphic to the subposet $\Delta_{\geq A} \subseteq \Delta$ via $B \mapsto B \cup A$ ($B \in \text{lk } A$); this subposet, however, is not a subcomplex (unless A is the empty simplex).

Note that the maximal simplices of $\text{lk } A$ are in 1-1 correspondence with the chambers of Δ having A as a face. But $\text{lk } A$ need not be a chamber complex. For it might not be possible to connect two chambers in $\Delta_{\geq A}$ by a gallery in $\Delta_{\geq A}$.

Proposition. *Let Δ be a labelled chamber complex such that the link of every vertex is again a chamber complex. Then Δ is determined up to isomorphism by its chamber system.*

PROOF: Given $J \subseteq I$, call two chambers *J-equivalent* if they can be connected by a gallery C_0, \dots, C_d such that any two consecutive chambers C_{k-1} and C_k are *j-adjacent* for some $j \in J$. This equivalence relation is defined entirely in terms of the chamber system associated to Δ . We will be interested in the case where $J = I - \{i\}$ for some $i \in I$, in which case our hypothesis on links gives the following interpretation of the equivalence

relation: Two chambers are $(I - \{i\})$ -equivalent if and only if they have the same vertex of type i .

It is now easy to reconstruct Δ from its chamber system: Δ has one vertex of type i for each $(I - \{i\})$ -equivalence class of chambers; a collection of vertices forms a simplex if and only if the corresponding equivalence classes have a non-empty intersection. \square

This proposition applies to the chamber complex Σ associated to a finite reflection group in view of the following result:

Proposition. *Let Σ be the poset of conical cells associated to a finite set of hyperplanes. For any $A \in \Sigma$ and any chambers $C, D \in \Sigma_{\geq A}$, every minimal gallery joining C to D lies entirely in $\Sigma_{\geq A}$. In particular, if Σ is simplicial (and hence a chamber complex), then $\text{lk}_{\Sigma} A$ is a chamber complex.*

PROOF: Let $\Gamma : C = C_0, \dots, C_d = D$ be a minimal gallery. Then the walls H_1, \dots, H_d crossed by Γ separate C from D (cf. §4E, Proposition 4). For each $i = 1, \dots, d$, it follows that A is contained in both closed half-spaces bounded by H_i , hence $A \subseteq H_i$. Assuming inductively that $A \leq C_{i-1}$, we conclude that $A \subseteq \overline{C}_{i-1} \cap H_i = \overline{C}_{i-1} \cap \overline{C}_i$, hence $A \leq C_i$. \square

II

Abstract Reflection Groups

The result of §I.5H above suggests the possibility of introducing geometry into abstract group theory: Let W be a group, possibly infinite, generated by a subset S consisting of elements of order 2. Define, as in Chapter I, a *special coset* to be a coset $w\langle S'\rangle$ with $w \in W$ and $S' \subseteq S$. Now *define* $\Sigma = \Sigma(W, S)$ to be the poset of special cosets, ordered by the opposite of the inclusion relation: $B \leq A$ in Σ if and only if $B \supseteq A$ as subsets of W , in which case we say that B is a *face* of A .

In case W is a finite reflection group and S is the set of reflections with respect to the walls of some fixed chamber, we know that Σ is actually a simplicial complex, which triangulates a sphere of dimension $n - 1$ (where $n = \text{card } S$). Moreover, Σ possesses a rich geometric theory, with walls, half-spaces, etc., and the elements of S act on Σ as reflections. This raises some natural questions: In the general case, is Σ simplicial? Does it contain “walls” which divide it into “half-spaces”? Do the elements of S act on Σ as “reflections”?

Two things will result from our attempts to answer these questions. First, we will discover some facts about the combinatorial group theory of finite reflection groups. Second, we will discover a much larger class of groups W which deserve to be called “reflection groups”, namely, those groups for which the questions above all have affirmative answers. The study of these groups W and their associated complexes Σ was initiated by Tits [51]. He called the groups “Coxeter groups” (and the complexes “Coxeter complexes”) for historical reasons that we will explain in §4.

1 In Search of Axioms

Let W be a group which is generated by a subset S consisting of elements of order 2, and let Σ be the poset defined above. Our strategy in this section is very simple: We will attempt to define geometric notions, using Chapter I as a guide. Eventually, we will run into difficulty and have to introduce some hypotheses on (W, S) . The present section, then, will consist essentially of a list of definitions, culminating in an axiom (A) that (W, S) ought to satisfy in order to deserve the name “reflection group”.

First, we need something to play the role of the set \mathcal{H} of Chapter I. We will call an element of W a *reflection* if it is conjugate to an element of S .

[You should be able to convince yourself that this is the “right” definition by looking at the proof of statement (d) in §I.5A.] Now let \mathcal{H} be an abstract set in 1-1 correspondence with the set of reflections, this correspondence being denoted by $H \mapsto s_H$ for $H \in \mathcal{H}$. The elements of \mathcal{H} will be called *walls*. Since W acts by conjugation on the set of reflections, we may use our 1-1 correspondence to define an action of W on \mathcal{H} . By definition, then, we have

$$s_w H = w s_H w^{-1} \text{ for } w \in W \text{ and } H \in \mathcal{H}.$$

The maximal elements of Σ will be called *chambers*. They are the minimal special cosets, i.e., the singletons $\{w\} \subset W$. Note that W acts on Σ by left translation and that this action is simply-transitive on the chambers. We will set $C = \{1\}$ and call it the *fundamental chamber*. Thus a typical chamber $\{w\}$ can be written as wC .

The elements $A \in \Sigma$ which are not chambers and which are maximal among the non-chambers are said to have codimension 1. They are the two-element special cosets $w\langle s \rangle = \{w, ws\}$. Such an element A is a face of precisely two chambers, namely, $\{w\}$ and $\{ws\}$.

Two chambers $\{w\}, \{w'\}$ will be called *adjacent* if they have a common codimension 1 face. If $w \neq w'$, this is the same as saying that $w' = ws$ for some $s \in S$. [We note in passing that we can refine the notion of adjacency by saying, for any $s \in S$, that two chambers $\{w\}, \{w'\}$ are *s-adjacent* if $w' = w$ or ws . The reader who has read about chamber systems in §D of the appendix to Chapter I will not be surprised by the existence of a family of adjacency relations, one for each $s \in S$.]

Now that we have an adjacency relation, we can define the notions of *gallery* and *combinatorial distance* exactly as in Chapter I. It follows at once that non-stuttering galleries $(C_i)_{0 \leq i \leq d}$ with $C_0 = C$ are in 1-1 correspondence with sequences s_1, \dots, s_d of elements of S , the correspondence being given by $C_i = \{w_i\}$, where $w_i = s_1 \cdots s_i$. Let H_i be the wall such that $s_i = s_{H_i}$; then the walls $w_{i-1}H_i$ ($i = 1, \dots, d$) will be called the *walls crossed* by the given gallery. The motivation for this can be found in the proof of statement (e) in §I.5A.

It is clear from these remarks that $d(C, wC)$ is the smallest integer d such that w can be expressed as a word $s_1 \cdots s_d$ in the generating set S . This minimal d is called the *length* of w with respect to S and is denoted $l(w)$ or $l_S(w)$. Thus we have

$$d(C, wC) = l(w).$$

Since combinatorial distance is invariant under the W -action, we obtain the following group-theoretic interpretation of the distance function:

$$d(wC, w'C) = d(C, w^{-1}w'C) = l(w^{-1}w').$$

Encouraged by our success so far, we move on to half-spaces. For each $H \in \mathcal{H}$ we would like to partition the chambers into two “halves” $U_+(H)$

and $U_-(H)$, where $U_+(H)$ is the half containing the fundamental chamber C . Thus the set of “half-spaces” should be in 1-1 correspondence with the set $\mathcal{H} \times \{\pm 1\}$.

We expect the W -action to take half-spaces to half-spaces and hence to induce an action of W on the set $\mathcal{H} \times \{\pm 1\}$. This action should have the form

$$w \cdot (H, \varepsilon) = (wH, \pm \varepsilon),$$

where the ambiguous sign is $+$ if and only if $wC \in U_+(wH)$. In other words, the sign should be $+$ if and only if wC and C are on the “same side” of the wall wH . Applying the action of w^{-1} , the condition for the sign to be $+$ becomes: “ C and $w^{-1}C$ are not separated by H ”. Thus the action of W on $\mathcal{H} \times \{\pm 1\}$ should satisfy:

$$w \cdot (H, \varepsilon) = (wH, -\varepsilon) \iff H \text{ separates } C \text{ from } w^{-1}C. \quad (*)$$

Apply this now to the case where $w = s \in S$. We expect there to be a unique $H \in \mathcal{H}$ which separates C from sC , namely, the H such that $s = s_H$. Hence s should act on $\mathcal{H} \times \{\pm 1\}$ as the involution ρ_s defined as follows:

$$\rho_s(H, \varepsilon) = \begin{cases} (H, -\varepsilon) & \text{if } s = s_H \\ (sH, \varepsilon) & \text{otherwise.} \end{cases}$$

Thus we have arrived at an easily stated condition that (W, S) ought to satisfy if it is to behave like a reflection group. We will call this condition **(A)** for “action”:

(A) *There is an action of W on the set $\mathcal{H} \times \{\pm 1\}$ such that, for every $s \in S$, s acts as the involution ρ_s defined above.*

It would be possible at this point to show that condition **(A)** is sufficient to enable one to give a reasonable definition of the desired half-spaces $U_{\pm}(H)$. But I prefer to postpone this until the next chapter, when we will have a conceptual definition of the notion of “half-space” in a simplicial complex. I would like to prove *something*, however, to convince you that the harmless-looking condition **(A)** is much more powerful than it appears. The following theorem should serve this purpose. It says, roughly speaking, that if **(A)** holds then the set of walls crossed by a gallery has the properties that Chapter I has led us to expect.

Theorem. *Suppose (W, S) satisfies **(A)**. Then one can associate to each $w \in W$ a finite subset $\mathcal{H}(w) \subseteq \mathcal{H}$ with the following properties:*

- (1) $\text{card } \mathcal{H}(w) = d(C, wC) = l(w)$.
- (2) *If Γ is a minimal gallery from C to wC , then the walls crossed by Γ are distinct and are precisely the elements of $\mathcal{H}(w)$.*
- (3) *Let Γ be an arbitrary non-stuttering gallery from C to wC . For any $H \in \mathcal{H}$, one has $H \in \mathcal{H}(w)$ if and only if H is crossed by Γ an odd number of times.*

PROOF: It should be clear that $\mathcal{H}(w)$ is supposed to be the set of walls which “separate C from wC ”. [The quotation marks are intended as a reminder that the phrase they enclose only has intuitive meaning at the moment, since we have not yet constructed the half-spaces.] Motivated by (*), we define $\mathcal{H}(w)$ to be the set of $H \in \mathcal{H}$ such that $w^{-1} \cdot (H, 1) = (w^{-1}H, -1)$.

Suppose now that $w = s_1 \cdots s_d$ with $s_i \in S$, and let Γ be the corresponding gallery $(w_i C)_{0 \leq i \leq d}$, where $w_i = s_1 \cdots s_i$. Since $w^{-1} = s_d \cdots s_1$, we can compute $w^{-1} \cdot (H, 1)$ by first applying s_1 , then applying s_2 , etc. After applying s_1, \dots, s_{i-1} , we will have an element of the form $(w_{i-1}^{-1}H, \varepsilon)$, to which we must apply s_i . In view of the definition of ρ_{s_i} , the application of s_i will change ε to $-\varepsilon$ if and only if $w_{i-1}^{-1}H = H_i$, where H_i is the wall such that $s_i = s_{H_i}$. Hence $w^{-1} \cdot (H, 1) = (w^{-1}H, (-1)^p)$, where p is the number of i such that $H = w_{i-1}H_i$.

By the definition of $\mathcal{H}(w)$, we have $H \in \mathcal{H}(w)$ if and only if p is odd. On the other hand, $w_{i-1}H_i$ is (by definition again) the i th wall crossed by Γ , so p is the number of times that Γ crosses H . This proves (3). (1) and (2) follow immediately from (3) and the following lemma, which does not require condition (A):

Lemma. *For any $w \in W$ and any minimal gallery Γ from C to wC , the walls crossed by Γ are distinct.*

PROOF: Let Γ correspond to a sequence of generators s_1, \dots, s_d . Then the minimality of Γ implies that $w = s_1 \cdots s_d$ has length d . Suppose that the walls crossed by Γ are not distinct. Then we have, with the notation above, $w_{i-1}H_i = w_{j-1}H_j$ for some $i < j$. Passing to the associated reflections, this becomes

$$w_{i-1}s_iw_{i-1}^{-1} = w_{j-1}s_jw_{j-1}^{-1},$$

or

$$w_iw_{i-1}^{-1} = w_jw_{j-1}^{-1},$$

or

$$w_{i-1}^{-1}w_{j-1} = w_i^{-1}w_j.$$

In terms of the s 's, this last equation says

$$s_i \cdots s_{j-1} = s_{i+1} \cdots s_j,$$

which implies that

$$w = s_1 \cdots s_d = s_1 \cdots \hat{s}_i \cdots \hat{s}_j \cdots s_d,$$

where the hats indicate deleted letters. This contradicts the fact that $d = l(w)$. \square

As a corollary of this proof, we can deduce an incredible consequence of condition (A), which we alluded to in the case of finite reflection groups in a remark at the end of §I.5A. In order to state it, we introduce a second condition which a general pair (W, S) may or may not satisfy:

(D) If $w = s_1 \cdots s_d$ with $d > l(w)$, then there are indices $i < j$ such that $w = s_1 \cdots \hat{s}_i \cdots \hat{s}_j \cdots s_d$.

We will call this the “deletion condition”. As a corollary of the theorem and the proof of the lemma, we have:

Corollary. *If (W, S) satisfies (A), then it satisfies (D).*

PROOF: Suppose $w = s_1 \cdots s_d$ with $d > l(w)$, and consider the corresponding gallery Γ from C to wC . Then Γ cannot cross d distinct walls, since this would imply (by part (3) of the theorem) $\text{card } \mathcal{H}(w) = d$, contradicting part (1). Hence there must be repetition among the walls crossed by Γ ; the proof of the lemma now shows that we can delete two letters s_i, s_j . \square

Remark. You wouldn’t know it from the roundabout proof of the corollary, but condition (D) actually has a direct geometric interpretation in terms of galleries. We will see this in the next chapter (§III.4B, statement (d) and its proof).

2 Examples

2A Finite reflection groups

It is essentially obvious that a finite reflection group as in Chapter I satisfies (A) (where S is the set of reflections in the walls of some fixed chamber). Indeed, we arrived at the formulation of (A) by writing down properties which finite reflection groups were known to satisfy.

Conversely, it is true (but *not* obvious) that every finite group satisfying (A) is a finite reflection group. We will prove this in §5 below. In view of this fact, our remaining examples will necessarily be infinite groups.

2B The infinite dihedral group

Let W be the *infinite dihedral group* D_∞ . By definition, this is the group defined by the presentation


$$W = \langle s, t ; s^2 = t^2 = 1 \rangle.$$

(In case you are not familiar with this notation for group presentations, it simply means that you start with the free group $F = F(s, t)$ on two generators s, t and then divide out by the smallest normal subgroup containing s^2 and t^2 .) Note that the finite dihedral groups D_{2m} are quotients of W . It follows that the images of s and t in W are distinct and non-trivial, so no confusion will result if we denote these images by s and t . It also follows that st has infinite order and hence that W is infinite.

Let $S = \{s, t\} \subset W$. We will explain from three different points of view why (W, S) satisfies (A).

(i) *Combinatorial group theory.* The definition of W via the presentation above makes it easy to define homomorphisms from W to another group. One need only specify two elements of the target group whose squares are trivial, and there is then a homomorphism taking s and t to these elements. In particular, if we want W to act on some set, it suffices to specify involutions ρ_s and ρ_t of that set, and then we can make s and t act as ρ_s and ρ_t , respectively. Condition (A) is now evident.

(ii) *Euclidean geometry.* We make W act as a group of isometries of the real line L by letting s act as the reflection about 0 ($x \mapsto -x$) and t act as the reflection about 1 ($x \mapsto 2 - x$). Note, then, that W acts as a group of *affine* transformations $x \mapsto ax + b$. This action has an associated “chamber geometry”, entirely analogous to what we saw in Chapter I for finite (linear) reflection groups. It is illustrated in the following picture, where C denotes the open unit interval:



The vertices in the picture are the integers. They are shown in two “colors”, \circ and \bullet , to indicate the two orbits under the action of W .

It is now easy to identify our abstract set \mathcal{H} with the set of integers and to identify our abstract $\mathcal{H} \times \{\pm 1\}$ with the set of half-lines whose endpoint is an integer. The action of W on L induces an action of W on this set of half-lines, and condition (A) follows easily.


(iii) *Linear algebra.* There is a standard method for “linearizing” affine things by embedding the affine space in question as an affine hyperplane (i.e., a translate of a linear hyperplane) in a vector space of one higher dimension. In the present case, we do this by identifying the line L above with the affine line $y = 1$ in the plane $V = \mathbb{R}^2$. The affine action of W on L extends to a linear action of W on V . Explicitly, since we want $s(x, 1) = (-x, 1)$, we can set $s(x, y) = (-x, y)$; in other words, we can make s act via the matrix

$$\begin{pmatrix} -1 & 0 \\ 0 & 1 \end{pmatrix}.$$

Similarly, to make $t(x, 1) = (2 - x, 1)$, we can set $t(x, y) = (2y - x, y)$; thus t acts via the matrix

$$\begin{pmatrix} -1 & 2 \\ 0 & 1 \end{pmatrix}.$$

The picture of W acting on V is shown below. It is simply the cone over the picture of W acting on L . (C now denotes the *cone* over the unit interval in the line $y = 1$.) We may identify \mathcal{H} with the set of walls of the chambers shown in the picture; these walls are linear hyperplanes in V . And we may identify $\mathcal{H} \times \{\pm 1\}$ with the set of half-planes determined by these walls. Condition (A) now follows easily from the action of W on these half-planes.



Before leaving this example, let's compare this situation with that of Chapter I. As in Chapter I, s and t act as linear reflections on V , provided we interpret this term suitably: If V is a real vector space, not necessarily endowed with an inner product, then a *linear reflection* on V is a linear map which is the identity on some (linear) hyperplane H and is multiplication by -1 on some complement of H , i.e., some 1-dimensional subspace H' such that $V = H \oplus H'$. The reflections considered in Chapter I, where V has an inner product and $H' = H^\perp$, will be called *orthogonal* reflections from now on to distinguish them from the more general linear reflections that we have just defined. Note that a linear reflection is not uniquely determined by its hyperplane H of fixed-points.

In the present example it is still true, as in Chapter I, that W is generated by linear reflections whose associated hyperplanes are the two walls of our "fundamental chamber" C . And it is still true that \overline{C} is a fundamental domain for the action of W on $\bigcup_{w \in W} w\overline{C}$. But this union is not the whole vector space V . It is, rather, the convex cone consisting of the open upper half-plane together with the origin.

Finally, we wish to consider the question of how we might have discovered this representation of W as a "linear reflection group" if we had not had the geometry provided by (ii) as a guide. Our discussion of this will be long-winded and will require the results of the following exercise:

EXERCISE

- (a) If s is a linear reflection on a 2-dimensional vector space V , show that the only s -invariant affine lines not passing through the origin are those parallel to the (-1) -eigenspace.
- (b) Deduce that two linear reflections s, t of V have a common invariant line not passing through the origin if and only if they have the same (-1) -eigenspace.
- (c) Suppose s and t have the same $(+1)$ -eigenspace. Show that the induced reflections s^* and t^* of $V^* = \text{Hom}(V, \mathbf{R})$ have the same (-1) -eigenspace.

Returning now to $W = D_\infty$, suppose we try to construct a linear representation of it by imitating the procedure used in §I.5D. Thus we now write s_1 and s_2 instead of s and t , and we introduce the Coxeter matrix

$M = (m_{ij})$, where m_{ij} is the order of $s_i s_j$. The resulting matrix, then, is

$$M = \begin{pmatrix} 1 & \infty \\ \infty & 1 \end{pmatrix}.$$

The corresponding Coxeter diagram is $\circ \overset{\infty}{\text{---}} \circ$.

Even though some of the m_{ij} are infinite, we can still make sense out of the matrix $A = (a_{ij})$, where $a_{ij} = -\cos(\pi/m_{ij})$; this is given in the present case by

$$A = \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}.$$

Let B be the associated bilinear form on \mathbf{R}^2 (with $B(e_i, e_j) = a_{ij}$). For reasons that will be clear shortly, we will write V' for \mathbf{R}^2 endowed with this bilinear form. Note that B is *degenerate*, by which we mean that its matrix is singular. In particular, it cannot possibly be an inner product. Nevertheless, we can still define reflections s'_1, s'_2 on V' by imitating the formula which we wrote down in §I.5D:

$$s'_i x = x - 2B(e_i, x)e_i.$$

This is indeed a linear reflection. For it takes e_i to $-e_i$ and it is the identity on the hyperplane $e_i^\perp = \{x \in V' : B(e_i, x) = 0\}$.

Finally, note that our two reflections have different (-1) -eigenspaces but the same $(+1)$ -eigenspace. In fact, the linear functions $B(e_i, -)$ for $i = 1, 2$ are negatives of one another since the two columns of A are negatives of one another, hence $e_1^\perp = e_2^\perp$. In view of the exercise above, then, this linear action of W is *not* the linearization of an affine action on a line. But the exercise also shows that if we pass to the dual space V of V' , then the induced reflections $s_i = (s'_i)^*$ do have the same (-1) -eigenspace and hence an invariant affine line L . It is not hard to continue the analysis and show that the resulting W -action on V has a “chamber geometry” like that pictured at the beginning of this discussion.

To summarize, we have a linear representation of W on a vector space V , in which there is a nice chamber decomposition not too different from what we saw for finite reflection groups. And the dual representation (W acting on V') has nice algebraic formulas not too different from what we saw for finite reflection groups. As we will see in §5 below, this is the typical situation for the “reflection groups” that we will be discussing. [You might wonder why, if this is supposed to be typical, we did not notice it in Chapter I. In that context, we had both the chamber geometry and the algebraic formulas in the *same* vector space V . The answer, briefly, is that we had a W -invariant non-degenerate bilinear form on V in Chapter I, namely, the inner product $\langle -, - \rangle$; and one can identify V with its dual whenever there is such a form.]

2C The group $\text{PGL}_2(\mathbf{Z})$

Let $\text{GL}_2(\mathbf{Z})$ be the group of 2×2 invertible matrices over the ring \mathbf{Z} of

integers. Let $\mathrm{PGL}_2(\mathbf{Z})$ be the quotient of $\mathrm{GL}_2(\mathbf{Z})$ by the central subgroup of order two generated by -1 (= the negative of the identity matrix). Thus $\mathrm{PGL}_2(\mathbf{Z})$ is obtained from $\mathrm{GL}_2(\mathbf{Z})$ by identifying a matrix with its negative. We denote a typical element of $\mathrm{GL}_2(\mathbf{Z})$ by

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix}$$

and its image in $\mathrm{PGL}_2(\mathbf{Z})$ by

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix}.$$

It is easy to check that the group $W = \mathrm{PGL}_2(\mathbf{Z})$ is generated by the set $S = \{s_1, s_2, s_3\}$ of elements of order 2 defined by

$$s_1 = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \quad s_2 = \begin{bmatrix} -1 & 1 \\ 0 & 1 \end{bmatrix} \quad s_3 = \begin{bmatrix} -1 & 0 \\ 0 & 1 \end{bmatrix}.$$

(You can see this by thinking about elementary row operations.) We will now show that condition **(A)** is satisfied. We will use three different methods, analogous to those used for D_∞ . In each of the three cases, however, we will have to use one or more non-trivial facts that will be stated without proof. If you are not familiar with these facts, you are advised to just read the discussion casually, getting whatever you can out of it.

(i) *Combinatorial group theory.* A simple computation shows that the products $s_1 s_2$, $s_1 s_3$, and $s_2 s_3$ have orders 3, 2, and ∞ , respectively. It is also true (but not obvious) that W admits a presentation in which the defining relations simply specify the orders of the pairwise products:

$$W = \langle s_1, s_2, s_3 ; s_1^2 = s_2^2 = s_3^2 = (s_1 s_2)^3 = (s_1 s_3)^2 = 1 \rangle.$$

[Some readers will be familiar with the fact that W has a subgroup $\mathrm{PSL}_2(\mathbf{Z})$ of index 2 which admits a presentation $\langle u, v ; u^3 = v^2 = 1 \rangle$, cf. [46], §I.4.2, or [35], §§IV.5H and VII.2F. It is not too hard to deduce the presentation for W stated above from this presentation for $\mathrm{PSL}_2(\mathbf{Z})$.]

To verify **(A)**, now, we simply have to check that the involutions $\rho_i = \rho_{s_i}$ which occur in the statement of **(A)** satisfy the defining relations for W . Consider, for instance, the relation $(\rho_1 \rho_2)^3 = 1$. Let $S' = \{s_1, s_2\}$ and let W' be the dihedral group of order 6 generated by S' . The reflections in W' (i.e., the W' -conjugates of s_1 and s_2), form a subset of the reflections in W ; hence the set \mathcal{H}' of W' -walls is a subset of the set \mathcal{H} of W -walls.

Suppose, now, that we apply $(\rho_1 \rho_2)^3$ to $(H, \varepsilon) \in \mathcal{H} \times \{\pm 1\}$. Clearly the only thing we have to worry about is the possibility of sign changes in the second factor as we successively apply the ρ_i . But no sign changes will ever occur unless $w'H$ is in \mathcal{H}' for some $w' \in W'$, in which case we have $H \in \mathcal{H}'$. Thus we are reduced to showing that $(\rho_1 \rho_2)^3$ is the identity on $\mathcal{H}' \times \{\pm 1\}$, which follows from the fact that W' is a finite reflection group and hence is already known to satisfy **(A)**. [Alternatively, we could complete the proof by doing an easy computation in the dihedral group D_6 .]

(ii) *Hyperbolic geometry.* There is a famous tessellation of the hyperbolic plane by ideal hyperbolic triangles (i.e., triangles having their vertices on the circle at infinity). Figure 1 below shows this tessellation in the unit disk model of the hyperbolic plane.




Figure 1

The group of symmetries of this tessellation is a group of hyperbolic isometries generated by reflections, and it is, in fact, precisely the group W . In order to explain this in slightly more detail, we switch to the upper half-plane model of the hyperbolic plane. This is shown in Figure 2, which is a picture of the barycentric subdivision of the tessellation in Figure 1. In order to relate the two models of the hyperbolic plane, you should think of the vertices of the big triangle in Figure 1 as corresponding to the points 0, 1, and ∞ in Figure 2. The barycenter of this big triangle is shown as a heavy dot in Figure 2.

The action of W on the upper half-plane is given by

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix} \cdot z = \begin{cases} \frac{az+b}{cz+d} & \text{if } ad-bc=1 \\ \frac{a\bar{z}+b}{c\bar{z}+d} & \text{if } ad-bc=-1 \end{cases}$$

where \bar{z} is the complex conjugate of z . You may be familiar with this action restricted to $\text{PSL}_2(\mathbf{Z})$, where, of course, complex conjugation does not arise. Complex conjugation is necessary for the full group W , however, because elements of negative determinant acting by linear fractional transformations interchange the upper and lower half-planes.

Now under this action, the generating set S of W is the set of reflections in the three sides of one of the “chambers” C , as indicated in Figure 2.




Figure 2

Moreover, it is known that \overline{C} is a fundamental domain for the action of W . A proof of this can be found in almost any book that discusses modular forms. [Actually it is more likely that the analogous fact about $\mathrm{PSL}_2(\mathbf{Z})$ is proved: $\overline{C} \cup s_3 \overline{C}$ is a fundamental domain for this group. See, for instance, [44], §VII.1.2, or [35], §IV.5H.]

If you have followed all of this, then you can probably complete the geometric proof that (W, S) satisfies condition (A). Just identify $\mathcal{H} \times \{\pm 1\}$ with the set of hyperbolic half-planes determined by the hyperbolic lines in Figure 2, and use the action of W on these half-planes.

(iii) *Linear algebra.* As was the case with the group D_∞ , the linear algebra approach will take the longest to explain. But it is quite instructive and worth at least reading through, even if you don't check all the details. It is based on a 3-dimensional linear representation of W which has been studied extensively, starting with Gauss.

The vector space V on which W acts is the space of real quadratic forms q in two variables, i.e., the space of functions $q : \mathbf{R}^2 \rightarrow \mathbf{R}$ given by $q(x) = ax_1^2 + 2bx_1x_2 + cx_2^2$. Note that we can also write $q(x) = \beta(x, x)$, where β is the bilinear form on \mathbf{R}^2 with matrix

$$A = \begin{pmatrix} a & b \\ b & c \end{pmatrix}.$$

Thus we can, when it is convenient, identify V with the space of symmetric bilinear forms on \mathbf{R}^2 , or, equivalently, with the space of real symmetric 2×2 matrices.

The group $G = \text{GL}_2(\mathbf{R})$ acts on V by

$$(g \cdot q)(x) = q(xg)$$

for $g \in G$, $q \in V$, and $x \in \mathbf{R}^2$, where x is viewed as a row vector on the right-hand side of the equation. This action is said to be by *change of variable*, since $g \cdot q$ is obtained from q by replacing x_1 and x_2 by linear functions of x_1 and x_2 (with coefficients given by the columns of g). In terms of the symmetric matrix A corresponding to q , the action of g is given by $A \mapsto gAg^t$, where g^t is the transpose of g .

The elements $q \in V$ fall into exactly six orbits under the action of G . First, there are three types of non-degenerate forms: positive definite (G -equivalent to $x_1^2 + x_2^2$); negative definite (G -equivalent to $-x_1^2 - x_2^2$); and indefinite (G -equivalent to $x_1^2 - x_2^2$). Next, there are the non-zero degenerate forms, which are either positive semi-definite (G -equivalent to x_1^2) or negative semi-definite (G -equivalent to $-x_1^2$). And, finally, there is the zero form.

It is easy to visualize this partition of V into G -orbits. Let $Q : V \rightarrow \mathbf{R}$ be given by

$$Q(q) = -\det A = b^2 - ac,$$

where A is the matrix corresponding to q as above. (Thus Q is a quadratic form on the 3-dimensional space V of quadratic forms.) Then the degenerate forms q are the points of the quadric surface $Q = 0$ in V . If we introduce new coordinates x, y, z in V by setting

$$\begin{aligned} b &= x \\ a &= z + y \\ c &= z - y, \end{aligned}$$

then Q becomes $x^2 + y^2 - z^2$, so the quadric surface of degenerate forms is the double cone $z^2 = x^2 + y^2$. [Draw a picture!] The exterior of the cone is given by $Q > 0$ and consists of the indefinite forms. And the interior $Q < 0$ has two components, the upper half ($z > 0$), consisting of the positive definite forms, and the lower half, consisting of the negative definite forms.

The action of G on V is really an action of the quotient $G/\{\pm 1\}$, so we may restrict the action to $W = \text{PGL}_2(\mathbf{Z}) \subset G/\{\pm 1\}$. This is the desired 3-dimensional representation of W . Here are the basic facts about this representation:

First, the W -action leaves the form Q invariant, i.e., $Q(wq) = Q(q)$ for $w \in W$ and $q \in V$. This follows from the fact that every $g \in \text{GL}_2(\mathbf{Z})$ has $\det g = \pm 1$, so that

$$\det gAg^t = \det^2 g \cdot \det A = \det A$$

for any symmetric 2×2 matrix A . So W also leaves invariant the symmetric bilinear form B on V such that $Q(q) = B(q, q)$. One can easily compute B explicitly; in terms of symmetric matrices, we have

$$B(A, A') = bb' - \frac{1}{2}(ac' + a'c),$$

where

$$A = \begin{pmatrix} a & b \\ b & c \end{pmatrix} \quad \text{and} \quad A' = \begin{pmatrix} a' & b' \\ b' & c' \end{pmatrix}.$$

The next observation is that the generators s_i of W act on V as linear reflections. In fact, if you compute the (± 1) -eigenspaces of s_i , you find that s_i has a 1-dimensional (-1) -eigenspace $\mathbf{R}e_i$ and that s_i fixes the hyperplane $H_i = e_i^\perp$ (where e_i^\perp is defined with respect to our bilinear form $B(-, -)$). One can take the e_i , which are determined up to scalar multiplication, to be the following symmetric matrices:

$$e_1 = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \quad e_2 = \begin{pmatrix} -1 & -1 \\ -1 & 0 \end{pmatrix} \quad e_3 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}.$$

And the fixed hyperplanes H_i are given, respectively, by $a = c$, $c = 2b$, and $b = 0$.

I chose the eigenvectors e_i above so that they would satisfy $Q(e_i) = 1$; this determines them up to sign. It then follows as a formal consequence that the reflections s_i are given by the usual formula:

$$s_i q = q - 2B(e_i, q)e_i;$$

for the map defined by this formula is the identity on e_i^\perp and sends e_i to $-e_i$.

We now focus on the action of W on the cone P of positive definite forms, and we look for a fundamental domain for this action. Concretely, this means that we are looking for canonical forms for the positive definite q 's under integral change of variable. Gauss found the following fundamental domain. Let C be the simplicial cone in V defined by the inequalities $a > c > 2b > 0$. Then $C \subset P$, and \overline{C} is (more or less) a fundamental domain for the action of W on P .

The “more or less” here refers to the fact that \overline{C} touches the boundary of P . For if you compute the vertices of C (i.e., the rays which are 1-dimensional faces of C), you find that they are represented by the forms x_1^2 , $x_1^2 + x_2^2$, and $x_1^2 + x_1x_2 + x_2^2$, the first of which is degenerate. So the correct statement is the following: Let U be the convex cone in V consisting of the positive definite forms together with the forms $\lambda(ax_1 + bx_2)^2$ with $\lambda \geq 0$ and $a, b \in \mathbf{Z}$. Then $U = \bigcup_{w \in W} w\overline{C}$, and \overline{C} is a fundamental domain for the action of W on U ; moreover, the open simplicial cones wC are disjoint from one another.

You have probably noticed that the walls of C are precisely the fixed hyperplanes H_i of the reflections s_i . So we have, once again, the usual sort of chamber geometry, and it is possible to verify condition (A) by identifying $\mathcal{H} \times \{\pm 1\}$ with the half-spaces in V determined by the walls of the chambers wC . Details are omitted.

One final comment: I normalized the e_i above so that we would have $B(e_i, -) > 0$ on C . In view of Chapter I and the infinite dihedral group

example, it is therefore to be expected that

$$B(e_i, e_j) = -\cos \frac{\pi}{m_{ij}},$$

where m_{ij} is the order of $s_i s_j$. This is indeed the case, as direct computation shows. Thus our representation of W acting on V is what we should now be ready to call the “canonical linear representation” of W . Note also, for future reference, that the bilinear form B in this example is non-degenerate, although not positive definite. Indeed, we showed above that Q could be written as $x^2 + y^2 - z^2$ after a change of coordinates in V , so B has signature $(2, 1)$ [2 plus signs and 1 minus sign].

EXERCISE

What is the connection between the points of view in (ii) and (iii)? [HINT: The upper sheet of the hyperboloid $Q = -1$ is one of the standard models of the hyperbolic plane; it is contained in the positive definite cone P and it cuts across the chambers wC . Incidentally, the degenerate forms $\lambda(ax_1 + bx_2)^2$ which we adjoined to P also appeared in the discussion (or at least the pictures) in (ii): They correspond to the cusps in Figures 1 and 2.]

3 Consequences of the Deletion Condition

We return now to the general theory, which is much easier than the examples. Thus W is an arbitrary group with a set S of generators of order 2. We saw at the end of §1 that if the “action condition” (A) holds then so does the deletion condition (D). Now this deletion condition must certainly look surprising to you if you have not seen it before, so we will spend a little time exploring its consequences. We begin by giving a couple of reformulations of it.

3A Equivalent forms of (D)

We will need to formalize the concept of “word”, which we have already used informally. By a *word* in the generating set S we mean a sequence $s = (s_1, \dots, s_d)$ of elements of S . We will often be less formal and simply say that the “expression” $s_1 \cdots s_d$ is a word; but we must distinguish a word s from the *element* $w = s_1 \cdots s_d \in W$ that it represents. Whenever there is danger of confusion, we will be more precise and revert to the sequence notation (s_1, \dots, s_d) .

The word (s_1, \dots, s_d) is called *reduced* if the corresponding element $w = s_1 \cdots s_d$ has length $l(w) = d$, i.e., if it cannot be represented by a shorter word. We will also say, in this situation, that the given word is a *reduced decomposition* of w , or, less formally, that the equation $w = s_1 \cdots s_d$ is a reduced decomposition of w .

We can now state the first consequence of condition (D). It is called the *exchange condition*:

(E) Given $w \in W$, $s \in S$, and any reduced decomposition $w = s_1 \cdots s_d$ of w , either $l(sw) = d + 1$ or else there is an index i such that

$$w = ss_1 \cdots \hat{s}_i \cdots s_d.$$

The proof that (D) implies (E) is immediate. For if $l(sw) < d + 1$, then (D) says that sw is equal to $ss_1 \cdots s_d$ with two letters deleted; multiplying by s (and remembering that $l(w) = d$), we obtain $w = ss_1 \cdots \hat{s}_i \cdots s_d$.

In order to put (E) into perspective, note that, for general (W, S) , we have the following three possibilities for $l(sw)$: (a) $l(sw) = l(w) + 1$; this happens if and only if we can get a reduced decomposition of sw by putting s in front of a reduced decomposition of w . (b) $l(sw) = l(w) - 1$; this happens if and only if w admits a reduced decomposition starting with s . (c) $l(sw) = l(w)$.

[Possibility (c) might seem counterintuitive at first, but easy examples show that it can happen. It happens, for instance, if W is the direct product of two groups of order 2 and S consists of the three non-trivial elements of W .]

The content of (E), then, is the following: First, possibility (c) is prohibited. Second, if (b) holds then we can always find a reduced decomposition of w starting with s by taking an arbitrary reduced decomposition $w = s_1 \cdots s_d$ and then “exchanging” a suitable letter s_i for an s in front.

Note that (E) seems to be asymmetric, in that it involves only *left* multiplication by elements of S ; but if (E) holds, then we can apply it to w^{-1} to deduce the analogous fact about right multiplication. We will use this observation without comment whenever it is convenient.

Next we record a consequence of (E). It will be called the *folding condition*, because, as we will see shortly, it is closely related to the existence of “foldings” of Σ .

(F) Given $w \in W$ and $s, t \in S$ such that $l(sw) = l(w) + 1$ and $l(wt) = l(w) + 1$, either $l(swt) = l(w) + 2$ or else $swt = w$.

To see that (E) implies (F), take a reduced decomposition $w = s_1 \cdots s_d$. Then the word $s_1 \cdots s_d t$ is a reduced decomposition of wt . Applying (E) to s and wt , we conclude that either $l(swt) = d + 2$ or else we can exchange one of the letters in $s_1 \cdots s_d t$ for an s in front. Now the letter exchanged for s cannot be an s_i , since that would contradict the assumption that $l(sw) = d + 1$; so the letter must be the final t . Thus $wt = sw$, hence $swt = w$.

Finally, we show that (F) implies (D), so that (D), (E), and (F) are in fact all equivalent:

Suppose $w = s_1 \cdots s_d$ with $d > l(w)$. Assuming (F), we will show by induction on d that we can delete two letters. If either of the elements $s_1 \cdots s_{d-1}$ or $s_2 \cdots s_d$ has length less than $d - 1$, then we are done by the induction hypothesis. So suppose they both have length $d - 1$ and let $w' = s_2 \cdots s_{d-1}$. (This makes sense, because we necessarily have $d \geq 2$.)

Then $l(s_1 w') = l(w') + 1 = l(w' s_d)$ and $l(s_1 w' s_d) < l(w') + 2$; so (F) implies that $s_1 w' s_d = w'$, i.e., that $w = \hat{s}_1 s_2 \cdots s_{d-1} \hat{s}_d$.

3B Construction of foldings

Although we are still not ready for a systematic treatment of the geometry of half-spaces in Σ , we will show, as an illustration of (F), how to construct maps which “fold” Σ along a wall. Specifically, we wish to construct for a fixed $s = s_H \in S$ the map ϕ which, intuitively, folds Σ onto the half-space Φ determined by H which contains the fundamental chamber C .

In order to figure out how ϕ and Φ should be defined, recall that there are two possibilities for an element $w \in W$: either $l(sw) = l(w) - 1$ or $l(sw) = l(w) + 1$. In the first case, w admits a reduced decomposition starting with s , so there is a minimal gallery of the form C, sC, \dots, wC . We therefore expect that H separates C from wC in this case and hence that $wC \notin \Phi$. In the second case, there is a minimal gallery of the form C, sC, \dots, swC . So we expect that swC is not in Φ but that its “mirror image” wC is in Φ . These considerations motivate the following proposition and its proof:

Proposition. *Suppose that (W, S) satisfies the equivalent conditions (D), (E), and (F). Fix $s \in S$, and let \mathcal{C} be the set of chambers of Σ . Then there is a function $\phi = \phi_s : \mathcal{C} \rightarrow \mathcal{C}$ with the following properties:*

- (1) ϕ is a retraction onto its image Φ , which consists of the chambers wC such that $l(sw) = l(w) + 1$.
- (2) Each chamber in Φ is the image under ϕ of exactly one chamber in the complement Φ' of Φ .
- (3) The action of s on \mathcal{C} interchanges the sets Φ and Φ' .
- (4) ϕ takes adjacent chambers to adjacent chambers.

PROOF: It is clear how we should define ϕ :

$$\phi(wC) = \begin{cases} wC & \text{if } l(sw) = l(w) + 1 \\ swC & \text{if } l(sw) = l(w) - 1. \end{cases}$$

And it is immediate from this definition that assertions (1)–(3) hold. The crucial thing, then, is to verify (4), which is what makes it plausible that we will eventually be able to extend ϕ to a map on all of Σ . We will actually prove a more precise version of (4), namely, that ϕ takes t -adjacent chambers to t -adjacent chambers for all $t \in S$. We may assume that the t -adjacent chambers we start with are distinct, say wC and wtC . We may also assume that $l(wt) = l(w) + 1$; for if this fails then we can replace w by wt . We now have two cases to consider:

(a) $l(sw) = l(w) + 1$. Then $\phi(wC) = wC$. If $l(swt) = l(wt) + 1$, then $\phi(wtC) = wtC$, which is t -adjacent to $\phi(wC)$. Otherwise, (F) says that $swt = w$; so we have $\phi(wtC) = swtC = wC$, which is equal to $\phi(wC)$ and hence t -adjacent to it.

(b) $l(sw) = l(w) - 1$. Then w admits a reduced decomposition starting with s , so wt does also. Thus $\phi(wtC) = swtC$, which is t -adjacent to $swC = \phi(wC)$. \square

3C The word problem

Now we give a purely algebraic consequence of the deletion condition. Namely, we will show, following Tits [55], how it leads to a simple solution to the word problem for (W, S) . The *word problem* is the following: Given two S -words $\mathbf{s} = (s_1, \dots, s_d)$ and $\mathbf{t} = (t_1, \dots, t_e)$, decide whether they represent the same element of W . Let's begin with the case of a dihedral group D_{2m} generated by two elements s, t such that st has order m ($2 \leq m \leq \infty$).

It is obvious, first of all, that we may confine our attention to the case where \mathbf{s} and \mathbf{t} are *alternating* words, i.e., where they have no consecutive s 's or t 's. Secondly, we may assume that both words have length at most m . For the relation $(st)^m = 1$ (if m is finite) can be rewritten as

$$stst \cdots = tsts \cdots,$$

where both sides have length m . So in any word of length $> m$, we can take a subword (s, t, \dots) of length m and replace it by the word (t, s, \dots) of length m , thereby creating an (s, s) or (t, t) that can be deleted. Finally, the word problem for alternating words of length $\leq m$ has the following simple solution: The two alternating words of length m (if m is finite) represent the same element of D_{2m} ; all other pairs of distinct alternating words of length $\leq m$ represent different group elements. The proof is an easy computation, which is left to the reader. [Alternatively, think about what galleries look like when the plane is divided into $2m$ chambers by m lines through the origin (if m is finite) or when the line is divided into infinitely many intervals (if m is infinite).]

Returning now to an arbitrary (W, S) , consider the Coxeter matrix

$$M = (m(s, t))_{s, t \in S}$$

where $m(s, t)$ is the order of st . By an *elementary M -operation* on a word we mean an operation of one of the following two types:

- (I) Delete a subword of the form (s, s) .
- (II) Given $s, t \in S$ with $s \neq t$ and $m(s, t) < \infty$, replace an alternating subword (s, t, \dots) of length $m = m(s, t)$ by the alternating word (t, s, \dots) of length m .

Call a word *M -reduced* if it cannot be shortened by any finite sequence of elementary M -operations. It is not hard to see that one can effectively enumerate all possible words obtainable from a given one by elementary M -operations. [You might want to try an example; see, for instance, Exercise 5 in §3D below.] In particular, one can decide whether a word is M -reduced. Similarly, one can decide whether a word \mathbf{s} can be converted to a given

word t by means of elementary M -operations. Consequently, the following theorem of Tits [55] solves the word problem when (D) holds:

Theorem. *Assume that (W, S) satisfies the deletion condition.*

- (1) *A word is reduced if and only if it is M -reduced.*
- (2) *If s and t are reduced, then they represent the same element of W if and only if s can be transformed to t by the application of elementary M -operations of type (II).*

PROOF: We begin with (2). Suppose $s = (s_1, \dots, s_d)$ and $t = (t_1, \dots, t_d)$ are reduced words representing the same element $w \in W$. We will show by induction on $d = l(w)$ that s can be changed to t by operations of type (II). Let $s = s_1$ and $t = t_1$. There are two possibilities:

(a) $s = t$. Then we can cancel the first letter from each side of the equation

$$s_1 \cdots s_d = t_1 \cdots t_d,$$

and we are done by the induction hypothesis.

(b) $s \neq t$. We will show, then, that $m = m(s, t)$ is finite and that w admits a third reduced decomposition u starting with the alternating word (s, t, s, t, \dots) of length m . Assuming this for the moment, let u' be the word obtained from u by replacing this initial segment of length m by the word (t, s, t, s, \dots) of length m . We can then get from s to t by

$$s \mapsto u \mapsto u' \mapsto t,$$

where the first and third arrows are given by case (a) and the second is an operation of type (II).

It remains, then, to prove the finiteness of m and the existence of u . We do this by repeatedly applying the exchange condition: Since w admits a reduced decomposition starting with t , we can find one by exchanging one of the letters in (s, s_2, \dots, s_d) for a t in front. Now the letter exchanged for t cannot be the initial s , since we could then cancel $s_2 \cdots s_d$ and conclude that $s = t$. So it must be one of the others, and we obtain a reduced decomposition of w starting with (t, s) . If $m = 2$, we are done [because the (t, s) can be replaced by (s, t)]. Otherwise, $m \geq 3$ and we continue:

Since w admits a reduced decomposition starting with s , we can find one by starting with the decomposition $w = ts \cdots$ just obtained and exchanging one of the letters for an s in front. Now the exchanged letter cannot come from the initial segment of length 2; for this would contradict our analysis of the word problem in the dihedral subgroup of order $2m$ generated by s and t . So it must come from a later part of the word. Hence $d \geq 3$, and we have a reduced decomposition of w starting with (s, t, s) . If $m = 3$, we are done; otherwise, do an exchange on this most recent reduced decomposition to get a t in front.

There is no obstruction to continuing this process as long as our initial alternating segment has length $< m$. Since this length must always be $\leq d$,

it follows that $m \leq d < \infty$ and that we can find an initial alternating segment of length m , as claimed. This completes the proof of (2).

Turning now to (1), the non-trivial implication to prove is that if $\mathbf{s} = (s_1, \dots, s_d)$ is not reduced, then it can be shortened by M -operations. We argue by induction on d . If the subword $\mathbf{s}' = (s_2, \dots, s_d)$ is not reduced, then we are done by the induction hypothesis. So assume that \mathbf{s}' is reduced and let $w' = s_2 \cdots s_d$. Since, $l(s_1 w') < l(w') + 1$, we can find a reduced decomposition of w' starting with s_1 , say $\mathbf{t}' = (s_1, t_1, \dots, t_{d-2})$. By part (2) of the theorem, which we have already proved, we can transform \mathbf{s}' to \mathbf{t}' by M -operations, hence we can transform \mathbf{s} to $(s_1, s_1, t_1, \dots, t_{d-2})$. But this can then be reduced to $\mathbf{t} = (t_1, \dots, t_{d-2})$ by an operation of type (I). \square

Note that this solution of the word problem gives a complete description of the elements of W in terms of the Coxeter matrix. Consequently, we have the following generalization of a result which we already knew for finite reflection groups:

Corollary 1. *If (W, S) satisfies the deletion condition, then W is determined up to isomorphism by its Coxeter matrix M .* \square

We can make this more precise, in a way that gives us new information even for finite reflection groups:

Corollary 2. *If (W, S) satisfies the deletion condition, then W admits the presentation*

$$W = \langle S ; (st)^{m(s,t)} = 1 \rangle,$$

where there is one relation for each pair s, t with $m(s, t) < \infty$.

PROOF: Let \widetilde{W} be the abstract group defined by this presentation, and consider the canonical surjection $\widetilde{W} \twoheadrightarrow W$. An element \tilde{w} in the kernel can be represented by a word \mathbf{s} which is reducible to the empty word by M -operations. But M -operations do not change the element of \widetilde{W} represented by a word, so $\tilde{w} = 1$. \square

Corollary 3. *Assume (W, S) satisfies the deletion condition. Then for any $w \in W$ there is a subset $S(w) \subseteq S$ such that all reduced decompositions of w involve precisely the letters in $S(w)$. Moreover, $S(w)$ is the smallest subset $S' \subseteq S$ such that $w \in \langle S' \rangle$.*

PROOF: The first assertion is immediate from the theorem, since operations of type (II) do not change the set of letters which occur in a word (although they can change the number of times a given letter occurs). Now suppose S' is an arbitrary subset of S with $w \in \langle S' \rangle$. Then we can get a reduced decomposition of w by starting with an S' -word representing w and repeatedly applying (D) until we get a reduced word; hence $S(w) \subseteq S'$. \square

3D Exercises

Now it's your turn to play with condition **(D)** a little. Exercises 1–3 give some consequences of **(D)**, and Exercise 4 should convince you that these consequences are reasonable. The last two exercises are for people interested in the word problem. Assume throughout the exercises that (W, S) satisfies **(D)**.

1. Let S' be a subset of S and let W' be the group $\langle S' \rangle$ generated by S' . Show that the length of an element of W' with respect to S' is the same as its length with respect to S .

2. With W' as in Exercise 1, show that every coset wW' has a unique representative w_0 of minimal length and that $l(w_0 w') = l(w_0) + l(w')$ for all $w' \in W'$.

3. Suppose that W is finite. Show that W contains a unique element w_0 of maximal length and that $l(w_0) = l(w) + l(w^{-1}w_0)$ for all $w \in W$. [In other words, you can find a reduced decomposition of w_0 that starts with any feasible initial segment whatsoever.] Show further that w_0 is of order 2 and that $\mathcal{H}(w_0) = \mathcal{H}$, i.e., that the minimal galleries from C to $w_0 C$ cross every wall.

4. Try to give geometric interpretations of the results of Exercises 1–3.

5. Let W be the reflection group of type A_3 [symmetric group on four letters]. Find all reduced decompositions of the element w_0 of maximal length, which is $s_1 s_3 s_2 s_1 s_3 s_2$. [List all words obtainable from the given word by M -operations. There are 16 of them.]

6. Recall from the classification of irreducible finite reflection groups that their Coxeter diagrams have a number of special properties, including the following: The graph is a tree; it branches at at most one vertex, which is then necessarily of order 3; if it branches, there are no labelled edges (i.e., there are no $m_{ij} > 3$); if it doesn't branch, there is at most one labelled edge. As I mentioned in Chapter I, these facts are proved in the course of proving the classification theorem. Show that they all follow from Tits's solution of the word problem. [For each of the facts, assume it is false and then write down a word which represents an element of infinite order, contradicting the hypothesis that the group is finite. Details can be found in §3 of [55].]

4 Coxeter Groups

We return now to an arbitrary (W, S) , where W is a group and S is a set of generators of W of order 2. We have seen (§3C, Corollary 2) that if (W, S) satisfies **(D)** then it admits a presentation in which the relations simply specify the orders of the pairwise products of elements of S . Tits [51] initiated the systematic study of groups with such a presentation. He called them *Coxeter groups*, since Coxeter [25] had earlier studied finite groups of this type. We will therefore call the following condition on (W, S) the *Coxeter condition*:

(C) W admits the presentation

$$\langle S ; (st)^{m(s,t)} = 1 \rangle,$$

where $m(s,t)$ is the order of st and there is one relation for each pair s, t with $m(s,t) < \infty$.

We have now introduced five conditions on (W, S) , which are related as follows:

$$(A) \implies (D) \iff (E) \iff (F) \implies (C).$$

On the other hand, a Coxeter presentation as in (C) is precisely what we used in §2 to give a proof by combinatorial group theory that $\mathrm{PGL}_2(\mathbf{Z})$ satisfies (A). This proof goes through with no change to show, in general, that (C) implies (A). Thus we have come full circle:

Theorem. *The five conditions (A), (C), (D), (E), and (F) are all equivalent.* \square

It now seems safe to conclude that we have found the right class of groups that deserve to be called “abstract reflection groups”. This terminology is not standard, however, so we will follow Tits and say that W is a *Coxeter group* (or, more precisely, that the pair (W, S) is a *Coxeter system*) if the equivalent conditions of the theorem are satisfied.

5 Loose Ends

Before leaving combinatorial group theory and moving on to combinatorial geometry (i.e., the properties of Σ), we will comment briefly on some natural questions which may have occurred to you:

(a) Which matrices M can occur as the Coxeter matrix of a Coxeter group? [We saw in Chapter I that very few matrices can occur as the Coxeter matrix of a finite reflection group, because an associated matrix A has to be positive definite in that case. Are there conditions that A must satisfy in the general case?]

(b) For which M can the corresponding group $W = W_M$ be represented as a “geometric reflection group”? [The finite reflection groups of Chapter I should be thought of as *spherical* reflection groups, since they act as isometries of the sphere and the associated complex Σ triangulates the sphere. And clearly the examples D_∞ and $\mathrm{PGL}_2(\mathbf{Z})$ of §2 should be thought of as *Euclidean* and *hyperbolic*, respectively. Can all Coxeter groups be classified as spherical, Euclidean, or hyperbolic?]

(c) For which M can the group W_M be represented as a “linear reflection group”? [All examples we have seen have admitted a canonical linear representation, with an associated chamber decomposition of some convex cone in the corresponding vector space (or perhaps its dual). Is this always true?]

The remainder of this section, which may be omitted, is devoted to giving partial answers to these questions.

Let I be an arbitrary index set and let $M = (m_{ij})_{i,j \in I}$ be a matrix with $m_{ij} \in \mathbb{Z} \cup \{\infty\}$. We will call M a *Coxeter matrix* if

$$m_{ii} = 1 \quad \text{and} \quad 2 \leq m_{ij} = m_{ji} \leq \infty \text{ for } i \neq j.$$

This terminology is justified by the following theorem, which gives as nice an answer to question (a) as one could hope for:

Theorem A. *Every Coxeter matrix is the Coxeter matrix of a Coxeter group.*

PROOF: There is no choice as to how to define the Coxeter group; if we want (C) to hold, we must set

$$W = \langle (s_i)_{i \in I} ; (s_i s_j)^{m_{ij}} = 1 \rangle,$$

where, as usual, the relation occurs only if $m_{ij} < \infty$. I claim, first, that the images in W of the s_i are distinct and non-trivial. [This is not obvious; and it is important, since otherwise I would not even be an appropriate index set for the Coxeter matrix of W .] Accepting this for the moment, we may identify s_i with its image in W and we may introduce the set $S = \{s_i\} \subset W$. The next claim is that $s_i s_j$ has order m_{ij} in W . [Again, this is not obvious; all we know is that the order of $s_i s_j$ is a divisor of m_{ij} .] If we prove both claims, the theorem will follow at once.

The proof of the claims will be based on a linear representation of W which will look familiar (cf. §I.5D, as well as §2 of the present chapter). Let V be a vector space with a basis $(e_i)_{i \in I}$. Let B be the bilinear form on V such that $B(e_i, e_j) = -\cos(\pi/m_{ij})$, and let $\sigma_i : V \rightarrow V$ be the reflection defined by the usual formula: $\sigma_i x = x - 2B(e_i, x)e_i$. It takes e_i to $-e_i$ and it fixes the hyperplane $H_i = e_i^\perp = \{x \in V : B(e_i, x) = 0\}$.

For any $i \neq j$ with $m_{ij} < \infty$, the bilinear form B is non-degenerate on the subspace $V_1 = \mathbb{R}e_i \oplus \mathbb{R}e_j \subseteq V$. It follows by elementary linear algebra that we have a decomposition $V = V_1 \oplus V_0$, where $V_0 = V_1^\perp = H_i \cap H_j$. This decomposition is invariant under σ_i and σ_j , which generate a dihedral group $D_{2m_{ij}}$, acting in the canonical way on V_1 (and acting trivially on V_0). In particular, $\sigma_i \sigma_j$ has order m_{ij} . In view of the defining presentation of W , we get a linear action of W on V with s_i acting as σ_i . Since the σ_i are distinct and non-trivial, it follows that the same is true of the s_i . This proves the first claim. And since $\sigma_i \sigma_j$ has order precisely m_{ij} , the order of $s_i s_j$ cannot be a proper divisor of m_{ij} . This proves the second claim when $m_{ij} < \infty$.

Finally, to prove that $s_i s_j$ has infinite order when $m_{ij} = \infty$, note that σ_i and σ_j still leave the plane V_1 invariant in this case; so it suffices to show that their product has infinite order when restricted to V_1 . This can be done by direct computation, or, better, by noting that V_1 (with the reflections σ_i and σ_j) can be identified with the vector space called V' (with reflections s'_1 and s'_2) in our discussion of D_∞ in §2. The result now follows from the

fact (stated in that discussion) that the duals of s'_1 and s'_2 generate an infinite dihedral group acting on the dual of V' . \square

Note, as a result of this theorem, that you can now make up as many examples as you want of “abstract reflection groups”—just write down an arbitrary Coxeter matrix.

We turn next to question (c), since the proof of Theorem A has already suggested the answer. If (W, S) is an arbitrary Coxeter system, then the proof of Theorem A yields a vector space V on which W acts, with the generators $s \in S$ acting as linear reflections. We will refer to this action as the *canonical linear representation* of W . The following theorem is due to Tits [51].

Theorem C. *The canonical linear representation of W is faithful, i.e., the corresponding homomorphism from W to the group $\text{Aut } V$ of linear automorphisms of V is injective. Hence W is isomorphic to a group of linear transformations generated by reflections.*

PROOF: We continue with the notation of the proof of Theorem A; in particular, we write $S = \{s_i : i \in I\}$. Let V^* be the dual of V . Then W acts on V^* , with s_i acting as σ_i^* . We will use “inner product notation” for the canonical bilinear pairing between V and V^* ; thus $\langle x, \xi \rangle = \xi(x)$ for $x \in V$ and $\xi \in V^*$. Let C be the “fundamental chamber” in V^* defined by the inequalities $\langle e_i, - \rangle > 0$ for $i \in I$. [We note, in passing, that it would not be reasonable to try to avoid V^* by defining C to be the subset of V given by the inequalities $B(e_i, -) > 0$; for the linear functions $B(e_i, -)$ might be linearly dependent, as we already saw in the case of D_∞ . But if B is non-degenerate, as in the case of finite reflection groups and the group $\text{PGL}_2(\mathbf{Z})$, then there is no need to introduce V^* ; indeed, B then induces an isomorphism $V \rightarrow V^*$ (provided S is finite), given by $x \mapsto B(x, -)$.]

To prove W acts faithfully on V , we will show that, under the action of W on V^* , $wC \cap C = \emptyset$ for $1 \neq w \in W$. This follows from the lemma below, whose statement should be no surprise, given the way we have learned to think geometrically about (W, S) . [See, in particular, the construction of the folding ϕ in §3B above.]

Lemma. *Fix $s = s_i \in S$ and let $U_+(s)$ and $U_-(s)$ be the open half-spaces in V^* defined, respectively, by $\langle e_i, - \rangle > 0$ and $\langle e_i, - \rangle < 0$. Then for any $w \in W$ we have $wC \subseteq U_+(s)$ if $l(sw) = l(w) + 1$ and $wC \subseteq U_-(s)$ if $l(sw) = l(w) - 1$.*

SKETCH OF THE PROOF: We argue by induction on $l(w)$. If $l(sw) < l(w)$, then we may apply the induction hypothesis to the element sw to get $swC \subseteq U_+(s)$; multiplying by s , we find $wC \subseteq sU_+(s) = U_-(s)$, as required. Suppose now that $l(sw) > l(w)$. We may assume $w \neq 1$, so there is a $t \in S$ (necessarily different from s) such that $l(tw) < l(w)$. Rip off from w a maximal factor w' in the dihedral subgroup W' generated by s and t . In other words, write $w = w'w''$ with $w' \in W'$, $l(w) = l(w') + l(w'')$,

$l(sw'') > l(w'')$, and $l(tw'') > l(w'')$. By the induction hypothesis, we have $w''C \subseteq C' = U_+(s) \cap U_+(t)$. Now C' is essentially the fundamental chamber for the dual of the canonical representation of W' . More precisely, the dual of the canonical representation of W' is a 2-dimensional quotient of V^* , and C' is the inverse image of the fundamental chamber. But we've studied the canonical representation of W' and its dual in detail, whether W' is finite or infinite, and we know that its chamber geometry behaves in the expected way. In particular, since $l(sw') > l(w')$, it follows that $w'C' \subseteq U_+(s)$; hence $wC = w'w''C \subseteq w'C' \subseteq U_+(s)$. \square

We can now easily prove that there is no difference between a finite “abstract reflection group” and a finite reflection group:

Corollary. *If W is a finite Coxeter group, then W can be faithfully represented as a finite reflection group, in the sense of Chapter I.*

PROOF: This is immediate from the theorem, except for the requirement in Chapter I that the reflections are supposed to be orthogonal with respect to some inner product on V . But this is no problem, for the finite group W acting on the vector space V of the theorem necessarily leaves an inner product invariant—just take an arbitrary inner product $(-, -)$ on V and construct a W -invariant inner product $\langle -, - \rangle$ from it by “averaging”:

$$\langle x, y \rangle = \sum_{w \in W} (wx, wy).$$

For each $s \in S$, the (± 1) -eigenspaces of s are orthogonal to one another with respect to our W -invariant inner product, so s indeed acts on V as an orthogonal reflection. \square

Remarks

1. Combining the corollary with the classification of finite reflection groups (cf. §I.3), we recover Coxeter's list [25] of the finite Coxeter groups.

2. With C as in the proof of the theorem, Tits [51] showed that the set $U = \bigcup_{w \in W} w\overline{C}$ is always a convex cone in V^* and that \overline{C} is a fundamental domain for the action of W on U . [But U is never all of V^* unless W is finite.] Published proofs can be found in Bourbaki [16] or Vinberg [62]. Vinberg's paper is particularly recommended; it contains a general treatment of linear reflection groups, which adds a great deal to what we have sketched here.

We turn, finally, to question (b). I will not state a Theorem B which answers the question precisely, since this would take us too far afield. In particular, I would feel compelled to carefully define the terms “spherical reflection group”, etc. Instead, I will give only the following brief indication; see [16] and [62] for more information.

Let (W, S) be an irreducible Coxeter system with S finite. (“Irreducible” means, as in Chapter I, that the Coxeter diagram is connected.) Let M be

the Coxeter matrix and let B be the associated symmetric bilinear form, as in the proof of Theorem A. Then (W, S) is spherical if and only if B is positive definite, Euclidean if and only if B is positive semi-definite but degenerate, and hyperbolic if and only if B has signature $(n - 1, 1)$, where $n = \text{card } S$. [*Warning*: This characterization of the hyperbolic case is true as stated only if we confine our attention to hyperbolic reflection groups with a simplex as fundamental domain. But the fundamental domain of a hyperbolic reflection group does not have to be a simplex. See [62] for a characterization that covers the general case.]

This suggests that the geometric reflection groups (i.e., those associated to spherical, Euclidean, or hyperbolic geometry), are somewhat special among Coxeter groups as a whole. In particular, it is *not* true that all Coxeter groups are spherical, Euclidean, or hyperbolic.

III

Coxeter Complexes

Assume throughout this chapter that (W, S) is a Coxeter system with S finite. Let $\Sigma = \Sigma(W, S)$ be the poset which was defined at the beginning of Chapter II. Following Tits, we will call Σ the *Coxeter complex* associated to (W, S) . The word “complex” will be justified below, when we prove that Σ is indeed a simplicial complex. The purpose of this chapter is to develop the geometric properties of Coxeter complexes.

The assumption that S is finite is not really necessary, but we make it in order to stay closer to the geometric intuition (where S is thought of as the set of reflections in the walls of a chamber). Everything would go through with no essential change if we dropped this assumption, but we would have to deal with “simplicial complexes” in which a simplex can have infinitely many vertices.

1 The Coxeter Complex is Simplicial

Lemma. *The function $S' \mapsto \langle S' \rangle$ is a poset isomorphism from the set of subsets of S to the set of special subgroups of W .*

PROOF: We define a map in the other direction by $W' \mapsto W' \cap S$. It is clear that $W' = \langle W' \cap S \rangle$ if W' is a special subgroup. It is also clear that $S' \subseteq \langle S' \rangle \cap S$ for any $S' \subseteq S$. To prove the opposite inclusion, suppose $s \in \langle S' \rangle \cap S$. Then we can express s as an S' -word and repeatedly apply the deletion condition until the word's length has been reduced to 1; thus $s \in S'$. Hence $S' = \langle S' \rangle \cap S$, and our two maps are inverses of one another. \square

At this point, you may need to refer to the appendix to Chapter I for the terminology regarding chamber complexes and labellings. Let's add one more bit of terminology: A chamber complex is called *thin* if every codimension 1 simplex is a face of exactly two chambers.

Theorem. *The poset Σ is a simplicial complex. Moreover, it is a thin, labellable chamber complex of rank $n = \text{card } S$, and the W -action on Σ is type-preserving.*

[The last assertion means that $\lambda(wA) = \lambda(A)$ for $w \in W$ and $A \in \Sigma$, where λ is any labelling. In view of the essential uniqueness of λ (cf. Appendix to Chapter I), this is independent of the choice of λ .]

PROOF: To show that Σ is simplicial, there are two things we must verify (cf. Appendix to Chapter I):

(a) Any two elements of Σ have a greatest lower bound.

Using the W -action on Σ , we may assume that one of the two elements is a face of the fundamental chamber $C = \{1\}$. What we must prove, then, is that a special subgroup $\langle S' \rangle$ and a special coset $w\langle S'' \rangle$ have a least upper bound in the set of special cosets (with respect to the ordering by inclusion). Now any special coset containing the two given ones contains the identity and hence is a special subgroup. Moreover, it contains w and hence also $\langle S'' \rangle = w^{-1}w\langle S'' \rangle$. So the upper bounds of our two special cosets are the special subgroups containing S' , S'' , and w . In view of Corollary 3 in §II.3C, there is indeed a smallest upper bound, namely, the special subgroup $\langle S' \cup S'' \cup S(w) \rangle$.

(b) For any $A \in \Sigma$, the poset $\Sigma_{\leq A}$ is isomorphic to the set of subsets of some finite set.

It suffices to prove this for $A = C$. In this case, $\Sigma_{\leq C}$ is the set of special subgroups of W (ordered by the opposite of the inclusion relation). Using the lemma, then, we obtain

$$\Sigma_{\leq C} \approx (\text{subsets of } S)^{\text{op}} \approx (\text{subsets of } S),$$

where the second isomorphism is given by $S' \mapsto S - S'$. This proves (b) and completes the proof that Σ is simplicial.

Next, all maximal simplices of Σ have rank equal to $\text{card } S$ since W acts transitively on them. And the discussion at the beginning of Chapter II implies that any two maximal simplices can be connected by a gallery and that any codimension 1 simplex is a face of exactly two chambers. So Σ is a thin chamber complex.

Finally, we can define a W -invariant labelling λ of Σ , with S as the set of labels, by setting $\lambda(w\langle S' \rangle) = S - S'$. \square

We will continue to denote by λ the labelling just constructed, and we will call it the *canonical labelling* of Σ .

EXERCISE

The canonical labelling yields a notion of *s-adjacency* for any $s \in S$. On the other hand, we gave an *ad hoc* definition of “*s-adjacency*” in Chapter II. Show that the two definitions coincide.

2 Local Properties of Coxeter Complexes

By “local properties” we mean properties of the links of simplices. [See §D of the appendix to Chapter I for the definition of “link”.] For example, it is of interest to know whether these links are chamber complexes. The following proposition shows that, in fact, these links are again Coxeter complexes.

Proposition. *Given $A \in \Sigma = \Sigma(W, S)$, let $S' = S - \lambda(A)$ and let $W' = \langle S' \rangle$. Then $\text{lk}_\Sigma A$ is isomorphic to the Coxeter complex $\Sigma(W', S')$ associated to the Coxeter system (W', S') . In particular, this link is a chamber complex.*


[You should convince yourself, before proceeding to the proof, that (W', S') is indeed a Coxeter system. This is easy if you use condition (A) as the definition of “Coxeter system”. Alternatively, you can use (D) together with Exercise 1 of §II.3D.]

PROOF: We may assume that A is a face of the fundamental chamber. Then A is the special subgroup W' defined in the statement of the proposition. Recall now that there is a poset isomorphism $\text{lk}_\Sigma A \approx \Sigma_{\geq A}$; hence the link of A is isomorphic to the set of special cosets in W that are contained in W' , ordered by the opposite of the inclusion relation. But the special cosets that are contained in W' are precisely the same as the special cosets associated to the Coxeter system (W', S') . Thus $\Sigma_{\geq A} = \Sigma(W', S')$. \square

This proposition has a simple interpretation in terms of Coxeter matrices. Recall, first, that the Coxeter system (W, S) is determined by its Coxeter matrix $M = (m(s, t))_{s, t \in S}$. So we may think of Σ as a simplicial complex associated to M . Next, note that the rows and columns of M are indexed by S , which is also the set of labels of Σ . What the proposition says, then, is that $\text{lk } A$ is the Coxeter complex associated to the matrix M' obtained from M by deleting the rows and columns corresponding to the labels $s \in \lambda(A)$.

This becomes even easier to use if we translate it into the language of Coxeter *diagrams*. Recall that the diagram of (W, S) has one vertex for each $s \in S$, with s joined to t if $m(s, t) \geq 3$, and with a label over that edge if $m(s, t) \geq 4$. The passage from M to M' above, and hence the passage from Σ to $\text{lk } A$, corresponds to the following operation on the diagram: For each $s \in \lambda(A)$, delete the vertex s (and all edges touching s) from the diagram.

Consider, for example, the group $W = \text{PGL}_2(\mathbf{Z})$ studied in §II.2C. Its diagram is



The Coxeter complex Σ has rank 3 (dimension 2), so there are three types of vertices. Let's compute the link of each type of vertex.

According to the recipe above, we must delete one vertex at a time from the Coxeter diagram of W . This yields the Coxeter diagrams of the dihedral groups D_{2m} , where $m = \infty, 2$, and 3 , respectively. Now it is easy to figure out what the Coxeter complex associated to D_{2m} looks like, and, in fact, we have already seen it in Chapters I and II. Namely, it is a $2m$ -gon, i.e., it is a triangulated circle with $2m$ edges if $m < \infty$, and it is a triangulated line if $m = \infty$. So our three links in this example are a line, a quadrilateral, and a hexagon.

EXERCISE

Look at Figure 2 in §II.2C. Can you find the three types of links in the picture? [HINT: To find $\text{lk } v$, locate all chambers having v as a vertex. The union of the closures of these chambers (including the cusps) is the cone over $\text{lk } v$, with v as cone point; so you can see $\text{lk } v$ as the “boundary” of this union.]

This example illustrates a general principle, valid for all Coxeter complexes: The link of a codimension 2 simplex of type $S - \{s, t\}$ is a $2m$ -gon, where $m = m(s, t)$. This fact yields a geometric interpretation of the Coxeter matrix M :

Corollary 1. *The Coxeter matrix M of (W, S) can be recovered from Σ as follows: For any $s, t \in S$ with $s \neq t$, $m(s, t)$ is the unique number m ($2 \leq m \leq \infty$) such that the link of a simplex of type $S - \{s, t\}$ is a $2m$ -gon.* \square

This shows, in particular, that the Coxeter group W is determined up to isomorphism by Σ . We’ll see this again in the next section, from a different point of view.

Remark. Note that a $2m$ -gon has diameter m , where the *diameter* of a chamber complex is the supremum of the combinatorial distances between its chambers. So we can also write the geometric interpretation of M as

$$m(s, t) = \text{diam}(\text{lk } A),$$

where $\lambda(A) = S - \{s, t\}$ as above. The result in this form is valid even when $s = t$. [In this case the link has exactly two chambers, which are adjacent, so the diameter is indeed $1 = m(s, s)$.]

We can use this corollary, together with Tits’s solution to the word problem for Coxeter groups, to give a simple answer to a question which might seem, *a priori*, to be very difficult: How can one describe the totality of minimal galleries connecting two given chambers? This is easy in the 1-dimensional case, where Σ is a $2m$ -gon: Minimal galleries are unique except when the two given chambers C_1 and C_2 are at maximum distance m from each other; in this case m is necessarily finite, and there are exactly two minimal galleries connecting C_1 to the opposite chamber C_2 .

Translating this result to the link of a simplex A of codimension 2 in an arbitrary Coxeter complex, we obtain a similar description of the minimal galleries in the subposet $\Sigma_{\geq A}$. [Visualize, for example, the case where Σ is 2-dimensional and A is a vertex v whose link is finite. Then $\Sigma_{\geq A}$ contains $2m$ chambers for some m , which form a solid $2m$ -gon centered at v . The only non-uniqueness of minimal galleries in this subposet arises from the fact that there are two ways of going around the $2m$ -gon to get from a given chamber to the opposite chamber.]

Since galleries correspond to words, we can use the solution to the word problem (§II.3C) to analyze the general case. The answer, stated in rough form, is:

Corollary 2. *The non-uniqueness of minimal galleries in a Coxeter complex can be explained entirely in terms of the obvious non-uniqueness that occurs in links of codimension 2 simplices.* \square

I leave it as an exercise for you to restate this more precisely, in terms of “elementary operations” on galleries. Similarly, you should be able to state a method for using elementary operations to decide whether a given gallery is minimal and, if it isn’t, to obtain a minimal gallery from it.

Finally, we can use our calculation of links to answer another question that may have occurred to you, especially if you have some familiarity with combinatorial topology: When is Σ a manifold? This question arises naturally because triangulated manifolds (without boundary) are the canonical examples of thin chamber complexes. You already know the answer if W is finite, at least if you have read the optional section at the end of Chapter II: In this case W is in fact a finite reflection group, so Σ is a sphere (hence a manifold) by the results of Chapter I.

What happens if W (hence also Σ) is infinite? There is an obvious necessary condition. Namely, manifolds are locally compact, hence locally finite, i.e., every non-empty simplex A is a face of only finitely many chambers. In other words, the link of A must be finite. Conversely, if the link of every non-empty simplex is finite, then it is in fact a sphere (since it is a finite Coxeter complex). I leave it as an exercise for the interested reader to deduce that Σ is then a manifold. This proves:

Corollary 3. *The following conditions are equivalent:*

- (1) Σ is a manifold.
- (2) Σ is locally finite.
- (3) Every proper special subgroup of W is finite.

\square

For example, the Coxeter complex associated to $\mathrm{PGL}_2(\mathbf{Z})$ is not a manifold. You can see the non-manifold points in the pictures in Chapter II: They are the cusps.

EXERCISE

If condition (3) holds and W is infinite, show that (W, S) is irreducible.

Remark. Condition (3) is quite restrictive. One can show that it holds only in the following three cases: (a) W is finite; (b) W is an irreducible Euclidean reflection group [in the sense hinted at in §II.5]; (c) W is a hyperbolic reflection group whose fundamental domain is a closed simplex contained entirely in the interior of the hyperbolic space. See the exercises in Bourbaki [16] for more information.

Even though we have not yet officially discussed Euclidean reflection groups, you probably have some intuition about them, and you might be wondering why *reducible* Euclidean reflection groups were excluded in the remark (and in the exercise above): Given Euclidean reflection groups W_1

and W_2 acting on Euclidean spaces E_1 and E_2 , isn't their product W a Euclidean reflection group acting on $E = E_1 \times E_2$, which is a Euclidean space and hence a manifold? And doesn't Σ triangulate this manifold? The answer is "yes" to the first question, but "no" to the second. Exercise 1 below explains what happens.

EXERCISES

1. Let (W', S') and (W'', S'') be Coxeter systems, and let (W, S) be their "sum" (with $W = W' \times W''$ and $S = S' \cup S''$). Show that

$$\Sigma(W, S) \approx \Sigma(W', S') * \Sigma(W'', S''),$$

where the asterisk denotes the join operation. [Recall that the *join* Δ of two simplicial complexes Δ' and Δ'' with vertex sets \mathcal{V}' and \mathcal{V}'' has vertex set equal to the disjoint union $\mathcal{V}' \amalg \mathcal{V}''$ and has one simplex $A' \cup A''$ for every $A' \in \Delta'$ and $A'' \in \Delta''$. From the poset point of view, then, Δ is simply the Cartesian product of Δ' and Δ'' . But its geometric realization $|\Delta|$ is *not* the Cartesian product $|\Delta'| \times |\Delta''|$; in fact, Δ doesn't even have the right dimension for this to be true.]

2. If you are still reading this and haven't skipped ahead to the next section, you have presumably figured out why we had to restrict ourselves to irreducible Euclidean reflection groups above: The join of two manifolds is in general not a manifold. Explain, now, why no such irreducibility restriction was necessary for finite reflection groups.

3 Construction of Chamber Maps

We continue to assume that (W, S) is a Coxeter system with S finite and that $\Sigma = \Sigma(W, S)$ is the associated Coxeter complex. In studying Σ , it is quite easy to work with the chambers and the adjacency relations. It is awkward, on the other hand, to work with the vertices. (If you unwind the definitions, you will find that they are the maximal proper special cosets $w\langle S - \{s\} \rangle$.) It would therefore be of interest to prove that we do not ever have to think about the vertices, i.e., that Σ is determined by its associated *chamber system*, consisting of the set of chambers [which correspond to the elements of W] together with the adjacency relations [given by right multiplication by elements of S].

If you have read about chamber systems in §D of the appendix to Chapter I, then you know that we have, in fact, already proven this. For according to that appendix, it suffices to show that the link of every vertex in Σ is a chamber complex. And we know by the previous section that this is indeed the case.

The specific consequence of this that we will need is that if we want to construct an endomorphism of Σ (i.e., a chamber map $\phi : \Sigma \rightarrow \Sigma$), then we need only give a function ϕ' on the chambers which is compatible with the adjacency relations.

Rather than rely on the appendix to Chapter I for this result, I prefer to give a direct proof. In order to motivate the precise statement, let's think about what "compatible" should mean in the rough statement given above. If we take this to mean "preserving s -adjacency for all s ", then we are only dealing with type-preserving endomorphisms of Σ . To handle the general case we must specify, in addition to ϕ' , a permutation ϕ'' of S which describes how ϕ mixes up the vertex labels. The compatibility condition, then, is that ϕ' takes s -adjacent chambers to $\phi''(s)$ -adjacent chambers.

Here, now, is the precise result:

Proposition. *Endomorphisms ϕ of Σ are in 1-1 correspondence with pairs (ϕ', ϕ'') , where ϕ' is a function $W \rightarrow W$, ϕ'' is a permutation of S , and $\phi'(ws) = \phi'(w)$ or $\phi'(w)\phi''(s)$ for all $w \in W$ and $s \in S$.*

PROOF: We begin with a general observation which we will have occasion to use again. Let $\phi : \Delta \rightarrow \Delta'$ be a chamber map between labellable chamber complexes of the same dimension. Assume that λ is a labelling of Δ by a set I and that λ' is a labelling of Δ' by a set I' . The observation, then, is that there is a bijection $\phi_* : I \rightarrow I'$ which describes the behavior of ϕ with respect to labels, in the sense that $\lambda'(\phi(A)) = \phi_*(\lambda(A))$ for all $A \in \Delta$. This follows from the essential uniqueness of labellings. For $\lambda' \circ \phi$ is a labelling of Δ by I' , so it must differ from the given labelling λ by a bijection from I to I' ; this bijection is the desired ϕ_* .

Apply this now to an endomorphism ϕ of Σ , with λ and λ' both equal to the canonical labelling. We obtain a well-defined bijection $\phi'' = \phi_* : S \rightarrow S$. We also obtain from ϕ a function $\phi' : W \rightarrow W$, which is essentially the restriction of ϕ to the chambers of Σ . (These are the singleton special cosets and hence can be identified with the elements of W .) Clearly ϕ takes s -adjacent chambers to $\phi''(s)$ -adjacent chambers, which says precisely that $\phi'(ws) = \phi'(w)$ or $\phi'(w)\phi''(s)$.

Note that ϕ is completely determined by the pair (ϕ', ϕ'') . For if $A = w\langle S' \rangle$ is an arbitrary simplex of Σ , then A is the face of type $S - S'$ of the chamber $\{w\}$; so $\phi(A)$ must be the face of $\{\phi'(w)\}$ of type $\phi''(S - S') = S - \phi''(S')$; in other words, $\phi(w\langle S' \rangle) = \phi'(w)\langle \phi''(S') \rangle$.

Finally, we must show that every pair (ϕ', ϕ'') as in the statement of the proposition arises from an endomorphism ϕ . To this end we simply define ϕ , as we must, by $\phi(w\langle S' \rangle) = \phi'(w)\langle \phi''(S') \rangle$. It is easy to check that ϕ is a well-defined chamber map which induces ϕ' on the chambers and ϕ'' on the labels. \square

We now give three corollaries to illustrate the proposition. The first two involve automorphisms and the third involves foldings.

Recall that the W -action on Σ is simply-transitive on the chambers; in particular, this action is faithful, in the sense that the corresponding homomorphism $W \rightarrow \text{Aut } \Sigma$ is injective. Here $\text{Aut } \Sigma$ denotes the group of simplicial automorphisms of Σ .

Corollary 1. *The image of $W \hookrightarrow \text{Aut } \Sigma$ is the normal subgroup $\text{Aut}_0 \Sigma$ consisting of the type-preserving automorphisms of Σ .*

(This shows, for the second time, that W is determined up to isomorphism by its Coxeter complex Σ .)

PROOF: We already know that W acts as a group of type-preserving automorphisms of Σ . Conversely, suppose ϕ is an arbitrary type-preserving automorphism, and let ϕ' and ϕ'' be its “components” as in the proposition. Then ϕ'' is the identity, so $\phi'(ws) = \phi'(w)s$ for all w and s . [The possibility $\phi'(ws) = \phi'(w)$ is excluded because ϕ is an automorphism.] It follows easily that $\phi'(w) = \phi'(1)w$ for all w , so ϕ' is left-multiplication by $w_0 = \phi'(1)$ and hence ϕ is given by the action of w_0 . This proves everything except the normality of $\text{Aut}_0 \Sigma$, which is left as an exercise. \square

There is a second obvious source of automorphisms of Σ . Namely, there is a homomorphism $\text{Aut}(W, S) \rightarrow \text{Aut } \Sigma$, where $\text{Aut}(W, S)$ is the group of automorphisms of W stabilizing S ; for such an automorphism takes special cosets to special cosets and hence induces an automorphism of Σ .

Corollary 2. *The homomorphism $\text{Aut}(W, S) \rightarrow \text{Aut } \Sigma$ just defined is injective, and its image is the group $\text{Aut}(\Sigma, C)$ consisting of the automorphisms of Σ which stabilize the fundamental chamber $C = \{1\}$.*

PROOF: Given $\alpha \in \text{Aut}(W, S)$, its image $\phi \in \text{Aut } \Sigma$ has components $\phi' = \alpha$ and $\phi'' = \alpha|_S$. This shows that the homomorphism is injective. And ϕ stabilizes C because $\alpha(1) = 1$. Conversely, suppose we are given $\phi \in \text{Aut}(\Sigma, C)$, and let ϕ', ϕ'' be its components. Then ϕ' is a bijection satisfying $\phi'(1) = 1$ and $\phi'(ws) = \phi'(w)\phi''(s)$. It follows that $\phi'(s_1 \cdots s_d) = \phi''(s_1) \cdots \phi''(s_d)$ for all $s_1, \dots, s_d \in S$. This implies that ϕ' is a homomorphism, hence an automorphism, and that $\phi'(s) = \phi''(s)$ for all $s \in S$. Thus ϕ' is in $\text{Aut}(W, S)$ and ϕ is its image in $\text{Aut}(\Sigma, C)$. \square

Remark. The group $\text{Aut}(W, S)$ is quite easy to understand, in view of the Coxeter presentation of W : An element of this group is determined by giving a permutation π of S which is compatible with the Coxeter matrix, in the sense that $m(\pi(s), \pi(t)) = m(s, t)$ for all $s, t \in S$. More concisely, $\text{Aut}(W, S)$ is simply the group of automorphisms of the Coxeter diagram of (W, S) .

EXERCISES

1. Show that the full automorphism group of Σ is the semi-direct product $\text{Aut}_0 \Sigma \rtimes \text{Aut}(\Sigma, C)$. Hence $\text{Aut } \Sigma \approx W \rtimes \text{Aut}(W, S)$.

2. Suppose W is an irreducible finite reflection group. By looking at the list given in Chapter I of possible Coxeter diagrams, show that, with one exception, $\text{Aut}(W, S)$ is either trivial or of order 2. [The exception is the group of type D_4 .] So, with one exception, W is either the full automorphism group of Σ or a subgroup of index 2.

3. Specialize now to the case where W is the group of symmetries of a regular solid X , and note (again by looking at the list) that $\text{Aut}(W, S)$ is of order 2 if and only if X is self-dual. Explain this geometrically. More precisely, explain why an isomorphism from X to its dual induces a “label-reversing” automorphism of Σ . [HINT: Σ is the barycentric subdivision of the boundary of X .]

Finally, we complete the construction of foldings begun in §II.3B.

Corollary 3. *Let C_1 and C_2 be distinct adjacent chambers of Σ . Then there is an endomorphism ϕ of Σ with the following properties:*

- (1) ϕ is a retraction onto its image Φ .
- (2) Every chamber in Φ is the image of exactly one chamber not in Φ .
- (3) $\phi(C_2) = C_1$.

PROOF: We may assume that C_1 is the fundamental chamber C , in which case C_2 is necessarily sC for some $s \in S$. In this case we already constructed the first component ϕ' of the desired ϕ in §II.3B. Moreover, we saw in the course of that construction that ϕ' takes t -adjacent chambers to t -adjacent chambers for all $t \in S$. Thus we can take ϕ'' to be the trivial permutation of S . Everything should be clear now, except perhaps for (1), which can be expressed by saying that ϕ is *idempotent*, i.e., that $\phi^2 = \phi$. But ϕ^2 and ϕ are type-preserving chamber maps which agree on chambers, hence they agree on all simplices. \square

4 Half-spaces

We are ready, finally, to complete the circle of ideas begun in Chapter II. We will first develop, following Tits [56], a theory of half-spaces and reflections in a thin chamber complex; this theory is based on the notion of “folding” that we have already introduced informally. Once the basic properties of foldings have been laid out, it will be evident that a Coxeter complex $\Sigma(W, S)$ does indeed possess a rich supply of half-spaces and that W is generated by reflections of Σ . Finally, we will prove a theorem of Tits that characterizes the Coxeter complexes as the thin chamber complexes with a “rich supply” of half-spaces.

4A Foldings

Let Σ be an arbitrary thin chamber complex. Recall that an endomorphism ϕ of Σ is called *idempotent* if $\phi^2 = \phi$, or, equivalently, if ϕ is a retraction onto its image. A *folding* of Σ is an idempotent endomorphism ϕ such that for every chamber $C \in \phi(\Sigma)$ there is exactly one chamber $C' \notin \phi(\Sigma)$ with $\phi(C') = C$.

Let ϕ be a folding and let Φ be its image $\phi(\Sigma)$. It is easy to see that Φ is a chamber complex in its own right, since ϕ takes galleries to galleries. Let

Φ' be the subcomplex of Σ generated by the chambers not in Φ ; thus Φ' consists of all such chambers and their faces. By the definition of “folding”, then, ϕ induces a bijection

$$\text{Ch } \Phi' \xrightarrow{\approx} \text{Ch } \Phi,$$

where $\text{Ch } \Phi$ (resp. $\text{Ch } \Phi'$) denotes the set of chambers in Φ (resp. Φ').

We now define a function ϕ' on $\text{Ch } \Sigma$ by taking $\phi'| \text{Ch } \Phi'$ to be the identity and $\phi'| \text{Ch } \Phi$ to be the inverse of the bijection above. Intuitively, ϕ' is the “folding opposite to ϕ ”; but ϕ' is not really a folding, since it is only defined on chambers. I do not know whether, in the present generality, ϕ' can be extended to an endomorphism of Σ . Nevertheless, the following is true.

Lemma 1. *ϕ' takes adjacent chambers to adjacent chambers.*

PROOF: We may assume that the two given chambers C and D are distinct. If they are both in Φ' , there is nothing to prove. So assume that at least one of them, say C , is in Φ . Then $\phi'(C)$ is the unique $C' \in \Phi'$ such that $\phi(C') = C$. Let $A = C \cap D$ be the common face of C and D , and let A' be the face of C' such that $\phi(A') = A$. Finally, let D' be the chamber distinct from C' and adjacent to C' along A' . The following schematic diagram should help you keep all this notation straight:

$$\Phi \quad C \big|_A D \quad \bigg| \quad D' \big|_{A'} C' \quad \Phi'$$

The diagram shows the picture we would expect if C and D are both in Φ ; the big vertical bar in the middle is intended to suggest the “wall separating Φ from Φ' ”. You should visualize ϕ as folding from right to left along this wall and ϕ' as folding from left to right.

Since $\phi(D')$ is a chamber having A as a face, we must have either $\phi(D') = C$ or $\phi(D') = D$. Suppose first that $D' \in \Phi'$, as suggested by the picture. Then we cannot have $\phi(D') = C$, since then C' and D' would be distinct chambers in Φ' mapping to C . So we must have $\phi(D') = D$, which implies that $D \in \Phi$ and that $\phi'(D) = D'$. Thus $\phi'(D)$ is adjacent to $\phi'(C)$ in this case.

The other possibility is that $D' \in \Phi$. In this case the correct picture is presumably

$$C = D' \big| D = C',$$

but we must prove this rigorously. Since D' is in Φ , so is its face A' . Hence $A = \phi(A') = A'$. Thus all four of our chambers have the common face A . The thinness of Σ now implies that $\{C, D\} = \{C', D'\}$. Since $C \neq C'$ [because one is in Φ and the other is in Φ'], the only possibility is that $C = D' \in \Phi$ and $D = C' \in \Phi'$. Thus $\phi'(D) = D = C' = \phi'(C)$, so adjacency is again preserved. \square

Note that, as a consequence of this lemma, ϕ' takes galleries to galleries. In particular, it follows that Φ' is a chamber complex.

We now proceed to develop the basic properties of our folding ϕ and the associated function ϕ' and subcomplexes Φ and Φ' .

Lemma 2. *There exists a pair C, C' of distinct adjacent chambers with $C \in \Phi$ and $C' \in \Phi'$. For any such pair, we have $\phi(C') = C$ and $\phi'(C) = C'$.*

PROOF: Since $\text{Ch } \Phi$ and $\text{Ch } \Phi'$ are both non-empty, there is a gallery Γ which starts in Φ and ends in Φ' . Then Γ must cross from Φ to Φ' at some point, whence the first assertion. Suppose, now, that C and C' are as in the statement of the lemma, and let $A = C \cap C'$. Then $A < C \in \Phi$, so A is fixed by ϕ and hence $\phi(C')$ has A as a face. By thinness, we must have $\phi(C') = C$ or $\phi(C') = C'$. But the second possibility would imply $C' \in \Phi$, so $\phi(C') = C$. It now follows from the definition of ϕ' that $\phi'(C) = C'$. \square

Lemma 3. *Φ and Φ' are convex subcomplexes of Σ , in the sense that if Γ is a minimal gallery in Σ with both extremities in Φ (resp. Φ'), then Γ lies entirely in Φ (resp. Φ').*

PROOF: Suppose Γ is a minimal gallery with both extremities in Φ . If Γ is not contained in Φ , then it must cross from Φ to Φ' at some point. Thus there is a pair of consecutive chambers in Γ to which we can apply Lemma 2. But then $\phi(\Gamma)$ stutters. We can therefore eliminate the repetitions and get a shorter gallery with the same extremities as Γ , contradicting the minimality. A similar argument (using ϕ') works for Φ' . \square

Lemma 4. *Let C and C' be as in Lemma 2. Then*

$$\text{Ch } \Phi = \{ D \in \text{Ch } \Sigma : d(D, C) < d(D, C') \}$$

and

$$\text{Ch } \Phi' = \{ D \in \text{Ch } \Sigma : d(D, C) > d(D, C') \}.$$

In particular, no chamber of Σ is equidistant from C and C' .

(You should convince yourself that the last assertion is not vacuous, i.e., that there are thin chamber complexes in which a chamber D is equidistant from two adjacent chambers C, C' . The intuitive reason why it cannot happen in the present context is that the “wall” separating C from C' would have to cut through D , contradicting the fact that our two “halves” Φ and Φ' are subcomplexes.)

PROOF: Note that the right-hand sides of the two equalities to be proved are disjoint sets of chambers. Consequently, since Φ and Φ' partition the chambers of Σ , it suffices to prove that the left-hand sides are contained in the right-hand sides. Suppose, then, that we are given a chamber $D \in \Phi$, and let Γ be a minimal gallery from D to C' . Then, as before, Γ must cross from Φ to Φ' at some point, so we may fold it (i.e., apply ϕ to it) to obtain a stuttering gallery from D to $\phi(C') = C$. Hence $d(D, C) < d(D, C')$, as required. A similar argument, using ϕ' , proves the second inclusion. \square

Lemma 5. *Suppose C and C' are distinct adjacent chambers such that $\phi(C') = C$. Then ϕ is the unique folding taking C' to C .*

PROOF: Note first that we have $C \in \phi(\Sigma) = \Phi$ and $C' \in \Phi'$ [because $\phi(C') \neq C'$]. So Lemma 4 is applicable and yields a description of the two “halves” Φ and Φ' of Σ determined by ϕ . If ψ is a second folding with $\psi(C') = C$, then we can similarly apply Lemma 4 to obtain the same description of the halves of Σ determined by ψ . In particular, it follows that ψ , like ϕ , is the identity on Φ and maps $\text{Ch } \Phi'$ bijectively to $\text{Ch } \Phi$. We must show that ψ agrees with ϕ on all vertices of Φ' .

To begin with, we know that the two foldings both take C' to C and fix all vertices of the codimension 1 face $C \cap C'$ of C' ; hence they agree pointwise on C' (i.e., they agree on all vertices of C'). We will complete the proof by showing that ϕ and ψ continue to agree pointwise as we move away from C' along a non-stuttering gallery Γ in Φ' . It suffices to show that if ϕ and ψ agree pointwise on a chamber $D \in \Phi'$ then they agree pointwise on any chamber $E \in \Phi'$ which is adjacent to D (and distinct from D).

Let A be the common face $D \cap E$. Let $D_1 = \phi(D) = \psi(D)$, let $A_1 = \phi(A) = \psi(A)$, and let E_1 be the unique chamber distinct from D_1 and having A_1 as a face:

$$\Phi \quad E_1 \Big|_{A_1} D_1 \quad \Big| \quad D \Big|_A E \quad \Phi'$$

Then necessarily $\phi(E) = E_1 = \psi(E)$; for the only other possibility is that ϕ or ψ maps E to D_1 , contradicting the injectivity of ϕ and ψ on $\text{Ch } \Phi'$. And ϕ and ψ must agree pointwise on E , since they are already known to agree on all but one vertex of E . \square

Remark. The argument used in the previous two paragraphs will be called the *standard uniqueness argument*. It will be used repeatedly as we proceed. For pedagogical reasons, I prefer not to formalize the argument, since I think it is useful for you to have to think it through several more times. The basic idea to remember is the following: If a chamber map is known on all the vertices of one chamber, then you can often figure out what it has to do as you move away from that chamber along a gallery.

We will say that the folding ϕ is *reversible* if the function ϕ' defined above on chambers extends to a folding. Note that if C and C' are as in Lemma 5, then we have $\phi'(C) = C'$; so if ϕ is reversible, then the extension of ϕ' to a folding is unique: It is the folding of Σ taking C to C' . We will use the same symbol ϕ' for this extension, and we will call it the folding *opposite* to ϕ .

Lemma 6. *Let C and C' be distinct adjacent chambers with $\phi(C') = C$. Then ϕ is reversible if and only if there exists a folding taking C to C' . In this case there is an automorphism s of Σ such that $s|_{\Phi} = \phi'$ and $s|_{\Phi'} = \phi$. This automorphism is of order 2, and it can be characterized as the unique*

non-trivial automorphism of Σ which fixes $C \cap C'$ pointwise. Finally, the set of simplices of Σ fixed by s is the subcomplex $\Phi \cap \Phi'$ of Σ .

PROOF: We have already seen that if ϕ is reversible then the opposite folding ϕ' takes C to C' . Conversely, suppose there is a folding ϕ_1 such that $\phi_1(C) = C'$. Then we can apply Lemma 4 to ϕ_1 to deduce that ϕ_1 determines the same “halves” Φ and Φ' as ϕ (but with their roles reversed, i.e., Φ' is the image of ϕ_1). In particular, ϕ and ϕ_1 are both the identity on $H = \Phi \cap \Phi'$, so there is a well-defined endomorphism s of Σ with $s|_{\Phi} = \phi_1$ and $s|_{\Phi'} = \phi$. Note that H is the full fixed-point set of s ; for if $A \notin H$, say $A \notin \Phi$, then $s(A) = \phi(A) \neq A$.

It is clear that s maps $\text{Ch } \Phi$ bijectively to $\text{Ch } \Phi'$, and vice versa, so s is bijective on $\text{Ch } \Sigma$. Hence s^2 is bijective on $\text{Ch } \Sigma$. Since s^2 fixes C pointwise, the standard uniqueness argument is applicable and shows that s^2 is the identity. In particular, s is an automorphism.

We now prove that $\phi_1|_{\text{Ch } \Sigma} = \phi'$, and hence that ϕ is reversible. Since ϕ' and ϕ_1 are both the identity on $\text{Ch } \Phi'$, it suffices to consider chambers $D \in \Phi$. For any such D we have $D = s^2(D) = \phi(\phi_1(D))$, so $\phi_1(D)$ is the (unique) chamber in Φ' which is mapped by ϕ to D . Hence $\phi_1(D) = \phi'(D)$ by the definition of the latter.

Finally, to prove the characterization of s stated in the lemma, suppose that t is another non-trivial automorphism fixing $C \cap C'$ pointwise. Then t must interchange C and C' ; for otherwise t would have to fix them pointwise, and the standard uniqueness argument would show that t is trivial. Thus t agrees with s (pointwise) on C , and both are bijective on $\text{Ch } \Sigma$. We can therefore apply the standard uniqueness argument yet again to deduce that $s = t$. \square

We now introduce geometric language and summarize some of the results above in this language. A *half-space* of Σ is a subcomplex Φ which is the image of a reversible folding ϕ . In view of Lemmas 2 and 5, the folding ϕ is uniquely determined by Φ . The subcomplex Φ' generated by the chambers not in Φ is again a half-space, being the image of the opposite folding ϕ' ; it is called the half-space *opposite* to Φ .

The intersection $H = \Phi \cap \Phi'$ of two opposite half-spaces will be called the *wall* bounding Φ (or Φ'). Note that we can recover the pair of half-spaces $\{\Phi, \Phi'\}$ from the wall H and, in fact, from any simplex $A \in H$ which is of codimension 1 in Σ . To see this, it suffices to describe the pair of foldings $\{\phi, \phi'\}$ in terms of A : Let C_1 and C_2 be the chambers having A as a face. Then there is a unique folding ϕ_1 (resp. ϕ_2) such that $\phi_1(C_2) = C_1$ (resp. $\phi_2(C_1) = C_2$), and $\{\phi, \phi'\} = \{\phi_1, \phi_2\}$.

A wall H determines an automorphism $s = s_H$ by Lemma 6, which fixes H pointwise and interchanges the two half-spaces determined by H . For any $A \in H$ as in the previous paragraph, we can characterize s as the unique non-trivial automorphism of Σ that fixes A pointwise; in particular, s is the unique non-trivial automorphism fixing every simplex of H . We

call s the *reflection* of Σ with respect to H .

Finally, two chambers $C, C' \in \Sigma$ will be said to be *separated* by the wall H if one is in Φ and the other is in Φ' . If the two chambers are adjacent, Lemmas 2 and 5 imply that H is then the *unique* wall separating them.

In case Σ is a Coxeter complex $\Sigma(W, S)$, Corollary 3 in the previous section shows that every pair C_1, C_2 of distinct adjacent chambers is separated by a wall. For we have a folding taking C_2 to C_1 and also one taking C_1 to C_2 ; these foldings are therefore opposite to one another by Lemma 6 and determine a wall separating C_1 from C_2 . If C_1 and C_2 are C and sC for some $s \in S$, where C is the fundamental chamber, it is easy to see that the reflection associated to this wall is given by the action of s . It follows easily that the reflections of Σ determined by all possible walls are precisely the elements of W that we called reflections in Chapter II. Consequently, the abstract set \mathcal{H} of “walls” used in that chapter can be identified with the set of walls of Σ .

You should now be thoroughly convinced that Coxeter complexes possess a good theory of half-spaces. We will complete the chapter by showing that this property characterizes the Coxeter complexes among the thin chamber complexes.

EXERCISES

1. Let Φ be a half-space and s the associated reflection. If C and C' are chambers in Φ , show that $d(C, sC') > d(C, C')$. [HINT: Argue as in the proof of Lemma 4.]
2. You have now seen the standard uniqueness argument applied several times. Try to write down a lemma which includes all of these applications. [*Warning:* Unless you have incredible foresight, you can expect to have to modify your lemma one or more times as you see further applications of the argument. In fact, this might even happen in the next few pages.]

4B Characterization of Coxeter complexes

Now that we have begun considering abstract chamber complexes that are not necessarily given to us as complexes $\Sigma(W, S)$, it is convenient to slightly expand our previous terminology: From now on we will use the term *Coxeter complex* for any abstract simplicial complex Σ which is isomorphic to $\Sigma(W, S)$ for some Coxeter system (W, S) with S finite. This differs from our previous use of the term in that we do not assume that we are given a specific isomorphism $\Sigma \approx \Sigma(W, S)$ as part of the structure of Σ . In particular, no chamber of Σ has been singled out as “fundamental”.

The following theorem of Tits says, roughly speaking, that the Coxeter complexes can be characterized as the thin chamber complexes with “enough” half-spaces.

Theorem. *A thin chamber complex Σ is a Coxeter complex if and only if every pair of distinct adjacent chambers is separated by a wall.*

(*Note:* We can restate the condition of the theorem as follows: For every ordered pair C, C' of distinct adjacent chambers, there is a folding ϕ of Σ with $\phi(C') = C$. We don't need to specify here that ϕ is reversible; for this follows, as we saw above in the case of $\Sigma(W, S)$, from the existence of a folding taking C' to C .)

PROOF: We have already proven the “only if” part. Conversely, assume that every pair of distinct adjacent chambers is separated by a wall. Choose an arbitrary chamber C and let S be the set of reflections determined by the codimension 1 faces of C . Let $W \subseteq \text{Aut } \Sigma$ be the subgroup generated by S . We will prove that (W, S) is a Coxeter system and that $\Sigma \approx \Sigma(W, S)$.

The first observation is that W acts transitively on the chambers of Σ ; the proof of this is identical to the proof given in Chapter I for finite reflection groups. It follows that every codimension 1 simplex of Σ is W -equivalent to a face of C , and hence every reflection of Σ is W -conjugate to an element of S . This shows that the “reflections” in W , in the sense of Chapter II, are precisely the reflections of Σ obtained from the theory of half-spaces.

We can therefore identify the set \mathcal{H} used in Chapter II with the set of walls of Σ , and we can identify $\mathcal{H} \times \{\pm 1\}$ with the set of half-spaces of Σ . It is now a routine matter to verify condition (A) of Chapter II by using the W -action on the set of half-spaces. Thus (W, S) is a Coxeter system.

[If this argument seemed a little too quick, don't worry; we will give below a completely independent proof that (W, S) is a Coxeter system.]

To prove that $\Sigma \approx \Sigma(W, S)$, the crucial step is to calculate the stabilizers of the faces of C . We could simply repeat, essentially verbatim, the arguments which led to the analogous calculation for finite reflection groups in Chapter I. For the sake of variety, however, I will use a different method. This is actually a little longer, but it adds some geometric insight that we would not get by repeating the previous arguments. In particular, it leads to a simple geometric explanation of the deletion condition.

We now proceed with a sequence of observations that will lead, ultimately, to the desired calculation of stabilizers.

(a) Σ is labellable.

PROOF: Let \overline{C} be the subcomplex $\Sigma_{\leq C}$. It suffices to show that \overline{C} is a retract of Σ . The idea for showing this is to construct a retraction ρ by folding and folding and folding ..., until the whole complex Σ has been folded up onto \overline{C} .

To make this precise, let C_1, \dots, C_n be the chambers adjacent to C and distinct from it, and let ϕ_1, \dots, ϕ_n be the foldings such that $\phi_i(C_i) = C$. Let ψ be the composite $\phi_n \circ \dots \circ \phi_1$. I claim that $d(C, \psi(D)) < d(C, D)$ for any chamber $D \neq C$. To prove this, let $\Gamma : C, C', \dots, D$ be a minimal gallery from C to D ; we will show that $\psi(\Gamma)$ stutters. If $\phi_1(\Gamma)$ stutters, we

are done. Otherwise, the standard uniqueness argument shows that ϕ_1 fixes all the chambers of Γ pointwise. In this case, repeat the argument with ϕ_2 , etc. Eventually we will be ready to apply the folding ϕ_i which takes C' to C . If the previous foldings did not already make Γ stutter, then they have fixed Γ pointwise and the application of ϕ_i yields a stuttering gallery. This proves the claim.

It follows that, for any chamber D , $\psi^k(D) = C$ for k sufficiently large. Since ψ fixes C pointwise, this implies that the “infinite iterate” $\rho = \lim_{k \rightarrow \infty} \psi^k$ is a well-defined chamber map which retracts Σ onto \overline{C} . \square

It will be convenient to choose a fixed labelling λ with S as the set of labels, analogous to the canonical labelling that we used earlier in the chapter. To this end we label the vertices of C by setting $\lambda(v)$ equal to the reflection $s \in S$ which fixes the face of C opposite v ; we then extend this labelling to all of Σ by means of a retraction ρ of Σ onto \overline{C} . Note that this labelling λ has a property which by now should be very familiar: For any $s \in S$, the chambers C and sC are s -adjacent.

(b) *Foldings and reflections are type-preserving, hence all elements of W are type-preserving. Consequently, wC and wsC are s -adjacent for any $w \in W$ and $s \in S$.*

PROOF: A folding ϕ fixes at least one chamber pointwise, hence the induced map ϕ_* on labels is the identity (cf. §3 above). This proves that foldings are type-preserving, and everything else follows from this. \square

If $\Gamma : C_0, \dots, C_d$ is a non-stuttering gallery and H_i is the wall separating C_{i-1} from C_i , then, as usual, we will say that H_1, \dots, H_d are the walls crossed by Γ .

(c) *If $\Gamma : C_0, \dots, C_d$ is a minimal gallery, then the walls crossed by Γ are distinct and are precisely the walls separating C_0 from C_d . Hence the distance between two chambers is equal to the number of walls separating them.*

PROOF: Suppose H is a wall separating C_0 from C_d . Let Φ and Φ' be the corresponding half-spaces, say with $C_0 \in \Phi$ and $C_d \in \Phi'$. Then there must be some i with $1 \leq i \leq d$ such that $C_{i-1} \in \Phi$ and $C_i \in \Phi'$. Since Φ and Φ' are convex (Lemma 3 above), it follows that we have $C_0, \dots, C_{i-1} \in \Phi$ and $C_i, \dots, C_d \in \Phi'$. In other words, Γ crosses H exactly once. Now suppose H is a wall that does not separate C_0 from C_d . Then C_0 and C_d are both in the same half-space Φ , so the convexity of Φ implies that Γ does not cross H . \square

The crux of this proof, obviously, is the convexity of half-spaces, which in turn was based on the idea of using foldings to shorten galleries. We can now use this same idea to prove a geometric analogue of the deletion condition. In order to state it, we need to talk about the “type” of a gallery. If $\Gamma : C_0, \dots, C_d$ is a non-stuttering gallery, then the *type* of Γ is the

sequence $\mathbf{s} = (s_1, \dots, s_d)$ of labels such that C_{i-1} is s_i -adjacent to C_i for $i = 1, \dots, d$. [This notion of “type of a gallery” makes sense in any labelled chamber complex; we will use it again in later chapters.]

(d) *Let Γ be a non-stuttering gallery of type $\mathbf{s} = (s_1, \dots, s_d)$. If Γ is not minimal, then there is a non-stuttering gallery Γ' with the same extremities as Γ , such that Γ' has type $\mathbf{s}' = (s_1, \dots, \hat{s}_i, \dots, \hat{s}_j, \dots, s_d)$ for some $i < j$.*

PROOF: Since Γ is not minimal, (c) implies that the number of walls separating C_0 from C_d is less than d . Hence the walls crossed by Γ cannot all be distinct; for if a wall is crossed exactly once by Γ , then it certainly separates C_0 from C_d . We can therefore find a half-space Φ and indices i, j , with $1 \leq i < j \leq d$, such that C_{i-1} and C_j are in Φ but C_k is in the opposite half-space Φ' for $i \leq k < j$:

$$\begin{array}{ccc} & C_{i-1} & C_i \\ \Phi & \left| \begin{array}{c} \vdots \\ C_j \end{array} \right| & \Phi' \\ & C_j & C_{j-1} \end{array}$$

Let ϕ be the folding with image Φ . If we modify Γ by applying ϕ to the portion C_i, \dots, C_{j-1} , we obtain a gallery with the same extremities which stutters exactly twice:

$$C_0, \dots, C_{i-1}, \phi(C_i), \dots, \phi(C_{j-1}), C_j, \dots, C_d.$$

So we can delete C_{i-1} and C_j to obtain a non-stuttering gallery Γ' of length $d - 2$. The type \mathbf{s}' of Γ' is $(s_1, \dots, \hat{s}_i, \dots, \hat{s}_j, \dots, s_d)$ because ϕ is type-preserving. \square

(e) *The action of W is simply-transitive on the chambers of Σ .*

PROOF: We have already noted that the action is transitive. To prove that the stabilizer of C is trivial, note that if $wC = C$ then w fixes C pointwise, since w is type-preserving. But then $w = 1$ by the standard uniqueness argument. \square

It follows from (e) that we have a bijection $W \rightarrow \text{Ch } \Sigma$, given by $w \mapsto wC$. This yields the familiar 1-1 correspondence between non-stuttering galleries starting at C and words $\mathbf{s} = (s_1, \dots, s_d)$, where the gallery (C_i) corresponding to \mathbf{s} is given by $C_i = s_1 \cdots s_i C$ for $i = 0, \dots, d$. In view of (b), the type of this gallery is the sequence \mathbf{s} that we started with. So a direct translation of (d) into the language of group theory yields the deletion condition for (W, S) . This gives, as promised, a new proof that (W, S) is a Coxeter system.

(f) *The subcomplex $\overline{C} = \Sigma_{\leq C}$ is a fundamental domain for the action of W on Σ , in the sense that every simplex of Σ is W -equivalent to a unique $A \in \overline{C}$. Moreover, the stabilizer of the face of C of type $S - S'$ is the special subgroup $\langle S' \rangle$ of W .*

PROOF: The first assertion follows from the transitivity of W on the chambers, together with the fact that W is type-preserving. To prove the second,

let A be a face of C and let $\lambda(A) = S - S'$. It follows from the definition of λ that S' is the set of elements of S that fix A pointwise. In particular, the subgroup $W' = \langle S' \rangle$ stabilizes A . To prove that W' is the full stabilizer, suppose $wA = A$. We will show by induction on $l(w)$ that $w \in W'$. We may assume $w \neq 1$, so we can write $w = sw'$ with $s \in S$ and $l(w') < l(w)$. Our correspondence between words and galleries now implies that there is a minimal gallery of the form C, sC, \dots, wC . By (c), then, the wall H corresponding to s separates C from wC .

Let Φ be the half-space bounded by H which contains C . Then wC is in the opposite half-space $s\Phi$, so we have $w'C \in \Phi$. The equation $wA = A$ now yields

$$w'A = sA \in \Phi \cap s\Phi = H,$$

hence $A \in H$ and $w'A = A$. We therefore have $s \in S'$ [because s fixes A pointwise] and $w' \in W'$ by induction; thus $w = sw' \in W'$. \square

The desired isomorphism $\Sigma \approx \Sigma(W, S)$ is an easy consequence of (f). This completes the proof of the theorem. \square

IV

Buildings

The definition of “building”, to be given in the first section below, involves three axioms which are quite easy to state. It is not so easy, however, to motivate this definition. In particular, you will probably wonder how someone (namely, Tits) came up with these axioms. I will not attempt to answer this question now, but I will make some historical remarks in the next chapter (§V.4) which should make the definition seem less mysterious.

The terminology used in this subject is attributed by Tits to Bourbaki. In order to understand where it comes from, you need to interpret the word “chamber” that we have been using as meaning “room”. Thus Coxeter complexes are divided up into rooms by walls, and they are therefore called “apartments”. Buildings, then, are complexes which are built by putting apartments together. We now state the axioms, which specify the rules for putting the apartments together.

1 Definition and First Properties

A *building* is a simplicial complex Δ which can be expressed as the union of subcomplexes Σ (called *apartments*) satisfying the following axioms:

- (B0) *Each apartment Σ is a Coxeter complex.*
- (B1) *For any two simplices $A, B \in \Delta$, there is an apartment Σ containing both of them.*
- (B2) *If Σ and Σ' are two apartments containing A and B , then there is an isomorphism $\Sigma \rightarrow \Sigma'$ fixing A and B pointwise.*

Note that we can take both A and B to be the empty simplex in (B2); hence any two apartments are isomorphic. Note also that Δ is a chamber complex. For if C and C' are maximal simplices, then they are also maximal simplices of some apartment Σ by (B1), so they have the same dimension and are connected by a gallery.

Any collection \mathcal{A} of subcomplexes Σ satisfying the axioms will be called a *system of apartments* for Δ . Thus a building is a simplicial complex which admits a system of apartments. Note that we do *not* require that a building be equipped, as part of its structure, with a specific system of apartments. The reason for this is that it turns out that a building always admits a

canonical system of apartments. And in the important special case where the apartments are *finite* Coxeter complexes, it is even true that there is a *unique* system of apartments. We will prove both of these assertions later in the chapter (§§4 and 5, respectively).

Remarks

1. The definition of “building” given above is not the only one that is found in the literature. The complexes we have called buildings are sometimes called *weak buildings*, the term “building” being reserved for the case where Δ is *thick*. This means, by definition, that every codimension 1 simplex is a face of at least three chambers. If we confine ourselves to the thick case, then axiom (B0) can be considerably weakened. Namely, we need only assume that the apartments Σ are thin chamber complexes, and it then follows from (B1) and (B2) that they are in fact Coxeter complexes. The proof of this will be given in §7.

2. Axiom (B2) can be replaced by the following weaker axiom, which is simply the special case of (B2) in which one of the two given simplices is a chamber:

(B2') *If Σ and Σ' are apartments containing a simplex A and a chamber C , then there is an isomorphism $\Sigma \rightarrow \Sigma'$ fixing A and C pointwise.*

For suppose that (B1) and (B2') are known and that we are given an arbitrary pair of simplices A, B contained in two apartments Σ and Σ' . Choose chambers C and D with $A \leq C \in \Sigma$ and $B \leq D \in \Sigma'$, and choose an apartment Σ'' containing C and D . In view of (B2'), we have isomorphisms

$$\Sigma \xrightarrow{\sim} \Sigma'' \xrightarrow{\sim} \Sigma',$$

where the first isomorphism fixes C and B pointwise and the second isomorphism fixes A and D pointwise. The composite is then an isomorphism $\Sigma \rightarrow \Sigma'$ fixing A and B pointwise, so (B2) holds.

3. Axiom (B2'), in turn, is equivalent to the following axiom, which appears at first glance to be stronger:

(B2'') *If Σ and Σ' are two apartments with a common chamber, then there is an isomorphism $\Sigma \rightarrow \Sigma'$ fixing every simplex in $\Sigma \cap \Sigma'$.*

For suppose that (B2'') holds and that Σ and Σ' are apartments with a common chamber C . Then we have, for each $A \in \Sigma \cap \Sigma'$, an isomorphism $\phi_A : \Sigma \rightarrow \Sigma'$ fixing A and C pointwise. But our standard uniqueness argument (cf. §III.4) shows that there is at most one isomorphism from Σ to Σ' fixing C pointwise. So all the ϕ_A are equal to a single isomorphism ϕ , which therefore fixes the entire intersection $\Sigma \cap \Sigma'$.

Assume, for the remainder of this section, that Δ is a building and that \mathcal{A} is a fixed system of apartments.

Proposition 1. Δ is labellable. Moreover, the isomorphisms $\Sigma \rightarrow \Sigma'$ in axiom (B2) can be taken to be label-preserving.

PROOF: Fix an arbitrary chamber C , and label its vertices by some set I . If Σ is any apartment containing C , then [since Coxeter complexes are labellable] there is a unique labelling λ_Σ of Σ which agrees with the chosen labelling on C . For any two such apartments Σ, Σ' , the labellings λ_Σ and $\lambda_{\Sigma'}$ agree on $\Sigma \cap \Sigma'$; this follows from the fact that $\lambda_{\Sigma'}$ can be constructed as $\lambda_\Sigma \circ \phi$, where $\phi : \Sigma' \rightarrow \Sigma$ is the isomorphism fixing $\Sigma \cap \Sigma'$ [cf. (B2'')]. The various labellings λ_Σ therefore fit together to give a labelling λ defined on the union of the apartments containing C . But this union is all of Δ by (B1), so the first assertion of the proposition is proved.

To prove the second assertion, it suffices to consider the isomorphisms which occur in axiom (B2'). But such an isomorphism is automatically label-preserving, since it fixes a chamber pointwise. \square

Choose a fixed labelling λ of Δ by a set I . In view of the essential uniqueness of labellings, nothing we do will depend in any serious way on this choice. The labelling λ yields, for any apartment Σ , a *Coxeter matrix* $M = (m_{ij})_{i,j \in I}$, defined by

$$m_{ij} = \text{diam}(\text{lk}_\Sigma A),$$

where A is any simplex in Σ of type $I - \{i, j\}$ (cf. §III.2). Now Proposition 1 implies that any two apartments are isomorphic in a label-preserving way. Consequently:

Proposition 2. All apartments have the same Coxeter matrix M . \square

We will therefore call M the *Coxeter matrix* of Δ . Similarly, we can speak of the *Coxeter diagram* of Δ ; it is a graph with one vertex for each $i \in I$. Strictly speaking, we should be talking about the Coxeter matrix and diagram of the pair (Δ, \mathcal{A}) ; but we will show in §3 below that the matrix and diagram are really intrinsically associated to Δ and do not depend on the system of apartments \mathcal{A} .

The importance of the Coxeter matrix, of course, is that it completely determines the isomorphism type of the apartments. Let's spell this out in detail: Let W_M be the Coxeter group associated to M , with generators s_i ($i \in I$) and relations $(s_i s_j)^{m_{ij}} = 1$. Let Σ_M be the Coxeter complex $\Sigma(W_M, \{s_i\})$. It has a canonical labelling as in Chapter III, with I as the set of labels. [More precisely, the set of labels is the set of generators $\{s_i\}$; but this set is in 1-1 correspondence with I , so we may view I as the set of labels.] We can now deduce from Proposition 2:

Corollary. For any apartment Σ , there is a label-preserving isomorphism $\Sigma \approx \Sigma_M$.

PROOF: We can replace the labelling λ of Δ by a different one without affecting the truth of the statement to be proved. So we may assume that

λ has been constructed as follows: Choose an isomorphism $\phi : \Sigma \rightarrow \Sigma(W, S)$ for some Coxeter system (W, S) . Use ϕ to transport the canonical labelling of $\Sigma(W, S)$ to a labelling λ_Σ of Σ , with S as the set of labels. And now extend λ_Σ to a labelling λ of Δ . [To see that this last step is possible, simply take λ to be a labelling of Δ which agrees with λ_Σ on the vertices of one chamber of Σ ; then λ necessarily agrees with λ_Σ on all of Σ by the uniqueness of labellings of Σ .]

With this choice of λ , the result is essentially obvious. For the Coxeter matrix M of Σ is now equal to the Coxeter matrix of (W, S) by §III.2. So we can identify W_M with W and Σ_M with $\Sigma(W, S)$, and our original isomorphism ϕ is the desired label-preserving isomorphism. \square

Finally, we record one more simple consequence of the axioms. Recall that the study of local properties of Coxeter complexes consisted of a single result, which said that the link of a simplex in a Coxeter complex is again a Coxeter complex. The situation for buildings is similar:

Proposition 3. *If Δ is a building, then so is $\text{lk } A$ for any $A \in \Delta$.*

PROOF: Choose a fixed system of apartments \mathcal{A} for Δ . Given $A \in \Delta$, let \mathcal{A}' be the family of subcomplexes of $\text{lk}_\Delta A$ of the form $\text{lk}_\Sigma A$, where Σ is an element of \mathcal{A} containing A . Any such subcomplex is a Coxeter complex by the result cited above. So it remains to verify (B1) and (B2). Given $B, B' \in \text{lk}_\Delta A$, we can join them with A to obtain simplices $A \cup B$ and $A \cup B'$ in Δ . Since Δ satisfies (B1), there is an apartment Σ containing both of these simplices. Hence $\text{lk}_\Sigma A$ is an element of \mathcal{A}' containing B and B' . This proves that \mathcal{A}' satisfies (B1), and the proof of (B2) is similar. \square

2 Examples

Almost all of the examples in this section will be defined as flag complexes, so you should review the definition of the latter before proceeding (cf. §B of the appendix to Chapter I).

Let P be a set with an “incidence” relation as in the appendix just cited. Assume, in addition, that P is partitioned into non-empty subsets P_0, P_1, \dots, P_{n-1} . Elements of P_i are said to have *type i* . Or, to use more intuitive language, elements of $P_0, P_1, P_2 \dots$ might be called points, lines, planes, etc. If the incidence relation has the property that two elements of the same type are never incident unless they are equal, then we will call P (together with the partition and incidence relation) an *n -dimensional incidence geometry*.


If $n = 1$, then the geometry just consists of a set of points, with no further structure. If $n = 2$, then P is a “plane geometry”, consisting of points and lines, with some points declared to be incident to some lines. If $n = 3$, there are points, lines, and planes. And so on.

In practice, of course, one is interested in incidence geometries which are subject to certain axioms, such as the axioms for projective geometry or some other kind of geometry. We will see below that different types of geometries correspond to different types of buildings (where the “type” of a building is determined by its Coxeter matrix).

We proceed now to the examples, starting with a case which is trivial but nonetheless instructive.

1. Suppose Δ is a building of rank 1 (dimension 0). Then every apartment must be a 0-sphere S^0 , since this is the only rank 1 Coxeter complex. In particular, Δ must have at least two vertices. Conversely, a rank 1 complex with at least two vertices is a building (with every 2-vertex subcomplex as an apartment). Thus the rank 1 buildings are precisely the flag complexes of the 1-dimensional incidence geometries with at least 2 points. [It is, of course, reasonable to demand that a 1-dimensional geometry, or “line”, have at least 2 points. In fact, one often even demands that there be at least 3 points, which is equivalent to requiring the flag complex Δ to be thick.]

2. Suppose Δ is a building of rank 2 (dimension 1). Then an apartment Σ must be a $2m$ -gon for some m ($2 \leq m \leq \infty$). We will draw the Coxeter diagram as



which should be interpreted as



if $m = 2$ and as



if $m = 3$.

Let's begin with the case $m = 2$. Then every apartment is a quadrilateral:




(The two types of vertices shown here indicate a labelling of Δ .) It follows easily from the building axioms that every vertex of type \bullet is connected by an edge to every vertex of type \circ . Hence Δ is the flag complex of a 2-dimensional incidence geometry in which every point is incident to every line. Conversely, the flag complex of such a geometry is always a rank 2 building (with $m = 2$), provided that the geometry has at least two points and at least two lines.

Note that we can also describe Δ as the join of two rank 1 buildings. This suggests a general fact, which you might want to prove before continuing:

EXERCISE

If Δ is a building whose Coxeter diagram is disconnected, show that Δ is canonically the join of lower-dimensional buildings, one for each component of the diagram.

Returning now to Example 2, suppose next that $m = 3$. Then every apartment is a hexagon, which we may draw as the barycentric subdivision of a triangle:



Now that we have gotten used to thinking about flag complexes of incidence geometries, this picture suggests a configuration of three lines in a plane (one line for each \circ), whose pairwise intersections yield three points of the plane (one for each \bullet). If you have studied projective geometry, then you know that three lines in “general position” in a projective plane always yield such a configuration. So it is reasonable to guess that Δ is the flag complex of a projective plane. Before proceeding further, let’s recall the definition of the latter:


A *projective plane* is a 2-dimensional incidence geometry satisfying the following three axioms:

- (1) Any two points are incident to a unique line.
- (2) Any two lines are incident to unique point.
- (3) There exist three non-collinear points.

With this definition, it is indeed the case that our building Δ is the flag complex of a projective plane. You should try to prove this as an exercise. [The exercise is not entirely routine; if you get stuck, you’ll see it again at the end of §3, at which point it should be easier.] Conversely, the flag complex of a projective plane is a building, with one apartment for every triangle in the projective plane (where a triangle is a configuration of three non-collinear points and the three lines they determine). This converse is a routine exercise, which you should do.

The most familiar example of a projective plane is the projective plane over a field k . By definition, the set P_0 of “points” is the set of 1-dimensional subspaces of the 3-dimensional vector space k^3 ; the set P_1 of “lines” is the set of 2-dimensional subspaces of k^3 ; and “incidence” is given by inclusion, i.e., a point $x \in P_0$ is incident to a line $L \in P_1$ if $x \subset L$ as subspaces of k^3 . If you haven’t seen this before, you should pause to verify the axioms.

It is now easy to construct concrete examples of buildings. Let P be the projective plane over \mathbf{F}_2 , for instance, where \mathbf{F}_2 is the field with two elements. Then P has 7 points (each on exactly 3 lines) and 7 lines (each containing exactly 3 points). The resulting flag complex Δ is a thick building with 14 vertices and 21 edges. We will see in Exercise 1 below that the points of P can be put in 1-1 correspondence with the 7th roots of unity ζ^j ($\zeta = e^{2\pi i/7}$, $j = 0, \dots, 6$) in such a way that the lines of P are the triples $\{\zeta^j, \zeta^{j+1}, \zeta^{j+3}\}$, $j = 0, \dots, 6$. This leads to the picture of Δ shown below. You might find it instructive to locate some of the apartments (there are 28 of them) and to verify some cases of the building axioms.



The flag complex of the projective plane over \mathbb{F}_2^*

Remark. This picture is misleading in one respect; namely, it fails to reveal how much symmetry Δ has. You can see from the picture that Δ admits an action of the dihedral group D_{14} , but in fact $\text{Aut } \Delta$ is of order 336. The subgroup $\text{Aut}_0 \Delta$ of type-preserving automorphisms is $\text{GL}_3(\mathbb{F}_2)$, which is the simple group of order 168.

Continuing with Example 2, one could analyze in a similar way the buildings corresponding to $m = 4, 5, 6, \dots$. Each value of m corresponds to a particular type of plane geometry. When $m = 4$, for example, the associated geometry is something called “polar geometry”, in which there do not exist triangles but there do exist lots of quadrilaterals. Every quadrilateral in the polar plane yields an apartment in the flag complex, this apartment being an octagon (or barycentrically subdivided quadrilateral).

Finally, the case $m = \infty$ also has a simple interpretation. Namely, buildings of this type are simply trees with no endpoints (where an endpoint of a tree is a vertex which is on only one edge). To see that such a tree is a building, simply take the apartments to be all possible subcomplexes which are lines (i.e., ∞ -gons); the verification of the building axioms is a routine matter. The converse, that every building of this type is a tree, is more challenging. We will prove it in the next section, but you might want to try it on your own first.

The two remaining examples are intended to give you a brief glimpse of some higher-dimensional buildings. Details, which can be found in Tits [56], will be omitted. We will, however, give some details about these and other examples in the next chapter, from the point of view of group theory rather

*Reprinted from p. 424 of “Self-dual Configurations and Regular Graphs,” H. S. M. Coxeter, *Bulletin of the American Mathematical Society* (1950), Volume 56, Pages 413–455, by permission of the American Mathematical Society. The figure appears also on p. 118 of Coxeter’s book “Twelve Geometric Essays” (Southern Illinois University Press, Carbondale, 1968), where the whole article is reprinted with corrections.

than incidence geometry. See also Exercise 2 below.

3. If P is an n -dimensional projective geometry, then its flag complex is a rank n building of type A_n , i.e., having Coxeter diagram

$$\circ \text{---} \circ \text{---} \circ \cdots \text{---} \circ \text{---} \circ \quad (n \text{ vertices}).$$

Every apartment is isomorphic to the barycentric subdivision of the boundary of an n -simplex, and there is one such apartment for every frame in the projective space (where a frame is a set of $n+1$ points in general position).

Conversely, every building of type A_n is the flag complex of a projective geometry.

When $n = 2$, this example reduces to the case $m = 3$ of Example 2.

4. If P is an n -dimensional polar geometry, then its flag complex is a rank n building of type B_n , i.e., having Coxeter diagram

$$\circ \text{---} \circ \text{---} \circ \cdots \text{---} \overset{4}{\circ} \text{---} \circ \quad (n \text{ vertices}).$$

Every apartment is isomorphic to the barycentric subdivision of the boundary of an n -cube (or n -dimensional hyperoctahedron), and there is one such apartment for every “polar frame” in the given polar space.

Conversely, every building of type B_n is the flag complex of a polar geometry.

When $n = 2$, this example reduces to the case $m = 4$ of Example 2.

EXERCISES

1. (a) Let V be a 3-dimensional vector space over \mathbf{F}_2 , and let $P = P(V)$ be the plane geometry in which the points are the elements of $V - \{0\}$ and the lines are the triples $\{u, v, w\}$ with $u + v + w = 0$. Show that P is isomorphic to the projective plane over \mathbf{F}_2 .

(b) Let $\Delta(V)$ be the flag complex of $P(V)$, with its canonical labelling. If V^* is the dual of V , show that the correspondence between subspaces of V and subspaces of V^* induces a label-reversing isomorphism $\Delta(V) \approx \Delta(V^*)$. Consequently, any isomorphism $V \rightarrow V^*$ induces a type-reversing automorphism of $\Delta(V)$. If the isomorphism $V \rightarrow V^*$ comes from a non-degenerate symmetric bilinear form on V , show that the resulting automorphism of $\Delta(V)$ is an involution (i.e., is of order 2).

(c) Let V be the field \mathbf{F}_8 , viewed as a vector space over \mathbf{F}_2 . Show that there is a 7th root of unity $\zeta \in \mathbf{F}_8$ such that the lines in $P = P(V)$ are the triples $L_i = \{\zeta^i, \zeta^{i+1}, \zeta^{i+3}\}$, $i \in \mathbf{Z}/7\mathbf{Z}$. [HINT: The polynomial $x^3 + x + 1$ is irreducible over \mathbf{F}_2 .]

(d) With $V = \mathbf{F}_8$ as in (c), recall that there is a non-degenerate symmetric bilinear form on V given by $\langle x, y \rangle = \text{tr}(xy)$, where $\text{tr} : \mathbf{F}_8 \rightarrow \mathbf{F}_2$ is the trace. This induces a type-reversing involution τ of $\Delta = \Delta(V)$ by (b). Show that τ is given on vertices by $\zeta^i \mapsto L_{6-i}$. [HINT: ζ , ζ^2 , and ζ^4 all have the same minimal polynomial, from which you can read off that they have trace 0.] Describe τ in terms of the picture of Δ above.

2. Let V be a vector space of finite dimension $n \geq 2$ over an arbitrary field. The *projective geometry* associated to V consists of the non-zero proper subspaces of V , two such being called *incident* if one is contained in the other. (This is an example of an $(n-1)$ -dimensional projective geometry.) The purpose of this exercise is to prove that the flag complex Δ of this geometry is a building.

By a *frame* in V we will mean a set $\mathcal{F} = \{L_1, \dots, L_n\}$ of 1-dimensional subspaces of V such that $V = L_1 \oplus \dots \oplus L_n$. Given such a frame, consider the set of subspaces $V' \subset V$ such that V' is spanned by a non-empty proper subset of \mathcal{F} . Let $\Sigma = \Sigma(\mathcal{F})$ be the subcomplex of Δ consisting of flags of such subspaces. Call a subcomplex Σ of this form an *apartment*.

(a) Show that each apartment is a Coxeter complex of type A_{n-1} . [HINT: See the exercise in §I.5H.]

(b) Let C and C' be maximal simplices of Δ with vertex sets

$$V_1 \subset \dots \subset V_{n-1} \quad \text{and} \quad V'_1 \subset \dots \subset V'_{n-1},$$

respectively. Set $V_0 = V'_0 = 0$ and $V_n = V'_n = V$, and view $\{V_i\}_{0 \leq i \leq n}$ and $\{V'_i\}_{0 \leq i \leq n}$ as composition series for V . According to the Jordan-Hölder theorem, there is a permutation π of $\{1, \dots, n\}$ such that $V'_i/V'_{i-1} \approx V_j/V_{j-1}$ if $j = \pi(i)$. This is of course a triviality in the present context of vector spaces; but we need to review how the proof of the Jordan-Hölder theorem yields a canonical π and canonical isomorphisms $V'_i/V'_{i-1} \approx V_j/V_{j-1}$.

For each $i \in \{1, \dots, n\}$, the composition series $\{V_j\}_{0 \leq j \leq n}$ induces a filtration of V'_i/V'_{i-1} . [First intersect with V'_i , then take images mod V'_{i-1} .] Since V'_i/V'_{i-1} is one-dimensional, this filtration must be trivial, i.e., only one of the successive quotients is non-trivial. Define $\pi(i)$ to be the index j such that the j th quotient is non-trivial. Equivalently, $j = \pi(i)$ is characterized by the property that

$$V'_{i-1} + (V'_i \cap V_k) = \begin{cases} V'_{i-1} & \text{for } k < j \\ V'_i & \text{for } k \geq j. \end{cases}$$

The resulting function $\pi = \pi(C, C') : \{1, \dots, n\} \rightarrow \{1, \dots, n\}$ is called the *Jordan-Hölder permutation* associated to the pair C, C' . Show that π is indeed a permutation. More precisely, show that $\pi(C, C')$ and $\pi(C', C)$ are inverses of one another and that, if $\pi(i) = j$, there are isomorphisms

$$\frac{V'_i}{V'_{i-1}} \xleftarrow{\approx} \frac{V'_i \cap V_j}{(V'_{i-1} \cap V_j) + (V'_i \cap V_{j-1})} \xrightarrow{\approx} \frac{V_j}{V_{j-1}}$$

induced by inclusions. [HINT: All of this can be extracted from the proof of the Jordan-Hölder theorem as given in many standard texts (e.g., [34], §3.3). But you might prefer to just prove it directly. To show, for instance, that $V'_{i-1} \cap V_j \subseteq V_{j-1}$, suppose this is false; then $V_j = V_{j-1} + (V'_{i-1} \cap V_j)$. Intersect both sides with V'_i to obtain $V'_i \cap V_j = (V'_i \cap V_{j-1}) + (V'_{i-1} \cap V_j) \subseteq V'_{i-1}$, contradicting the definition of j .]

(c) Deduce from (b) that axiom (B1) holds. [HINT: Given maximal simplices C and C' , find a frame \mathcal{F} by choosing, for each i, j as above, a suitable $L_j \subseteq V'_i \cap V_j$.]

(d) Complete the proof that Δ is a building by verifying (B2''). [HINT: Any chamber in $\Sigma(\mathcal{F})$ determines a canonical ordering of the n elements of \mathcal{F} . So a chamber in two apartments determines a canonical isomorphism between the apartments.]

(e) Let W be the symmetric group on n letters, i.e., the Coxeter group of type A_{n-1} . As a byproduct of the proof that Δ is a building, we have obtained a function $\text{Ch } \Delta \times \text{Ch } \Delta \rightarrow W$, given by $(C, C') \mapsto \pi(C, C')$. Can you guess what the geometric meaning of this function is? [HINT: What is $\pi(C, C')$ if C and C' are i -adjacent, i.e., if $V_j = V'_j$ for $j \neq i$?] We will come back to this in Exercise 3(e) of §4.

3 Retractions

Retractions, as we saw in Chapter IV, can be quite useful technical tools. In this section we will establish the existence and formal properties of retractions of a building onto its apartments.

Assume throughout this section that Δ is a building and that \mathcal{A} is an arbitrary system of apartments.

Proposition 1. *Every apartment Σ is a retract of Δ .*

PROOF: This is very similar to the proof of labellability: Fix a chamber C of the given apartment Σ , and consider all the apartments Σ' which contain C . For any such Σ' there is a unique isomorphism $\phi_{\Sigma'} : \Sigma' \rightarrow \Sigma$ which fixes C pointwise, where the existence follows from axiom (B2) and the uniqueness is proved by the standard argument. For any two such apartments Σ', Σ'' , the isomorphisms $\phi_{\Sigma'}$ and $\phi_{\Sigma''}$ agree on $\Sigma' \cap \Sigma''$. This follows from the fact that we can construct $\phi_{\Sigma''}$ by composing $\phi_{\Sigma'}$ with the isomorphism $\Sigma'' \rightarrow \Sigma'$ which fixes every simplex of $\Sigma' \cap \Sigma''$ [cf. (B2'')]. The various isomorphisms $\phi_{\Sigma'}$ therefore fit together to give a chamber map $\rho : \Delta \rightarrow \Sigma$, and ρ is a retraction since ϕ_{Σ} is the identity. \square

One useful consequence of this is that combinatorial distances between chambers of Δ can be computed in terms of the distance functions on apartments, which we understand reasonably well:

Corollary 1. *Let C and D be chambers of Δ , and let Σ be any apartment containing C and D . Then $d_{\Delta}(C, D) = d_{\Sigma}(C, D)$. Consequently, the diameter of Δ is equal to the diameter of any apartment.*

PROOF: Suppose Γ is a minimal gallery in Σ from C to D . Then Γ is also minimal in Δ ; for if there were a shorter gallery in Δ , then we could get a shorter one in Σ by applying a retraction. This proves the first assertion. As an immediate consequence, we have $\text{diam } \Sigma \leq \text{diam } \Delta$. To prove the opposite inequality, let C' and D' be arbitrary chambers of Δ and let Σ' be an apartment containing them. Then we have

$$d_{\Delta}(C', D') = d_{\Sigma'}(C', D') \leq \text{diam } \Sigma'.$$

But $\Sigma \approx \Sigma'$, so $\text{diam } \Sigma' = \text{diam } \Sigma$ and hence $\text{diam } \Delta \leq \text{diam } \Sigma$. \square

Corollary 1 can be used to prove the important fact that the Coxeter matrix $M = (m_{ij})_{i,j \in I}$ of Δ , as defined in §1, really is an invariant of Δ :

Corollary 2. *The Coxeter matrix M depends only on Δ , not on the system of apartments. It is given by*

$$m_{ij} = \text{diam}(\text{lk}_\Delta A),$$

where A is any simplex with $\lambda(A) = I - \{i, j\}$.

PROOF: Suppose $\lambda(A) = I - \{i, j\}$, and let Σ be any apartment containing A . Then we have $m_{ij} = \text{diam}(\text{lk}_\Sigma A)$ by definition. But $\text{lk}_\Sigma A$ is an apartment in the building $\text{lk}_\Delta A$, so it has the same diameter as the latter by Corollary 1. This proves the second assertion, and the first assertion follows at once. \square

As an example, suppose that Δ is 1-dimensional and that the apartments are $2m$ -gons as in Example 2 of the previous section. Then *every* system of apartments in Δ must consist of $2m$ -gons (for the same $m = \text{diam } \Delta$).

Returning now to the general study of retractions, note that the proof of Proposition 1 actually yields, for any apartment Σ and any chamber $C \in \Sigma$, a *canonical* retraction $\rho = \rho_{\Sigma, C} : \Delta \rightarrow \Sigma$. We call ρ the retraction onto Σ *centered* at C . It can be characterized as the unique chamber map $\Delta \rightarrow \Sigma$ which fixes C pointwise and maps every apartment containing C isomorphically onto Σ .

This characterization makes it appear that ρ depends on the apartment system \mathcal{A} . But part (3) of the following proposition gives a different characterization, which shows that ρ depends only on Σ and C , not on \mathcal{A} .

Proposition 2. *The retraction $\rho = \rho_{\Sigma, C}$ has the following properties.*

- (1) *For any face $A \leq C$, $\rho^{-1}(A) = \{A\}$.*
- (2) *ρ preserves distances from C , i.e., $d(C, \rho(D)) = d(C, D)$ for any chamber $D \in \Delta$.*
- (3) *ρ is the unique chamber map $\Delta \rightarrow \Sigma$ which fixes C pointwise and preserves distances from C .*


PROOF: (1) Suppose $B \in \Delta$ is a simplex such that $\rho(B) = A \leq C$. Choose an apartment Σ' containing both B and C . Then $\rho|_{\Sigma'}$ is an isomorphism, and it maps both A and B to A . Hence $B = A$.

(2) Let D be a chamber and let Σ' be an apartment containing C and D . Since $\rho|_{\Sigma'}$ is an isomorphism, we have $d_\Sigma(C, \rho(D)) = d_{\Sigma'}(C, D)$. In view of Corollary 1 above, we can delete the subscripts to obtain $d(C, \rho(D)) = d(C, D)$, as required.

(3) Suppose $\phi : \Delta \rightarrow \Sigma$ is another chamber map which fixes C pointwise and preserves distances from C . Then ϕ and ρ must both take any minimal gallery in Δ starting at C to a minimal gallery in Σ . In particular, the image galleries must be non-stuttering. But this is all that is needed to make the standard uniqueness argument go through, hence $\phi = \rho$. \square

Property (2) above enables one, in practice, to figure out what ρ looks like. Suppose, for example, that Δ is a tree as in Example 2 of the previous

section. Then Σ is a triangulated line and C is an edge of Σ . Let v and w be the vertices of C . Then it follows easily from (2) that ρ simply “flattens Δ out” onto Σ , with the part of Δ closer to v than to w going to the corresponding part of Σ , and similarly for the part closer to w :



We close this section by remarking that the retractions ρ make it easy to complete the discussion of 1-dimensional buildings that we began in Example 2 of the previous section. Let's prove, for instance, that a building Δ of type $\circ\text{---}\infty\text{---}\circ$ is a tree.

Suppose, to the contrary, that Δ contains a cycle, i.e., a subcomplex Z which is a k -gon for some k with $3 \leq k < \infty$. Let C be a chamber in Z , with vertices v and w , let Σ be any apartment containing C , and let ρ be the retraction $\rho_{\Sigma, C}$. As we traverse Z starting at C (thought of as oriented from v to w , say), the image under ρ is a closed curve in Σ passing through the vertices v, w, \dots and never passing through w again before returning to v [cf. part (1) of Proposition 2]. But Σ is a line, so there is no way for the curve to get back to v without passing through w . This contradiction shows that Z cannot exist.

EXERCISES

1. The *girth* of a 1-dimensional simplicial complex Δ is the smallest integer $k \geq 3$ such that Δ contains a k -gon. (Or, if Δ is a tree, the girth is defined to be ∞ .) Arguing as above, show that a building Δ of type $\circ\text{---}^m\text{---}\circ$ with $m < \infty$ has girth $2m$. In case $m = 3$, deduce that Δ is the flag complex of a projective plane.

2. Let Δ be a connected labellable 1-dimensional simplicial complex in which every vertex is a face of at least two edges. Show that Δ is a building if and only if Δ has diameter m and girth $2m$ for some m with $2 \leq m \leq \infty$. [“Diameter” here means the supremum of the distances between vertices.]

4 The Complete System of Apartments

We have seen a number of cases where something that seemed *a priori* to depend on a choice of apartment system \mathcal{A} turned out to be independent of \mathcal{A} . We now prove the ultimate result of this type; it can be viewed as saying that all systems of apartments in a given building are compatible with one another:

Theorem. *If Δ is a building, then the union of any family of apartment systems is again an apartment system. Consequently, Δ admits a largest system of apartments.*

PROOF: It is obvious that (B0) and (B1) hold for the union, so the only problem is to prove (B2). We will work with the variant (B2''). Suppose,

then, that Σ and Σ' are apartments in different apartment systems and that $\Sigma \cap \Sigma'$ contains at least one chamber. We must find an isomorphism $\Sigma' \rightarrow \Sigma$ which fixes every simplex of $\Sigma \cap \Sigma'$.

Choose an arbitrary chamber $C \in \Sigma \cap \Sigma'$. There are then two obvious candidates for the desired isomorphism $\Sigma' \rightarrow \Sigma$. On the one hand, we know by the previous section that Σ and Σ' have the same Coxeter matrix M , so we can find a label-preserving isomorphism $\phi : \Sigma' \rightarrow \Sigma$ by the Corollary to Proposition 2 of §1. And we can certainly choose ϕ so that $\phi(C) = C$, since the group of type-preserving automorphisms of Σ is transitive on the chambers. It then follows that ϕ fixes C pointwise. Unfortunately, it is not obvious that ϕ fixes every simplex of $\Sigma \cap \Sigma'$.

The other candidate is provided by the theory of retractions. Namely, let ρ be the retraction $\rho_{\Sigma, C}$ and let $\psi : \Sigma' \rightarrow \Sigma$ be the restriction of ρ to Σ' . Then ψ obviously fixes every simplex of $\Sigma \cap \Sigma'$, simply because ρ is a retraction onto Σ . But it is not obvious that ψ is an isomorphism. [If you're tempted to say that ψ is an isomorphism by the construction of ρ , recall that we don't know that Σ' is part of an apartment system containing Σ ; indeed, that's what we're trying to prove!]

To complete the proof, we will show by the standard uniqueness argument that ϕ and ψ are in fact the same map, which therefore has all the required properties. Since ϕ and ψ both fix C pointwise, the standard argument will go through if we can show that $\phi(\Gamma)$ and $\psi(\Gamma)$ are non-stuttering for any minimal gallery Γ in Σ' starting at C . This is clear for $\phi(\Gamma)$ since ϕ is an isomorphism. And it is true for $\psi(\Gamma)$ because of two facts proved in the previous section: (a) Γ is still minimal when viewed as a gallery in Δ ; and (b) ρ preserves distances from C . \square

The maximal apartment system will be called the *complete* system of apartments. It consists, then, of all subcomplexes $\Sigma \subseteq \Delta$ such that Σ is in some apartment system \mathcal{A} .

Remark. This description of the complete apartment system is not very informative. Here are two characterizations that are more useful. Let Σ be a chamber subcomplex of Δ , i.e., a subcomplex which is a chamber complex in its own right and has the same dimension as Δ (so that $\text{Ch } \Sigma \subseteq \text{Ch } \Delta$). Let \mathcal{A} be the complete system of apartments. Then:

- (1) $\Sigma \in \mathcal{A}$ if and only if Σ is isomorphic in a label-preserving way to Σ_M .
- (2) Assume Σ is thin. Then $\Sigma \in \mathcal{A}$ if and only if Σ is convex in Δ , in the sense that any minimal gallery in Δ with both extremities in Σ is entirely contained in Σ .

The sufficiency of each of the stated conditions will be proved in the exercises below. The necessity has already been proved in the case of (1) and will now be proved for (2):

Proposition. *Every apartment Σ is convex.*

PROOF: Let $\Gamma : C_0, \dots, C_d$ be a minimal gallery with $C_0, C_d \in \Sigma$. If Γ is not contained in Σ , then there is an index i with $C_{i-1} \in \Sigma$ and $C_i \notin \Sigma$. Let C be the chamber of Σ distinct from C_{i-1} and having $C_{i-1} \cap C_i$ as a face, and let ρ be the retraction $\rho_{\Sigma, C}$. Then $\rho(C_i) = C_{i-1}$ [why?], so $\rho(\Gamma)$ stutters and has the same extremities as Γ . This contradicts the minimality of Γ . \square

Remark. The same proof yields a stronger result. Recall that the distance $d(C, A)$ between a chamber C and a simplex A is defined to be the minimal length d of a gallery $\Gamma : C_0, \dots, C_d$ with $C_0 = C$ and $C_d \geq A$. Any gallery Γ which achieves this minimum is said to be *stretched* from C to A . [You should draw some low-dimensional pictures to get a feeling for what this means.] Suppose, now, that A and C are in an apartment Σ . Then one can prove, exactly as above, that every gallery stretched from C to A is contained in Σ .

EXERCISES

Most of the results in these exercises are taken from Tits [59]. For the benefit of readers who want to see proofs of these results but do not want to attempt them as exercises, I have provided substantial “hints”, which in many cases are almost complete solutions, immediately following the set of exercises.

1. Let Δ be a labelled building with Coxeter matrix M . Recall that the labelling enables one to speak of the *type* of any non-stuttering gallery $\Gamma : C_0, \dots, C_d$ (cf. §III.4B); it is a sequence $\mathbf{i} = (i_1, \dots, i_d)$ of labels. It will be convenient to identify \mathbf{i} with the corresponding word $\mathbf{s} = (s_{i_1}, \dots, s_{i_d})$ in the generators of W_M . In particular, it makes sense to ask whether \mathbf{i} is *reduced*, or is a *reduced decomposition* of an element $w \in W_M$. Prove that Γ is minimal if and only if its type \mathbf{i} is reduced.

2. Prove the characterization (1) of the complete apartment system.

3. Recall that if (W, S) is a Coxeter system, then the combinatorial distance function on $\Sigma = \Sigma(W, S)$ is given by $d(\{w\}, \{w'\}) = l(w^{-1}w')$. This suggests that we consider the “ W -valued distance function” $\delta : \text{Ch } \Sigma \times \text{Ch } \Sigma \rightarrow W$ defined by $\delta(\{w\}, \{w'\}) = w^{-1}w'$. The function δ turns out to be an extremely useful refinement of d . The purpose of this exercise is to explore its geometric meaning and define an analogous function on any building Δ .

(a) Let Σ be a labelled Coxeter complex with Coxeter matrix M . Show that there is a function $\delta : \text{Ch } \Sigma \times \text{Ch } \Sigma \rightarrow W_M$, characterized by the following property: If $\Gamma : C_0, \dots, C_d$ is a non-stuttering gallery of type \mathbf{i} , then $\delta(C_0, C_d)$ is the element of W_M represented by the word \mathbf{i} .

Assume, for the remainder of this exercise, that Δ is a labelled building with Coxeter matrix M .

(b) Show that the functions δ defined in (a) on the apartments of Δ are compatible with one another and hence give a well-defined function

$$\delta : \text{Ch } \Delta \times \text{Ch } \Delta \rightarrow W_M.$$

This function is characterized by the following property: If $\Gamma : C_0, \dots, C_d$ is a minimal gallery of type \mathbf{i} , then $\delta(C_0, C_d)$ is the element of W_M represented by \mathbf{i} .

(c) Suppose C and D are chambers of Δ with $\delta(C, D) = w \in W_M$. Given any reduced decomposition \mathbf{i} of w , show that there is a minimal gallery from C to D of type \mathbf{i} .

(d) For any apartment Σ and chamber $C \in \Sigma$, show that the retraction $\rho = \rho_{\Sigma, C}$ satisfies $\delta(C, \rho(D)) = \delta(C, D)$ for any chamber $D \in \Delta$.

(e) Suppose Δ is the complex of flags of proper non-zero subspaces of an n -dimensional vector space, as in Exercise 2 at the end of §2 above. Show that W_M can be identified with the symmetric group on n letters and that δ associates to any pair C, C' of chambers the Jordan–Hölder permutation $\pi(C, C')$.

4. Let Δ and δ be as in Exercise 3. By “apartment” in what follows, we will always mean an apartment in the complete apartment system. Given two subsets $\mathcal{C}, \mathcal{D} \subseteq \text{Ch } \Delta$, a *strong isometry* from \mathcal{C} into \mathcal{D} is a function $\alpha : \mathcal{C} \rightarrow \mathcal{D}$ such that $\delta(\alpha(C), \alpha(C')) = \delta(C, C')$ for all $C, C' \in \mathcal{C}$. We will also say, in this situation, that \mathcal{C} is *strongly isometric* to its image $\alpha(\mathcal{C})$. The purpose of this exercise is to prove:

Theorem. *If \mathcal{C} is strongly isometric to a subset of an apartment, then \mathcal{C} is contained in an apartment.*

It will be convenient to introduce some canonical chamber maps ρ which are slight variants of the retractions onto apartments. Given an apartment Σ , a chamber $C \in \Sigma$, and a chamber $D \in \Delta$, there is a unique type-preserving chamber map $\rho : \Delta \rightarrow \Sigma$ such that $\rho(D) = C$ and ρ maps every apartment containing D isomorphically onto Σ . [Uniqueness is clear; for existence, take the retraction $\rho_{\Sigma', D}$, where Σ' is any apartment containing D , and follow it by the unique type-preserving isomorphism $\Sigma' \rightarrow \Sigma$ taking D to C .] For lack of a better name, we will simply call ρ the *canonical map* $\Delta \rightarrow \Sigma$ such that $\rho(D) = C$.

(a) Let \mathcal{C} be as in the theorem and let Σ be an apartment. Show that a strong isometry $\alpha : \mathcal{C} \rightarrow \text{Ch } \Sigma$ is completely determined once $\alpha(C_0)$ is known for one chamber $C_0 \in \mathcal{C}$. Moreover, α is necessarily the restriction to \mathcal{C} of the canonical map $\rho : \Delta \rightarrow \Sigma$ taking C_0 to $\alpha(C_0)$.

(b) Let $\mathcal{D} = \alpha(\mathcal{C}) \subseteq \text{Ch } \Sigma$, and let $\beta = \alpha^{-1} : \mathcal{D} \rightarrow \mathcal{C}$. If D and D' are adjacent chambers of Σ with $D \in \mathcal{D}$ and $D' \notin \mathcal{D}$, show that β extends to a strong isometry from $\mathcal{D} \cup \{D'\}$ into Δ .

(c) Show by repeated applications of (b) that β can be extended to a strong isometry defined on all of $\text{Ch } \Sigma$.

(d) Deduce the theorem from (c).

5. Go back through the proof in Exercise 4, and get the following additional information: Suppose the set \mathcal{C} in the theorem is contained in a convex chamber subcomplex $\Delta' \subseteq \Delta$ in which every codimension 1 simplex is a face of at least two chambers; then the apartment containing \mathcal{C} can be constructed inside Δ' . Deduce that Δ' is an apartment if it is thin. This proves the characterization (2) of the complete apartment system.

HINTS

1. If Γ is minimal, then it is contained in an apartment; \mathbf{i} is then reduced by the connection between words in a Coxeter group and galleries in the associated Coxeter complex. Conversely, suppose \mathbf{i} is reduced. We may assume by

induction that the subgallery C_1, \dots, C_d is minimal and hence is contained in an apartment Σ . Then $\rho_{\Sigma, C_1}(\Gamma)$ is a non-stuttering gallery in Σ with the same type i as Γ , so it is minimal. It follows that Γ is minimal, since its image under a chamber map is minimal.

2. Suppose Σ is a chamber subcomplex which is isomorphic to Σ_M in a label-preserving way. It suffices to show that if Σ is adjoined to an apartment system \mathcal{A} , then axiom (B2'') still holds. The proof is essentially the same as the proof that a union of apartment systems is again an apartment system. The given complex Σ plays the role of the complex Σ' that occurred in that proof, and the only extra ingredient required is that one needs to use Exercise 1 to show that every minimal gallery in Σ is still minimal in Δ .

3. (a) We may assume that $\Sigma = \Sigma(W, S)$ with its canonical labelling; we can then take δ to be the difference function $(\{w\}, \{w'\}) \mapsto w^{-1}w'$ discussed above.

(c) Work in an apartment containing C and D .

(e) Work in the apartment Σ corresponding to a frame $\{L_1, \dots, L_n\}$. The symmetric group $W = W_M$ acts on Σ in an obvious way, and one need only check that $\pi(C, wC) = w$ for any $w \in W$ if C is the standard flag $L_1 \subset L_1 \oplus L_2 \subset \dots \subset L_1 \oplus \dots \oplus L_{n-1}$.

4. (b) Let $C = \beta(D)$. It suffices to find a chamber C' such that the canonical map $\rho : \Delta \rightarrow \Sigma$ with $\rho(C') = D'$ induces $\alpha : \mathcal{C} \rightarrow \mathcal{D}$; for then ρ will induce a strong isometry $\mathcal{C} \cup \{C'\} \rightarrow \mathcal{D} \cup \{D'\}$ whose inverse extends β (cf. Exercise 3(d) above). Suppose, for the moment, that C' is *any* chamber distinct from C and i -adjacent to it, where i is the label such that D and D' are i -adjacent. Let ρ be as above. Given $D'' \in \mathcal{D}$, let $C'' = \beta(D'')$ and let $w = \delta(D, D'') = \delta(C, C'')$. Thus $D'' = wD$ if we identify Σ with Σ_M in such a way that D is the fundamental chamber. Similarly, $D' = sD$, where s is the generator s_i of W_M . Let Φ and Φ' be the half-spaces of Σ corresponding to the reflection s , with $D \in \Phi$ and $D' \in \Phi'$. There are two cases to consider.

Case 1: $D'' \in \Phi$, i.e., $l(sw) = l(w) + 1$. Take any minimal gallery Γ from C to C'' . Its type i is a reduced decomposition of w . The gallery (C', Γ) consisting of C' followed by Γ is then non-stuttering and of reduced type (i, i) , hence it is minimal. It follows that ρ takes this gallery to a non-stuttering gallery in Σ of the same type, whence $\rho(C) = D$ and $\delta(D, \rho(C'')) = w$. This implies that $\rho(C'') = D''$, as required.

Case 2: $D'' \in \Phi'$, i.e., $l(sw) = l(w) - 1$. Then there is a reduced decomposition i of w starting with i . By Exercise 3(c), we can find a minimal gallery Γ from C to C'' of type i . If the second chamber C_1 of Γ happens to be our chosen C' , then $\delta(D', \rho(C'')) = sw$ and hence $\rho(C'') = D''$. Otherwise there is a non-stuttering gallery C', C_1, \dots, C'' of type i , obtained by replacing C by C' in Γ . This gallery is still minimal, so $\delta(D', \rho(C'')) = w$ and $\rho(C'') = swD = sD''$.

Let $f = \rho\beta : \mathcal{D} \rightarrow \text{Ch } \Sigma$. We have just seen that $f(D'') = D''$ or sD'' for all $D'' \in \mathcal{D}$, and that the possibility $f(D'') = sD''$ can only occur for $D'' \in \Phi'$. Moreover, it is clear from the discussion of Case 2 that we can choose C' so that $f(D'') = D''$ for at least one $D'' \in \mathcal{D} \cap \Phi'$ [unless $\mathcal{D} \cap \Phi' = \emptyset$, in which case we're already done]. Note further that f is distance-decreasing, in the sense that $d(f(D_1), f(D_2)) \leq d(D_1, D_2)$ for all $D_1, D_2 \in \mathcal{D}$; for β preserves distances,

and ρ , like all chamber maps, takes galleries to galleries and hence is distance-decreasing. The following simple observation, which is an immediate consequence of Exercise 1 in §III.4A, now shows that $f(D'') = D''$ for all D'' :

Let \mathcal{E} be any subset of $\text{Ch } \Phi'$ and let $f : \mathcal{E} \rightarrow \text{Ch } \Sigma$ be a distance-decreasing function such that $f(E) = E$ or sE for all $E \in \mathcal{E}$. If $f(E) = E$ for one $E \in \mathcal{E}$, then $f(E) = E$ for all $E \in \mathcal{E}$.

(d) The extended isometry, still called β , preserves the adjacency relations and hence extends further to a label-preserving chamber map $\beta : \Sigma \rightarrow \Delta$. This chamber map is an isomorphism onto its image, with inverse given by a canonical map ρ as in (a). The image is then an apartment containing \mathcal{C} by Exercise 2.

5 The Spherical Case

A Coxeter complex Σ is called *spherical* if it is isomorphic to the complex associated to a finite reflection group. In view of a result in the optional §II.5, this is equivalent to saying that Σ is finite; but we will not make any use of this equivalence, except in another optional section (§6 below). It is not hard to see that Σ is finite if and only if it has finite diameter, so we can also characterize the spherical Coxeter complexes as those of finite diameter.

A building Δ is called *spherical* if its apartments are spherical. In this case Δ has finite diameter, equal to the diameter of any apartment. [And, conversely, if Δ has finite diameter, then it is spherical by the last sentence of the previous paragraph.] Two chambers C, C' in a spherical building Δ are said to be *opposite* if $d(C, C') = \text{diam } \Delta$. This terminology is motivated by the corollary to Proposition 4 in §I.4E. As an easy consequence of that corollary, we will prove:

Lemma. *Let C and C' be opposite chambers and let Σ be any apartment containing C and C' . Then every chamber $D \in \Sigma$ occurs in some minimal gallery from C to C' .*

[This is the combinatorial analogue of the following geometric fact: Given two opposite points x, x' of a sphere, the geodesics (great semi-circles) from x to x' cover the entire sphere.]

PROOF: Let $d = d(C, C') = \text{diam } \Sigma$. The result of Chapter I cited above says that d is also equal to the number of walls of Σ , all of which separate C from C' . For any chamber $D \in \Sigma$ and any wall H , D is either in the half-space containing C or the half-space containing C' . Hence H separates D from either C or C' , but not both. Since the combinatorial distance between two chambers equals the number of separating walls, it follows that

$$d = d(C, C') = d(C, D) + d(D, C').$$

Hence we can construct a minimal gallery from C to C' by juxtaposing a minimal gallery from C to D with one from D to C' . \square

Since we know that any apartment Σ is convex, it follows from the lemma that Σ is the convex hull of $\{C, C'\}$ for any pair of opposite chambers $C, C' \in \Sigma$, i.e., Σ is the smallest convex chamber subcomplex of Δ containing C and C' . This simple observation leads to the following theorem, which shows that the nature of apartment systems in a spherical building is considerably simpler than in the general case:

Theorem 1. *A spherical building Δ admits a unique system of apartments. The apartments are precisely the convex hulls of pairs C, C' of opposite chambers.*

PROOF: Let \mathcal{A} be an arbitrary system of apartments. Every apartment in \mathcal{A} contains a pair of opposite chambers and hence is their convex hull. Conversely, given a pair C, C' of opposite chambers in Δ , there is an apartment $\Sigma \in \mathcal{A}$ containing them both, and Σ is then equal to their convex hull; hence this convex hull is indeed in \mathcal{A} . \square

Remark. In non-spherical buildings there can definitely exist apartment systems other than the complete one. We will see in the next chapter that such apartment systems arise naturally from group theory. But there is an easy example available now, namely, the case where Δ is a tree. The complete apartment system \mathcal{A} in this case consists of all lines in Δ , but it is easy to see that one usually does not need to take all of the lines as apartments in order to satisfy the building axioms. If you know about ends of trees (cf. [46], §§I.2.2 and II.1.1), then you can understand the situation as follows: The set \mathcal{A} of lines is in 1-1 correspondence with the set of pairs of distinct ends of Δ ; in particular, \mathcal{A} has a natural topology. To get a system of apartments, one need only take a dense subset of \mathcal{A} .

The notion of “opposite chamber” that we introduced above has other uses apart from the characterization of apartments. It arises, for instance, if one attempts to analyze the homotopy type of a spherical building Δ (i.e., the homotopy type of the geometric realization $|\Delta|$). Here’s a sketch of how that can be done, following [47].

Fix a chamber C of Δ and let Δ' be the subcomplex obtained by deleting all chambers opposite to C . I claim that Δ' is contractible. Now a contractible subcomplex can be collapsed to a point without affecting the homotopy type (cf. [48], 7.1.5 and 7.6.2). So the claim yields the following theorem of Solomon and Tits [47]:

Theorem 2. *If Δ is a spherical building of rank n , then $|\Delta|$ has the homotopy type of a bouquet of $(n - 1)$ -spheres, where there is one sphere for every apartment containing C .* \square

It remains to say something about the claim. Note first that every apartment Σ containing C admits a canonical label-preserving isomorphism to Σ_M , with C going to the fundamental chamber of Σ_M . Now $|\Sigma_M|$ can be identified with the unit sphere in the vector space V on which the reflection

group W_M acts. Hence Σ admits a “spherical geometry”; in particular, the punctured sphere $|\Sigma \cap \Delta'|$ admits a canonical contracting homotopy which contracts it to the barycenter of $|C|$ along geodesics, i.e., arcs of great circles. One can show that the various homotopies, one for each apartment containing C , are compatible with one another; hence they fit together to give a well-defined contracting homotopy on Δ' .

Remarks

1. The idea of introducing geodesics (and other geometric notions) into the study of buildings is extremely useful. We will not do any more with it in connection with spherical buildings, but in Chapter VI we will carry out in detail the analogous program for Euclidean buildings, i.e., buildings whose apartments can be identified with Euclidean space.

2. The next section contains a purely combinatorial proof, also based on [47], of the contractibility of Δ' . That proof, while not as close to the geometric intuition as the proof sketched above, has the advantage that the method extends to non-spherical buildings and enables one to analyze their homotopy type as well. The result is that every non-spherical building is contractible.

3. If you know what a Cohen–Macaulay complex is (cf. [40], §8), then you can easily deduce from our study of the homotopy type of a building that every building is a Cohen–Macaulay complex. [The point here is that links in buildings are again buildings, so we also understand the homotopy type of any link.]

EXERCISE

Let Δ be the complex of flags of proper non-zero subspaces of an n -dimensional vector space, as in Exercise 2 of §2 and Exercise 3(e) of §4. Let C and C' be chambers with vertex sets $V_1 \subset \cdots \subset V_{n-1}$ and $V'_1 \subset \cdots \subset V'_{n-1}$, respectively. Show that C and C' are opposite if and only if $V = V_i \oplus V'_{n-i}$ for all i . [HINT: Begin by noting that the length of a permutation π (with respect to the standard generating set S for the symmetric group on n letters) is the number of pairs $i < j$ such that $\pi(i) > \pi(j)$. You can check this directly or you can deduce it from §I.4E, Proposition 4, part (1).]

6 The Homotopy Type of a Building

This section is optional. Its purpose, as stated in Remark 2 above, is to give Solomon’s combinatorial proof of the Solomon–Tits theorem and of the contractibility of non-spherical buildings. The idea of this proof is to start with a fixed chamber C and then keep track of the homotopy type as you successively adjoin the chambers adjacent to C , then the chambers at distance 2 from C , etc. The proposition below enables one to figure out

what happens each time a new chamber is adjoined (along with its faces). Recall that if A is a simplex in an abstract simplicial complex Δ , then \bar{A} denotes the subcomplex $\Delta_{\leq A}$, whose geometric realization is the closed simplex in $|\Delta|$ corresponding to A .

Proposition. *Let Δ be an arbitrary building. Fix a chamber C and an integer $d \geq 1$, and let \mathcal{D} be a set of chambers with the following two properties:*

- (1) $d(C, D) \leq d$ for every $D \in \mathcal{D}$.
- (2) \mathcal{D} contains every chamber $D \in \Delta$ with $d(C, D) < d$.

Let Δ' be the subcomplex of Δ generated by \mathcal{D} , and let D be a chamber of Δ such that $d(C, D) = d$ and $D \notin \mathcal{D}$. Then

$$\bar{D} \cap \Delta' = \bigcup_{A \in \mathcal{F}} \bar{A},$$

where \mathcal{F} is a non-empty set of codimension 1 faces of D . The set \mathcal{F} contains all the codimension 1 faces of D if and only if Δ is spherical and of diameter d .

It is a routine matter to use this proposition to complete the analysis of the homotopy type of Δ . The details of this are left as an exercise.

The rest of this section will be devoted to a proof of the proposition.

Lemma 1. *Given simplices $C, A \in \Delta$ with C a chamber, there is a unique chamber $D \geq A$ such that $d(C, D) = d(C, A)$.*

PROOF: Let Σ be an apartment containing C and A . Then, as we noted in §4, Σ contains any gallery Γ stretched from C to A . So Σ contains every chamber $D \geq A$ with $d(C, D) = d(C, A)$. We are therefore reduced to proving the lemma for the Coxeter complex Σ . We may assume that $\Sigma = \Sigma(W, S)$ for some Coxeter system (W, S) and that C is the fundamental chamber. Then A corresponds to a special coset in W , and chambers having A as a face correspond to representatives of this coset. The lemma now follows from the first assertion of Exercise 2 in §II.3D, which says that a special coset has a unique representative of minimal length. Alternatively, you might find it instructive to give a geometric proof of the lemma, using foldings; see Proposition 2.29 of Tits [56] if you get stuck. [It is also instructive to think about why the lemma is true when Σ is the poset of cells associated to a set of hyperplanes, as in Chapter I.] \square

Lemma 2. *Let C, A , and D be as in Lemma 1. For any chamber $D' \geq A$, one can find a minimal gallery from C to D' which consists of a gallery from C to D followed by a gallery which is contained in $\Delta_{\geq A}$.*

PROOF: By working in an apartment containing C and D' , we are again reduced to the case of a Coxeter complex. The result this time follows from the second assertion of the same exercise cited in the proof of Lemma 1. Once again there is a geometric proof, for which you can see Propositions 2.30.6 and 2.7 of [56]. \square

Lemma 3. *If Δ is spherical and C and D are opposite chambers, then $d(C, A) < d(C, D)$ for every codimension 1 face A of D . Conversely, suppose Δ is arbitrary and C and D are chambers such that $d(C, A) < d(C, D)$ for every codimension 1 face A of D ; then Δ is spherical and C and D are opposite.*

PROOF: As usual, it suffices to consider $\Sigma(W, S)$ and to assume that C is the fundamental chamber. The first assertion, then, simply states the obvious fact that if w is the unique element of maximal length in a finite reflection group W , then $l(ws) < l(w)$ for all $s \in S$. This is equally obvious geometrically: If D is the unique chamber at maximal distance from C , then any chamber D' adjacent to D and distinct from it is closer to C than D is.

For the converse, we are given a Coxeter system (W, S) with S finite and an element $w \in W$ such that $l(ws) < l(w)$ for all $s \in S$. We must show that W is finite and that w is the element of maximal length. It suffices to show that $l(w) = l(ww'^{-1}) + l(w')$ for all $w' \in W$. [This implies that w has maximal length, and the finiteness of W then follows from the fact that the length function is bounded.] The proof is by induction on $l(w')$, which may be assumed positive.

Take a reduced decomposition $w' = s_1 \cdots s_d$. Then the induction hypothesis applies to the element $w'' = s_1 \cdots s_{d-1}$ and shows that w admits a reduced decomposition ending with $s_1 \cdots s_{d-1}$. Since $l(ws_d) < l(w)$ by hypothesis, the exchange condition implies that we may exchange one of the letters in our reduced decomposition of w for an s_d at the end. The exchanged letter cannot come from the “ w'' -part” of w , since that would contradict the fact that $l(w') = d$. So we obtain a new reduced decomposition of w ending with $s_1 \cdots s_d$, as required. \square

Lemma 4. *With the hypotheses and notation of the proposition,*

$$\overline{D} \cap \Delta' = \{ B < D : d(C, B) < d \}.$$

PROOF: The right-hand side is contained in the left-hand side by hypothesis (2) of the proposition. To prove the opposite inclusion, suppose $B \in \overline{D} \cap \Delta'$. Then there is a chamber $D' \in \mathcal{D}$ with $B < D'$, and we have

$$d(C, B) \leq d(C, D') \leq d(C, D) = d$$

by hypothesis (1). Since $D \neq D'$, Lemma 1 implies that $d(C, B) < d$, as required. \square

PROOF OF THE PROPOSITION: Given $B < D$ with $d(C, B) < d$, it follows from Lemma 2 that there is a minimal gallery $\Gamma : C = C_0, \dots, C_d = D$ with $B < C_{d-1}$. Setting $A = C_{d-1} \cap D$, we then have $B \leq A$ and $d(C, A) < d$. Thus

$$\{ B < D : d(C, B) < d \} = \bigcup_{A \in \mathcal{F}} \bar{A},$$

where \mathcal{F} is the set of codimension 1 faces A of D such that $d(C, A) < d$. Note that the left-hand side of this equation is non-empty [it contains the empty face of D]; hence \mathcal{F} is non-empty. The proposition now follows from Lemmas 4 and 3. \square

7 The Axioms for a Thick Building

The purpose of this final section of the chapter is to show that axiom **(B0)** can essentially be eliminated if Δ is thick. We won't actually need this result in what follows, since it will always be clear in our examples that the purported apartments are in fact Coxeter complexes. But the proof is very instructive, being based on a clever use of retractions and the standard uniqueness argument.

Theorem. *Let Δ be a thick chamber complex with a family \mathcal{A} of thin chamber subcomplexes Σ satisfying axioms **(B1)** and **(B2)**. Then every $\Sigma \in \mathcal{A}$ is a Coxeter complex, so Δ is a building and \mathcal{A} is a system of apartments.*

PROOF: Note first that much of the theory of retractions developed in §3 did not use axiom **(B0)**, but only the fact that the apartments Σ are thin chamber complexes. In particular, Propositions 1 and 2 of that section remain valid, as does Corollary 1 of Proposition 1. We will also need to know that the retraction $\rho = \rho_{\Sigma, C}$ preserves distances from any face A of C :

$$d(A, \rho(D)) = d(A, D) \quad (*)$$

for any chamber $D \in \Delta$ and any $A \leq C$. This was proved in §3 for the case $A = C$, and the proof in general is identical.

We now show that any $\Sigma \in \mathcal{A}$ is a Coxeter complex by constructing foldings (cf. §III.4B). Given distinct adjacent chambers $C, C' \in \Sigma$, we must find a folding $\phi : \Sigma \rightarrow \Sigma$ such that $\phi(C') = C$. Let A be the common face $C \cap C'$, let C'' be a third chamber of Δ having A as a face, and let Σ' be an apartment containing C and C'' . Let $\phi : \Sigma \rightarrow \Sigma$ be the restriction to Σ of $\rho_{\Sigma, C'} \circ \rho_{\Sigma', C}$. Then ϕ fixes C pointwise and satisfies $\phi(C') = C$. We will prove that ϕ is a folding. [You should draw a picture, of the tree case for instance, to see why this is plausible.]

In view of (*), ϕ preserves distances from A , i.e., $d(A, \phi(D)) = d(A, D)$ for any chamber $D \in \Sigma$. [Distances here may be computed either in Σ or in Δ , but we will be thinking about Σ -distances when we apply this.] In other words, if Γ is a gallery in Σ stretched from A to D , then $\phi(\Gamma)$ is stretched from A to $\phi(D)$. In particular, $\phi(\Gamma)$ is non-stuttering, and this will enable us to apply the standard uniqueness argument.

A first such application shows that if D is a chamber of Σ with $d(A, D) = d(C, D)$ (i.e., if there is a gallery stretched from A to D which starts with C),

then ϕ fixes D pointwise. Thus ϕ is the identity on the subcomplex Φ generated by $\{ D \in \text{Ch } \Sigma : d(A, D) = d(C, D) \}$. And this subcomplex Φ is precisely the image of ϕ . For suppose D is any chamber of Σ and Γ is a gallery stretched from A to D ; then $\phi(\Gamma)$ is stretched from A to $\phi(D)$ and starts with C , so $\phi(D) \in \Phi$. Thus ϕ is a retraction of Σ onto Φ .

Everything we have done so far can also be done with the roles of C and C' reversed. Hence there is an endomorphism ϕ' of Σ with $\phi'(C) = C'$, such that ϕ' preserves distances from A and retracts Σ onto the subcomplex Φ' generated by $\{ D \in \text{Ch } \Sigma : d(A, D) = d(C', D) \}$.

We show next that Φ and Φ' have no chamber in common: Suppose D is a chamber in $\Phi \cap \Phi'$. Then D is fixed pointwise by both ϕ and ϕ' . If Γ is a gallery stretched from D to A , it follows by the standard uniqueness argument that ϕ and ϕ' fix every chamber in Γ pointwise. But this is absurd, since Γ ends with either C or C' .

We now have $\text{Ch } \Sigma = \text{Ch } \Phi \amalg \text{Ch } \Phi'$. The proof that ϕ is a folding will be complete if we can show that ϕ maps $\text{Ch } \Phi'$ bijectively to $\text{Ch } \Phi$. To this end, consider the composites $\phi\phi'$ and $\phi'\phi$. The first takes C to C and fixes A pointwise, so it fixes C pointwise; it is therefore the identity on Φ by the standard uniqueness argument. Similarly, the other composite is the identity on Φ' . Hence ϕ induces an isomorphism $\Phi' \rightarrow \Phi$ with inverse induced by ϕ' . \square

V

Buildings and Groups

In this chapter we will develop the group theory that goes along with the theory of buildings, in much the same way that the theory of Coxeter groups goes along with the theory of Coxeter complexes. In particular, we will discover a class of groups G for which we can construct an associated building Δ , on which G acts as a group of type-preserving simplicial automorphisms.

We will only consider thick buildings. It would be possible to treat the general case, but the thickness assumption leads to some simplifications and suffices for most applications to group theory.

We begin by assuming that we have a group G which acts in a nice way on a thick building Δ . This imposes some conditions on G , and we will then take these conditions as axioms for the class of groups we are seeking.

1 Strongly Transitive Automorphism Groups

1A Definitions

Let Δ be a thick building and \mathcal{A} a system of apartments. Suppose a group G acts on Δ as a group of type-preserving simplicial automorphisms leaving \mathcal{A} invariant. Thus if Σ is an apartment, then so is its image $g\Sigma$. This is automatic, of course, if \mathcal{A} is the complete system of apartments; but we want to allow \mathcal{A} to be arbitrary.

We will say that the G -action is *strongly transitive* if G acts transitively on the set of pairs (Σ, C) consisting of an apartment $\Sigma \in \mathcal{A}$ and a chamber $C \in \Sigma$. This is equivalent to saying that G is transitive on $\text{Ch } \Delta$ and that the stabilizer of a given chamber C is transitive on the set of apartments containing C . Alternatively, it is equivalent to saying that G is transitive on \mathcal{A} and that the stabilizer of a given apartment Σ is transitive on $\text{Ch } \Sigma$.

Assume throughout the remainder of this section that the G -action is strongly transitive, and choose an arbitrary pair (Σ, C) as above. We will often refer to Σ as the “fundamental apartment” of Δ and to C as the “fundamental chamber” of Σ . Let W be the group of type-preserving automorphisms of Σ , and let $S \subset W$ be the set of reflections associated to the codimension 1 faces of C . Then (W, S) is a Coxeter system, and Σ can be identified with $\Sigma(W, S)$.


We now introduce three subgroups of G :

$$B = \{ g \in G : gC = C \}$$

$$N = \{ g \in G : g\Sigma = \Sigma \}$$

$$T = \{ g \in G : g \text{ fixes } \Sigma \text{ pointwise} \}.$$


Note that T is a normal subgroup of N , being the kernel of the homomorphism $\alpha : N \rightarrow W$ induced by the action of N on Σ . Note also that α is surjective, so that $W \approx N/T$. For if we are given $w \in W$, then we can find $n \in N$ such that $nC = wC$; since n and w are both type-preserving, they agree pointwise on C , hence $\alpha(n) = w$ by the standard uniqueness argument. Note, finally, that $T = B \cap N$; for if $n \in B \cap N$, then n fixes C pointwise and hence acts trivially on Σ . The following diagram summarizes the notation:



Let's pause now to see what all of this looks like in a simple example.

1B Example

Let P be the projective plane over a field k , as defined in §IV.2, and let Δ be its flag complex. It is a rank 2 building, with one vertex for every proper non-zero subspace of k^3 and one edge for each pair consisting of a 1-dimensional subspace contained in a 2-dimensional subspace. The unique apartment system for Δ has one apartment for every triple $\{L_1, L_2, L_3\}$ of 1-dimensional subspaces such that $k^3 = L_1 \oplus L_2 \oplus L_3$. Given any subset $X \subseteq k^3$, let's denote by $[X]$ the subspace spanned by X . Then a triple as above has the form $\{[e_1], [e_2], [e_3]\}$ for some basis e_1, e_2, e_3 of k^3 , and the corresponding apartment is the following subcomplex of Δ :



As fundamental apartment Σ we take the apartment associated to the standard basis of k^3 . And as fundamental chamber C we take the edge $[e_1] \text{ --- } [e_1, e_2]$.

Let G be the group $\text{GL}_3(k)$ of linear automorphisms of k^3 . Then any $g \in G$ takes subspaces to subspaces and induces a type-preserving automorphism of Δ . It is easy to check that this action of G is strongly transitive. Let's compute B , N , T , W , and S .

The subgroup $B \subset G$ consists of all automorphisms of k^3 which leave the subspaces $[e_1]$ and $[e_1, e_2]$ invariant; hence B is the *upper-triangular group*

$$\begin{pmatrix} * & * & * \\ 0 & * & * \\ 0 & 0 & * \end{pmatrix}$$

The subgroup $N \subset G$ consists of all automorphisms which permute the three subspaces $[e_1], [e_2], [e_3]$. Hence N is the *monomial group*, consisting of all matrices with exactly one non-zero element in every row and every column.

Given $n \in N$, the action of n on Σ is determined by the permutation of $\{[e_1], [e_2], [e_3]\}$ induced by n ; so T consists of the diagonal matrices (which induce the trivial permutation), and $W = N/T$ can be identified with the symmetric group on 3 letters, or, equivalently with the group of 3×3 permutation matrices. [Thus we have a splitting $N = T \rtimes W$.]

Finally, it is easy to check that the set S of “fundamental reflections” consists of the permutations represented by

$$\begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}.$$

As a mnemonic aid, you might find it helpful to know that B is what is called a Borel subgroup of G in the theory of matrix groups, T is a maximal torus, N is the normalizer of T , and W is the Weyl group.

Remarks

1. Instead of taking $G = \text{GL}_3(k)$ above, we could equally well have taken G to be the subgroup $\text{SL}_3(k)$ consisting of matrices of determinant 1. The groups B , N , and T would then be the intersections with $\text{SL}_3(k)$ of the groups B , N , and T above. The quotient $W = N/T$ would still be the symmetric group on 3 letters (as it has to be, since $W = \text{Aut}_0(\Sigma)$, independent of G). The set $S \subset W$ consists of the same two permutations as above, which can be represented by the monomial matrices

$$\begin{pmatrix} 0 & -1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & -1 \\ 0 & 1 & 0 \end{pmatrix}$$

of determinant 1.

2. Still another variation on this example is obtained by replacing $\text{GL}_3(k)$ by its quotient $\text{PGL}_3(k) = \text{GL}_3(k)/Z$, where $Z \subset \text{GL}_3(k)$ is the central subgroup consisting of scalar multiples of the identity matrix. The subgroup Z acts trivially on Δ , so we obtain an action of $\text{PGL}_3(k)$ on Δ ; and clearly this action is still strongly transitive. Similarly, $\text{PSL}_3(k) = \text{SL}_3(k)/(\text{SL}_3(k) \cap Z)$ acts strongly transitively on Δ .

3. If you have done Exercise 2 at the end of §IV.2, then you can easily generalize all this from GL_3 to GL_n (or SL_n , or PGL_n , or PSL_n).

1C The chamber system associated to Δ

Let G and Δ be as in §1A. Recall the notation that we introduced there:

$$\begin{array}{ccc} & G & \\ B & \swarrow \searrow & N \\ & T & \end{array} \rightarrow W = \langle S \rangle$$

We will identify W with N/T ; in particular, we will often find it convenient to view an element $w \in W$ as a subset of G , namely, the coset $\tilde{w}T$, where $\tilde{w} \in N$ is any element of N representing w .


Let's try now to describe Δ in terms of G and see, in the process, how properties of Δ translate into properties of G .

Give Δ the labelling λ whose restriction to Σ is the canonical labelling, with S as the set of labels. [Here Σ is identified with $\Sigma(W, S)$.] Then Δ gives rise to a chamber system, consisting of the set $\text{Ch } \Delta$ together with s -adjacency relations, one for each $s \in S$. Since G acts transitively on $\text{Ch } \Delta$ with B as the stabilizer of C , we have a bijection

$$G/B \xrightarrow{\sim} \text{Ch } \Delta,$$

which takes a coset gB to the chamber gC . Given $s \in S$, we need to figure out the s -adjacency relation induced on the set G/B of cosets.

Suppose first that $h \in G$ is an element such that hC is s -adjacent to C . Then $C \cap hC$ is the face $A = C \cap sC$ of C of type $S - \{s\}$:



Since h is type-preserving, it must take A to the face of hC of the same type. But A itself is the face of hC of this type, so $hA = A$. Thus hC is s -adjacent to C if and only if h is in the stabilizer P_s of the face A of C of type $S - \{s\}$. Applying the G -action, we conclude that gC is s -adjacent to $g'C$ if and only if $g' = gh$ for some $h \in P_s$. In other words, s -adjacency of chambers corresponds to the following relation on G/B :

$$gB \sim g'B \iff gP_s = g'P_s.$$

Note the analogy here with the Coxeter group situation, where chambers correspond to cosets of the trivial subgroup, and two are s -adjacent if and only if they represent the same coset of $\langle s \rangle$.

Next, we will make this analogy even better by showing that P_s is the subgroup $\langle B, s \rangle$ generated by B and s . [Recall that s is viewed as a coset $\tilde{s}T \subset N$, so the symbol $\langle B, s \rangle$ makes sense. Moreover, $\langle B, s \rangle = \langle B, \tilde{s} \rangle$ for any representative \tilde{s} .] Given $h \in P_s$, choose an apartment Σ' containing C and hC . By strong transitivity we can find $b \in B$ such that $b\Sigma' = \Sigma$. The chamber bhC of Σ is then s -adjacent to $bC = C$, so it is either C

or sC . It follows that bh is either in B or in $sB [= \tilde{s}B]$. Thus we have $h \in B \cup b^{-1}sB$. This proves that $P_s = \langle B, s \rangle$, and, in fact, it proves the following much more precise result:

$$P_s = B \cup BsB.$$

Digression. The set $BsB [= B\tilde{s}B]$ which just arose is a *double coset*. Most people don't learn about double cosets until they need to, so let's take a moment to review them. Given a group G and a subgroup B , you are certainly familiar with the partition of G into left cosets gB , which are minimal subsets invariant under right multiplication by B . Similarly, there are right cosets Bg , which are closed under left multiplication by B and which give a different partition of G . Double cosets BgB provide a third partition of G , this time into subsets which are closed under both left and right multiplication by B .

The three types of cosets all coincide if B is normal in G , in which case the set of cosets inherits a group structure. In general, however, the set G/B of left cosets is just a set with left G -action, the set $B \backslash G$ of right cosets is a set with right G -action, and the set $B \backslash G / B$ of double cosets is a set with no further structure. What we saw above is a situation (namely, $B \subset P_s$) where there are exactly two double cosets.

Example. Let $G = \text{GL}_3(k)$ as in §1B above. The two subgroups P_s are the stabilizers of $[e_1]$ and $[e_1, e_2]$, respectively; hence they are the subgroups

$$\begin{pmatrix} * & * & * \\ 0 & * & * \\ 0 & * & * \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} * & * & * \\ * & * & * \\ 0 & 0 & * \end{pmatrix}.$$

I recommend as an exercise that you verify by direct matrix computation the double coset decomposition $P_s = B \cup BsB$ for each of these groups P_s .

We now have a description, entirely in terms of (G, B, N, S) , of the chamber system associated to Δ . This gives, in principle, a group-theoretic description of Δ . But we will make this much more explicit below by extending the stabilizer calculation to the case of an arbitrary face $A \leq C$.

1D Stabilizers

Given $S' \subseteq S$, let A be the face of C of type $S - S'$. Equivalently, A is the face whose stabilizer in W is the special subgroup $W' = \langle S' \rangle$. Let $P_{S'}$ be the stabilizer of A in G . Given $g \in P_{S'}$, choose an apartment Σ' containing C and gC , and choose (by strong transitivity) an element $b \in B$ such that $b\Sigma' = \Sigma$. Then $bgC = wC$ for some $w \in W$, whence $bg \in wB$ and $g \in BwB$. Since $g \in P_{S'}$ and $B \subseteq P_{S'}$, we have $wA = A$ and hence $w \in W'$. Thus

$$P_{S'} = \bigcup_{w \in W'} BwB,$$

or, more concisely,

$$P_{S'} = BW'B.$$

When S' is a singleton $\{s\}$, we recover the calculation of P_s above.

This formula for $P_{S'}$ raises an obvious question: Are the various double cosets BwB all different as w ranges over W ? We will give an affirmative answer to this question shortly. Accepting this result for the moment, one concludes that the function $S' \mapsto P_{S'}$ is a poset isomorphism from the set of subsets of S to the set of subgroups of G of the form $P_{S'}$. To see this, note that we can recover S' from the group $P = P_{S'}$ by

$$S' = \{s \in S : BsB \subseteq P\}.$$

For if $BsB \subseteq P$, then $BsB = BwB$ for some $w \in \langle S' \rangle$, hence $s = w$ by the result we are accepting; but then $s \in S'$ by the lemma in §III.1.

It now follows formally that Δ , as a poset, is isomorphic to the set of cosets gP , where P is a subgroup of the form $P_{S'}$ and the cosets are ordered by the opposite of the inclusion relation. Thus we have entirely reconstructed Δ from G , B , N , and S .

1E The Bruhat decomposition

The stabilizer calculation above is of interest even when A is the empty face of C . In this case $S' = S$, $P_{S'} = G$, and $W' = W$. So the result is that

$$G = BWB = \bigcup_{w \in W} BwB.$$

In particular, G is generated by B and N . We now prove, as promised, the following more precise result, which is known as the *Bruhat decomposition*:

$$G = \coprod_{w \in W} BwB. \quad (*)$$

To prove (*), we take a closer look at the stabilizer calculation for the case $A = \emptyset$. Given $g \in G$, choose Σ' containing C and gC as before, and choose $b \in B$ such that $b\Sigma' = \Sigma$. Then the action of b induces the unique isomorphism $\Sigma' \rightarrow \Sigma$ fixing C pointwise. So we have

$$bgC = \rho(gC),$$

where ρ is the retraction $\rho_{\Sigma, C}$. What we proved above, then, is that $g \in BwB$, where w is the unique element of W such that $\rho(gC) = wC$. In other words, if we define a function $\bar{\rho} : G \rightarrow W$ by setting $\bar{\rho}(g)$ equal to the element $w \in W$ such that $\rho(gC) = wC$, then we have $g \in B\bar{\rho}(g)B$.

The Bruhat decomposition will follow if we show for any $w \in W$ that $\bar{\rho}$ maps the entire double coset BwB to w . Choose $\tilde{w} \in N$ representing w , and consider an arbitrary element $g = b\tilde{w}b' \in BwB$, where $b, b' \in B$. Then $gC = b\tilde{w}C = bwC$; hence gC and C are both in the apartment $b\Sigma$. Now b^{-1} maps this apartment back to Σ , whence $\rho(gC) = b^{-1}gC = wC$. Thus $\bar{\rho}(g) = w$, as required.

EXERCISES

1. Suppose $G = \text{GL}_n(k)$ and W is the symmetric group on n letters, as in Remark 3 of §1B above. Show that the Bruhat decomposition has the following interpretation in terms of Jordan–Hölder permutations (cf. Exercise 2 of §IV.2): Given $g \in G$, let $w = \pi(C, gC)$, where C is the fundamental chamber of Δ ; then $g \in BwB$. [HINT: *Method 1.* Show that $\pi(C, gC)$ depends only on the double coset containing g ; this reduces you to the case where g is a permutation matrix. *Method 2.* Use the geometric interpretation of the Bruhat decomposition in terms of retractions, and apply Exercise 3(e) of §IV.4.]

2. Let $G = \text{GL}_n(k)$ again. Prove by direct matrix calculations (row and column operations) that $G = BWB$. If you're ambitious, try to prove (*) by matrix calculations also.

1F Products of double cosets

The proof of the Bruhat decomposition leads to a remarkable fact about products of double cosets. Note first that a product $BgB \cdot Bg'B$ of two double cosets in a group is a subset which contains gg' and is closed under left and right multiplication by B ; hence it is a union of double cosets, one of which is $Bgg'B$. This is all that can be said in general. In the present situation, however, we can say much more, at least when one of the factors is BsB for some $s \in S$. Namely, we will show that

$$BwB \cdot BsB \subseteq BwB \cup BwsB \quad (**)$$

for all $w \in W$ and $s \in S$, so that the product consists either of two double cosets or one. [If the product is one double coset, of course, then it is necessarily $BwsB$.]

To prove (**), we need only recall that $\rho : \Delta \rightarrow \Sigma$ is a type-preserving chamber map and hence preserves s -adjacency for all $s \in S$. So the function $\bar{\rho} : G \rightarrow W$ defined above must satisfy $\bar{\rho}(gh) = \bar{\rho}(g)$ or $\bar{\rho}(g)s$ for $g \in G$ and $h \in P_s$. Taking $g \in BwB$ and $h \in BsB \subseteq P_s$, we conclude that $\bar{\rho}(gh) = w$ or ws , i.e., that $gh \in BwB$ or $BwsB$. This proves (**).

Remarks

1. Since $BwB \cdot BsB = BwBsB = B(ws)B$, the content of (**) is that $wBs \subseteq BwB \cup BwsB$.

2. Taking inverses in (**), we can equally well write the result in the form

$$BsB \cdot BwB \subseteq BwB \cup BswB$$

for $s \in S$ and $w \in W$; equivalently, $sBw \subseteq BwB \cup BswB$.

It will be convenient to write $C(w) = BwB$ for $w \in W$. Then, as we noted above, (**) implies that

$$C(w)C(s) = C(ws) \quad \text{or} \quad C(w)C(s) = C(w) \cup C(ws).$$

To finish this discussion, we will show that the first case occurs at least half of the time:

$$C(w)C(s) = C(ws) \text{ if } l(ws) = l(w) + 1. \quad (***)$$

It suffices to show that $wBs \subseteq BwsB$ if $l(ws) = l(w) + 1$. Choose elements $\tilde{w}, \tilde{s} \in N$ representing w and s , respectively. Given $g = \tilde{w}b\tilde{s} \in wBs$, we must show that $\bar{\rho}(g) = ws$, i.e., that $\rho(gC) = wsC$. To compute $\rho(gC)$, we need a minimal gallery from C to $gC = \tilde{w}bsC$. Now $\tilde{w}bsC$ is s -adjacent to $\tilde{w}bC = wC$ and is distinct from it. So if Γ is a minimal gallery in Σ from C to wC , then there is a non-stuttering gallery $\Gamma' = (\Gamma, gC)$, consisting of Γ followed by gC . If we can show that Γ' is minimal, then it will follow that $\rho(gC) = wsC$, as required; for otherwise we would have $\rho(gC) = wC$ and $\rho(\Gamma')$ would stutter, contradicting the fact that ρ preserves distances from C .

To prove the minimality of Γ' , we could simply apply Exercise 1 of §IV.4. But here is a direct proof, which consists, essentially, of one of the steps in the solution to that exercise. Let $\rho' = \rho_{\Sigma, wC}$. Then $\rho'(gC) \neq wC$, hence $\rho'(gC) = wsC$. So $\rho'(\Gamma') = (\Gamma, wsC)$, which is minimal by the correspondence between words in W and galleries starting at C in Σ . But then Γ' must be minimal since its image under the chamber map ρ' is minimal. This completes the proof of (***)

EXERCISE

If $l(ws) = l(w) - 1$, show that $C(w)C(s) = C(w) \cup C(ws)$.

1G Thickness

Finally, we spell out the group-theoretic meaning of our assumption that Δ is thick. (We haven't used this assumption yet, except in the exercise above.)

Since Δ is thick, every codimension 1 face of C is a face of at least three chambers. In other words, for each $s \in S$ there is a chamber C' s -adjacent to C and not equal to C or sC . Such a C' necessarily has the form hC for some $h \in P_s$, and the condition $hC \neq C$ says that $h \notin B$; so we must have $h \in BsB$. Similarly, the condition $hC \neq sC$ says that $h \notin sB$. Thus $BsB \not\subseteq sB$, which implies that $Bs \not\subseteq sB$. This can be rewritten in various ways, such as

$$s^{-1}Bs \not\subseteq B,$$

or, since s has order 2,

$$sBs \not\subseteq B.$$

1H Conclusion

Our rambling discussion has led to a long list of conditions that G must satisfy if it admits a strongly transitive action on a thick building. It's time to put these conditions together and look at them from the point of view of group theory.

2 BN-Pairs

2A The axioms

Suppose we are given a quadruple (G, B, N, S) , where G is a group; B and N are subgroups which generate G ; N normalizes the intersection $T = B \cap N$; and S is a set of generators of the quotient group $W = N/T$. Thus we have the usual setup

$$\begin{array}{ccc} & G & \\ B & & N \\ & T & \end{array} \rightarrow W = \langle S \rangle$$

Let $C(w) = BwB$ for $w \in W$, and let $BW'B = \bigcup_{w \in W'} C(w)$ for any subset W' of W . As the notation suggests, we will be primarily interested in the case where W' is a special subgroup $\langle S' \rangle$ of W . Consider the following conditions on (G, B, N, S) , all of which are known to hold if (G, B, N, S) arises as in §1 from a strongly transitive G -action on a thick building:

- (1) S consists of elements of order 2 and (W, S) is a Coxeter system.
- (2) $B \cup C(s)$ is a subgroup of G for every $s \in S$.
- (3) $BW'B$ is a subgroup of G for every special subgroup $W' \subseteq W$.
- (4) $G = \coprod_{w \in W} C(w)$.
- (5) $C(s)C(w) \subseteq C(w) \cup C(sw)$ for every $s \in S$ and $w \in W$.
- (6) $C(s)C(w) = C(sw)$ if $l(sw) \geq l(w)$.
- (7) $C(s)C(w) = C(w) \cup C(sw)$ if $l(sw) \leq l(w)$.
- (8) For every $s \in S$, $sBs^{-1} \not\subseteq B$.

This long list would make a rather unwieldy system of axioms. Fortunately, it turns out that there is a great deal of redundancy in (1)–(8):

Theorem. *If (5) and (8) hold, then all of the properties (1)–(8) hold.*

The proof will be broken up into several steps:

Lemma 1. *Assume that S consists of elements of order 2. If (5) holds, then all of the properties (2)–(6) hold.*

PROOF: Property (2) is a special case of (3), so we begin with the latter. We must show that $C(w)C(w') \subseteq BW'B$ for every $w, w' \in W'$. Write w as a word $s_1 \cdots s_d$ in the generators S' of W' . Then an easy induction on d shows that $C(w)C(w') \subseteq \bigcup C(w''w')$, where w'' ranges over the 2^d elements of W obtained from the word $s_1 \cdots s_d$ by deleting zero or more letters. [When $d = 1$, this is precisely our hypothesis (5).] Since each w'' is in W' , this proves (3).

We now know, in particular, that BWB is a subgroup of G . Since this subgroup contains B and N , which generate G , we have $G = BWB$. To prove (4), then, it remains to show that

$$C(w) = C(w') \implies w = w'.$$

The proof is by induction on $d = \min\{l(w), l(w')\}$. We may assume $d = l(w')$. If $d = 0$, then $w' = 1$ and the hypothesis $C(w) = C(w')$ says that $C(w) = B$. This clearly implies that $w = 1$ in $W = N/(B \cap N)$, as required. So suppose $d > 0$, and write $w' = sw''$ with $s \in S$ and $l(w'') = d - 1$. Then the hypothesis says that $sw''B \subseteq BwB$. Multiplying by s and using (5), we conclude that

$$w''B \subseteq sBwB \subseteq C(w) \cup C(sw),$$

hence $C(w'') = C(w)$ or $C(sw)$. This implies, by induction, that $w'' = w$ or sw . Now the first case cannot occur, since $l(w'') < d \leq l(w)$. So $w'' = sw$, whence $w' = w$. This proves (4).

Finally, we prove (6) by induction on $l(w)$. The result is trivial if $l(w) = 0$, so suppose $l(w) > 0$ and write $w = w't$ with $t \in S$ and $l(w') = l(w) - 1$. Assume that $C(s)C(w) \neq C(sw)$. Then (5) implies that sBw meets BwB and hence that sBw' meets $BwBt$. Applying (5) again [or, rather, the result obtained from (5) by taking inverses], we conclude that

$$sBw' \text{ meets } C(w) \cup C(wt) = C(w) \cup C(w').$$

If we assume now that $l(sw) \geq l(w)$, then $l(sw') \geq l(w')$; for otherwise we would have

$$l(sw) = l(sw't) \leq l(sw') + 1 < l(w') + 1 = l(w).$$

So the induction hypothesis implies that $C(s)C(w') = C(sw')$, and the result of the previous paragraph now says that $C(sw') = C(w)$ or $C(w')$. In view of (4), it follows that $sw' = w$ or w' . The second possibility is absurd since s has order 2 and hence is non-trivial. But the first possibility is also absurd, since it would imply $l(sw) = l(w') < l(w)$. This contradiction completes the inductive proof of (6) and hence the proof of the lemma. \square

Lemma 2. *Assume that S consists of elements of order 2. If (2), (6), and (8) hold, then so does (7).*

PROOF: By (2) we have $C(s)C(s) \subseteq B \cup C(s)$, so $C(s)C(s) = B$ or $B \cup C(s)$. The first possibility would contradict (8), whence

$$C(s)C(s) = B \cup C(s).$$

This is actually a special case of (7), and this special case, together with (6), easily yields the general case. Indeed, if $l(sw) \leq l(w)$, then $l(s \cdot sw) \geq l(sw)$. So (6) implies

$$C(w) = C(s \cdot sw) = C(s)C(sw),$$

hence

$$\begin{aligned}
C(s)C(w) &= C(s)C(s)C(sw) \\
&= (B \cup C(s))C(sw) \\
&= C(sw) \cup C(s)C(sw) \\
&= C(sw) \cup C(w),
\end{aligned}$$

as required. \square

Lemma 3. *Assume that S consists of elements of order 2. If (4), (6), and (7) hold, then so does (1).*

PROOF: We will verify the folding condition **(F)** of Chapter II. Suppose $l(sw) = l(w) + 1 = l(wt)$ but $l(swt) < l(w) + 2$. Then we have

$$C(s)C(wt) = C(wt) \cup C(swt) = C(wt) \amalg C(swt),$$

where the first equality follows from (7) and the second follows from (4). Since $C(wt) = C(w)C(t)$ by (6), it follows that

$$C(s)C(w)C(t) = C(wt) \amalg C(swt).$$

Similarly, computing $C(sw)C(t)$ by (7) and $C(s)C(w)$ by (6), we find

$$C(s)C(w)C(t) = C(sw) \amalg C(swt).$$

Hence $C(sw) = C(wt)$, and so $sw = wt$ by (4). This proves the folding condition. \square

For practical purposes, Lemmas 1, 2, and 3 prove the theorem. More precisely, they prove the theorem under the extra hypothesis that S consists of elements of order 2. Now in all examples I know of, it is trivial to verify by inspection that S consists of elements of order 2, so it would be harmless to add this as a hypothesis. On the other hand, it is interesting that (5) and (8) actually force this to be true. So here, for your amusement, is the final step in the proof of the theorem:

Lemma 4. *If (5) and (8) hold, then every $s \in S$ has order 2.*

PROOF: Take $w = s^{-1}$ in (5) to get $C(s)C(s^{-1}) \subseteq C(s^{-1}) \cup B$. Since $C(s)C(s^{-1}) \neq B$ by (8), we must have

$$C(s)C(s^{-1}) = C(s^{-1}) \amalg B.$$

Taking inverses, we obtain $C(s)C(s^{-1})$ again on the left, but $C(s) \amalg B$ on the right. This implies that $C(s^{-1}) = C(s)$, and the equality above becomes

$$C(s)C(s) = C(s) \amalg B.$$

On the other hand, if we take $w = s$ in (5) and use the fact that $C(s)C(s)$ is known to consist of two double cosets, then we find

$$C(s)C(s) = C(s) \amalg C(s^2).$$

Hence $C(s^2) = B$ and $C(s) \neq B$, so s has order 2 in W . \square

Remark. It follows from the theorem that we were very inefficient in §1. For instance, there was no need to prove (6) geometrically after having proved (5), nor was there any need for the exercise in §1F. The advantage of having been inefficient, however, is that we have acquired some geometric intuition to go along with each of the properties (1)–(8).

We now formally state our axioms, which have been boiled down to (5) and (8): Following Tits, we say that a pair of subgroups B and N of a group G is a *BN-pair* if B and N generate G , the intersection $T = B \cap N$ is normal in N , and the quotient $W = N/T$ admits a set of generators S such that the following two conditions hold for all $s \in S$ and $w \in W$:

$$(BN1) \quad C(s)C(w) \subseteq C(w) \cup C(sw).$$

$$(BN2) \quad sBs^{-1} \not\subseteq B.$$

One also says, in this situation, that the quadruple (G, B, N, S) is a *Tits system*. Such a quadruple, then, has all of the properties (1)–(8). The group W will be called the *Weyl group* associated to the BN-pair.

It might seem strange that we don't take S as part of the structure in the definition of "BN-pair" (i.e., that we don't define a "BNS-triple"). The reason for omitting S from the notation is that S turns out to be uniquely determined by B and N . We will see this below as a byproduct of our study of the subgroup structure of G . Then, finally, we will be ready to construct a building associated to a group with a BN-pair.

2B Parabolic subgroups

Assume that G is a group with a BN-pair and that S is as in the definition above. Recall that every subset $S' \subseteq S$ gives rise to a subgroup $BW'B$ of G , where $W' = \langle S' \rangle$. As in §1D above, one shows easily that the function $S' \mapsto BW'B$ is a poset isomorphism from the set of subsets of S to the set of subgroups of G of the form $BW'B$. We will call a subgroup of G of this form *special*. This notion of "special subgroup" seems to depend on B , N , and S . Surprisingly, it turns out to depend only on B :

Theorem 1. *The special subgroups of G are precisely the subgroups containing B .*

The crux of the proof is provided by the first assertion of the following lemma:

Lemma. *Let $w \in W$ admit a reduced decomposition $w = s_1 \cdots s_d$. Then the subgroup of G generated by $C(w)$ contains $C(s_i)$ for $i = 1, \dots, d$. Moreover, this subgroup is generated by B and wBw^{-1} .*

PROOF OF THE LEMMA: The subgroup generated by $C(w)$ contains w and B . We therefore have

$$\langle B, wBw^{-1} \rangle \subseteq \langle C(w) \rangle \subseteq \langle C(s_1), \dots, C(s_d) \rangle.$$

So both assertions will follow if we can show that the subgroup $P = \langle B, wBw^{-1} \rangle$ contains $C(s_i)$ for each i . We argue by induction on d . Since $l(s_1w) < l(w)$, we know [cf. (7) above] that s_1Bw meets BwB . Hence s_1B meets $BwBw^{-1}$, which implies that $C(s_1) \subseteq P$. It follows that P also contains $s_1wBw^{-1}s_1$. We can now apply the induction hypothesis to s_1w to conclude that P contains $C(s_i)$ for $i = 2, \dots, d$, whence the lemma. \square

PROOF OF THEOREM 1: The special subgroups obviously contain $B = C(1)$. Conversely, suppose P is a subgroup containing B . Then P is a union of double cosets, hence $P = BW'B$, where W' is the subset of W defined by

$$W' = \{w \in W : C(w) \subseteq P\}.$$

Now W' is a subgroup of W , since $C(w^{-1}) = C(w)^{-1}$ and $C(ww') \subseteq C(w)C(w')$. And the first assertion of the lemma implies that W' contains, for each of its elements w , the generators $s \in S$ which occur in any reduced decomposition of w . Hence W' is the special subgroup of W generated by $W' \cap S$, and P is therefore a special subgroup of G . \square

It is now easy to prove the assertion above that S is uniquely determined by B and N :

EXERCISE

Deduce from Theorem 1 (or directly from the lemma) that S consists of all non-trivial elements $w \in W$ such that $B \cup C(w)$ is a subgroup of G .

By a *special coset* in G we will mean a coset gP such that P is a special subgroup. Equivalently, a left coset in G is special if and only if it contains a left coset of B . Motivated by §1D above, we introduce the poset $\Delta(G, B)$ of special cosets, ordered by the opposite of the inclusion relation. The upshot of §1D, then, is that if our BN-pair arises from a strongly transitive action of G on a building Δ , then we can reconstruct Δ as $\Delta(G, B)$.

Notice that this description of Δ does not refer to N . We need N , however, if we want to describe the apartment system \mathcal{A} in terms of G . Namely, the fundamental apartment $\Sigma \subseteq \Delta$ corresponds to the set of special cosets wP with $w \in W$, or, equivalently, to the set of special cosets with a representative in N ; and the remaining apartments are gotten from Σ by using the G -action.

Returning now to our arbitrary group with a BN-pair, here is a consequence of the second assertion of the lemma:

Theorem 2. *Every special subgroup is equal to its own normalizer, and no two special subgroups are conjugate.*

PROOF: Let P and P' be special subgroups, and suppose $gP'g^{-1} = P$ for some $g \in G$. Let $C(w)$ be the double coset containing g . Then $wP'w^{-1} = P$, so P contains B and wBw^{-1} . The lemma now implies that $C(w) \subseteq P$, so we have $g \in P$ and $P' = P$. This proves the theorem. \square

A subgroup $Q \subseteq G$ is called *parabolic* if Q contains a conjugate of B , or, equivalently, if Q is conjugate to a special subgroup. It follows at once from Theorem 2 that there is a bijection from the set of special cosets to the set of parabolic subgroups, given by $gP \mapsto gPg^{-1}$. This bijection is easily seen to be compatible with the inclusion relation. Consequently, we obtain a new description of the poset $\Delta(G, B)$:

Corollary. *The poset $\Delta(G, B)$ is isomorphic to the set of parabolic subgroups of G , ordered by the opposite of the inclusion relation.* \square

3 The Building Associated to a BN-Pair

Let G be a group with a BN-pair, and assume that the set S of generators of W is finite. [As we remarked at the beginning of Chapter III, this assumption is not really necessary.] We already have our candidate $\Delta = \Delta(G, B)$ for the associated building. So let's check that it really is a building.

Lemma. *The poset Δ is a simplicial complex.*

PROOF: We have the usual two things to check: (a) Any two elements of Δ have a greatest lower bound; and (b) for any $A \in \Delta$, the poset $\Delta_{\leq A}$ is isomorphic to the set of subsets of a finite set.

Given two special cosets gP and $g'P'$, there is certainly a smallest left coset containing them, namely, gP'' , where $P'' = \langle P, P', g^{-1}g' \rangle$. Since P'' contains B , it is special. This proves (a).

As to (b), we have already noted that the function $S' \mapsto B\langle S' \rangle B$ is a poset isomorphism from the set of subsets of S to the set of special subgroups of G . So if we denote by C the special coset B , then we have

$$\Delta_{\leq C} = (\text{subgroups } \supseteq B)^{\text{op}} \approx (\text{subsets of } S)^{\text{op}} \approx (\text{subsets of } S),$$

where the last isomorphism is given by complementation. Using the G -action on Δ , we immediately conclude that (b) holds. \square

Now let's try to construct a system of apartments in Δ . Let Σ be the subcomplex of Δ consisting of the special cosets of the form wP with $w \in W$, and let \mathcal{A} be the set of transforms $g\Sigma$ of Σ by elements of G .

Theorem. *The complex Δ is a thick building, and \mathcal{A} is a system of apartments. The action of G is type-preserving and strongly transitive.*

PROOF: There is a map $\iota : \Sigma(W, S) \rightarrow \Delta(G, B)$ given by $\iota(wW') = wP$, where W' is a special subgroup of W and $P = BW'B$ is the corresponding special subgroup of G . It is easy to check that ι is well-defined (i.e., independent of the choice of coset representative w) and is a simplicial map. Note that the image of ι is the subcomplex Σ defined above.

There is also a map $\rho : \Delta(G, B) \rightarrow \Sigma(W, S)$, given by $\rho(gP) = wW'$ if $g \in C(w)$ and $P = BW'B$. To see that ρ is well-defined, suppose that

$g' = gh$ is another representative of the coset gP . Then $h \in C(w')$ for some $w' \in W'$. Write w' as a word in the generating set $S' = W \cap S$ of W' . Then we have

$$g' = gh \in C(w)C(w') \subseteq \bigcup C(ww''),$$

where w'' ranges over the elements of W' obtained by deleting zero or more letters from the word representing w' [cf. proof of Lemma 1 in §2A above]. So $g' \in C(ww'')$ for some such w'' , and we obtain the same special coset $ww''W' = wW'$ if we use g' instead of g in the definition of ρ . This shows that ρ is well-defined, and it is easy to check that it is a simplicial map.

Clearly $\rho\iota = \text{id}_{\Sigma(W,S)}$. In particular, ι maps $\Sigma(W,S)$ isomorphically onto Σ , and the latter is therefore a Coxeter complex. It follows that (Δ, \mathcal{A}) satisfies axiom **(B0)**. The proof of this has also given us a retraction $\rho' = \iota\rho : \Delta \rightarrow \Sigma$, given by $\rho'(gP) = wP$ if $g \in C(w)$.

To verify **(B1)**, we may assume that one of the two given simplices is a special subgroup P . The other one is then gP' for some $g \in G$ and some special subgroup P' . Writing $g = bnb'$ with $n \in N$ and $b, b' \in B$, we have $gP' = bnP' \in b\Sigma$, so $b\Sigma$ is an apartment containing P and gP' . This proves **(B1)**.

It follows from what we have done so far that Δ is a chamber complex. Moreover, the proof of the lemma essentially constructed a G -invariant labelling of Δ , given by $gP \mapsto S - S'$ if $P = B\langle S' \rangle B$. So Δ is labellable and G is type-preserving. Strong transitivity is also immediate; for G is transitive on \mathcal{A} , and the subgroup N stabilizes Σ and is transitive on $\text{Ch } \Sigma$.

To complete the proof that Δ is a building, we will prove the variant **(B2'')** of axiom **(B2)**. Thus we are given two apartments with a common chamber C , and we must construct an isomorphism between them fixing their intersection. By strong transitivity, we may assume that one of the two apartments is Σ and that C is the special coset B . Let Σ' be the other apartment. Since the stabilizer of C is precisely the subgroup B , we can apply strong transitivity again to find an isomorphism $\phi : \Sigma' \rightarrow \Sigma$ given by the action of some element $b \in B$.

I claim that $\phi = \rho'|_{\Sigma'}$, where ρ' is the retraction constructed above. Indeed, every simplex of $\Sigma' = b^{-1}\Sigma$ has the form $b^{-1}wP$ for some $w \in W$ and some special subgroup P ; the definition of ρ' now gives

$$\rho'(b^{-1}wP) = wP = b \cdot b^{-1}wP = \phi(b^{-1}wP),$$

as claimed. It follows that ϕ fixes $\Sigma' \cap \Sigma$, so **(B2'')** holds and Δ is a building.

Finally, the thickness of Δ is an easy consequence of axiom **(BN2)**. In fact, if you go back to §1G above (where we proved that **(BN2)** is a consequence of thickness), you will see that all the steps in that proof are reversible. \square

EXERCISE

Extract from the first part of the proof the following generalization of the Bruhat decomposition: For any special subgroup $P = BW'B$ of G , there is a bijection

$B \backslash G / P \approx W / W'$, where $B \backslash G / P$ is the set of double cosets of the form BgP . Still more generally, show for any two special subgroups $P' = BW'B$ and $P'' = BW''B$ that $P' \backslash G / P'' \approx W' \backslash W / W''$.

Remarks

1. As we noted in the proof, the fundamental apartment Σ has a chamber C [namely, the special coset B] whose stabilizer is B . The stabilizer of Σ , however, might be strictly bigger than N . So we do not quite have a 1-1 correspondence between BN-pairs and strongly transitive actions on thick buildings.

For a simple example of this, let's go back to the group $G = \mathrm{GL}_3(k)$. We constructed the BN-pair in that example by starting with an action on a building Δ and taking B (resp. N) to be the stabilizer of a fundamental chamber (resp. apartment). But suppose we now replace N by any subgroup $N' < N$ which surjects onto W ; for instance, we could take N' to be the group of permutation matrices. Then it is easy to check that we still have a BN-pair (B, N') and that we still get the same building $\Delta(G, B)$; but now our N' is not the full stabilizer of Σ .

So if we want the group theory to precisely reflect the geometry, we need to add a new axiom which guarantees that N is "big enough". The appropriate axiom turns out to be:

$$(BN3) \quad T = \bigcap_{w \in W} wBw^{-1}.$$

One says that the BN-pair is *saturated* if (BN3) holds.

To see why this is the appropriate axiom, consider an arbitrary BN-pair and the associated building Δ . Let N^* be the stabilizer of the fundamental apartment Σ . Then it is easy to see that $N^* = NT^*$, where $T^* = \{g \in G : g \text{ fixes } \Sigma \text{ pointwise}\}$. Now an element of G fixes Σ pointwise if and only if it stabilizes every chamber of Σ [why?]; so we have $T^* = \bigcap_{w \in W} wBw^{-1}$. Thus (BN3) simply says that $T^* = T$, which implies that $N^* = N$.

In practice, there is no reason to impose (BN3). For if we want to apply geometry to group theory, the important thing is to be able to construct a building associated to a given group.

2. Let $G_0 = \bigcap_{g \in G} gBg^{-1}$; this is the normal subgroup of G consisting of the elements which act trivially on Δ . Let $\bar{G} = G/G_0$. By analogy with the situation for Coxeter groups and their associated complexes, one might expect to be able to recover \bar{G} from Δ as the group $\mathrm{Aut}_0 \Delta$ of type-preserving automorphisms. This turns out to be false in general; counterexamples will be given in §8 below (cf. Remark 3 in §8B).

3. Since the simplices of $\Delta = \Delta(G, B)$ correspond to the parabolic subgroups of G , one can use geometric language to express properties of the parabolic subgroups. Consider, for example, the minimal parabolics (conjugates of B). These correspond to the chambers of Δ , so we can talk about

the distance between two minimal parabolics, and we can even apply the refined distance function δ , which has values in W (cf. Exercise 3 of §IV.4). In the spherical case, we can ask whether two minimal parabolics are opposite to one another. In $GL_3(k)$, for example, the upper triangular group and the lower triangular group are opposite to one another. (To see this, take a look at the picture of the fundamental apartment at the beginning of §1B, and compute the stabilizer of the chamber opposite to C in this apartment.)

EXERCISE

Show that the upper and lower triangular groups are opposite in $GL_n(k)$ for arbitrary n . [HINT: See the exercise at the end of §IV.5.]

4 Historical Remarks

I began Chapter IV by writing down, with no motivation, the strange-looking axioms for buildings. I then showed in the present chapter how these lead in a fairly natural way to some equally strange-looking axioms for groups with a BN-pair. In this brief section I will attempt to put both of these axiom systems in their historical context. They may still seem strange when I'm done, but at least you will have some idea of where they came from.

Our starting point is a 1954 paper of F. Bruhat [21] on the representation theory of complex Lie groups. Bruhat was especially interested in the four classical families A_n , B_n , C_n , D_n of simple matrix groups G . If you're not familiar with these, you can just think about the group $G = SL_n(\mathbb{C})$; this is the classical group of type A_{n-1} . [I'm being a little sloppy here. The group $SL_n(\mathbb{C})$ is not really simple, but it is "almost simple". More precisely, its center Z is finite, and the quotient $PSL_n(\mathbb{C}) = SL_n(\mathbb{C})/Z$ is simple.]

At the time of Bruhat's work, it had been known for a long time how to associate to G a finite reflection group W , called the Weyl group of G . It is given by $W = N/T$, where T is a "maximal torus" and N is its normalizer. And people were becoming aware of the importance of a certain subgroup $B \subset G$ (which eventually became known as the "Borel subgroup" of G as a result of the fundamental work of Borel [12]). What was not yet known, however, was the connection between B and W provided by the Bruhat decomposition $G = \coprod_{w \in W} BwB$.

Bruhat discovered this while studying so-called "induced representations". Questions about these led him to ask whether the set $B \backslash G / B$ of double cosets was finite. He was apparently surprised to discover, by a separate analysis for each of the four families of classical simple groups, that the set of double cosets was not only finite but was in 1-1 correspondence with the finite reflection group W .

The Bruhat decomposition was a fundamental fact that had previously gone unnoticed. Chevalley picked up on it immediately, and it became a basic tool in his work on the construction and classification of simple algebraic groups ([23], [24]). He replaced Bruhat's case-by-case proof by a unified proof that applied not only to the classical groups (types A–D) but also to the five exceptional groups (types E_6 , E_7 , E_8 , F_4 , and G_2). Moreover, he worked over an arbitrary field k , not just $k = \mathbb{C}$. In particular, since k could be finite, one now had for each of the types A–G examples of *finite* simple groups which admitted a Bruhat decomposition, with the Weyl group W being the finite reflection group of the given type. Finally, Chevalley's work included a study of the basic properties of the parabolic subgroups of his matrix groups.

In the early 1960's Tits analyzed Chevalley's methods and extracted the two axioms (BN1) and (BN2). He showed that these two axioms were sufficient to imply Chevalley's results on the Bruhat decomposition and parabolic subgroups. This axiomatization can be found in a short 1962 paper [52], which contains most of the results of §2 above. The only serious omission from this paper is the proof that the Weyl group W associated to a BN-pair is necessarily a Coxeter group; this fact was discovered a year or two later by Tits (cf. [54]) and, independently, by Matsumoto [37].

Meanwhile, Tits had been engaged since the mid 1950's in a different project, in which he was giving geometric interpretations of algebraic matrix groups. "Geometric" here refers to incidence geometry. Thus, in the same way that projective plane geometry is closely related to the group SL_3 [cf. §1B above], Tits constructed incidence geometries associated to very general matrix groups, even more general than those considered by Chevalley. See [58] for Tits's own account of this project and of his early attempts to provide an axiomatic framework for the geometries he was constructing.

By the early 1960's, then, Tits was thinking about axioms for geometries as well as axioms for groups with a Bruhat decomposition. It was natural for him to combine these two lines of thought, and he did this in [53]. In §4.2 of that paper one finds three axioms which look very much like the three axioms for buildings, except that they are stated in terms of incidence geometries instead of simplicial complexes. And in §5 Tits indicates how a group with a BN-pair gives rise to a geometry satisfying his three axioms. In other words, he essentially constructs the building $\Delta(G, B)$. One also finds in this paper some of the fundamental ideas in the theory of buildings, such as retractions onto apartments (phrased in the language of incidence geometry).

Now flag complexes did not actually appear explicitly in [53], but the paper did use flags extensively. So it was only a matter of time before Tits focused on the flag complexes themselves and restated his axioms in terms of simplicial complexes. The first published account of this was given in a 1965 Bourbaki Seminar exposé [54], where buildings were called "complexes with Weyl structure". This paper contains, among other things, an outline

of much of the basic theory of Coxeter complexes and buildings that we gave in Chapters III and IV. It also contains the correspondence between BN-pairs and strongly transitive actions on buildings (§§1 and 3 above).

This completes my highly condensed account of the origin of buildings. I hope it gives you some idea, admittedly vague, as to how Tits discovered buildings by combining (a) years of work on incidence geometries associated to matrix groups and (b) ideas inspired by Chevalley's treatment of the Bruhat decomposition.

There is an interesting footnote to this story. I mentioned that the types A–G of finite reflection groups all arose in Chevalley's work as the Weyl group W of a finite simple group with a BN-pair. What about the remaining types H_3 , H_4 , and $I_2(m)$ ($m = 5$ or $m \geq 7$)? The type $I_2(8)$ (dihedral group of order 16) was observed fairly early; it arises, for instance, from a BN-pair in a finite simple group constructed by Ree, of order $17,971,200 = 2^{11} \cdot 3^3 \cdot 5^2 \cdot 13$. But it turns out that this is the only "unusual" Weyl group that can arise from a finite group (simple or not) with a BN-pair. This is a consequence of a theorem of Feit and Higman [27], which can be restated as follows in the language of buildings:

Theorem. *If Δ is a finite thick building, then every connected component of its Coxeter diagram is of type A_n , B_n , D_n , E_n , F_4 , G_2 , or $I_2(8)$.*

I'll say a few words about the proof. First, it suffices to consider the case where Δ is irreducible, by which we mean that its Coxeter diagram is connected. For in the general case, Δ can be decomposed as a join of irreducible buildings, one for each component of the diagram. Next, it suffices to consider the case where Δ is of rank 2. For the only other cases to worry about are H_3 and H_4 ; and if Δ had either of these types, then a suitable link in Δ would be a finite thick building of the prohibited type $I_2(5)$.

Now a rank 2 building Δ is necessarily the flag complex of a plane incidence geometry P , simply because Δ is labellable. [It's actually true that *all* buildings are flag complexes, but this is not immediately obvious. It is proved in [56], Proposition 3.16; see also §VI.5 below, where we will prove it in a special case.] If one unwinds the characterization of rank 2 buildings (§IV.3, Exercise 2), one finds that P is what is known as a *generalized m -gon*, where m is the integer such that Δ has type $\circ \xrightarrow{m} \circ$. This terminology comes from the fact that P has formal properties analogous to those of the geometry consisting of the vertices and edges of an m -gon; see [58] or [27] for the precise definition.

We are therefore reduced to the following question: For which integers $m \geq 3$ do there exist finite generalized m -gons in which every point is on at least 3 lines and every line contains at least 3 points? The bulk of the Feit–Higman paper is devoted to answering this question, and what they show is that the only possibilities for m are 3, 4, 6, and 8. These correspond, respectively, to the types A_2 , B_2 , G_2 , and $I_2(8)$, whence the theorem. See [27] or [54] for more details.

5 Example: The General Linear Group

Let $G = \mathrm{GL}_n(k)$ ($n \geq 2$), where k is an arbitrary field. [Note: Everything we are about to do goes through if we instead take G to be $\mathrm{SL}_n(k)$, or $\mathrm{PGL}_n(k)$, or $\mathrm{PSL}_n(k)$.] As we have already noted in various remarks and exercises, it is possible to proceed exactly as in the case $n = 3$ (cf. §1B). In other words, one can obtain a BN-pair by using the action of G on the complex Δ of flags of proper non-zero subspaces of k^n . This approach relies, of course, on the fact (proved in Exercise 2 of §IV.2) that Δ is a building.

In the present section we will give an alternate approach, which does not depend on the exercise just cited. Namely, we will simply verify the BN-pair axioms by direct matrix computations. As a byproduct, we will obtain a new proof that the flag complex Δ is indeed a building.

Let $B \subset G$ be the upper triangular group, i.e., the stabilizer of the standard flag

$$[e_1] \subset [e_1, e_2] \subset \cdots \subset [e_1, \dots, e_{n-1}],$$

where e_1, \dots, e_n is the standard basis of k^n . Let $N \subset G$ be the monomial group, i.e., the stabilizer of the set of lines $\{[e_1], \dots, [e_n]\}$. Then N acts as a group of permutations of this set, and we obtain a surjection from N onto the symmetric group on n letters. The kernel of this homomorphism is $T = B \cap N$, which is the diagonal subgroup of G . So N normalizes T , and $W = N/T$ can be identified with the symmetric group on n letters.

To see that B and N generate G , we need only note that the subgroup $\langle B, N \rangle$ contains the lower triangular group, which is wBw^{-1} for a suitable $w \in W$, hence it contains all elementary matrices. [Recall that an elementary matrix is one which has 1's on the diagonal and exactly one non-zero off-diagonal entry; left multiplication (resp. right multiplication) by such a matrix corresponds to an elementary row (resp. column) operation.] It now follows from elementary linear algebra that $\langle B, N \rangle = G$.

Let $S \subset W$ be the standard set of generators $\{s_1, \dots, s_{n-1}\}$, where s_i is the transposition which interchanges i and $i+1$. To simplify the notation, we will verify the axioms (BN1) and (BN2) only for $s = s_1$; the other elements of S are treated similarly. Our s , then, is represented by any monomial matrix of the form

$$\begin{pmatrix} 0 & * & & & \\ * & 0 & & & \\ & & * & & \\ & & & \ddots & \\ & & & & * \end{pmatrix},$$

where the blank regions are understood to be filled with zeroes.

Axiom (BN1) says that $sBw \subseteq BwB \cup BsB$. Multiplying on the right by w^{-1} , we can rewrite this as

$$sB \subseteq BB' \cup BsB',$$

where $B' = wBw^{-1}$. In other words, we must show that any matrix in sB is reducible to either 1 or s via left multiplication by B and right multiplication by B' . It turns out that we will only need to use the elementary matrices in B and B' , so that we will simply be doing some elementary row and column operations (also known as “pivoting”).

Note first that left multiplication by upper triangular elementary matrices allows us to pivot upwards, i.e., to add a multiple of a row to any higher row. Now a typical element of sB has the form

$$\begin{pmatrix} 0 & * & * & \dots & * \\ * & * & * & \dots & * \\ & & * & \dots & * \\ & & & \ddots & \vdots \\ & & & & * \end{pmatrix},$$

which is easily reduced to

$$\begin{pmatrix} 0 & * & & & \\ * & * & & & \\ & & * & & \\ & & & \ddots & \\ & & & & * \end{pmatrix}$$

by pivoting upwards. If the (2,2)-entry is zero, then we have already reduced the matrix to s . So we may assume that all three $*$'s in the upper left 2×2 block are non-zero.

Now let's use right multiplication by B' . To avoid messing up what we have already achieved, we will only use $B' \cap \mathrm{GL}_2$; here GL_2 is identified with the subgroup

$$\{g \in \mathrm{GL}_n : g[e_1, e_2] = [e_1, e_2] \text{ and } ge_i = e_i \text{ for } i > 2\}.$$

Note that $B' = wBw^{-1}$ is the stabilizer of the flag

$$[e_{w(1)}] \subset [e_{w(1)}, e_{w(2)}] \subset \dots \subset [e_{w(1)}, \dots, e_{w(n-1)}].$$

It follows easily that $B' \cap \mathrm{GL}_2$ is the stabilizer in GL_2 of the line spanned by either e_1 or e_2 , whichever occurs first in the list $e_{w(1)}, \dots, e_{w(n)}$. In other words, $B' \cap \mathrm{GL}_2$ is the upper triangular subgroup of GL_2 if $w^{-1}(1) < w^{-1}(2)$ and the lower triangular subgroup otherwise.

Looking at the elementary matrices in $B' \cap \mathrm{GL}_2$, we see that we now have a column operation available: If $w^{-1}(1) < w^{-1}(2)$, then we can add a multiple of column 1 to column 2, and otherwise we can add a multiple of column 2 to column 1. In the first case, we pivot on the (2,1)-entry of our matrix above in order to clear out the (2,2)-entry; this reduces the matrix to s . In the second case, we pivot on the (2,2)-entry in order to clear out the (2,1)-entry; the resulting matrix is in B , and we are done. (Note: The proof of (BN1) actually showed that $sBw \subseteq BswB$ for half of the elements $w \in W$. This should not surprise you.)

Finally, it is trivial to check that (BN2) holds; for we have

$$sBs = \begin{pmatrix} * & 0 & * & \dots & * \\ * & * & * & \dots & * \\ & & * & \dots & * \\ & & & \ddots & \vdots \\ & & & & * \end{pmatrix} \not\subseteq B.$$

Having verified the BN-pair axioms [with remarkably little effort], we obtain a building $\Delta(G, B)$. Let's show that this building is isomorphic to the complex Δ of flags of proper non-zero subspaces of k^n . Consider the action of G on Δ . If C is the standard flag, it is immediate that $\overline{C} = \Delta_{\leq C}$ is a fundamental domain for the action. Moreover, the stabilizers of the faces of C are precisely the special subgroups of G . Indeed, they are special since they contain B , and it is trivial to verify that they are all distinct. Since C has 2^{n-1} faces and G contains only 2^{n-1} special subgroups (one for each subset of S), the stabilizers must exhaust the special subgroups. The desired isomorphism now follows easily. In particular, we have obtained a group-theoretic proof, independent of Exercise 2 of §IV.2, that the flag complex Δ is a building.

6 Example: The Symplectic Group

Let k continue to be an arbitrary field. Let $\langle -, - \rangle$ be the bilinear form on k^{2n} ($n \geq 1$) defined as follows on the standard basis vectors:

$$\langle e_i, e_j \rangle = \begin{cases} 0 & \text{if } i + j \neq 2n + 1 \\ 1 & \text{if } i + j = 2n + 1 \text{ and } i < j \\ -1 & \text{if } i + j = 2n + 1 \text{ and } i > j. \end{cases}$$

If we denote the standard basis vectors by $e_1, e_2, \dots, e_n, f_n, f_{n-1}, \dots, f_1$, then the non-zero “inner products” above can be written more simply as

$$\langle e_i, f_i \rangle = 1 = -\langle f_i, e_i \rangle.$$

The bilinear form $\langle -, - \rangle$ is *alternating* by which we mean that $\langle v, v \rangle = 0$ for all v . [This implies skew-symmetry: $\langle v, v' \rangle = -\langle v', v \rangle$. Conversely, skew-symmetry of a bilinear form implies that the form is alternating, provided $\text{char } k \neq 2$.]

It is easy to explicitly compute $\langle -, - \rangle$ in terms of coordinates: If we write a typical element of k^{2n} as a pair (X, Y) with $X, Y \in k^n$, then we have

$$\langle (X, Y), (Z, W) \rangle = X \cdot W' - Y \cdot Z',$$

where the prime means “reverse the coordinates” and the dot denotes the ordinary dot product of vectors in k^n :

$$(x_1, \dots, x_n) \cdot (y_1, \dots, y_n) = \sum_{i=1}^n x_i y_i.$$

When $n = 1$, for instance, $\langle v, w \rangle$ is simply the determinant of the 2×2 matrix with v and w as columns.

We now define the *symplectic group* $\mathrm{Sp}_{2n}(k)$ to be the group of automorphisms g of k^{2n} which preserve $\langle -, - \rangle$, i.e., which satisfy $\langle gv, gw \rangle = \langle v, w \rangle$ for all $v, w \in k^{2n}$. It is enough to check this equation when v and w are basis vectors. So an element of $\mathrm{Sp}_{2n}(k)$ is a $2n \times 2n$ matrix g whose columns v_1, \dots, v_{2n} satisfy the same inner product relations as the standard basis vectors e_1, \dots, e_{2n} . When $n = 1$, this simply says that $\det g = 1$; thus $\mathrm{Sp}_2 = \mathrm{SL}_2$.

For each $i = 1, \dots, n$ there is a copy of $\mathrm{Sp}_2 [= \mathrm{SL}_2]$ in Sp_{2n} , which stabilizes the plane $[e_i, f_i]$ and fixes all basis vectors other than e_i and f_i . Taking $n = 2$ and $i = 1$, for instance, we obtain a copy of SL_2 in Sp_4 that looks like this:

$$\begin{pmatrix} * & & & * \\ & 1 & 0 & \\ & 0 & 1 & \\ * & & & * \end{pmatrix}$$

In addition, there are various ways to embed GL_2 in Sp_{2n} . Namely, given $1 \leq i < j \leq n$, there is a copy of GL_2 which stabilizes $[e_i, e_j]$ and $[f_i, f_j]$ and fixes all basis vectors other than these four; an automorphism g of this type can do anything at all on $[e_i, e_j]$, but its effect on $[f_i, f_j]$ is then forced by the requirement that g be symplectic. Suppose, for example, that we take $n = 2$ again and try to construct an element $g \in \mathrm{Sp}_4$ which is given by an elementary matrix $\begin{pmatrix} 1 & a \\ 0 & 1 \end{pmatrix}$ on $[e_1, e_2]$ and which stabilizes $[f_1, f_2] = [e_4, e_3]$. Then g must have the form

$$\begin{pmatrix} 1 & a & & \\ 0 & 1 & & \\ & & * & * \\ & & * & * \end{pmatrix},$$

and a simple computation shows that this will be symplectic if and only if the lower right 2×2 block is $\begin{pmatrix} 1 & -a \\ 0 & 1 \end{pmatrix}$.

Finally, for each $i < j$ as above there is also a copy of GL_2 in Sp_{2n} which stabilizes $[e_i, f_j]$ and $[e_j, f_i]$.

Call a symplectic matrix *elementary* if it is the image of a 2×2 elementary matrix under one of the embeddings described in the three previous paragraphs. You might find it a useful exercise to explicitly write down all the types of elementary matrices in Sp_4 . (There are two copies of SL_2 and two copies of GL_2 , hence 8 types of elementary matrices.) You can also check as an exercise that Sp_{2n} is generated by elementary matrices. The idea is to interpret multiplication by elementary matrices in terms of row and column operations. You will then find it easy to use these operations to reduce any symplectic matrix to the form

$$\begin{pmatrix} I & \\ & A \end{pmatrix},$$

where I and A are $n \times n$ matrices and I is the identity. But then A is forced to be the identity also since the matrix is symplectic.

We need one last bit of terminology before constructing a BN-pair in $\mathrm{Sp}_{2n}(k)$. A subspace $V \subset k^{2n}$ is called *totally isotropic* if $\langle v, v' \rangle = 0$ for all $v, v' \in V$. This is equivalent to saying that $V \subseteq V^\perp$, where the orthogonal subspace V^\perp is defined in the usual way. Note, for instance, that a subset of the standard basis spans a totally isotropic subspace provided it contains no pair $\{e_i, f_i\}$. The chain of totally isotropic subspaces

$$[e_1] \subset [e_1, e_2] \subset \cdots [e_1, \dots, e_n]$$

will be called the *standard isotropic flag* in k^{2n} . Note that the subspaces orthogonal to these totally isotropic subspaces form a descending chain

$$[e_1, \dots, e_{2n-1}] \supset [e_1, \dots, e_{2n-2}] \supset \cdots \supset [e_1, \dots, e_n];$$

so if we take the standard isotropic flag together with the orthogonal subspaces, we get the standard ordinary flag in k^{2n} (with the subspace $[e_1, \dots, e_n]$ counted twice). Incidentally, the set of non-zero totally isotropic subspaces, with inclusion as the incidence relation, is an example of what is called a *polar geometry*.

Now let B be the group of upper triangular symplectic matrices, i.e., the stabilizer in $G = \mathrm{Sp}_{2n}(k)$ of the standard flag in k^{2n} . If a symplectic matrix stabilizes a subspace V , then it stabilizes V^\perp too; hence B can also be described as the stabilizer in G of the standard isotropic flag. Let $N \subset G$ be the group of symplectic monomial matrices, i.e., the stabilizer in G of the set of lines $\{L_1, \dots, L_n, L'_n, \dots, L'_1\}$, where $L_i = [e_i]$ and $L'_i = [f_i]$. Then $T = B \cap N$ is the group of diagonal symplectic matrices, i.e., the group of matrices of the form $\mathrm{diag}(\lambda_1, \dots, \lambda_n, \lambda_n^{-1}, \dots, \lambda_1^{-1})$. In particular, N normalizes T . Note that T is isomorphic to the product of n copies of k^* ; in the language of the theory of algebraic groups, T is a torus of rank n .

The quotient $W = N/T$ can be identified with a group of permutations of the set of $2n$ lines above. We will show that W is equal to the group W' consisting of all permutations which map each pair $\{L_i, L'_i\}$ to another such pair. The inclusion $W \subseteq W'$ is immediate, since W preserves the orthogonality relations among the given lines. To prove the opposite inclusion, note first that W' is generated by the following set $S = \{s_1, \dots, s_n\}$ of permutations: s_i for $i < n$ is the product of the two transpositions $L_i \leftrightarrow L_{i+1}$ and $L'_i \leftrightarrow L'_{i+1}$; and s_n is the transposition $L_n \leftrightarrow L'_n$. So the inclusion $W' \subseteq W$ follows from the easy observation that each $s \in S$ can be represented by a symplectic monomial matrix in one of our embedded SL_2 's or GL_2 's.

One can now use elementary row and column operations, exactly as in the case of GL_n , to complete the proof that we have a BN-pair. Details are left to the reader. [You might want to check (BN1) for $n = 2$ and $s = s_1$; for instance, just to convince yourself that the same method really does work.] The crucial thing that keeps the proof from becoming unpleasant is that each s_i is in a GL_2 or SL_2 ; one is thereby able to reduce (BN1)

to a 2×2 computation before ever having to think about what the group $B' = wBw^{-1}$ looks like.

Finally, one can check that the associated building $\Delta = \Delta(G, B)$ is isomorphic to the flag complex of the set of non-zero totally isotropic subspaces. The proof is essentially the same as the proof of the analogous statement for GL_n , provided that you know the basic linear algebra of alternating forms. See, for instance, Artin [8], Chapter III.

EXERCISE

Draw a picture of the fundamental apartment in Δ when $n = 2$; it is a barycentrically subdivided quadrilateral, whose 8 vertices consist of the totally isotropic subspaces spanned by subsets of the standard basis.

Remarks

1. You probably noticed that the Weyl group W is isomorphic to the “signed permutation group on n letters”, which is the finite reflection group of type B_n or C_n , cf. §I.3. [Strictly speaking, type B_n was only defined for $n \geq 2$ in Chapter I; but we make the convention that $B_1 = A_1$.] This calculation of W is consistent with the fact that Sp_{2n} is said to be of type C_n in the classification of matrix groups.

2. The “inner product” $\langle -, - \rangle$ that we worked with may have seemed arbitrary. But, in fact, one can show that it is the typical non-degenerate alternating bilinear form, in the following sense: If V is a finite dimensional vector space with a non-degenerate alternating bilinear form $\langle -, - \rangle$, then $\dim V$ is even and V has a basis $e_1, \dots, e_n, f_n, \dots, f_1$ whose inner products look like those of our example. A proof can be found in the book of Artin cited above.

7 Example: The Orthogonal Group

We assume now that $\text{char } k \neq 2$, although much of what we do would work in characteristic 2 also. Let $\langle -, - \rangle$ be the symmetric bilinear form on k^m ($m \geq 2$) defined as follows on the standard basis vectors e_1, \dots, e_m :

$$\langle e_i, e_j \rangle = \begin{cases} 1 & \text{if } i + j = m + 1 \\ 0 & \text{otherwise.} \end{cases}$$

We will write $m = 2n$ (resp. $m = 2n + 1$) if m is even (resp. odd), and we will denote the last n basis vectors by f_n, \dots, f_1 . The non-zero inner products, then, are

$$\langle e_i, f_i \rangle = 1 = \langle f_i, e_i \rangle$$

and, if $m = 2n + 1$,

$$\langle e_{n+1}, e_{n+1} \rangle = 1.$$

It is easy to explicitly compute $\langle -, - \rangle$ in terms of coordinates: Write a typical vector $v \in k^m$ as (X, Y) if m is even and as (X, λ, Y) if m is odd, with $X, Y \in k^n$ and $\lambda \in k$; then we have

$$\begin{aligned}\langle (X, Y), (Z, W) \rangle &= X \cdot W' + Y \cdot Z' \\ \langle (X, \lambda, Y), (Z, \mu, W) \rangle &= X \cdot W' + Y \cdot Z' + \lambda\mu,\end{aligned}$$

where the prime and the dot product have the same meaning as in §6. In particular, the associated quadratic form Q is given by

$$Q(X, Y) = 2X \cdot Y' \quad \text{or} \quad Q(X, \lambda, Y) = 2X \cdot Y' + \lambda^2.$$

This form is equivalent, under change of coordinates, to the form Q' given by

$$Q'(z_1, \dots, z_m) = -z_1^2 - \dots - z_n^2 + z_{n+1}^2 + \dots + z_m^2.$$

And if k contains $\sqrt{-1}$, then it is equivalent to the “standard” quadratic form $\sum_{i=1}^m z_i^2$.

We now define the *orthogonal group* $O_m(k, Q)$ to be the group of automorphisms g of k^m which preserve $\langle -, - \rangle$ or, equivalently, Q . The *special orthogonal group* $SO_m(k, Q)$ is defined to be the group of orthogonal matrices of determinant 1. We will often suppress Q from the notation and simply write $O_m(k)$ and $SO_m(k)$, since the quadratic form Q defined above is the only one we will consider.

Note that SO_2 is the “rank 1 torus” consisting of diagonal matrices $\text{diag}(\lambda, \lambda^{-1})$. It has index 2 in O_2 , the non-trivial coset being the set of matrices

$$\begin{pmatrix} 0 & \lambda^{-1} \\ \lambda & 0 \end{pmatrix},$$

which have determinant -1 .

Let's focus now on $G = SO_m$, returning to the case of O_m afterwards. For each $i = 1, \dots, n$ we have a copy of SO_2 in G which stabilizes $[e_i, f_i]$ and fixes all the other basis vectors. In case $m = 2n + 1$, we can extend this to an embedding $O_2 \hookrightarrow G$ by using the “extra” basis vector $v = e_{n+1}$: For if $g \in O_2$ has determinant -1 , then we can copy g on $[e_i, f_i]$ and then send v to $-v$ in order to make the determinant 1.

Next, given $1 \leq i < j \leq n$, there are two ways of embedding GL_2 in $G = SO_m$, exactly analogous to the two embeddings used for Sp_{2n} . In particular, this gives us lots of elementary matrices to work with.

We now construct the BN-pair in the usual way: B is the upper triangular subgroup of G and N is the monomial subgroup. Then $T = B \cap N$ consists of the diagonal elements of G . If $m = 2n$, these elements necessarily have the form

$$\text{diag}(\lambda_1, \dots, \lambda_n, \lambda_n^{-1}, \dots, \lambda_1^{-1}),$$

and if $m = 2n + 1$ they have the form

$$\text{diag}(\lambda_1, \dots, \lambda_n, 1, \lambda_n^{-1}, \dots, \lambda_1^{-1}).$$

In both cases, T is a “rank n torus”. The Weyl group $W = N/T$ can be identified with a group of permutations of the $2n$ lines $L_i = [e_i]$ and $L'_i = [f_i]$, $i = 1, \dots, n$. This is clear if $m = 2n$, but it is also true if $m = 2n+1$; for e_{n+1} is the only non-isotropic basis vector, so W necessarily fixes the line it spans.

If $m = 2n+1$, then W is the same permutation group W' as in the case of Sp_{2n} . One proves this by exhibiting elements of N which represent the generators s_i of W' ($i = 1, \dots, n$) constructed in §6. For $i < n$, the required element of N can be found in the embedded GL_2 acting on $[e_i, e_{i+1}]$ and $[f_i, f_{i+1}]$. And for $i = n$, the required element can be found in the embedded O_2 acting on $[e_n, e_{n+1}, f_n]$. This calculation of W is consistent with the fact that the matrix group SO_{2n+1} is said to be of type B_n .

If $m = 2n$, however, it is impossible to represent s_n by an orthogonal monomial matrix of determinant 1. The group W in this case turns out to be a subgroup of index 2 in W' . This subgroup is generated by the s_i for $i < n$ together with one additional element t , which is the product of the transpositions $L_{n-1} \leftrightarrow L'_n$ and $L_n \leftrightarrow L'_{n-1}$. Note that t is in the embedded GL_2 acting on $[e_{n-1}, f_n]$ and $[e_n, f_{n-1}]$. (I am assuming here that $n \geq 2$; if $n = 1$, then $G = \mathrm{SO}_2$, and we have already said everything there is to say about that group.)

In both cases, we now have a set of n generators for W , and it is a routine (although somewhat tedious) matter to verify the BN-pair axioms, using the same methods as in §§5 and 6. We have already identified the Weyl group in case $m = 2n+1$. In case $m = 2n$, you can check, by computing orders of products of generators, that W is of type D_n . [Strictly speaking, D_n was only defined for $n \geq 4$ in Chapter I; but the appropriate convention is that $D_3 = A_3$ and that the diagram of type D_2 is the union of two copies of the diagram of type A_1 .] This calculation of W is consistent with the fact that the matrix group SO_{2n} is said to be of type D_n .

Remark. Did the definition of t above seem *ad hoc*? Was there a different choice of t that seemed more natural to you? If so, you would have struggled in vain to verify the BN-pair axioms for your choice. Indeed, we know from the general theory that, given B and N , there can only be one set S for which the BN-pair axioms hold (cf. §2B, Exercise).

Let's try now to figure out what the building $\Delta(G, B)$ is. The naïve guess is that it is the flag complex Δ of non-zero totally isotropic subspaces of k^m . This guess is correct if m is odd, and the proof is the same as the proofs of the analogous assertions in §§5 and 6. But it is wrong if m is even, as one can see in a variety of ways. For one thing, the flag complex Δ has the wrong type. [It is the flag complex of a polar geometry and hence has type B_n ; but we've already seen that the Weyl group of SO_{2n} has type D_n .] For another thing, one can show that the action of SO_{2n} on $\mathrm{Ch} \Delta$ is not transitive. [There are precisely two orbits.] Yet a third thing that goes wrong is that Δ is not thick. [A simple computation shows that there are

only two ways of extending the flag

$$[e_1] \subset \cdots \subset [e_1, \dots, e_{n-1}]$$

to a maximal flag.]

The correct answer, when m is even, turns out to be that $\Delta(G, B)$ is the flag complex of the following so-called “oriflamme geometry” P : The elements of P are the non-zero totally isotropic subspaces of k^{2n} of dimension $\neq n - 1$; two such subspaces are called incident if one is contained in the other or if both have dimension n and their intersection has dimension $n - 1$. As an example of a flag in P we have the following chamber C :

$$[e_1] \subset \cdots \subset [e_1, \dots, e_{n-2}] \begin{array}{l} \subset [e_1, \dots, e_{n-1}, e_n] \\ \subset [e_1, \dots, e_{n-1}, f_n] \end{array}$$

We can get some feeling for this by thinking about the case $n = 2$ and describing P in the language of projective geometry: An isotropic line in k^4 is simply a point in the 3-dimensional projective space \mathbf{P}^3 over k that lies on the quadric surface X defined by $Q = 0$. And an isotropic plane in k^4 is simply a line in \mathbf{P}^3 which is contained in X . Our geometry P , then, consists of lines in the surface X , two such being called incident if they intersect. If you believe, as asserted above, that the flag complex of P is a building of type $D_2 = A_1 \amalg A_1$, then it must be true that there are two types of lines in X , and that every line of one type intersects every line of the other type (cf. §IV.2, Example 2). One can actually see this directly, by exhibiting an isomorphism between X and the direct product $\mathbf{P}^1 \times \mathbf{P}^1$ of two copies of the projective line; the two types of lines in X , then, are simply the two types of slices of the product. Details can be found in van der Waerden [61], §I.7, which contains an interesting discussion of the groups SO_m for $3 \leq m \leq 6$.

Finally, what happens if we look at the full orthogonal group O_m instead of SO_m ? For m odd, everything goes through with no essential change. In particular, we again get a BN-pair, with the associated building being the same as the building for SO_m . This is not surprising, since $O_m = SO_m \times \{\pm 1\}$ if m is odd, so there is virtually no difference between the two groups.

When m is even, on the other hand, the situation is more complicated. One still has an action of $G = O_{2n}$ on the oriflamme complex Δ , but the action is not type-preserving. For it is easy to give examples of orthogonal matrices which stabilize the flag C constructed above but do not fix it pointwise. So the action of G on this building does not yield a BN-pair in G . We obtain, instead, something called a *generalized BN-pair*. See Bourbaki [16], §IV.2, Exercise 8, if you want to know precisely what this means. I’ll say more about it in the next section, in connection with a different group.

One last comment about $G = O_{2n}$: In addition to the building Δ of type D_n , one still has the non-thick building of type B_n , consisting of flags

of totally isotropic subspaces. The action of G on this building is type-preserving and strongly transitive, so all of the results of §1 are applicable except those which used thickness. In particular, properties (1)–(6) of §2A all hold, with B equal to the upper triangular subgroup of O_{2n} [which happens to be the same as the upper triangular subgroup of SO_{2n}] and with the Weyl group being of type B_n . One thus has a choice of two geometries associated to G , one of type D_n [which yields a generalized BN-pair] and one of type B_n [which yields what could be called a “weak” BN-pair]. The former might seem more natural since G is classically considered to be of type D_n , but the latter has had applications also; see, for instance, Vogtmann [63].

8 Example: The Special Linear Group Over a Field With Discrete Valuation

Up to now, all of our examples of BN-pairs have had finite Weyl groups (and hence spherical buildings). It turns out that the same matrix groups that occurred in those examples admit a second BN-pair structure whenever the ground field comes equipped with a discrete valuation. This was first noticed by Iwahori and Matsumoto [33] and was later generalized to a much larger class of groups by Bruhat and Tits [22]. The Weyl group for this second BN-pair is an infinite Euclidean reflection group, and the associated building therefore has apartments which are Euclidean spaces. We will illustrate this by treating the groups SL_n . But first we must review discrete valuations.

8A Discrete valuations

Let K be a field and K^* its multiplicative group of non-zero elements. A *discrete valuation* on K is a surjective homomorphism $v : K^* \rightarrow \mathbf{Z}$ satisfying the following inequality:

$$v(x + y) \geq \min\{v(x), v(y)\}$$

for all $x, y \in K^*$ with $x + y \neq 0$. It is convenient to extend v to a function defined on all of K by setting $v(0) = +\infty$; the inequality then remains valid for all $x, y \in K$. Note that we necessarily have $v(-1) = 0$ since \mathbf{Z} is torsion-free; hence $v(-x) = v(x)$. It follows from this and the inequality above that the set $A = \{x \in K : v(x) \geq 0\}$ is a subring of K ; it is called the *valuation ring* associated to K . And any ring A which arises in this way from a discrete valuation is called a *discrete valuation ring*.

The group A^* of units of A is precisely the kernel $v^{-1}(0)$ of v . So if we pick an element $\pi \in K$ with $v(\pi) = 1$, then every element $x \in K^*$ is uniquely expressible in the form $x = \pi^n u$ with $n \in \mathbf{Z}$ and $u \in A^*$. In particular, K is the field of fractions of A .

The principal ideal πA generated by π can be described in terms of v as $\{x \in K : v(x) > 0\}$. It is a maximal ideal, since every element of A not in πA is a unit. The quotient ring $k = A/\pi A$ is therefore a field, called the *residue field* associated to the valuation v .

Example. Let K be the field \mathbf{Q} of rational numbers, and let p be a prime number. The *p-adic valuation* on \mathbf{Q} is defined by setting $v(x)$ equal to the exponent of p in the prime factorization of x . More precisely, given $x \in \mathbf{Q}^*$, write $x = p^n u$, where n is a (possibly negative) integer and u is a rational number whose numerator and denominator are not divisible by p ; then $v(x) = n$. The valuation ring A is the ring of fractions a/b with $a, b \in \mathbf{Z}$ and b not divisible by p . [The ring A happens to be the localization of \mathbf{Z} at p , but we will not make any use of this.] The residue field k is the field \mathbf{F}_p of integers mod p ; one sees this by using the homomorphism $A \rightarrow \mathbf{F}_p$ given by $a/b \mapsto (a \bmod p)(b \bmod p)^{-1}$, where a and b are as above.

The valuation ring A in this example can be described informally as the largest subring of \mathbf{Q} on which reduction mod p makes sense. It is thus the natural ring to introduce if one wants to relate the field \mathbf{Q} to the field \mathbf{F}_p . This illustrates our point of view toward valuations: We will be interested in studying things (namely, matrix groups) defined over a field K , and we wish to “reduce” to a simpler field k as an aid in this study; a discrete valuation makes this possible by providing us with a nice ring A to serve as intermediary between K and k :

$$\begin{array}{ccc} A & \hookrightarrow & K \\ \downarrow & & \\ k & & \end{array}$$

Returning now to the general theory, we note that the study of the arithmetic of A (e.g., ideals and prime factorization) is fairly trivial:

Proposition 1. *A discrete valuation ring A is a principal ideal domain, and every non-zero ideal is generated by π^n for some $n \geq 0$. In particular, πA is the unique non-zero prime ideal of A .*

PROOF: Let I be a non-zero ideal and let $n = \min\{v(a) : a \in I\}$. Then I contains π^n , and every element of I is divisible by π^n ; hence $I = \pi^n A$. \square

One consequence of this is that we can apply the basic facts about modules over a principal ideal domain (e.g., a submodule of a free module is free). Let’s recall some of these facts, in the form in which we’ll need them later. Let V be the vector space K^n . By a *lattice* (or *A-lattice*) in V we will mean an A -submodule $L \subset V$ of the form $L = Ae_1 \oplus \cdots \oplus Ae_n$ for some basis e_1, \dots, e_n of V . In particular, L is a free A -module of rank n . If we take e_1, \dots, e_n to be the standard basis of V , then the resulting lattice is A^n , which we call the *standard lattice*.

If L' is a second lattice in V , then we can choose our basis e_1, \dots, e_n for L in such a way that L' admits a basis of the form $\lambda_1 e_1, \dots, \lambda_n e_n$ for some scalars $\lambda_i \in K^*$. This fact should be familiar to you for the case $L' \subseteq L$, and the general case follows easily. [Choose a large integer M such that $\pi^M L' \subseteq L$, and apply the usual theory to $L'' = \pi^M L'$.] The scalars λ_i can be taken to be powers of π , and they are then unique up to order. They are called the *elementary divisors* of L' with respect to L .

All of this follows from well-known results about modules over principal ideal domains. But I will sketch the proof of part of it (namely, the existence of the e_i and λ_i) in the case at hand, where A is a discrete valuation ring; for the proof involves ideas that will be needed later anyway.

Start with arbitrary bases of L and L' , and express the basis elements of L' as linear combinations of those of L ; this yields an element of $GL_n(K)$. It is easy to see that this matrix can be reduced to a monomial matrix by means of integral row and column operations, where “integral” means that the operation is given by multiplication by an elementary matrix in $SL_n(A)$. [In other words, when we add a scalar multiple of one row or column to another, the scalar is required to be in A .] To see that this is possible, choose a matrix entry a_{ij} with $v(a_{ij})$ minimal. Then pivot to clear out everything other than a_{ij} in the i th row and j th column, noting that this pivoting only requires integral row and column operations. Now ignore the i th row and j th column and repeat the process, using an element of minimal valuation in the rest of the matrix. It is clear that we will eventually obtain a monomial matrix by continuing in this way.

The row and column operations above correspond to changes of basis in L and L' . So what we have just done is to replace the given bases of L and L' by new ones, such that the new basis elements of L' are scalar multiples of the new basis elements of L . This completes the proof.

Note that if the matrix in $GL_n(K)$ that we started with above happened to be in $SL_n(K)$, then the same would be true of the monomial matrix that we ended with. So we obtain, as a byproduct of the proof:

Proposition 2. $SL_n(K)$ is generated by its monomial subgroup together with the elementary matrices in $SL_n(A)$. \square

We end this review of discrete valuations by commenting briefly on the notion of *completeness*. A discrete valuation v induces a real-valued absolute value on K , defined by

$$|x| = e^{-v(x)}.$$

We then have

$$|xy| = |x| \cdot |y| \quad \text{and} \quad |x + y| \leq \max\{|x|, |y|\}.$$

This inequality is a very strong form of the triangle inequality. In particular, we get a metric on K by setting $d(x, y) = |x - y|$. It therefore makes sense to ask whether K is complete, in the sense that every Cauchy sequence

converges. If not, then one can form the *completion* \hat{K} of K by formally adjoining limits of Cauchy sequences, in exactly the same way that one constructs \mathbf{R} in elementary analysis by completing \mathbf{Q} . The only difference is that the construction is actually easier in the present context, because of the strong form of the triangle inequality. In fact, one can build the completion purely algebraically by using inverse limits; see, for instance, Atiyah–MacDonald [9], Chapter 10.

The field operations and the function v extend to \hat{K} by continuity, and \hat{K} is again a field with a discrete valuation. Its valuation ring is the completion \hat{A} of A (i.e., the closure of A in \hat{K}), and its residue field is the same as that of K . In case the residue field k is finite, one can show that \hat{A} is compact; since \hat{A} is the closed unit ball in \hat{K} , the latter is locally compact in this case.

The canonical example for all this is the p -adic valuation on \mathbf{Q} discussed above. The completion is the field \mathbf{Q}_p of p -adic numbers. It is a complete, locally compact, discretely valued field, with residue field \mathbf{F}_p . Its valuation ring is called the ring of p -adic integers.

8B The group $\mathrm{SL}_n(K)$

Let K continue to denote a field with a discrete valuation v , and let A , π , and k be as in §8A. We then have a diagram of matrix groups

$$\begin{array}{ccc} \mathrm{SL}_n(A) & \hookrightarrow & \mathrm{SL}_n(K) \\ \downarrow & & \\ \mathrm{SL}_n(k) & & \end{array}$$

which we will use to construct a BN-pair in $\mathrm{SL}_n(K)$ by “lifting” the BN-pair in $\mathrm{SL}_n(k)$ that we studied in §5. More precisely, we will take B to be the inverse image in $\mathrm{SL}_n(A)$ of the upper triangular subgroup of $\mathrm{SL}_n(k)$, but we will take N , as before, to be the monomial subgroup of $\mathrm{SL}_n(K)$. [It wouldn’t make sense to also construct N as an inverse image, since B and N would then both be subgroups of $\mathrm{SL}_n(A)$ and hence could not possibly generate $\mathrm{SL}_n(K)$.]

Note that B contains the upper triangular subgroup of $\mathrm{SL}_n(A)$; the subgroup generated by B and N therefore contains both the upper triangular and lower triangular subgroups of $\mathrm{SL}_n(A)$ and hence all elementary matrices in $\mathrm{SL}_n(A)$. This subgroup is therefore the whole group $\mathrm{SL}_n(K)$ by Proposition 2 of §8A. The intersection $T = B \cap N$ is the diagonal subgroup of $\mathrm{SL}_n(A)$, which is easily checked to be normalized by N ; in fact, the conjugation action of N on T simply permutes the diagonal entries of a matrix in T .

We need some notation in order to describe the group $W = N/T$. For any commutative ring R , let $N(R)$ (resp. $T(R)$) denote the monomial (resp. diagonal) subgroup of $\mathrm{SL}_n(R)$. Then our N and T above are $N(K)$ and $T(A)$, so $W = N(K)/T(A)$. Let $\overline{W} = N(K)/T(K)$, identified as usual with the

symmetric group on n letters. Then \overline{W} is a quotient of W , and we have a short exact sequence

$$1 \rightarrow T(K)/T(A) \rightarrow W \rightarrow \overline{W} \rightarrow 1.$$

This sequence splits, since the subgroup $N(A)/T(A) \subset W$ maps isomorphically to the quotient \overline{W} ; so we have

$$W \approx F \rtimes \overline{W},$$

where $F = T(K)/T(A) \approx (K^*/A^*)^{n-1}$. Note that the valuation v induces an isomorphism $K^*/A^* \xrightarrow{\approx} \mathbf{Z}$, so the normal subgroup F above is free abelian of rank $n - 1$. In order to understand the action of \overline{W} on this free abelian group, identify F with $\{(x_1, \dots, x_n) \in \mathbf{Z}^n : \sum_{i=1}^n x_i = 0\}$. [The isomorphism of F with this group is obtained by applying v to the n diagonal entries of an element of $T(K)$.] The action of \overline{W} on F , then, simply permutes the n coordinates. We now need to find a suitable set of generators for W and verify the BN-pair axioms. Let's start with the case $n = 2$, deferring the general case to §8C below.

When $n = 2$, we have $W \approx \mathbf{Z} \rtimes \{\pm 1\}$, with the non-trivial action of $\{\pm 1\}$ on \mathbf{Z} ; this is the infinite dihedral group. It is generated by s_1 and u , where s_1 is the non-trivial element of the $\{\pm 1\}$ factor and u is a generator of the infinite cyclic normal subgroup $F = T(K)/T(A) \subset W$. So W is generated by the set $S = \{s_1, s_2\}$ of elements of order 2, where $s_2 = s_1 u$. Let's take u to be represented by the element $\mathrm{diag}(\pi, \pi^{-1}) \in T(K)$; then s_1 is represented by

$$\begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$$

and s_2 is represented by

$$\begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} \pi & 0 \\ 0 & \pi^{-1} \end{pmatrix} = \begin{pmatrix} 0 & -\pi^{-1} \\ \pi & 0 \end{pmatrix}.$$

Note that it is by no means clear, *a priori*, that we have made the right choice of u —we could replace u by u^{-1} and still get a set of two generators of W of order 2. But, as we noted when discussing SO_{2n} in §7, only one of these choices can be “right” (in the sense that the BN-pair axioms hold). We have, in fact, made the right choice, and one can easily verify the axioms. The verification of (BN1) is slightly tedious since it requires consideration of several cases. As an example, let's take $s = s_2$ and $w = (s_1 s_2)^r$, where $r > 0$. In this case we'll show that $sBw \subseteq BswB$, which is what should be true since $l(sw) = l(w) + 1$.

Note first that the subgroup $B \subset \mathrm{SL}_2(K)$ can be described by the following conditions on the valuations of the matrix entries:

$$\begin{pmatrix} v = 0 & v \geq 0 \\ v \geq 1 & v = 0 \end{pmatrix}.$$

Computing sBw , we find that its elements satisfy

$$\begin{pmatrix} v \geq r & v = -r - 1 \\ v = r + 1 & v \geq -r + 1 \end{pmatrix}.$$

The unique entry with minimal valuation is in the upper right-hand corner, so we pivot at this position to clear out the diagonal entries. This requires multiplication by elementary matrices of the form

$$\begin{pmatrix} 1 & 0 \\ \pi^2 a & 1 \end{pmatrix}$$

with $a \in A$. Elementary matrices of this form are in B , so the pivoting operations are legal and we have reduced our matrix to a monomial matrix in $\mathrm{SL}_2(K)$ whose upper right-hand corner has valuation $-r - 1$. This monomial matrix is equivalent mod T to the matrix

$$\begin{pmatrix} 0 & -\pi^{-r-1} \\ \pi^{r+1} & 0 \end{pmatrix},$$

which represents sw . This completes the verification of our special case of (BN1). The other cases are equally easy.

Now let's try to describe the building $\Delta(G, B)$. Note first that one of the special subgroups is $\mathrm{SL}_2(A)$, since this is a subgroup containing B . Since $\mathrm{SL}_2(A)$ is the stabilizer in $\mathrm{SL}_2(K)$ of the standard A -lattice A^2 in K^2 , this suggests that the vertices of $\Delta(G, B)$ should correspond to lattices in K^2 . [There is, of course, an obvious action of $\mathrm{SL}_2(K)$, and even $\mathrm{GL}_2(K)$, on the set of lattices.] On the other hand, our experience in §5 suggests that Δ should admit an action of $\mathrm{PGL}_2(K) = \mathrm{GL}_2(K)/Z$, where Z is the group of scalar multiples of the identity. So it is more reasonable to expect the vertices to correspond to Z -orbits of lattices. With this as motivation, we proceed to describe the building.

Call two A -lattices L, L' in K^2 equivalent if $L = \lambda L'$ for some $\lambda \in K^*$. Note that the scalar λ can then be taken to be a power of π . Let $[L]$ denote the equivalence class of a lattice L . If L is given as $Af_1 \oplus Af_2$ for some basis f_1, f_2 of K^2 , then we will also write $[[f_1, f_2]]$ for the class $[L]$.

I want to assign a "type" to a lattice class. To this end, consider the obvious action of $\mathrm{GL}_2(K)$ on the set of lattice classes. This action is transitive, and the stabilizer of $[A^2]$ is $Z \cdot \mathrm{GL}_2(A)$, where Z is as above. It follows that $v(\det g)$ is an even integer for every g in this stabilizer. We can now say that a lattice class Λ is of type 0 (resp. type 1) if $v(\det g)$ is even (resp. odd) for every $g \in \mathrm{GL}_2(K)$ such that $g[A^2] = \Lambda$. In other words, the type of $[[f_1, f_2]]$ is $v(\det(f_1, f_2)) \bmod 2$, where $\det(f_1, f_2)$ is the determinant of the matrix with f_1 and f_2 as columns.

Call two distinct lattice classes Λ, Λ' *incident* if they have representatives L, L' which satisfy

$$\pi L \subset L' \subset L.$$

Note that the representatives $\pi L, L'$ then satisfy $\pi L' \subset \pi L \subset L'$, so the incidence relation is symmetric. Note also that, in this situation, the ele-

mentary divisors of L' with respect to L are necessarily 1 and π , so we have $\Lambda = [[f_1, f_2]]$ and $\Lambda' = [[f_1, \pi f_2]]$ for some basis f_1, f_2 of K^2 . It follows that Λ and Λ' are of different types, so we have a plane incidence geometry.

Let's show now that the flag complex Δ of this geometry is isomorphic to $\Delta(G, B)$ for $G = \mathrm{SL}_2(K)$ as above. We will prove this, as usual, by finding a fundamental domain and computing stabilizers. Let C be the edge with vertices $[[e_1, e_2]]$ and $[[e_1, \pi e_2]]$, where e_1, e_2 is the standard basis of K^2 . Let $C' = \{\Lambda, \Lambda'\}$ be an arbitrary edge, with Λ of type 0. Then there is a basis f_1, f_2 such that $\Lambda = [[f_1, f_2]]$, $\Lambda' = [[f_1, \pi f_2]]$, and $\det(f_1, f_2) = \pi^{2r}u$ for some $r \in \mathbb{Z}$ and $u \in A^*$. Replacing f_1 by $\pi^{-r}u^{-1}f_1$ and f_2 by $\pi^{-r}f_2$, we still have $\Lambda = [[f_1, f_2]]$ and $\Lambda' = [[f_1, \pi f_2]]$, but now $\det(f_1, f_2) = 1$. So the matrix g with f_1 and f_2 as columns is an element of G such that $gC = C'$. Since the action of G is type-preserving, it follows easily that \bar{C} is a fundamental domain.

The stabilizer of $[A^2]$ in G is $\mathrm{SL}_2(K) \cap (Z \cdot \mathrm{GL}_2(A)) = \mathrm{SL}_2(A)$. And the stabilizer of $[[e_1, \pi e_2]]$ is the conjugate $g\mathrm{SL}_2(A)g^{-1}$, where $g = \mathrm{diag}(1, \pi)$. This conjugate is the subgroup of G defined by

$$\begin{pmatrix} v \geq 0 & v \geq -1 \\ v \geq 1 & v \geq 0 \end{pmatrix}.$$

The stabilizer of C , then, which is the intersection of the stabilizers of its two vertices, is precisely B . The desired isomorphism $\Delta \approx \Delta(G, B)$ follows easily.


The fundamental apartment Σ is obtained by applying the elements of the monomial group N to C ; it is a line, with vertices $[[\pi^a e_1, \pi^b e_2]]$, $a, b \in \mathbb{Z}$. An arbitrary apartment $g\Sigma$ is the same sort of line, but with e_1, e_2 replaced by an arbitrary basis of K^2 .

Remarks

1. Since $\Delta(G, B)$ is a building of type $\circ \xrightarrow{\infty} \circ$, we know from §IV.3 that the flag complex Δ is a tree. We could have simply proven this directly (cf. Serre [46], Chapter II) and deduced the BN-pair structure in G without any matrix computations. (We would then have had to check strong transitivity, but this is easy.)

2. This example shows that incomplete apartment systems arise naturally. In fact, it is not hard to see that the apartment system $\mathcal{A} = \{g\Sigma\}$ is complete if and only if the field K is complete with respect to the valuation v . One can show further that if an arbitrary K is replaced by its completion \hat{K} , then the building Δ remains the same—all that changes is the apartment system, which gets completed. These assertions, and their analogues for SL_n with $n > 2$, will be proved in the next chapter (§VI.9F).

Consider, for example, the case where K is the field \mathbb{Q}_2 of 2-adic numbers. Then one can show that Δ is the tree pictured below, with uncountably many apartments. If we instead take $K = \mathbb{Q}$ (with the 2-adic valuation),

The building for $SL_2(\mathbb{Q}_2)$

then we get the same tree, equipped with a certain countable apartment system.

3. The isomorphism type of the tree Δ associated to $SL_2(K)$ depends only on the cardinality of the residue field $k = A/\pi A$. If $k = \mathbb{F}_2$, for instance, then every vertex is on exactly 3 edges, so the tree is necessarily the one pictured above. This shows that many different choices of (K, v) can yield the same tree Δ , even if we stick to the case where K is complete. Thus there is no hope of recovering the group G , or even the quotient \bar{G} defined in Remark 2 of §3 above, from the tree Δ . In particular, one cannot expect that \bar{G} is the group of type-preserving automorphisms of Δ . I recommend as an exercise that you consider the case where K has residue field \mathbb{F}_2 and show directly that the tree pictured above admits type-preserving automorphisms not in \bar{G} (which is $PSL_2(K)$).

This discussion might seem to suggest that there is a very poor correspondence between groups with a BN-pair and buildings, unlike the situation for Coxeter groups and Coxeter complexes. But the tree case is atypical in this regard, and the correspondence is much better for some other classes of buildings; see Tits ([56], [57], [60]).

Before moving on to SL_n for $n > 2$, we look briefly at what happens if we replace $SL_2(K)$ by $GL_2(K)$. It is clear that $GL_2(K)$ acts on the flag complex Δ ; but the action does not preserve types. This is the same situation that we saw at the end of §7, so we obtain some kind of generalized BN-pair structure in $GL_2(K)$. I'll illustrate this by proving one result about double cosets in $GL_2(K)$.

Let (G, B, N, S) continue to have the same meaning as above, with $G =$

$\mathrm{SL}_2(K)$, etc. The result to be proved is that

$$sBm \subseteq BmB \cup BsmB$$

for any $s \in S$ and any monomial matrix $m \in \mathrm{GL}_2(K)$. It would be quite easy to deduce this by direct matrix calculations from the axiom **(BN1)** for G . But I prefer to imitate the geometric argument given in §1, since this will illustrate how one can get information about a group from an action which is not type-preserving. As in §1, the geometric argument will prove the following equivalent form of the assertion:

$$mBs \subseteq BmB \cup BmsB.$$

Let $\tilde{s} \in N$ be a representative of s , and consider an arbitrary element $g = m\tilde{s} \in mBs$. Note that the three chambers mC , msC , and gC have a common codimension 1 face. Now the first two of these chambers are in Σ , since m stabilizes Σ . Hence $\rho(gC) = mC$ or msC , where $\rho = \rho_{\Sigma, C}$. On the other hand, we know from §1 that $\rho(gC) = b'gC$ for some $b' \in B$, so $gC = b'^{-1}mC$ or $b'^{-1}msC$. Thus $g \in BmB' \cup BmsB'$, where B' is the stabilizer of C in $\mathrm{GL}_2(K)$. But all elements of $Bm \cup Bms$ have the same determinant as g ; so we can replace B' by $B' \cap \mathrm{SL}_2(K) = B$ to get $g \in BmB \cup BmsB$, as required.

8C The group $\mathrm{SL}_n(K)$, concluded

We continue with the notation of §8B, but we now assume $n \geq 3$. Recall that $W = F \rtimes \overline{W}$, where $F = T(K)/T(A) \approx \mathbf{Z}^{n-1}$ and $\overline{W} = N(A)/T(A)$. We need a set $S = \{s_1, \dots, s_n\}$ of generators of W . For the first $n-1$ of these we use the standard generators of the symmetric group \overline{W} ; thus s_i for $i < n$ is represented by a monomial matrix in the embedded $\mathrm{SL}_2(A) \subset \mathrm{SL}_2(K) \hookrightarrow \mathrm{SL}_n(K)$ acting on $[e_i, e_{i+1}]$. And for s_n we take the element of W represented by

$$\tilde{s}_n = \begin{pmatrix} 0 & & & & -\pi^{-1} \\ & 1 & & & \\ & & \ddots & & \\ & & & 1 & \\ \pi & & & & 0 \end{pmatrix}.$$

This monomial matrix is in the embedded $\mathrm{SL}_2(K)$ acting on $[e_n, e_1]$. To see that S generates W , note first that the subgroup $W' = \langle S \rangle$ contains \overline{W} . Multiplying s_n by a suitable element of \overline{W} , we conclude that W' also contains the element of F represented by $\mathrm{diag}(\pi, 1, \dots, 1, \pi^{-1})$. Conjugating this element by \overline{W} , we obtain a set of generators for F , so $W' = W$.

The verification of **(BN2)** presents no problem. To check **(BN1)**, one proceeds as in §§5–7, the idea being to use row operations to reduce to a 2×2 matrix computation. As an illustration, here are the details for the case where $n = 3$ and $s = s_3$.

The statement to be proved is that $sB \subseteq BB' \cup BsB'$, where $B' = mBm^{-1}$ for some monomial matrix m of determinant 1. Motivated by what we just did at the end of §8B, let's prove this inclusion more generally for an arbitrary monomial matrix. Now an element of sB has the form

$$\begin{pmatrix} v \geq 0 & v \geq 0 & v = -1 \\ v \geq 1 & v = 0 & v \geq 0 \\ v = 1 & v \geq 1 & v \geq 1 \end{pmatrix},$$

so we start by pivoting at the upper right-hand corner to clear out the two entries below it. (The row operations required for this are given by left multiplication by lower triangular elements of B .) This leaves us with

$$\begin{pmatrix} v \geq 0 & v \geq 0 & v = -1 \\ v \geq 1 & v = 0 & 0 \\ v = 1 & v \geq 1 & 0 \end{pmatrix}.$$

We can also make the middle entry (the one with $v = 0$) equal to 1; for this can be achieved by multiplication by an element of T .

Now pivot on this middle entry to clear out the entries above and below it. This yields

$$\begin{pmatrix} v \geq 0 & 0 & v = -1 \\ v \geq 1 & 1 & 0 \\ v = 1 & 0 & 0 \end{pmatrix}.$$

Pivoting at the lower left-hand corner, finally, reduces us to a matrix in the copy of SL_2 which contains $s = s_3$:

$$\begin{pmatrix} v \geq 0 & 0 & v = -1 \\ 0 & 1 & 0 \\ v = 1 & 0 & 0 \end{pmatrix}.$$

For the rest of the proof we ignore the middle row and middle column and work entirely in the SL_2 that remains. We need some notation. Let $G_3 = SL_3(K)$ and let G_2 be the embedded $SL_2(K)$ that we have just reduced to, i.e., $G_2 = \{g \in G_3 : ge_2 = e_2, g[e_1, e_3] = [e_1, e_3]\}$. Let B_3 and B_2 be the corresponding B 's. The matrix above, then, is an element of $sB_3 \cap G_2$, and we wish to reduce it to 1 or s by multiplying on the left by $B_3 \cap G_2$ and on the right by $B' \cap G_2$.

It is easy to check that $B_3 \cap G_2 = B_2$ and that $sB_3 \cap G_2 = \tilde{s}B_2$, where \tilde{s} is the monomial matrix \tilde{s}_3 that we wrote down above. So we will have reduced our problem about G_3 to the same sort of problem for G_2 [which we solved in §8B], provided $B' \cap G_2$ is a " B' -type" subgroup of G_2 . To finish the proof, then, we need to understand what kind of subgroup of G_3 can arise as a B' , and we need to show that $B' \cap G_2$ is the "same kind" of subgroup of G_2 .

Let L be the standard lattice in K^3 , with basis e_1, e_2, e_3 . Then we can identify $L/\pi L$ with the vector space k^3 over the residue field k . Note that

any $g \in SL_3(A)$ stabilizes L and hence acts on $L/\pi L = k^3$; indeed, this is one way of describing the homomorphism $SL_3(A) \rightarrow SL_3(k)$ that we wrote down at the beginning of §8B. We can now describe $B = B_3$ as follows: Given $g \in G_3$, we have

$$g \in B_3 \iff gL = L \text{ and } gC = C,$$

where C is the standard flag in k^3 . Consequently, $B' = mBm^{-1}$ admits a similar description:

$$g \in B' \iff gL' = L' \text{ and } gC' = C',$$

where $L' = mL$ and C' is the flag in $L'/\pi L'$ corresponding to C under the isomorphism $L/\pi L \rightarrow L'/\pi L'$ induced by m .

Now the lattice L' has a basis $\pi^a e_1, \pi^b e_2, \pi^c e_3$ for some $a, b, c \in \mathbb{Z}$. So the k -vector space $L'/\pi L'$ comes equipped with a “standard basis”, and C' is simply the “permuted standard flag” obtained from a permutation of that standard basis. [The permutation that arises is the one corresponding to the monomial matrix m .] It is now easy to describe $B' \cap G_2$: Let L'_2 be the lattice in $[e_1, e_3]$ with basis $\pi^a e_1, \pi^c e_3$; given $g \in G_2$, we have


$$g \in B' \cap G_2 \iff gL'_2 = L'_2 \text{ and } gC'_2 = C'_2,$$

where C'_2 is a certain permuted standard flag in $L'_2/\pi L'_2$. This characterization of $B' \cap G_2$ as a subgroup of G_2 is the exact analogue of the characterization of B' as a subgroup of G_3 , so we are done.

Remark. If we had simply tried to prove (BN1) instead of a generalization of it, then we would have assumed $\det m = 1$. The only difference this would have made is that we would have had $a + b + c = 0$ in the description of L' . But the analogous sum $a + c$ for L'_2 would not necessarily have been 0, so we would still have needed the generalized (BN1) for the 2×2 case. In other words, we needed to understand GL_2 , and not just SL_2 , in order to deal with SL_n for $n \geq 3$.

The building Δ associated to $SL_n(K)$ is a flag complex as in the case $n = 2$: One considers classes of lattices in K^n , one assigns a type to any such class by taking the valuation of a determinant and reducing mod n , and one defines incidence exactly as before. The fundamental chamber is the simplex with vertices $[[e_1, \dots, e_i, \pi e_{i+1}, \dots, \pi e_n]]$, $i = 1, \dots, n$. Further details are left to the interested reader.

There are a number of interesting things to say about this example, and they motivate much of what we will do in the next chapter. Consider first the Weyl group $W = \mathbb{Z}^{n-1} \rtimes \overline{W}$. It is not one that we have seen before. Computing the orders of the products $s_i s_j$, one finds that its Coxeter diagram is



where there are n vertices all together. [We are still assuming $n \geq 3$; the diagram for $n = 2$ is different.] We will see in the next chapter (§VI.1F) that W is a Euclidean reflection group acting on \mathbf{R}^{n-1} . If $n = 3$, for instance, W is the group of isometries of the plane generated by the reflections in the sides of an equilateral triangle. The apartments in this case are planes tiled by equilateral triangles, and these planes are somehow glued together to form the building Δ . This can be viewed as a 2-dimensional analogue of a tree, which is constructed by gluing lines together.

If we delete a vertex from the Coxeter diagram above, we obtain the diagram of type A_{n-1} . So the link of a vertex in our “Euclidean building” Δ is a spherical building of type A_{n-1} . (This is true if $n = 2$ also.) If we take, for instance, the vertex $[[e_1, \dots, e_n]]$ whose stabilizer is $SL_n(A)$, this link is some building of type A_{n-1} which comes equipped with an action of $SL_n(A)$. The obvious guess is that the link is the spherical building associated to $SL_n(k)$ as in §5, and it is easy to check that this guess is correct.

There is, of course, another building of type A_{n-1} that one would naturally think of, namely, the building Δ' obtained by forgetting that K has a valuation and applying §5 to $SL_n(K)$. This building admits an action of the full group $SL_n(K)$. Is it somehow related to our Euclidean building Δ also? The answer is that every Euclidean building gives rise in a canonical way to a spherical building “at infinity”, obtained by attaching a sphere at infinity to each apartment. When this procedure is applied to Δ , it yields Δ' . Details will be given at the end of the next chapter.

This connection between Δ and Δ' provides a nice geometric explanation of the fact that both of our BN-pairs in $SL_n(K)$ used the same N . For N is the stabilizer of the fundamental apartment in both Δ and Δ' ; the use of the same N therefore yields a 1-1 correspondence between the apartments in Δ and those in Δ' . The previous paragraph explains this correspondence geometrically.

All of this is easy to understand when $n = 2$ if you know about ends of trees. The building Δ' in this case is 0-dimensional, so it is simply the discrete set G/B [with B equal to the upper triangular group], in which the two-element subsets have been called apartments. This set can be identified with a set of ends of the tree Δ (cf. Serre [46], §II.1.3), so Δ' is clearly a building “at infinity” associated to Δ . And the 1-1 correspondence between apartment systems simply reflects the fact that a line in a tree gives rise to a pair of ends and that, conversely, a pair of ends determines a unique line.

VI

Euclidean Buildings

The last example in Chapter V motivates a systematic study of buildings in which the apartments are Euclidean spaces. We need to begin by understanding the apartments themselves. This requires that we go back and re-do much of Chapter I in a more general setting.

1 Euclidean Reflection Groups

Let V be a real vector space of finite dimension $n \geq 1$. All of the geometric notions introduced in Chapter I treated the origin of V as a special point. Our hyperplanes, for instance, were required to go through the origin. Our reflection groups therefore fixed the origin, and our cells were all cones with the origin as cone point. By the end of that chapter it had become clear that we weren't doing Euclidean geometry at all, but rather spherical geometry. In Euclidean geometry, there is nothing special about the origin. So let's introduce the appropriate language for talking about reflections whose fixed hyperplane does not necessarily pass through the origin.

1A Affine concepts

An *affine subspace* of V is a subset of the form $x + V_0$ with $x \in V$ and V_0 a linear subspace of V . In other words, it is a coset of a linear subspace. The *dimension* of $x + V_0$ is defined to be the dimension of V_0 . If the dimension is $n - 1$ (i.e., if V_0 is a linear hyperplane), then $x + V_0$ is called an *affine hyperplane*. Equivalently, an affine hyperplane is a subset defined by a linear equation of the form $f = c$, where $f : V \rightarrow \mathbf{R}$ is a non-zero linear map and c is a constant.

For any non-empty subset $X \subseteq V$, there is a smallest affine subspace containing X , called the *affine span* of X . An *affine frame* for V is a subset X such that V is the affine span of X but not of any proper subset of X . Such a frame necessarily has exactly $n + 1$ elements; for we may assume that one of the elements is the origin, in which case the remaining elements form a basis for V . When $n = 2$, for instance, an affine frame is simply a set of 3 non-collinear points.

An *affine map* from V to a vector space V' is a map α of the form $\alpha(x) = g(x) + v'$, where $g : V \rightarrow V'$ is linear and v' is a vector in V' . In other words,

α is the composite $\tau_{v'}g$, where $\tau_{v'}$ is the translation $x' \mapsto x' + v'$. We will mainly be interested in the case where $V' = V$ and g is an automorphism, in which case we will call α an *affine automorphism* of V . Such an α is uniquely expressible as $\tau_v g$ with $v \in V$ and $g \in \text{GL}(V)$, where $\text{GL}(V)$ is the group of linear automorphisms of V . One deduces easily that the group $\text{Aff}(V)$ of affine automorphisms of V is the semi-direct product

$$\text{Aff}(V) = V \rtimes \text{GL}(V),$$

where V is identified with the (normal) subgroup consisting of translations. The $\text{GL}(V)$ -component g of an element $\alpha \in \text{Aff}(V)$ will be called the *linear part* of α .

Assume, now, that V is *Euclidean*, by which we mean that it comes equipped with a positive definite inner product $\langle -, - \rangle$. We then have a distance function on V given by $d(x, y) = \|x - y\|$, where $\|v\| = \sqrt{\langle v, v \rangle}$, and we will be interested in affine maps which are isometries. With the notation above, the affine map α is an isometry if and only if its linear part g is in the orthogonal group $\text{O}(V) \subset \text{GL}(V)$, consisting of automorphisms which preserve the inner product. Thus the group of affine isometries of V is $V \rtimes \text{O}(V)$. It is worth noting that the word “affine” is redundant here; see Exercise 3 below.

Let H be an affine hyperplane and let H_0 be the linear hyperplane parallel to H , i.e., satisfying $H = x + H_0$ for some x . Let s_0 be the orthogonal reflection s_{H_0} , and let $s = \tau_x s_0 \tau_{-x}$; in other words, s is the conjugate of s_0 by some translation taking H_0 to H . Explicitly, we have

$$sy = x + s_0(y - x) = s_0y + (1 - s_0)x$$

for any $y \in V$. In particular, s is an affine isometry whose linear part is s_0 . It is easy to check that s depends only on H , and not on the choice of the representative x ; this follows, for instance, from the formula above together with the observation that x is unique mod $H_0 = \ker(1 - s_0)$. We can therefore write $s = s_H$ and call s the *reflection* with respect to H .

EXERCISES

1. Let x_0, x_1, \dots, x_n be an affine frame. Show that a point $x \in V$ is completely determined by the $n + 1$ numbers $d(x, x_i)$, i.e., if $d(x, x_i) = d(x', x_i)$ for all i , then $x = x'$. [HINT: We may assume $x_0 = 0$. Then if we know the numbers $d(x, x_i)$, we can compute the inner products $\langle x, x_i \rangle$ via the identity $d^2(x, y) = \|x\|^2 + \|y\|^2 - 2\langle x, y \rangle$, which is essentially the law of cosines.]

2. Let x_0, \dots, x_n be an affine frame and let y_0, \dots, y_n be points such that $d(y_i, y_j) = d(x_i, x_j)$ for all i, j . Show that there is an affine isometry α such that $\alpha(x_i) = y_i$ for all i . In particular, y_0, \dots, y_n is an affine frame. [HINT: Law of cosines again.]

3. Deduce from Exercises 1 and 2 that every isometry $\alpha : V \rightarrow V$ is affine.

4. Show that the reflection s_H can be characterized as the unique non-trivial isometry of V that fixes H pointwise.

5. Let X and X' be isometric subsets of V . Show that every isometry $X \rightarrow X'$ extends to an isometry of V . [We will only need this result when V is the affine span of X . In this case one can find an affine frame x_0, \dots, x_n in X and argue as in Exercise 3. The proof in the general case is similar but takes slightly more work.]

1B Affine reflection groups

We continue to assume that V is a Euclidean vector space of dimension $n \geq 1$. Let W be a group of affine isometries generated by reflections s_H , where H ranges over a set \mathcal{H} of affine hyperplanes. It is clear that we can enlarge \mathcal{H} , if necessary, to make it W -invariant. We will say that W is an *affine reflection group* if we can find such a W -invariant family \mathcal{H} which is locally finite, in the sense that every point of V has a neighborhood which meets only finitely many $H \in \mathcal{H}$. Much of what we did in Chapter I then goes through with little or no change, although some of the arguments get slightly longer since we now only have local finiteness of \mathcal{H} instead of finiteness. I will sketch the theory, including proofs only for those results that are really special to the affine case. All omitted details can be found in Bourbaki [16], Chapter V.

The hyperplanes $H \in \mathcal{H}$ yield a partition of V into convex *cells*, these being non-empty sets A defined by linear equalities or strict inequalities, one for each $H \in \mathcal{H}$. More precisely, if H is defined by a linear equality $f = c$, then the definition of A will involve either the same equality or else one of the inequalities $f > c$ or $f < c$. A cell A has a *support*, defined as in Chapter I; the support is an affine subspace of V , and A is open relative to its support. The *dimension* of A is the dimension of its support. The cells of maximal dimension n are called *chambers*; they are the connected components of the complement of \mathcal{H} in V . Cells have faces, with properties similar to those proven in Chapter I.

The supports of the codimension 1 faces of a chamber C are called the *walls* of C , and C is defined by the inequalities corresponding to its walls. [Warning: This is the first place where it is not entirely a routine matter to generalize the proof given in Chapter I; the basic idea remains the same, but the proof needs to be rearranged.]

Everything we have said so far applies to any locally finite collection of affine hyperplanes. Now let's bring W into the picture. Choose a chamber C and let S be the set of reflections with respect to the walls of C . Note that S , *a priori*, might be infinite; we'll return to this question below. Let H_s , for $s \in S$ be the hyperplane fixed by s , and let e_s be the unit vector perpendicular to H_s and pointing toward C . Thus e_s^\perp is the linear hyperplane parallel to H_s , and one of the defining inequalities for C has the form $\langle e_s, - \rangle > c$. Let $m(s, t)$ for $s, t \in S$ be the (possibly infinite) order of st .

The following basic facts proved in the finite case remain valid:

- (a) W is simply-transitive on the chambers.
- (b) W is generated by S .
- (c) \mathcal{H} necessarily consists of all affine hyperplanes H with $s_H \in W$.
- (d) (W, S) is a Coxeter system.
- (e) $\langle e_s, e_t \rangle = -\cos(\pi/m(s, t))$ for all $s, t \in S$. Moreover, $m(s, t) = \infty$ if and only if H_s and H_t are parallel.
- (f) \overline{C} is a fundamental domain for the action of W on V , and the stabilizer of a point $x \in \overline{C}$ is the special subgroup of W generated by $\{s \in S : sx = x\}$.

Everything except (d) is proved as in Chapter I. For (d), one can argue as in §II.2A, or one can verify the exchange condition as a byproduct of the proof of (a). Incidentally, the possibility $m(s, t) = \infty$ mentioned in (e) hardly ever occurs. In fact, we will see in §1D below that the infinite dihedral group provides the only irreducible example whose Coxeter matrix involves ∞ . [Alternatively, instead of appealing to §1D, one can apply Corollary 3 in §III.2; for we will see below that the Coxeter complex $\Sigma(W, S)$ triangulates V if W is infinite and irreducible.]

1C Finiteness results

Let's now settle the question of the finiteness of S , along with some related questions:

Theorem.

- (1) C has only finitely many walls, hence S is finite.
- (2) The hyperplanes $H \in \mathcal{H}$ fall into finitely many classes under the relation of parallelism; in other words, there are only finitely many linear hyperplanes H_0 such that \mathcal{H} contains a translate of H_0 .
- (3) Let $\overline{W} \subset \text{GL}(V)$ be the set of linear parts of the elements $w \in W$, i.e., the image of W under the projection $\text{Aff}(V) \rightarrow \text{GL}(V)$. Then \overline{W} is a finite reflection group.

PROOF: (1) The inner product formula in (e) above shows that the angle $\angle(e_s, e_t)$ between e_s and e_t satisfies $\angle(e_s, e_t) \geq \pi/2$ for $s \neq t$. But if S were infinite, then the e_s would have a cluster point on the unit sphere and hence there would be s, t with $\angle(e_s, e_t)$ very small.

(2) Let $R = \{\pm e_H : H \in \mathcal{H}\}$, where e_H is the unit vector perpendicular to H and pointing toward C . We must show that R is finite. As in the proof of (1), it suffices to show that $\angle(e, e')$ is bounded away from 0 for $e \neq e'$ in R . We will show that, in fact, there are only finitely many possibilities for this angle. Let H and H' be elements of \mathcal{H} perpendicular to e and e' , respectively. If H and H' are parallel, then $\angle(e, e') = \pi$ [since $e \neq e'$]. Otherwise, choose $x \in H \cap H'$ and choose $w \in W$ such that $wx \in \overline{C}$; this is possible by (f) above. Then wH and wH' are elements of \mathcal{H} which

meet \overline{C} , and they are perpendicular to the vectors $\bar{w}e, \bar{w}e'$, where \bar{w} is the linear part of w . Since $\angle(e, e') = \angle(\bar{w}e, \bar{w}e')$, the proof will be complete if we show that only finitely many elements of \mathcal{H} meet \overline{C} . Now C has only finitely many walls by (1); so it is defined by finitely many inequalities and hence has only finitely many faces. And each face meets only finitely many elements of \mathcal{H} by local finiteness and the definition of “cell”. The union \overline{C} of the faces therefore meets only finitely many elements of \mathcal{H} .

(3) The set R defined in the proof of (2) is \overline{W} -invariant. Since it was proven to be finite, \overline{W} is a finite reflection group by the proposition in §I.1. \square

1D The structure of C

Call W *essential* if the associated finite reflection group \overline{W} is essential. It is easy to reduce the general case to the essential case, as in Chapter I. Similarly, we can decompose V according to the irreducible components of the Coxeter diagram of (W, S) and thereby reduce to the irreducible case. In this case we will prove:

Theorem. *Assume that W is essential and irreducible. Then C is either a simplex or a simplicial cone. More precisely, one of the following holds:*

(a) *W is finite and has a fixed point. In this case C has exactly n walls, where $n = \dim V$, and C is a simplicial cone.*

(b) *C has exactly $n + 1$ walls, and any n of the $n + 1$ vectors e_s are linearly independent. The essentially unique linear relation $\sum_{s \in S} \lambda_s e_s = 0$ among the e_s has all of its coefficients λ_s non-zero and of the same sign. The chamber C is a simplex, and W is infinite.*

Remarks

1. If W has a fixed point x , then we can always assume that $x = 0$; for we can replace W by its conjugate $\tau_{-x}W\tau_x$. But if W fixes 0, then W is linear. Thus case (a), for practical purposes, is precisely the “spherical” case treated in Chapter I. Case (b), then, describes the “genuinely Euclidean” irreducible reflection groups.

2. It follows from the theorem that the numbers $m(s, t)$ are always finite when W is irreducible, unless $n = 1$ and $W = D_\infty$; for an n -simplex with $n \geq 2$ cannot have parallel walls.

3. If we drop the assumption that W is essential and irreducible, then C is a product of a vector space [corresponding to the inessential part of V], a simplicial cone [corresponding to the product of the finite irreducible factors of W], and simplices [one for each infinite irreducible factor of W].

PROOF OF THE THEOREM: Let H_1, \dots, H_r be the walls of C , and let e_1, \dots, e_r be the corresponding unit vectors (pointing toward C). Then e_1, \dots, e_r span V ; for $[e_1, \dots, e_r]^\perp$ is the fixed point set of \overline{W} , hence it is the trivial subspace. We therefore have $r \geq n$.

Suppose $r = n$, so that the vectors e_1, \dots, e_r form a basis. Then C is defined by n inequalities $\langle e_i, - \rangle > c_i$, hence it is a simplicial cone; its cone point is the unique $x \in V$ such that $\langle e_i, x \rangle = c_i$ for all i . Replacing W by a conjugate, we may assume that $x = 0$. Then the generating reflections of W are linear, so $W = \overline{W}$ and we are in case (a).

Now suppose $r > n$, so that e_1, \dots, e_r are linearly dependent. Choose a linear relation

$$\sum_{i \in I} \lambda_i e_i = 0$$

with $\emptyset \neq I \subseteq \{1, \dots, r\}$ and $\lambda_i \neq 0$ for all $i \in I$. Since $\langle e_i, e_j \rangle \leq 0$ for $i \neq j$, we can replace I by a subset, if necessary, to get a relation with $\lambda_i > 0$ for all $i \in I$; this follows from the first paragraph of the proof in §I.4D. We can now deduce from the irreducibility assumption that I is the entire set of indices $\{1, \dots, r\}$. For suppose it is not, and let J be the complementary set. Then for any $j \in J$ we have $\sum_{i \in I} \lambda_i \langle e_i, e_j \rangle = 0$, which implies that $\langle e_i, e_j \rangle = 0$ for all $i \in I$. But then the parts of the Coxeter diagram corresponding to I and J are disjoint, contradicting irreducibility.

Our relation now has the form $\sum_{i=1}^r \lambda_i e_i = 0$, with $\lambda_i > 0$ for all i . But we arrived at this relation by starting with an arbitrary relation among a subset of the e_i and then possibly passing to a further subset. It follows that every proper subset of the e_i is linearly independent, hence $r = n + 1$.

Since e_1, \dots, e_n form a basis for V , the intersection $\bigcap_{i=1}^n H_i$ consists of a single point, which we may take to be the origin as in the discussion of case (a). The chamber C is therefore defined by inequalities $\langle e_i, - \rangle > 0$ for $i = 1, \dots, n$ and $\langle e_{n+1}, - \rangle > c$ for some constant c . Since e_{n+1} is a negative linear combination of e_1, \dots, e_n , the last inequality can be rewritten in the form $\sum_{i=1}^n \mu_i \langle e_i, - \rangle < c'$ with $\mu_i > 0$. The constant c' is necessarily positive [and hence our original c was negative], since otherwise C would be empty. So we may multiply the inequality by a scalar in order to arrange that $c' = 1$. Thus C is defined by inequalities $f_i > 0$ for $i = 1, \dots, n$ and $\sum f_i < 1$, where $f_i = \mu_i \langle e_i, - \rangle$. Since the f_i form a basis for the dual space V^* , these inequalities define an open n -simplex.

Everything in (b) has now been proved except for the assertion that W is infinite. But this follows from the fact that W has compact fundamental domain \overline{C} ; for if W were finite, then $V = \bigcup_{w \in W} w\overline{C}$ would be compact. \square

By a *Euclidean reflection group* we will mean an essential irreducible infinite affine reflection group, as in case (b). In this case we will denote by Σ , or $\Sigma(W, V)$, the poset consisting of the cells in V together with the empty set. Since the chambers are simplices, it is easy to see (as in §I.5G) that Σ is an abstract simplicial complex and that there is a canonical bijection $|\Sigma| \approx V$. If you know how to topologize the geometric realization of an infinite simplicial complex (cf. [48], §3.1), then you can easily check that this bijection is a homeomorphism. But we will not need this fact; for the only topology on $|\Sigma|$ that we will use is the one that is *defined* so as to

make the bijection above a homeomorphism.

Finally, we record the following consequence of property (f) of §1B:

Proposition. *If W is a Euclidean reflection group, then the complex $\Sigma = \Sigma(W, V)$ is isomorphic to the Coxeter complex $\Sigma(W, S)$.* \square

1E The structure of W

Let W be a Euclidean reflection group as above. Recall that we have a surjection $W \twoheadrightarrow \overline{W}$, where \overline{W} is the finite reflection group consisting of the linear parts of the elements of W . The first result about the structure of W is that this surjection always splits:

Proposition 1. *There exist points $x \in V$ such that the stabilizer W_x maps isomorphically onto \overline{W} .*

PROOF: Let $\overline{\mathcal{H}}$ be the set of linear hyperplanes H such that H is parallel to some element of \mathcal{H} . Then \overline{W} is generated by $\{s_H : H \in \overline{\mathcal{H}}\}$. Since \overline{W} is essential, it follows from Chapter I that \overline{W} is actually generated by n such reflections s_H , whose hyperplanes form the walls of a simplicial cone. Choose $H_1, \dots, H_n \in \mathcal{H}$ parallel to these walls, and let $s_i = s_{H_i}$. Then $\bigcap_{i=1}^n H_i$ consists of a single point x , which is fixed by each s_i , and the linear parts of the s_i generate \overline{W} . This shows that W_x surjects onto \overline{W} . But W_x also injects into \overline{W} ; for $\ker\{W_x \rightarrow \overline{W}\}$ consists of translations which fix x , hence it is trivial. \square

Replacing W by a conjugate, if necessary, we may assume that $x = 0$. The elements of $W_x = W_0$ are then linear, and the isomorphism $W_0 \rightarrow \overline{W}$ above is the identity map. Thus W contains the linear part of each of its elements. It follows that W also contains the translation component of each of its elements, so we have

$$W = L \rtimes \overline{W} \subset V \rtimes \mathrm{GL}(V),$$

where $L = \{v \in V : \tau_v \in W\}$. We will complete our analysis of the structure of W by showing that L is a *lattice* in V , by which we mean a subgroup of the form $\mathbb{Z}e_1 \oplus \dots \oplus \mathbb{Z}e_n$ for some \mathbb{R} -basis e_1, \dots, e_n of V .

Lemma 1. *L is a discrete subgroup of the additive group of V , i.e., there is a neighborhood U of 0 in V such that $U \cap L = \{0\}$. The quotient group V/L , with the quotient topology, is compact.*

PROOF: Pick a chamber C and a point $y \in C$. Since the transforms wC for $w \in W$ are all disjoint, the same is true of the translates $C + l$ for $l \in L$. So if we set $U = \{v \in V : y + v \in C\}$, then U is a neighborhood of 0 in V such that $U \cap L = \{0\}$. Recall now that the closed simplex \overline{C} is a fundamental domain for the action of W . It follows that every point of V is equivalent mod L to a point of the compact set $\bigcup_{w \in \overline{W}} w\overline{C}$, hence V/L is compact. \square

Lemma 2. *If L is a discrete subgroup of a finite dimensional vector space V , then $L = \mathbb{Z}e_1 \oplus \cdots \oplus \mathbb{Z}e_r$ for some linearly independent vectors e_1, \dots, e_r . If, in addition, V/L is compact, then L is a lattice.*

PROOF: The second assertion follows immediately from the first. The following proof of the first assertion is taken from Pontryagin [39], Chapter 3, §19, Example 33, where a more general result is proved. We argue by induction on $\dim V$. If $L = 0$ there is nothing to prove, so assume $L \neq 0$. Give V an arbitrary inner product and choose (by discreteness) a non-zero vector $e \in L$ of minimal length δ . Then any $l \in L$ which is not in $\mathbb{Z}e$ has distance at least $\delta/2$ from the line $\mathbb{R}e$. For if $y \in \mathbb{R}e$, then we can find $l' \in \mathbb{Z}e$ with $d(y, l') \leq \delta/2$; since $d(l, l') = \|l - l'\| \geq \delta$, the triangle inequality implies that $d(l, y) \geq \delta/2$, as claimed. Consider now the subgroup $L/(L \cap \mathbb{R}e) = L/\mathbb{Z}e \subset V/\mathbb{R}e$. If we give $V/\mathbb{R}e$ the metric induced by the canonical isomorphism $e^\perp \approx V/\mathbb{R}e$, then what we have just proven is that every non-zero element of $L/\mathbb{Z}e$ has distance at least $\delta/2$ from the origin of $V/\mathbb{R}e$. So $L/\mathbb{Z}e$ is a discrete subgroup of $V/\mathbb{R}e$. The lemma now follows easily from the induction hypothesis. \square

Returning now to our group $L = \{v \in V : \tau_v \in W\}$, the lemmas yield:

Proposition 2. *L is a lattice in V . In particular, the Euclidean reflection group W is isomorphic to a semi-direct product $\mathbb{Z}^n \rtimes \overline{W}$.* \square

Remarks

1. There is a fairly obvious way to topologize the affine group $\text{Aff}(V)$, and the proof of Lemma 1 shows that an affine reflection group W is a discrete subgroup of $\text{Aff}(V)$. Conversely, any discrete subgroup W of $\text{Aff}(V)$ generated by reflections is an affine reflection group. For one can use the discreteness assumption to prove that the set of hyperplanes H with $s_H \in W$ is locally finite. Details are left to the interested reader.

2. Proposition 2 shows that there is a non-trivial condition satisfied by the finite reflection group \overline{W} , namely, it leaves a lattice invariant. One says that \overline{W} is *crystallographic*. It turns out that this condition characterizes the groups \overline{W} among the essential irreducible finite reflection groups. More precisely, the following conditions on an essential irreducible finite reflection group W_1 are equivalent:

- (1) W_1 is the group \overline{W} associated to some Euclidean reflection group.
- (2) W_1 is crystallographic.
- (3) W_1 is the Weyl group of a root system.
- (4) W_1 is not of type H_3 , H_4 , or $I_2(m)$.

A proof can be found in Bourbaki [16], VI.2.5. The equivalence of (1) and (3) suggests that there might be a connection between Euclidean reflection groups and root systems. We will return to this at the end of the next subsection, after looking at an example.

1F Example

Let \overline{W} be the symmetric group on n letters ($n \geq 2$). Recall that \overline{W} is an essential irreducible finite reflection group acting on the $(n-1)$ -dimensional space

$$V = \{(x_1, \dots, x_n) \in \mathbf{R}^n : \sum x_i = 0\}.$$

The associated linear hyperplanes in V are defined by the equations

$$x_i - x_j = 0 \quad (i \neq j).$$

We wish to construct an affine analogue of this example by introducing a lattice of translations.

Let $L = \mathbf{Z}^n \cap V$. Then L is a \overline{W} -invariant lattice in V . It is generated, as an abelian group, by the vectors $e_i - e_j$ ($i \neq j$), where $\{e_1, \dots, e_n\}$ is the standard basis for \mathbf{R}^n . We now set

$$W = L \rtimes \overline{W} \subset V \rtimes \mathrm{GL}(V) = \mathrm{Aff}(V).$$

Note that W , as an abstract group, is the Weyl group that arose in §V.8. To see that W is a Euclidean reflection group, let \mathcal{H} be the set of affine hyperplanes in V of the form $x_i - x_j = k$ with $i \neq j$ and $k \in \mathbf{Z}$. It is easy to check that \mathcal{H} is locally finite and W -invariant. It is also easy to compute the reflection with respect to the hyperplane $x_i - x_j = k$; one finds that it is given by

$$x \mapsto s_{ij}x + k(e_i - e_j),$$

where s_{ij} is the transposition which interchanges the i th and j th coordinates. Thus W contains the reflections s_H for $H \in \mathcal{H}$.

The subgroup generated by these reflections contains \overline{W} , and hence it contains the translations $x \mapsto x + k(e_i - e_j)$. It follows that this subgroup is the whole group W , which is therefore an affine reflection group. We will compute its Coxeter diagram below and see that W is irreducible; alternatively, the irreducibility of W follows from that of \overline{W} . Since W is obviously infinite, it is a Euclidean reflection group.

As fundamental chamber C we take the subset of V defined by


$$x_1 < \dots < x_n < x_1 + 1.$$

This is an intersection of half-spaces associated to n of the elements of \mathcal{H} , and it lies on one side of every $H \in \mathcal{H}$. So it is indeed a chamber. The reflections with respect to the walls $x_i = x_{i+1}$ are the basic transpositions $s_i = s_{i,i+1}$ ($i = 1, \dots, n-1$). And the reflection s_n with respect to the wall $x_n = x_1 + 1$ is the map $x \mapsto s_{n,1}x + (e_n - e_1)$. The canonical unit vectors f_1, \dots, f_n associated to C are given by

$$f_i = \begin{cases} \frac{e_{i+1} - e_i}{\sqrt{2}} & \text{if } i \leq n-1 \\ \frac{e_1 - e_n}{\sqrt{2}} & \text{if } i = n. \end{cases}$$

Notice that they satisfy the linear relation $\sum f_i = 0$, which has positive coefficients.

One can now find the Coxeter diagram of W , either by computing the orders of the products $s_i s_j$ or by computing the inner products $\langle f_i, f_j \rangle$. The diagram is $\circ \text{---} \infty \text{---} \circ$ if $n = 2$ and



(with n vertices) if $n \geq 3$. In case $n = 3$, the diagram shows that the fundamental triangle C has all of its angles equal to $\pi/3$. Thus the Coxeter complex Σ in this case can be visualized as the Euclidean plane tiled by equilateral triangles.

We finish the discussion of this example by indicating briefly how to get an explicit isomorphism between $\Sigma(W, V)$ and the fundamental apartment for $\text{SL}_n(K)$ (cf. §V.8). We will need such an isomorphism in §9F below. For this purpose it is convenient to replace V by the canonically isomorphic vector space $\mathbf{R}^n/V^\perp = \mathbf{R}^n/\mathbf{R} \cdot (1, 1, \dots, 1)$. One can check that the set of vertices of $\Sigma(W, V)$ is the subset $\mathbf{Z}^n/\mathbf{Z} \cdot (1, \dots, 1)$ of this quotient. A further calculation yields the following description of $\Sigma(W, V)$:

Given $u = (u_1, \dots, u_n)$ and $v = (v_1, \dots, v_n)$ in \mathbf{Z}^n , write $u \preceq v$ if

$$u_i \leq v_i \leq u_i + 1$$

for all i . Call two elements of $\mathbf{Z}^n/\mathbf{Z} \cdot (1, \dots, 1)$ *incident* if they admit representatives u, v with $u \preceq v$. This incidence relation is symmetric, and one can check that the resulting flag complex is $\Sigma(W, V)$. Recall now that we gave a description of the building for $\text{SL}_n(K)$ in terms of classes of A -lattices in K^n . The fundamental apartment Σ of this building has as vertices the classes $[[\pi^{a_1} e_1, \dots, \pi^{a_n} e_n]]$, where e_1, \dots, e_n is the standard basis for K^n . In view of the definition of the incidence relation on the set of lattice classes, it is now evident that there is an isomorphism $\Sigma \rightarrow \Sigma(W, V)$ which sends the vertex $[[\pi^{a_1} e_1, \dots, \pi^{a_n} e_n]]$ to the class of $(a_1, \dots, a_n) \bmod (1, \dots, 1)$.

Remarks

1. The group W we have been discussing is called the *affine Weyl group* of the root system of type A_{n-1} . The other root systems (types B_n , C_n , etc.) also give rise to affine Weyl groups, and one obtains all Euclidean reflection groups in this way. Details can be found in Bourbaki [16], Chapter VI. Incidentally, the affine Weyl groups W of types B_n and C_n are not the same, even though they have the same associated finite reflection group \overline{W} .

2. We have already seen that the affine Weyl group of the root system of type A_{n-1} arises from a BN-pair in $\text{SL}_n(K)$ (which is the matrix group of type A_{n-1}), where K is a field with a discrete valuation. The other Euclidean reflection groups arise similarly from matrix groups over K .

3. In view of Remark 1, there are two 2-dimensional Euclidean reflection groups besides the example we have given. They correspond to the root systems B_2 and G_2 , and they have Coxeter diagrams

$$\circ \xrightarrow{4} \circ \xrightarrow{4} \circ \quad \text{and} \quad \circ \xrightarrow{6} \circ,$$

respectively. The chambers of the first group have angles $\pi/2$, $\pi/4$, and $\pi/4$, so its Coxeter complex Σ corresponds to the tiling of the plane by isosceles right triangles. The chambers of the second have angles $\pi/2$, $\pi/3$, and $\pi/6$; its complex Σ is the barycentric subdivision of the honeycomb tessellation of the plane (by regular hexagons).

2 Euclidean Coxeter Complexes

Suppose Σ is an abstract Coxeter complex which is isomorphic to the complex $\Sigma(W, V)$ associated to a Euclidean reflection group (W, V) . We then say that Σ is a *Euclidean Coxeter complex*. Choose such a (W, V) and an isomorphism $\Sigma \approx \Sigma(W, V)$. Since there is a canonical bijection between $|\Sigma(W, V)|$ and V (cf. §1D), we obtain a bijection $|\Sigma| \approx V$. We wish to use this bijection to transport to $|\Sigma|$ the notions of Euclidean geometry. The lemma below will enable us to show that these notions are independent of the choice of (W, V) and the choice of isomorphism $\Sigma \approx \Sigma(W, V)$.

The intuitive content of the lemma is that one can reconstruct the Euclidean space V (up to a dilation of its metric) from the abstract simplicial complex $\Sigma(W, V)$. Here is the precise statement:

Lemma. *Let (W, V) and (W', V') be Euclidean reflection groups. Let $\phi : \Sigma(W, V) \rightarrow \Sigma(W', V')$ be a simplicial isomorphism. Then the composite bijection*

$$V \approx |\Sigma(W, V)| \xrightarrow{|\phi|} |\Sigma(W', V')| \approx V'$$

is a similarity map, i.e., an affine isomorphism whose linear part g satisfies $\langle gv, gv' \rangle = \lambda \langle v, v' \rangle$ for some positive constant λ .

PROOF: Choose a chamber C in V , with faces A_i ($i = 1, \dots, n+1$). Let H_i be the support of A_i , let e_i be the canonical unit normal to H_i (pointing toward C), and let s_i be the reflection with respect to H_i . As in the proof in §1D above, we may conjugate W by a translation in order to arrange that H_i is defined by $\langle e_i, - \rangle = 0$ if $i \leq n$ and by $\langle e_{n+1}, - \rangle = c$ if $i = n+1$, where $c < 0$. Conjugating W by a dilation, we can further arrange that $c = -1$. Let $C' = \phi(C)$ and $A'_i = \phi(A_i)$. Let H'_i , e'_i , and s'_i be the associated wall, unit vector, and reflection, respectively. We may assume that H'_i is defined by $\langle e'_i, - \rangle = 0$ if $i \leq n$ and by $\langle e'_{n+1}, - \rangle = -1$ if $i = n+1$.

Recall now that the Coxeter matrix $M = (m_{ij})_{1 \leq i, j \leq n+1}$ of W is a combinatorial invariant of the Coxeter complex (cf. §III.2); namely, m_{ij} is

the diameter of the link of $A_i \cap A_j$. In view of the isomorphism ϕ , it follows that M is also the Coxeter matrix of W' . Consequently, we have

$$\langle e_i, e_j \rangle = \langle e'_i, e'_j \rangle$$

for all i, j . We can now construct a linear isometry $\alpha : V \rightarrow V'$ such that $\alpha(e_i) = e'_i$ for all i . For if we take α to be the linear map such that $\alpha(e_i) = e'_i$ for $i \leq n$, then α preserves inner products and satisfies $\langle \alpha(e_{n+1}), - \rangle = \langle e'_{n+1}, - \rangle$, whence $\alpha(e_{n+1}) = e'_{n+1}$.

Note that $\alpha(H_i) = H'_i$, so that $\alpha s_i \alpha^{-1} = s'_i$ and $\alpha W \alpha^{-1} = W'$. Thus α induces an isomorphism of pairs $(W, V) \rightarrow (W', V')$, and hence a simplicial isomorphism $\Sigma(W, V) \rightarrow \Sigma(W', V')$. This isomorphism takes C to C' and A_i to A'_i , so it coincides with our original isomorphism ϕ by the standard uniqueness argument. The lemma now follows from the commutative diagram

$$\begin{array}{ccc} |\Sigma(W, V)| & \longrightarrow & |\Sigma(W', V')| \\ \downarrow & & \downarrow \\ V & \longrightarrow & V' \end{array}$$

where the vertical arrows denote the canonical bijections and the horizontal arrows are induced by α . \square

We now have a “Euclidean structure” on $|\Sigma|$ for any Euclidean Coxeter complex Σ , by which we mean that we can apply to $|\Sigma|$ any notion of Euclidean geometry which is invariant under similarity maps. In particular, $|\Sigma|$ has a well-defined equivalence class of metrics, where two metrics are equivalent if one is a positive scalar multiple of the other. It will be convenient in what follows to have a canonical representative of this equivalence class. We can achieve this in many ways; for definiteness, let’s agree to choose the representative which makes the chambers have diameter 1. Thus $|\Sigma|$ is now a metric space, and any abstract isomorphism $\phi : \Sigma \rightarrow \Sigma'$ of Euclidean Coxeter complexes induces an isometry $|\Sigma| \rightarrow |\Sigma'|$.

We close this section by introducing one last bit of terminology: By a *Euclidean space* we will mean a metric space E which is isometric to \mathbf{R}^n for some n . What we have done above, then, is to give $|\Sigma|$ a canonical Euclidean space structure.

Note that we are making a somewhat pedantic distinction between the notions of “Euclidean space” and “Euclidean vector space”; recall that we have defined the latter to mean “vector space with an inner product”. As a practical matter, the only difference between the two notions is that a Euclidean vector space comes equipped with a preferred origin. More precisely, suppose E is a Euclidean space and x_0 is an arbitrary point in E . Then we can give E the structure of Euclidean vector space with x_0 as origin by choosing an isometry $\alpha : \mathbf{R}^n \rightarrow E$ with $\alpha(0) = x_0$ and using α to transport the vector space structure and inner product from \mathbf{R}^n to E . It follows from Exercise 3 of §1A that this structure is independent of the choice of α .

EXERCISE

Convince yourself that the following assertion is meaningful and true: If Σ is a Euclidean Coxeter complex, then the Euclidean space $E = |\Sigma|$ contains a canonical locally finite collection \mathcal{H} of hyperplanes, from which one can recover the decomposition of E into simplices.

3 Euclidean Buildings as Metric Spaces

This section, as well as most of the rest of the chapter, is based on the paper of Bruhat and Tits [22].

Let Δ be a building, equipped with an arbitrary system of apartments \mathcal{A} . Nothing we do in this section will depend on the choice of \mathcal{A} . We will say that Δ is *Euclidean* if its apartments are Euclidean Coxeter complexes. One also says that Δ is an *affine* building, or a building of *affine type*. In view of the previous section, we then have a Euclidean structure on every apartment Σ . The purpose of the present section is to see what sort of geometry these Euclidean structures impose on the building as a whole.

Assume throughout this section that Δ is a Euclidean building and that \mathcal{A} is a fixed system of apartments.

3A Construction of a metric on $|\Delta|$

Let X be the geometric realization $|\Delta|$ of the Euclidean building Δ . For the moment, X is just a set, with no topology. It is the union of open simplices $|A|$, one for each $A \in \Delta$ (cf. Appendix to Chapter I). To avoid cumbersome notation, we will omit the vertical bars and simply denote by A this open simplex and by \bar{A} the corresponding closed simplex. Thus \bar{A} now denotes the geometric realization of the subcomplex $\Delta_{\leq A}$ that we have sometimes called \bar{A} in earlier chapters.

It will be convenient to apply to X terminology that we have previously used for the abstract complex Δ . In particular, we will refer to X itself as a building and to the subsets $E = |\Sigma|$ as apartments ($\Sigma \in \mathcal{A}$). For any such apartment E and any chamber C of E , the geometric realization of $\rho_{\Sigma, C} : \Delta \rightarrow \Sigma$ is a retraction $X \rightarrow E$, denoted $\rho_{E, C}$.

In view of §2, each apartment E of X is a Euclidean space, with a metric d_E . Moreover, the isomorphisms between apartments given by the building axiom (B2) are isometries. We now wish to piece the metrics d_E together to make the entire building X a metric space.

Given two points $x, y \in X$, axiom (B1) implies that there is an apartment E containing both x and y . Choose such an E and set

$$d(x, y) = d_E(x, y).$$

If E' is another apartment containing x and y , then (B2) gives us an isometry $E \rightarrow E'$ fixing x and y , so $d(x, y)$ is independent of the choice of

apartment. We therefore have a well-defined function

$$d : X \times X \rightarrow \mathbf{R}.$$

This should not be confused with the combinatorial distance function that we have used before (defined via galleries), although, as we will see, the two kinds of distance functions have some similar properties. To avoid confusion, we will write $d(-, -)$ from now on for the combinatorial distance function.

Theorem.

- (1) *The function $d : X \times X \rightarrow \mathbf{R}$ is a metric.*
- (2) *The metric space X is complete.*
- (3) *The retraction $\rho = \rho_{E,C} : X \rightarrow E$ is distance-decreasing for any apartment E and chamber C of E , i.e.,*

$$d(\rho(x), \rho(y)) \leq d(x, y)$$

for all $x, y \in X$. Equality holds if $x \in \overline{C}$.

- (4) *For any $x, y \in X$, choose an apartment E containing x and y and let $[x, y]$ be the line segment joining them in the Euclidean space E . Then $[x, y]$ is independent of the choice of E and can be characterized by*

$$[x, y] = \{ z \in X : d(x, y) = d(x, z) + d(z, y) \}.$$

- (5) *Given $x, y \in X$ and $t \in [0, 1]$, let $(1 - t)x + ty$ denote the point z in $[x, y]$ such that $d(x, z) = td(x, y)$. Then the function $(x, y, t) \mapsto (1 - t)x + ty$ is a continuous map $X \times X \times [0, 1] \rightarrow X$, where X is topologized by means of the metric d . In particular, the metric space X is contractible.*

PROOF: We begin by proving (3), which makes sense even before we know that d is a metric. The second assertion of (3) follows immediately from the fact that ρ maps every apartment containing C isometrically onto E . This fact also implies that, for any chamber C' , ρ maps $\overline{C'}$ isometrically onto its image. Suppose now that x and y are arbitrary points of X . Choose an apartment E' containing them, and let $[x, y]$ be the line segment joining them in E' . It is easy to see that we can subdivide this segment in such a way that each subinterval is contained in a closed chamber. [Use the fact that the decomposition of E' into simplices is induced by a locally finite collection \mathcal{H} of hyperplanes.] Let the subdivision points be $x = x_0, x_1, \dots, x_m = y$. We then have

$$d(\rho(x), \rho(y)) \leq \sum_{i=1}^m d(\rho(x_{i-1}), \rho(x_i)) = \sum_{i=1}^m d(x_{i-1}, x_i) = d(x, y),$$

where the inequality follows from the triangle inequality in the Euclidean space E , and the first equality follows from the fact that ρ is an isometry on each closed chamber. This proves the first assertion of (3).

It is now easy to prove (1), the content of which is that d satisfies the triangle inequality: Given $x, y, z \in X$, choose an apartment E containing x and y , and let $\rho = \rho_{E,C}$ for some chamber C of E . Using (3) and the triangle inequality in E , we find

$$d(x, y) \leq d(x, \rho(z)) + d(\rho(z), y) \leq d(x, z) + d(z, y),$$

as required.

Continuing with the same notation, suppose $d(x, y) = d(x, z) + d(z, y)$. Then both inequalities above must be equalities. From the first equality and elementary Euclidean geometry we conclude that $\rho(z) \in [x, y]$, the latter being the line segment joining x and y in E . And from the second equality and (3) we conclude that

$$d(x, \rho(z)) = d(x, z) \quad \text{and} \quad d(\rho(z), y) = d(z, y).$$

Consequently, $\rho(z)$ must be the point $p_t = (1 - t)x + ty$ of $[x, y]$, where $t = d(x, z)/d(x, y)$. Recall that ρ here is $\rho_{E,C}$ for *any* chamber C of E . In particular, we can take C to be a chamber with $p_t \in \overline{C}$; this is legitimate, since the definition of p_t is independent of the choice of C . Then the equation $\rho(z) = p_t$ implies, in view of the second assertion of (3), that $z = p_t$. This proves the characterization of $[x, y]$ stated in (4), and the rest of (4) follows at once.

The proof of (5) will be based on the following formula from Euclidean geometry: Given points x, y in a Euclidean space E , and given $t \in [0, 1]$, let $p_t = (1 - t)x + ty \in [x, y]$; then for any $z \in E$,

$$d^2(z, p_t) = (1 - t)d^2(z, x) + td^2(z, y) - t(1 - t)d^2(x, y). \quad (*)$$

(You have probably seen this, or a slight variant of it, in case $t = 1/2$; the result in this case is essentially the parallelogram law.) To prove (*) we may assume that E is a Euclidean vector space with z as the origin. Then

$$d^2(z, p_t) = \|p_t\|^2 = (1 - t)^2\|x\|^2 + t^2\|y\|^2 + 2t(1 - t)\langle x, y \rangle;$$

one obtains (*) from this by using the formula

$$d^2(x, y) = \|x\|^2 + \|y\|^2 - 2\langle x, y \rangle$$

to eliminate $\langle x, y \rangle$.

Now suppose x, y, z are points of our building X , and define $p_t \in [x, y]$ as above. Choose an apartment E containing x and y (and hence $[x, y]$), and choose a closed chamber \overline{C} in this apartment containing p_t . Let $\rho = \rho_{E,C}$. Applying (*) to the points $x, y, \rho(z) \in E$, we obtain

$$d^2(\rho(z), p_t) = (1 - t)d^2(\rho(z), x) + td^2(\rho(z), y) - t(1 - t)d^2(x, y).$$

In view of (3), it follows that

$$d^2(z, p_t) \leq (1 - t)d^2(z, x) + td^2(z, y) - t(1 - t)d^2(x, y). \quad (**)$$

To prove (5), we will apply this inequality to $z = (1 - t')x' + t'y'$ for (x', y', t') close to (x, y, t) . Since $d(z, x)$ is close to $d(z, x') = t'd(x', y')$, it is

clear that $d(z, x) \rightarrow td(x, y)$ as $(x', y', t') \rightarrow (x, y, t)$; hence the first term on the right side of (**) approaches $(1-t)t^2d^2(x, y)$. Similarly, the second term approaches $t(1-t)^2d^2(x, y)$. The right side of (**) therefore approaches

$$(1-t)t^2d^2(x, y) + t(1-t)^2d^2(x, y) - t(1-t)d^2(x, y) = 0,$$

whence $d(z, p_t) \rightarrow 0$ and (5) is proved.

Finally, we must prove (2). Fix a chamber C and let $\lambda : X \rightarrow \overline{C}$ be the geometric realization of the unique retraction $\Delta \rightarrow \Delta_{\leq C}$ (which exists by the labellability of Δ). Then λ maps any closed chamber $\overline{C'}$ of X isometrically onto \overline{C} . To see this, choose an apartment $E = |\Sigma|$ containing C and C' and let w be the unique type-preserving automorphism of Σ such that $wC' = C$. Then $\lambda|_{\overline{C'}}$ is induced by the restriction of w to C' and its faces, and the assertion now follows from the fact that w is an isometry of E .

It now follows easily that λ is distance-decreasing; the proof is the same as the proof of the analogous fact about retractions onto apartments. So if we are given a Cauchy sequence $(x_m)_{m \geq 1}$ in X , then the image sequence $(\lambda(x_m))$ is a Cauchy sequence in \overline{C} . The latter being a closed subset of a Euclidean space, it follows that there is a point $y \in \overline{C}$ such that $\lambda(x_m) \rightarrow y$ as $m \rightarrow \infty$. Choose for each m a chamber C_m with $x_m \in \overline{C_m}$, and let y_m be the unique point in $\overline{C_m}$ such that $\lambda(y_m) = y$. (We will then say that y_m is of *type* y .) Since $\lambda|_{\overline{C_m}}$ is an isometry, we have

$$d(x_m, y_m) = d(\lambda(x_m), y) \rightarrow 0 \quad \text{as } m \rightarrow \infty,$$

hence (y_m) is also a Cauchy sequence. On the other hand, we will prove below that the set of points of a given type y is discrete. [Draw a picture of the tree case to see why this is intuitively plausible.] Hence the Cauchy sequence (y_m) is eventually constant, and the fact that $d(x_m, y_m)$ tends to 0 now says that $x_m \rightarrow y'$, where $y' = y_m$ for large m . This completes the proof of the theorem, except for the discreteness assertion.

Recall that the *star* of a point $x \in X$, denoted $\text{st } x$ or $\text{st}_X x$, is the union of the closed simplices \overline{A} containing x . If our metric on X is at all reasonable, we expect the star of x to be a neighborhood of x . In fact, a more precise statement is true:

Lemma. *Given $y \in \overline{C}$ there is a $\delta > 0$ with the following property: For any $x \in X$ of type y , $\text{st } x$ contains the closed ball of radius δ centered at x .*

Now $\text{st } x$ contains no point distinct from x and having the same type as x . So the lemma implies that $d(x, x') > \delta$ for any two distinct points x, x' of type y . This proves the discreteness assertion, modulo the lemma.

PROOF OF THE LEMMA: Choose an apartment E containing C , and let \mathcal{H} be the locally finite collection of walls that defines the simplicial decomposition of E . Let δ be the minimum distance from y to a wall $H \in \mathcal{H}$ not containing y . Then for any $y' \in E$ with $d(y, y') \leq \delta$, the open segment

(y, y') does not cross any wall. We therefore have $(y, y') \subseteq A$ for some open cell A , hence $y, y' \in \bar{A}$ and $y' \in \text{st}_E y$.

Now suppose x and x' are points of X with x of type y and $d(x, x') \leq \delta$. We can find an apartment E' containing x and x' and an isomorphism $\phi : E' \rightarrow E$ such that $\phi(x) = y$. Then $d(y, \phi(x')) \leq \delta$, so $\phi(x') \in \text{st}_E y$ by the previous paragraph and hence $x' \in \text{st}_{E'} x \subseteq \text{st}_X x$. \square

EXERCISES


1. The *open star* of x is the union of the open simplices A such that $x \in \bar{A}$. State and prove an analogue of the lemma for open stars.

2. Let X' be a subcomplex of X , i.e., $X' = |\Delta'|$ for some subcomplex Δ' of Δ . [Equivalently, X' is a subset of X which is a union of closed simplices.] Deduce from Exercise 1 that X' is a closed subset of X .

3. Deduce from Exercise 1 or 2 that any chamber C is an open subset of X .

3B Negative curvature

The segment $[x, y]$ defined in the theorem will be called the *geodesic* joining x and y . Comparing the inequality (**) with the equality (*) above, we see that a typical point $z \in X$ is closer to points of $[x, y]$ than it “ought” to be. This suggests the picture



which will be familiar to readers who have studied Riemannian manifolds of negative curvature (e.g., the hyperbolic plane). For this reason we will call (**) the *negative curvature inequality*. [Note: One should interpret “negative” here as meaning “non-positive”, since equality might hold; e.g., X might consist of a single apartment and hence be a Euclidean space.] If you know enough Riemannian geometry, you might enjoy the following exercise:

EXERCISE

Show that the negative curvature inequality holds in a complete simply-connected manifold M of negative curvature. [HINT: Let V be the tangent space to M at p_t and let $\rho : M \rightarrow V$ be the inverse of the exponential map. Then ρ is distance-decreasing and preserves distances from p_t .]

In what follows we will only need the special case $t = 1/2$ of the negative curvature inequality; letting m be the midpoint $p_{1/2}$ of $[x, y]$, we can write this special case as

$$d^2(z, m) \leq \frac{1}{2} (d^2(z, x) + d^2(z, y)) - \frac{1}{4} d^2(x, y). \quad (\text{NC})$$

For brevity we will say that a metric space X has property (NC) if for any two points $x, y \in X$ there is a point m such that the inequality (NC) holds for all z . Thus Euclidean buildings and complete simply-connected Riemannian manifolds of negative curvature are examples of spaces with property (NC).

Note that any point m as in the definition of property (NC) necessarily satisfies

$$d(x, m) = d(y, m) = \frac{1}{2}d(x, y);$$

this follows from (NC) applied to $z = x$ and $z = y$. Moreover, m is the only point satisfying these equations; for if m' satisfies the equations, then (NC) applied with $z = m'$ implies that $d(m', m) = 0$. We will call m the *midpoint* of the pair $\{x, y\}$.

A final comment about the intuition behind negative curvature: One often interprets negative curvature as meaning that two geodesics emanating from a given point separate faster than they ought to. The first exercise below gives a precise version of this statement.

EXERCISES

Assume throughout these exercises that X is a metric space with property (NC).

1. Given three points $a, b, c \in X$, let m_1 be the midpoint of $\{a, b\}$ and let m_2 be the midpoint of $\{a, c\}$. Show that $d(b, c) \geq 2d(m_1, m_2)$. [HINT: Start by using (NC) to say something about $d^2(m_1, m_2)$.]

2. Fix $x, y \in X$ and let $t \in [0, 1]$ be a dyadic rational number, i.e., a rational number whose denominator is a power of 2. Show that there is a point $p_t \in X$ such that the inequality (**) of §3A holds for all $z \in X$. Show further that any such p_t can be characterized as the unique point satisfying the equations

$$\begin{aligned} d(x, p_t) &= td(x, y) \\ d(y, p_t) &= (1 - t)d(x, y). \end{aligned}$$

Extend all this to arbitrary $t \in [0, 1]$ if X is complete.

3. A subset Y of X will be called *midpoint convex* if it contains the midpoint of any pair of its points. Suppose that Y is midpoint convex, let $x \in X$ be arbitrary, and let $d = d(x, Y) = \inf_{y \in Y} d(x, y)$. If there is a point $y \in Y$ such that $d(x, y) = d$, show that $d(y, y') \leq d(x, y')$ for all $y' \in Y$. [HINT: First draw a picture to see why the assertion is plausible. For the proof, consider the points p_t “between” y and y' as in Exercise 2, where t ranges over the dyadic rationals in $[0, 1]$. Then

$$d^2(x, p_t) \leq (1 - t)d^2 + td^2(x, y') - t(1 - t)d^2(y, y').$$

The right-hand side of this inequality would be a decreasing function of t for small t if $d(y, y') > d(x, y')$.

4 The Bruhat–Tits Fixed-Point Theorem

If G is a compact group of isometries of a complete simply-connected Riemannian manifold M of negative curvature, then a famous theorem of E. Cartan says that G fixes a point of M . Cartan’s fixed-point theorem is a fundamental tool in the theory of Lie groups. In this section we will prove a generalization of Cartan’s theorem to complete metric spaces with property (NC). This generalization, due to Bruhat and Tits [22], then applies to Euclidean buildings as well as to complete simply-connected manifolds of negative curvature. Like Cartan’s theorem, the Bruhat–Tits theorem has applications to group theory; we will explore these in the next section.

Theorem 1. *Let G be a group of isometries of a complete metric space X with property (NC). If G stabilizes a non-empty bounded subset of X , then G has a fixed point.*

Remarks

1. In the situation of Cartan’s theorem, where G is compact, any orbit Gx is a bounded set stabilized by G . So Cartan’s theorem is indeed a special case of the Bruhat–Tits theorem.

2. The geometric realization of a tree is a metric space with property (NC). The theorem in this case can be found in Serre’s book on trees ([46], §I.4.3, Proposition 19), which also contains applications to group theory.

3. Theorem 1 is of interest even when X is a Euclidean space or, more generally, a Hilbert space (real or complex, possibly infinite dimensional). The result in this case has a cohomological interpretation, which is often used in representation theory; see Exercise 2 below.

The Bruhat–Tits proof of Theorem 1 consists of associating to every non-empty bounded subset $A \subseteq X$ a point $c = c(A) \in X$ which, intuitively, is some sort of “center” of A . The construction of c depends only on the metric on X , so it is compatible with isometries. In particular, if A is invariant under a group G of isometries of X , then c is fixed by G .

We will give a variant of this proof due to Serre. The basic idea remains the same, but Serre’s definition of $c(A)$ is different from that of Bruhat and Tits. Namely, $c(A)$ is defined to be the center of the sphere circumscribed about A . Here are the details.

Let X be an arbitrary metric space and A a non-empty bounded subset. For any $x \in X$, let $r(x, A)$ be the smallest real number r such that A is contained in the closed ball of radius r centered at x ; equivalently,

$$r(x, A) = \sup_{a \in A} d(x, a).$$

The *circumradius* of A , denoted $r(A)$, is defined by

$$r(A) = \inf_{x \in X} r(x, A).$$

If $r(A) = r(x, A)$ for some $x \in X$, then any such x will be called a *circumcenter* of A . Note, for example, that the midpoint m discussed in §3B above is a circumcenter (and even the unique circumcenter) of the two-point set $\{x, y\}$.

If X is a sphere, which is a space of positive curvature, then circumcenters always exist but are not necessarily unique (see Exercise 1 below). In Euclidean space, however, it is known that circumcenters exist and are unique. More generally, we have the following observation of Serre:

Theorem 2. *If X is a complete metric space with property (NC), then every non-empty bounded subset A admits one and only one circumcenter.*

As we explained above, Theorem 1 follows immediately from Theorem 2; for the circumcenter of A will clearly be fixed by any group of isometries of X that stabilizes A . It remains to prove Theorem 2.

PROOF OF THEOREM 2: For any two points $x, y \in X$, we can apply the inequality (NC) with $z \in A$ to get

$$r^2(m, A) \leq \frac{1}{2} (r^2(x, A) + r^2(y, A)) - \frac{1}{4} d^2(x, y),$$

where m is the midpoint of $\{x, y\}$. Hence

$$d^2(x, y) \leq 2 (r^2(x, A) + r^2(y, A) - 2r^2(m, A)).$$

Since $r(m, A) \geq r(A)$, this implies

$$d^2(x, y) \leq 2 (r^2(x, A) + r^2(y, A) - 2r^2(A)).$$

Uniqueness of the circumcenter is now immediate; for if x and y are both circumcenters, then the right-hand side is 0, hence $x = y$. To prove existence, take a sequence of points $x_n \in X$ such that $r(x_n, A) \rightarrow r(A)$, and apply the inequality above with $x = x_n$ and $y = x_m$. Then the right-hand side can be made arbitrarily small by taking n and m sufficiently large, so (x_n) is a Cauchy sequence. Hence (x_n) has a limit $x \in X$, and it is easy to check that $r(x, A) = r(A)$. \square

We close this section by proving one more result about circumcenters. This will not be needed in what follows, but it provides a nice illustration of the negative curvature inequality. In Euclidean geometry, it is known that the circumcenter of a bounded set A is contained in the closure of the convex hull of A . We will show that this too generalizes to spaces with property (NC). The precise statement uses the notion of “midpoint convexity” defined in Exercise 3 of §3B.

Theorem 3. *Under the hypotheses of Theorem 2, let Y be the smallest closed, midpoint convex subset of X which contains A . Then the circumcenter of A is contained in Y .*

At first glance it might seem that this is a formal consequence of Theorem 2. For Y , in its own right, is a complete metric space with property

(NC), so Theorem 2 implies that A has a circumcenter in Y . What is not obvious, however, is that the circumcenter of A in Y is the same as its circumcenter in X . In other words, it is conceivable that the circumradius $r_Y(A)$ of A in Y is bigger than the circumradius $r_X(A)$ of A in X . Theorem 3 will follow as soon as we prove that this cannot happen:

Lemma. *Let X be a metric space with property (NC), and let Y be a midpoint convex subset. Then $r_Y(A) = r_X(A)$ for any non-empty bounded subset $A \subseteq Y$.*

PROOF: Given any $x \in X$ and any $r > r(x, A)$, we must find $y \in Y$ such that $r(y, A) \leq r$. This is easy if there is a $y \in Y$ with $d(x, y) = d(x, Y)$; in this case we have $r(y, A) \leq r(x, A) < r$ by Exercise 3 of §3B. In the general case, let $d = d(x, Y)$, and choose a sequence of points $y_n \in Y$ such that $d_n = d(x, y_n) \rightarrow d$. We will show that $r(y_n, A) \leq r$ for some n . Suppose this is false. Then for each n we can find a point $a_n \in A$ such that $d(y_n, a_n) > r$. As in the exercise just cited, consider the points p_t between y_n and a_n , where t ranges over the dyadic rationals in $[0, 1]$. These points are in Y , and we have

$$d^2 \leq d^2(x, p_t) \leq (1-t)d_n^2 + tr^2(x, A) - t(1-t)r^2.$$

Fixing t and letting $n \rightarrow \infty$, we conclude that

$$d^2 \leq (1-t)d^2 + tr^2(x, A) - t(1-t)r^2 = d^2 + \alpha t + r^2 t^2,$$

where $\alpha = -d^2 + r^2(x, A) - r^2$. But this is absurd; for α is negative, so $\alpha t + r^2 t^2 < 0$ for small $t > 0$. This contradiction shows that $r(y_n, A) \leq r$ for some n , as required. \square

EXERCISES

1. (a) If X is a compact metric space, show that every non-empty subset admits a circumcenter.

(b) If X is a sphere (of any dimension ≥ 0), show that X has subsets with more than one circumcenter. In fact, there is even a subset such that *every* point of X is a circumcenter. (More generally, this happens whenever X is a metric space of finite diameter which admits a transitive group of isometries.)

2. Let V be a real or complex Hilbert space on which a group G acts by linear isometries. A *1-cocycle* on G with values in V is a function $c : G \rightarrow V$ such that $c(gh) = c(g) + gc(h)$ for all $g, h \in G$. It is called a *coboundary* if there is a vector $v \in V$ such that $c(g) = gv - v$ for all $g \in G$. Deduce from Theorem 1 that a cocycle is a coboundary if and only if it is bounded. [HINT: Use the given cocycle to define an action of G on V by affine isometries.]

5 Application: Bounded Subgroups

There is a classical application of Cartan's fixed-point theorem to the study of compact subgroups of a Lie group G : Under suitable hypotheses on G ,

one constructs a complete simply-connected Riemannian manifold X of negative curvature, on which G acts as a group of isometries; the fixed-point theorem then implies that any compact subgroup of G must be contained in the stabilizer G_x of some point $x \in X$. If $G = \mathrm{SL}_n(\mathbf{R})$, for instance, then G acts transitively on the associated X and has the special orthogonal group $\mathrm{SO}_n(\mathbf{R})$ as one of the stabilizers. The conclusion, then, is that every compact subgroup of $\mathrm{SL}_n(\mathbf{R})$ is conjugate to a subgroup of $\mathrm{SO}_n(\mathbf{R})$. [Note: The group SO_n here is defined with respect to the standard inner product on \mathbf{R}^n , not the one used in §V.7.] In this section we will use the Bruhat–Tits fixed-point theorem to prove similar results for groups acting on Euclidean buildings. These results then apply to certain “ p -adic Lie groups” such as $\mathrm{SL}_n(\mathbf{Q}_p)$.

Let G be a group with a BN-pair and let Δ be the associated building. We will say that the BN-pair is *Euclidean* if Δ is Euclidean. It is then immediate from the definition of the metric on $X = |\Delta|$ that G acts as a group of isometries of X . In many cases G has a natural topology, so that the notion of compact subgroup makes sense. In general, however, it is more convenient to deal with “bounded” subgroups. We begin by figuring out what that should mean.

Lemma. *The following conditions on a subset $F \subseteq G$ are equivalent:*

- (1) F is contained in a finite union of double cosets BwB .
- (2) For some $x \in X$, the set $Fx = \{gx : g \in F\}$ is a bounded subset of the metric space X .
- (3) For every bounded set $Y \subseteq X$, the set $FY = \bigcup_{y \in Y} Fy$ is a bounded subset of X .

PROOF: (1) \implies (2): It suffices to consider the case where F is a double coset BwB . Let C be the fundamental chamber of X ; it is fixed pointwise by B . Let \tilde{w} be a representative of w in N . Then for any $g = b\tilde{w}b' \in F$ and any $x \in \overline{C}$, we have

$$d(x, gx) = d(bx, gx) = d(x, \tilde{w}b'x) = d(x, wx).$$

Hence Fx is contained in the sphere of radius $r = d(x, wx)$ centered at x .

(2) \implies (3): This is left as an exercise; it is valid for any set of isometries of any metric space.

(3) \implies (1): By (3) applied with Y equal to the fundamental chamber C , the set FC is bounded. Let E be the fundamental apartment, and let $\rho = \rho_{E,C} : X \rightarrow E$. Since ρ is distance-decreasing, $\rho(FC)$ is a bounded subset of E . In view of the local finiteness of the set \mathcal{H} of hyperplanes defining the simplicial decomposition of E , a bounded subset meets only finitely many chambers; hence $\rho(FC)$ is a finite union of chambers. The geometric interpretation of the Bruhat decomposition (cf. §V.1E) now implies that F is contained in a finite union of double cosets. \square

We will call F *bounded* if it satisfies the equivalent conditions of the

lemma. The following exercises should convince you that this definition is reasonable.

EXERCISES

1. Suppose $G = \mathrm{SL}_n(K)$ as in §V.8. Show that a set F is bounded if and only if there is an upper bound on the absolute values of the matrix entries of the elements of F . If K is complete and the residue field k is finite (e.g., $K = \mathbf{Q}_p$), show further that F is bounded if and only if it is relatively compact. Here G is topologized as a subset of the vector space of $n \times n$ matrices, and a set is called *relatively compact* if its closure is compact.

2. Show that our notion of “bounded set” satisfies the following conditions, which are the axioms for a *bornology* on a set:

- (i) Every singleton is bounded.
- (ii) If $F' \subset F$ and F is bounded, then F' is bounded.
- (iii) A finite union of bounded sets is bounded.

Show further that the following two axioms for a *bornological group* are satisfied:

- (iv) If F_1 and F_2 are bounded, then so is their product $F_1 F_2$.
- (v) If F is bounded, then so is F^{-1} .

We are now ready to apply the fixed-point theorem. Note that the stabilizers of the points of X are the same as the stabilizers of the non-empty simplices; hence they are the proper parabolic subgroups. In particular, the maximal elements among these stabilizers are the maximal (proper) parabolic subgroups, which are the stabilizers of the vertices. We will omit the word “proper” in what follows, since the notion of “maximal parabolic subgroup” would be of no interest otherwise.

Theorem. *The following conditions on a subgroup $H \subseteq G$ are equivalent:*

- (1) H is bounded.
- (2) H fixes a point of X .
- (3) H fixes a vertex of X .
- (4) H is contained in a maximal parabolic subgroup.

PROOF: It is immediate that $(4) \iff (3) \iff (2) \implies (1)$. The content of the theorem, then, is that $(1) \implies (2)$, and this follows from the fixed-point theorem. \square

Corollary. *Every bounded subgroup is contained in a maximal bounded subgroup, and the maximal bounded subgroups are the maximal parabolic subgroups. G contains precisely $n+1$ conjugacy classes of maximal bounded subgroups, where $n = \dim X$; they are represented by the special subgroups $BW'B$ with $W' = \langle S - \{s\} \rangle$ for some $s \in S$.* \square

Remark. Suppose we are in the situation where G is a topological group and “bounded” is the same as “relatively compact” (e.g., $G = \mathrm{SL}_n(\mathbf{Q}_p)$). Then a maximal bounded subgroup is necessarily compact, since otherwise its closure would be a bigger bounded subgroup. Consequently, the corollary remains valid with “bounded” replaced by “compact”.

EXERCISE

Describe the n conjugacy classes of maximal bounded subgroups of $SL_n(K)$. [HINT: Recall that the building can be described in terms of lattices. What is the stabilizer of a vertex?]

We close this section by proving that the building Δ can be entirely reconstructed from the group G , viewed simply as a bornological group. In particular, in the situation of the remark above, Δ can be reconstructed from G as a topological group. The precise statement will be given below, after a sequence of lemmas.

Lemma 1. *Every apartment Σ is a flag complex.*

PROOF: The proof given in §B of the Appendix to Chapter I for spherical Coxeter complexes goes through verbatim in the Euclidean case. [In fact, essentially the same proof works for *any* Coxeter complex, cf. Tits [56], Corollary 2.28.] \square

Lemma 2. *Every apartment Σ is a full subcomplex of Δ . In other words, if a simplex $A \in \Delta$ has all of its vertices in Σ , then $A \in \Sigma$.*

PROOF: Let $\rho : \Delta \rightarrow \Sigma$ be a retraction. Then ρ fixes all the vertices of A , so $A = \rho(A) \in \Sigma$. \square

Lemma 3. *Δ is a flag complex.*

PROOF: We must show that if v_1, \dots, v_k are pairwise joinable vertices, then $\{v_1, \dots, v_k\}$ is a simplex of Δ . Arguing by induction on k , we may assume that $\{v_1, \dots, v_{k-1}\}$ is a simplex. Choose an apartment Σ containing $\{v_1, \dots, v_{k-1}\}$ and v_k . Then v_1, \dots, v_k are pairwise joinable vertices of Σ by Lemma 2, applied to each of the 1-simplices $\{v_i, v_j\}$. Lemma 1 now implies that $\{v_1, \dots, v_k\}$ is a simplex of Σ and hence also of Δ . \square

Recall that Δ can be identified with the poset of parabolic subgroups of G , ordered by the opposite of the inclusion relation. Lemma 3 therefore implies that Δ is the flag complex of the incidence geometry consisting of the maximal parabolic subgroups, with two maximal parabolics P, Q incident if and only if $P \cap Q$ contains a parabolic subgroup. [This says precisely that the corresponding vertices of Δ are joinable.] Since any subgroup of G containing a parabolic subgroup is itself parabolic, we can state this more simply: P and Q are incident if and only if $P \cap Q$ is parabolic.

Lemma 4. *If P and Q are distinct maximal parabolics, then $P \cap Q$ is parabolic if and only if $P \cap Q$ is a maximal (proper) subgroup of P .*

PROOF: Let x (resp. y) be the vertex fixed by P (resp. Q). If $P \cap Q$ is parabolic, then x and y are joinable and $P \cap Q$ is the stabilizer of the edge A that they determine. Any subgroup P' with $P > P' > P \cap Q$ would be parabolic and would therefore correspond to a simplex A' with $x < A' < A$. So no such P' can exist, i.e., $P \cap Q$ is maximal in P .

Conversely, suppose that $P \cap Q$ is a maximal subgroup of P , and consider the geodesic $[x, y]$. It is fixed pointwise by $P \cap Q$, since the latter is a group of isometries fixing x and y . For any $x' \in (x, y]$ sufficiently close to x , the segment $(x, x']$ is contained in some simplex A of positive dimension having x as a vertex; hence the stabilizer P' of x' (which is the same as the stabilizer of A) is properly contained in P . We therefore have $P > P' \geq P \cap Q$, which implies that $P \cap Q$ is equal to the parabolic subgroup P' . \square

We have now proved:

Theorem. *The building Δ associated to a group G with a Euclidean BN-pair is isomorphic to the flag complex of the incidence geometry consisting of the maximal bounded subgroups of G , where two distinct such subgroups P, Q are incident if and only if $P \cap Q$ is a maximal subgroup of P .* \square

There is an analogous theorem about Lie groups. Under suitable hypotheses on a Lie group G , the associated manifold X of negative curvature can be identified with the set of maximal compact subgroups of G , and the structure of Riemannian manifold on X depends only on G as a topological group.

6 Bounded Subsets of Apartments

We return to the study of a general Euclidean building $X = |\Delta|$. The theorem of this section is the analogue for Euclidean buildings of the fact that a spherical building admits a unique system of apartments, consisting of the convex hulls of pairs of opposite chambers (cf. §IV.5).

Given chambers C, C' of X , let $\mathcal{B}(C, C')$ be the union of all closed chambers $\overline{C''}$ such that C'' occurs in some minimal gallery from C to C' .

Theorem. *Let \mathcal{B} be the collection of bounded subsets $Y \subset X$ such that Y is contained in an apartment. Then \mathcal{B} is independent of the system of apartments \mathcal{A} . In fact, \mathcal{B} consists of all subsets $Y \subset X$ such that Y is contained in $\mathcal{B}(C, C')$ for some pair C, C' of chambers.*

Remark. It can be shown that $\mathcal{B}(C, C')$ is the smallest convex subcomplex containing C and C' ; here “convex” can be interpreted either combinatorially (in terms of minimal galleries), or geometrically (in terms of geodesics $[x, y]$). If you want to try to prove this as an exercise, start by reviewing Exercise 2 in §I.4E.

The proof of the theorem requires a result about Euclidean Coxeter complexes analogous to the lemma in §IV.5. We need some terminology before we can state it.

Let $E = |\Sigma|$ be the geometric realization of a Euclidean Coxeter complex, and let \mathcal{H} be the associated set of hyperplanes in E . Fix $x \in E$ and let $\overline{\mathcal{H}}$ be the set of hyperplanes through x and parallel to some element of \mathcal{H} . Then


$\overline{\mathcal{H}}$ is finite (cf. §1C above) and defines a decomposition of E into conical cells \mathfrak{A} with x as the cone point. These cells will simply be referred to as *conical cells based at x* . Here is another description of them that we will often use:

Choose an identification of Σ with the complex $\Sigma(W, V)$ associated to a Euclidean reflection group. This yields an identification of E with V . Let \overline{W} be the finite reflection group consisting of the linear parts of the elements of W . The conical cells based at x , then, are simply the translates $\mathfrak{A} = x + \mathfrak{D}$, where \mathfrak{D} is a cell associated to \overline{W} . We will call \mathfrak{D} the *direction* of \mathfrak{A} .

Suppose the \overline{W} -cell \mathfrak{D} is a chamber (hence a simplicial cone). Then the conical cell $x + \mathfrak{D}$ will be called a *sector* (“quartier” in [22]). If the n walls of \mathfrak{D} are defined by linear equations $f_i = 0$, where $f_i > 0$ on \mathfrak{D} , then a sector \mathfrak{C} with direction \mathfrak{D} is given by linear inequalities of the form $f_i > c_i$ ($i = 1, \dots, n$). It is clear from this that the intersection of two sectors with direction \mathfrak{D} is again a sector with direction \mathfrak{D} .


If \mathfrak{C} and \mathfrak{C}' are sectors with $\mathfrak{C}' \subseteq \mathfrak{C}$, then we will say that \mathfrak{C}' is a *subsector* of \mathfrak{C} . Note that \mathfrak{C} and \mathfrak{C}' then necessarily have the same direction. For suppose $\mathfrak{C} = x + \mathfrak{D}$ and $\mathfrak{C}' = x' + \mathfrak{D}'$. Letting \mathfrak{D} be defined by inequalities $f_i > 0$ as above, we conclude that the f_i are bounded below on \mathfrak{D}' , hence no f_i can be negative on the cone \mathfrak{D}' . Thus $f_i > 0$ on \mathfrak{D}' for all i , which implies that $\mathfrak{D}' \subseteq \mathfrak{D}$ and hence $\mathfrak{D}' = \mathfrak{D}$.

Consider now two sectors $\mathfrak{C}_1 = x + \mathfrak{D}$ and $\mathfrak{C}_2 = y - \mathfrak{D}$ having opposite directions \mathfrak{D} and $-\mathfrak{D}$. Let $\overline{\mathfrak{C}}_1$ and $\overline{\mathfrak{C}}_2$ be the closures $x + \overline{\mathfrak{D}}$ and $y - \overline{\mathfrak{D}}$. Assume that $x \in \overline{\mathfrak{C}}_2$ and $y \in \overline{\mathfrak{C}}_1$, so that the two closed sectors $\overline{\mathfrak{C}}_1, \overline{\mathfrak{C}}_2$ overlap, as in the following picture:



We will show that if C_1 and C_2 are chambers which are “sufficiently far out” in \mathfrak{C}_1 and \mathfrak{C}_2 , respectively, then $\mathcal{B}(C_1, C_2)$ contains $\overline{\mathfrak{C}}_1 \cap \overline{\mathfrak{C}}_2$. Let \mathfrak{C}'_1 be the subsector of \mathfrak{C}_1 based at y and let \mathfrak{C}'_2 be the subsector of \mathfrak{C}_2 based at x , as indicated by the dotted lines in the picture below; in other words, $\mathfrak{C}'_1 = y + \mathfrak{D}$ and $\mathfrak{C}'_2 = x - \mathfrak{D}$.

Lemma. *With the notation above, suppose C_1 and C_2 are chambers in E such that \overline{C}_1 meets \mathfrak{C}'_1 and \overline{C}_2 meets \mathfrak{C}'_2 . If C is any chamber of E such that \overline{C} meets $\overline{\mathfrak{C}}_1 \cap \overline{\mathfrak{C}}_2$, then C occurs in some minimal gallery from C_1 to C_2 .*



(Note: Since sectors are open sets, the hypothesis that \overline{C}_i meets \mathfrak{C}'_i implies that C_i meets \mathfrak{C}'_i .)

PROOF: As in the proof of the analogous lemma in §IV.5, we must prove that no wall (i.e., element of \mathcal{H}) separates C from both C_1 and C_2 . Let H be a wall, defined by a linear equation $f = c$. We may choose f so that $f > 0$ on \mathfrak{D} , in which case we will say that the closed half-space $f \geq c$ (resp. $f \leq c$) is the positive (resp. negative) side of H . (In the pictures above, you should think of \mathfrak{C}_1 and \mathfrak{C}'_1 as opening in the positive direction.)

The closed chamber \overline{C} is on one side of H . Suppose first that it is on the positive side. Then y must be on the positive side of H . For if y were strictly on the negative side, then \overline{C}_2 would be strictly on the negative side, contradicting the fact that \overline{C}_2 meets \overline{C} . It follows that \mathfrak{C}'_1 is strictly on the positive side of H , hence C_1 is on the positive side. Thus H does not separate C from C_1 . A similar argument shows that H does not separate C from C_2 if C is on the negative side of H . \square

EXERCISE

Take $x = y$ in the lemma, so that \overline{C}_1 and \overline{C}_2 meet only at the basepoint x . Deduce that $\mathcal{B}(C_1, C_2)$ contains a neighborhood of x if C_1 meets \mathfrak{C}_1 and C_2 meets \mathfrak{C}_2 .

PROOF OF THE THEOREM: Suppose Y is a bounded subset of an apartment E . Take an arbitrary direction \mathfrak{D} . Then we can find a pair of sectors $\mathfrak{C}_1, \mathfrak{C}_2$ as in the lemma, with $Y \subseteq \mathfrak{C}_1 \cap \mathfrak{C}_2$. In fact, with the notation above, we need only choose constants c_i, c'_i ($i = 1, \dots, n$) such that $c_i < f_i < c'_i$ on Y for all i . Consequently, there is a pair of chambers C_1, C_2 in E such that $Y \subseteq \mathcal{B}(C_1, C_2)$. Conversely, given chambers C, C' of X , choose an apartment E containing C and C' . Then the combinatorial convexity of apartments (cf. §IV.4) implies that E contains $\mathcal{B}(C, C')$; the latter is therefore a bounded subset of E [since there are only finitely many minimal galleries from C to C' in E], hence so is any subset of it. \square

7 A Metric Characterization of the Apartments

Assume from now on that \mathcal{A} is the complete system of apartments in our Euclidean building $X = |\Delta|$. In view of §6, of course, anything we say

concerning bounded subsets of apartments will then apply to an arbitrary apartment system. Recall that we gave two combinatorial characterizations of the apartments in §IV.4. We now characterize them from the metric space point of view.

Theorem 1. *Let $n = \dim X$. Then a subset $E \subseteq X$ is an apartment if and only if E is isometric to \mathbf{R}^n .*

Another way to say this is that (a) a subset isometric to \mathbf{R}^n is necessarily a subcomplex, and (b) the collection of all such subcomplexes is a system of apartments. These assertions can be viewed as generalizations to arbitrary X of elementary facts about trees.

We will deduce Theorem 1 from the more precise Theorem 2 below. A subset of X is called *convex* if it contains the geodesic $[x, y]$ connecting any two of its points x, y .

Theorem 2. *Let Y be a subset of X . Assume either that Y is convex or that Y has non-empty interior. If Y is isometric to a subset of \mathbf{R}^n , then Y is contained in an apartment.*

To deduce Theorem 1 from Theorem 2, suppose E is isometric to \mathbf{R}^n . Then E is easily seen to be convex; this follows from the characterization of geodesics in terms of the metric (cf. part (4) of the theorem in §3A). So Theorem 2 implies that E is contained in an apartment E' . But E' cannot be isometric to a proper subset of itself, so E must be the entire apartment E' .

The rest of this section will be devoted to the proof of Theorem 2. The proof will use some of the ideas introduced in the exercises of §IV.4 (and in the “hints” following them). The first lemma is the analogue of Exercise 4(a) of §IV.4. For any chambers C, D of X , let $\lambda_{C,D} : \overline{C} \rightarrow \overline{D}$ be the unique type-preserving simplicial isomorphism. It is an isometry; this follows from the discussion of the map called λ in the proof in §3A above.

Lemma 1. *Suppose Y contains a non-empty open subset U of some chamber C . Let E be any apartment and D any chamber of E . If Y is isometric to a subset of \mathbf{R}^n , then there is a unique isometry α from Y into E such that $\alpha|U = \lambda_{C,D}|U$. Moreover, $\alpha = \rho|Y$, where ρ is the canonical map $X \rightarrow E$ taking C to D .*

(The notion of “canonical map” was introduced in Exercise 4 of §IV.4.)

PROOF: Suppose first that there exists an isometry α from Y into E such that $\alpha|U = \lambda|U$, where $\lambda = \lambda_{C,D}$. Then $\rho\alpha^{-1} : \alpha(Y) \rightarrow E$ fixes the open set $\lambda(U)$ pointwise and preserves distances from points of $\lambda(U)$. Hence $\rho\alpha^{-1} = \text{id}_{\alpha(Y)}$ by Exercise 1 of §1A. This proves the last assertion of the lemma, as well as the uniqueness of α . To prove the existence, start with an arbitrary isometry β from Y into E . Then $\beta(U)$ and $\lambda(U)$ are isometric subsets of E , and the isometry $\lambda\beta^{-1} : \beta(U) \rightarrow \lambda(U)$ extends to an isometry $\gamma : E \rightarrow E$ by Exercise 5 of §1A. So we may take $\alpha = \gamma\beta$. \square

The next lemma is the crucial step in the proof. It is the analogue of Exercise 4(b) of §IV.4.

Lemma 2. *Suppose that Y contains a closed chamber \overline{C} and that α is an isometry from Y into an apartment E such that α maps \overline{C} onto a closed chamber \overline{D} by the map $\lambda_{C,D}$. Let D' be a chamber of E adjacent to D . Then there is a chamber C' adjacent to C such that α extends to an isometry from $Y \cup \overline{C'}$ into E taking $\overline{C'}$ to $\overline{D'}$ by the map $\lambda_{C',D'}$.*

PROOF: We know from Lemma 1 that $\alpha = \rho|_Y$, where $\rho : X \rightarrow E$ is the canonical map taking C to D . We wish to find a chamber C' adjacent to C such that α is also equal to $\rho'|_Y$, where $\rho' : X \rightarrow E$ is the canonical map taking C' to D' . If we can find such a C' we will be done. For then $\rho'|_Y$ will extend α , will be $\lambda_{C',D'}$ on $\overline{C'}$, and will be an isometry because ρ' satisfies $d(\rho'(x), \rho'(y)) = d(x, y)$ for any $x, y \in X$ with $x \in \overline{C'}$.

We may assume $D \neq D'$. Choose a labelling of Δ , and let i be the label such that D and D' are i -adjacent. Let C' be any chamber distinct from C and i -adjacent to it, and let $\rho' : X \rightarrow E$ be the canonical map taking C' to D' . The hint to Exercise 4(b) of §IV.4 essentially contains a complete analysis of the relation between ρ and ρ' . Namely, for any chamber C'' of X , we have the following three possibilities:

(a) $\rho(C'')$ is on the same side of H as D , where H is the wall of E separating D from D' . In this case there is a minimal gallery of the form C', C, \dots, C'' , and we have $\rho'(C'') = \rho(C'')$. [This equality was proved in the hint cited above via a distance-function δ with values in the reflection group W associated to E ; alternatively, one could use the standard uniqueness argument.]

(b) $\rho(C'')$ is on the same side of H as D' , and there is a minimal gallery of the form C, C', \dots, C'' . Once again, $\rho'(C'') = \rho(C'')$ in this case.

(c) $\rho(C'')$ is on the same side of H as D' , and there is a chamber C_1 which is i -adjacent to both C and C' and satisfies $d(C, C'') = d(C', C'') = d(C_1, C'') + 1$. Thus there are minimal galleries of the form (C, Γ) and (C', Γ) , where Γ is a minimal gallery from C_1 to C'' . In this case $\rho'(C'') = s\rho(C'')$, where s is the reflection of E with respect to H . [As before, this was proved via δ in the hint to Exercise 4(b) of §IV.4, but one could also apply the standard uniqueness argument to see that $\rho' = s\rho$ along Γ .]

It follows that for any $y \in Y$ we have $\rho'(y) = \alpha(y)$ except possibly if $\alpha(y)$ is in the open half-space U bounded by H and containing D' , in which case we might have $\rho'(y) = s\alpha(y)$. Let $Z = \alpha(Y) \cap U$ and let $f = \rho'\alpha^{-1} : Z \rightarrow E$. Then f is distance-decreasing, and for all $z \in Z$ we have $f(z) = z$ or sz . We must show that C' can be chosen so that $f(z) = z$ for all $z \in Z$. Note first that we can choose C' so that $f(z) = z$ for at least one $z \in Z$ (unless $Z = \emptyset$, in which case we're already done). For let z be a point of Z , and let $\overline{C''}$ be a closed chamber containing $\alpha^{-1}(z)$. Then C'' satisfies either (b) or (c). If (b) holds, then $f(z) = z$. If (c) holds, then we can simply change the choice of C' and use the chamber C_1 instead; with this new C' , then,

we are in case (b) and hence $f(z) = z$.

We now use the fact that $d(z, sz') > d(z, z')$ for any $z, z' \in U$. [This has a proof similar to that of the analogous combinatorial fact, given in Exercise 1 of §III.4A; namely, consider the line segment $[z, sz']$, and fold it back onto \bar{U} to obtain a path from z to z' which has the same length but is not straight.] Since $f(z) = z$ for some $z \in Z$, this fact implies that $f(z') = z'$ for all $z' \in Z$, as required; for otherwise we would have $f(z') = sz'$, contradicting the fact that f is distance-decreasing. \square

Lemma 3. *If Y contains a closed chamber and is isometric to a subset of \mathbf{R}^n , then Y is contained in an apartment.*

(This is a special case of the theorem we are trying to prove; for a closed chamber has non-empty interior by Exercise 3 of §3A.)

PROOF: Suppose $\bar{C} \subseteq Y$, and let E be any apartment containing C . By Lemma 1 there is an isometry α from Y into E which fixes C pointwise. By repeated application of Lemma 2, we can successively adjoin closed chambers to $\alpha(Y)$ and extend α^{-1} to an isometry β from E into X which is simplicial on each closed chamber of E . But then β is a type-preserving simplicial map, and its image $\beta(E)$ is therefore an apartment containing Y by Exercise 2 of §IV.4. \square

PROOF OF THEOREM 2: In view of Lemma 3, it suffices to show that we can enlarge the given Y to a set which contains a closed chamber and is still isometric to a subset of \mathbf{R}^n . Suppose first that Y has non-empty interior. Then Y contains a non-empty open subset of a chamber C . Let E be an apartment containing C and let $\rho = \rho_{E,C} : X \rightarrow E$. Lemma 1 implies that ρ maps Y isometrically into E . But then ρ also maps $Y \cup \bar{C}$ isometrically into E , since ρ is the identity on \bar{C} and preserves distances from points of \bar{C} . So we are done in this case.

Suppose now that Y is convex. Choose a simplex A which is maximal among the simplices meeting Y . Let C be a chamber having A as a face, and let E be an apartment containing C . We will show that $\rho = \rho_{E,C}$ maps $Y \cup \bar{C}$ isometrically into E . As above, it suffices to show that $d(\rho(y), \rho(z)) = d(y, z)$ for all $y, z \in Y$. We may assume $y, z \notin \bar{C}$. Choose $x \in Y \cap A$, and let $T \subseteq Y$ be the convex hull of $\{x, y, z\}$. Note that any isometry α from Y into \mathbf{R}^n must take T to the convex hull of $\{\alpha(x), \alpha(y), \alpha(z)\}$; this follows from the metric characterization of geodesics. So T is, in an obvious sense, a *Euclidean triangle*. Since $y \neq x$ and $z \neq x$, it follows that there is a well-defined angle θ at the vertex x , with $0 \leq \theta \leq \pi$.

If we take any $y' \in (x, y]$ and $z' \in (x, z]$, then the triangle T' determined by $\{x, y', z'\}$ has the same angle at x . In particular, we will take y' and z' close enough to x that they are in A and hence in \bar{C} . [This is possible because of the maximality of A ; for if $[x, y]$, say, does not stay in A for a little while, then it enters a simplex having A as a proper face.] Now ρ maps $[x, y]$ (resp. $[x, z]$) isometrically onto $[x, \rho(y)]$ (resp. $[x, \rho(z)]$) and fixes T' .

The angle θ is therefore equal to the angle between $[x, \rho(y)]$ and $[x, \rho(z)]$. By elementary geometry ("two sides and the included angle") we conclude that T is congruent to the triangle in E with vertices $x, \rho(y), \rho(z)$. Hence $d(\rho(y), \rho(z)) = d(y, z)$, as required. \square

EXERCISES

1. Extract the following fact from the proof above: Given $x, y, z \in X$ with $x \neq y$ and $x \neq z$, there is a well-defined angle θ between $[x, y]$ and $[x, z]$. [HINT: Take $y' \in (x, y)$ and $z' \in (x, z)$ close enough to x that $[x, y']$ and $[x, z']$ are each contained in a closed simplex. Then $[x, y'] \cup [x, z']$ is contained in an apartment, so the convex hull of $\{x, y', z'\}$ is a Euclidean triangle.]

2. With x, y, z as in Exercise 1, prove the following *cosine inequality*:

$$d^2(y, z) \geq d^2(x, y) + d^2(x, z) - 2d(x, y)d(x, z)\cos\theta,$$

with equality if and only if $\{x, y, z\}$ is contained in an apartment. This reinforces the negative curvature intuition again: Two geodesics emanating from x tend to separate faster than they ought to. [HINT: To prove the inequality, take an apartment E containing $\{x, y', z'\}$ as in the hint to Exercise 1, and consider $\rho = \rho_{E, C}$, where $x \in \overline{C}$. If equality holds, then ρ is an isometry on $\{y, z\} \cup \overline{C}$.]

3. Let E be an apartment and U a half-space of E bounded by a wall H . Let C be a chamber of X having a codimension 1 face in H . Show that there is an apartment E' containing U and C . [HINT: For any chamber D in U , there is a minimal gallery from D to C of the form (Γ, C) , where Γ is a gallery in U . Deduce that $U \cup C$ is isometric to a subset of E .]

4. If Δ is thick, show that the simplicial decomposition of $X = |\Delta|$ is completely determined by the metric. [HINT: With the notation of Exercise 3, what is $E \cap E'$?]

8 Construction of Apartments

We continue to denote by X a Euclidean building, equipped with its complete system of apartments. As an illustration of Theorem 2 above, we will prove two results (parts (1) and (2) of the theorem below) asserting the existence of apartments containing given subsets of X .

By a *sector* in X we will mean a subset \mathfrak{C} which is contained in some apartment E and is a sector in E (cf. §6). \mathfrak{C} is then a sector in any apartment E' that contains it. To prove this, note first that $E \cap E'$ is a subcomplex with non-empty interior, hence it contains a chamber; the general theory of buildings therefore implies that there is an isomorphism $E' \rightarrow E$ fixing $E \cap E'$ pointwise, and our assertion follows easily.

Remark. Note, for future reference, that we can similarly treat arbitrary conical cells \mathfrak{A} in X , not just sectors (which are conical cells of maximal dimension). The point is that there is still an isomorphism $E' \rightarrow E$ fixing

$E \cap E'$ even if E and E' do not have a common chamber. In fact, it suffices to take an isomorphism $\phi : E' \rightarrow E$ which fixes pointwise a maximal simplex A of the subcomplex $E \cap E'$. One then proves that ϕ fixes $E \cap E'$ pointwise by a geometric analogue of the standard uniqueness argument. [Choose any $x \in A$. Given $y \neq x$ in $E \cap E'$, consider the geodesic $[x, y]$. It is contained in $E \cap E'$, and a non-trivial initial segment of it is contained in A by maximality; hence $[x, y]$ and its image $[x, \phi(y)]$ have a common initial segment. But these are geodesics of the same length in the Euclidean space E , so they must coincide. In particular, $\phi(y) = y$.]

Let's return now to sectors, which are the only conical cells that will concern us in this section.

Theorem.

- (1) *Given a sector \mathfrak{C} and a chamber C in X , there is a subsector $\mathfrak{C}' \subseteq \mathfrak{C}$ such that \mathfrak{C}' and C are contained in some apartment.*
- (2) *Given two sectors \mathfrak{C}_1 and \mathfrak{C}_2 , there are subsectors $\mathfrak{C}'_1 \subseteq \mathfrak{C}_1$ and $\mathfrak{C}'_2 \subseteq \mathfrak{C}_2$ such that \mathfrak{C}'_1 and \mathfrak{C}'_2 are contained in some apartment.*

(You should think about trees to see why the theorem is plausible; a sector, in this case, is simply a ray tending toward an "end" of the tree.)

PROOF OF (1): Choose an apartment E containing \mathfrak{C} . In view of the previous section, it suffices to find a subsector $\mathfrak{C}' \subseteq \mathfrak{C}$ and a chamber C_0 in E such that the retraction $\rho = \rho_{E, C_0} : X \rightarrow E$ is an isometry on $\mathfrak{C}' \cup C$. Note first that there is a bounded subset Z of E such that, for any choice of C_0 , we will have $\rho(C) \subseteq Z$. In fact, let z be any point of E and let Y be any ball in X centered at z and containing C ; then we can take $Z = Y \cap E$. [This works because our retractions are distance-decreasing.] Now let \mathfrak{C}'' be a sector in E containing Z and having direction opposite to that of \mathfrak{C} . We can choose \mathfrak{C}'' so that its basepoint x is in \mathfrak{C} . The desired \mathfrak{C}' and C_0 are now obtained as follows: \mathfrak{C}' is the subsector of \mathfrak{C} based at x , and C_0 is any chamber of E whose closure contains x .

To prove that these choices work, we will show that $\rho = \rho_{E, C_0}$ is an isometry on $\overline{C'} \cup C$ for any chamber C' which meets \mathfrak{C}' . Now we know that $\rho_{E, C'}$ is an isometry on $\overline{C'} \cup C$, so it suffices to show that $\rho(C) = \rho_{E, C'}(C)$. Let Γ be a minimal gallery from C_0 to C and let Γ' be a minimal gallery in E from C' to C_0 . Then $\rho(\Gamma)$ is a minimal gallery from C_0 to $\rho(C)$. Now apply the lemma of §6 to the two opposite sectors $\mathfrak{C}', \mathfrak{C}''$, with $y = x$. Recalling that $\rho(C) \subseteq Z \subseteq \mathfrak{C}''$, we conclude that $\Gamma' * \rho(\Gamma)$ is a minimal gallery from C' to $\rho(C)$, where the star denotes composition of galleries. Since this gallery is the image under ρ of $\Gamma' * \Gamma$, the latter must be a minimal gallery from C' to C . But then its image under $\rho_{E, C'}$ is also minimal, so the standard uniqueness argument shows that $\rho(C) = \rho_{E, C'}(C)$. This completes the proof of (1). \square


As a consequence of (1) we can define a new kind of retraction onto an apartment, which will be useful in the proof of (2). Given an apartment E

and a sector \mathfrak{C} in E , (1) implies that X is the union of the apartments E' which contain a subsector of \mathfrak{C} . We now define $\rho = \rho_{E,\mathfrak{C}} : X \rightarrow E$ to be the map whose restriction to any such E' is the isomorphism $\phi_{E'} : E' \rightarrow E$ which fixes $E \cap E'$ pointwise. It is easy to check, as in the construction of the “ordinary” retractions $\rho_{E,C}$ in Chapter IV, that ρ is well-defined. Note, for future reference, that for any chamber C of X there is a subsector \mathfrak{C}' of \mathfrak{C} such that $\rho_{E,\mathfrak{C}}(C) = \rho_{E,\mathfrak{C}'}(C)$ for any chamber C' of E which meets \mathfrak{C}' ; indeed, if we take \mathfrak{C}' to be a subsector such that \mathfrak{C}' and C are contained in an apartment E' , then both sides of the equality to be proved are equal to $\phi_{E'}(C)$.

EXERCISES

1. Describe $\rho_{E,\mathfrak{C}}$ if X is a tree.
2. Prove that for any bounded subset Y of X there is a subsector \mathfrak{C}' of \mathfrak{C} such that $\rho_{E,\mathfrak{C}}|Y = \rho_{E,\mathfrak{C}'}|Y$ for any chamber C of E which meets \mathfrak{C}' . [HINT: Use the method of proof of (1) to find a \mathfrak{C}' such that $\rho_{E,\mathfrak{C}}|Y$ is the same for all chambers C which meet \mathfrak{C}' .]
3. Prove or disprove the following purported generalization of (1): *Given a sector \mathfrak{C} and a bounded subset Y of an apartment, there is a subsector $\mathfrak{C}' \subseteq \mathfrak{C}$ such that \mathfrak{C}' and Y are contained in some apartment.*

We will use this new kind of retraction in the proof of part (2) of the theorem. More precisely, we will choose an apartment E_1 containing \mathfrak{C}_1 , and we will find subsectors \mathfrak{C}'_i of \mathfrak{C}_i ($i = 1, 2$) such that $\rho = \rho_{E_1,\mathfrak{C}_1}$ is an isometry on $\mathfrak{C}'_1 \cup \mathfrak{C}'_2$. The key step in the proof is to figure out how to choose \mathfrak{C}'_2 . Consider, for example, the case where X is a tree. The following picture shows a typical configuration:



Note that $\rho(\mathfrak{C}_2)$ heads toward the same end of E_1 as \mathfrak{C}_1 at first, but then it reverses direction and heads toward the opposite end. So if we want to find a subsector \mathfrak{C}'_2 on which ρ is an isometry, we need only start \mathfrak{C}'_2 at any chamber C in \mathfrak{C}_2 such that $\rho(C)$ is already heading in the opposite direction from \mathfrak{C}_1 .

It turns out that a similar idea works in higher dimensions also. Here are the details:

PROOF OF (2): Choose apartments E_1 and E_2 containing \mathfrak{C}_1 and \mathfrak{C}_2 , respectively, and let $\rho = \rho_{E_1,\mathfrak{C}_1}$ as above. Identify E_1 with a vector space V as in §6, so that \mathfrak{C}_1 has a direction \mathfrak{D} ; the latter is a chamber in the Coxeter complex associated to a finite reflection group \overline{W} . We will use combinatorial distances $d(-, -)$ in this Coxeter complex in order to compare directions.

The intuitive idea to keep in mind is that the bigger $d(\mathfrak{D}, \mathfrak{D}')$ is, the more nearly opposite \mathfrak{D} and \mathfrak{D}' are.

We wish to associate to any chamber C of E_2 a \overline{W} -chamber \mathfrak{D}' in E_1 , which we think of as the *direction of $\rho(\mathfrak{C}_2)$ at C* . It is defined as follows: Choose a directed line segment \overrightarrow{xy} in C which is *parallel* to \mathfrak{C}_2 , in the sense that it is a translate of a segment going from the cone point of \mathfrak{C}_2 to some point of \mathfrak{C}_2 . Then \mathfrak{D}' is defined to be the unique \overline{W} -chamber such that $\rho(\overrightarrow{xy})$ is parallel to \mathfrak{D}' . One can easily check (by using a suitable isomorphism $E_2 \rightarrow E_1$) that \mathfrak{D}' is independent of the choice of \overrightarrow{xy} .

Let's focus now on those chambers of E_2 which meet \mathfrak{C}_2 . Choose among these a chamber C_0 such that the resulting direction \mathfrak{D}' makes $d(\mathfrak{D}, \mathfrak{D}')$ as big as possible. Such a C_0 exists because $d(\mathfrak{D}, \mathfrak{D}')$ is bounded by the diameter of the spherical Coxeter complex associated to \overline{W} . Let x_0 be any point in C_0 , and let \mathfrak{C}'_2 be the subsector of \mathfrak{C}_2 based at x_0 . As we noted above while defining ρ , we can find a subsector $\mathfrak{C}'_1 \subseteq \mathfrak{C}_1$ such that $\rho(C_0) = \rho_{E_1, C}(C_0)$ for any chamber C which meets \mathfrak{C}'_1 . Passing to a further subsector if necessary, we can also arrange that $\mathfrak{C}'_1 \subseteq \rho(x_0) + \mathfrak{D}$. We will show that ρ is an isometry on $\mathfrak{C}'_1 \cup \mathfrak{C}'_2$. In view of §7, this will complete the proof.

Let C' be any chamber of E_2 which meets \mathfrak{C}'_2 . Choose $x \in C' \cap \mathfrak{C}'_2$, and consider the directed line $\overrightarrow{x_0x}$. It crosses exactly those walls H_1, \dots, H_l of E_2 which separate C_0 from C' . By moving x slightly, if necessary, we can make sure that $\overrightarrow{x_0x}$ does not simultaneously cross two walls. For if $\overrightarrow{x_0x}$ meets $H_i \cap H_j$ for some $i \neq j$, then x is in the affine span of x_0 and $H_i \cap H_j$; this affine span is a hyperplane, so we need only choose x so as to miss finitely many hyperplanes. With such a choice of x , then, $\overrightarrow{x_0x}$ passes through chambers $C_0, \dots, C_l = C'$ which form a minimal gallery from C_0 to C' .

I claim that ρ maps C_0, \dots, C_l to a non-stuttering gallery in E_1 and that, in addition, $\rho(C_i) = \rho_{E_1, C}(C_i)$ for each $i = 1, \dots, l$ and any chamber C that meets \mathfrak{C}'_1 . Once the claim is proved, we will be done. For the first assertion of the claim implies [by the standard uniqueness argument] that ρ coincides on \mathfrak{C}'_2 with the type-preserving isomorphism $E_2 \rightarrow E_1$ taking C_0 to $\rho(C_0)$. Hence ρ is an isometry on \mathfrak{C}'_2 . And the second assertion of the claim implies that ρ preserves distances between points of \mathfrak{C}'_1 and points of \mathfrak{C}'_2 . It remains to prove the claim.


Arguing by induction on l , we may assume that $l > 0$ and that the claim is known for the subgallery C_0, \dots, C_{l-1} . Hence ρ is an isometry on $\mathfrak{C}'_1 \cup C_0 \cup \dots \cup C_{l-1}$, and the latter is therefore contained in an apartment E' . Note, then, that ρ maps E' isomorphically onto E_1 . Moreover, $\rho|_{E'} = \rho_{E_1, C}|_{E'}$ for any chamber C of $E_1 \cap E'$ and hence, in particular, for any chamber C which meets \mathfrak{C}'_1 . Let A be the common face of C_{l-1} and C_l , and let C'_l be the chamber of E' distinct from C_{l-1} and adjacent to it along A . Let H be the support of A in E' , i.e., the wall of E' separating

C_{l-1} from C'_l . As in the proof of the lemma in §6, it will be convenient to refer to the two closed half-spaces of E' determined by H as the positive and negative sides, the positive side being defined by a linear inequality $f \geq c$ with f bounded below on \mathfrak{C}'_1 . We can similarly define the positive and negative sides of the wall $\rho(H)$ in E_1 .

We now consider three cases:

(a) $C'_l = C_l$. In other words, E' contains C_l . Since $\rho|_{E'}$ is an isomorphism and coincides with $\rho_{E_1, C}|_{E'}$ for any C meeting \mathfrak{C}'_1 , the claim is trivial in this case.


(b) $C'_l \neq C_l$, and C_0 is on the positive side of H . The following picture illustrates this case when X is a tree:



Since x_0 is on the positive side of H , $\rho(x_0)$ is on the positive side of $\rho(H)$; hence all of $\rho(x_0) + \mathfrak{D}$ is on the positive side of $\rho(H)$. In view of our choice of \mathfrak{C}'_1 , it follows that \mathfrak{C}'_1 is on the positive side of $\rho(H)$ in E_1 , whence \mathfrak{C}'_1 is on the positive side of H in E' .

Suppose now that C is any chamber which meets \mathfrak{C}'_1 . Then C is on the positive side of H , so there is a minimal gallery in E' of the form C, \dots, C_{l-1}, C'_l . Replacing C'_l by C_l , we obtain a gallery which is still minimal; you can prove this either by considering types of galleries (cf. Exercise 1 in §IV.4) or by using a suitable retraction of X onto E' . So $\rho_{E_1, C}$ maps this gallery to a minimal gallery in E_1 ; in particular, $\rho_{E_1, C}$ maps our original gallery C_0, \dots, C_l to a non-stuttering gallery. Since $\rho_{E_1, C}(C_i) = \rho(C_i)$ for $i < l$, it follows that $\rho_{E_1, C}(C_l)$ is independent of C . This common value of $\rho_{E_1, C}(C_l)$ for C meeting \mathfrak{C}'_1 must equal $\rho(C_l)$, and the claim is now proved in case (b).

(c) $C'_l \neq C_l$, and C_0 is on the negative side of H . The following picture illustrates this case when X is a tree; it suggests that ρ changes direction as one goes from C_{l-1} to C_l .



We will show that our choice of C_0 prohibits this case from occurring. Let $\mathfrak{C}''_1 \subseteq \mathfrak{C}'_1$ be the subsector based at some point of \mathfrak{C}'_1 on the positive side of H . Then \mathfrak{C}''_1 lies entirely on the positive side of H . Considering galleries as in case (b), we conclude that $\rho_{E, C}(C_l) = \rho_{E, C}(C_{l-1})$ for any chamber C that meets \mathfrak{C}''_1 . Consequently, $\rho(C_l) = \rho(C_{l-1})$.

Recall now that we have a directed line segment $\overrightarrow{x_0x}$ which is parallel to \mathfrak{C}_2 and which passes through C_0, \dots, C_l . The initial portion of this segment is mapped by ρ to a directed segment in E parallel to a \overline{W} -chamber that we called \mathfrak{D}' . Let x_1 be the point where $\overrightarrow{x_0x}$ crosses A , and let y_0, y_1 , and y be, respectively, $\rho(x_0)$, $\rho(x_1)$, and $\rho(x)$. Let $z = sy$, where s is the reflection of E_1 with respect to $\rho(H)$. Then ρ maps $\overrightarrow{x_0x}$ to the path obtained from $\overrightarrow{y_0z}$ by folding it onto the negative side of $\rho(H)$. In particular, $\overrightarrow{y_1z}$ has the same direction as $\overrightarrow{y_0y_1} = \rho(\overrightarrow{x_0x_1})$, hence it is parallel to \mathfrak{D}' . Thus $y_1 + \mathfrak{D}'$ is on the positive side of $\rho(H)$, i.e., on the same side as $y_1 + \mathfrak{D}$. Moreover, since $\rho(\overrightarrow{x_1x}) = s(\overrightarrow{y_1z})$, the direction of $\rho(\mathfrak{C}_2)$ at C_l is $\bar{s}\mathfrak{D}'$, where \bar{s} is the linear part of s . But $d(\mathfrak{D}, \bar{s}\mathfrak{D}') > d(\mathfrak{D}, \mathfrak{D}')$ by Exercise 1 in §III.4A, contradicting the maximality of $d(\mathfrak{D}, \mathfrak{D}')$. This contradiction completes the proof of the claim and hence of the theorem. \square

Remark. It is not really necessary to use the existence of the apartment E' above, although this does make it easier to follow the proof. You might want to try, as an exercise, to carry out the inductive step without using E' . If you get stuck, see Bruhat–Tits [22], Proposition 2.9.5.

9 The Spherical Building at Infinity

At the end of Chapter V we suggested the possibility of constructing a spherical building by attaching a “sphere at infinity” to every apartment of a Euclidean building. In this section we carry out the details of that construction. X continues to denote a Euclidean building, and “apartment” continues to refer to the complete apartment system unless the contrary is explicitly stated.

9A Ideal points and ideal simplices

A *ray* in X is a subset \mathfrak{r} which is isometric to the half-line $[0, \infty)$. The point $x \in \mathfrak{r}$ which corresponds to 0 under the unique such isometry will be called the *origin* or *basepoint* of \mathfrak{r} . We will also say that \mathfrak{r} *emanates* from x . A ray is easily seen to be convex, so §7 implies that it is contained in some apartment E . As a subset of E , it is necessarily a ray in the usual sense, i.e., a subset of the form $\mathfrak{r} = \{(1-t)x + ty : t \geq 0\}$ for some $y \neq x$.

We will say that two rays $\mathfrak{r}, \mathfrak{s}$ are *parallel* if the sets of real numbers

$$\{d(y, \mathfrak{s}) : y \in \mathfrak{r}\} \quad \text{and} \quad \{d(z, \mathfrak{r}) : z \in \mathfrak{s}\}$$

are bounded. In other words, we require that there be a number M such that for each $y \in \mathfrak{r}$ there is a $z \in \mathfrak{s}$ with $d(y, z) < M$, and similarly with the roles of \mathfrak{r} and \mathfrak{s} reversed. If \mathfrak{r} and \mathfrak{s} are subsets of some apartment E , one can easily check that they are parallel if and only if there is a translation of E taking one to the other. It is also easy to check that the relation of

parallelism is an equivalence relation. An equivalence class of rays will be called an *ideal point* of X .

If τ is a ray emanating from x and representing an ideal point e , one thinks of e as sitting “at infinity”, or at the “end” of τ . To reinforce this intuition, we will write $\tau = [x, e)$; this notation is justified by the following lemma, which shows that an ideal point admits a unique representative ray emanating from a given point x .

Lemma 1. *Given a point x and a ray s , there is a unique ray τ which is based at x and parallel to s .*

PROOF: To prove existence, let E be an apartment containing s . Let \mathcal{C} be a sector in E , based at the origin of s , such that the closure of \mathcal{C} contains s . By §8 we can find a subsector \mathcal{C}' of \mathcal{C} such that \mathcal{C}' and x are contained in some apartment E' . Since \mathcal{C}' is a translate of \mathcal{C} in E , its closure contains a ray s' parallel to s ; we can now translate s' in E' to obtain the desired τ .

We must now show that there cannot exist distinct parallel rays with the same origin. Suppose, to the contrary, that τ_1 and τ_2 are distinct parallel rays based at x . Then $\tau_1 \cap \tau_2$, being closed and convex, must be an interval $[x, x']$. Replacing τ_1 and τ_2 by the subrays based at x' , we are reduced to the case where $\tau_1 \cap \tau_2 = \{x\}$. By the exercises in §7, there is then a well-defined angle $\theta > 0$ between τ_1 and τ_2 at x , and the cosine inequality holds: For any $s, t \geq 0$ let p_s (resp. q_t) be the point of τ_1 (resp. τ_2) at distance s (resp. t) from x ; then

$$d^2(p_s, q_t) \geq s^2 + t^2 - 2st \cos \theta.$$

Fix s and consider the right-hand side of this inequality as t varies. If $\theta \geq \pi/2$, then the minimum value of the right-hand side is s^2 , which is achieved when $t = 0$. If $\theta < \pi/2$, then the minimum value is $s^2 \sin^2 \theta$, which is achieved when $t = s \cos \theta$. In either case, $\lim_{s \rightarrow \infty} d(p_s, \tau_2) = \infty$, contradicting the assumption that τ_1 and τ_2 are parallel. \square

Let X_∞ be the set of ideal points. We wish to decompose X_∞ into “ideal simplices”. Let \mathfrak{A} be a conical cell in X , as defined at the beginning of §8. The *face of \mathfrak{A} at infinity*, denoted \mathfrak{A}_∞ , is defined to be the set of ideal points e such that \mathfrak{A} contains the open ray $(x, e) = [x, e) - \{x\}$, where x is the cone point of \mathfrak{A} . Note that we can recover \mathfrak{A} from its cone point x and its face at infinity $F = \mathfrak{A}_\infty$; namely, \mathfrak{A} is the “open join” $x * F$, where the latter is defined as follows:

$$x * F = \begin{cases} \{x\} & \text{if } F = \emptyset \\ \bigcup_{e \in F} (x, e) & \text{otherwise.} \end{cases}$$

We now define an *ideal simplex* of X to be a subset F of X_∞ such that $F = \mathfrak{A}_\infty$ for some conical cell \mathfrak{A} .

Lemma 2. *If F is an ideal simplex and x is an arbitrary point of X , then there is a conical cell \mathfrak{A} based at x such that $F = \mathfrak{A}_\infty$. Consequently, there*

is a 1-1 correspondence between the set of ideal simplices of X and the set of conical cells based at any given point $x \in X$.

PROOF: The proof is similar to that of Lemma 1: Write $F = \mathfrak{B}_\infty$ for some conical cell \mathfrak{B} , let E be an apartment containing \mathfrak{B} , and choose a sector \mathfrak{C} in E (based at the cone point of \mathfrak{B}) such that \mathfrak{B} is a face of \mathfrak{C} . Replacing \mathfrak{C} by a subsector and \mathfrak{B} by a translate, we may assume that \mathfrak{B} and x are contained in an apartment E' . The desired \mathfrak{A} is then a translate of \mathfrak{B} in E' . This proves the first assertion, and the second follows at once. \square

Lemma 3. *The ideal simplices partition X_∞ .*

PROOF: Any open ray (x, e) is contained in an apartment E . It is therefore contained in some conical cell \mathfrak{A} in E based at x , whence $e \in \mathfrak{A}_\infty$. This shows that X_∞ is the union of the ideal simplices. Suppose now that we have two distinct ideal simplices $F = \mathfrak{A}_\infty$ and $F' = \mathfrak{A}'_\infty$. Then \mathfrak{A} (resp. \mathfrak{A}') is a face of a sector \mathfrak{C} (resp. \mathfrak{C}') in some apartment E (resp. E'). By §8, there is an apartment E'' containing subsectors of \mathfrak{C} and \mathfrak{C}' . We may therefore replace \mathfrak{A} and \mathfrak{A}' by translates (in E and E'), in order to reduce to the case where F and F' are represented by conical cells in E'' . But now it is evident that F and F' are disjoint; in fact, we can represent F and F' by conical cells $x * F$ and $x * F'$ with $x \in E''$, so our assertion follows from the fact that the conical cells in E'' based at a given point x partition E'' . \square

In the theory of trees, one usually defines an *end* of a tree X to be an equivalence class of rays, where two rays are equivalent if they have a common subray. If you are familiar with this theory, then you have probably already thought about how it relates to the theory being developed here. The answer is that if our Euclidean building X is a tree, then the ideal points of X are the same as its ends. In other words, two rays represent the same ideal point if and only if they have a common subray. Here is a generalization of this fact to Euclidean buildings of arbitrary dimension:

Lemma 4. *Two sectors of X have the same face at infinity if and only if they have a common subsector.*

PROOF: It is obvious that a sector has the same face at infinity as any subsector, whence the "if" part. Conversely, suppose \mathfrak{C}_1 and \mathfrak{C}_2 are sectors with the same face at infinity. Let E be an apartment containing subsectors \mathfrak{C}'_1 and \mathfrak{C}'_2 . Then these subsectors have the same face at infinity, so they have the same direction \mathfrak{D} (defined with respect to some vector space structure on E). The intersection $\mathfrak{C}'_1 \cap \mathfrak{C}'_2$ is then a sector with direction \mathfrak{D} , so it is a common subsector of \mathfrak{C}_1 and \mathfrak{C}_2 . \square

9B Construction of the building at infinity

Let Δ_∞ be the set of ideal simplices of X . The first step is to define a face relation on this set. Recall that there is a face relation on the set of conical cells in an apartment E based at a given point x . We extend this to conical

cells in X based at x by saying that \mathfrak{A}' is a *face* of \mathfrak{A} if \mathfrak{A}' is contained in the closure of \mathfrak{A} and is a face of \mathfrak{A} in some apartment containing \mathfrak{A} . In this case \mathfrak{A}' is a face of \mathfrak{A} in every apartment containing \mathfrak{A} .

We can now use the 1-1 correspondence in Lemma 2 to define a face relation on Δ_∞ . Thus we say that F' is a *face* of F if $x * F'$ is a face of $x * F$ for some $x \in X$. A glance at the proof of Lemma 2 shows that $x * F'$ is then a face of $x * F$ for every $x \in X$. For any apartment $E = |\Sigma|$ of X , let Σ_∞ be the set of ideal simplices F such that $F = \mathfrak{A}_\infty$ for some conical cell \mathfrak{A} in E . Note that Σ_∞ is a subset of Δ_∞ closed under passage to faces. (So we can call it a subcomplex, as soon as we have proven that Δ_∞ is a simplicial complex.) Note further that Σ_∞ , with the face relation that it inherits from Δ_∞ , is a finite Coxeter complex. In fact, if we identify Σ with $\Sigma(W, V)$ as we have done before, then Σ_∞ is isomorphic to $\Sigma(\overline{W}, V)$. We will call Σ_∞ an *apartment* of Δ_∞ .

Theorem. Δ_∞ is a spherical building. Its apartments are in 1-1 correspondence with those of X .

PROOF: The proof of Lemma 3 showed that any two elements F, F' of Δ_∞ are contained in an apartment Σ_∞ . Since the latter is a simplicial complex and is closed under passage to faces, it follows that (a) for any $F \in \Delta_\infty$, the poset $(\Delta_\infty)_{\leq F}$ is isomorphic to the set of subsets of a finite set; and (b) any two elements of Δ_∞ have a greatest lower bound. Thus Δ_∞ is a simplicial complex (cf. Appendix to Chapter I), and each Σ_∞ is a subcomplex. Moreover, in the course of proving this we have already verified the building axioms (B0) and (B1). To complete the proof that Δ_∞ is a building, we will prove the variant (B2'') of (B2).

Suppose $E = |\Sigma|$ and $E' = |\Sigma'|$ are two apartments of X such that Σ_∞ and Σ'_∞ have a common chamber. This means that there are sectors $\mathfrak{C} \subset E$ and $\mathfrak{C}' \subset E'$ such that $\mathfrak{C}_\infty = \mathfrak{C}'_\infty$. By Lemma 4 above, \mathfrak{C} and \mathfrak{C}' have a common subsector; in particular, $E \cap E' \neq \emptyset$. Let $\phi : E' \rightarrow E$ be the isomorphism which fixes this intersection pointwise. Then ϕ induces an isomorphism $\phi_\infty : \Sigma'_\infty \rightarrow \Sigma_\infty$, and we will show that ϕ_∞ fixes every simplex $F \in \Sigma_\infty \cap \Sigma'_\infty$. Choose any $x \in E \cap E'$. Then $x * F$ is the unique conical cell in E based at x with F as its face at infinity, and similarly for E' . So $x * F \subseteq E \cap E'$. Hence ϕ fixes $x * F$ pointwise, and ϕ_∞ therefore fixes F . This completes the proof that Δ_∞ is a spherical building.

Continuing with the same notation, suppose that $\Sigma_\infty = \Sigma'_\infty$. Then the previous paragraph shows that $x * F \subseteq E \cap E'$ for any $F \in \Sigma_\infty = \Sigma'_\infty$, hence $E = E'$. The function $E \mapsto \Sigma_\infty$ is therefore a bijection from the set of apartments of X to the set of apartments of Δ_∞ . \square

We will call Δ_∞ the *building at infinity*. Its geometric realization is a union of spheres, one for each apartment E of X .

EXERCISES

1. Show that there is a bijection $|\Delta_\infty| \approx X_\infty$. [HINT: For any apartment

$E = |\Sigma|$, you can use “radial projection” from any point $x \in E$ to get a bijection from the unit sphere centered at x to the “sphere at infinity” E_∞ , where the latter is the set of ideal points e with a representative ray $\tau \subset E$. Deduce that there is a bijection $|\Sigma_\infty| \approx E_\infty$, which is independent of x .]

2. The following questions are deliberately stated in a vague way, in order to give you something to think about. We will return to them in the next chapter (§VII.2A), where a reference will be given for some partial answers.

Is there a reasonable way to topologize $X \amalg X_\infty$? Under what hypotheses will this yield a compactification of X , with X as an open subset? Should one expect the bijection in Exercise 1 to be a homeomorphism? [Start by thinking about the case of a tree, such as the one pictured in §V.8B. That tree was drawn with its edges getting smaller and smaller, so as to suggest the possibility of compactifying it.]

9C Type-preserving maps

Since buildings are labellable, it makes sense to ask whether a map between subcomplexes of X or of Δ_∞ is type-preserving. The following result will be used when we look at type-preserving automorphism groups below.

Proposition. *Let $\phi : E \rightarrow E'$ be a type-preserving isomorphism between apartments $E = |\Sigma|$ and $E' = |\Sigma'|$ of X . Then the induced isomorphism $\phi_\infty : \Sigma_\infty \rightarrow \Sigma'_\infty$ is type-preserving. In particular, a type-preserving automorphism of Δ induces a type-preserving automorphism of Δ_∞ .*

PROOF: Assume first that E and E' have a common sector \mathfrak{C} and that ϕ is the isomorphism that fixes $E \cap E'$ pointwise. Then ϕ_∞ fixes \mathfrak{C}_∞ and all its faces, so it is type-preserving. For arbitrary E, E' , choose an apartment E'' such that $E \cap E''$ and $E' \cap E''$ each contain a sector. Let $\psi : E \rightarrow E'$ be the composite of the isomorphisms $E \rightarrow E'' \rightarrow E'$ of the type just considered. Then ψ_∞ is type-preserving. Now the automorphism $w = \psi^{-1}\phi$ of E is in the Coxeter group $W = \text{Aut}_0(\Sigma)$ of type-preserving automorphisms of Σ , and w_∞ is simply the image of w in the finite reflection group $\overline{W} = \text{Aut}_0(\Sigma_\infty)$. So w_∞ is type-preserving and hence so is $\phi_\infty = \psi_\infty w_\infty$. \square

9D Incomplete apartment systems

If \mathcal{A} is now an incomplete apartment system, it is reasonable to try to define a building at infinity by using only the apartments in \mathcal{A} . This is of some interest because incomplete apartment systems arise quite often (e.g., from BN-pairs). Let \mathcal{A}_∞ be the set of apartments Σ_∞ in Δ_∞ with $\Sigma \in \mathcal{A}$. Let $\Delta_\infty(\mathcal{A})$ be the union of these apartments Σ_∞ . It is a subcomplex of Δ_∞ , and we would like to know whether it is a building, with \mathcal{A}_∞ as system of apartments. Since Δ_∞ is already known to be a building, the only issue is whether axiom (B1) holds. The following result is immediate:

Proposition. *$\Delta_\infty(\mathcal{A})$ is a building with \mathcal{A}_∞ as system of apartments if and only if \mathcal{A} has the following property: Given apartments $E_1, E_2 \in \mathcal{A}$*

and sectors $\mathfrak{C}_i \subset E_i$ ($i = 1, 2$), there is an apartment $E \in \mathcal{A}$ containing subsectors of \mathfrak{C}_1 and \mathfrak{C}_2 . \square

We will say that \mathcal{A} is *good* if it satisfies the condition of the proposition. By thinking about trees, you can easily convince yourself that there are many good apartment systems besides the complete one, but that not all apartment systems are good.

9E BN-pairs

Suppose that X is the building $|\Delta(G, B)|$ associated to a group with a Euclidean BN-pair. Let \mathcal{A} be the corresponding apartment system (with respect to which G is strongly transitive). If \mathcal{A} is good, then we have a spherical building $\Delta_\infty(\mathcal{A})$, and the action of G on X induces an action of G on $\Delta_\infty(\mathcal{A})$. Let $E = |\Sigma|$ be the fundamental apartment of X , let \mathfrak{C} be a fixed sector in E , and let $\mathfrak{B} \subseteq G$ be the stabilizer of \mathfrak{C}_∞ .

Proposition 1. *Assume that \mathcal{A} is good. Then the action of G on the building $\Delta_\infty(\mathcal{A})$ is type-preserving and strongly transitive. Consequently, the pair of subgroups (\mathfrak{B}, N) is a BN-pair in G whose associated building is $\Delta_\infty(\mathcal{A})$.*

PROOF: The action of G on Δ_∞ is type-preserving (cf. §9C), so the same is true of the action on the subcomplex $\Delta_\infty(\mathcal{A})$. Clearly G acts transitively on \mathcal{A}_∞ . The subgroup N stabilizes Σ_∞ and acts on the latter via the quotient map $N \twoheadrightarrow W \twoheadrightarrow \overline{W}$; so N is transitive on the chambers of Σ_∞ , and G is therefore strongly transitive on $\Delta_\infty(\mathcal{A})$. In view of Chapter V, we now have a BN-pair (\mathfrak{B}, N^*) in G , where N^* is the stabilizer of Σ (and hence also of Σ_∞). Since N surjects onto the associated Weyl group \overline{W} , it is easy to check that the BN-pair axioms still hold if we replace N^* by N . \square

Two obvious questions remain: (a) How can we decide whether \mathcal{A} is good, or, equivalently, whether the set of apartments \mathcal{A}_∞ in $\Delta_\infty(\mathcal{A})$ satisfies axiom (B1)? (b) How can we compute \mathfrak{B} ? Question (a) reduces to the question of whether, for any $g \in G$, there is an apartment in \mathcal{A}_∞ containing \mathfrak{C}_∞ and $g\mathfrak{C}_\infty$. Now it is clear from the proof above that the G -action on $\Delta_\infty(\mathcal{A})$ is strongly transitive (which makes sense whether (B1) holds or not). The methods of Chapter V therefore show that \mathfrak{C}_∞ and $g\mathfrak{C}_\infty$ are contained in an apartment if and only if $g \in \mathfrak{B}N\mathfrak{B}$. Consequently:

Proposition 2. *\mathcal{A} is good if and only if $G = \mathfrak{B}N\mathfrak{B}$.* \square

To answer question (b), suppose we are given $g \in \mathfrak{B}$. Then $g\mathfrak{C}$ and \mathfrak{C} have a common subsector. In particular, gE and E have a common chamber. Since G is strongly transitive and type-preserving on X , it follows that there is an element $g' \in G$ which maps gE to E by the unique isomorphism which fixes $E \cap gE$ pointwise. Then g' is in the subgroup $\mathfrak{B}' \subset \mathfrak{B}$ consisting of those elements of G which fix some subsector of \mathfrak{C} pointwise. Now $g'g$ stabilizes E and \mathfrak{C}_∞ , so its action on E is given by an element $w \in W' =$

$\ker\{W \rightarrow \overline{W}\}$. Let T^* be the set of elements of G that fix E pointwise, and let \tilde{w} be a representative of w in N . Then we have $g'g\tilde{w}^{-1} \in T^* \subseteq \mathfrak{B}'$, whence $g \in \mathfrak{B}'\tilde{w}$. As in Chapter V, we will write $\mathfrak{B}'w$ instead of $\mathfrak{B}'\tilde{w}$ (this being independent of the choice of \tilde{w}). So what we have proven so far may be written as

$$\mathfrak{B} = \mathfrak{B}'W'. \quad (*)$$

Now let's look more carefully at \mathfrak{B}' . Let $\mathfrak{B}_0 \subseteq \mathfrak{B}'$ be the set of elements of G that fix \mathfrak{C} pointwise. We will show that

$$\mathfrak{B}' = \bigcup_{w \in W'} w\mathfrak{B}_0w^{-1}. \quad (**)$$

Note first that $w\mathfrak{B}_0w^{-1}$ is the set of elements of G which fix $w\mathfrak{C}$ pointwise. So $w\mathfrak{B}_0w^{-1}$ fixes the subsector $\mathfrak{C} \cap w\mathfrak{C}$ of \mathfrak{C} , and the right-hand side of (**) is therefore contained in the left-hand side. For the opposite inclusion, it suffices to show that any subsector \mathfrak{C}' of \mathfrak{C} contains a translate $w\mathfrak{C}$ for some $w \in W'$. Now it is easy to check that \mathfrak{C}' contains a fundamental domain for the (translation) action of W' on E ; this follows from the fact that there is a bounded fundamental domain. In particular, we can find a $w \in W'$ which maps the cone point of \mathfrak{C} into \mathfrak{C}' , whence $w\mathfrak{C} \subseteq \mathfrak{C}'$. This completes the proof of (**). As an immediate consequence of (*) and (**) we have:

Proposition 3. \mathfrak{B} is generated by \mathfrak{B}_0 and any set of representatives for W' in N . \square

EXERCISES

1. Show that there is a short exact sequence $1 \rightarrow \mathfrak{B}' \rightarrow \mathfrak{B} \rightarrow W' \rightarrow 1$.
2. Let W'' be the submonoid of W' consisting of those $w \in W'$ such that $w\mathfrak{C} \subseteq \mathfrak{C}$. [Equivalently, if we identify E with a vector space V and W' with a lattice L in V , then $W'' = L \cap \overline{\mathfrak{D}}$, where \mathfrak{D} is the direction of \mathfrak{C} .] Show that (**) remains valid if W' is replaced by W'' .

Remark. If you know about ascending HNN extensions (cf. [49], §1.2), then the situation in these exercises should look familiar. This suggests that \mathfrak{B} is, in some sense which I won't make precise, a "generalized ascending HNN extension with base group \mathfrak{B}_0 ".

9F Example

Let K be a field with a discrete valuation, and consider the Euclidean BN-pair in $G = \mathrm{SL}_n(K)$ constructed in §V.8. Its Weyl group W is the Euclidean reflection group studied in §1F of the present chapter. We already know that there is a second BN-pair in G , obtained by forgetting that K has a valuation and applying §V.5; its Weyl group is the symmetric group on n letters, which is the finite reflection group \overline{W} associated to W . The following result is therefore not surprising:

Proposition. *Let X be the Euclidean building $|\Delta(G, B)|$ associated to $G = \mathrm{SL}_n(K)$.*

- (1) *There is a sector \mathfrak{C} in the fundamental apartment $E = |\Sigma|$ such that the stabilizer \mathfrak{B} of \mathfrak{C}_∞ is the upper triangular subgroup of G .*
- (2) *The apartment system \mathcal{A} associated to (G, B, N) is good. The subcomplex $\Delta_\infty(\mathcal{A})$ of Δ_∞ is therefore isomorphic to the spherical building associated to G in §V.5.*
- (3) *\mathcal{A} is the complete apartment system if and only if K is complete with respect to the given valuation.*

SKETCH OF THE PROOF: Identify the fundamental apartment Σ with the complex $\Sigma(W, V)$ studied in §1F above. [Recall that we gave an explicit way of making this identification.] Then $E = |\Sigma|$ is identified with V . As “fundamental sector” \mathfrak{C} we take the subset of V defined by $x_1 < \cdots < x_n$. Its closure is a subcomplex of E whose vertices, from the point of view of A -lattices, are the classes $[[\pi^{a_1} e_1, \dots, \pi^{a_n} e_n]]$ with $a_1 \leq \cdots \leq a_n$. The group \mathfrak{B}_0 which fixes \mathfrak{C} pointwise is therefore given by

$$\mathfrak{B}_0 = \bigcap_{d \in D} d \cdot \mathrm{SL}_n(A) \cdot d^{-1},$$

where D is the set of matrices in $\mathrm{GL}_n(K)$ of the form $\mathrm{diag}(\pi^{a_1}, \dots, \pi^{a_n})$ with $a_1 \leq \cdots \leq a_n$. An easy computation shows that this intersection is the upper triangular subgroup of $\mathrm{SL}_n(A)$.

Now apply the formula (**) above to get $\mathfrak{B}' = \bigcup_{t \in T} t \mathfrak{B}_0 t^{-1}$, where T is the diagonal subgroup of G . Another easy computation then shows that this union consists of all upper triangular matrices in G whose diagonal entries are units in A . Finally, (*) says that we get \mathfrak{B} by adjoining T . Thus \mathfrak{B} is indeed the full upper triangular subgroup of G .

Statement (1) is now proved, and (2) follows immediately from (1) and Propositions 1 and 2 of §9E. Turning now to (3), suppose first that K is not complete. Let \hat{K} be the completion of K , and let \hat{G} , \hat{B} , and \hat{N} be the analogues of G , B , and N over \hat{K} . Then G is dense in \hat{G} and \hat{B} is an open subgroup of \hat{G} ; it follows that $\Delta(G, B) = \Delta(\hat{G}, \hat{B})$. On the other hand, it is easy to see that \hat{G}/\hat{N} is strictly bigger than G/N , so we definitely get more apartments using $(\hat{G}, \hat{B}, \hat{N})$ than we get from (G, B, N) . The apartment system associated to (G, B, N) is therefore not complete.

Finally, suppose K is complete, and let E' be an arbitrary apartment in the complete system. We will show that $E' \in \mathcal{A}$ by constructing a $g \in G$ such that $E' = gE$, where E is the fundamental apartment. We may assume that E' contains the fundamental chamber C , in which case we will find the desired g in B .

Let $\phi : E \rightarrow E'$ be the isomorphism which fixes $E \cap E'$. In view of §6, every bounded subset of E' is contained in an apartment in \mathcal{A} . So if we exhaust E by an increasing sequence of bounded sets F_i containing C , then we can find $b_i \in B$ such that b_i maps F_i into E' by the map $\phi|_{F_i}$. I claim

that the F_i and b_i can be chosen so that $b_{i+1} \equiv b_i \pmod{\pi^i}$. Accepting this for the moment, we can easily complete the proof. For the completeness of K implies that b_i converges to some $b \in B$ as $i \rightarrow \infty$, whence $E' = bE$ and we are done. It remains to prove the claim.

By looking at the stabilizers of the vertices of E , one sees first that the F_i can be chosen so that any element of G that fixes F_i pointwise is in $\mathrm{SL}_n(A)$ and is congruent to a diagonal matrix mod π^i . In particular, for any choice of the b_i we will have $b_{i+1}^{-1}b_i$ congruent to a diagonal matrix mod π^i . Now assume inductively that b_1, \dots, b_i have been chosen and that they satisfy the required congruences. Let b be any element of B such that $b|_{F_{i+1}} = \phi|_{F_{i+1}}$. Then there is a diagonal matrix $t \in \mathrm{SL}_n(A)$ such that $b^{-1}b_i \equiv t \pmod{\pi^i}$. Since t fixes E pointwise, we can complete the inductive step by setting $b_{i+1} = bt$. \square

VII

Applications to Group Cohomology

This final chapter is a survey, without proofs, of a few of the applications of buildings to the cohomology theory of groups. A prerequisite for this chapter is a familiarity with the basic facts about group cohomology, as given for instance in [17]. I will also use some algebraic topology (fundamental group, covering spaces, homology theory of manifolds, etc.).

A less serious prerequisite involves the theory of algebraic groups. In order to make accurate statements, I will need to use standard terminology about linear algebraic groups. But I hope that it is possible to get the flavor of the results by thinking of familiar examples. For example, if you see “Let G be a linear algebraic group” you can think “Let $G = \mathrm{SL}_n$ ”. Symbols like $G(\mathbf{Q})$ or $G(\mathbf{R})$ can then be interpreted as $\mathrm{SL}_n(\mathbf{Q})$ or $\mathrm{SL}_n(\mathbf{R})$. Any technical assumptions about G (semisimplicity, simple connectivity, etc.) can be ignored, since they are all satisfied by the example $G = \mathrm{SL}_n$.

For the benefit of the reader who is not content to think about SL_n , there is an appendix to this chapter which defines most of the terms that are used.

1 Arithmetic Groups Over the Rationals

1A Definition

An arithmetic group, roughly speaking, is a group of integral matrices defined by polynomial equations. For example, $\mathrm{SL}_n(\mathbf{Z})$ is an arithmetic group, defined by the single equation $\det(a_{ij}) = 1$. For the precise definition of “arithmetic group”, start with a linear algebraic group G defined over \mathbf{Q} (e.g., $G = \mathrm{SL}_n$). We can think of G as a subgroup of GL_n for some n , defined by polynomial equations (with rational coefficients) in the n^2 matrix entries. The rational matrices satisfying the given equations then form a group $G(\mathbf{Q})$ (e.g., $\mathrm{SL}_n(\mathbf{Q})$). And for any extension field $K \supseteq \mathbf{Q}$, the matrices in $\mathrm{GL}_n(K)$ satisfying the defining equations form a group $G(K)$ (e.g., $\mathrm{SL}_n(\mathbf{R})$, $\mathrm{SL}_n(\mathbf{C})$, $\mathrm{SL}_n(\mathbf{Q}_p)$, etc.).

We can also consider the invertible integral matrices satisfying the defining equations. These form a group $G(\mathbf{Z}) = G(\mathbf{Q}) \cap \mathrm{GL}_n(\mathbf{Z})$, which is said

to be an *arithmetic group*. More generally, suppose $\Gamma \subseteq G(\mathbf{Q})$ is a subgroup which is *commensurable* with $G(\mathbf{Z})$, by which we mean that the intersection $\Gamma \cap G(\mathbf{Z})$ is of finite index in both Γ and $G(\mathbf{Z})$; then Γ is also said to be arithmetic. For example, if $\Gamma \subseteq \mathrm{SL}_n(\mathbf{Z})$ is the subgroup consisting of matrices which are congruent to the identity matrix mod m for some integer $m \geq 2$, then Γ has finite index in $\mathrm{SL}_n(\mathbf{Z})$ and hence is arithmetic. For future reference, we remark that Γ is torsion-free if $m \geq 3$ (cf. [17], §II.4, Exercise 3).

Technical Remark. If you have read the appendix, you might legitimately object to the notation $G(\mathbf{Z})$. For G is only assumed to be defined over \mathbf{Q} , and \mathbf{Z} is not a \mathbf{Q} -algebra. Indeed, we were only able to define the group $G(\mathbf{Z})$ above because we assumed we were given a specific embedding of G in a linear group GL_n , and a different embedding can lead to a different group $G(\mathbf{Z})$. It turns out, however, that the new $G(\mathbf{Z})$ is commensurable with the old one. So the notion of “arithmetic subgroup of $G(\mathbf{Q})$ ” is well-defined in spite of our abuse of notation.

1B The symmetric space

The way to get homological information about an arithmetic group Γ is to view it as a subgroup of $L = G(\mathbf{R})$. The latter is a closed subgroup of $\mathrm{GL}_n(\mathbf{R})$, hence it is a locally compact topological group. In fact, L is known to be a Lie group and to have only finitely many connected components. And Γ is a discrete subgroup. (It suffices to verify this assertion for $\Gamma = G(\mathbf{Z})$, in which case it is an immediate consequence of the fact that \mathbf{Z} is discrete in \mathbf{R} .) The significance of having Γ embedded as a discrete subgroup of a Lie group is that it enables us to exhibit an Eilenberg–MacLane space of type $K(\Gamma, 1)$ for computing the cohomology of Γ , provided Γ is torsion-free.

The starting point for constructing a $K(\Gamma, 1)$ is the existence of a contractible manifold X associated to L , on which L acts by diffeomorphisms. If $G = \mathrm{SL}_2$, for example, then X is the hyperbolic plane, which we may take to be the upper half-plane with $\mathrm{SL}_2(\mathbf{R})$ acting by linear fractional transformations (cf. §II.2C). In general, X can be constructed as the homogeneous space L/H , where H is a maximal compact subgroup of L . (Such an H exists and is unique up to conjugacy.)

Remark. We will soon specialize to the case where the algebraic group G is connected and semisimple. The space X is then a complete simply connected Riemannian manifold of negative curvature, and L acts by isometries. X is called the *symmetric space* associated to L . The fact that all maximal compact subgroups are conjugate to H follows from Cartan’s fixed-point theorem in this case (cf. §VI.4). One also knows that H is equal to its own normalizer in L . Hence X can be identified with the set of maximal compact subgroups of L , with L acting by conjugation.

The compactness of the subgroup H implies that the action of L on X is *proper*. This means that for every compact subset $C \subseteq X$, the set $\{g \in L : gC \cap C \neq \emptyset\}$ is a compact subset of L . The action remains proper if we restrict it to any closed subgroup of L . In particular, the discrete subgroup Γ of L acts properly on X . But then the compact subsets of Γ which occur in the definition of “proper” are finite. One easily deduces that the Γ -action satisfies the following condition, which is sometimes taken as the definition of properness for a discrete group action: For every $x \in X$, the stabilizer Γ_x of x in Γ is finite, and x has a neighborhood U such that $gU \cap U = \emptyset$ for all $g \in \Gamma - \Gamma_x$.

Suppose now that the arithmetic group Γ is torsion-free. [This assumption is relatively harmless, since we can always achieve it by passing to a subgroup of finite index; we have already seen this for $\mathrm{SL}_n(\mathbf{Z})$.] The finite stabilizers Γ_x are then trivial, and properness reduces to a familiar condition from covering space theory. Thus if we form the quotient space $Y = \Gamma \backslash X$, then X is a regular covering space of Y , with Γ as group of deck transformations. Since X is contractible, it follows that Y is an Eilenberg–MacLane space of type $K(\Gamma, 1)$. Consequently, the homology and cohomology groups of Γ are the same as those of the manifold Y .

Now it is no easy matter to actually calculate the cohomology of the manifold Y . But it is at least possible to get some qualitative results. For example, since Y is finite-dimensional, we immediately conclude that the *cohomological dimension* $\mathrm{cd} \Gamma$ is finite:

$$\mathrm{cd} \Gamma \leq d, \quad (*)$$

where $d = \dim Y = \dim L - \dim H$. This means that $H^i(\Gamma) = 0$ for $i > d$ and any coefficient module.

If Y happens to be compact, we can say a lot more. We will spell this out now, for motivation, before returning to the more typical non-compact case.

1C The compact case

Assume that Y is compact, in which case Γ is said to be *cocompact*. Then Y is a closed manifold, so it has non-zero cohomology in the top dimension d (with $\mathbf{Z}/2\mathbf{Z}$ coefficients, for instance). Thus equality holds in (*).

Another consequence of compactness is that the groups $H^i(\Gamma, M)$ are finitely generated whenever the coefficient module M is finitely generated as an abelian group; for one knows that $H^i(Y, M)$ is finitely generated. And if we use the triangulability of Y , then we can deduce a stronger homological finiteness property of Γ . Namely, the $\mathbf{Z}\Gamma$ -module \mathbf{Z} admits a free resolution (F_i) such that F_i is finitely generated for all i and zero for all sufficiently large i . One expresses this by saying that Γ is of “type FL”.

Note next that there is a *Poincaré duality* isomorphism between the homology and cohomology of Γ . More precisely,

$$H^i(\Gamma, M) \approx H_{d-i}(\Gamma, \Omega \otimes M)$$

for any Γ -module M , where Ω is the orientation module of X , i.e., Ω is a free abelian group whose two generators correspond to the two orientations of X . The tensor product above is over \mathbf{Z} and is given the diagonal Γ -action. We can get rid of Ω by replacing Γ by its subgroup (of index 1 or 2) consisting of the elements whose action on X is orientation-preserving.

Finally, the compactness of Y implies that Γ is a finitely presented group. This is not really a homological result, but it is usually discussed along with homological finiteness properties such as the FL property.

Unfortunately, it is relatively rare that the results above are applicable (i.e., that Γ is cocompact). If $G = \mathrm{SL}_n$, for instance, then Γ is not cocompact except in the trivial case $n = 1$. In the case of SL_2 you can see the non-cocompactness directly from the discussion in §II.2C. For if Γ were cocompact, then $W \backslash X$ would be compact, where $W = \mathrm{PGL}_2(\mathbf{Z})$; but we know that $W \backslash X$ is homeomorphic to a closed 2-simplex with one vertex removed.

1D The general case

It is remarkable that all of the properties mentioned above generalize to the case where Γ is not cocompact. Most surprising, perhaps, is that there is a generalization of the duality theorem. It is in proving this that buildings come into the picture.

To avoid uninteresting technicalities, we will state the results under the assumption that the algebraic group G is connected and semisimple. We denote by l the \mathbf{Q} -rank of G , i.e., the rank of a maximal \mathbf{Q} -split torus (cf. Appendix). The number l is significant for us for two reasons: (a) Γ is cocompact if and only if $l = 0$; thus §1C was really a discussion of the very special case $l = 0$. (b) The group $G(\mathbf{Q})$ has a BN-pair and an associated spherical building, and l is the rank of that building. In other words, l is the number of vertices of a chamber, or, equivalently, the number of generating reflections of the Weyl group W . If $G = \mathrm{SL}_n$, for example, then $l = n - 1$.

The first step in dealing with the general case is to prove that the manifold Y can be compactified by the adjunction of a boundary, i.e., Y is diffeomorphic to the interior of a compact manifold \bar{Y} with boundary. This was first proved by Raghunathan [41]. The inclusion $Y \hookrightarrow \bar{Y}$ is a homotopy equivalence, so \bar{Y} is still a $K(\Gamma, 1)$ manifold. This implies as above that Γ is finitely presented and of type FL. Raghunathan's proof, however, yields no information about the boundary $\partial\bar{Y}$ that is adjoined to Y , so we get no further homological properties of Γ . In particular, we do not get a calculation of $\mathrm{cd} \Gamma$ or a duality theorem yet.

Borel and Serre [14] give a more explicit construction of \bar{Y} . They in fact work directly with X (independent of any particular Γ) and adjoin a boundary to it. The construction is canonical enough that the action of $G(\mathbf{Q})$ on X extends to the resulting manifold \bar{X} with boundary, and the action of any arithmetic subgroup Γ is still proper and yields a compact

quotient. This quotient is then the desired \bar{Y} when Γ is torsion-free. If $G = \mathrm{SL}_2$, for example, \bar{X} is obtained from X by adjoining a disjoint union of lines, one for each cusp that you see in the pictures in §II.2C. If you have trouble visualizing this, see [45], p. 216, for a picture. [If you have trouble visualizing *any* example of a 2-dimensional manifold whose boundary is a disjoint union of infinitely many lines, you can draw one yourself: First draw a picture of an infinite tree which branches at every vertex. Now trace over that tree using a marker with a very wide tip. You will then have a picture of a surface whose boundary consists of infinitely many lines.]

The crucial feature of the Borel–Serre construction is that one is able to understand the algebraic topology of the boundary $\partial\bar{X}$: It is homotopy equivalent to the spherical building Δ associated to the BN-pair in $G(\mathbf{Q})$. The idea behind the proof of this is that $\partial\bar{X}$ is constructed as a disjoint union of contractible pieces e_P , indexed by the proper parabolic subgroups $P \subset G(\mathbf{Q})$. These pieces fit together in a manner that reflects the inclusion relations among the parabolic subgroups, and the desired homotopy equivalence then follows from a consideration of nerves of covers.

In view of the Solomon–Tits theorem (§§IV.5 and IV.6), we now know that $\partial\bar{X}$ has the homotopy type of a bouquet of $(l-1)$ -spheres. This leads to the calculation of $\mathrm{cd} \Gamma$ and to the duality theorem. A detailed explanation of the method can be found in [17], §§VIII.7–10, so I will be brief. Let H_c^* denote cohomology with compact supports and let \tilde{H}_* denote reduced homology. We take \mathbf{Z} coefficients in both cases. Combining Poincaré–Lefschetz duality and the contractibility of \bar{X} , one finds

$$H_c^i(\bar{X}) \approx H_{d-i}(\bar{X}, \partial\bar{X}) \approx \tilde{H}_{d-i-1}(\partial\bar{X}).$$

Hence $H_c^i(\bar{X}) = 0$ unless $i = d - l$, and $H_c^{d-l}(\bar{X})$ is free abelian. Since $H_c^*(\bar{X}) \approx H^*(\Gamma, \mathbf{Z}\Gamma)$, one concludes, first, that

$$\mathrm{cd} \Gamma = d - l.$$

The point here is that the cohomological dimension of a group Γ of type FL can be computed as the top dimension in which $H^*(\Gamma, \mathbf{Z}\Gamma)$ is non-trivial.

In the present case, the top dimension is the *only* dimension in which $H^*(\Gamma, \mathbf{Z}\Gamma)$ is non-trivial. Using this, together with the fact that the non-trivial cohomology group is \mathbf{Z} -torsion-free, one deduces that Γ satisfies “Bieri–Eckmann duality”:

$$H^i(\Gamma, M) \approx H_{d-l-i}(\Gamma, D \otimes M)$$

for any Γ -module M and any i . Here D , the “dualizing module”, is a fixed Γ -module, independent of M . In the present situation, D is simply the Γ -module $H_c^{d-l}(\bar{X})$; it is isomorphic to $\tilde{H}_{l-1}(\Delta) \otimes \Omega$, where Ω is the orientation module of X as in §1C.

If $l = 0$, then $\tilde{H}_{l-1}(\Delta) = \tilde{H}_{-1}(\emptyset) = \mathbf{Z}$, so $D = \Omega$ and Bieri–Eckmann duality reduces to Poincaré duality. If $l > 0$, on the other hand, D is a free abelian group of infinite rank.

To summarize, we have:

Theorem. *Let G be a connected semisimple linear algebraic group defined over \mathbf{Q} . Let d be the dimension of the symmetric space associated to $G(\mathbf{R})$, and let l be the \mathbf{Q} -rank of G . Let Γ be a torsion-free arithmetic subgroup of $G(\mathbf{Q})$. Then Γ is finitely presented and of type FL and is a $(d - l)$ -dimensional duality group. It is a Poincaré duality group if and only if $l = 0$, i.e., if and only if Γ is cocompact.*

Remarks

1. I have said practically nothing about the actual construction of \bar{X} , which is extremely difficult. Grayson ([29], [30]) has given an alternate approach which avoids some of the technical problems faced by Borel and Serre. Instead of explicitly attaching a boundary to X , he finds his \bar{X} inside of X . In other words, he constructs the sort of manifold one would get from the Borel–Serre \bar{X} by removing an open collar neighborhood of the boundary.

2. The theorem generalizes to an arbitrary linear algebraic group defined over \mathbf{Q} , but one has to define the integers d and l slightly differently in the general case.

1E Virtual notions

One says that a group “virtually” has a certain property if a subgroup of finite index has that property. It is sometimes convenient to use this language in order to avoid the assumption that Γ is torsion-free. For example, any arithmetic subgroup $\Gamma \subseteq G(\mathbf{Q})$ is “virtually torsion-free”, and we can speak of its “virtual cohomological dimension” $\text{vcd } \Gamma$; this is the common cohomological dimension $d - l$ of its torsion-free subgroups of finite index. Similarly, we say that Γ is “virtually of type FL”, or that Γ is of “type VFL”, because it has a subgroup of finite index which is of type FL. Finally, we say that Γ is a “virtual duality group”. Note that we can dispense with “virtual” when talking about finite presentation: Γ is itself finitely presented since it has a finitely presented subgroup of finite index.

For our canonical example of SL_n , one has $d = n(n + 1)/2 - 1$ and $l = n - 1$, so

$$\text{vcd}(\text{SL}_n(\mathbf{Z})) = \frac{n(n - 1)}{2}.$$

There is an easy way to remember this result—it says that $\text{vcd}(\text{SL}_n(\mathbf{Z}))$ is equal to the “obvious” lower bound on this vcd that one gets by looking at the strict upper triangular subgroup of $\text{SL}_n(\mathbf{Z})$.

In order to explain this, we need to recall some facts about solvable groups. Let Γ be an abstract solvable group. Choose a normal series

$$1 = \Gamma_0 \triangleleft \Gamma_1 \triangleleft \cdots \triangleleft \Gamma_r = \Gamma$$

with abelian quotients Γ_i/Γ_{i-1} , and set

$$h = \sum_{i=1}^r \dim_{\mathbf{Q}} \mathbf{Q} \otimes (\Gamma_i/\Gamma_{i-1}).$$

Then h is independent of the choice of normal series. It is called the *Hirsch rank* of Γ , and it is closely related to the homological and cohomological dimension of Γ . In particular,

$$h \leq \text{cd } \Gamma \leq h + 1$$

if Γ is torsion-free (cf. [11], §7.3).

Returning now to the strict upper triangular group, it is a torsion-free nilpotent group of Hirsch rank $n(n-1)/2$, whence the “obvious” inequality $\text{vcd}(\text{SL}_n(\mathbf{Z})) \geq n(n-1)/2$.

Here’s another example to illustrate this principle. Consider the group $\text{SL}_n(\mathbf{Z}[1/p])$, where p is a prime number. (This is not arithmetic, but we are about to enlarge our framework so as to allow groups like this.) Its full upper triangular subgroup is a virtually torsion-free solvable group of Hirsch rank $n(n-1)/2 + n - 1$, so we get

$$\text{vcd}(\text{SL}_n(\mathbf{Z}[1/p])) \geq \frac{n(n-1)}{2} + n - 1.$$

We will see in the next section how to use a Euclidean building to prove that equality holds.

2 S-Arithmetic Groups

Let S be a finite set of prime numbers, and let $\mathbf{Z}_S \subset \mathbf{Q}$ be the ring of rational numbers a/b ($a, b \in \mathbf{Z}$) such that the primes dividing b are in S . Thus the elements of \mathbf{Z}_S are “integral except possibly at S ”. If we go back to the definition of “arithmetic group” and replace \mathbf{Z} by \mathbf{Z}_S everywhere, then we get the notion of “ S -arithmetic group”. For example, $\text{SL}_n(\mathbf{Z}[1/p])$ is an S -arithmetic subgroup of $\text{SL}_n(\mathbf{Q})$, with $S = \{p\}$. In this section we will indicate how Borel and Serre [15] extend their results about arithmetic groups to the S -arithmetic case.

2A A p -adic analogue of the symmetric space

Fix a prime number p , and let L now denote the “ p -adic Lie group” $G(\mathbf{Q}_p)$. Assume that the algebraic group G is simply connected and absolutely almost simple. These assumptions guarantee that L admits a Euclidean BN-pair, analogous to the one we have studied for $G = \text{SL}_n$. We therefore have a Euclidean building X on which L acts by type-preserving simplicial automorphisms. X is a locally finite simplicial complex and is contractible (by §IV.6 or §VI.3). Its dimension is the \mathbf{Q}_p -rank of G (or, more precisely, of the linear algebraic group over \mathbf{Q}_p obtained from G by extension of scalars). Recall that this dimension is $n - 1$ for the case $G = \text{SL}_n$. The stabilizers of the simplices are compact open subgroups of L , and it follows easily that the action of L on X is proper.

For applications to the cohomology of discrete groups, we will want to know $H_c^*(X)$. This is computed by Borel and Serre [15], using a method remarkably similar to the method used for the symmetric space associated to $G(\mathbf{R})$. The first step is to embed X as a dense open subspace of a compact contractible space $\bar{X} = X \cup \partial X$. The compact space ∂X that is adjoined to X is, as a set, the geometric realization of the spherical building at infinity. The exercises in §VI.9B hinted at the possibility of doing this and also suggested that the topology on ∂X should *not* be expected to be the usual simplicial topology.

If X is 1-dimensional, for example, then it is a tree and \bar{X} is its endpoint compactification. Thus ∂X is the space of ends of X in this case; it is a Cantor set, whose points are in 1-1 correspondence with the vertices of the (0-dimensional) spherical building at infinity. If this spherical building were given the simplicial topology, however, then it would be discrete.

The significance of the Borel–Serre compactification \bar{X} is that it enables one to compute $H_c^*(X)$. Using a suitable cohomology theory (e.g., Alexander–Spanier cohomology), one finds

$$H_c^i(X) \approx H^i(\bar{X}, \partial X) \approx H^{i-1}(\partial X).$$

Borel and Serre go on to prove an analogue of the Solomon–Tits theorem for the Alexander–Spanier cohomology of the compactly topologized spherical building ∂X : This cohomology is zero except in the top dimension (which is $\dim X - 1$), and it is free abelian in that dimension. The end result, then, is that $H_c^i(X)$ vanishes for $i \neq \dim X$ and is free abelian for $i = \dim X$.

2B Cohomology of S -arithmetic groups: Method 1

In §1 our emphasis was on using *proper* actions of discrete groups to get homological information. We will return to that point of view in §2C below. But first, for the sake of variety, we will show how to get the same kind of information from an action that is not proper. The method we will follow here is based on [43] and [20]. To keep the discussion as simple as possible, we begin with the familiar case $G = \mathrm{SL}_n$, and we assume that S is a singleton $\{p\}$. Thus an S -arithmetic group, for the moment, is simply a subgroup of $\mathrm{SL}_n(\mathbf{Q})$ commensurable with $\mathrm{SL}_n(\mathbf{Z}[1/p])$.

Let $\Gamma = \mathrm{SL}_n(\mathbf{Z}[1/p])$, viewed as a subgroup of $L = \mathrm{SL}_n(\mathbf{Q}_p)$. Note that Γ is not discrete in L ; in fact, it is dense in L . But we can still consider the (non-proper) simplicial action of Γ on the Euclidean building X associated to L . Because of the density of Γ (and the fact that the stabilizers of the simplices are open in L), a fundamental domain for the L -action on X will still be a fundamental domain for the Γ -action. Hence Γ has a closed chamber \bar{C} as fundamental domain. Moreover, the stabilizers of the faces of C are commensurable with $\mathrm{SL}_n(\mathbf{Z})$. For example, if v is the vertex corresponding to the standard lattice, then we know that the stabilizer of v in L is $\mathrm{SL}_n(\mathbf{Z}_p)$, where \mathbf{Z}_p is the ring of p -adic integers, hence the stabilizer of v in Γ is $\mathrm{SL}_n(\mathbf{Z}[1/p]) \cap \mathrm{SL}_n(\mathbf{Z}_p) = \mathrm{SL}_n(\mathbf{Z})$. Thus the stabilizers are not

finite as they would be in a proper action, but they are groups which are known to have good finiteness properties.

Now let Γ be a torsion-free subgroup of $\mathrm{SL}_n(\mathbf{Q})$ commensurable with $\mathrm{SL}_n(\mathbf{Z}[1/p])$, e.g., a torsion-free subgroup of $\mathrm{SL}_n(\mathbf{Z}[1/p])$ of finite index. Then Γ acts on X with compact quotient and torsion-free arithmetic stabilizers. Since torsion-free arithmetic groups are finitely presented and of type FL, it follows that Γ is finitely presented (cf. [18]) and of type FL ([43], Proposition 11).

To calculate $\mathrm{cd} \Gamma$ and prove duality, we use the equivariant cohomology spectral sequence for (Γ, X) with coefficients in $\mathbf{Z}\Gamma$ ([17], VII.7.10). All the stabilizers are duality groups of the same dimension $m = \mathrm{vcd}(\mathrm{SL}_n(\mathbf{Z}))$, so the spectral sequence is concentrated on the line $q = m$. Moreover, one can calculate $E_1^{*,m}$ by the method of [20], §§2 and 3, and one finds that it is $C_c^*(X) \otimes D$; here C_c^* denotes simplicial cochains with compact support, and D is the dualizing module for the torsion-free arithmetic subgroups of $\mathrm{SL}_n(\mathbf{Z})$, i.e., $D = \check{H}_{l-1}(\Delta) \otimes \Omega$ in the notation of §1. In view of the calculation of $H_c^*(X)$ stated in §2A, the spectral sequence collapses at E_2 and gives the following result: $H^*(\Gamma, \mathbf{Z}\Gamma)$ is concentrated in dimension $\mathrm{vcd}(\mathrm{SL}_n(\mathbf{Z})) + \dim X$, and in that dimension it is the Γ -module $H_c^{\dim X}(X) \otimes D$. Thus Γ is a duality group, and we have calculated its dimension and its dualizing module. In particular,

$$\mathrm{vcd}(\mathrm{SL}_n(\mathbf{Z}[1/p])) = \frac{n(n-1)}{2} + n - 1.$$

The method works equally well if S consists of more than one prime. Simply pick some $p \in S$, let Γ act on the corresponding X , and note that the stabilizers are $(S - \{p\})$ -arithmetic. So the analysis can be done by induction on the number of primes in S . The method also works equally well if SL_n is replaced by any G which is simply connected and absolutely almost simple. To state the result, let X_p be the Euclidean building associated to $G(\mathbf{Q}_p)$, let $d_p = \dim X_p$, and let $D_p = H_c^{d_p}(X_p)$. It is convenient to introduce a fictitious prime ∞ and to set X_∞ equal to the symmetric space associated to $G(\mathbf{R})$. Let $d_\infty = \dim X_\infty - l$, where l is the \mathbf{Q} -rank of G , and let $D_\infty = H_c^{d_\infty}(\bar{X}_\infty)$. Let $S' = S \cup \{\infty\}$, and set

$$d = \sum_{p \in S'} d_p \quad \text{and} \quad D = \bigotimes_{p \in S'} D_p.$$

Then we have:

Theorem. *Any torsion-free S -arithmetic subgroup of $G(\mathbf{Q})$ is finitely presented and of type FL and is a duality group of dimension d with dualizing module D .*

Note that the proof shows more than what was stated—it gives a way of describing each D_p in terms of a spherical building.

2C Cohomology of S -arithmetic groups: Method 2

We now sketch the method actually used by Borel and Serre [15] to prove the theorem stated above. Instead of letting the torsion-free S -arithmetic group Γ act on the various X_p one at a time, they let it act on them simultaneously. More precisely, let $L_p = G(\mathbf{Q}_p)$ for $p \in S'$, where $\mathbf{Q}_\infty = \mathbf{R}$. Let $L = \prod_{p \in S'} L_p$. Then Γ can be embedded diagonally in the locally compact group L , and it is discrete in L . The point here is that \mathbf{Z}_S is a discrete subring of $\prod_{p \in S'} \mathbf{Q}_p$, since a sequence of non-zero elements of \mathbf{Z}_S which converges to 0 p -adically for all $p \in S$ will not converge to 0 in \mathbf{R} . Now Γ acts properly on the contractible space $X = \bar{X}_\infty \times \prod_{p \in S} X_p$. As in the arithmetic case, the quotient $Y = \Gamma \backslash X$ is a compact $K(\Gamma, 1)$ -space.

A suitable triangulation theorem now implies that Γ is finitely presented and of type FL. Moreover, letting d and D be as above, we can apply the Künneth theorem to calculate that $H_c^*(X) = H^*(\Gamma, \mathbf{Z}\Gamma)$ is concentrated in dimension d and is isomorphic to D in that dimension. Thus Γ is a d -dimensional duality group with dualizing module D .

Remark. Borel and Serre prove a more general theorem than the one stated above. First, they work over an arbitrary algebraic number field F , not just \mathbf{Q} . Their L involves the groups $G(\hat{F})$ for various completions \hat{F} of F , which may include several copies of \mathbf{R} , several copies of \mathbf{C} , and several p -adic completions. Secondly, their hypothesis on G is weaker than the one stated above. All they assume is that G is a linear algebraic group (defined over F) such that the connected component of the identity is reductive.

2D The non-reductive case

The finiteness properties proven by Borel and Serre in the reductive case hold for some non-reductive groups, but not for all. Consider, for example, the following subgroups of the 2×2 upper triangular group:

$$G = \begin{pmatrix} 1 & * \\ & * \end{pmatrix}, \quad G_0 = \begin{pmatrix} 1 & * \\ & 1 \end{pmatrix}.$$

Then G is not reductive because it has the connected unipotent group G_0 as a normal subgroup. (See the appendix, §§G and I, for the relevant definitions.) Nevertheless, it is not hard to show that $G(\mathbf{Z}[1/p])$ is finitely presented and of type VFL. On the other hand, $G_0(\mathbf{Z}[1/p])$ is isomorphic to the additive group $\mathbf{Z}[1/p]$, and so it is not even finitely generated. Another interesting example is the 3×3 group

$$G_1 = \begin{pmatrix} 1 & * & * \\ & * & * \\ & & 1 \end{pmatrix};$$

one can show that $G_1(\mathbf{Z}[1/p])$ is finitely generated but not finitely presented.

The groups G_0 and G_1 are part of an infinite sequence of groups whose study was initiated by H. Abels [1]. The next one in the sequence is

$$G_2 = \begin{pmatrix} 1 & * & * & * \\ & * & * & * \\ & & * & * \\ & & & 1 \end{pmatrix}.$$

In general, G_n is the subgroup of GL_{n+2} consisting of upper triangular matrices such that the diagonal entries in the upper left-hand corner and lower right-hand corner are 1. (The subscript n indicates the rank of the torus consisting of the diagonal matrices in G_n .)

We have already noted that $G_0(\mathbf{Z}[1/p])$ is not finitely generated, whereas $G_1(\mathbf{Z}[1/p])$ is finitely generated but not finitely presented. And Abels [1] proved that $G_2(\mathbf{Z}[1/p])$ is finitely presented. In order to describe the situation for arbitrary n , we need to introduce finiteness conditions F_n that generalize finite generation and finite presentation.

We will say that a group Γ is of type F_1 if it is finitely generated. For any $n \geq 2$, we will say that Γ is of type F_n if it is finitely presented and if the $\mathbf{Z}\Gamma$ -module \mathbf{Z} (with trivial Γ -action) admits a free resolution (P_i) with P_i finitely generated for $i \leq n$. If Γ is of type F_n for all n , then we say that Γ is of type F_∞ ; this is equivalent to saying that Γ is finitely presented and that there is a resolution as above with P_i finitely generated for all i ([17], VIII.4.5). Finally, we agree that every group Γ is of type F_0 . Let $\phi(\Gamma)$ be the largest n ($0 \leq n \leq \infty$) such that Γ is of type F_n . We call $\phi(\Gamma)$ the *finiteness length* of Γ . It is easy to see that any group of type VFL is of type F_∞ ([17], VIII.5.1). Hence $\phi(\Gamma) = \infty$ if Γ is arithmetic, and the same is true if Γ is S -arithmetic and G is reductive.

Let's return now to our sequence of non-reductive S -arithmetic groups $\Gamma_n = G_n(\mathbf{Z}[1/p])$. The results stated above can be expressed by saying that $\phi(\Gamma_0) = 0$, $\phi(\Gamma_1) = 1$, and $\phi(\Gamma_2) \geq 2$. Abels and Brown [4] generalized these results by showing that $\phi(\Gamma_n) = n$ for all n . A slightly different proof was later given by Brown [19]. Both proofs involve an analysis of the action of Γ_n on the Euclidean building X associated to $SL_{n+2}(\mathbf{Q}_p)$. (Recall that $GL_{n+2}(\mathbf{Q}_p)$ acts on this building, so Γ_n also acts.) As in §2B above, the stabilizers are arithmetic groups. The problem, however, is that the quotient is not compact; so it takes some work to deduce finiteness properties (or the lack thereof) from the action.

At the moment, this application of buildings is not understood in any systematic way. In other words, one does not know how to find a suitable building to use for the study of the finiteness properties of an arbitrary S -arithmetic group.

Remark. Given the action of Γ_n on the $(n+1)$ -dimensional building X , we can interpret the result that $\phi(\Gamma_n) = n$ as saying that Γ_n just barely fails to be of type F_∞ . For one has the following general result, which is a consequence of [19], Theorems 2.2 and 3.2: Suppose a group Γ acts on a

d -dimensional contractible complex X . If the stabilizer of every simplex is of type F_∞ , then Γ is of type F_∞ if and only if it is of type F_d .

It should be mentioned, finally, that the F_1 and F_2 conditions (i.e., finite generation and finite presentation) are understood for an arbitrary S -arithmetic group Γ . The results, which are due to Kneser, Borel–Tits, and Abels, are too complicated to state here. See the introduction to [2] for a survey.

3 Cohomological Dimension of Linear Groups

A special case of the results of §2 is that $\text{vcd}(\text{SL}_n(\mathbf{Z}_S)) < \infty$ for any n and S . The proof of this does not require the full force of the arguments sketched in §2, as long as we do not care about the precise value of the virtual cohomological dimension. In particular, we need the proper action of $\text{SL}_n(\mathbf{Z}_S)$ on the contractible finite-dimensional space $X = \prod_{p \in S'} X_p$, but we do not need the spaces \bar{X}_p . As a consequence of this result, we have the following theorem, first pointed out by Serre ([43], Théorème 5):

Theorem. *Let F be an algebraic extension of \mathbf{Q} and let Γ be an arbitrary finitely generated subgroup of $\text{GL}_n(F)$. Then $\text{vcd } \Gamma < \infty$.*

To prove this, we may assume that F is finite over \mathbf{Q} , and then we can easily reduce to the case where $F = \mathbf{Q}$. [An n -dimensional vector space over F is a finite-dimensional vector space over \mathbf{Q} .] Then $\Gamma \subseteq \text{GL}_n(A)$ for some finitely generated subring $A \subset \mathbf{Q}$, hence $\Gamma \subseteq \text{GL}_n(\mathbf{Z}_S)$ for some S . Finally, we may replace Γ by $\Gamma \cap \text{SL}_n(\mathbf{Z}_S)$ since $\det(\Gamma)$ is a finitely generated abelian group. The theorem now follows from the result about $\text{SL}_n(\mathbf{Z}_S)$ stated above.

It is natural to ask what can be said if F is not assumed to be algebraic. For example, what if F is a rational function field $\mathbf{Q}(X)$? Easy examples show that finitely generated subgroups do not necessarily have finite vcd in this case. Suppose, for instance, that $\Gamma = G(\mathbf{Z}[X])$, where G is the 2×2 matrix group defined at the beginning of §2D. Then Γ is finitely generated, but $\text{vcd } \Gamma = \infty$ because the unipotent subgroup $G_0(\mathbf{Z}[X])$ is free abelian of infinite rank.

It turns out that this example is essentially the only kind of counterexample. In other words, the failure of $\text{vcd } \Gamma$ to be finite can always be explained in terms of the unipotent subgroups of Γ . This is the content of the following theorem of Alperin and Shalen [7]:

Theorem. *Let Γ be a finitely generated subgroup of $\text{GL}_n(F)$, where F is a field of characteristic 0. Then $\text{vcd } \Gamma < \infty$ if and only if there is an upper bound on the Hirsch ranks of the unipotent subgroups of Γ .*

Recall that any unipotent subgroup U of $\text{GL}_n(F)$ is torsion-free and nilpotent by Kolchin's theorem (cf. Appendix, §G). So the Hirsch rank

of U is indeed defined and differs from $\text{cd } U$ by at most 1 (cf. §1E above). Thus “Hirsch rank” could be replaced by “cohomological dimension” in the statement of the theorem. The “only if” part is now obvious. I will say a few words about the proof of the “if” part, in order to indicate how buildings enable one to reduce the question of finite vcd to the question of finding a bound on $\text{cd } U$, where U ranges over the unipotent subgroups of Γ .

As in the proof of Serre’s theorem, the finite generation of Γ guarantees that $\Gamma \subseteq \text{GL}_n(A)$ for some finitely generated subring $A \subset F$. We may assume that F is the field of fractions of A and that $\Gamma \subseteq \text{SL}_n(A)$. The first step in the proof is pure commutative algebra. One shows that there is a finite collection $\{v_i\}$ of discrete valuations on F that can be used to test integrality of elements of A , in the following sense: Let A_i be the valuation ring of v_i and let B be the ring of algebraic integers in F ; then $A \cap \bigcap_i A_i \subseteq B$.

Let X_i be the Euclidean building associated to $\text{SL}_n(F)$ and the valuation v_i . Then Γ acts on X_i for all i . We can either analyze one of these actions at a time as in §2B or we can let Γ act on the product as in §2C. Either way, we are reduced to finding a bound on $\text{vcd } \Gamma'$, where Γ' ranges over the subgroups of Γ which stabilize a vertex in each X_i . Now the stabilizer of a vertex of X_i stabilizes an A_i -lattice in F^n , hence its characteristic polynomial has coefficients in A_i . Consequently, the characteristic polynomial of each element of Γ' has coefficients in the ring of integers B . One says that Γ' has *integral characteristic*.

Intuitively, Γ' resembles a subgroup of the arithmetic group $\text{SL}_n(B)$, and so one might hope to be able to bound $\text{vcd } \Gamma'$ by embedding Γ' as a discrete subgroup of a product $\text{SL}_n(\mathbf{R})^{r_1} \times \text{SL}_n(\mathbf{C})^{r_2}$. Now this is certainly not possible in general—unipotent groups again provide counter-examples. (Note that any unipotent group has integral characteristic.) But Alperin and Shalen, using techniques introduced by Bass for studying groups of integral characteristic, show that unipotent groups are the only obstruction. More precisely, there is a unipotent normal subgroup $U \triangleleft \Gamma'$ such that Γ'/U is a discrete subgroup of a Lie group as above. Since $\text{vcd } \Gamma' \leq \text{cd } U + \text{vcd } \Gamma'/U$, we are done by the hypothesis on the unipotent subgroups.

4 S-Arithmetic Groups Over Function Fields

We close with a discussion of the finiteness properties of some matrix groups in characteristic $\neq 0$. Let K be the function field of an irreducible projective smooth curve C defined over a finite field $k = \mathbf{F}_q$. If you don’t know what this means, you can think of C as something like a Riemann surface and K as the field of meromorphic functions on C . The canonical example is the rational function field $K = k(t)$, which corresponds to the case where C is the projective line (analogue of the Riemann sphere).

Let S be a finite non-empty set of (closed) points of C , and let $\mathcal{O}_S \subset K$ be the ring of functions which have no poles except possibly at points in S . We can also describe \mathcal{O}_S in terms of discrete valuations. Each point p of the curve C gives rise to a discrete valuation v_p on K such that $v_p(f)$ is the order to which f vanishes at p . Thus $v_p(f) < 0$ if and only if f has a pole at p , hence the valuation ring \mathcal{O}_p associated to v_p is the set of functions which do not have a pole at p . The definition of \mathcal{O}_S can now be rewritten as

$$\mathcal{O}_S = \bigcap_{p \notin S} \mathcal{O}_p.$$

Suppose, for example, that $K = k(t)$. Then the curve has a point at infinity together with one point for every irreducible polynomial in $k[t]$. If S consists only of the point at infinity, then \mathcal{O}_S is the polynomial ring $k[t]$.

Let G be a linear algebraic group defined over K . We can then define the notion of “ S -arithmetic subgroup” of $G(K)$ exactly as in §2, with \mathbb{Z}_S replaced by \mathcal{O}_S . For example, $\mathrm{SL}_n(k[t])$ is an S -arithmetic subgroup of $\mathrm{SL}_n(k(t))$ when $S = \{\infty\}$ as above.

Assume now that G is simply connected and absolutely almost simple. For each $p \in S$ we have a locally compact group $L_p = G(K_p)$, where K_p is the completion of K with respect to v_p , and we have a Euclidean building X_p on which L_p acts properly. Set

$$L = \prod_{p \in S} L_p \quad \text{and} \quad X = \prod_{p \in S} X_p.$$

As in the number field case, the group $\Gamma = G(\mathcal{O}_S)$ is a discrete subgroup of L . It therefore acts properly on the contractible space X . And, as in the number field case again, the quotient $\Gamma \backslash X$ is compact if and only if $l = 0$, where l is the K -rank of G . One can deduce that Γ is of type VFL, and hence of type F_∞ , when $l = 0$. [Note: Part of what has to be proved here is that Γ is virtually torsion-free, which is not automatic in characteristic $\neq 0$. The proof in the present case uses the compactness of $\Gamma \backslash X$, cf. [43], Théorème 4.]

If $l > 0$, on the other hand, the situation becomes different from that in the number field case, at least as far as the F_n properties are concerned. Indeed, Γ need not even be finitely generated. The simplest example is $\Gamma = \mathrm{SL}_2(k[t])$. The space X is a tree in this case, and there is a half-line which is a fundamental domain for the action. By analyzing the stabilizers along this half-line (cf. [46], §II.1.6) one can see that Γ is not finitely generated. [This is a theorem of Nagao, which had been proved earlier without the aid of the tree.] More generally, $\mathrm{SL}_n(k[t])$ acts on its $(n-1)$ -dimensional building with a sector as fundamental domain, and one suspects that this group is of type F_{n-2} but not of type F_{n-1} . We will return to this example below.

For arbitrary G , K , and S there are very few results about the finiteness properties of $\Gamma = G(\mathcal{O}_S)$. The rest of this section is a summary of what is

known. Assume throughout this discussion that $l > 0$, and set

$$d = \dim X = \sum_{p \in S} d_p,$$

where $d_p = \dim X_p$ (= the K_p -rank of G). Note that $d_p \geq l$, so we always have $d \geq 1$.

Theorem 1. *Γ is finitely generated if and only if $d \geq 2$.*

Thus finite generation fails if and only if X is a tree. For example, $\mathrm{SL}_2(\mathcal{O}_S)$ is finitely generated if and only if $\mathrm{card} S \geq 2$. For $n \geq 3$, on the other hand, $\mathrm{SL}_n(\mathcal{O}_S)$ is finitely generated for any S . See Behr [10] for further discussion of this theorem and for references.

Turning next to finite presentation, there is a lot of evidence to suggest:

Conjecture. *Γ is finitely presented if and only if $d \geq 3$.*

See Behr [10] for a list of the cases where this has been verified. They include $G = \mathrm{SL}_n$ (for arbitrary K and S). Thus $\mathrm{SL}_2(\mathcal{O}_S)$ is finitely presented if and only if $\mathrm{card} S \geq 3$, $\mathrm{SL}_3(\mathcal{O}_S)$ is finitely presented if and only if $\mathrm{card} S \geq 2$, and $\mathrm{SL}_n(\mathcal{O}_S)$ is finitely presented for any S if $n \geq 4$.

For SL_2 , one not only knows when Γ is finitely generated or finitely presented, but one in fact knows the precise finiteness length, i.e., the largest m such that Γ is of type F_m . This is given by the following theorem of Stuhler [50]:

Theorem 2. *$\mathrm{SL}_2(\mathcal{O}_S)$ has finiteness length $d - 1 = \mathrm{card} S - 1$.*

Suppose, finally, that $G = \mathrm{SL}_n$ for arbitrary n and that we are in the simplest possible case: $\mathcal{O}_S = k[t]$. Then $\Gamma = \mathrm{SL}_n(k[t])$, and X has only one factor, which is a building of dimension $d = n - 1$. We know from the results stated above that the finiteness length of $\Gamma = \mathrm{SL}_n(k[t])$ is 0 if $n = 2$, is 1 if $n = 3$, and is at least 2 if $n \geq 4$. The following result, due independently to Abels [3] and Abramenko [5], almost settles the question for arbitrary n :

Theorem 3. *For any n there is an integer N such that $\mathrm{SL}_n(\mathbb{F}_q[t])$ has finiteness length $n - 2$ if $q \geq N$.*

If $n \leq 5$, Abramenko has shown that one can take $N = 2$, i.e., there is no restriction on q . If $n \geq 6$, however, the best known value of N is $N = \max_{1 \leq i \leq n-2} \binom{n-2}{i}$, again due to Abramenko. In particular, one does not know the finiteness length of $\mathrm{SL}_6(\mathbb{F}_q[t])$ for $q = 2, 3, 4, 5$.

For a general G, K, S (with $l > 0$), the theorems and conjecture stated in this section suggest the following question:

Question. *Is the finiteness length $\phi(\Gamma)$ always equal to $d - 1$?*

It would be interesting to at least have the inequality $\phi(\Gamma) < d$, i.e., to know that Γ is not of type F_d . In view of a remark near the end of §2D above, this is equivalent to saying that Γ is not of type F_∞ .

Remark. In this section we have focused on a particular kind of finiteness property, for which the function field case seems very different from the number field case. But if one asks slightly different questions, then the two cases do not seem quite so different. In fact, Grayson's version of the Borel–Serre construction for number fields ([29], [30]), was inspired by homological finiteness results in the function field case proved by Serre for SL_2 and Quillen for SL_n . See [46], §§II.2.8 and II.2.9, and [28].

Appendix. Linear Algebraic Groups

A typical example of what we want to talk about in this appendix is the group $SL_n(k)$, where k is a field. As a set, this is defined by the equation $\det(a_{ij}) = 1$, which is a polynomial equation of degree n in the n^2 variables a_{ij} . And the group structure is given by the matrix multiplication map $SL_n(k) \times SL_n(k) \rightarrow SL_n(k)$, which is also describable by polynomials. In fact, each of the n^2 components of this map is a quadratic function of $2n^2$ variables. As this example suggests, we will be looking at groups G of the following form: G is a set defined by polynomial equations, with a group structure defined by a polynomial map.

We will see later that there are many properties of such groups which are only revealed when we pass from k to a bigger field k' . Examples are given in §§F–I below. Thus we will want to take the defining equations for G and look at their solutions over various extensions k' of k . More generally, there are reasons (although you won't see them in this appendix) for looking at solutions over k -algebras R which are not necessarily fields. [By a *k-algebra* here we mean a commutative ring with identity which comes equipped with a homomorphism $k \rightarrow R$. Since k is a field, the given homomorphism is necessarily 1-1, and we may think of R as a ring which contains k as a subring.] These considerations lead to the notion of “group scheme”.

A Group schemes

Suppose we are given a collection of polynomial equations in m variables with coefficients in a field k . For any k -algebra R , let $G(R) \subseteq R^m$ be the solution set of the given equations. Assume further that we are given m polynomials in $2m$ variables such that the map $R^m \times R^m \rightarrow R^m$ which they define sends $G(R) \times G(R)$ into $G(R)$ for every R and makes $G(R)$ a group. What we have, then, is not just a single group, but rather a group-valued functor on the category of k -algebras. [You should take a moment to verify this assertion; the essential point is that the formulas defining the group structure are compatible with k -algebra maps.] A functor G defined in this way is called an *algebraic affine group scheme* over k , and the group $G(R)$ is called the group of *R-points* of G . The canonical example is $G = SL_n$, viewed now as the functor $R \mapsto SL_n(R)$.

Remark. If you are familiar with topological groups, you might be surprised that we did not require the inversion map $g \mapsto g^{-1}$ to be a polynomial map. The reason for not requiring it is that it turns out to be a formal consequence of our definition (cf. Waterhouse [64], Chapter 1). In our SL_n example, for instance, Cramer's rule provides a polynomial formula for the inverse.

Examples

1. Fix a matrix $g \in \mathrm{SL}_n(k)$. For any k -algebra R , let $G(R)$ be the centralizer of g in $\mathrm{SL}_n(R)$. Then G is defined by polynomial equations (which depend on the given g), and the group law is given by matrix multiplication. Thus G is an algebraic affine group scheme over k .

2. The *multiplicative group* is the functor $R \mapsto R^*$, the latter being the group of invertible elements of R . To describe it by an equation, note that we have $x \in R^*$ if and only if there is a $y \in R$ with $xy = 1$. Since y is uniquely determined by x , we can identify R^* with the “hyperbola” $xy = 1$ in the plane R^2 . The group structure is given by

$$(x, y) \cdot (x', y') = (xx', yy').$$

[Note, incidentally, that we have the polynomial formula $(x, y)^{-1} = (y, x)$ for inversion.] Another way to describe this group by equations is to identify it with the diagonal subgroup of SL_2 ; the latter is defined by the equations $a_{12} = a_{21} = 0$, $\det(a_{ij}) = 1$. The group law is of course matrix multiplication.

3. The general linear group GL_n can be treated similarly. Namely, we can identify it with the set of solutions of the equation $\det(a_{ij}) \cdot y = 1$ in $n^2 + 1$ variables a_{ij}, y . [EXERCISE: Write down a polynomial formula for inversion.] Alternatively, we can identify GL_n with the matrix group

$$\begin{pmatrix} * & \cdots & * \\ \vdots & \ddots & \vdots \\ * & \cdots & * \\ & & & * \end{pmatrix} \subset \mathrm{SL}_{n+1},$$

which is defined by adding $2n$ equations to the determinant equation defining SL_{n+1} . When $n = 1$, this example reduces to Example 2.

4. The *additive group* is defined by $G(R) = R$, with addition as the group law. It is the set of solutions of the empty set of equations in one variable, and the group structure is clearly given by a polynomial map. Alternatively, G can be identified with the matrix group $\begin{pmatrix} 1 & * \\ 0 & 1 \end{pmatrix}$.

5. The *circle group* is the curve $x^2 + y^2 = 1$ in the plane, with group structure given by imitating the familiar rule for multiplying complex numbers of norm 1:

$$(x, y) \cdot (x', y') = (xx' - yy', xy' + x'y).$$

Once again, our group can be identified with a matrix group; namely, it is isomorphic to the *rotation group*, consisting of matrices $\begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \mathrm{SL}_2$ with $a = d$ and $b = -c$.

6. For any integer $n \geq 2$, there is a group scheme μ_n , called the *group of n th roots of unity*, defined by $\mu_n(R) = \{x \in R : x^n = 1\}$ with group structure given by multiplication. This is a group of 1×1 matrices.

Remark. It is no accident that we were able to represent every example as a matrix group. In fact, one of the first theorems of the subject is that every affine algebraic group scheme is isomorphic to a closed subgroup of some GL_n , i.e., a subgroup defined by polynomial equations in the n^2 matrix entries. See [64], §3.4.

B The affine algebra of G

Let G be an algebraic affine group scheme defined, as above, by a collection of polynomial equations in m variables. Write the equations in the form $f(x_1, \dots, x_m) = 0$, and let I be the ideal in the polynomial ring $k[X_1, \dots, X_m]$ generated by the given f 's. The *affine algebra* of G is the quotient $A = k[X_1, \dots, X_m]/I$. For example, the affine algebra of the circle is $k[X, Y]/(X^2 + Y^2 - 1)$.

The affine algebra A of G “represents” G in the following sense: For any k -algebra R , the set $G(R)$ is in 1-1 correspondence with the set $\mathrm{Hom}(A, R)$ of algebra homomorphisms $A \rightarrow R$. More concisely, $G = \mathrm{Hom}(A, -)$. Of course this only describes G as a set-valued functor. To describe the group structure on G we need to impose some extra structure on A , consisting of a “comultiplication” $c : A \rightarrow A \otimes A$ satisfying certain axioms. The algebra A with this extra structure is called a *Hopf algebra*. See [64], Chapter 1, for details.

C Extension of scalars

Suppose we have a field extension $k' \supseteq k$. Then any polynomial with coefficients in k also has coefficients in k' . So if G is a group scheme defined as above by polynomials with coefficients in k , then G yields a group scheme G' over k' defined by the same formulas. In other words, we simply “restrict” G from the category of k -algebras to the category of k' -algebras. [Any k' -algebra R' can be viewed as a k -algebra.] The group scheme G' is said to be obtained from G by *extension of scalars* from k to k' . If G is represented by a Hopf algebra A over k , then G' is represented by the Hopf algebra $A' = k' \otimes_k A$ over k' .

Here is an example to show what can happen when one extends scalars. Let G be the circle group over \mathbf{R} . After extending scalars to \mathbf{C} , the resulting G' is still the circle group, viewed now as a group scheme over \mathbf{C} . But the defining equation for the circle can be written as $(x + iy)(x - iy) = 1$ over \mathbf{C} , and it follows easily that there is an isomorphism $G' \rightarrow \mathrm{GL}_1$ given

by $(x, y) \mapsto x + iy$. Thus G becomes isomorphic to the multiplicative group after extension of scalars, but the two group schemes are easily seen to be non-isomorphic over \mathbf{R} .

D Group schemes from groups

Let's go back to the naïve point of view, as in the first paragraph of the appendix. Thus we assume that we are given a group $G_0 \subseteq k^m$ which is the solution set of a collection of polynomial equations and which has a group law $G_0 \times G_0 \rightarrow G_0$ defined by a polynomial map. Assume further that the inversion map $G_0 \rightarrow G_0$ is a polynomial map. There is then a canonical way to "extend" G_0 to a group scheme G , with G_0 as the group of k -points $G(k)$. Namely, consider *all* polynomial equations which are satisfied by G_0 , and define $G(R)$ to be the set of solutions of the same equations in R^m . It is not hard to show that the polynomial formula defining the group structure on G_0 works for arbitrary R and makes G a group scheme (cf. [64], §4.4).

This passage from groups to group schemes has a simple interpretation in terms of Hopf algebras: Given G_0 , let A be the ring of functions $G_0 \rightarrow k$ given by polynomials. Equivalently, $A = k[X_1, \dots, X_m]/I$, where I is the ideal consisting of all polynomials which vanish on G_0 . Then the group structure on G_0 yields a Hopf algebra structure on A , and the group scheme G is simply the functor $\text{Hom}(A, -)$ represented by A .

A group scheme G over k which arises from a group G_0 in this way will be said to be *determined by its k -points*. For example, one can show that GL_n and SL_n are determined by their k -points as long as k is infinite ([64], §4.5). On the other hand, the group scheme μ_3 over \mathbf{Q} is not determined by its group of \mathbf{Q} -points, which is the trivial group.

It is easy to characterize the group schemes G which are determined by their k -points: If A is the affine algebra of G , then G is determined by its k -points if and only if no non-zero element of A goes to zero under all k -algebra homomorphisms $A \rightarrow k$. In case k is algebraically closed, Hilbert's Nullstellensatz allows us to restate the criterion as follows (cf. [64], §4.5): G is determined by its k -points if and only if A is reduced, i.e., if and only if A has no non-zero nilpotent elements.

E Linear algebraic groups

We are ready, finally, for the main definition. Let k be a field and let \bar{k} be its algebraic closure. Let G be an algebraic affine group scheme over k , and let \bar{G} be the group scheme over \bar{k} obtained from G by extension of scalars. We say that G is a *linear algebraic group defined over k* if the group scheme \bar{G} is determined by its group of \bar{k} -points. For example, GL_n and SL_n are linear algebraic groups over k for any k , and μ_3 is a linear algebraic group over k unless k has characteristic 3. In characteristic 3, on the other hand, μ_3 over \bar{k} is not determined by its group of \bar{k} -points, which

is the trivial group; so μ_3 is not a linear algebraic group defined over k in this case.

The following remarks should help you digest the definition.

Remarks

1. Our primary interest here is in actual groups rather than group functors. This is why we insist that we should get an actual group (i.e., the group scheme associated to an actual group) after extension of scalars. But it would be too restrictive to demand that G itself be the group scheme associated to a group, since that would exclude such examples as μ_3 over \mathbf{Q} or SL_n over a finite field.

2. The word “linear” in the definition above serves as a reminder of the fact, mentioned at the end of §A, that G is isomorphic to a closed subgroup of a general linear group.

3. In view of Hilbert’s Nullstellensatz, we can restate the definition of “linear algebraic group defined over k ” in terms of the affine algebra A of G (cf. the last paragraph of §D above): The group scheme G is a linear algebraic group defined over k if and only if $\bar{k} \otimes_k A$ is reduced. This is equivalent to a condition called *smoothness* (cf. [64], Chapter 11), and there are techniques for checking it. In characteristic 0 it is known that *all* algebraic affine group schemes are smooth, so there is nothing to check. In characteristic p , however, we have already seen that smoothness can fail (e.g., μ_3 in characteristic 3).

F Tori

Let G be the “ n -dimensional torus” over \mathbf{R} , i.e., the product of n copies of the circle group. The example in §C above shows that G becomes isomorphic to the direct product $(\mathrm{GL}_1)^n$ of n copies of the multiplicative group after extension of scalars to \mathbf{C} . This motivates the following terminology: A linear algebraic group G is a *torus* of rank n if G becomes isomorphic to $(\mathrm{GL}_1)^n$ after extension of scalars to \bar{k} . The torus is said to be *split* (or *k -split*) if it is already isomorphic to $(\mathrm{GL}_1)^n$ over k . We saw in Chapter V the canonical examples where split tori arise “in nature”; namely, the diagonal groups called T in §§V.5–7 are split tori.

G Unipotent groups

An element $g \in \mathrm{GL}_n(k)$ is called *unipotent* if $g - 1$ is nilpotent. This is equivalent to saying that g is conjugate to an element of $U_n(k)$, where U_n is the strict upper triangular group (i.e., the group of upper triangular matrices with 1’s on the diagonal). A group of $n \times n$ matrices is called *unipotent* if each of its elements is unipotent. This is equivalent, by a theorem of Kolchin (cf. [64], §8.1), to saying that the group is conjugate to a subgroup of $U_n(k)$. Finally, if G is a linear algebraic group over k , choose an

embedding of G as a closed subgroup of some GL_n , and call G *unipotent* if $G(\bar{k})$ is a unipotent subgroup of $\mathrm{GL}_n(\bar{k})$. This is equivalent to saying that there is an element of $\mathrm{GL}_n(k)$ which conjugates G into U_n ([64], §8.3). Moreover, this notion is independent of the choice of embedding of G in a general linear group.

H Connected groups

There is a topology on k^m , called the *Zariski topology*, in which the closed sets are the subsets defined by polynomial equations. The subset $G(k) \subseteq k^m$ inherits a Zariski topology, and we can therefore apply topological concepts, such as connectivity, to $G(k)$. More useful for us is the Zariski topology on $G(\bar{k}) \subseteq \bar{k}^m$. In particular, we will say that the linear algebraic group G is *connected* if $G(\bar{k})$ is connected in the Zariski topology. For example, GL_n and SL_n are connected (for any k), but μ_3 over \mathbf{Q} is not. Note, however, that $\mathrm{GL}_n(k)$ is disconnected if k is a finite field, whereas the disconnected group μ_3 has the property that $\mu_3(\mathbf{Q})$ is connected. Thus it is important to look at $G(\bar{k})$ rather than $G(k)$ in order to get the “right” answer.

I Reductive, semisimple, and simple groups

Let G be a connected linear algebraic group over k . Then G is called *reductive* if $G(\bar{k})$ contains no non-trivial connected normal unipotent subgroup. For example, GL_n and SL_n are reductive.

G is called *semisimple* if $G(\bar{k})$ contains no non-trivial connected normal solvable subgroup. For example, SL_n is semisimple but GL_n is not (because of its center). Note that any semisimple group is reductive, since unipotent matrix groups are solvable (and even nilpotent) by Kolchin’s theorem.

If G is semisimple, then $G(\bar{k})$ is “almost” a finite direct product of simple groups. More precisely, $G(\bar{k})$ has the following properties: (i) it has only finitely many minimal non-trivial closed connected normal subgroups N_i ; (ii) the N_i commute and generate $G(\bar{k})$; (iii) the canonical surjection $\prod_i N_i \twoheadrightarrow G(\bar{k})$ has finite kernel; and (iv) each N_i is *almost simple*, which means that its center Z_i is finite and that the quotient N_i/Z_i is simple as an abstract group. Proofs can be found in [32], §§27.5 and 29.5. If there is only one N_i , i.e., if $G(\bar{k})$ is almost simple, then G is said to be *absolutely almost simple*. For example, SL_n is absolutely almost simple.

J BN-pairs and spherical buildings

If G is reductive, then the group $G(k)$ has a BN-pair whose associated building is spherical. I’ll give a brief description of this in the semisimple case. For more details, see [56] and the references cited there. See also [32], §28.3, for the case where k is algebraically closed.

A *Borel subgroup* of $G(\bar{k})$ is a maximal connected solvable subgroup of $G(\bar{k})$. Borel subgroups exist and are unique up to conjugacy. If k is algebraically closed, any Borel subgroup can serve as the B of the BN-pair

in $G(k) = G(\bar{k})$. Let T be a maximal torus in G . These also exist and are unique up to conjugacy. Since $T(k)$ is connected and solvable, we can choose the Borel subgroup B to contain $T(k)$. Still assuming that $k = \bar{k}$, we can then take the N of the BN-pair to be the normalizer of $T(k)$ in $G(k)$. The resulting spherical building has rank l (dimension $l - 1$), where l is the rank of T . Its chambers are in 1-1 correspondence with the Borel subgroups of G , and its apartments are in 1-1 correspondence with the maximal tori in G . We have seen this construction in Chapter V for the semisimple groups SL_n , Sp_{2n} , and SO_n .

It is immediate from the definitions above that the parabolic subgroups with respect to this BN-pair are the subgroups of $G(k)$ which contain a Borel subgroup. There is another characterization of them, whose statement involves concepts that we have not defined (and will not define): They are the subgroups of the form $P(k)$, where P is a closed subgroup of G such that G/P is a projective variety. For example, let $G = SL_n$ and let P be the subgroup defined by $a_{i1} = 0$ for $i > 1$ (i.e., P is the stabilizer of the line $[e_1]$); then G/P is $(n - 1)$ -dimensional projective space.

When $k \neq \bar{k}$, the situation is more complicated, the problem being that $G(\bar{k})$ might not have a Borel subgroup which is defined over k (i.e., which is the group of \bar{k} -points of a linear algebraic subgroup of G defined over k). Orthogonal groups provide examples of this phenomenon. We did not see it in §V.7 because we only considered the orthogonal groups of some very special quadratic forms.

To get a BN-pair, in general, one has to forget about Borel subgroups and instead take B to be the group $P(k)$ for some minimal parabolic subgroup P of G , where now “parabolic” is *defined* by the property that G/P is a projective variety. B is again unique up to conjugacy. We can choose B to contain $T(k)$, where T is now a maximal k -split torus in G , and we take N to be the normalizer of $T(k)$ in $G(k)$. The rank of the resulting spherical building is again equal to the rank l of T . This rank l is also called the *k-rank* of G . It can be strictly smaller than the rank of the building associated to $G(\bar{k})$. In fact, it can be zero, in which case the building is empty (i.e., it consists only of the empty simplex).

K BN-pairs and Euclidean buildings

Here I will be even more brief. Assume that G is absolutely almost simple and is defined over a field K with a discrete valuation. Assume further that G is simply connected. (This is another term that I have not defined; an example is SL_n .) Then there is a Euclidean BN-pair in $G(K)$, analogous to the one we have seen for $SL_n(K)$. The associated Euclidean building has dimension l , where l is again the K -rank of G ; thus the dimension of this building exceeds by 1 the dimension of the spherical building associated to $G(K)$. When K is locally compact (e.g., $K = \mathbf{Q}_p$), the parabolic subgroups of $G(K)$ with respect to this BN-pair are open and compact. The theorem at the end of §VI.5 therefore provides a description of the Euclidean building

in terms of the maximal compact subgroups of $G(K)$. Here, of course, we use the locally compact topology that comes from the valuation, not the Zariski topology. See Bruhat–Tits [22] for more information.

L Group schemes versus groups

I have insisted on thinking of algebraic groups as functors because I find this point of view useful. But, in so doing, I may have given the misleading impression that a “typical” linear algebraic group G defined over k is not determined by its group of k -points. I will therefore close by stating a theorem which says that G is determined by its k -points much more often than one might expect. Let G be a linear algebraic group defined over k . Then G is determined by its k -points whenever the following three conditions are satisfied: (i) G is connected; (ii) k is infinite; and (iii) either k is perfect or G is reductive. (Note that (ii) and (iii) hold automatically in characteristic 0.) For a proof of this theorem see Borel [13], Corollary 18.3, where the result is stated in the following equivalent form: If (i), (ii), and (iii) hold, then $G(k)$ is Zariski-dense in $G(\bar{k})$.

Suggestions for Further Reading

One way to continue the study of buildings would be to learn more about their applications to group theory. In particular, there is a very close connection between buildings and algebraic groups, which I have only hinted at in this book. To learn more about this and other applications, start by looking at the survey by Tits [57]. You will find many additional references there.

Proofs of some of the results announced in [57] can be found in [60]. The latter also contains a very interesting generalization of the notion of Euclidean building. Roughly speaking, Tits considers metric spaces with “building-like” properties. In the 1-dimensional case, these metric spaces are essentially the same as “**R**-trees”, which are the object of much current research, cf. [6] and [38].

You might also enjoy browsing through the other papers in the volume that contains [60]. You will discover, for instance, that some finite group theorists, hoping to find geometric interpretations for the sporadic simple groups, are currently quite interested in buildings.

Finally, there is the forthcoming book by Ronan [42]. Although that book has some overlap with the present one, the point of view is quite different. In addition, Ronan treats a number of topics that I have not touched on.

References

- [1] H. Abels, *An example of a finitely presented solvable group*, in "Homological Group Theory," London Math. Soc. Lecture Notes 36 (edited by C. T. C. Wall), Cambridge University Press, Cambridge, 1979, pp. 205–211.
- [2] H. Abels, "Finite Presentability of S -Arithmetic Groups; Compact Presentability of Solvable Groups," Lecture Notes in Mathematics 1261, Springer-Verlag, Berlin, 1987.
- [3] H. Abels, *Finiteness properties of certain arithmetic groups in the function field case*, preprint, 1988.
- [4] H. Abels and K. S. Brown, *Finiteness properties of solvable S -arithmetic groups: An example*, J. Pure Appl. Algebra **44** (1987), 77–83.
- [5] P. Abramenko, *Endlichkeitseigenschaften der Gruppen $SL_n(\mathbb{F}_q[t])$* , Dissertation, Frankfurt, 1987.
- [6] R. C. Alperin and H. Bass, *Length functions of group actions on Λ -trees*, in "Combinatorial Group Theory and Topology," Ann. of Math. Studies 111 (edited by S. M. Gersten and J. R. Stallings), Princeton University Press, Princeton, 1987, pp. 265–378.
- [7] R. C. Alperin and P. B. Shalen, *Linear groups of finite cohomological dimension*, Invent. Math. **66** (1982), 89–98.
- [8] E. Artin, "Geometric Algebra," Wiley Interscience, New York, 1957.
- [9] M. F. Atiyah and I. G. MacDonald, "Introduction to Commutative Algebra," Addison-Wesley, Reading, 1969.
- [10] H. Behr, *Finite presentability of arithmetic groups over global function fields*, Proc. Edinburgh Math. Soc. (2) **30** (1987), 23–39.
- [11] R. Bieri, "Homological Dimension of Discrete Groups," 2nd edition, Queen Mary College Mathematics Notes, London, 1981.
- [12] A. Borel, *Groupes linéaires algébriques*, Ann. of Math. **64** (1956), 20–82.
- [13] A. Borel, "Linear Algebraic Groups," Notes by H. Bass, Benjamin, New York, 1969.
- [14] A. Borel and J-P. Serre, *Corners and arithmetic groups*, Comment. Math. Helv. **48** (1974), 244–297.
- [15] A. Borel and J-P. Serre, *Cohomologie d'immeubles et de groupes S -arithmétiques*, Topology **15** (1976), 211–232.
- [16] N. Bourbaki, "Groupes et Algèbres de Lie," Chapitres 4–6, Masson, Paris, 1981.
- [17] K. S. Brown, "Cohomology of Groups," Graduate Texts in Mathematics 87, Springer-Verlag, New York, 1982.
- [18] K. S. Brown, *Presentations for groups acting on simply-connected complexes*, J. Pure Appl. Algebra **32** (1984), 1–10.
- [19] K. S. Brown, *Finiteness properties of groups*, J. Pure Appl. Algebra **44** (1987), 45–75.

- [20] K. S. Brown and R. Geoghegan, *Cohomology with free coefficients of the fundamental group of a graph of groups*, Comment. Math. Helv. **60** (1985), 31–45.
- [21] F. Bruhat, *Représentations des groupes de Lie semi-simples complexes*, C. R. Acad. Sci. Paris **238** (1954), 437–439.
- [22] F. Bruhat and J. Tits, *Groupes réductifs sur un corps local. I. Données radicielles valuées*, Inst. Hautes Études Sci. Publ. Math. **41** (1972), 5–251.
- [23] C. Chevalley, *Sur certains groupes simples*, Tôhoku Math. J. (2) **7** (1955), 14–66.
- [24] C. Chevalley, “Classification des groupes de Lie algébriques,” Séminaire École Norm. Sup., Paris, 1956–58. Mimeographed notes, 2 volumes.
- [25] H. S. M. Coxeter, *The complete enumeration of finite groups of the form $R_i^2 = (R_i R_j)^{k_{ij}} = 1$* , J. London Math. Soc. **10** (1935), 21–25.
- [26] H. S. M. Coxeter, “Regular Polytopes,” 2nd edition, Macmillan, New York, 1963.
- [27] W. Feit and G. Higman, *The nonexistence of certain generalized polygons*, J. Algebra **1** (1964), 114–131.
- [28] D. R. Grayson, *Finite generation of K -groups of a curve over a finite field [after Daniel Quillen]*, in “Algebraic K -Theory,” Proceedings of a June 1980 Oberwolfach conference, Part I, Lecture Notes in Mathematics 966, Springer-Verlag, Berlin, 1982, pp. 69–90.
- [29] D. R. Grayson, *Reduction theory using semistability*, Comment. Math. Helv. **59** (1984), 600–634.
- [30] D. R. Grayson, *Reduction theory using semistability, II*, Comment. Math. Helv. **61** (1986), 661–676.
- [31] L. C. Grove and C. T. Benson, “Finite Reflection Groups,” 2nd edition, Graduate Texts in Mathematics 99, Springer-Verlag, New York, 1985.
- [32] J. E. Humphreys, “Linear Algebraic Groups,” Graduate Texts in Mathematics 21, Springer-Verlag, New York, 1975.
- [33] N. Iwahori and H. Matsumoto, *On some Bruhat decomposition and the structure of the Hecke rings of p -adic Chevalley groups*, Inst. Hautes Études Sci. Publ. Math. **25** (1965), 5–48.
- [34] N. Jacobson, “Basic Algebra II,” W. H. Freeman and Co., San Francisco, 1980.
- [35] J. Lehner, “Discontinuous Groups and Automorphic Functions,” Mathematical Surveys no. 8, American Mathematical Society, Providence, 1964.
- [36] R. C. Lyndon, “Groups and Geometry,” London Mathematical Society Lecture Note Series 101, Cambridge University Press, Cambridge, 1985.
- [37] H. Matsumoto, *Générateurs et relations des groupes de Weyl généralisés*, C. R. Acad. Sci. Paris **258** (1964), 3419–3422.
- [38] J. W. Morgan and P. B. Shalen, *Valuations, trees, and degenerations of hyperbolic structures. I*, Ann. of Math. **120** (1984), 401–476.
- [39] L. S. Pontryagin, “Topological Groups,” 2nd edition, Gordon and Breach, New York, 1966.
- [40] D. Quillen, *Homotopy properties of the poset of non-trivial p -subgroups of a group*, Advances in Math. **28** (1978), 101–128.
- [41] M. S. Raghunathan, *A note on quotients of real algebraic groups by arithmetic subgroups*, Invent. Math. **4** (1968), 318–335.

- [42] M. A. Ronan, "Lectures on Buildings," Academic Press, Boston, 1989.
- [43] J-P. Serre, *Cohomologie des groupes discrets*, in "Prospects in Mathematics," Ann. of Math. Studies 70, 1971, pp. 77–169.
- [44] J-P. Serre, "A Course in Arithmetic," Graduate Texts in Mathematics 7, Springer-Verlag, New York, 1973.
- [45] J-P. Serre, *Arithmetic groups*, in "Homological Group Theory," London Math. Soc. Lecture Notes 36 (edited by C. T. C. Wall), Cambridge University Press, Cambridge, 1979, pp. 105–136.
- [46] J-P. Serre, "Trees," Springer-Verlag, Berlin, 1980. Translation of "Arbres, Amalgames, SL_2 ," *Astérisque* 46, 1977.
- [47] L. Solomon, *The Steinberg character of a finite group with BN-pair*, in "Theory of Finite Groups," Proc. Sympos. Harvard 1968 (edited by R. Brauer and C. H. Sah), Benjamin, New York, 1969, pp. 213–221.
- [48] E. H. Spanier, "Algebraic Topology," McGraw-Hill, New York, 1966.
- [49] R. Strebel, *Finitely presented soluble groups*, in "Group Theory: Essays for Philip Hall," edited by K. W. Gruenberg and J. E. Roseblade, Academic Press, New York, 1984, pp. 257–314.
- [50] U. Stuhler, *Homological properties of certain arithmetic groups in the function field case*, Invent. Math. 57 (1980), 263–281.
- [51] J. Tits, *Groupes et géométries de Coxeter*, unpublished manuscript, 1961.
- [52] J. Tits, *Théorème de Bruhat et sous-groupes paraboliques*, C. R. Acad. Sci. Paris 254 (1962), 2910–2912.
- [53] J. Tits, *Géométries polyédriques et groupes simples*, in "Atti della II Riunione del Groupement de Mathématiciens d'Expression Latine," Florence (1961), Edizioni Cremonese, Rome, 1963, pp. 66–88.
- [54] J. Tits, *Structures et groupes de Weyl*, Séminaire Bourbaki 1964-65, exposé no. 288, February 1965.
- [55] J. Tits, *Le problème des mots dans les groupes de Coxeter*, Symposia Mathematica (INDAM, Rome, 1967/68), vol. 1, pp. 175–185, Academic Press, London, 1969.
- [56] J. Tits, "Buildings of Spherical Type and Finite BN-Pairs," Lecture Notes in Mathematics 386, Springer-Verlag, Berlin, 1974.
- [57] J. Tits, *On buildings and their applications*, in "Proceedings of the International Congress of Mathematicians," Vancouver 1974, vol. 1, Canad. Math. Congress, Montreal, 1975, pp. 209–220.
- [58] J. Tits, *Buildings and Buekenhout geometries*, in "Finite Simple Groups, II," Proc. sympos. Durham 1978 (edited by M. J. Collins), Academic Press, London, 1980, pp. 309–320.
- [59] J. Tits, *A local approach to buildings*, in "The Geometric Vein," Coxeter Festschrift (edited by C. Davis, B. Grünbaum, and F. A. Sherk), Springer, New York, 1981, pp. 519–547.
- [60] J. Tits, *Immeubles de type affine*, in "Buildings and the Geometry of Diagrams," Proc. CIME Como 1984 (edited by L. A. Rosati), Lecture Notes in Mathematics 1181, Springer-Verlag, Berlin, 1986, pp. 159–190. [Note: The author has written a correction to this paper, which is not yet published, in which he has modified axiom (A5).]
- [61] B. L. van der Waerden, "Gruppen von Linearen Transformationen," Springer-Verlag, Berlin, 1935.

- [62] E. B. Vinberg, *Discrete linear groups generated by reflections*, Math. USSR Izvestija **5** (1971), 1083–1119.
- [63] K. Vogtmann, *Spherical posets and homology stability for $O_{n,n}$* , Topology **20** (1981), 119–132.
- [64] W. C. Waterhouse, “Introduction to Affine Group Schemes,” Graduate Texts in Mathematics 66, Springer-Verlag, New York, 1979.

Notation Index

$\langle - \rangle$ 3	F_4 5, 23
\leq 7	(F) 47
$*$ 63, 175	FL 185
$[-]$ 100, 132	F_n 193
$[[-]]$ 132	$\phi(\Gamma)$ 193
A_n 5, 22	G_2 5, 23
$ A $ 28	$GL(V)$ 140
$A \cap B$ 12, 27	H_n 5, 23
\bar{A} 8, 72, 95, 151	\mathcal{H} 15, 34, 141
$\text{Aff}(V)$ 140	$I_2(m)$ 6, 23
$\text{Aut}_0 \Sigma$ 65	\hat{K} 130
\mathcal{A}_∞ 178	$l(-, -)$ 13
\mathfrak{A}_∞ 175	$l(w)$ 34
(A) 35	$\lambda_{C,D}$ 166
B_n 5, 22	$\text{lk } A$ 31
$\mathcal{B}(C, C')$ 163	(NC) 155
$(B0)$ 76	(NC) 156
$(B1)$ 76	O_m 124
$(B2)$ 76	$O(V)$ 140
$(B2')$ 77	$\pi(-, -)$ 84
$(B2'')$ 77	PGL_n 41, 101
$(BN1)$ 110	PSL_n 41, 101
$(BN2)$ 110	\mathbf{Q}_p 130
$(BN3)$ 114	$\rho_{\Sigma, C}$ 86
C_n 5, 22	$\rho_{E, C}$ 151
(C) 53	$\rho_{E, e}$ 171
cd 185	$r(x, A)$ 157
Ch 67	$r(A)$ 157
D_{2m} 3	$\Sigma(W, S)$ 33
D_∞ 37	$\Sigma(W, V)$ 144
D_n 5, 22	Σ_M 78
$d(-, -)$ 14, 15, 89, 151	Σ_∞ 177
$d(-, -)$ 152	SO_m 124
$\delta(-, -)$ 89	Sp_{2n} 121
$ \Delta $ 28	st x 154
$\Delta(G, B)$ 111	s_H 1, 34, 70, 140
Δ_∞ 176	τ_v 140
$\Delta_\infty(\mathcal{A})$ 178	vcd 188
(D) 37	VFL 188
E_n 5, 22, 23	W_M 78
(E) 47	X_∞ 175

Subject Index

- absolutely almost simple 203
- action condition 35
- adjacent 13, 29, 31, 34, 59
- affine algebra 200
- affine building 151
- affine concepts 139–140
- affine reflection group 141
- affine Weyl group 148
- algebraic group 201
- almost simple 203
- Alperin–Shalen theorem 194
- alternating form 120
- apartment 76
- arithmetic group 183–184

- barycentric subdivision 28
- BN-pair 110
- Borel subgroup 203
- Borel–Serre compactification 186, 190
- bornological group 161
- bornology 161
- bounded set 160–161
- bounded subgroups 159ff
- Bruhat decomposition 104, 115
- building 76
- building at infinity 174ff

- canonical labelling 59
- canonical map 90
- canonical representation 55
- cd 185
- cell 6ff, 141
- chamber 9, 29, 34, 141
- chamber complex 29
- chamber map 30
- chamber system 31, 63
- circumcenter 158
- circumradius 157
- closed cell 8
- coboundary 159
- cocompact 185
- cocycle 159
- Cohen–Macaulay complex 94
- cohomological dimension 185
- combinatorial distance 14, 34
- commensurable 184
- compact subgroups 159ff
- complete system of apartments 88
- completion 130
- conical cell 164, 169
- connected group 203
- convex hull 93
- convex subcomplex 15, 68, 88
- convex subset 166
- cosine inequality 169
- Coxeter complex 33, 58, 71
- Coxeter condition 52
- Coxeter diagram 22, 60, 78
- Coxeter group 33, 52–53
- Coxeter matrix 21, 54, 78
- Coxeter system 53
- crystallographic 146

- deletion condition 37, 73–74
- diameter 15, 61
- dihedral group 3, 37
- dimension 27
- direction 164
- discrete valuation 127
- discrete valuation ring 127
- distance 14, 15, 34, 89, 151
- dual of a regular solid 4

- elementary divisors 129
- elementary M -operation 49
- end 176
- essential 2, 11, 143
- Euclidean BN-pair 160
- Euclidean building 151
- Euclidean Coxeter complex 149
- Euclidean reflection group 144
- Euclidean space 150
- exchange condition 46
- extension of scalars 200

- face 7, 27, 33, 177

- face at infinity 175
- Feit–Higman theorem 117
- finite reflection group 1
- finiteness conditions 193
- finiteness length 193
- fixed-point theorem 157
- FL 185
- flag 28
- flag complex 28
- folding 48, 66
- folding condition 47
- frame 84, 139
- fundamental chamber 34, 99
- fundamental apartment 99

- gallery 13, 29, 34
- generalized BN-pair 126, 134–135
- generalized m -gon 117
- geodesic 94, 155
- geometric realization 28
- girth 87
- good apartment system 179
- group scheme 198

- half-space 35, 70
- Hirsch rank 189
- Hopf algebra 200
- hyperplane 1, 139

- ideal point 175
- ideal simplex 175
- idempotent 66
- incidence geometry 79
- irreducible 4, 23, 56
- irreducible components 23
- isotropic subspace 122

- join 63
- joinable simplices 29
- Jordan–Hölder permutation 84, 105

- labelling 29
- lattice 128, 145
- length 13, 34
- linear algebraic group 201
- linear part 140
- link 31

- midpoint 156
- midpoint convex 156
- minimal gallery 14
- monomial group 101
- monomial matrix 4
- M -reduced 49

- negative curvature 155ff

- open cell 8
- opposite chamber 92
- opposite folding 69
- opposite half-space 70
- opposite parabolics 115
- oriflamme geometry 126
- orthogonal group 124

- p -adic integers 130
- p -adic numbers 130
- p -adic valuation 128
- parabolic subgroup 112
- parallel 172, 174
- polar geometry 82–83, 122
- poset 12, 27
- projective geometry 81ff
- projective plane 81
- proper action 185
- property (NC) 156

- quartier 164

- rank 27, 30, 204
- ray 174
- reduced decomposition 46
- reduced word 46
- reductive group 203
- reflection 1, 33, 39, 71, 140
- regular solid 2
- residue field 128
- retraction 30, 66, 85ff, 170–171
- reversible folding 69
- root system 2
- R -tree 206

- s -adjacent 34
- S -arithmetic group 189, 196
- saturated BN-pair 114
- sector 164, 169
- semisimple group 203
- simplicial complex 24, 27
- simplicial cone 11

- simplicial map 30
- simply connected group 204
- Solomon-Tits theorem 93ff
- special coset 26, 33, 111
- special orthogonal group 124
- special subgroup 26, 110
- spherical building 92
- spherical Coxeter complex 92
- standard uniqueness argument 69, 71
- star 154–155
- stretched 89
- strong isometry 90
- strongly transitive 99
- stuttering gallery 13
- subcomplex 15, 27
- subsector 164
- support 7, 141
- symmetric space 184, 189
- symplectic group 121
- system of apartments 76
- thick chamber complex 77
- thin chamber complex 58
- Tits system 110
- torus 202
- totally isotropic subspace 122
- type F_n 193
- type FL 185
- type of a gallery 73–74
- type of a simplex 30
- type VFL 188
- type-preserving 58
- unipotent 202
- valuation 127
- valuation ring 127
- vcd 188
- VFL 188
- virtual cohomological dimension 188
- wall 10, 34, 70, 141
- walls crossed 14, 34
- weak BN-pair 127
- weak building 77
- Weyl group 2, 110, 115
- word problem 49
- word 46
- Zariski topology 203

