

THE SHEAF HOMOLOGY OF THE CAYLEY MODULE FOR $G_2(q)$

by

MARK WILLIAM BUTLER

A thesis submitted to
The University of Birmingham
for the degree of
PHD PURE MATHEMATICS

School of Mathematics
The University of Birmingham
31st March 2022

UNIVERSITY OF
BIRMINGHAM

University of Birmingham Research Archive

e-theses repository

This unpublished thesis/dissertation is copyright of the author and/or third parties. The intellectual property rights of the author or third parties in respect of this work are as defined by The Copyright Designs and Patents Act 1988 or as modified by any successor legislation.

Any use made of information contained in this thesis/dissertation must be in accordance with that legislation and must be properly acknowledged. Further distribution or reproduction in any format is prohibited without the permission of the copyright holder.

ABSTRACT

In the 1980s, Ronan and Smith developed a homology theory to describe modules for groups of Lie type, using sheaves constructed on geometries associated with such groups. For a finite field k of characteristic π , they establish a one-to-one correspondence between certain ‘fixed-point’ sheaves \mathcal{F}_V and the irreducible kG -modules for G a Chevalley group defined over k . This correspondence is given by the irreducible kG -module V being a unique irreducible quotient of the zero-homology module $H_0(\mathcal{F}_V)$. In certain cases, this homology module $H_0(\mathcal{F}_V)$ is in fact isomorphic to the irreducible module V . The question explored in this thesis is whether or not the Cayley module \overline{C} for the group $G_2(k)$ is isomorphic to $H_0(\mathcal{F}_{\overline{C}})$, as was speculated by Segev and Smith in 1986.

ACKNOWLEDGEMENTS

As I come towards the end of my time as a student at the University of Birmingham, there are many people who I would like to thank. First and foremost are my two PhD supervisors, Chris and Hoff. Hoff introduced me to the wonderful world of sheaves on buildings, for which I am incredibly grateful, and our discussions were always enlightening.

Chris took over supervision halfway through my PhD, and proposed the research topic which eventually ended up as this thesis. His infinite patience with my endless stream of questions never ceased to amaze me. This thesis would certainly not have been possible without all of his support, guidance and brilliant insight.

A large debt of gratitude is owed to my fellow research students. In particular, I should mention Martin, Jack and Andrea, my friends and colleagues from the Algebra research group who have been a constant source of support. The lockdown made me realise more than ever the benefits of doing maths collaboratively — without friends in the same room to share ideas with, progress became much slower. I have also been lucky enough to work in an office alongside some great friends. By attempting to name them all I would surely forget some, and I fear this thesis is long enough already, but they know who they are and the shared memories I will cherish forever.

I would also like to thank Michael Grove for his kindness and support whilst I was working firstly as a PGTA in the department, and then as a Teaching Fellow. In particular, an enormous thank you is owed for allowing me the time to finish off this thesis.

I am very grateful for the financial and practical support provided by the EPSRC, especially for the additional support offered due to the impact of the COVID-19 pandemic on research students including myself.

A big thank you to my parents and my brother Iain for supporting me throughout this journey — they probably didn't expect, when driving me to Birmingham for the first time for my undergraduate studies, that I would still be here eight years later. Finally, I must thank my partner Louise for all her loving support, encouragement, and for making the experience of writing a thesis during lockdown not just bearable but somewhat enjoyable.

CONTENTS

1	Introduction	1
1.1	Sheaves on buildings	2
1.2	The sheaf of the Cayley module for $G_2(k)$	3
1.3	Related results in the literature	4
1.4	Chapters in this thesis	7
 2	 Buildings	 9
2.1	Simplicial complexes	10
2.2	Coxeter groups and chamber systems	11
2.2.1	Coxeter groups	11
2.2.2	Chamber systems	14
2.2.3	The geometric realisation of a chamber system	16
2.3	Apartments and buildings	18
2.3.1	Buildings	18
2.3.2	Roots and the Moufang property	19
2.4	The finite groups of Lie type	22
2.4.1	The subgroups U , B , T and N	23
2.4.2	Parabolic subgroups	24
2.5	The classification of finite, thick, irreducible buildings	26
 3	 $G_2(k)$: The Cayley Algebra and the Generalised Hexagon	 28
3.1	The Cayley algebra	28
3.1.1	A bilinear form on C^+	30
3.1.2	The Cayley module	31
3.2	The group $G = G_2(k)$	32
3.3	The generalised hexagon Δ and the incidence graph Γ	38
3.3.1	The construction of Δ and Γ	39
3.3.2	Transitivity on points and lines	40
3.3.3	Arc transitivity	43
3.3.4	More properties of Γ	46
3.3.5	Opposite vertices	48
3.3.6	The generalised hexagon Δ is the building of G	49
3.3.7	Subspaces of \overline{C}	53

3.3.8	3-transitivity results	56
3.3.9	Ideal lines	57
4	Some Weight Theory of Groups of Lie Type	59
4.1	Preliminary material	59
4.1.1	The Weyl group	61
4.1.2	Theory of weights	62
4.1.3	The Weyl module $V(\lambda)$ and its unique simple quotient $L(\lambda)$	65
4.2	An example for the group $\mathrm{SL}_2(k)$	67
4.3	An example for the group $\mathrm{G}_2(k)$	71
5	Sheaves on Buildings	76
5.1	Sheaves on buildings: definition and examples	77
5.1.1	Definition	77
5.1.2	The constant sheaf	79
5.1.3	The fixed-point sheaf	79
5.2	Sheaf morphisms	83
5.2.1	Kernels and images	83
5.3	Sheaf homology	84
5.3.1	Results of Ronan and Smith	88
6	Sheaves on the Building of $\mathrm{G}_2(k)$	95
6.1	Data for some small fields	95
6.2	Setting up notation	96
6.2.1	Subspaces of \underline{C}_0, H and \overline{C}	97
6.3	Segev and Smith's Theorem	100
7	Restriction to the Prime Subfield	102
8	A Lower Bound on the Dimension of H	106
8.1	Calculating the highest weight of K/Y	107
8.2	A lower bound on the dimension of K/Y	111
9	A Cubic Upper Bound on the Dimension of H	113
9.1	Preliminaries	113
9.2	Forming a cubic bound	115
9.2.1	Points opposite both p and p^\dagger : Brouwer's theorem and rogues . . .	117
10	A Quadratic Upper Bound on the Dimension of H	121
10.1	Spines, pages and books	123
10.2	Diagonals	124
10.3	Crossbraces	127
10.4	Locating an interior point with three incident lines contained in L_X	136
10.5	Transitivity arguments	141

10.5.1	All interior points have three of their incident lines contained in L_X	142
10.5.2	All interior points have all their incident lines contained in L_X	145
10.6	Points on spines	147
10.6.1	Transitivity on the set of points like z	148
10.6.2	Transitivity on the q non-spine lines through z	150
10.7	Proof that $H = X$	151
10.7.1	A note about the condition $3 \nmid q - 1$	152
11	Proof of the Main Theorem	154
A	Magma Programs	156
	List of References	157

CHAPTER 1

INTRODUCTION

In this thesis we present a partial answer to a speculation of Segev and Smith found in their 1986 paper *Apartments and the Cayley-Algebra Module for $G_2(k)$* [26]. Our main result is as follows:

Theorem A *Let $k = \mathbb{F}_q$, where $q = \pi^a$ (π a prime) with $\pi > 3$ and $3 \nmid \pi - 1$. Set $G = G_2(k)$, and let Δ be the building of G . Denote by \overline{C} the Cayley module for G , and let $\mathcal{F}_{\overline{C}}$ be the corresponding fixed-point sheaf on Δ . Then $H_0(\Delta, \mathcal{F}_{\overline{C}}) \cong \overline{C}$ and so $\dim H_0(\Delta, \mathcal{F}_{\overline{C}}) = 7$.*

In this chapter we will explain as briefly as possible the above statement, before reviewing some related results from the literature. Finally, we will give a short summary of the contents of each chapter to provide an overview of the thesis as a whole.

The Segev and Smith paper [26] builds on earlier work of Ronan and Smith. In a series of three papers beginning with *Sheaves on Buildings and Modular Representations of Chevalley Groups* [23] published in 1985, Ronan and Smith developed techniques to study the representation theory of Chevalley groups using sheaf homology.

All groups of Lie type, including Chevalley groups, have an associated geometric structure called a *building*, on which they act in a nice manner (see Chapter 2 for details). The

theory of buildings was developed by Tits from the late 1950s onward, with the first thorough treatment being given in *Buildings of Spherical Type and Finite BN-pairs* in 1974 [32]. Buildings can be considered as simplicial complexes: these are structures consisting of simplices of various dimensions, along with *face relations* which identify one simplex as being a face of another.

1.1 Sheaves on buildings

Ronan and Smith's contribution begins by defining a *sheaf* on the building. Fix a finite field $k = \mathbb{F}_q$ of characteristic π .¹ If G is a Chevalley group with building $\Delta = \Delta(G)$, then the stabiliser $P_\sigma \leq G$ of a simplex $\sigma \in \Delta$ is called a *parabolic subgroup* of G . A sheaf \mathcal{F} is a *coefficient system* on Δ , in which each simplex $\sigma \in \Delta$ is assigned a coefficient in the form of a kP_σ -module \mathcal{F}_σ . The sheaf \mathcal{F} also contains *connecting maps* $\mathcal{F}_\sigma \rightarrow \mathcal{F}_\tau$ whenever τ is a face of σ . Finally, the G -action on Δ is extended to a G -action on the sheaf \mathcal{F} . (See Chapter 5 for a much more complete exposition on Ronan and Smith's sheaves.)

One sheaf of particular interest is the *fixed-point sheaf* \mathcal{F}_V for a kG -module V . This has coefficients given by

$$(\mathcal{F}_V)_\sigma := C_V(U_\sigma) = \{v \in V \mid vg = v \text{ for all } g \in U_\sigma\},$$

where $U_\sigma = O_\pi(P_\sigma)$ is the unipotent radical of the parabolic P_σ . The connecting maps in \mathcal{F}_V are simply restrictions of the identity map id_V , since $(\mathcal{F}_V)_\sigma \subseteq (\mathcal{F}_V)_\tau$ whenever τ is a face of σ . The motivation for the construction of this sheaf is Smith's Lemma [28], which implies that if V is an irreducible kG -module then each term $(\mathcal{F}_V)_\sigma$ is an irreducible kP_σ -module.

¹We use π for the characteristic of the field, rather than the more traditional p , because the letter p will be used throughout the thesis to denote a *point* in a geometry.

A *chain complex* is, in our context, a sequence of kG -modules C_i along with boundary maps $\delta_i: C_i \rightarrow C_{i-1}$ from each module into the previous one, such that the image of a homomorphism δ_i lies inside the kernel of δ_{i-1} . *Homology* is the study of exactly how each $\text{im}(\delta_i)$ embeds into $\ker(\delta_{i-1})$.

We can form a chain complex from \mathcal{F} by taking each chain space C_i to be the direct sum of the coefficients \mathcal{F}_σ over all simplices σ of dimension i . The boundary maps δ_i are then based on the connecting maps of the sheaf, with some sign modifiers (+/-) included. The homology module $H_i(\Delta, \mathcal{F})$ is given by $\ker(\delta_i)/\text{im}(\delta_{i-1})$.

This brings us to the first of the main results of Ronan and Smith. Theorem 2.3 of [24] states that, if V is an irreducible module, then $H_0(\Delta, \mathcal{F}_V)$ contains a unique maximal submodule K and $H_0(\Delta, \mathcal{F}_V)/K \cong V$.

For some choices of V , the kernel K is $\{0\}$ and so we get an isomorphism $H_0(\Delta, \mathcal{F}_V) \cong V$. Corollary 2.4 of [23] states that if V is the Steinberg module for G (the irreducible module $L(\lambda)$ with weight $\lambda = (\pi - 1, \dots, \pi - 1)$), then $H_0(\Delta, \mathcal{F}_V) \cong V$. The same is true for the trivial module by [23, p.324, Lemma 1.1]. (To apply the lemma we first must note that the ‘constant sheaf’ \mathcal{K}_V described by Ronan and Smith is isomorphic to the fixed-point sheaf \mathcal{F}_V when V is the trivial module [23, p.322].)

1.2 The sheaf of the Cayley module for $G_2(k)$

In their paper [26], Segev and Smith investigated the case $G = G_2(k)$ and $V = \overline{C}$, the 7-dimensional Cayley module for G . Denote $\Delta = \Delta(G)$ and $H = H_0(\Delta, \mathcal{F}_{\overline{C}})$. A certain subspace $\overline{D}_p \leq \overline{C}$ of dimension 3 is defined [26, p.495]. This subspace corresponds to a particular subcomplex of the building, and the submodule of H spanned by the sheaf terms on this subcomplex is denoted $D_p \leq H$. Since \overline{C} is a quotient of H , we have that D_p maps onto \overline{D}_p . Segev and Smith show the following [26, p.495, Theorem]:

Theorem *Suppose that W is a quotient of H by some $X \leq H$, such that $\dim(D_p +$*

$X)/X = 3$. Then $\dim W \leq 7$.

If the characteristic of k is not 2 then the module \overline{C} is irreducible, and so $\dim W \leq 7$ implies $W \cong \overline{C}$ [26, p.495]. Thus if $\dim D_p = 3$ and the characteristic of k is not 2 then $H \cong \overline{C}$.

Ronan and Smith did some computations for $G = G_2(k)$ [23, p.335, Example 3.3]. In particular, they showed that if $k = \mathbb{F}_2$ and V is the 6-dimensional irreducible quotient of \overline{C} , then $\dim H_0(\Delta, \mathcal{F}_V) = 14$; and if $k = \mathbb{F}_3$ and $V = \overline{C}$ then $\dim H_0(\Delta, \mathcal{F}_V) = 14$ as well.

Segev and Smith note without proof that if k is a prime field \mathbb{F}_π , then $\dim D_p \in \{3, \pi + 2\}$. We provide a proof of this in Lemma 6.2.4. The consequence of this fact is that if $\dim D_p > 3$ then $\dim H \geq \dim D_p = \pi + 2$. This is true over the fields \mathbb{F}_2 and \mathbb{F}_3 by the computations of Ronan and Smith, but Segev and Smith speculate that it is not the case for an arbitrary finite field \mathbb{F}_q , saying: ‘if it holds in general, the situation would run counter to the intuition of many geometers’ [26, p.497].

Indeed, our proof of Theorem A shows that their speculation was correct. Although we have not been able to prove the result for every $q = \pi^a$, we have shown that $H \cong \overline{C}$ whenever $\pi > 3$ with $3 \nmid \pi - 1$. It seems plausible that the latter condition is unnecessary. By computation using the computer algebra system MAGMA [5], we have shown that $H \cong \overline{C}$ if $q \in \{5, 7, 11\}$, the case $q = 7$ providing some evidence that the condition $3 \nmid \pi - 1$ may be superfluous.

1.3 Related results in the literature

We give a brief survey of the literature concerning sheaves on buildings. Various authors have explored the question of for which pairs (G, V) we have $H_0(\Delta(G), \mathcal{F}_V) \cong V$. The final section of the first Ronan and Smith paper [23, p.337, Section 4] concerns *minimal-weight modules*. A dominant weight λ is *minimal* if there does not exist any

other dominant weight $\mu \neq \lambda$ such that $\lambda - \mu$ is a non-negative linear combination of fundamental roots. It turns out that a minimal weight λ is always equal to some fundamental dominant weight λ_i [15, p.72, Exercise 13].¹ Write $L(\lambda)$ for the unique irreducible kG -module of highest weight λ .

For minimal weight modules, we get an isomorphism in nearly all cases [23, p.338, Theorem 4.1]:

Theorem *Suppose that G is a Chevalley group defined over k . If G is of type C_n then suppose further that k does not have characteristic 2 (otherwise, we may take k to be any finite field \mathbb{F}_q). If M is a minimal weight kG -module then $H_0(\Delta(G), \mathcal{F}_M) \cong M$.*

The root systems of types G_2 , F_4 and E_8 have no minimal weights, so this theory does not help us in those cases.

Ronan shows that if V is the 14-dimensional adjoint module for $G = G_2(k)$ then $H_0(\Delta(G), \mathcal{F}_V) \cong V$ in [20, p.184, Example 4].

Let us specialise for a moment to the case $G = \mathrm{SL}_{n+1}(k)$ of type A_n . Here, all of the fundamental dominant weights are minimal [15, p.72, Exercise 13], and so by the above theorem of Ronan and Smith we have $H_0(\Delta(G), \mathcal{F}_{L(\lambda_i)}) \cong L(\lambda_i)$ for all $1 \leq i \leq n$. In 1991, Fisher published an upper bound for the dimension of $H_0(\Delta(G), \mathcal{F}_{L(\lambda)})$ in the case that λ is a sum of two fundamental weights (with some exceptions) [12, p.119, Theorem]:

Theorem *Let $G = \mathrm{SL}_{n+1}(\mathbb{F}_q)$, with fundamental dominant weights $\{\lambda_1, \dots, \lambda_n\}$. Suppose $1 \leq i < j \leq n$, such that $i + j \notin \{n, n + 1, n + 2\}$. Then*

$$\dim H_0(\Delta(G), \mathcal{F}_{L(\lambda_i + \lambda_j)}) \leq \dim L(\lambda_i) \cdot \dim L(\lambda_j).$$

In 1990, Cohen and Smith provided an answer to the isomorphism question when V is the 26-dimensional natural module for the group $G = F_4(\mathbb{F}_q)$ [9, p.474, Theorem]. The

¹See also Bourbaki [6, p.239-241, Exercises 24 and 5] where these are called *minuscule* weights.

module V is irreducible whenever the characteristic of \mathbb{F}_q is not 3, and in this situation we have $H_0(\Delta(G), \mathcal{F}_V) \cong V$. On the other hand, if the characteristic of the field is 3 then V has a 25-dimensional irreducible quotient V' , and we have $H_0(\Delta(G), \mathcal{F}_{V'}) \cong V'$.

Work has also been undertaken generalising the theory to fields other than \mathbb{F}_q . The adjoint module A for the group $G = \mathrm{SL}_3(k)$ of type A_2 is a minimal weight module, and so by [23, p.338, Theorem 4.1] we have $H = H_0(\Delta(G), \mathcal{F}_A) \cong A$ when $k = \mathbb{F}_q$. Smith and Volklein go further, finding the structure of H for any field k ; in particular, we have $H \cong A$ if and only if k is either an algebraic field extension of \mathbb{Q} , or a perfect field of prime characteristic [29, p.128, Corollary].

Ronan and Smith wrote two further joint papers on the subject. The first, *Universal Presheaves on Group Geometries and Modular Representations* [24], generalises the theory to geometries Γ other than the building of G . It also introduces the concept of a *universal presheaf*: if \mathcal{F}' is a sheaf defined only on a certain subgeometry of Γ (a *stalk*), an *extension* \mathcal{U} of \mathcal{F}' can be constructed with some universal properties. The construction involves recursively computing H_0 of the partial sheaf to ‘fill in’ more of \mathcal{U} . If \mathcal{F}' is restriction of a full sheaf \mathcal{F} to some stalk $\Gamma' \subseteq \Gamma$, and \mathcal{U} is its universal extension, then $H_0(\Gamma, \mathcal{F})$ is a quotient of $H_0(\Gamma, \mathcal{U})$ [24, p.140, Theorem 2.1].

The final paper, *Computation of 2-Modular Sheaves and Representations of $L_4(2)$, A_7 , $3S_6$, and M_{24}* [25] contains a lot of homology computations over a variety of different geometries for the aforementioned groups defined over \mathbb{F}_2 .

In [21], Ronan defines a *dual sheaf* \mathcal{F}^* of a sheaf \mathcal{F} , and presents a proof that the top cohomology module $H^{n-1}(\Delta, \mathcal{F}^*)$ of the dual sheaf is isomorphic to the bottom homology module $H_0(\Delta, \mathcal{F})$ of \mathcal{F} [21, p.266, Theorem 2]. This result, applied to a universal presheaf constructed from data at chambers (maximal simplices) and panels (submaximal simplices), was employed when writing the MAGMA programs which we have used to compute various homology modules throughout the writing of this thesis.

Sheaves on buildings continue to be an active topic of research. In 2015, Ward [34] computed the zero-homology modules over $k = \mathbb{F}_2$ of all the ‘panel-irreducible’ sheaves for the symmetric group S_6 and for the Mathieu groups M_{11} and M_{22} , on a particular type of geometry called a *minimal parabolic system*. The irreducible quotients of the zero-homology modules for Mathieu groups M_{12} , M_{23} and M_{24} were also computed.

1.4 Chapters in this thesis

In Chapter 2 we present the required background material on buildings, discussing Coxeter groups, chamber systems, apartments and buildings themselves. We also give our working definition of a *group of Lie type* as a particular subgroup of the automorphism group of a building.

Chapter 3 introduces the group $G = G_2(k)$. The group is given as the automorphism group of the 8-dimensional Cayley algebra C^+ ; we show that G preserves a 7-dimensional subspace $\overline{C} \subseteq C^+$ called the *Cayley module*. This is the kG -module which we will use to build our fixed-point sheaves. This chapter also introduces the building $\Delta = \Delta(G)$ in the form of a *generalised hexagon*, a point-line geometry constructed using C^+ . We prove some transitivity results and introduce the concept of an *ideal line* which will be key to our proof of the main theorem.

In Chapter 4 we take a brief diversion to talk about weight theory. We introduce the concept of a root system and present some of the more fundamental results, such as the correspondence between the q -restricted weights λ and the irreducible kG -modules $L(\lambda)$. A result of Premet regarding the weights appearing in the module $L(\lambda)$ will be crucial in a later chapter when we want to bound the dimension of $H_0(\Delta, \mathcal{F}_{\overline{C}})$ from below.

Chapter 5 covers in detail the sheaves on buildings introduced in [23]. The relevant objects are defined, we discuss how their homology is computed, and then we present the important results from Ronan and Smith.

We finally reach the problem at hand in Chapter 6, where the result of Segev and Smith is discussed. Since Segev and Smith have proved that $H_0(\Delta, \mathcal{F}_{\overline{C}}) \cong \overline{C}$ if $\dim D_p = 3$, we introduce the standing assumption that $\dim D_p > 3$ for the remainder of the thesis. We also take the opportunity here to present some dimension data for small fields computed using MAGMA.

The next five chapters contain the proof of the main theorem. We use a combination of a weight-theoretic argument to determine which modules may be involved in H , and combinatorial arguments in the building to bound the dimension of H from above. The proof begins with an argument in Chapter 7 which allows us to concern ourselves only with the case where k is a prime field; the proposition presented here gives us all the cases $q = \pi^a$ for free, once $q = \pi$ is dealt with.

In Chapter 8 we give a lower bound on the dimension of $H_0(\Delta, \mathcal{F}_{\overline{C}})$ when k is a prime field, in the case that $\dim D_p > 3$. We begin by finding the weight of some quotient of H , and then applying Premet's result to bound its dimension from below.

Chapters 9 and 10 both contain upper bounds on the dimension of H . Chapter 9 presents a cubic bound, which is then refined to a quadratic bound in Chapter 10. Here we introduce new terminology for parts of the generalised hexagon, defining *pages*, *spines*, *diagonals* and *crossbraces*. The bound in Chapter 9 holds for all $\pi > 3$; it is in Chapter 10 that we hit a problem when $3 \mid \pi - 1$.

Finally, Chapter 11 contains the proof of the main theorem, which merely amounts to noting that the upper and lower bounds which we obtain when $3 \nmid \pi - 1$ cannot hold simultaneously if $\pi > 3$.

CHAPTER 2

BUILDINGS

One of the most informative ways to learn about a group is to study its action on some kind of geometric structure. The types of geometries studied include graphs, as well as their higher-dimensional analogues such as a simplicial complexes.

An important class of finite groups are the *groups of Lie type*. The geometries usually associated with groups of Lie type are called *buildings*. The theory of buildings was developed from the late 1950s onward by French mathematician Jacques Tits, culminating in his 1974 book *Buildings of Spherical Type and Finite BN-pairs* [32]. In this chapter we will introduce the buildings first, and then construct the group of Lie type G as a certain subgroup of the automorphism group of the building.

We borrow notations from a variety of sources. *Buildings* by Abramenko and Brown [1] is a very thorough treatment covering both Tits' original simplicial approach, as well as a more modern definition using a distance function taking values from a Weyl group W . We follow Ronan's *Lectures on Buildings* [22] for the preliminary material on chamber systems, as well as the definitions of G and its subgroups. *Points and Lines* by Shult [27] takes an approach based on *point-line geometries* and is our main reference for generalised polygons. Other excellent books include *Buildings and Classical Groups* by Garrett [13], and *The Structure of Spherical Buildings* by Weiss [35].

2.1 Simplicial complexes

In Tits' original definition, a building is described as a type of structure called a *simplicial complex*. A simplicial complex is often considered as the union of a collection of objects called *simplices*, which are polytopes embedded into a Euclidean space. Initially, however, we will take a purely combinatorial approach where we do not worry about the Euclidean space in which the simplices lie, merely recording what the simplices are and how they relate to one another. This definition is often called an *abstract simplicial complex* in the literature.

Definition 2.1.1 (Simplicial complex) *An (abstract) **simplicial complex** Δ on a vertex set V is a set of non-empty subsets of V such that:*

(S1) *If $\sigma \in \Delta$ and τ is a non-empty subset of σ , then $\tau \in \Delta$.*

(S2) *$\{v\} \in \Delta$ for every $v \in V$.*

*The elements of Δ are called **simplices**; a simplex with $n + 1$ elements is called an n -simplex.*

If $\sigma, \tau \in \Delta$ with $\tau \subseteq \sigma$ then we say that ' τ is a face of σ ', and we write $\tau \prec \sigma$. The maximal faces of Δ (those simplices which are not a face of any other simplex) are called **facets**. The facets of a simplicial complex completely determine the entire complex, since simplicial complexes are closed under taking subsets by axiom (S1).

We will borrow terminology from the traditional definition of a simplicial complex, where an n -simplex is an n -dimensional analogue of a triangle. Hence we will often refer to 0-simplices $\{v\}$ as **vertices**, 1-simplices $\{v_1, v_2\}$ as **edges** and 2-simplices $\{v_1, v_2, v_3\}$ as **triangles**. (We refer to a simplex containing $n + 1$ points as an ' n -simplex' because it is an n -dimensional polytope — so the terminology is less unusual than it looks at first.)

Also, when drawing figures we will often represent the simplices in this way, for improved clarity over trying to depict subsets of a vertex set.

Since we will be investigating groups G acting on a simplicial complex Δ , we need to define a simplicial isomorphism. First, a simplicial map:

Definition 2.1.2 (Simplicial map) *Let Δ_1 and Δ_2 be simplicial complexes with vertex sets $V(\Delta_1)$ and $V(\Delta_2)$ respectively. A map $f : V(\Delta_1) \rightarrow V(\Delta_2)$ is a **simplicial map** if for every k -simplex $\sigma \in \Delta_1$ spanned by vertices $\{v_0, \dots, v_k\}$, the vertices $\{f(v_0), \dots, f(v_k)\}$ span a k -simplex $f(\sigma)$ in Δ_2 .*

A simplicial map takes simplices to other simplices; an isomorphism of simplicial complexes requires this property in both directions:

Definition 2.1.3 (Simplicial isomorphism) *Let $f : V(\Delta_1) \rightarrow V(\Delta_2)$ be a simplicial map. If f is a bijection and f^{-1} is also a simplicial map then we say that f is a **simplicial isomorphism**.*

A simplicial isomorphism from Δ to itself is called a **simplicial automorphism**. The set of simplicial automorphisms of Δ forms a group $\text{Aut}(\Delta)$ under composition.

Let G be a group. We say that G acts (on the right) on a simplicial complex Δ if each $g \in G$ corresponds to some simplicial automorphism $\tilde{g} : V(\Delta) \rightarrow V(\Delta)$, such that $\tilde{g}_1 \tilde{g}_2 = (g_1 g_2)^\sim$ for all $g_1, g_2 \in G$.

2.2 Coxeter groups and chamber systems

Buildings are comprised of substructures called *apartments*, which are constructed from groups called *Coxeter groups*.

2.2.1 Coxeter groups

We begin with a definition from [22, p.9].

Definition 2.2.1 (Coxeter group) A **Coxeter group** is a finitely-presented group which admits a presentation of the following form:

$$W = \langle s_1, s_2, \dots, s_n \mid (s_i s_j)^{m_{ij}} = 1 \rangle,$$

where each $m_{ii} = 1$ (so that all generators are involutions) and $m_{ij} \geq 2$ for $i \neq j$. To avoid redundancy, we assume that $m_{ij} = m_{ji}$ for all $1 \leq i, j \leq n$.

Examples of Coxeter groups include the finite dihedral group of order $2m$

$$D_{2m} := \langle s, t \mid s^2 = t^2 = (st)^m = 1 \rangle, \quad (2.2.2)$$

and the infinite dihedral group

$$D_\infty := \langle s, t \mid s^2 = t^2 = 1 \rangle, \quad (2.2.3)$$

in which the product (st) has infinite order. The symmetric group S_n (for $n \geq 2$) is also a Coxeter group, since

$$S_{n+1} \cong \langle s_1, \dots, s_n \mid s_i^2 = 1, (s_i, s_{i+1})^3 = 1 \forall i < n, (s_i s_j)^2 = 1 \forall |i - j| > 1 \rangle, \quad (2.2.4)$$

by the isomorphism $(i, i + 1) \mapsto s_i$.

A **Coxeter system** is a pair (W, S) consisting of a Coxeter group W and its identified set of generating elements S [35, p.9]. (The choice of S is not necessarily unique for a given Coxeter group W , but is specified in a Coxeter system.) The **rank** of the Coxeter system (W, S) is the number of generators n in S .

This information can be encoded in a **Coxeter diagram**, which is a graph with one vertex for each generator s_i , and the following rules for edges:

- If $m_{ij} \geq 3$, then join vertices i and j with an edge.
- If $m_{ij} \geq 4$, then label the edge ij with m_{ij} .


Thus the Coxeter system for S_{n+1} as shown in Equation 2.2.4 can be encoded with the following diagram



where there are n nodes in total. We say that this Coxeter diagram is of *type* A_n . Similarly, the diagrams of type B_n (for $n \geq 2$)¹



and of type D_n (for $n \geq 4$)



both contain n nodes. In general, the subscript number refers to the rank of the Coxeter system. The last example we will give is the Coxeter diagram of type $I_2(m)$ for $m \geq 3$



which corresponds to the Coxeter group D_{2m} (as per Equation 2.2.2).

The diagrams corresponding to *finite* Coxeter groups consist of three infinite families A_n, B_n and D_n whose rank depends on n , one further infinite family $I_2(m)$ in which all groups have rank 2, and six diagrams corresponding to exceptional groups E_6, E_7, E_8, F_4, H_3 and H_4 [36, p.33, Table 2.1].

In Chapter 4 we will meet a kind of finite reflection group called a *Weyl group*, which is a *crystallographic* Coxeter group (the stabiliser of a lattice). All Weyl groups are Coxeter

¹The Coxeter type B_n is also called C_n in some sources. This due to the fact that two non-isomorphic root systems B_n and C_n can be constructed (for $n > 2$), with identical angles between roots (and hence identical Coxeter diagrams) but differing root lengths. We discuss root systems in more detail in Chapter 4.

groups, although the converse is not true; the Coxeter groups of type H_3 and H_4 are not Weyl groups, and neither is the Coxeter group of type $I_2(m)$ for $m \notin \{3, 4, 6\}$ [36, p.34, Section 2.8.4].

Note that $I_2(3) = A_2$ and $I_2(4) = B_2$, so only it is only $I_2(6)$ which is not included in one of the other infinite families. In the language of Weyl groups and groups of Lie type, the diagram $I_2(6)$ is given the special name G_2 . The group of Lie type $G_2(k)$ is the main focus of this thesis.

2.2.2 Chamber systems

Another way to graphically represent a Coxeter system is a **Cayley graph**. Here, the graph vertices are the group elements, and two vertices g and h are joined by an edge labelled i if $g = s_i h$. (Note that if this is the case then $h = s_i g$ as well, since the generators are involutions. So the edges are bi-directional.)

A Cayley graph is an example of a more general type of structure called a chamber system [22, p.1], which is a set along with some indexed adjacency relations:

Definition 2.2.5 (Chamber system) *Let I be an index set, and Σ be a set. We say that Σ is a **chamber system over I** if each element $i \in I$ determines a partition of Σ . We refer to the elements of Σ as the **chambers**. If $x, y \in \Sigma$ lie in the same i -partition then we say that they are ‘ i -adjacent’, and write $x \sim_i y$.*

The following example is very important.

Example 2.2.6 *The Cayley graph of a Coxeter system (W, S) is a chamber system over the index set S ; take $\Sigma = W$, and say that g and h are i -adjacent if $g = s_i h$. Here, the vertices represent chambers and the edges represent adjacency. (In this example each $1 \leq i \leq n$ partitions W into pairs.) We call this chamber system $\Sigma(W, S)$.*

Example 2.2.7 *Suppose that $W = \langle s_1, s_2 \mid s_i^2 = 1, (s_1 s_2)^6 = 1 \rangle$ so that W has type $I_2(6) = G_2$. Then the Cayley graph of W is as depicted in Figure 2.1.*




Figure 2.1: The Cayley graph of $W = \langle s_1, s_2 \mid s_i^2 = 1, (s_1s_2)^6 = 1 \rangle$, the Coxeter group of type $I_2(m) = G_2$. The vertex opposite id is the element $s_2s_1s_2s_1s_2s_1 = s_1s_2s_1s_2s_1s_2$.

Residues and connectedness

We say that two chambers are **adjacent** if they are i -adjacent for some $i \in I$. A **gallery** is a sequence c_1, \dots, c_{n+1} of chambers such that c_j is adjacent to c_{j+1} for all $1 \leq j \leq n$. The gallery has **type** (i_1, \dots, i_n) , where c_j is i_j -adjacent to c_{j+1} . (In an arbitrary chamber system a gallery may have more than one type, but in the chamber system $\Sigma(W, S)$ from Example 2.2.6 this is not the case as it would imply $s_i = s_j$ for some $i \neq j$.) We say that a chamber system is **connected** if there exists a gallery between any two chambers x and y in C . (This is equivalent to Σ being a connected graph.) A **minimal gallery** is a shortest gallery between two chambers.

We can also consider galleries using only adjacencies from some subset $J \subseteq I$. This gives us the notion of a chamber system being ‘ J -connected’. The J -connected components are called **J -residues**, or **residues of type J** . The **cotype** of a J -residue is $I \setminus J$.

The **rank** of a chamber system is $|I|$; the rank of a J -residue is $|J|$, and its **corank** is $|I \setminus J|$. Note that a residue of rank 0 (*i.e.* a residue of type \emptyset) is a chamber. We refer to

a residue of rank 1 as a **panel**. (If it has type i , we may call it an i -panel.)

2.2.3 The geometric realisation of a chamber system

Given a chamber system Σ , we can construct a simplicial complex $\Delta(\Sigma)$ using its residues. Let R and S be residues of Σ of types J and K respectively. We say that S is a **face** of R if $S \supset R$ and $K \supset J$. A residue R of corank r has r faces S of corank $r - 1$, as we obtain one face from adding each element of the cotype of R into the type of S . Similarly, R has r faces S of corank 1; these are obtained by adding all but one elements of the cotype of R into the type of S .

We use the following algorithm to obtain $\Delta = \Delta(\Sigma)$ from Σ .

Algorithm 2.2.8 (A1) *Let the vertex set $V = V(\Delta)$ contain one vertex v for each residue R of corank 1. Then since Δ is a simplicial complex, we have a 0-simplex $\{v\} \in \Delta$ corresponding to each R by axiom (S2).*

(A2) *For each residue S of corank 2, let $v_1, v_2 \in V$ be the vertices corresponding to the two faces of S of corank 1, and add the simplex $\{v_1, v_2\}$ to Δ to correspond to S .*

(A3) *Continue in this way: for each residue T of corank r , add the $(r - 1)$ -simplex $\{v_1, \dots, v_r\}$ to Δ , where v_1, \dots, v_r are the vertices corresponding to the faces of T of corank 1.*

We give two examples of this algorithm being used in practice.

Example 2.2.9 *Figure 2.2 shows the result of Algorithm 2.2.8 applied to the Cayley graph of S_4 with Coxeter generators $s_1 = (1, 2)$, $s_2 = (2, 3)$ and $s_3 = (3, 4)$, where an n -simplex is represented in the figure by an n -dimensional polytope. The vertices marked with stars and circles correspond to residues of cotypes 1 and 3 respectively; the vertices marked with squares correspond to residues of cotype 2. A maximal simplex is highlighted, and then Figure 2.3 shows the stabiliser of each face of this maximal simplex in the group S_4 .*




Figure 2.2: The Coxeter complex of S_4 with a maximal simplex (chamber) highlighted. Vertices and edges on the underside and rear surfaces have been omitted for clarity.




Figure 2.3: The stabilisers of each face of one of the maximal simplices (chambers) of the Coxeter complex of S_4 .




Figure 2.4: The geometric realisation of the Cayley graph of $W = \langle s_1, s_2 \mid s_i^2 = 1, (s_1 s_2)^6 = 1 \rangle$.

Example 2.2.10 *Figure 2.4 shows the result of Algorithm 2.2.8 applied to the Cayley graph of type G_2 from Example 2.2.7. Note that in rank 2, the geometric realisation is merely a graph and can look very similar to the Cayley graph; however, the role of edges and vertices are swapped, with the chambers being edges and the panels being vertices.*

2.3 Apartments and buildings

A simplicial complex Δ is called a **Coxeter complex** if it is isomorphic to $\Delta(\Sigma(W, S))$ for some Coxeter system (W, S) .

2.3.1 Buildings

We can now give our definition of a building:

Definition 2.3.1 (Building) *A **building** Δ is a simplicial complex which is the union of subcomplexes called **apartments**, satisfying the following conditions:*

(B1) *Each apartment is a Coxeter complex.*

(B2) *Any two simplices in Δ lie in some common apartment \mathcal{A} .*

(B3) For any two simplices σ_1 and σ_2 in Δ , and two apartments \mathcal{A}_1 and \mathcal{A}_2 containing them both, there exists a simplicial complex isomorphism $\mathcal{A}_1 \rightarrow \mathcal{A}_2$ which fixes σ_1 and σ_2 pointwise. (By fixing a simplex pointwise, we mean that every vertex of the simplex is fixed.)

Notice that the definition of a building does not specify the choice of system of apartments; it is simply a simplicial complex Δ which *admits* a system of apartments such that the axioms hold. However, for any building Δ there is a canonical choice of apartments, and in the case where each apartment is a finite Coxeter complex, this choice is in fact unique [1, p.174].

Condition (B3) implies that the apartments are pairwise isomorphic. Hence there is a particular Coxeter complex (and thus Coxeter system (W, S)) associated with a building Δ .

A **thick** building is one where every simplex of codimension 1 is a face of at least three maximal simplices. (Equivalently, every panel lies on at least three chambers.) All other buildings are **thin**, and consist of a single apartment [22, p.29]. A building is **spherical** if its apartments are finite Coxeter complexes. (The name comes from the fact that such apartments are triangulations of a sphere [8, p.1245].) Finally, a building is **irreducible** if the Coxeter diagram of its associated Coxeter system (W, S) is connected.

2.3.2 Roots and the Moufang property

Suppose that Δ is a Coxeter complex for the Coxeter system (W, S) , and let $r = s_i^w$ be a W -conjugate of one of the elements of S . Since the generators s_i all have order 2, we have that r is an involution. Define the **wall** corresponding to r as $M_r := \{\sigma \in \Delta \mid \sigma \cdot r = \sigma\}$. A wall contains no chambers (since only the identity of W fixes chambers), but does contain panels because the involution $r = s_i^w$ swaps the two chambers w and $s_i w$ which form an i -panel. Thus in the Coxeter complex Δ , the subcomplex M_r has codimension 1

[22, p.13].

Say that a gallery c_1, \dots, c_n *crosses the wall* M_r if r swaps consecutive chambers c_i and c_{i+1} . A gallery may cross a wall multiple times, but a minimal gallery may not; furthermore, the number of times a minimal gallery from c to c' crosses M_r is independent of the choice of gallery [22, p.13, Lemma 2.5 (i)]. Therefore, the wall M_r divides the Coxeter complex Δ into two halves called **roots**; choosing any initial chamber c , we have the root containing those chambers c' for which a minimal gallery c, \dots, c' does not cross the wall M_r , and the root containing chambers for which it does. If α is one root, let the root on the other side of the wall be denoted $-\alpha$. Roots are **convex**, meaning that any minimal gallery between two chambers of a root α lies entirely within α [22, p.14, Proposition 2.6 (i)].

Opposite chambers and opposite simplices

Let c, c' be chambers in a Coxeter complex Δ . We say that c and c' are **opposite** if the length of a minimal gallery from c to c' is equal to the diameter of Δ (here, we mean the graph diameter in the associated chamber system). If c and c' are opposite then they do not lie in any common root α , and so every wall M_r separates c and c' . Furthermore, for any chamber $d \in \Delta$ there exists a minimal gallery γ from c to c' containing d . Finally, every chamber of a Coxeter complex has a unique opposite [22, p.20, Theorem 2.15 (iii)]. Denote the map sending every chamber of Δ to its opposite by op_Δ .

We would like to extend the notion of an ‘opposite’ to simplices of any dimension, rather than just chambers. In fact, for any k -simplex $\{v_0, \dots, v_k\}$ we have that $\{\text{op}_\Delta(v_0), \dots, \text{op}_\Delta(v_k)\}$ is another k -simplex, and so we can define the map $\text{op}_\Delta : \Delta \rightarrow \Delta$ for simplices of any dimension.

In the more general case, when Δ is a building, we write $\text{op}_\mathcal{A} : \mathcal{A} \rightarrow \mathcal{A}$ for the map sending each simplex to its opposite in the apartment \mathcal{A} . We say that two simplices σ_1 and σ_2 are **opposite** if they are opposite in some apartment \mathcal{A} of Δ .

Define $\Phi(\mathcal{A})$ to be the set of roots contained in \mathcal{A} . For c a chamber in \mathcal{A} , let $\Phi_c^+(\mathcal{A})$ be the set of roots of \mathcal{A} containing the chamber c . These are called the **positive roots with respect to** c . Then we define $\Phi_c^-(\mathcal{A}) := \Phi_{c'}^+(\mathcal{A})$, where $c' = \text{op}_{\mathcal{A}}(c)$ is the unique chamber opposite c in \mathcal{A} . Thus we have that

$$\Phi(\mathcal{A}) = \Phi_c^+(\mathcal{A}) \cup \Phi_c^-(\mathcal{A})$$

is a disjoint union.

The Moufang property

Now suppose that Δ is an irreducible, spherical building whose apartments are isomorphic to the Coxeter complex for (W, S) . The following definition comes from [22, p.66]:

Definition 2.3.2 (Root group) *Let \mathcal{A} be an apartment of Δ , and let α be a root of \mathcal{A} corresponding to a reflection $r \in W$. Define the **root group***

$$X_\alpha := \{g \in \text{Aut}(\Delta) \mid g \text{ fixes every chamber with a panel in } \alpha \setminus M_r\}.$$

It follows from the definition that X_α fixes the root α pointwise [1, p.385]. We say that Δ is **Moufang** if, for every root α , the root group U_α is transitive on the set of apartments containing α [22, p.66]. The following theorem is due to Tits [32, p.274, (1)] (see also [22, p.67, Corollary 6.7]).

Theorem 2.3.3 *Suppose that Δ is a thick, irreducible, spherical building of rank ≥ 3 . Then Δ is Moufang.*

In the rank 2 case, the situation is more complex. Rank 2 thick, irreducible, spherical buildings are equivalent to *generalised polygons*, which will be defined in Chapter 3. The Moufang generalised polygons have been classified by Tits and Weiss in [33]. In particular,

Feit and Higman showed in [11] that they only exist for W of type $I_2(3) = A_2$ (the *generalised projective planes*), $I_2(4) = B_2$ (the *generalised quadrangles*), $I_2(6) = G_2$ (the *generalised hexagons*) and $I_2(8)$ (the *generalised octagons*).

2.4 The finite groups of Lie type

The precise definition of a *group of Lie type* varies between sources, with certain groups being included by some authors but not by others. The collection of *finite groups of Lie type*, however, is more widely agreed upon. We follow the exposition from [22, p.76], and take the following definition from there.

Definition 2.4.1 (Finite group of Lie type) *Suppose that Δ is a thick, irreducible, spherical building which satisfies the Moufang property, and fix once and for all an apartment \mathcal{A}_0 of Δ . We define*

$$G := \langle X_\alpha \mid \alpha \in \Phi(\mathcal{A}_0) \rangle \leq \text{Aut}(\Delta),$$

*and call G a **finite group of Lie type**.*

Notice that by axioms (B2) and (B3), the group $\text{Aut}(\Delta)$ is transitive on the set of apartments. Therefore the choice of apartment \mathcal{A}_0 does not change the isomorphism class of the group G .

Lemma 2.4.2 *The action of G on Δ is type-preserving.*

Proof. By [1, p.175, Proposition 4.6], the building Δ admits a *type function*, which means that we can assign each vertex of Δ a type $i \in I$ such that every chamber has exactly one vertex of each type. (The proof of this essentially uses the axiom (B3) to extend the types of some apartment \mathcal{A} to the entire building Δ .) A type function on a chamber complex is unique up to some reordering of the index set I , since once the types are fixed on one

chamber, these choices propagate automatically to any adjacent chamber [1, p.665]. So if an automorphism $g \in \text{Aut}(\Delta)$ fixes one chamber $c \in \Delta$ pointwise, then it fixes the type of every vertex $v \in \Delta$. Every root group X_α fixes a root α (and thus a chamber $c \in \alpha$) pointwise, and so G is type-preserving. \square

2.4.1 The subgroups U , B , T and N

Fix a chamber $c_0 \in \mathcal{A}_0$, and define

$$U^\varepsilon := \langle X_\alpha \mid \alpha \in \Phi_{c_0}^\varepsilon(\mathcal{A}_0) \rangle,$$

for $\varepsilon \in \{+, -\}$. We will often write U for U^+ . Notice that $U = U^+$ fixes the chamber c_0 and U^- fixes $\text{op}_{\mathcal{A}_0}(c_0)$, the unique chamber opposite c_0 in \mathcal{A}_0 . The following theorem tells us about the structure of U^ε when the building is finite [22, p.107, Theorem 8.7]:

Theorem 2.4.3 *Suppose Δ is a finite building. Then there is some fixed prime π such that U^ε is a Sylow π -subgroup of G .*

Let T be the subgroup of G which fixes every chamber of \mathcal{A}_0 . Now we define

$$B^\varepsilon := \langle T, U^\varepsilon \rangle,$$

for $\varepsilon \in \{+, -\}$. We will often write B for B^+ .

Theorem 2.4.4 $B^\varepsilon = U^\varepsilon \rtimes T$.

Proof. See [22, p.77, Theorem 6.17]. \square

By [22, p.76, Proposition 6.16], the group $B = B^+$ is the stabiliser in G of the chamber c_0 .

Theorem 2.4.5 *Suppose Δ is finite. Then $U^\varepsilon = O_\pi(B^\varepsilon)$.*

Proof. By Theorem 2.4.3, we have that U^ε is a Sylow π -subgroup of G and hence of B^ε . Since $U^\varepsilon \trianglelefteq B^\varepsilon$ by Theorem 2.4.4, we have $U^\varepsilon = O_\pi(B^\varepsilon)$. \square

For α a root of \mathcal{A}_0 , let $r(\alpha)$ be the reflection corresponding to the wall M_r separating α and $-\alpha$. Then we define

$$N := \langle T, r(\alpha) \mid \alpha \in \Phi(\mathcal{A}_0) \rangle.$$

The names B and T come from the language of Lie groups, in which the corresponding subgroups are called the *Borel* subgroup and the *torus* respectively. The subgroups B and N generate G , and form a what is called a *BN-pair* [22, p.76, Proposition 6.16], or a *Tits system*. This pair actually encodes all of the information required to construct the building, and any group with a *BN-pair* has a corresponding building; see [22, Chapter 5] for details.

2.4.2 Parabolic subgroups

A subgroup of G which stabilises some simplex of Δ is called a **parabolic subgroup**. We write the parabolic subgroup stabilising a simplex $\sigma \in \Delta$ as P_σ .

Lemma 2.4.6 *Suppose that $\tau \prec \sigma$. Then $P_\sigma \subseteq P_\tau$.*

Proof. The action of G is type-preserving by Lemma 2.4.2, so if an element $g \in G$ fixes a simplex σ then it also fixes any $\tau \prec \sigma$, and so $P_\sigma \subseteq P_\tau$. \square

The minimal parabolics, therefore, are the stabilisers of maximal simplices (chambers)¹. Hence they are conjugates of the subgroup B .

¹The use of the term ‘minimal parabolic’ in the literature is not consistent — sometimes it refers to the stabiliser of a chamber, as we use it here, and sometimes to the stabiliser of a panel, which we will call a ‘rank 1 parabolic’.

Suppose that $\sigma \in \mathcal{A}_0$ is a face of c_0 . Let $\Theta_\sigma = \{c \in \mathcal{A}_0 \mid c \succ \sigma, c \text{ a chamber}\}$ and $\Phi_\sigma = \{\alpha \in \Phi(\mathcal{A}) \mid \sigma \in M_{r(\alpha)}\}$. We define

$$U_\sigma := \langle X_\alpha \mid \Theta_\sigma \subset \alpha \in \Phi(\mathcal{A}) \rangle$$

and

$$L_\sigma := \langle T, X_\alpha \mid \alpha \in \Phi_\sigma \rangle.$$

Note that $c_0 \in \Theta_\sigma$, so if α is a root of \mathcal{A}_0 containing every chamber in Θ_σ , then certainly α contains c_0 . Hence $U_\sigma \leq U$. Furthermore, we have $\langle U_\sigma \mid \sigma \prec c_0, \sigma \text{ a panel} \rangle = U$.

Also, note that the sets of root subgroups generating the groups U_σ and L_σ are disjoint: if X_α is one of the generators of U_σ , then the root α contains every chamber on σ ; but if X_α is among the generators of L_σ , then σ is in the wall $M_{r(\alpha)}$ corresponding to α , and hence there exist chambers on σ which lie in α as well as chambers on σ which lie in $-\alpha$.

Corresponding subgroups $U_\sigma, L_\sigma \leq P_\sigma$ for arbitrary simplices $\sigma \in \Delta$ can be found via conjugation of the parabolics containing c_0 . Then we have the following theorem [22, p.78, Theorem 6.18]

Theorem 2.4.7 *Let Δ be an irreducible, spherical building which satisfies the Moufang property. Then for each simplex $\sigma \in \Delta$, we have $P_\sigma = U_\sigma \rtimes L_\sigma$.*

We call L_σ a **Levi complement** of the parabolic P_σ , and U_σ the **unipotent radical** of P_σ . In general, a Levi complement of a parabolic P_σ is the intersection $P_\sigma \cap P_{\sigma'}$, where σ' is any simplex opposite σ . We then have $P_\sigma = U_\sigma \rtimes L_\sigma$ for any Levi complement L_σ of P_σ .

Lemma 2.4.8 *Suppose that $\tau \prec \sigma$, so that $P_\sigma \subseteq P_\tau$. Then $U_\sigma \supseteq U_\tau$.*

Proof. Since $\tau \prec \sigma$, if a chamber $c \succ \sigma$ then $c \succ \tau$ as well. Hence we have $\Theta_\tau \supseteq \Theta_\sigma$. Now, if $\alpha \supset \Theta_\tau$ then certainly we have $\alpha \supset \Theta_\sigma$. Therefore $U_\sigma \supseteq U_\tau$. \square

2.5 The classification of finite, thick, irreducible buildings

If Δ is a finite building then its apartments must also be finite, so Δ is spherical and its apartments are isomorphic to the Coxeter complex of a finite Coxeter group W . Moreover, a theorem of Feit and Higman restricts the choice of W further. The version below is a generalisation as stated in [1, p.337, Theorem 6.94] (see also [11]):

Theorem 2.5.1 (Feit-Higman Theorem) *Suppose Δ is a finite, thick, irreducible building, in which the apartments are Coxeter complexes for a Coxeter group W . Then W is a Weyl group, and so must have type $A_n, B_n, D_n, G_2, F_4, E_6, E_7$ or E_8 .*

The subgroup $G \leq \text{Aut}(\Delta)$ from Definition 2.4.1 has an associated *root system* Φ , which may be of a different type to that of the building. (Root systems will be introduced later on, in Chapter 4.) Indeed, a given Coxeter type may give rise to multiple isomorphism classes of buildings, all with a different corresponding group of Lie type G .

The classification of finite, thick, irreducible buildings of rank ≥ 3 was given by Tits in 1974 [32]. The rank 2 case is much harder, and indeed a full classification may not be possible (see [1, p.500] for a discussion on this matter); but the *Moufang* buildings of rank 2 have been classified by Tits and Weiss [33]. Note that a finite building of rank 1 is merely a finite set of vertices (that is, a permutation representation of G).

Table 2.5, condensed from [22, p.191, Appendix 6], contains all finite, thick, irreducible buildings of rank ≥ 3 , as well as all finite *Moufang* buildings of rank 2, and two further buildings of rank 1. The second column shows the root system type of the corresponding group of Lie type $G \leq \text{Aut}(\Delta)$ for each building.

The groups appearing in the table consist of the **Chevalley groups**, which include the *classical groups* of types A_n, B_n, C_n and D_n as well as the *exceptional groups* of type

Type of Δ	Type of G
A_n	$A_n(q)$
B_n	$B_n(q)$
B_n	$C_n(q)$
B_n	${}^2A_{2n-1}(q)$
B_n	${}^2A_{2n}(q)$
B_n	${}^2D_{n+1}(q)$
D_n	$D_n(q)$
E_6	$E_6(q)$
E_7	$E_7(q)$
E_8	$E_8(q)$
F_4	$F_4(q)$
F_4	${}^2E_6(q)$
G_2	$G_2(q)$
G_2	${}^3D_4(q)$
$I_2(8)$	${}^2F_4(q), q = 2^{2k+1}$
A_1	${}^2B_2(q), q = 2^{2k+1}$
A_1	${}^2G_2(q), q = 3^{2k+1}$

Table 2.1: A partial classification of finite, thick, irreducible buildings, with the corresponding group of Lie type $G \leq \text{Aut}(\Delta)$. Condensed from [22, p.191, Appendix 6].

G_2, F_4, E_6, E_7 and E_8 ; the **Steinberg groups** of types ${}^2A_n, {}^2D_n, {}^3D_4$ and 2E_6 ; and the **Suzuki-Ree** groups of type ${}^2B_2, {}^2G_2$ and 2F_4 . Together, these are the *finite groups of Lie type*.

The superscript 2 or 3 before a root system type indicates a modification to the construction of the group, involving an automorphism (or order 2 or 3 respectively) of either the Coxeter diagram (in the case of the Steinberg groups), or the field over which the group is defined (in the case of the Suzuki-Ree groups). Details of the construction of these groups are beyond the scope of this thesis; an excellent reference is *The Finite Simple Groups* by Wilson [36, Chapters 3, 4 and 5].

CHAPTER 3

$G_2(k)$: THE CAYLEY ALGEBRA AND THE GENERALISED HEXAGON

In this chapter we will introduce a construction of the group of Lie type $G_2(k)$, where k is a field, as the automorphism group of an 8-dimensional algebra over k . We will also use the algebra to construct the building of $G_2(k)$ in the form of a point-line geometry. This construction is well-known, although the details seem to be somewhat hard to find all in one place. (We borrow notation mostly from two sources, Wilson [36] and Segev and Smith [26].)

3.1 The Cayley algebra

We begin with the definition of an *alternative algebra*; alternativity is a weakened version of associativity.

Definition 3.1.1 (Alternative algebra) An *algebra* over a field k is a k -vector space equipped with a bilinear product. An *alternative algebra* is an algebra A which further satisfies the relations

$$(x \cdot x) \cdot y = x \cdot (x \cdot y)$$

and

$$(x \cdot y) \cdot y = x \cdot (y \cdot y)$$

for all $x, y \in A$.

We see that associative algebras are indeed alternative, since both alternativity relations follow immediately from the associativity condition $(x \cdot y) \cdot z = x \cdot (y \cdot z)$.

Let k be a finite field of order $q = \pi^a$, for π a prime, and C^+ be the Cayley algebra defined over k (our description of this algebra below uses the basis from Wilson [36, p.123]). This is an 8-dimensional alternative algebra with basis $\{e_1, e_2, \dots, e_8\}$, for which we often use the shorthand $\{\bar{1}, \dots, \bar{8}\}$, and multiplication table as follows:

	$\bar{1}$	$\bar{2}$	$\bar{3}$	$\bar{4}$	$\bar{5}$	$\bar{6}$	$\bar{7}$	$\bar{8}$
$\bar{1}$					$\bar{1}$	$\bar{2}$	$-\bar{3}$	$-\bar{4}$
$\bar{2}$			$-\bar{1}$	$\bar{2}$			$-\bar{5}$	$\bar{6}$
$\bar{3}$		$\bar{1}$		$\bar{3}$		$-\bar{5}$		$-\bar{7}$
$\bar{4}$	$\bar{1}$			$\bar{4}$		$\bar{6}$	$\bar{7}$	
$\bar{5}$		$\bar{2}$	$\bar{3}$		$\bar{5}$			$\bar{8}$
$\bar{6}$	$-\bar{2}$		$-\bar{4}$		$\bar{6}$		$\bar{8}$	
$\bar{7}$	$\bar{3}$	$-\bar{4}$			$\bar{7}$	$-\bar{8}$		
$\bar{8}$	$-\bar{5}$	$-\bar{6}$	$\bar{7}$	$\bar{8}$				

Here, a blank entry indicates that the product is 0. Notice that, whilst an arbitrary element of C^+ is a linear combination $\sum_{i=1}^8 \lambda_i e_i$, this particular choice of basis gives us a very neat multiplication table in which all the products are given by $\pm e_i$ for some $1 \leq i \leq 8$, rather than by linear combinations of multiple basis elements.

The light grey lines on the table are there as a visual aid, to highlight another key property of this basis: adding together the fourth and fifth rows, or the fourth and fifth columns, gives us the basis elements $\bar{1}$ to $\bar{8}$ in order. Hence the element $e := \bar{4} + \bar{5}$ acts as

the multiplicative identity from both sides. An algebra such as C^+ with a multiplicative identity is called **unital**. The identity element is necessarily unique; if \bar{e} and \bar{f} are both multiplicative identity elements then $\bar{e} = \bar{e}\bar{f} = \bar{f}$.

We can immediately see from the multiplication table that C^+ is not associative; for example, we have

$$\bar{6} = (\bar{1} \cdot \bar{6}) \cdot \bar{8} \neq \bar{1} \cdot (\bar{6} \cdot \bar{8}) = 0.$$

3.1.1 A bilinear form on C^+

We will equip C^+ with a symmetric bilinear form, as per [26, p.498]¹. If $\bar{x} = \sum_{i=1}^8 \lambda_i e_i$ is an element of C^+ , we write \bar{x}_i for the coefficient λ_i for each $1 \leq i \leq 8$.

Definition 3.1.2 (Bilinear form $b(\cdot, \cdot)$) Let $\bar{x}, \bar{y} \in C^+$. We define

$$b(\bar{x}, \bar{y}) := (\bar{x}_4 \bar{y}_4 + \bar{x}_5 \bar{y}_5) - \sum_{i=1}^3 (\bar{x}_i \bar{y}_{9-i} + \bar{x}_{9-i} \bar{y}_i).$$

Studying the multiplication table reveals that $b(\bar{x}, \bar{y})$ is equal to the sum of the coefficients of $\bar{4}$ and $\bar{5}$ in the product $\bar{x}\bar{y}$. Thus

$$b(\bar{x}, \bar{y}) = (\bar{x}\bar{y})_4 + (\bar{x}\bar{y})_5.$$

We have the following lemma:

Lemma 3.1.3 Let $\bar{x}, \bar{y} \in C^+$. Then $\bar{x}\bar{y} = 0$ implies that $b(\bar{x}, \bar{y}) = 0$.

Proof. Suppose $\bar{x}\bar{y} = 0$. Then we have $b(\bar{x}, \bar{y}) = (\bar{x}\bar{y})_4 + (\bar{x}\bar{y})_5 = 0_4 + 0_5 = 0$. □

The form $b(\cdot, \cdot)$ is **non-degenerate**, meaning that there is no element $\bar{x} \in C^+ \setminus \{0\}$ such that $b(\bar{x}, \bar{y}) = 0$ for all $\bar{y} \in C^+$. Check this by setting $\bar{x} := \sum_{i=1}^8 \lambda_i e_i$, an arbitrary

¹The indices used in the form $b(\cdot, \cdot)$ have been altered from the version found in Segev and Smith's paper to match the basis we are using from Wilson's book.

element of C^+ . Then using the multiplication table, we can verify that for $i \in \{1, 2, 3\} \cup \{6, 7, 8\}$ we have $b(\bar{x}, e_i) = -\lambda_{9-i}$, and for $i \in \{4, 5\}$ we have $b(\bar{x}, e_i) = \lambda_i$. If $b(\bar{x}, e_i) = 0$ for all $1 \leq i \leq 8$, then every λ_i is 0 and so $\bar{x} = 0$.

Definition 3.1.4 (Isotropic vectors and totally isotropic subspaces) *We say that a vector $\bar{x} \in C^+$ is **isotropic** if $b(\bar{x}, \bar{x}) = 0$, and a subspace $S \leq C^+$ is **totally isotropic** if $b(\bar{x}, \bar{y}) = 0$ for all $\bar{x}, \bar{y} \in S$.*

Definition 3.1.5 (Automorphism group of an algebra) *Let A be an algebra over k . A k -linear bijection $\theta: A \rightarrow A$ which preserves multiplication is called an **automorphism** of A . The group of all such maps under composition is the **automorphism group of A** .*

Lemma 3.1.6 *Every automorphism of C^+ fixes \bar{e} .*

Proof. Let G be the automorphism group of C^+ , and let $\theta \in G$. Then for all $\bar{x} \in C^+$ we have

$$\bar{x}\theta = (\bar{x}\bar{e})\theta = \bar{x}\theta \cdot \bar{e}\theta.$$

Since θ is a bijection, the set $\{\bar{x}\theta \mid \bar{x} \in C^+\}$ contains every element of C^+ , and so $\bar{e}\theta$ acts as a multiplicative identity. Hence $\bar{e}\theta = \bar{e}$, and so every automorphism of C^+ fixes the identity. \square

3.1.2 The Cayley module

With respect to the bilinear form $b(\cdot, \cdot)$, the orthogonal complement of $\langle \bar{e} \rangle$ is given by

$$\bar{C} := \langle \bar{1}, \bar{2}, \bar{3}, \bar{4} - \bar{5}, \bar{6}, \bar{7}, \bar{8} \rangle.$$

Since $C^+ = \overline{C} \oplus \langle \bar{e} \rangle$, and every automorphism fixes $\langle \bar{e} \rangle$, the group G also fixes \overline{C} . Thus \overline{C} is in fact a 7-dimensional kG -module; it is called the **Cayley module**. (Note that \overline{C} is not a subalgebra of C^+ since it is not closed under multiplication.) We will show in the next section that the form is preserved by the action of G .

3.2 The group $G = G_2(k)$

We define a matrix group $G = G_2(k)$ of automorphisms of C^+ ; later, we will see that this is isomorphic to the Chevalley group $G_2(k)$ as defined in Chapter 3. We will give matrix generators for G taken from [36, p.124]. Firstly, we define:

$$r_p := \begin{pmatrix} -1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & -1 & 0 & 0 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0 & 0 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & -1 \end{pmatrix}, \quad r_\ell := \begin{pmatrix} 0 & -1 & 0 & 0 & 0 & 0 & 0 & 0 \\ -1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & -1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & -1 \\ 0 & 0 & 0 & 0 & 0 & 0 & -1 & 0 \end{pmatrix}.$$

Both r_p and r_ℓ have order 2.

Let k^\times be the multiplicative group of the field k , containing all non-zero elements of k . We also define

$$T_1(\lambda) := \text{diag}(\lambda, 1, \lambda, 1, 1, \lambda^{-1}, 1, \lambda^{-1})$$

and

$$T_2(\lambda) := \text{diag}(1, \lambda, \lambda^{-1}, 1, 1, \lambda, \lambda^{-1}, 1)$$

for each $\lambda \in k^\times$. By inspecting the multiplication table of C^+ , we can see that

$$T := \langle T_i(\lambda) \mid \lambda \in k^\times, i \in \{1, 2\} \rangle$$

is precisely the set of diagonal elements which preserve the multiplication of C^+ . (For example, the equation $\bar{3} \cdot \bar{2} = \bar{1}$ implies that if we scale $\bar{2}$ by a factor of λ then we either have to scale $\bar{3}$ by a factor of λ^{-1} , as in $T_2(\lambda)$, or scale $\bar{1}$ by a factor of λ , as in $T_1(\lambda)$.) The subgroup T has order $(q-1)^2$ and is abelian, being a direct product of two cyclic groups of order $q-1$. Also let

$$N := \langle r_p, r_\ell, T \rangle,$$

which has order $12(q-1)^2$.

Again following [36, p.124], for each $\lambda \in k$ we define the following elements of G :

$$\begin{aligned}
A(\lambda) &= \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ -\lambda & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & \lambda & 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix}, & B(\lambda) &= \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ -\lambda & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & \lambda & 0 & 0 & 0 & 0 & 1 \end{pmatrix}, \\
C(\lambda) &= \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ -\lambda & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ \lambda & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & -\lambda & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & \lambda & 0 & 0 & 0 & 1 & 0 \\ \lambda^2 & 0 & 0 & -\lambda & \lambda & 0 & 0 & 1 \end{pmatrix}, & D(\lambda) &= \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ -\lambda & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & -\lambda & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & \lambda & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & \lambda^2 & 0 & -\lambda & \lambda & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & \lambda & 0 & 1 \end{pmatrix}, \\
E(\lambda) &= \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & -\lambda & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & \lambda & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix}, & F(\lambda) &= \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ -\lambda & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & \lambda & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & -\lambda & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & \lambda^2 & \lambda & -\lambda & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \lambda & 1 \end{pmatrix}.
\end{aligned}$$

Then we have

$$G := \langle T, r_{\mathbf{p}}, r_{\mathbf{\ell}}, A(\lambda), B(\lambda), \dots, F(\lambda) \mid \lambda \in k \rangle.$$

A **point** is a 1-dimensional subspace of C^+ on which multiplication restricts to 0. Similarly, a **line** is a 2-dimensional subspace of C^+ on which multiplication restricts to 0. We will fix for the entirety of this thesis a designated point $\mathbf{p} := \langle \bar{1} \rangle$ and a designated line $\mathbf{\ell} := \langle \bar{1}, \bar{2} \rangle$. We use the letters p and ℓ to refer to other points and lines throughout, but appearing in boldface \mathbf{p} and $\mathbf{\ell}$ refer to these designated subspaces.

Definition 3.2.1 Let $U := \langle A(\lambda), B(\lambda), \dots, F(\lambda) \mid \lambda \in k \rangle$ and $B := \langle U, T \rangle$.

Definition 3.2.2 Define $G_{\mathbf{p}} := \langle B, r_{\mathbf{p}} \rangle$ which is the stabiliser of the 1-space \mathbf{p} , and $G_{\mathbf{\ell}} := \langle B, r_{\mathbf{\ell}} \rangle$ which is the stabiliser of the 2-space $\mathbf{\ell}$.

Definition 3.2.3 We define the **Weyl group of G** by $W := N/T$.

The Weyl group W is isomorphic to the dihedral group D_{12} of order 12. One element of the Weyl group which we will use in particular is

$$w_0 := \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix},$$

which is given by $w_0 = (r_{\mathbf{p}} r_{\mathbf{\ell}})^3$.

We can extract the following lemmas from the text in Wilson's book [36, p.124-125, Equation (4.34) onwards]:

Lemma 3.2.4 *The subgroup U is a Sylow π -subgroup of G , of order q^6 . The subgroup B has order $q^6(q-1)^2$, and we have $B = N_G(U)$. Finally, $T \trianglelefteq N$.*

Lemma 3.2.5 *Both $G_{\mathbf{p}}$ and G_{ℓ} have order $q^6(q^2 - 1)(q - 1)$. We have $\langle G_{\mathbf{p}}, G_{\ell} \rangle = G$, and G has order $q^6(q^6 - 1)(q^2 - 1)$.*

Now we can prove the following lemma:

Lemma 3.2.6 *G preserves the bilinear form $b(\cdot, \cdot)$.*

Proof. Let $g \in G$ and $x \in C^+$. We will show that the sum of the coefficients of $\bar{4}$ and $\bar{5}$ are the same in \bar{x} and $\bar{x}g$.

It suffices to check that this is true when g is one of the generators of G . The generators r_{ℓ} , $r_{\mathbf{p}}$, $T_1(\lambda)$ and $T_2(\lambda)$ leave the coefficients of $\bar{4}$ and $\bar{5}$ unchanged, as do the generators $A(\lambda)$, $B(\lambda)$ and $E(\lambda)$. The generator $C(\lambda)$ maps $\bar{8} \mapsto \lambda^2\bar{1} - \lambda\bar{4} + \lambda\bar{5} + \bar{8}$, so the sum of the coefficients of $\bar{4}$ and $\bar{5}$ is unchanged. The generators $D(\lambda)$ and $F(\lambda)$ behave similarly, altering the coefficients of $\bar{4}$ and $\bar{5}$ when x has non-zero coefficients on $\bar{7}$ and $\bar{6}$ respectively, but not affecting the sum. Hence $\bar{x}_4 + \bar{x}_5 = (\bar{x}g)_4 + (\bar{x}g)_5$.

Now, for $\bar{y}, \bar{z} \in C^+$ we have

$$b(\bar{y}g, \bar{z}g) = (\bar{y}g\bar{z}g)_4 + (\bar{y}g\bar{z}g)_5 = (\bar{y}\bar{z}g)_4 + (\bar{y}\bar{z}g)_5 = (\bar{y}\bar{z})_4 + (\bar{y}\bar{z})_5 = b(\bar{y}, \bar{z}),$$

as required. □

The **orthogonal group** corresponding to a non-degenerate symmetric bilinear form is the group of invertible linear transformations preserving the form. Thus Lemma 3.2.6 shows that G is a subgroup of the orthogonal group corresponding to the form $b(\cdot, \cdot)$.

Lemma 3.2.7 *G acts transitively on the totally isotropic 1-spaces of \bar{C} .*

Proof. In a non-degenerate orthogonal space of dimension $2m + 1$ (such as \bar{C} , for $m = 3$) there are $q^{2m} - 1$ non-zero isotropic vectors [36, p.71]. Since a scalar multiple of an isotropic vector is also isotropic, the 7-dimensional space \bar{C} contains $(q^6 - 1)/(q - 1)$

distinct 1-spaces containing only isotropic vectors; and since the form $b(\cdot, \cdot)$ is bilinear, these subspaces are in fact totally isotropic. (There cannot be any more totally isotropic subspaces than these, as there are no more isotropic vectors.)

Recall that $G_{\mathbf{p}}$ is the stabiliser of the (totally isotropic) 1-space $\mathbf{p} = \langle \bar{1} \rangle$. Now, the index of $G_{\mathbf{p}}$ in G is given by

$$(q^6(q^6 - 1)(q^2 - 1)) / (q^6(q^2 - 1)(q - 1)) = (q^6 - 1)/(q - 1).$$

Therefore, G acts transitively on the totally isotropic 1-spaces of \bar{C} as required. \square

Lemma 3.2.8 *The totally isotropic 1-spaces of \bar{C} are precisely the 1-spaces on which multiplication restricts to 0.*

Proof. Suppose $S \leq \bar{C}$ and multiplication restricts to 0 on S . Then certainly S is totally isotropic, because for any $\bar{x}, \bar{y} \in S$ we have

$$\begin{aligned} b(\bar{x}, \bar{y}) &= (\bar{xy})_4 + (\bar{xy})_5 \\ &= 0_4 + 0_5 \\ &= 0. \end{aligned}$$

Now suppose that $S \leq \bar{C}$ is a totally isotropic 1-space. Then $S = \langle \bar{x} \rangle$ for some $\bar{x} \in \bar{C}$.

Write

$$\bar{x} = \lambda_1 \bar{1} + \lambda_2 \bar{2} + \lambda_3 \bar{3} + \lambda_{4-5}(\bar{4} - \bar{5}) + \lambda_6 \bar{6} + \lambda_7 \bar{7} + \lambda_8 \bar{8}.$$

We can modify the multiplication table of C^+ to obtain the following table, which shows how elements of \bar{C} multiply with each other. (This is not a ‘multiplication table of \bar{C} ’,

since \overline{C} is not closed under multiplication.)

	$\overline{1}$	$\overline{2}$	$\overline{3}$	$\overline{4} - \overline{5}$	$\overline{6}$	$\overline{7}$	$\overline{8}$
$\overline{1}$				$-\overline{1}$	$\overline{2}$	$-\overline{3}$	$-\overline{4}$
$\overline{2}$			$-\overline{1}$	$\overline{2}$		$-\overline{5}$	$\overline{6}$
$\overline{3}$		$\overline{1}$		$\overline{3}$	$-\overline{5}$		$-\overline{7}$
$\overline{4} - \overline{5}$	$\overline{1}$			$\overline{4} + \overline{5}$	$\overline{6}$	$\overline{7}$	
$\overline{6}$	$-\overline{2}$		$-\overline{4}$	$-\overline{6}$		$\overline{8}$	
$\overline{7}$	$\overline{3}$	$-\overline{4}$		$-\overline{7}$	$-\overline{8}$		
$\overline{8}$	$-\overline{5}$	$-\overline{6}$	$\overline{7}$	$\overline{8}$			

(3.2.9)

Now, we have

$$\begin{aligned}
\overline{x} \cdot \overline{x} &= \sum_{i=1}^8 \sum_{j=1}^8 \lambda_i \lambda_j (\overline{ij}) \\
&= \sum_{i=1}^8 \lambda_i^2 (\overline{ii}) + \sum_{i=1}^8 \sum_{j=i+1}^8 \lambda_i \lambda_j (\overline{ij} + \overline{ji}).
\end{aligned}$$

Note that, other than the product $(\overline{4} - \overline{5}) \cdot (\overline{4} - \overline{5})$, all the products on the leading diagonal of Table 3.2.9 are 0. Note also that the table is *almost* skew-symmetric; we have that $\overline{ij} = -\overline{ji}$ everywhere except for on the antidiagonal. Thus we can rewrite the product as

$$\begin{aligned}
\overline{x} \cdot \overline{x} &= \lambda_{\overline{4-5}}^2 (\overline{4} + \overline{5}) + (\lambda_1 \lambda_8 + \lambda_2 \lambda_7 + \lambda_3 \lambda_6) (-\overline{4} - \overline{5}) \\
&= (\lambda_{\overline{4-5}}^2 - (\lambda_1 \lambda_8 + \lambda_2 \lambda_7 + \lambda_3 \lambda_6)) (\overline{4} + \overline{5}).
\end{aligned}$$

This shows that the only basis elements appearing in a square $\overline{x} \cdot \overline{x}$, for $\overline{x} \in \overline{C}$, are $\overline{4}$ and $\overline{5}$. Furthermore, the coefficients of $\overline{4}$ and $\overline{5}$ are equal. Since we have assumed that S is totally isotropic, therefore, we must have $\overline{x} \cdot \overline{x} = 0$. □

3.3 The generalised hexagon Δ and the incidence graph

Γ

We can use the algebra C^+ to construct a point-line geometry called a **generalised hexagon**. This is in fact a $G_2(k)$ building, as we will see in this section.

Definition 3.3.1 (Point-line geometry) A *point-line geometry* is a structure consisting of a set of **points** P and a set of **lines** L , along with an incidence relation $I \subseteq P \times L$. We interpret I as being the set of point-line pairs (p, ℓ) for which ℓ ‘goes through’ p (or equivalently, p ‘lies on’ ℓ), and we write $p \in \ell$ as shorthand for this.

A point-line geometry can be considered as a rank 2 simplicial complex, in which the 0-simplices are the elements of $P \cup L$ and the 1-simplices are the elements of I , with the face relations given by $p, \ell \prec (p, \ell)$.

Definition 3.3.2 (Incidence graph of a point-line geometry) The *incidence graph* of a point-line geometry Δ is a bipartite graph $\Gamma = \Gamma(\Delta)$ with vertex set $P \cup L$, where a graph edge indicates incidence in the geometry. That is, the edge set of Γ is given by the incidence relation I .

Let $d(a, b)$ be the standard distance function on Γ . If a is a vertex in Γ and $i \in \mathbb{N}$, then we define $\Gamma_i(a)$ to be the set of vertices at distance i from a . Finally, we write $a \sim b$ to mean that there is a graph edge from a to b , and sometimes we will write $a \sim_n b$ to indicate that there is a shortest path of length n from a to b (so, for example, if $a, b \in P$ then $a \sim_2 b$ means that a and b are collinear).

Recall that the **girth** of a graph is the length of a shortest cycle, and the **diameter** is the length of a shortest path between two vertices of maximum distance apart.

Definition 3.3.3 (Generalised polygon) A *generalised polygon* is a point-line geometry for which the incidence graph Γ has diameter n and girth $2n$.

Definition 3.3.4 (Generalised hexagon) A *generalised hexagon* is a generalised polygon with $n = 6$; that is, it has diameter 6 and girth 12.

A point-line geometry Δ is in fact a rank 2 chamber system over index set $\{P, L\}$, where the chambers are the pairs $(p, \ell) \in I$, the P -adjacency is determined by chambers sharing a point and the L -adjacency by chambers sharing a line. That is, we have $(p, \ell) \underset{P}{\sim} (p, \ell')$ and $(p, \ell) \underset{L}{\sim} (p', \ell)$. Considered in this way, the incidence graph $\Gamma = \Gamma(\Delta)$ is just the geometric realisation of Δ , as per Algorithm 2.2.8. For more details about this correspondence, see [27, p.335, Theorem 9.4.10].

3.3.1 The construction of Δ and Γ

We use the Cayley algebra C^+ to construct a point-line geometry Δ as follows: the points are the 1-subspaces of \overline{C} for which multiplication restricts to 0, and the lines are the 2-subspaces of \overline{C} for which multiplication restricts to 0. Say that a point $p = \langle \overline{x} \rangle$ lies on a line $\ell = \langle \overline{y}, \overline{z} \rangle$ if and only if $\langle \overline{x} \rangle \subset \langle \overline{y}, \overline{z} \rangle$. Let $\Gamma = \Gamma(\Delta)$ be the incidence graph of the geometry Δ . We will show soon that Δ is indeed a generalised hexagon.

Figure 3.1 shows a part of Δ displayed as a **collinearity hypergraph**. Here, graph vertices represent points and graph hyperedges represent lines (although we do not always display all points on a given hyperedge). Note that this is *not* a portion of the graph Γ , because lines are represented by edges rather than vertices.

As the figures we are drawing get increasingly complex, drawing the incidence graph Γ would result in a cluttered diagram, so this approach becomes clearer. However, care must be taken to interpret these figures because each line has more than two incident points, so potential ambiguity arises when interpreting a straight line segment with three or more points on (is it one line, or multiple?) — later, we will introduce a convention to mitigate this possible confusion.

In this thesis, we will always use the distance function $d(\cdot, \cdot)$ to mean distance in the




Figure 3.1: A 12-circuit in Γ , represented as part of the collinearity hypergraph of Δ .

incidence graph Γ . Therefore the distance between $\langle \bar{1} \rangle$ and $\langle \bar{2} \rangle$ in Figure 3.1 is 2, via the path $\langle \bar{1} \rangle, \langle \bar{1}, \bar{2} \rangle, \langle \bar{1} \rangle$.

We will often use the following notation:

Definition 3.3.5 ($\langle x_1, x_2 \rangle$ for $x_1, x_2 \in P$) *Let $x_1, x_2 \in P$ with $d(x_1, x_2) = 2$. Then we define $\langle x_1, x_2 \rangle$ to be the unique line collinear to both x_1 and x_2 .*

We continue with some basic lemmas regarding Γ .

Lemma 3.3.6 *Let x_1 and x_2 be two distinct points in Γ , with $x_1 = \langle \bar{v}_1 \rangle$ and $x_2 = \langle \bar{v}_2 \rangle$. Then $d(x_1, x_2) = 2$ if and only if $\bar{v}_1 \cdot \bar{v}_2 = 0$.*

Proof. Suppose $d(x_1, x_2) = 2$. Then there is a path $x_1, \langle x_1, x_2 \rangle, x_2$, so we have $\langle \bar{v}_1 \rangle \subseteq \langle \bar{v}_1, \bar{v}_2 \rangle \supseteq \langle \bar{v}_2 \rangle$. Thus multiplication restricts to 0 on the subspace $\langle \bar{v}_1, \bar{v}_2 \rangle$, and so $\bar{v}_1 \cdot \bar{v}_2 = 0$.

Now suppose that $\bar{v}_1 \cdot \bar{v}_2 = 0$. Then the points x_1 and x_2 span a 2-space on which multiplication restricts to zero, so $\langle x_1, x_2 \rangle$ is a line. Hence there is a path $x_1, \langle x_1, x_2 \rangle, x_2$. \square

3.3.2 Transitivity on points and lines

We have the following lemmas.

Lemma 3.3.7 G acts transitively on the set of points of Γ .

Proof. This follows immediately from Lemma 3.2.7 and Lemma 3.2.8. \square

Lemma 3.3.8 Γ is $(q + 1)$ -regular.

Proof. The valency of a line x is $q + 1$ because every 1-space contained in the 2-space corresponding to x is a point; since multiplication restricts to 0 on the 2-space, it also does so on each 1-dimensional subspace. There are $q + 1$ such subspaces in each 2-space, and so each line is incident to $q + 1$ points.

To show that each point lies on $q + 1$ lines, we first consider the point $\mathbf{p} = \langle \bar{1} \rangle$. The lines incident to \mathbf{p} correspond to the 2-spaces in \bar{C} containing $\langle \bar{1} \rangle$ on which multiplication restricts to 0. Using the multiplication table of C^+ , we can deduce that any such 2-space must be contained within the span $\langle \bar{1}, \bar{2}, \bar{3} \rangle$; the products $\bar{1} \cdot (\bar{4} - \bar{5}) = \bar{1}$, $\bar{1} \cdot \bar{6} = -\bar{2}$, $\bar{1} \cdot \bar{7} = \bar{3}$ and $\bar{1} \cdot \bar{8} = -\bar{5}$ are all non-zero and linearly independent, so any vector forming a 2-space with $\langle \bar{1} \rangle$, on which multiplication restricts to 0, cannot have any component from the set $\{\bar{4} - \bar{5}, \bar{6}, \bar{7}, \bar{8}\}$.

Now, the span $\langle \bar{2}, \bar{3} \rangle$ contains $q^2 - 1$ non-zero vectors and thus $(q^2 - 1)/(q - 1) = q + 1$ unique 1-spaces, each spanned by some vector $\bar{v} = \alpha\bar{2} + \beta\bar{3}$. All of these 1-spaces are in fact points, because $\bar{v}^2 = \alpha^2(\bar{2} \cdot \bar{2}) + \alpha\beta(\bar{2} \cdot \bar{3} + \bar{3} \cdot \bar{2}) + \beta^2(\bar{3} \cdot \bar{3}) = \alpha\beta(\bar{1} - \bar{1}) = 0$, and all of the points form a line with \mathbf{p} since $\bar{1} \cdot \bar{2} = \bar{2} \cdot \bar{1} = 0$ and $\bar{1} \cdot \bar{3} = \bar{3} \cdot \bar{1} = 0$. Thus \mathbf{p} lies on $q + 1$ lines.

Finally, the fact that G acts transitively on points due to Lemma 3.3.7 shows that *all* points lie on $q + 1$ lines. \square

Corollary 3.3.9 $|P| = |L|$.

Proof. By Lemma 3.3.8, we have that Γ is a regular bipartite graph with vertex partition $\{P, L\}$. Hence $|P| = |L|$. \square

Lemma 3.3.10 G acts transitively on the set of lines of Γ .

Proof. The stabiliser of the line $\ell = \langle \bar{1}, \bar{2} \rangle$ is G_ℓ and the stabiliser of the point $\mathbf{p} = \langle \bar{1} \rangle$ is $G_{\mathbf{p}}$. Furthermore, we have $|G_\ell| = |G_{\mathbf{p}}|$ as per Lemma 3.2.5. We use the orbit-stabiliser theorem; we have

$$|P| = [G : G_{\mathbf{p}}],$$

since G is transitive on points. Then the size of the orbit of ℓ under G is given by

$$|\{\ell g \mid g \in G\}| = [G : G_\ell] = [G : G_{\mathbf{p}}] = |P| = |L|,$$

where the final equality is due to Lemma 3.3.9. Hence all lines lie in a single G -orbit, proving transitivity. \square

Lemmas 3.3.7 and 3.3.10, showing the transitivity of G on points and lines, allow us to define **parabolic subgroups** of G , which are G -conjugates of $G_{\mathbf{p}}$ and G_ℓ . These two families of parabolic subgroups are the stabilisers of points and lines respectively. As in Chapter 2, we write the Levi decompositions of the parabolics as $G_{\mathbf{p}} = \langle U_{\mathbf{p}}, L_{\mathbf{p}} \rangle$ and $G_\ell = \langle U_\ell, L_\ell \rangle$, where $U_x = O_\pi(G_x)$ and $L_x = G_x \cap G_{x'}$ for some x' opposite x .

For the point \mathbf{p} , we have

$$U_{\mathbf{p}} = O_\pi(G_{\mathbf{p}}) = \langle A(\lambda), \dots, D(\lambda), F(\lambda) \mid \lambda \in k \rangle$$

of order q^5 , and

$$L_{\mathbf{p}} = \langle T, r_{\mathbf{p}}, E(\lambda) \mid \lambda \in k \rangle = G_{\mathbf{p}} \cap G_{\langle \bar{8} \rangle}$$

which has order $q(q+1)(q-1)^2$ is a Levi complement of $G_{\mathbf{p}}$. Similarly, for the line ℓ , we have

$$U_\ell = O_\pi(G_\ell) = \langle A(\lambda), \dots, E(\lambda) \mid \lambda \in k \rangle$$

of order q^5 , and

$$L_\ell = \langle T, r_\ell, F(\lambda) \mid \lambda \in k \rangle = G_\ell \cap G_{\langle \bar{7}, \bar{8} \rangle}$$

which has order $q(q+1)(q-1)^2$ is a Levi complement of G_ℓ . Furthermore we have $L_{\mathbf{p}} \cong L_\ell \cong \text{GL}_2(k)$; a proof of this appears in Section 3.3.8.

We set up notation for intersections of parabolic subgroups as follows:

Definition 3.3.11 (Notation for stabilisers in G) *Let G_x denote the stabiliser of x in G , where x is either a point or a line. Furthermore, let $G_{x_1, x_2, \dots}$ denote the intersection $\bigcap_i G_{x_i}$.*

3.3.3 Arc transitivity

Let $s \geq 0$. An s -**arc** is a sequence of $s+1$ vertices x_0, \dots, x_s such that $x_i \sim x_{i+1}$ for each $i = 0, \dots, s-1$, and $x_{i-1} \neq x_{i+1}$ for all $i = 1, \dots, s-1$ (that is, the sequence never immediately returns to the previous vertex). We say that a group action on a graph is s -**arc transitive** if there exists a group element mapping any s -arc to any other.

Clearly the action of G on Γ cannot be s -arc transitive for any value of s , since G is not transitive on vertices (preserving the partition $P \cup L$). But it may satisfy a modified, weaker condition: a group G acts **locally s -arc transitively** if, for any given vertex y and any two s -arcs γ, γ' starting at y , there exists $g \in G$ mapping γ to γ' . (See Figure 3.2.) Note that since Γ is bipartite, such an element maps points to points and lines to lines, so it is conceivable that G could act with this property.

Lemma 3.3.12 *The action of G on Γ is locally 7-arc transitive.*

Proof. Since G is transitive on points and lines, it is enough to show that G is transitive on the 7-arcs beginning at the point $\mathbf{p} = \langle \bar{1} \rangle$, and on the 7-arcs beginning at the line $\ell = \langle \bar{1}, \bar{2} \rangle$.




Figure 3.2: A part of the graph Γ , with an element $g \in G$ mapping a 2-arc starting at a vertex y to another 2-arc starting at y .

We begin with the 7-arcs starting at $\langle \bar{1} \rangle$, where we have $G_{\langle \bar{1} \rangle} = G_{\mathbf{p}}$. Studying the multiplication table of C^+ , we see that any point collinear with $\langle \bar{1} \rangle$ is a linear combination of $\langle \bar{1} \rangle$, $\langle \bar{2} \rangle$ and $\langle \bar{3} \rangle$, and thus the lines incident to $\langle \bar{1} \rangle$ are given by $\langle \bar{1}, \sigma \bar{2} + \mu \bar{3} \rangle$ for $\sigma, \mu \in k$ not both equal to 0. A glance at the matrices $r_{\mathbf{p}}$ and $E(\lambda)$ shows us that $G_{\langle \bar{1} \rangle}$ acts transitively on these lines; the action of $r_{\mathbf{p}}$ allows us to swap $\bar{2}$ for $\bar{3}$ if necessary to ensure that the coefficient of $\bar{3}$ is non-zero, and then multiplication by $E(\lambda)$ allows us to fix any ratio of $\bar{2}$ and $\bar{3}$, all whilst fixing the point $\langle \bar{1} \rangle$. Without loss of generality, therefore, we may assume that our 7-arc emanating from $\langle \bar{1} \rangle$ begins with the path $\langle \bar{1} \rangle, \langle \bar{1}, \bar{2} \rangle$.

The stabiliser $G_{\langle \bar{1} \rangle, \langle \bar{1}, \bar{2} \rangle}$ is the intersection $G_{\mathbf{p}} \cap G_{\ell} = B$. An s -arc is not allowed to double back on itself, so the point following the line $\langle \bar{1}, \bar{2} \rangle$ is not permitted to be the point $\langle \bar{1} \rangle$. Therefore we must show that $G_{\langle \bar{1} \rangle, \langle \bar{1}, \bar{2} \rangle}$ is transitive on the remaining q points of $\langle \bar{1}, \bar{2} \rangle$. These are precisely $\langle \bar{2} + \mu \bar{1} \rangle$ for $\mu \in k$. We see that

$$\langle \bar{2} + \mu \bar{1} \rangle \cdot F(\lambda) = \langle \bar{2} + (\mu - \lambda) \bar{1} \rangle,$$

so by careful choice of λ we can map any such point to any other. Without loss of generality, therefore, we may assume that the next point is $\langle \bar{2} \rangle$.

Now, it is easily checked that

$$G_{\langle \bar{1} \rangle, \langle \bar{1}, \bar{2} \rangle, \langle \bar{2} \rangle} = \langle T, A(\lambda), B(\lambda), \dots, E(\lambda) \mid \lambda \in k \rangle,$$

and has order $q^5(q-1)^2$. We continue in this manner; the lines incident to $\langle \bar{2} \rangle$ which are distinct from $\langle \bar{1}, \bar{2} \rangle$ are given by $\langle \bar{2}, \sigma \bar{1} + \bar{6} \rangle$ where $\sigma \in k$. Then $G_{\langle \bar{1} \rangle, \langle \bar{1}, \bar{2} \rangle, \langle \bar{2} \rangle}$ acts transitively on the 1-spaces of the form $\langle \sigma \bar{1} + \bar{6} \rangle$ since we can multiply by $B(\lambda)$, and so acts transitively on these lines.

Without loss, we choose the line $\langle \bar{2}, \bar{6} \rangle$ to continue our arc, and so our new stabiliser is given by

$$G_{\langle \bar{1} \rangle, \langle \bar{1}, \bar{2} \rangle, \langle \bar{2} \rangle, \langle \bar{2}, \bar{6} \rangle} = \langle T, A(\lambda), C(\lambda), D(\lambda), E(\lambda) \mid \lambda \in k \rangle,$$

and has order $q^4(q-1)^2$. As before, the elements $C(\lambda)$ give us transitivity on the remaining points of $\langle \bar{2}, \bar{6} \rangle$. Again, we may make any choice here, so we continue the arc with $\langle \bar{6} \rangle$. Then

$$G_{\langle \bar{1} \rangle, \langle \bar{1}, \bar{2} \rangle, \langle \bar{2} \rangle, \langle \bar{2}, \bar{6} \rangle, \langle \bar{6} \rangle} = \langle T, A(\lambda), D(\lambda), E(\lambda) \mid \lambda \in k \rangle,$$

of order $q^3(q-1)^2$.

In the same fashion, we continue with the line $\langle \bar{6}, \bar{8} \rangle$ (losing the generators $A(\lambda)$ from the stabiliser), then with the point $\langle \bar{8} \rangle$ (losing the generators $D(\lambda)$), and finally with the line $\langle \bar{7}, \bar{8} \rangle$ (losing the generators $E(\lambda)$). We are left with the stabiliser

$$G_{\langle \bar{1} \rangle, \langle \bar{1}, \bar{2} \rangle, \langle \bar{2} \rangle, \langle \bar{2}, \bar{6} \rangle, \langle \bar{6} \rangle, \langle \bar{6}, \bar{8} \rangle, \langle \bar{8} \rangle, \langle \bar{8}, \bar{7} \rangle} = T,$$

so we can go no further at this point. This completes the proof of local 7-arc transitivity of 7-arcs emanating from $\langle \bar{1} \rangle$ (and therefore from all points by transitivity on points).

The other case to be considered, 7-arcs originating from lines rather than from points,

can be treated almost identically. We can construct the stabiliser

$$G_{\langle \bar{2}, \bar{1} \rangle, \langle \bar{1} \rangle, \langle \bar{1}, \bar{3} \rangle, \langle \bar{3} \rangle, \langle \bar{3}, \bar{7} \rangle, \langle \bar{7} \rangle, \langle \bar{7}, \bar{8} \rangle, \langle \bar{8} \rangle} = T,$$

step-by-step as before; in doing so we lose the generators $A(\lambda), \dots, F(\lambda)$ in the reverse order. \square

Corollary 3.3.13 *G acts transitively on ordered pairs of opposite points, and on ordered pairs of opposite lines.*

Proof. Suppose that $x_0, y_0, x_1, y_1 \in P$ are points with x_0 opposite y_0 and x_1 opposite y_1 . We need to show that there exists some $g \in G$ such that $x_0g = x_1$ and $y_0g = y_1$. By Lemma 3.3.7, G is transitive on the points of Γ so there exists some $h \in G$ such that $x_0h = x_1$.

The points y_0h and y_1 are both opposite x_1 , and so the length of a geodesic from x_1 to either of them is 6. Therefore since G acts locally 7-arc transitively on Γ by Lemma 3.3.12, there exists an $f \in G_{x_1}$ such that $(y_0h)f = y_1$. Hence we may take $g = hf$, and then $(x_0, y_0)g = (x_1, y_1)$ as required.

For pairs of opposite lines the proof is identical, using Lemma 3.3.10 for transitivity on lines. \square

3.3.4 More properties of Γ

We can now deduce some more information about the graph Γ .

Lemma 3.3.14 *The diameter of Γ is 6 and the girth of Γ is 12.*

Proof. We begin by showing that $d(\langle \bar{1} \rangle, \langle \bar{8} \rangle) = 6$. Figure 3.1 shows one path of length 6 from $\langle \bar{1} \rangle$ to $\langle \bar{8} \rangle$, so certainly $d(\langle \bar{1} \rangle, \langle \bar{8} \rangle) \leq 6$. Assume for a contradiction that there is a shorter path. Then either $\langle \bar{1} \rangle$ and $\langle \bar{8} \rangle$ are collinear, which is not true by Lemma 3.3.6

since $\bar{1} \cdot \bar{8} \neq 0$, or $d(\langle \bar{1} \rangle, \langle \bar{8} \rangle) = 4$ and there is a point collinear to both. But the points collinear with $\langle \bar{1} \rangle$ lie in the space $\langle \bar{1}, \bar{2}, \bar{3} \rangle$, and the points collinear with $\langle \bar{8} \rangle$ lie in the space $\langle \bar{6}, \bar{7}, \bar{8} \rangle$, and these have trivial intersection. Therefore $d(\langle \bar{1} \rangle, \langle \bar{8} \rangle) = 6$ and thus the diameter of Γ is at least 6.

Suppose the diameter of Γ is greater than 6. Then by transitivity on points (Lemma 3.3.7), there must be some line at distance 7 from the point $\langle \bar{1} \rangle$.

Consider the 7-arc given by graph vertices $\langle \bar{1} \rangle, \langle \bar{1}, \bar{2} \rangle, \langle \bar{2} \rangle, \langle \bar{2}, \bar{6} \rangle, \langle \bar{6} \rangle, \langle \bar{6}, \bar{8} \rangle, \langle \bar{8} \rangle, \langle \bar{7}, \bar{8} \rangle$. The final vertex of the arc, which is the line $\langle \bar{7}, \bar{8} \rangle$, has distance at most 5 from $\langle \bar{1} \rangle$, via the path $\langle \bar{1} \rangle, \langle \bar{1}, \bar{3} \rangle, \langle \bar{3} \rangle, \langle \bar{3}, \bar{7} \rangle, \langle \bar{7} \rangle, \langle \bar{7}, \bar{8} \rangle$ (see Figure 3.1). By local 7-arc transitivity (Lemma 3.3.12), this statement holds for the final line of *any* 7-arc beginning at $\langle \bar{1} \rangle$. This contradicts the existence of a line at distance 7 from the point $\langle \bar{1} \rangle$, and so the diameter of Γ is 6.

There is a 12-cycle in Figure 3.1, so the girth of Γ is at most 12. Suppose there is a cycle γ of length less than 12 in Γ . Since the graph is bipartite, γ must alternate between points and lines and thus includes at least one point. By the transitivity of G on points, we may assume that γ contains the point $\langle \bar{1} \rangle$ (if it does not, apply a graph automorphism $g \in G$ mapping one of the points on γ to the point $\langle \bar{1} \rangle$).

Now, starting with $\langle \bar{1} \rangle$ and moving round in a fixed direction, the first 8 graph vertices in γ form a 7-arc γ' . Since G is locally 7-arc transitive there is an automorphism mapping γ' to the 'standard' 7-arc from $\langle \bar{1} \rangle$ which goes through points $\langle \bar{2} \rangle, \langle \bar{6} \rangle$ and $\langle \bar{8} \rangle$ and ends at $\langle \bar{7}, \bar{8} \rangle$, traversing a little over half-way around the 12-cycle from Figure 3.1.

This shows firstly that all the points and lines of γ' are distinct and so the length of γ is at least 8. But we may use the same argument as before to note that the distance from $\langle \bar{7}, \bar{8} \rangle$ back to $\langle \bar{1} \rangle$ is at least 5, because no point on $\langle \bar{7}, \bar{8} \rangle$ is collinear with $\langle \bar{1} \rangle$. The same must hold for the final line of γ' , and so γ cannot be any smaller than a 12-cycle. Thus the girth of Γ is 12. □

Corollary 3.3.15 Δ is a generalised hexagon.

Proof. This follows immediately from Definition 3.3.4 and Lemma 3.3.14. \square

3.3.5 Opposite vertices

We say that x and y in $P \cup L$ are **opposite** if $d(x, y) = 6$. Since Γ is bipartite and has diameter 6, if we have two distinct points $x, y \in \Gamma$ then $d(x, y) \in \{2, 4, 6\}$. The same is true for any pair of distinct lines.

Lemma 3.3.16 Suppose x and y are points (resp. lines) and $d(x, y) = 4$. Then there is a unique path of length 4 from x to y .

Proof. Suppose not. Then there are two paths γ and γ' of length 4 from x to y . The concatenation $\gamma \circ \gamma'$ is therefore a closed walk of length 8, and since $\gamma \neq \gamma'$ it contains a cycle of length at most 8; a contradiction since the girth of Γ is 12. \square

Therefore we can introduce the following notation:

Definition 3.3.17 ($x * y$) Let x and y be points (resp. lines) such $d(x, y) = 4$. We define $x * y$ to be the unique point (resp. line) in $\Gamma_2(x) \cap \Gamma_2(y)$.

Lemma 3.3.18 Let x and x^* be opposite vertices in Γ and y^* be a vertex incident to x^* . Then there is a unique neighbour r of y^* with distance $d(r, x) = 4$.

Proof. Since x and x^* are opposite, there is a path γ of length 6 from x^* to x . Denote the two vertices following x^* in the path γ by y^* and r , so that $\gamma = (x^*, y^*, r, \dots, x)$ and $d(r, x) = 4$. Let γ_0 be the subpath of γ from r to x , of length 4.

Suppose that s is a neighbour of y^* with $s \neq r$ and $d(s, x) = 4$, and write ψ for the path (r, y^*, s) of length 2. Then there is a path ρ from x to s of length 4, and therefore the concatenation $\rho \circ \psi \circ \gamma_0$ is a closed walk in Γ containing a circuit of length at most 10; a contradiction since the girth of Γ is 12. So the choice of $r \in y^*$ must be unique.

Lemma 3.3.12 then implies that we can find such an r incident to any of the neighbours y^* of x^* , which concludes the proof. \square

This motivates the following definition:

Definition 3.3.19 *Let x, x^* and y^* be as in Lemma 3.3.18. We define $r(x, y^*)$ to be the unique neighbour r of y^* with distance $d(r, x) = 4$.*

For example, if p and p^* are a pair of opposite points with $\ell^* \ni p^*$ then $r(p, \ell^*)$ is the unique point on ℓ^* at distance 4 from p .

Corollary 3.3.20 *Let x, x^* and y^* be as in Lemma 3.3.18. Then there is a unique path γ of length 5 from y^* to x .*

Proof. By Lemma 3.3.18, $r = r(x, y^*)$ is the unique neighbour of y^* at distance 4 from x , so γ must start with (y^*, r, \dots) . Then by Lemma 3.3.16, there is a unique path of length 4 from r to x , so γ is unique. \square

3.3.6 The generalised hexagon Δ is the building of G

In Corollary 3.3.15 we showed that the geometry we have constructed is indeed a generalised hexagon according to Definition 3.3.4. It remains to show that it is in fact the building of the Chevalley group $G_2(k)$.

Firstly we will show that Γ satisfies the axioms to be a building as per Definition 2.3.1.

Theorem 3.3.21 *The graph Γ , considered as a simplicial complex, is a building whose apartments are precisely the 12-cycles of Γ .*

Before we prove this theorem, we require the following lemmas.

Lemma 3.3.22 *Suppose $e = (a_1, a_2)$ and $e' = (b_1, b_2)$ are two distinct edges of Γ , with $a_i, b_i \in P \cup L$. Then either e and e' are the first and last edges of a geodesic of Γ , or they are antipodal edges in a 12-circuit in Γ .*

Proof. Set $d = \min\{d(a_i, b_j) \mid 1 \leq i, j \leq 2\}$. For ease of notation, we re-number the vertices such that $d(a_1, b_1) = d$, if this is not already the case. Since Γ is a bipartite graph, we must have $d(a_2, b_1) = d(a_1, b_2) = d + 1$. There are two options for $d(a_2, b_2)$; either it is d or $d + 2$.

- **Case 1:** $d(a_2, b_2) = d$. Then there are geodesic paths γ_1 from a_1 to b_1 and γ_2 from a_2 to b_2 , both of length d . This means that $\omega = (a_2, a_1) \circ \gamma_1 \circ (b_1, b_2) \circ \gamma_2^{-1}$ is a closed walk with at least two edges (namely e and e') which are only used once, and thus contains a cycle of positive length. Since the girth of Γ is 12, this cycle must have length at least 12. Furthermore, the length of ω is given by $2d + 2$. Hence $2d + 2 \geq 12$, so $d \geq 5$.

Conversely, we have $d(a_1, b_2) = d + 1 \leq 6$, since 6 is the diameter of the Γ , so $d \leq 5$. This forces $d = 5$ and $|\omega| = 12$, so ω is a circuit of length 12 and the edges e and e' are antipodal edges of ω .

- **Case 2:** $d(a_2, b_2) = d + 2$. Let γ be a geodesic path from a_1 to b_1 . Then $(a_2, a_1) \circ \gamma \circ (b_1, b_2)$ is a path of length $d + 2$ from a_2 to b_2 , and thus is a geodesic path with the edges e and e' at either end. □

Lemma 3.3.23 *Any geodesic in Γ of length $d < 6$ can be extended to a geodesic of length $d + 1$.*

Proof. Let γ be a geodesic of length $d < 6$ from u to v , and suppose that γ cannot be extended to a geodesic of length $d + 1$. Let x_1 be the penultimate vertex of γ . Since v has $q + 1$ neighbours, we may choose another vertex x_2 adjacent to v with $x_1 \neq x_2$, and $d(u, x_1) = d(u, x_2) = d - 1$ (or else γ could have been extended to visit x_2 after v).

Let ε be a geodesic from u to x_2 . Then the closed walk $\omega = \gamma \circ (v, x_2) \circ \varepsilon^{-1}$ includes at least two distinct edges (those incident to v), and therefore contains a circuit of positive

length. The girth of Γ is 12, so $|\omega| = 2d \geq 12$ and thus $d \geq 6$, a contradiction. This proves the lemma. \square

Lemma 3.3.24 *Any two distinct edges of Γ lie in a common 12-circuit.*

Proof. Let e and e' be two distinct edges of Γ as before. By Lemma 3.3.22, either e and e' are antipodal edges in a 12-circuit (in which case we are done), or they are the first and last edges of some geodesic path γ . We suppose that we are in the second case.

By repeated application of Lemma 3.3.23, we may extend γ to a geodesic γ' of length 6. Let the initial vertex of γ' be denoted a_2 , followed by vertex a_1 (so that $e = (a_1, a_2)$). (The unusual choice of notation will soon make sense.) Denote the final vertex of γ' by b_1 .

Since b_1 has $q + 1$ neighbours, we may choose a vertex b_2 adjacent to b_1 such that $b_2 \notin \gamma'$. Let $f = (b_1, b_2)$. Now, we have $d(a_1, b_1) = 5$ and $d(a_2, b_1) = d(a_1, b_2) = 6$. As in the proof of Lemma 3.3.22, we have $d(a_2, b_2) \in \{5, 7\}$; but it cannot be 7 since the diameter of Γ is 6. The remainder of the proof therefore follows Case 1 in the proof of Lemma 3.3.22, and we find a 12-circuit of Γ containing γ' (and therefore both e and e'). \square

Now we can prove the theorem.

Proof (Proof of Theorem 3.3.21). We have to show that the axioms from Definition 2.3.1 are satisfied. By Example 2.2.10, the Coxeter complex of type G_2 is a 12-circuit, so axiom (B1) is satisfied.

Now, let x and y be two simplices (that is, vertices or edges) of Γ . We need to show that there is some apartment (12-circuit) of Γ containing both x and y . If x and y are edges, then this follows immediately from Lemma 3.3.24. If either x or y (or both) is a vertex, then replace it with one of its $q + 1$ incident edges, and then apply Lemma 3.3.24. A 12-cycle containing an edge e certainly contains any vertex contained in e , so axiom (B2) is satisfied.

We turn our attention to axiom $(B3)$. Let $\mathcal{A}, \mathcal{A}'$ be two apartments both containing simplices x and y (which again can be vertices or edges). The relative positions of x and y must be the same in the apartments \mathcal{A} and \mathcal{A}' ; if a shorter path from x to y existed in either one of the apartments, then this would result in Γ possessing a cycle shorter than 12, a contradiction since the girth of Γ is 12. Hence we can form an isomorphism $\mathcal{A} \rightarrow \mathcal{A}'$ by ‘pairing up’ the vertices and edges of each 12-cycle in order, starting from x . This isomorphism fixes x and y pointwise.

Hence all of the axioms are satisfied, and so Γ is indeed a building. \square

In particular, the building Γ we have formed is the building for the group of Lie type $G_2(k)$; the following is all by construction.

Theorem 3.3.25 *The graph Γ is isomorphic (as a simplicial complex) to the building for the group of Lie type $G_2(k)$. The chambers (1-simplices) in the building correspond to edges in Γ , and are stabilised by conjugates of B . The panels (0-simplices) correspond to vertices in Γ , and are stabilised by conjugates of the parabolic subgroups G_p and G_ℓ . The apartments of the building are precisely the 12-cycles in Γ , and they are stabilised by conjugates of N .*

Recall that we defined a *group of Lie type* as the subgroup of the automorphism group of a building generated by the *root groups* X_α (see Definition 2.4.1). The group G is a subgroup of the automorphism group of the generalised hexagon, which is the building for the group of Lie type $G_2(k)$; in order to demonstrate that these groups are isomorphic (and thus that the two definitions of ‘ $G_2(k)$ ’ are equivalent), it suffices to show that G is also generated by the root groups. We refer to Wilson [36, Section 4.3.5] for this, noting also that the group G defined in this chapter is simple for $q \neq 2$ [36, Section 4.3.7], as is the Chevalley group $G_2(k)$ when $|k| \neq 2$ [30, p.33, Theorem 5].

3.3.7 Subspaces of \overline{C}

For a subspace $\overline{X} \leq \overline{C}$, the orthogonal complement with respect to $b(\cdot, \cdot)$ is defined in the usual way:

$$\overline{X}^\perp := \{\overline{v} \in \overline{C} \mid b(\overline{x}, \overline{v}) = 0, \forall \overline{x} \in \overline{X}\}.$$

We have the following lemma about pairs of opposite points:

Lemma 3.3.26 *Let $x_1 = \langle \overline{v}_1 \rangle, x_2 = \langle \overline{v}_2 \rangle$ be distinct points of Γ . Then x_1 is opposite x_2 if and only if $b(\overline{v}_1, \overline{v}_2) \neq 0$.*

Proof. By the transitivity of G on points and the local 7-arc transitivity of G , it suffices to consider $x_1 = \langle \overline{1} \rangle, x_2 = \langle \overline{2} \rangle, x_3 = \langle \overline{6} \rangle$ and $x_4 = \langle \overline{8} \rangle$, which form a path $x_1 \sim_2 x_2 \sim_2 x_3 \sim_2 x_4$ (see Figure 3.1). By direct calculation we have $b(\overline{1}, \overline{2}) = 0 = b(\overline{1}, \overline{6})$, and $b(\overline{1}, \overline{8}) = -1 \neq 0$. By Lemma 3.2.6, the action of G preserves the form $b(\cdot, \cdot)$, so the result follows immediately. \square

In view of Lemma 3.3.26, therefore, we can define the orthogonal complement of a point:

Definition 3.3.27 (x^\perp) *If $x \in \Gamma$ is a point then let x^\perp be the set of all points which are not opposite x . For a set of points $S = \{x_1, x_2, \dots\}$, define*

$$S^\perp := \bigcap_{i \geq 1} x_i^\perp.$$

(It is a slight abuse of notation that we use the same notation for the orthogonal complement in \overline{C} as we do for the set of points not opposite a point x . It should always be clear from the context which we mean.)

Lemma 3.3.28 *Suppose S is a set of points and $T \subseteq S$. Then $T^\perp \supseteq S^\perp$.*

Proof. This follows immediately from the definition. \square

Example 3.3.29 We calculate \mathbf{p}^\perp and ℓ^\perp . Note from Table 3.2.9 that the only basis element x of \overline{C} such that $b(\mathbf{p}, x) \neq 0$ is $x = \langle \overline{8} \rangle$. Hence $\mathbf{p}^\perp = \langle \overline{1}, \overline{2}, \overline{3}, \overline{4} - \overline{5}, \overline{6}, \overline{7} \rangle$.

The line ℓ is spanned by the points \mathbf{p} and $\langle \overline{2} \rangle$, so $\ell^\perp = \mathbf{p}^\perp \cap \langle \overline{2} \rangle^\perp = \langle \overline{1}, \overline{2}, \overline{3}, \overline{4} - \overline{5}, \overline{6}, \overline{7} \rangle \cap \langle \overline{1}, \overline{2}, \overline{3}, \overline{4} - \overline{5}, \overline{6}, \overline{8} \rangle = \langle \overline{1}, \overline{2}, \overline{3}, \overline{4} - \overline{5}, \overline{6} \rangle$.

We define some useful subspaces of \overline{C} , which arise from subgraphs of Γ .

Definition 3.3.30 (\overline{C}_p , \overline{C}_ℓ , \overline{D}_p , \overline{E}_ℓ and \overline{F}_p) Let $p \in P$ be a point, and $\ell \in L$ be a line.

- Define \overline{C}_p to be the 1-space of \overline{C} corresponding to p .
- Define \overline{C}_ℓ to be the 2-space of \overline{C} corresponding to ℓ .
- Set $\overline{D}_p := \langle \overline{C}_\ell \mid \ell \ni p \rangle$. (In Γ , this corresponds to the subgraph $\Gamma_2(p)$.)
- Set $\overline{E}_\ell := \langle \overline{D}_p \mid p \in \ell \rangle$. (In Γ , this corresponds to the subgraph $\Gamma_3(\ell)$.)
- Set $\overline{F}_p := \langle \overline{E}_\ell \mid \ell \ni p \rangle$. (In Γ , this corresponds to the subgraph $\Gamma_4(p)$.)

Due to the transitivity of the group action on points and lines, the dimensions of these spaces do not depend on the choice of p and ℓ . By definition, $\dim \overline{C}_p = 1$ and $\dim \overline{C}_\ell = 2$.

The rest we can calculate.

Lemma 3.3.31 The dimension of \overline{D}_p is 3.

Proof. We have already shown that all the lines incident to $\langle \overline{1} \rangle$ span the subspace $\langle \overline{1}, \overline{2}, \overline{3} \rangle$, so by transitivity on points the result follows. \square

Lemma 3.3.32 The dimension of \overline{E}_ℓ is 5.

Proof. We calculate in \overline{C} . By the transitivity of G on lines, it suffices to work with $\ell = \langle \overline{1}, \overline{2} \rangle$. By Example 3.3.29 we have $\ell^\perp = \langle \overline{1}, \overline{2}, \overline{3}, \overline{4} - \overline{5}, \overline{6} \rangle$ which has dimension 5. Take $\mathbf{p} = \langle \overline{1} \rangle \in \ell$. Then $\overline{D}_p = \langle \overline{1}, \overline{2}, \overline{3} \rangle$ which is a subspace of ℓ^\perp , and because G_ℓ acts

transitively on the points in ℓ by Lemma 3.3.10, and leaves ℓ^\perp invariant, we have $\overline{D_p} \leq \ell^\perp$ for any $p \in \ell$. Hence $\overline{E_\ell} \leq \ell^\perp$.

Since $\overline{E_\ell}$ contains $\overline{D_{\langle \bar{1} \rangle}} = \langle \bar{1}, \bar{2}, \bar{3} \rangle$, $\overline{D_{\langle \bar{2} \rangle}} = \langle \bar{1}, \bar{2}, \bar{6} \rangle$ and $\overline{D_{\langle \bar{1} + \bar{2} \rangle}} = \langle \bar{1}, \bar{2}, \bar{3} + \bar{6} - (\bar{4} - \bar{5}) \rangle$, we see that $\dim \overline{E_\ell} \geq 5$. Hence $\overline{E_\ell} = \ell^\perp$ and $\dim \overline{E_\ell} = 5$. \square

Lemma 3.3.33 *The dimension of $\overline{F_p}$ is 6.*

Proof. Again by transitivity on points it suffices to show that $\dim \overline{F_p} = 6$, where $p = \langle \bar{1} \rangle$. By definition, $\overline{F_p}$ contains precisely all those points which are not opposite to p . By Lemma 3.3.26, this consists of all points $\langle \bar{x} \rangle$ for which $b(\bar{1}, \bar{x}) = 0$. The subspace $p^\perp = \langle \bar{x} \in \overline{C} \mid b(\bar{1}, \bar{x}) = 0 \rangle$ is a hyperplane of \overline{C} and thus has dimension $7 - 1 = 6$. \square

The following two observations will motivate the bounds we present in Chapters 9 and 10 respectively.

Lemma 3.3.34 *Suppose that p^\dagger is opposite p . Then $\langle \overline{F_p}, \overline{C_{p^\dagger}} \rangle = \overline{C}$.*

Proof. By Lemma 3.3.26 and Lemma 3.3.33 we have that $\overline{F_p}$ is the hyperplane of \overline{C} spanned by all points not opposite p . Furthermore, $\overline{C_{p^\dagger}}$ is a 1-space of \overline{C} which is not contained in $\overline{F_p}$. Hence $\langle \overline{F_p}, \overline{C_{p^\dagger}} \rangle = \overline{C}$. \square

Lemma 3.3.35 *Suppose that p^\dagger is opposite p . Then $\langle \overline{E_\ell}, \overline{E_{\ell^\dagger}} \rangle = \overline{C}$.*

Proof. By transitivity on ordered pairs of opposite lines (Corollary 3.3.13), it suffices to take $\ell = \langle \bar{1}, \bar{2} \rangle$ and $\ell^\dagger = \langle \bar{7}, \bar{8} \rangle$. From Figure 3.1 we see that $\langle \bar{1}, \bar{2}, \bar{3}, \bar{6} \rangle \leq \overline{E_\ell}$ and $\langle \bar{3}, \bar{6}, \bar{7}, \bar{8} \rangle \leq \overline{E_{\ell^\dagger}}$. Therefore $\langle \overline{E_\ell}, \overline{E_{\ell^\dagger}} \rangle$ contains the 6-space $\langle \bar{1}, \bar{2}, \bar{3}, \bar{6}, \bar{7}, \bar{8} \rangle$; it remains to show that it also contains the vector $\bar{4} - \bar{5}$.

Let $v = \bar{3} - \bar{4} + \bar{5} + \bar{6}$ and consider the point $x = \langle v \rangle$. We can see that x lies in both $\overline{E_\ell}$ and $\overline{E_{\ell^\dagger}}$, since it is collinear with both of the points $\langle \bar{1} + \bar{2} \rangle \in \ell$ and $\langle \bar{7} - \bar{8} \rangle \in \ell^\dagger$. Hence $\bar{4} - \bar{5} \in \langle \overline{E_\ell}, \overline{E_{\ell^\dagger}} \rangle = \overline{C}$. \square

3.3.8 3-transitivity results

Let $x \in P \cup L$. The parabolic G_x acts 3-transitively on the $q + 1$ lines or points incident to x , as we will demonstrate in the following lemmas.

Lemma 3.3.36 *Let $\ell \in L$. The subgroup G_ℓ acts as $\text{GL}_2(k)$ on the 2-space $\overline{C_\ell}$. In particular, G_ℓ operates 3-transitively on the $q + 1$ points in ℓ .*

Proof. Without loss of generality, we take $\ell = \boldsymbol{\ell}$. Let $L_\ell = \langle T, r_\ell, F(\lambda) \mid \lambda \in k \rangle$, which is a Levi complement of G_ℓ , being the stabiliser of $\boldsymbol{\ell} = \langle \overline{1}, \overline{2} \rangle$ and an opposite line $\ell^\dagger = \langle \overline{7}, \overline{8} \rangle$. Then $L_\ell \leq G_\ell$ and $G_\ell = \langle L_\ell, U_\ell \rangle$, where $U_\ell = O_\pi(G_\ell) = \langle A(\lambda), \dots, E(\lambda) \mid \lambda \in k \rangle$.

We have that U_ℓ fixes every vector in $\overline{C_\ell}$. The generators of $L_\ell|_{\overline{C_\ell}}$ are found by extracting a 2×2 block corresponding to row and column indices $\{1, 2\}$ from each of the generators of L_ℓ . Thus, L_ℓ acts as $\text{GL}_2(k)$ on $\overline{C_\ell}$ with $F(\lambda)|_{\overline{C_\ell}} = \begin{pmatrix} 1 & 0 \\ -\lambda & 1 \end{pmatrix}$ for $\lambda \in k$, $T|_{\overline{C_\ell}} = \langle \text{diag}(\gamma, \mu) \mid \gamma, \mu \in k^\times \rangle$ and $r_\ell|_{\overline{C_\ell}} = \begin{pmatrix} 0 & -1 \\ -1 & 0 \end{pmatrix}$. Hence, by [4, Theorem 2.6.2], G_ℓ acts 3-transitively on the points in ℓ . \square

Lemma 3.3.37 *Let $p \in P$. The subgroup G_p acts as $\text{GL}_2(k)$ on the 2-space $\overline{D_p}/\overline{C_p}$. In particular, G_p acts 3-transitively on the $q + 1$ lines containing p .*

Proof. Again we may just consider the case $p = \boldsymbol{p}$ by transitivity on points. For a line $h \ni \boldsymbol{p}$ we have that $\overline{C_h}$ is contained in $\overline{D_p}$, where $\overline{D_p} = \langle \overline{1}, \overline{2}, \overline{3} \rangle$. Moreover, $\overline{C_h}/\overline{C_p}$ is a 1-subspace of the 2-space $\overline{D_p}/\overline{C_p}$.

Let $L_p = \langle T, r_p, E(\lambda) \mid \lambda \in k \rangle$, which is a Levi complement of G_p being the stabiliser of $\boldsymbol{p} = \langle \overline{1} \rangle$ and $p^\dagger = \langle \overline{8} \rangle$, a point opposite \boldsymbol{p} . The generators of $L_p|_{\overline{D_p}/\overline{C_p}}$ are found by extracting a 2×2 block corresponding to row and column indices $\{2, 3\}$ from each of the generators of L_p . Hence $L_p|_{\overline{D_p}/\overline{C_p}} \cong \text{GL}_2(k)$ is generated by $E(\lambda)|_{\overline{D_p}/\overline{C_p}} = \begin{pmatrix} 1 & 0 \\ -\lambda & 1 \end{pmatrix}$ for $\lambda \in k$, $T|_{\overline{D_p}/\overline{C_p}} = \langle \text{diag}(\gamma, \gamma^{-1}\mu) \mid \gamma, \mu \in k^* \rangle = \langle \text{diag}(\gamma, \mu) \mid \gamma, \mu \in k^* \rangle$, and $r_p|_{\overline{D_p}/\overline{C_p}} = \begin{pmatrix} 0 & -1 \\ -1 & 0 \end{pmatrix}$.

Furthermore, $G_{\mathbf{p}} = \langle L_{\mathbf{p}}, U_{\mathbf{p}} \rangle$ where $U_{\mathbf{p}} = O_{\pi}(G_{\mathbf{p}}) = \langle A(\lambda), \dots, D(\lambda), F(\lambda) \mid \lambda \in k \rangle$ which fixes every vector in $\overline{D_{\mathbf{p}}}/\overline{C_{\mathbf{p}}}$. It follows that $G_{\mathbf{p}}$ acts on $\overline{D_{\mathbf{p}}}/\overline{C_{\mathbf{p}}}$ as $\mathrm{GL}_2(k)$. Again using [4, Theorem 2.6.2], $G_{\mathbf{p}}$ acts 3-transitively on the $q + 1$ lines containing \mathbf{p} . \square

Corollary 3.3.38 *Let $G_x = \langle U_x, L_x \rangle$ where $x \in P \cup L$. Then the Levi complement $L_x \cong \mathrm{GL}_2(k)$.*

Proof. We showed in Lemma 3.3.36 and Lemma 3.3.37 that both L_{ℓ} and L_p act as $\mathrm{GL}_2(k)$ on a 2-space. Furthermore, the subgroups U_{ℓ} and U_p fix these subspaces, and $|L_{\ell}| = |L_p| = |\mathrm{GL}_2(k)|$. \square

3.3.9 Ideal lines

In this section we define an *ideal line* and prove some results which will become important later on.

Definition 3.3.39 (Ideal line $I_x(y)$) *Let x, y be a pair of opposite points. We define*

$$I_x(y) := \Gamma_2(x) \cap \Gamma_4(y).$$

By Lemma 3.3.18 we have $|I_x(y)| = q + 1$, since every line incident to x contains exactly one point at distance 4 from y . Like an ordinary line, the points of an ideal line span a 2-space in \overline{C} as we will demonstrate shortly. However, multiplication on an ideal line does not restrict to 0, preventing it from being an ordinary line of Γ .

Lemma 3.3.40 *Suppose that $x = \langle \overline{x} \rangle$ and $y = \langle \overline{y} \rangle$ are opposite points in Γ . Then $\langle I_x(y) \rangle = \overline{D_x} \cap \overline{y}^{\perp}$ and has dimension 2.*

Proof. We have $I_x(y) \subseteq \Gamma_2(x)$ and therefore $\langle I_x(y) \rangle \leq \overline{D_x}$. Also we have $\langle I_x(y) \rangle \leq \Gamma_4(y) \leq \overline{y}^{\perp}$. Since $\dim \overline{D_x} = 3$ by Lemma 3.3.31, and $b(\overline{x}, \overline{y}) \neq 0$, we have $\dim(D_x \cap \overline{y}^{\perp}) \leq$

2. On the other hand, $|I_x(y)| = q + 1$ and so $\dim\langle I_x(y) \rangle \geq 2$. Together these bounds prove the lemma. \square

Theorem 3.3.41 *Suppose that $x, y, z \in P$ with x opposite both y and z . If $|I_x(y) \cap I_x(z)| \geq 2$ then $I_x(y) = I_x(z)$.*

Proof. We have $\dim\langle I_x(y) \cap I_x(z) \rangle \geq 2$ because it contains at least two points. Since

$$\langle I_x(y) \rangle \geq \langle I_x(y) \cap I_x(z) \rangle \leq \langle I_x(z) \rangle,$$

with $\dim\langle I_x(y) \rangle = \dim\langle I_x(z) \rangle = 2$ and $\dim\langle I_x(y) \cap I_x(z) \rangle \geq 2$, we have that

$$\langle I_x(y) \rangle = \langle I_x(y) \cap I_x(z) \rangle = \langle I_x(z) \rangle.$$

But both $I_x(y)$ and $I_x(z)$ contain exactly $q + 1$ points, and so $I_x(y) = I_x(z)$. \square

CHAPTER 4

SOME WEIGHT THEORY OF GROUPS OF LIE TYPE

We begin with some basic weight theory, consisting of a series of definitions and lemmas which hold in the general case, for G any group of Lie type. Later on in the chapter we will specialise to the particular case where $G = G_2(k)$, with k a finite field of order π^a where π is a prime.

The representation theory of Chevalley groups over algebraically closed fields of characteristic π is governed by the abstract theory of weights, as described in Bourbaki ([6, Chapter VI]). The representation theory of Lie algebras is also built upon this abstract weight theory, and hence many of the results and proofs also appear in Humphreys' book on this subject [15] in a readily accessible form.

4.1 Preliminary material

Let G be a group of Lie type with a BN -pair, and let $T = B \cap N$. The subgroup T is abelian and has rank ℓ , and thus we can represent elements of T with vectors (t_1, \dots, t_ℓ) where each $t_i \in k^\times$. Let $X(T)$ be the set of all homomorphisms from T to k^\times ; we call these homomorphisms **characters**. A character λ , being a homomorphism from T to

k^\times , must map $(t_1, \dots, t_\ell) \mapsto \prod_{i=1}^\ell t_i^{a_i}$ for some choice of exponents $a_i \in \mathbb{Z}$. Hence we can express λ as a vector $\lambda = (a_1, \dots, a_\ell)$ in a free \mathbb{Z} -module of rank ℓ .

For any kG -module V and any $\lambda \in X(T)$, we define the subspace

$$V_\lambda := \{v \in V \mid v \cdot t = (t\lambda)v, \text{ for all } t \in T\}.$$

If λ and V are such that $V_\lambda \neq \{0\}$, then we call λ a **weight** of V , and V_λ a **weight space** [10, p.214]. The **multiplicity** of a weight λ is the dimension of the weight space V_λ .

We form an ℓ -dimensional Euclidean space $E := \mathbb{R} \otimes_{\mathbb{Z}} X(T)$. Every group of Lie type G has an associated **Lie algebra**, which can be considered as a kG -module, and its weights are called the **roots** of G [10, p.214]. These roots form a **root system** [14, (A.2), p.229], which is a set of vectors Φ , and a corresponding set of reflections

$$\sigma_\alpha: E \rightarrow E, \text{ given by } (\beta)\sigma_\alpha := \beta - \frac{2(\beta, \alpha)}{(\alpha, \alpha)}\alpha,$$

such that

(R1) Φ is finite, $\langle \Phi \rangle = E$, and $0 \notin \Phi$,

(R2) $\alpha \in \Phi$ implies $-\alpha \in \Phi$, and no other multiples of α are contained in Φ ,

(R3) $\Phi\sigma_\alpha = \Phi$ for all $\alpha \in \Phi$, and

(R4) $\langle \alpha, \beta \rangle \in \mathbb{Z}$ for all $\alpha, \beta \in \Phi$.

We can choose a **base** Δ , which is a subset of Φ such that

(B1) Δ is a basis of E , and

(B2) each root $\beta \in \Phi$ can be expressed as an integral linear combination of the elements of Δ , such that either all coefficients are non-negative, or all are non-positive.

Every root system has a base [14, (A.4), p.229]. We fix a base Δ and refer to the ℓ roots contained in Δ as the **simple roots**. With Δ fixed, we call the reflections $\{\sigma_\alpha \mid \alpha \in \Delta\}$ the **simple reflections**. The simple roots are the basis for a lattice $\Lambda_r := \mathbb{Z}\Delta$ called the **root lattice**.

(Recall the *roots of an apartment* \mathcal{A} defined in Section 2.3.2. These are in one-to-one correspondence with the roots of the root system Φ associated with the building. For a root (half-apartment) $\gamma \in \mathcal{A}$, the reflection σ_γ fixes the wall M_γ and maps γ to $\text{op}_{\mathcal{A}}(\gamma)$. Similarly, the root $\gamma \in E$ is mapped by σ_γ to the root $-\gamma \in E$.)

4.1.1 The Weyl group

The reflections σ_α , for $\alpha \in \Phi$, generate a group called the **Weyl group**:

$$W := \langle \sigma_\alpha \mid \alpha \in \Phi \rangle.$$

(Recall the *Weyl group of* $G = \text{G}_2(k)$ from Chapter 3, which was given by N/T where $N = \langle r_p, r_\ell, T \rangle$. Here, the generating reflections σ_1 and σ_2 correspond to the elements $r_p T$ and $r_\ell T$ in some permutation – we will determine which way round these are in Chapter 8. In general, a Weyl group can always be formed as a quotient of its corresponding group of Lie type in this way.)

In fact, for any choice of base Δ , the group W is generated by just the elements $\{\sigma_\alpha \mid \alpha \in \Delta\}$ [14, (A.5), p.229]. Each σ_α is a reflection and thus fixes a hyperplane $P_\alpha \subset E$ which is perpendicular to the root α . The collection of hyperplanes $\{P_\alpha \mid \alpha \in \Phi\}$ partitions E into **Weyl chambers**; they are the connected components of $E \setminus \{P_\alpha \mid \alpha \in \Phi\}$. For any $\gamma \in E$ which does not lie on any hyperplane P_α , we define $\mathfrak{C}(\gamma)$ to be the (unique) Weyl chamber containing γ .

Two points φ_1 and φ_2 are in the same Weyl chamber if and only if they lie on the same side of each hyperplane P_α , or in other words if and only if (φ_1, α) and (φ_2, α) have

the same sign (+/−) for each $\alpha \in \Phi$. With Δ fixed, there is one Weyl chamber, called the **fundamental Weyl chamber**, for which all the inner products are positive:

$$\mathfrak{C}(\Delta) := \{\gamma \in E \mid (\gamma, \alpha) > 0, \text{ for all } \alpha \in \Delta\}.$$

The Weyl group acts transitively on the set of Weyl chambers [14, (A.4), p.229]. Furthermore, the closure of the fundamental chamber $\overline{\mathfrak{C}(\Delta)}$, which is defined as

$$\overline{\mathfrak{C}(\Delta)} := \{\gamma \in E \mid (\gamma, \alpha) \geq 0, \text{ for all } \alpha \in \Delta\},$$

is a *fundamental domain* for the action of the Weyl group W on E ; that is, it contains exactly one point from each W -orbit [6, p.166, Theorem 2 (ii)].

Note that no simple reflection fixes any $\gamma \in \mathfrak{C}(\Delta)$, since $(\gamma, \alpha) > 0$ for all $\alpha \in \Delta$. Thus the stabiliser of a point $\gamma \in \mathfrak{C}(\Delta)$ is trivial, and its W -orbit has length $|W|$. However, points in $\overline{\mathfrak{C}(\Delta)} \setminus \mathfrak{C}(\Delta)$ do have non-trivial stabilisers, and thus have shorter orbits. Clearly, the point $0 \in E$ is fixed by all σ_α and so is in a W -orbit of length 1. In between these examples are points lying on some, but not all, of the hyperplanes $\{P_\alpha \mid \alpha \in \Delta\}$. The stabiliser of such a point is a parabolic subgroup of W , *i.e.* a subgroup generated by some subset of the generators $\{\sigma_\alpha \mid \alpha \in \Delta\}$.

4.1.2 Theory of weights

Let

$$\Lambda := \{\lambda \in E \mid \frac{2(\lambda, \alpha)}{(\alpha, \alpha)} \in \mathbb{Z} \text{ for all } \alpha \in \Phi\}.$$

(Note that the quantity $\frac{2(\lambda, \alpha)}{(\alpha, \alpha)}$ is often denoted $\langle \lambda, \alpha \rangle$. However, we will not adopt this convention to avoid introducing an excessive amount of notation, as it will only be used a limited number of times in this section.)

We call Λ the **weight lattice**, and the elements of Λ are **abstract weights**. The

connection to the weights of a kG -module V is as follows: it turns out that, for any choice of V , if $V_\lambda \neq \{0\}$ then $\frac{2(\lambda, \alpha)}{(\alpha, \alpha)}$ is an integer for all roots α [10, p.214]. Hence all weights are in fact elements of Λ . For brevity, we will use the term *weights* to mean elements of Λ ; when we use the term instead to describe the weights of a particular kG -module V , we will make this explicit.

The root lattice Λ_r is a sublattice of the weight lattice Λ . A weight λ is **dominant** if (λ, α_i) is non-negative for all $1 \leq i \leq \ell$, and the set of all dominant weights is denoted Λ^+ . Notice that Λ^+ is the set of weights which lie in $\overline{\mathfrak{C}(\Delta)}$.

The simple roots $\{\alpha_1, \dots, \alpha_\ell\}$ form a basis of E , but we can form a dual basis $\{\lambda_1, \dots, \lambda_\ell\}$ such that

$$\frac{2(\lambda_i, \alpha_j)}{(\alpha_j, \alpha_j)} = \delta_{ij}.$$

From the above relation we can see that the λ_i are all dominant weights, since $(\alpha_j, \alpha_j) > 0$ and $\delta_{ij} \geq 0$. We call λ_i the **fundamental dominant weights**. Any weight $\lambda \in \Lambda$ can be written as $\lambda = \sum_i m_i \lambda_i$ for some choice of coefficients m_i , and λ is dominant if and only if all $m_i \geq 0$.

The fundamental dominant weights generate the weight lattice Λ , on which the Weyl group W acts. To demonstrate this, note that

$$\begin{aligned} \lambda_i \sigma_{\alpha_j} &= \lambda_i - \frac{2(\lambda_i, \alpha_j)}{(\alpha_j, \alpha_j)} \alpha_j \\ &= \lambda_i - \delta_{ij} \alpha_j, \end{aligned}$$

and so elements of the Weyl group map weights to weights.

For a group of Lie type defined over a field k of order $q = \pi^a$, where π is a prime, we say that a dominant weight λ is **π -restricted** if $\lambda = \sum_i m_i \lambda_i$ with $0 \leq m_i \leq \pi - 1$ for all $1 \leq i \leq \ell$. Similarly, λ is **q -restricted** if $0 \leq m_i \leq q - 1$ for all $1 \leq i \leq \ell$.

Definition 4.1.1 ($\alpha \prec \beta$) We define a partial order \prec on the set of weights. Let $\alpha, \beta \in \Lambda$. We say that $\alpha \prec \beta$ if and only if $\beta - \alpha$ is a sum of simple roots with non-negative integer coefficients (including the case that $\alpha = \beta$).

The set of dominant weights forms a set of orbit representatives of the action of W on Λ [15, p.68, Lemma 13.2A] (see also [6, p.181]):

Lemma 4.1.2 Each W -orbit of weights contains precisely one dominant weight. Furthermore, if λ is the dominant weight of a W -orbit then $\lambda w \prec \lambda$ for all $w \in W$.

Definition 4.1.3 (Saturated set of weights) Let $\Theta \subset \Lambda$. We say that Θ is a **saturated set of weights** if, for every weight $\lambda \in \Theta$, we have

$$\left\{ \lambda - i\alpha \mid \alpha \in \Phi \text{ and } 0 \leq i \leq \frac{2(\lambda, \alpha)}{(\alpha, \alpha)} \right\} \subseteq \Theta.$$

If $\lambda \in \Theta$ and $\mu \prec \lambda$ for all $\mu \in \Theta$ then we call λ the **highest weight** of Θ .

A saturated set of weights is necessarily stable under W . Each reflection σ_α is equivalent to subtracting $\frac{2(\lambda, \alpha)}{(\alpha, \alpha)}\alpha$ from a weight λ , the resulting weight being contained in the saturated set by definition; these reflections generate W . Humphreys gives two lemmas about saturated sets of weights [15, p.70, Lemmas 13.3A and 13.3B] (see also [6, p.239, Ex. 23]:

Lemma 4.1.4 If Θ is a saturated set of weights with a highest weight, then Θ is finite.

Lemma 4.1.5 Suppose Θ is a saturated set of weights with highest weight λ . Then any dominant weight $\mu \in \Lambda^+$ such that $\mu \prec \lambda$ is in Θ .

With these lemmas in hand, we can completely describe a saturated set of weights possessing a highest weight.

Theorem 4.1.6 *Suppose Θ is a saturated set of weights with highest weight λ . Then Θ consists precisely of all the weights $\mu \in \Lambda^+$ with $\mu \prec \lambda$, and the conjugates of these weights under W .*

Furthermore, the cardinality of Θ is equal to the sum of the orbit lengths of all $\mu \in \Lambda^+$ such that $\mu \prec \lambda$.

Proof. By Lemma 4.1.5, all such $\mu \in \Lambda^+$ with $\mu \prec \lambda$ lie in Θ , and since Θ is closed under the action of W , all of their conjugates do as well.

Now let $\mu \in \Theta$. By Lemma 4.1.2, there is some $w \in W$ such that μw is dominant, and since Θ is closed under W , we have $\mu w \in \Theta$. But since λ is the highest weight of Θ , we have $\mu w \prec \lambda$. □

4.1.3 The Weyl module $V(\lambda)$ and its unique simple quotient $L(\lambda)$

For any kG -module M , the subgroup $B \leq G$ stabilises some 1-dimensional subspace of M [10, p.215]. Let v be a non-zero vector in this subspace. The subgroup T , being a subgroup of B , also fixes the 1-space. Thus for any $t = (t_1, \dots, t_\ell) \in T$, we have $v \cdot t = \lambda(t)v$ for some scalar $\lambda(t) \in k$. Since M is a kG -module, and therefore a kT -module, the map $t \mapsto \lambda(t)$ is a homomorphism $T \rightarrow k^\times$, and is therefore a character. Hence v lies in the weight space V_λ for some weight $\lambda = (a_1, \dots, a_\ell)$. We call v a **maximal vector** of weight λ .

Suppose there exists a maximal vector v of weight λ such that $M = \langle v \rangle_G$. Then we say that M is a **highest weight module**, and it has **highest weight** λ [10, p.215]. In this case, we find that for all other weights μ of M , we have $\mu \prec \lambda$ (hence the terminology ‘highest weight’) [14, p.189, Proposition 31.2].

Consider the set of isomorphism classes of kG -modules with highest weight λ . There is a unique irreducible module denoted $L(\lambda)$ which is a quotient of all the others, and a unique module $V(\lambda)$ (the **Weyl module**) of which all the others are quotients [10, p.215].

If k has characteristic 0 then the Weyl module itself is irreducible, and so $W(\lambda) = L(\lambda)$. However, in positive characteristic $V(\lambda)$ is often not irreducible, but does always contain a unique simple quotient $L(\lambda)$ [16, p.22]. The following statement can be found in Humphreys [16, p.17, Theorem], but is originally due to Steinberg.

Theorem 4.1.7 (Isomorphism classes of kG -modules) *Each of the q -restricted dominant weights λ (of which there are q^ℓ) gives rise to a simple kG -module $L(\lambda)$. These modules are pairwise non-isomorphic, and every simple kG -module is isomorphic to $L(\lambda)$ for some dominant, q -restricted weight λ .*

Weights of the Weyl module and its irreducible quotient

Define $\Pi(M)$ to be the set of all weights of M (with their multiplicities ignored). The weights of $V(\lambda)$ are of the form

$$\mu = \lambda - \sum_{i=1}^{\ell} a_i \alpha_i \quad (a_i \in \mathbb{Z}^+).$$

Therefore $\mu \prec \lambda$ for all $\mu \in \Pi(V(\lambda))$. Conversely, the following proposition of Humphreys tells us that, for a Weyl module $V(\lambda)$, the set $\Pi(V(\lambda))$ is as large as it can be [16, p.22, Section 3.2] (see also [15, p.114, Proposition 21.3]):

Proposition 4.1.8 *Let $\lambda \in \Lambda^+$ be a dominant weight, and $V(\lambda)$ be the corresponding Weyl module. Then $\Pi(V(\lambda))$ is saturated.*

This allows us to find a lower bound for the dimension of $V(\lambda)$, since every weight space must have dimension at least 1 and so $\dim V(\lambda) \geq |\Pi(V(\lambda))|$. Furthermore, if we do not know the highest weight λ of a Weyl module $V(\lambda)$, but we do know another weight $\mu \prec \lambda$, we can use μ to determine a lower bound for $\dim V(\lambda)$ by the following observation:

Lemma 4.1.9 *Let $\lambda \in \Lambda^+$ and assume that $\mu \in \Lambda^+$ with $\mu \prec \lambda$. Then*

$$\Pi(V(\mu)) \subseteq \Pi(V(\lambda)).$$

Proof. Every W -orbit of weights contains precisely one dominant weight by Lemma 4.1.2. It suffices, therefore, to show that

$$\{\gamma \in \Lambda^+ \mid \gamma \prec \mu\} \subseteq \{\gamma \in \Lambda^+ \mid \gamma \prec \lambda\}.$$

But this is trivial since \prec is a partial order, so $\gamma \prec \mu$ implies $\gamma \prec \lambda$. □

All results thus far have concerned the weights appearing in $V(\lambda)$, but we are often interested in its unique simple quotient $L(\lambda)$ instead. We will make use of a result from Premet [19] which says that (in nearly all cases) every weight of $V(\lambda)$ is also a weight of $L(\lambda)$. There are conditions on the characteristic of the field k , which are dependent on the group of Lie type in question. For $G = G_2(k)$, the theorem can be stated as follows:

Theorem 4.1.10 (Premet) *Let $G = G_2(k)$, with k of order $q = \pi^a$ and $\pi > 3$ (π a prime). Then for all π -restricted dominant weights, $\Pi(V(\lambda)) = \Pi(L(\lambda))$.*

(See also [17, p.139, Theorem 4.1].) This result tells us nothing about the multiplicities of the weights, but since every weight space has dimension at least 1, it does allow us to form a lower bound on $\dim L(\lambda)$.

4.2 An example for the group $SL_2(k)$

Let $k = \mathbb{F}_\pi$ for π a prime and set $G = SL_2(k)$. In this situation, we have one simple root α , and so E is 1-dimensional. We calculate the fundamental dominant weight λ using the formula $\frac{2(\lambda, \alpha)}{(\alpha, \alpha)} = 1$, giving $\lambda = \alpha/2$. The weight lattice Λ is spanned by λ , so the π -restricted weights are $m\lambda$ where $0 \leq m < \pi$. Thus by Theorem 4.1.7 the number of

isomorphism classes of irreducible kG -modules is exactly π , one for each $0 \leq m < \pi$. We will calculate the weights of each of the the Weyl modules for G .

Notice that $W = \langle \sigma_\alpha \rangle$, where σ_α is the map given by $\lambda \mapsto -\lambda$. Therefore the W -orbits of weights have length 2, except for the orbit $\{0\}$ of length 1. We can see that $0 \prec m\lambda$ if and only if $m \in \mathbb{Z}^+$ is even (and therefore $m - 0 = 2n \cdot \lambda = n \cdot \alpha$ for some $n \in \mathbb{Z}$). Hence $V(m\lambda)$ contains an odd number of weights when m is even, and an even number of weights when m is odd.

We have that $V(0)$ is 1-dimensional and is the trivial module, $V(\lambda)$ is 2-dimensional and is the standard $\mathrm{SL}_2(k)$ -module, $V(2\lambda) = V(\alpha)$ has dimension 3, and, in general, $V(m\lambda)$ is irreducible of dimension $m + 1$ for $0 \leq m < \pi$.

In these cases, for π -restricted weights, we have $V(m\lambda) = L(m\lambda)$. (We can use the polynomial ring construction to make such a module of each dimension — see [2, p.14-16] for details.)

The following technical lemma will become important later on.

Lemma 4.2.1 *Suppose $\pi > 3$ and $k = \mathbb{F}_\pi$. Let $H = \mathrm{SL}_2(k)$ and $B = N_H(S)$ where $S \in \mathrm{Syl}_\pi(H)$. Let W be a 1-dimensional kB -module and $V = \mathrm{Ind}_B^H(W)$. Assume that V has a quotient N isomorphic to the irreducible kH -module $V(\lambda)$ of dimension 2. Then V is a non-split extension of M by N , where $M \cong V((\pi - 2)\lambda)$ is the unique irreducible kH -module of dimension $\pi - 1$.*

Proof. We begin by showing that V is an indecomposable kH -module.

The index $[H : B] = (\pi(\pi + 1)(\pi - 1))/\pi(\pi - 1) = \pi + 1$ so a transversal $\{g_0 = 1, g_1, \dots, g_\pi\}$ of the right cosets $\{Bg \mid g \in H\}$ has $\pi + 1$ elements. By definition, the induced module $\mathrm{Ind}_B^H(W)$ acts on the vector space

$$V = \bigoplus_{i=0}^{\pi} Wg_i, \tag{4.2.2}$$

which we consider as an internal direct sum. Since $\dim W = 1$, therefore, we have $\dim V = \pi + 1$. The crucial point is that H acts transitively on the summands. Hence no proper kH -submodule $M \subset V$ can contain W , since $\langle W \rangle_{kH} = V \supset M$.

Consider the natural action of H on N . We may choose S to be the Sylow π -subgroup $S = \langle \begin{pmatrix} 1 & \alpha \\ 0 & 1 \end{pmatrix} \mid \alpha \in k \rangle$, so that $B = \langle \begin{pmatrix} \gamma & \alpha \\ 0 & \gamma^{-1} \end{pmatrix} \mid \alpha \in k, \gamma \in k^\times \rangle$. Then B (and hence $S \leq B$) fixes the 1-space $A := \langle (0, x) \mid x \in k \rangle$. Since only the scalar matrices of $\text{GL}_2(k)$ leave invariant every 1-space and the matrices in S are not scalar (other than the identity), using $|S| = \pi$ we obtain that S has orbits of length 1 and π on the 1-spaces in N .

Consider the restriction $V|_S$. Since $S \leq B$, the first summand $W = Wg_0$ is a kS -submodule of $V|_S$. By the above argument, S acts transitively on the remaining π summands $\{Wg_i \mid 1 \leq i \leq \pi\}$. Therefore $V|_S \cong W \oplus R$ where W is the trivial 1-dimensional kS -module and $R = kS$ is the regular representation of S .

By hypothesis, there is a kH -submodule $U \subset V$ such that $V/U \cong N$. We know that $W \not\subseteq U$ since W is not contained in any proper submodule of V . Hence $(W + U)/U$ is a 1-dimensional subspace of V/U . As W is a kB -module, B leaves $(W + U)/U$ invariant and so $(W + U)/U$ is in fact a 1-dimensional kB -submodule of V/U . By the Second Isomorphism Theorem,

$$(W + U)/U \cong_{kB} W/(W \cap U) = W.$$

Now, in the 2-dimensional module $N \cong V/U$, the subgroup B acts as $\langle \begin{pmatrix} \gamma & \alpha \\ 0 & \gamma^{-1} \end{pmatrix} \mid \alpha \in k, \gamma \in k^\times \rangle$. There are no non-zero subspaces of N centralised by B , so B does not centralise $(W + U)/U$. Hence B does not centralise W .

Now suppose that U_1 is a kH -submodule of V such that V/U_1 is 1-dimensional. Again we have $W \not\subseteq U_1$, and so $V \cong_{kB} U_1 \oplus W$. Thus $W \cong_{kB} V/U_1$. But V/U_1 is a kH -module of dimension 1, so is centralised by H and thus by B . This implies that W is centralised

by B , a contradiction, so there is no such submodule U_1 .

Next, suppose that there exists U_2 , a kH -submodule of V such that $V/U_2 \cong N \oplus X$ for some kH -submodule X . Using the correspondence theorem, let $\tilde{X} \supseteq U_2$ and $\tilde{N} \supseteq U_2$ be kH -submodules of V such that $\tilde{X}/U_2 \cong X$ and $\tilde{N}/U_2 \cong N$. Then

$$\dim V/\tilde{X} = \dim N = 2, \quad (4.2.3)$$

and $\dim V/\tilde{N} = \dim X$. Furthermore, $\dim X \geq 2$, as otherwise V/\tilde{N} is 1-dimensional which contradicts the non-existence of a submodule U_1 above.

Furthermore, $W \not\subseteq \tilde{X}$ since $\tilde{X} \neq V$ and $W \not\subseteq \tilde{N}$ since $\tilde{N} \neq V$. Now, $(W + \tilde{X})/W$ and $(W + \tilde{N})/W$ are both kS -modules, and in fact are kS -submodules of $V/W \cong R$.

As S is cyclic, the regular kS -module R is uniserial by [2, p. 26]. This means that for any two submodules $A_1, A_2 \subseteq R$ we have either $A_1 \subseteq A_2$ or vice-versa. In addition, R has precisely one submodule of each dimension $0 \leq d \leq \dim R$. Now,

$$\dim(W + \tilde{X})/W = \dim \tilde{X} \geq \dim \tilde{N} = \dim(W + \tilde{N})/W,$$

so $(W + \tilde{X})/W \supseteq (W + \tilde{N})/W$ since R is uniserial. Therefore,

$$W + \tilde{X} = W + \tilde{X} + \tilde{N} = V.$$

So $\dim V = 1 + \dim \tilde{X}$, implying that $\dim V/\tilde{X} = 1$; a contradiction since we have already shown that $\dim V/\tilde{X} = 2$ in (4.2.3). Therefore no such U_2 exists. In particular, V has a unique maximal kH -submodule M .

If V is decomposable, then $V \cong A \oplus B$ for some proper non-zero kH -submodules A and B . Then, as M is the unique maximal submodule, $A \subseteq M$ and $B \subseteq M$. Hence $V = A + B \subseteq M \subset V$, a contradiction. Hence V is indecomposable and $V/M \cong_{kH} N$.

Therefore, by [2, p.33, Lemma 5.5], we have $V \cong_{kH} P_2/Y$ where P_2 is the principal indecomposable kH -module with $P_2/\text{Rad}(P_2) \cong N$ and Y is a kH -submodule of P_2 . The structure of P_2 is described in [2, p.78]. It has dimension 2π , and $\text{Rad}(P_2)/\text{Soc}(P_2) \cong C \oplus D$, where C is irreducible of dimension $\pi - 1$ and D is irreducible of dimension $\pi - 3$.

As $\dim P_2 = 2\pi$ and $\dim V = \pi + 1$, we have $\dim Y = \pi - 1$. Also, Y contains $\text{Soc}(P_2)$ which has dimension 2 by [2, p.43, Theorem 6.6]. Thus, as $\dim Y = \pi - 1 = (\pi - 3) + 2$, we see that Y is an extension of $\text{Soc}(P_2)$ by D . It follows that $M \cong C$ of dimension $\pi - 1$ as claimed. □

4.3 An example for the group $G_2(k)$

Let $k = \pi^a$ for π a prime and set $G = G_2(k)$. Here, the fundamental weights are given by $\lambda_1 = (1, 0)$ and $\lambda_2 = (0, 1)$, and the simple roots are $\alpha_1 = (2, -1)$ and $\alpha_2 = (-3, 2)$ [15, p.69, Table 1]. Rearranging, we obtain $\lambda_1 = 2\alpha_1 + \alpha_2$ and $\lambda_2 = 3\alpha_1 + 2\alpha_2$. For brevity, we write σ_1 and σ_2 instead of σ_{α_1} and σ_{α_2} .

The following example, in which we calculate the number of distinct weights appearing in the module $V(t\lambda_2)$, will prove very useful later on.

Example 4.3.1 *Let G be as above, and let $V(t\lambda_2)$ be the Weyl module of highest weight $t\lambda_2$ for G . Then*

$$\dim V(t\lambda_2) \geq |\Pi(V(t\lambda_2))| = 1 + 6 \left\lfloor \frac{5t}{2} \right\rfloor + 12 \sum_{i=1}^{t-1} \left\lfloor \frac{3i}{2} \right\rfloor.$$

Proof. Since every Weyl group orbit on the weights in $\Pi(V(t\lambda_2))$ has a unique dominant weight representative by Lemma 4.1.2, to construct $\Pi(V(t\lambda_2))$ we should calculate the dominant weights $\theta \in \Lambda^+$ with $\theta \prec t\lambda_2$.

Write $\theta = m\lambda_1 + n\lambda_2$. We have $\theta \prec t\lambda_2$ if and only if

$$\begin{aligned} t\lambda_2 - \theta &= t\lambda_2 - (m\lambda_1 + n\lambda_2) \\ &= (t - n)\lambda_2 - m\lambda_1 \end{aligned}$$

is a non-negative integer combination of α_1 and α_2 .

Substituting the expressions for λ_1 and λ_2 we get

$$\begin{aligned} (t - n)\lambda_2 - m\lambda_1 &= (t - n)(3\alpha_1 + 2\alpha_2) - m(2\alpha_1 + \alpha_2) \\ &= (3(t - n) - 2m)\alpha_1 + (2(t - n) - m)\alpha_2. \end{aligned}$$

So we need to determine the pairs (m, n) of non-negative integers such that

$$3(t - n) - 2m \geq 0 \quad \text{and} \quad 2(t - n) - m \geq 0.$$

Rearranging these in terms of m , we see that the the second condition is in fact redundant:

$$m \leq \frac{3}{2}(t - n) \quad \text{and} \quad m \leq 2(t - n).$$

Figure 4.1 shows the pairs satisfying this condition for $t \in \{2, 3, 4\}$. The weight 0, corresponding to pair $(m, n) = (0, 0)$ and denoted by a star on the figure, is fixed by the Weyl group W . Therefore it is in a W -orbit of length 1 and only contributes 1 to the cardinality of $\Pi(V(t\lambda_2))$.

Weights of the form $m\lambda_1$ or $n\lambda_2$ are marked with a white dot in the figure, and are fixed by reflections σ_1 and σ_2 respectively. The stabilisers of these points are parabolic subgroups of W , and thus both have order 2 being given by $\langle \sigma_1 \rangle$ and $\langle \sigma_2 \rangle$. These weights are therefore in W -orbits of length 6, and each contribute 6 to the cardinality of $\Pi(V(t\lambda_2))$.




Figure 4.1: Points labelled with a star lie in a W -orbit of length 1, those labelled with a white dot lie in an orbit of length 6, and those with a black dot lie in an orbit of length 12.

All other weights, marked by black dots in the figure, are not fixed by either of the simple reflections. They are therefore in W -orbits of length 12, each contributing 12 to the cardinality of $\Pi(V(t\lambda_2))$.

The white dots are located at $\{(0, n) : 1 \leq n \leq t\}$ and $\{(m, 0) : 1 \leq m \leq 3t/2\}$, for a total of $\lfloor \frac{5t}{2} \rfloor$ dots.

The black dots consist of $t - 1$ columns, which have heights given (from right-to-left) by the sequence $\lfloor \frac{3i}{2} \rfloor = (1, 3, 4, 6, 7, \dots)$. The total number of black dots is therefore

$$\sum_{i=1}^{t-1} \left\lfloor \frac{3i}{2} \right\rfloor.$$

This means that the total number of weights is given by

$$B_t^{(1)} := 1 + 6 \left\lfloor \frac{5t}{2} \right\rfloor + 12 \sum_{i=1}^{t-1} \left\lfloor \frac{3i}{2} \right\rfloor.$$

Every weight space has dimension at least 1, and so $\dim V(t\lambda_2) \geq B_t^{(1)}$ as required. \square

Making use of Theorem 4.1.10, we can find a bound for the dimension of $L((\pi - 2)\lambda_2)$.

Example 4.3.2 *Let $G = G_2(k)$, with k of order $q = \pi^a$ and $\pi > 3$ (π a prime), and let $L((\pi - 2)\lambda_2)$ be the unique simple quotient of the Weyl module for G with highest weight $(\pi - 2)\lambda_2$. Then*

$$\dim L((\pi - 2)\lambda_2) \geq |\Pi(V((\pi - 2)\lambda_2))| = 1 + 6 \left\lfloor \frac{5(\pi - 2)}{2} \right\rfloor + 12 \sum_{i=1}^{\pi-3} \left\lfloor \frac{3i}{2} \right\rfloor.$$

Proof. The weight $(\pi - 2)\lambda_2$ is π -restricted, so we can apply Theorem 4.1.10. Thus

$$\Pi(L((\pi - 2)\lambda_2)) = \Pi(V((\pi - 2)\lambda_2)).$$

Using the bound from Example 4.3.1, and setting $t = \pi - 2$, we obtain

$$|\Pi(V((\pi - 2)\lambda_2))| = 1 + 6 \left\lfloor \frac{5(\pi - 2)}{2} \right\rfloor + 12 \sum_{i=1}^{\pi-3} \left\lfloor \frac{3i}{2} \right\rfloor = B_{\pi-2}^{(1)}.$$

Every weight space has dimension at least 1, and so $\dim L((\pi - 2)\lambda_2) \geq |\Pi(L((\pi - 2)\lambda_2))|$ and the result follows. \square

CHAPTER 5

SHEAVES ON BUILDINGS

In Chapter 2 we defined a *building* Δ and a *group of Lie type* $G \leq \text{Aut}(\Delta)$. From this chapter going forward, we now assume in addition that G is a *Chevalley group*, meaning that G has type $A_n, B_n, C_n, D_n, G_2, F_4, E_6, E_7$ or E_8 . This allows us to directly cite the theorems of Ronan and Smith, who also make this restriction (although they note that many of their results hold in more general contexts, including for the twisted finite groups of Lie type [23, p.321, (i)]). For our main theorems later on in this thesis we will specialise further, taking $G = G_2(\mathbb{F}_q)$.

We can make use of the building Δ when studying the representation theory of G . By associating to each simplex $\sigma \in \Delta$ a representation of its stabiliser, the parabolic subgroup P_σ , we form a *sheaf* \mathcal{F} on the building Δ . Then we can use a technique called *sheaf homology* to obtain a representation of the group G . This approach was first used by Ronan and Smith in [23], with the theory developed further in [24] and [25].

To quote Segev and Smith [26, p.493], this technique is a ‘local approach’ to group representation theory. Using sheaf homology it is possible to construct irreducible representations of a group G algorithmically, from known irreducible representations of the parabolic subgroups of G [23, p.320].




Figure 5.1: The coefficient system must satisfy $\varphi_{\rho\sigma} \circ \varphi_{\sigma\tau} = \varphi_{\rho\tau}$ for all $\rho \succ \sigma \succ \tau$.

5.1 Sheaves on buildings: definition and examples

Let $k = \mathbb{F}_q$ be a finite field of characteristic π . Take a universal Chevalley group G defined over k and let Δ be the building of G , considered as a simplicial complex on which G acts on the right.

5.1.1 Definition

The following setup and notation is all taken from Ronan and Smith [23]. We define a **coefficient system** \mathcal{F} to be a set of k -vector spaces \mathcal{F}_σ , one associated with each simplex $\sigma \in \Delta$, along with a linear map $\varphi_{\sigma\tau}: \mathcal{F}_\sigma \rightarrow \mathcal{F}_\tau$ for every face-relation $\sigma \succ \tau$ in Δ . For any triple $\rho \succ \sigma \succ \tau$, our maps must satisfy the relation

$$\varphi_{\rho\sigma} \circ \varphi_{\sigma\tau} = \varphi_{\rho\tau}, \quad (5.1.1)$$

as shown in Figure 5.1. By a slight abuse of notation, we also write \mathcal{F} for the direct sum of all the vector spaces \mathcal{F}_σ .

We define a right G -action on \mathcal{F} , which therefore becomes a representation of G . Denote the image of g in this representation by \tilde{g} . We require the action to satisfy the following property: the restriction of \tilde{g} to some \mathcal{F}_σ is a linear map $\tilde{g}_\sigma: \mathcal{F}_\sigma \rightarrow \mathcal{F}_{\sigma g}$, for

all $g \in G$ and $\sigma \in \Delta$. In other words, the action of G on the coefficient system \mathcal{F} must follow the action of G on the simplicial complex Δ .

Finally, we require that the maps \tilde{g}_σ must commute with the connecting maps $\varphi_{\sigma\tau}$, so that this diagram commutes whenever $\sigma \succ \tau$:

$$\begin{array}{ccc}
 \mathcal{F}_\sigma & \xrightarrow{\varphi_{\sigma\tau}} & \mathcal{F}_\tau \\
 \tilde{g}_\sigma \downarrow & & \downarrow \tilde{g}_\tau \\
 \mathcal{F}_{\sigma g} & \xrightarrow{\varphi_{\sigma g, \tau g}} & \mathcal{F}_{\tau g}
 \end{array} \tag{5.1.2}$$

This property is called **G -equivariance**.

Again, following the nomenclature of Ronan and Smith [23], we will refer to such a coefficient system \mathcal{F} along with connecting maps $\varphi_{\sigma\tau}$ and associated G -action as a **sheaf**.

Recall from Chapter 2 the subgroups $B, N, T \leq G$. Let $\sigma \in \Delta$ and denote the stabiliser of σ in G as P_σ . We call P_σ a *parabolic subgroup*; if σ is a maximal simplex (chamber) then P_σ is a conjugate of B and we call it a *minimal* or *rank 0* parabolic; otherwise, P_σ is a parabolic subgroup of higher rank and contains a conjugate of B . Note that:

- (i) The set of elements $\{\tilde{g}_\sigma : g \in P_\sigma\}$ form a representation of P_σ , since each $\tilde{g}_\sigma \in \text{End}(\mathcal{F}_\sigma)$.
- (ii) If $h \in P_\sigma$ then $(h^g)\tilde{g}_{\sigma g} = (g^{-1}hg)\tilde{g}_{\sigma g} \in \text{End}(\mathcal{F}_{\sigma g})$. This is because the action of g^{-1} maps $\mathcal{F}_{\sigma g}$ to \mathcal{F}_σ , the action of h leaves \mathcal{F}_σ invariant, and then the action of g maps \mathcal{F}_σ back to $\mathcal{F}_{\sigma g}$.

Definition 5.1.3 (Subsheaf) Suppose \mathcal{G} is a sheaf on Δ such that each $\mathcal{G}_\sigma \subseteq \mathcal{F}_\sigma$, and the connecting maps $\psi_{\sigma\tau}$ of \mathcal{G} are the restrictions of the connecting maps $\varphi_{\sigma\tau}$ of \mathcal{F} , so $\psi_{\sigma\tau} = \varphi_{\sigma\tau}|_{\mathcal{G}_\sigma}$. Then we say \mathcal{G} is a **subsheaf** of \mathcal{F} and write $\mathcal{G} \subseteq \mathcal{F}$. Also, we say that \mathcal{G} is a **proper subsheaf** of \mathcal{F} if $\mathcal{G} \subseteq \mathcal{F}$ and $\mathcal{G} \neq \mathcal{F}$.

Let V be some kG -module. As we will see over the next sections, we can create a

sheaf on Δ by taking certain ‘compatible’ submodules of V at each simplex, and using corresponding restrictions of the identity map $\text{id}: V \rightarrow V$ as the connecting maps.

5.1.2 The constant sheaf

Our first example is the **constant sheaf** denoted \mathcal{K}_V , which we define to be the sheaf with the module V at each simplex and where each homomorphism $\varphi_{\sigma\tau}$ is the identity map $V \rightarrow V$.

The G -action on \mathcal{K}_V can be formed as follows. Fix a chamber (maximal simplex) c . Like any maximal simplex in a building, c contains one face of each cotype. We use c and its faces as orbit representatives; every G -orbit of simplices contains either c or one of its faces. Suppose $\sigma \prec c$ with stabiliser P_σ . Let $\{g_1, \dots, g_r\}$ be a set of representatives for the right cosets of G/P_σ , and induce the module V from P_σ up to G . The induced representation acts on the space $\bigoplus_{i=1}^r Vg_i$, so we have a copy of V for each simplex in the orbit, and the G -action on the direct sum follows the action on the building.

We do the same for all the other faces, including c itself. Then \mathcal{K}_V is the direct sum of the induced modules over all simplex cotypes.

5.1.3 The fixed-point sheaf

The fixed-point sheaf \mathcal{F}_V is a particular subsheaf of the constant sheaf \mathcal{K}_V , and will be the main object of study in this thesis. In order to describe it, we need some preliminary definitions.

Preliminary definitions and lemmas

For M a kG -module and $H \leq G$, there are two vector subspaces of M which we will use frequently.

Definition 5.1.4 (Centraliser $C_M(\mathbf{H})$) *Let M be a kG -module and $H \leq G$. Then we*

define the **centraliser**

$$C_M(H) := \{v \in M \mid vg = v \text{ for all } g \in H\}.$$

Definition 5.1.5 (Commutator $[M, H]$) Let M be a kG -module and $H \leq G$. The **commutator** $[M, H]$ is defined as the vector span

$$[M, H] := \langle vg - v \mid v \in M, g \in H \rangle.$$

Some standard theory of commutators and centralisers follows:

Lemma 5.1.6 Let M be a kG -module and U be a subspace of M . Then U is a kG -submodule of M if and only if $[U, G] \leq U$.

Proof. Assume U is a kG -submodule of M . Let $g \in G$. Then $ug \in U$, so $ug - u \in U$, and therefore $[U, G] \leq U$.

Now assume that $[U, G] \leq U$ and let $g \in G$. Then $w := ug - u \in U$, so $ug = u + w \in U$. Thus U is a kG -submodule of M . □

Lemma 5.1.7 Suppose $U \leq M$, $H \leq G$ and $[U, H] = \{0\}$. Then $U \leq C_M(H)$.

Proof. Since $[U, H] = \{0\}$, we have $ug - u = 0$ for all $u \in U, g \in H$. Thus $ug = u$ for all $u \in U, g \in H$, and so $U \leq C_M(H)$. □

Lemma 5.1.8 Suppose $\{M_1, \dots, M_n\}$ is a set of kG -submodules of M and $H \leq G$. Then

$$\left[\sum_{i=1}^n M_i, H \right] = \sum_{i=1}^n [M_i, H].$$

Proof. We have

$$\begin{aligned} \left[\sum_{i=1}^n M_i, H \right] &= \langle vh - v \mid v \in \sum_{i=1}^n M_i, h \in H \rangle \\ &= \sum_{i=1}^n \langle vh - v \mid v \in M_i, h \in H \rangle. \end{aligned} \quad \square$$

In [28, Proposition and Corollary, p.286], Smith gives a version of the following result which is now known as Smith's Lemma. For our specific statement and a more concise proof, in particular of the second part, see Timmesfeld [31] or Meierfrankenfeld and Stellmacher [18, p.21, Lemma 4.1]. Recall that for X a finite group,

$$O^{\pi'}(X) := \langle \text{Syl}_{\pi}(X) \rangle.$$

In the case that L is a Levi complement in a Chevalley group G , we have that $O^{\pi'}(L)$ is a Chevalley group as well.

Lemma 5.1.9 (Smith's Lemma) *Let G be a Chevalley group over a finite field k , P a parabolic subgroup with unipotent radical U and Levi complement L , and V a finite-dimensional irreducible kG -module. Then the centraliser $C_V(U)$ affords an irreducible kL -module as well as an irreducible $kO^{\pi'}(L)$ -module.*

Furthermore, if V is the irreducible kG -module of highest weight $\lambda: T \rightarrow k^\times$, then $C_V(U)$ (considered as a $kO^{\pi'}(L)$ -module) is the unique irreducible $kO^{\pi'}(L)$ -module with highest weight $\lambda|_{T \cap O^{\pi'}(L)}$.

Definition of the fixed-point sheaf

Let V be an irreducible kG -module. Then the **fixed-point sheaf** \mathcal{F}_V is the subsheaf of \mathcal{K}_V having sheaf terms given by $\mathcal{F}_\sigma := C_V(U_\sigma)$, with connecting maps $\varphi_{\sigma\tau}$ given by inclusion. We will check that \mathcal{F}_V is indeed a sheaf as per our definition.

- (i) Firstly, we need to verify that the subspace $C_V(U_\sigma)$ is indeed a kP_σ -module, so that $\tilde{g}_\sigma \in \text{End}(\mathcal{F}_\sigma)$. By definition, $C_V(U_\sigma)$ is fixed pointwise by U_σ . By Smith's Lemma (Lemma 5.1.9), $C_V(U_\sigma)$ affords an irreducible kL_σ -module, and is therefore also fixed (not necessarily pointwise) by L_σ . Since $P_\sigma = \langle U_\sigma, L_\sigma \rangle$ we have that $C_V(U_\sigma)$ is fixed by P_σ , so is a kP_σ -module. Hence for $g \in P_\sigma$ we have $\tilde{g}_\sigma \in \text{End}(\mathcal{F}_\sigma)$, as required.
- (ii) Also we must check that the inclusion behaves as desired; that is, if $\tau \prec \sigma$ then $\mathcal{F}_\sigma \subseteq \mathcal{F}_\tau$ so that $\varphi_{\sigma\tau}: \mathcal{F}_\sigma \rightarrow \mathcal{F}_\tau$ given by $v \mapsto v$ makes sense. By Lemma 2.4.6, $\tau \prec \sigma$ implies that $P_\sigma \subseteq P_\tau$. Then by Lemma 2.4.8, $P_\sigma \subseteq P_\tau$ implies that $U_\sigma \supseteq U_\tau$, and so $C_V(U_\sigma) \subseteq C_V(U_\tau)$ as required.
- (iii) Finally, the G -equivariance property follows immediately from the fact that all the maps $\varphi_{\sigma\tau}$ are just restrictions of the identity map.

In relation to the fixed-point sheaves we have defined, Smith's Lemma tells us that if V is an irreducible kG -module, then the kP_σ -modules we have defined at each simplex are irreducible when considered as kL_σ -modules (and thus as kP_σ -modules).

The following theorem of Steinberg [16, p.42, Theorem] allows us to prove a lemma regarding the dimension of terms at chambers in a fixed-point sheaf.

Theorem 5.1.10 *Suppose that G is a Chevalley group defined over k . Then k is a splitting field for every irreducible kG -module.*

When we say that k is a *splitting field* for an irreducible kG -module V , we mean that V remains irreducible as a KG -module, where K is any field extension of k .

Lemma 5.1.11 *Let V be an irreducible kG -module and $\sigma \in \Delta$ be a chamber. Then $(\mathcal{F}_V)_\sigma$ is an irreducible kT -module. In particular, $\dim(\mathcal{F}_V)_\sigma = 1$.*

Proof. The stabiliser of a chamber is a conjugate of B , and so the term $(\mathcal{F}_V)_\sigma$ is given by $C_V(U)$ where $U = O_\pi(B)$. By Smith's Lemma (Lemma 5.1.9), we have that $(\mathcal{F}_V)_\sigma$ is an irreducible representation of T , the Levi complement of B .

The subgroup T is itself a Chevalley group defined over k , so k is a splitting field for the kT -module $(\mathcal{F}_V)_\sigma$ by Theorem 5.1.10, and thus $(\mathcal{F}_V)_\sigma$ is irreducible as a KT -module where K is the algebraic closure of k . Hence, as T is abelian, $\dim(\mathcal{F}_V)_\sigma = 1$ by Schur's Lemma. \square

5.2 Sheaf morphisms

Let \mathcal{F} and \mathcal{G} be sheaves on Δ with connecting maps $\varphi_{\sigma\tau}$ and $\psi_{\sigma\tau}$ respectively. We can define a map $m: \mathcal{F} \rightarrow \mathcal{G}$ by choosing a linear map $m_\sigma: \mathcal{F}_\sigma \rightarrow \mathcal{G}_\sigma$ for each $\sigma \in \Delta$ such that each $m_\sigma \in \text{End}(\mathcal{F}_\sigma, \mathcal{G}_\sigma)$. This map is called a *sheaf morphism* if the following diagrams commute:

$$\begin{array}{ccc} \mathcal{F}_\sigma & \xrightarrow{m_\sigma} & \mathcal{G}_\sigma \\ \tilde{g}_\sigma \downarrow & & \downarrow \tilde{g}_\sigma \\ \mathcal{F}_{\sigma g} & \xrightarrow{m_{\sigma g}} & \mathcal{G}_{\sigma g} \end{array} \quad (5.2.1)$$

for all $g \in G$, and

$$\begin{array}{ccc} \mathcal{F}_\sigma & \xrightarrow{m_\sigma} & \mathcal{G}_\sigma \\ \varphi_{\sigma\tau} \downarrow & & \downarrow \psi_{\sigma\tau} \\ \mathcal{F}_\tau & \xrightarrow{m_\tau} & \mathcal{G}_\tau \end{array} \quad (5.2.2)$$

for all $\sigma \succ \tau \in \Delta$.

Since sheaf maps $\mathcal{F} \rightarrow \mathcal{G}$ are defined as collections of k -linear maps from each term \mathcal{F}_σ to \mathcal{G}_σ , we can add them and multiply by scalars from k in the natural way. Hence we may form a k -vector space $\text{Hom}_k(\mathcal{F}, \mathcal{G})$ consisting of all the sheaf maps from \mathcal{F} to \mathcal{G} .

5.2.1 Kernels and images

Let \mathcal{F} and \mathcal{G} be sheaves on Δ and $m: \mathcal{F} \rightarrow \mathcal{G}$ be a sheaf morphism. We can define the *kernel* of m , denoted \mathcal{K}_m , to be the sheaf consisting of terms $\ker(m_\sigma)$, where the

connecting maps and G -action are inherited from \mathcal{F} . Similarly, we define the *image* \mathcal{I}_m to consist of terms $\text{im}(m_\sigma)$, where connecting maps and G -action are inherited from \mathcal{G} .

We have the usual notions of *injectivity* ($\ker(m)$ has a 0-dimensional vector space at each simplex) and *surjectivity* ($\text{im}(m) = \mathcal{G}$); an injective and surjective sheaf morphism is called a *sheaf isomorphism*.

5.3 Sheaf homology

Suppose Δ has rank n . In the same way that we consider \mathcal{F} as the formal direct sum of all terms \mathcal{F}_σ , we can define a **chain complex** consisting of **chain spaces** $C_r(\Delta, \mathcal{F})$, for $0 \leq r < n$, which are the formal direct sum of terms at all simplices of dimension r , and **boundary maps** between adjacent chain spaces, going from each $C_r(\Delta, \mathcal{F})$ into $C_{r-1}(\Delta, \mathcal{F})$. Having done this, we we can write

$$\mathcal{F} = \bigoplus_{r=0}^n C_r(\Delta, \mathcal{F}).$$

Thus we define $C_0(\Delta, \mathcal{F})$ to be the formal direct sum of all terms at 0-simplices (vertices) in Δ ; define $C_1(\Delta, \mathcal{F})$ to be the formal direct sum of all terms at 1-simplices (edges) in Δ , and so on, until we reach $C_{n-1}(\Delta, \mathcal{F})$ which is the formal direct sum of all terms at maximal simplices (chambers) of Δ .

To form the boundary maps, we make use of the connecting maps $\varphi_{\sigma\tau}$. We modify these slightly by choosing a sign (+/−) for each in the following manner.

Recall that each 0-simplex (vertex) in Δ is of cotype $\{i\}$ for some $i \in \{1, \dots, n\}$; every 1-simplex (edge) is of cotype $\{i_0, i_1\} \subseteq \{1, \dots, n\}$ with $i_0 < i_1$, and so on, up to (maximal) $(n - 1)$ -simplices with cotype $\{1, \dots, n\}$.

In general, if a simplex σ has dimension r then its cotype has $r + 1$ elements; $J := \{i_0, \dots, i_r\}$. Furthermore, for each face $\tau \prec \sigma$ of dimension $r - 1$, the cotype of τ is

precisely the cotype of σ with one element removed: $J \setminus \{i_j\}$. Note that σ has exactly $r + 1$ such maximal faces, one corresponding to each i_j . Let us denote the maximal face corresponding to the removal of i_j by τ_j , and define

$$\delta_{\sigma\tau_j} := (-1)^j \varphi_{\sigma\tau_j}.$$

Take the sum of all these $r + 1$ maps emanating from σ , to form a map

$$\delta_\sigma := \sum_{j=0}^r \delta_{\sigma\tau_j}.$$

Then, take the sum of all of *these* maps over every simplex of dimension r , to obtain

$$\delta_r := \sum_{\dim \sigma = r} \delta_\sigma.$$

For $0 < r < n$, this is the boundary map $\delta_r: C_r(\Delta, \mathcal{F}) \rightarrow C_{r-1}(\Delta, \mathcal{F})$. Finally, define δ_0 to be the map sending every vector to 0. We write δ for the sum of all the boundary maps, so

$$\delta = \sum_{r=0}^{n-1} \delta_r$$

maps a vector in any chain space to a vector in the chain space below.

A chain complex must have the property $\delta^2 = 0$; that is, the composition of two consecutive boundary maps gives the zero map. (In other words, the image of each boundary map is in the kernel of the next one.) We need to check that the complex we have defined has this property; it suffices to show this holds locally, so that for any $v \in \mathcal{F}_\sigma$ (where $\dim \sigma = r$) we have $v\delta_\sigma\delta_{r-1} = 0$.

As an example, consider some 2-simplex ρ , with 1-faces σ_1, σ_2 and σ_3 and 0-faces τ_1, τ_2 and τ_3 , with face-relations as shown in Figure 5.2. Suppose ρ has cotype $I = \{1, 2, 3\}$, each 1-simplex σ_i has cotype $I \setminus \{i\}$ and each 0-simplex τ_j has cotype $\{j\}$. The δ -maps




Figure 5.2: Example of simplex cotypes on a 2-simplex.

$\rho \rightarrow \sigma_i$	$\sigma_i \rightarrow \tau_j$
$\delta_{\rho\sigma_1} = \varphi_{\rho\sigma_1}$	$\delta_{\sigma_1\tau_2} = \varphi_{\sigma_1\tau_2}$ $\delta_{\sigma_1\tau_3} = -\varphi_{\sigma_1\tau_3}$
$\delta_{\rho\sigma_2} = -\varphi_{\rho\sigma_2}$	$\delta_{\sigma_2\tau_1} = \varphi_{\sigma_2\tau_1}$ $\delta_{\sigma_2\tau_3} = -\varphi_{\sigma_2\tau_3}$
$\delta_{\rho\sigma_3} = \varphi_{\rho\sigma_3}$	$\delta_{\sigma_3\tau_1} = \varphi_{\sigma_3\tau_1}$ $\delta_{\sigma_3\tau_2} = -\varphi_{\sigma_3\tau_2}$

Table 5.1: The δ maps for Figure 5.2.

are therefore as shown in Table 5.1. Then the restriction of the boundary map to the simplex ρ is given by

$$\delta_\rho = \varphi_{\rho\sigma_1} \circ (\varphi_{\sigma_1\tau_2} - \varphi_{\sigma_1\tau_3}) - \varphi_{\rho\sigma_2} \circ (\varphi_{\sigma_2\tau_1} - \varphi_{\sigma_2\tau_3}) + \varphi_{\rho\sigma_3} \circ (\varphi_{\sigma_3\tau_1} - \varphi_{\sigma_3\tau_2}).$$

We can compose the maps using (5.1.1) to obtain

$$\delta_\rho = (\varphi_{\rho\tau_1} - \varphi_{\rho\tau_1}) + (\varphi_{\rho\tau_2} - \varphi_{\rho\tau_2}) + (\varphi_{\rho\tau_3} - \varphi_{\rho\tau_3}) = 0,$$

as required.

It is not difficult to see that this works in general. For each $(r - 2)$ -dimensional face τ of an r -simplex ρ , there are two elements i_a and i_b (with $a < b$) in the cotype of ρ but not of τ . We can either remove i_a first, and follow the path corresponding to cotypes $J \rightarrow J \setminus \{i_a\} \rightarrow J \setminus \{i_a, i_b\}$; or we can remove i_b first, following the path corresponding to $J \rightarrow J \setminus \{i_b\} \rightarrow J \setminus \{i_a, i_b\}$. Suppose the first path goes through $(r - 1)$ -simplex σ_1 and the second through σ_2 . If i_a and i_b both have an even-numbered position in J (indexed counting from 0), then $\delta_{\rho\sigma_1}$ has positive sign (*i.e.* $\delta_{\rho\sigma_1} = +\varphi_{\rho\sigma_1}$); but with i_a removed, i_b now has odd index in the cotype of σ_1 , so $\delta_{\sigma_1\tau}$ has negative sign, and so does the composition. The map $\delta_{\rho\sigma_2}$ has positive sign, and i_a has even index in the cotype of σ_2 (since $i_b > i_a$ so the order is unaffected), so this composition has positive sign. Thus the compositions sum to the zero map.

The other three parity cases can be computed similarly and all result in zero maps, which is a nice exercise.

Satisfied that $\delta^2 = 0$, we are now able to define the **cycles** $Z_r(\Delta, \mathcal{F}) := \ker(\delta_r)$ and **boundaries** $B_r(\Delta, \mathcal{F}) := \text{im}(\delta_{r+1})$, and the **homology spaces** which are the quotients

$$H_r(\Delta, \mathcal{F}) := Z_r(\Delta, \mathcal{F})/B_r(\Delta, \mathcal{F}).$$

Note that each $g \in G$ acts on the simplicial complex whilst preserving simplex cotype, and hence rank. Thus $C_r(\Delta, \mathcal{F}) \cdot g = C_r(\Delta, \mathcal{F})$ for all $g \in G$ and $0 \leq r < n$. Since the G -action commutes with each $\varphi_{\sigma\tau}$, it must also commute with the boundary maps δ_r . Hence the kernels $Z_r(\Delta, \mathcal{F})$, images $B_r(\Delta, \mathcal{F})$ and homology spaces $H_r(\Delta, \mathcal{F})$ are all stabilised by G , which implies that they are all kG -modules.

Notation: Where the building Δ in question is clear, we may omit it from the notation and simply write $H_0(\mathcal{F})$, for example.

5.3.1 Results of Ronan and Smith

We now present some key results of Ronan and Smith. The following lemma is one of the most fundamental results about sheaf homology on Δ .

Lemma 5.3.1 *Let V be some kG -module. Then $H_0(\mathcal{K}_V) \cong V$.*

Proof. See [23, p.324, Lemma 1.1]. □

The rest of the results we present are working towards the following theorem ([23, p.331, Theorem 2.3]) which will be crucial later on in this thesis.

Theorem 5.3.2 *Suppose V is an irreducible kG -module. Then $H := H_0(\mathcal{F}_V)$ contains a unique maximal kG -submodule K , and we have $H/K \cong V$. In particular, H is indecomposable.*

We require three intermediate results before we are able to prove Theorem 5.3.2. Recall the Frobenius-Nakayama formula, relating the operations of module induction and restriction: if $H \leq G$, for any kH -module M and any kG -module N , we have

$$\mathrm{Hom}_{kG}(N, \mathrm{Ind}_H^G(M)) \cong \mathrm{Hom}_{kH}(\mathrm{Res}_H^G(N), M)$$

[3, p.60]. Theorem 5.3.3, due to Ronan and Smith [23, p.325, Theorem 1.2], is an analogue of this correspondence. In place of induction, we use the operation of forming the constant sheaf \mathcal{K}_W from a module W . Taking the place of restriction we have H_0 .

Theorem 5.3.3 *Let W be a kG -module, \mathcal{K}_W be the constant sheaf of W on Δ , and \mathcal{F} be an arbitrary sheaf on Δ . Then $\text{Hom}(\mathcal{F}, \mathcal{K}_W) \cong_k \text{Hom}(H_0(\mathcal{F}), W)$.*

Proof. We begin by defining a map $\varphi_\sigma: \mathcal{F}_\sigma \rightarrow H_0(\mathcal{F})$ for each simplex $\sigma \in \Delta$. If σ is a vertex, then we define φ_σ to be the restriction to \mathcal{F}_σ of the natural map $\varphi: C_0(\mathcal{F}) \rightarrow H_0(\mathcal{F})$. If σ is not a vertex, we choose any of its vertices (0-faces) τ and compose φ_τ through the connecting map $\varphi_{\sigma\tau}$.

This map is well-defined because the choice of τ does not matter: if τ' is another vertex of σ then there is an edge (1-simplex) e connecting τ with τ' , since they are both faces of σ . The map φ has kernel $B_0(\mathcal{F})$, and the image of δ_e is contained in $B_0(\mathcal{F})$, so $\delta_e \circ \varphi = 0$. Recall that $\delta_e = \pm(\varphi_{e\tau} - \varphi_{e\tau'})$, the sign depending on the respective cotypes of τ and τ' . In either case we obtain $\varphi_{e\tau} \circ \varphi_\tau = \varphi_{e\tau'} \circ \varphi_{\tau'}$, and so

$$\begin{aligned} \varphi_{\sigma\tau} \circ \varphi_\tau &= \varphi_{\sigma e} \circ (\varphi_{e\tau} \circ \varphi_\tau) \\ &= \varphi_{\sigma e} \circ (\varphi_{e\tau'} \circ \varphi_{\tau'}) \\ &= \varphi_{\sigma\tau'} \circ \varphi_{\tau'}. \end{aligned}$$

Hence we may set $\varphi_\sigma := \varphi_{\sigma\tau} \circ \varphi_\tau$, and this is well-defined. These maps have two nice properties.

- (i) Firstly, if $\rho \succ \sigma$ then we have $\varphi_{\rho\sigma} \circ \varphi_\sigma = \varphi_\rho$. If σ is a vertex then this is by definition, so assume not. Then there is a vertex τ such that $\rho \succ \sigma \succ \tau$, and

$$\varphi_\rho = \varphi_{\rho\tau} \circ \varphi_\tau = (\varphi_{\rho\sigma} \circ \varphi_{\sigma\tau}) \circ \varphi_\tau = \varphi_{\rho\sigma} \circ \varphi_\sigma,$$

as required.

- (ii) Secondly, for any $g \in G$ and any $\sigma \in \Delta$, we have $g \circ \varphi_{\sigma g} = \varphi_{\sigma} \circ g$. This is due to the G -equivariance (see Diagram 5.1.2) of the connecting maps and hence of the homology spaces.

With this notation in hand, we can prove the theorem. We will begin by taking an arbitrary kG -module homomorphism $\alpha: H_0(\mathcal{F}) \rightarrow W$ and using it form a sheaf morphism $\bar{\alpha}: \mathcal{F} \rightarrow \mathcal{K}_W$. For each $\sigma \in \Delta$, set $\bar{\alpha}_{\sigma} = \varphi_{\sigma} \circ \alpha$. Since the mapping from α to $\bar{\alpha}$ is k -linear, it just remains to check the two conditions for a sheaf morphism from Section 5.2. The first condition, that $\bar{\alpha}$ commutes with G (Diagram 5.2.1) is satisfied by property (ii) of the φ_{σ} maps above. The second condition, that $\bar{\alpha}$ commutes with the connecting maps $\varphi_{\sigma\tau}$ (Diagram 5.2.2) follows due to property (i) above.

Going in the other direction, we will start with an arbitrary sheaf morphism $\beta: \mathcal{F} \rightarrow \mathcal{K}_W$, and use it to form a kG -module homomorphism $\widehat{\beta}: H_0(\mathcal{F}) \rightarrow W$. Note that for $\sigma \in \Delta$, the restriction β_{σ} is a kG -module homomorphism $\mathcal{F}_{\sigma} \rightarrow W$. Hence any $v \in \mathcal{F}_{\tau}$ where τ is a vertex we define $(v)\widehat{\beta} := (v)\beta_{\tau}$. This is in fact a map from $C_0(\mathcal{F})$ to W , but it is easy to check that it vanishes on $B_0(\mathcal{F})$. Consider an edge e and $v \in \mathcal{F}_e$; we have

$$(v)\delta_e \circ \widehat{\beta} = (v)\varphi_{e\tau} \circ \beta_{\tau} - (v)\varphi_{e\tau'} \circ \beta_{\tau'} = (v)\beta_e - (v)\beta_e = 0,$$

and so $(B_0(\mathcal{F}))\widehat{\beta} = 0$ as required. Clearly $\widehat{\beta}$ is k -linear, and for a vertex σ , a vector $v \in \mathcal{F}_{\sigma}$, and any $g \in G$ we have

$$(v)g\widehat{\beta} = (v)g\beta_{\sigma g} = (v)\beta_{\sigma}g = (v)\widehat{\beta}g,$$

so $\widehat{\beta}: H_0(\mathcal{F}) \rightarrow W$ is a kG -module homomorphism as required.

That $\widehat{\alpha} = \alpha$ and $\widetilde{\beta} = \beta$ follows from their definitions, which completes the proof. \square

The following definitions are also due to Ronan and Smith [23, p.323-324].

Definition 5.3.4 (Chamber-generated sheaf) *Recall that a chamber σ is a maximal simplex of Δ . We say that a sheaf \mathcal{F} is **chamber-generated** if, for any simplex $\tau \in \Delta$, we have*

$$\mathcal{F}_\tau = \sum_{\substack{\sigma \text{ a chamber} \\ \sigma > \tau}} (\mathcal{F}_\sigma) \varphi_{\sigma\tau}.$$

Definition 5.3.5 (Irreducible sheaf) *Suppose \mathcal{F} is a chamber-generated sheaf. We say that \mathcal{F} is **irreducible** if it does not contain a non-zero, proper, chamber-generated subsheaf.*

The next two results appear as Theorem 2.1 and Lemma 2.2 in Ronan and Smith [23, p.330].

Theorem 5.3.6 *Suppose V is an irreducible kG -module. Then \mathcal{F}_V is an irreducible sheaf.*

Proof. Let $\mathcal{F} = \mathcal{F}_V$. We begin by showing that \mathcal{F} is chamber generated. This is equivalent to showing that

$$\sum_{\substack{\sigma \in \Delta \\ \sigma \text{ a chamber}}} C_V(U_\sigma) = V.$$

Indeed, let σ be a chamber of Δ . Then σg is a chamber of Δ and $C_V(U_\sigma)g = C_V(U_{\sigma g})$. Hence this sum is G -invariant. Since V is irreducible, and $C_V(U_\sigma) \neq \{0\}$, the sum is equal to V .

Now suppose $\mathcal{G} \subseteq \mathcal{F}$ with \mathcal{G} a non-zero, chamber-generated sheaf. Let $\tau \in \Delta$ with stabiliser $P_\tau = \langle U_\tau, L_\tau \rangle$. By Lemma 5.1.9, each module $\mathcal{F}_\tau = C_V(U_\tau)$ is an irreducible kL_τ -module, and therefore is also irreducible as a kP_τ -module because $L_\tau \leq P_\tau$. Since $\mathcal{G} \subseteq \mathcal{F}$ then $\mathcal{G}_\tau \subseteq \mathcal{F}_\tau$, and so $\mathcal{G}_\tau \in \{\mathcal{F}_\tau, 0\}$.

Suppose that $\mathcal{G}_\sigma = 0$ for some chamber (maximal simplex) $\sigma \in \Delta$. Then *all* chambers have a 0-dimensional module in \mathcal{G} , since they all lie in the same orbit under G . Therefore for any $\tau \in \Delta$, we have

$$\sum_{\substack{\sigma \text{ a chamber} \\ \sigma \succ \tau}} (\mathcal{G}_\sigma) \varphi_{\sigma\tau} = 0,$$

so since \mathcal{G} is chamber-generated, all coefficients \mathcal{G}_τ are 0. But this means that \mathcal{G} is the zero sheaf, a contradiction to the choice of \mathcal{G} .

Now suppose that $\mathcal{G}_\tau = 0$ for some (non-maximal) simplex $\sigma \in \Delta$. We know that $\mathcal{G}_\sigma \subseteq \mathcal{G}_\tau$ for all $\tau \prec \sigma$, because the connecting maps are restrictions of the identity map inherited from \mathcal{F}_V , and they could not be defined otherwise. So therefore $\mathcal{G}_\sigma = 0$ for any $\sigma \succ \tau$, which includes at least one chamber. Thus all chambers have 0-dimensional modules in \mathcal{G} and we reach the same contradiction as before. So we must have $\mathcal{G} = \mathcal{F}$, and so \mathcal{F} is irreducible. \square

Lemma 5.3.7 *Suppose V and W are irreducible kG -modules and $\mathcal{F}_V \cong \mathcal{F}_W$. Then $V \cong W$.*

Proof. Both V and W are irreducible kG -modules, so they each correspond to a q -restricted weight λ_v and λ_w . Suppose for a contradiction that V and W are not isomorphic; then $\lambda_v \neq \lambda_w$ by Theorem 4.1.7.

Suppose that G has simple roots $\{\alpha_1, \dots, \alpha_n\}$. Now, if we let $\lambda_v = \sum_{i=1}^n \beta_i \lambda_i$ and $\lambda_w = \sum_{i=1}^n \gamma_i \lambda_i$ then there is some $1 \leq i \leq n$ such that $\beta_i \neq \gamma_i$.

Recall that $P_i = \langle U_i, L_i \rangle$ is a parabolic subgroup stabilising an i -panel of Δ . By Lemma 5.1.9, the module $(\mathcal{F}_V)_\sigma = C_V(U_\sigma)$ is the irreducible kL_i -module with weight β_i , and the module $(\mathcal{F}_W)_\sigma = C_W(U_\sigma)$ is the irreducible kL_i -module with weight γ_i . Since $\beta_i \neq \gamma_i$, these irreducible kL_i modules cannot be isomorphic. But by assumption, since the sheaves are isomorphic, we have $(\mathcal{F}_V)_\sigma \cong (\mathcal{F}_W)_\sigma$, a contradiction. \square

With these results in hand, we are now able to prove Theorem 5.3.2, which we restate below for convenience.

Theorem *Suppose V is an irreducible kG -module. Then $H := H_0(\mathcal{F}_V)$ contains a unique maximal kG -submodule K , and we have $H/K \cong V$. In particular, H is indecomposable.*

Proof. Let $W \neq 0$ be an irreducible quotient of $H_0(\mathcal{F}_V)$ by some maximal submodule X , so we have a surjective kG -homomorphism $\alpha: H_0(\mathcal{F}_V) \rightarrow W$. Using the procedure from the proof of Theorem 5.3.3, we form the associated sheaf morphism $\bar{\alpha}: \mathcal{F}_V \rightarrow \mathcal{K}_W$.

Now, fix a simplex $\sigma \in \Delta$ with stabiliser $P_\sigma = \langle U_\sigma, L_\sigma \rangle$, and let $g \in U_\sigma$ so that $v\tilde{g} = v$ for all $v \in (\mathcal{F}_V)_\sigma$, and $\sigma g = \sigma$. By Diagram 5.2.2, we have that $v\bar{\alpha}_\sigma\tilde{g}$ must be equal to $v\tilde{g}\bar{\alpha}_{\sigma g} = v\bar{\alpha}_\sigma$; and so in fact $\text{im}(\bar{\alpha})$ must lie in \mathcal{F}_W , since each $v\bar{\alpha}_\sigma$ is fixed by all $\{\tilde{g} \mid g \in U_\sigma\}$. So we have $\bar{\alpha}: \mathcal{F}_V \rightarrow \mathcal{F}_W$, where \mathcal{F}_W is a quotient of \mathcal{F}_V .

By Theorem 5.3.6, the fact that W is irreducible implies that \mathcal{F}_W is an irreducible sheaf, so $\text{im}(\bar{\alpha})$ is either \mathcal{F}_W or the zero sheaf. But since the image of α is non-zero, the image of $\bar{\alpha}$ is non-zero by construction, and so $\text{im}(\bar{\alpha}) = \mathcal{F}_W$ and hence $\bar{\alpha}: \mathcal{F}_V \rightarrow \mathcal{F}_W$ is surjective.

Next we want to show that the kernel of $\bar{\alpha}$ is the zero sheaf. This follows from the fact that \mathcal{F}_V and \mathcal{F}_W are irreducible sheaves; they are both chamber-generated, so $\ker \bar{\alpha}$ must be a chamber-generated subsheaf of \mathcal{F}_V , and thus equal to either the zero sheaf or \mathcal{F}_V itself. But if $\ker \bar{\alpha} = \mathcal{F}_V$ then $\bar{\alpha}$ is not surjective, a contradiction. So $\ker \bar{\alpha} = \{0\}$.

Therefore $\bar{\alpha}$ is an isomorphism and so $\mathcal{F}_V \cong \mathcal{F}_W$. Thus by Lemma 5.3.7 we have $V \cong W$. So any non-zero irreducible quotient of $H_0(\mathcal{F}_V)$ is isomorphic to V itself. It remains to show that the choice of maximal submodule $X \subset H_0(\mathcal{F}_V)$ is unique.

The terms of \mathcal{F}_V at chambers are 1-dimensional by Lemma 5.1.11, and so $\text{End}_{kP_\sigma}((\mathcal{F}_V)_\sigma)$ is 1-dimensional for any chamber σ . But by G -equivariance, this implies that $\text{End}_k(C_{n-1}(\mathcal{F}_V))$

is 1-dimensional. Then since \mathcal{F}_V is chamber-generated we have that $\text{End}_k(\mathcal{F}_V)$ is 1-dimensional as well.

The proof above that a sheaf map from $\mathcal{F}_V \rightarrow \mathcal{K}_W$ has \mathcal{F}_W as its image also holds for maps $\mathcal{F}_V \rightarrow \mathcal{K}_V$; hence $\text{Hom}_k(\mathcal{F}_V, \mathcal{K}_V) = \text{End}_k(\mathcal{F}_V)$ is 1-dimensional. Thus by Theorem 5.3.3, $\text{Hom}_k(H_0(\mathcal{F}_V), V)$ is 1-dimensional, and therefore the maximal kG -submodule X is unique. \square

A consequence of Theorem 5.3.2 is that it facilitates a recursive construction of any irreducible kG -module V from modules for the parabolic subgroups, without prior knowledge of V itself [23, p.320]. By Smith's Lemma (Lemma 5.1.9), the sheaf terms $(\mathcal{F}_V)_\sigma$ are all irreducible — therefore we can enumerate all possible sheaves, since there are a finite number of choices for each simplex type. An algorithm for obtaining V from $H_0(\Delta, \mathcal{F}_V)$ is given in [23, p.334, Algorithm 3.2]. It is preceded by a discussion on how the connecting maps $\varphi_{\sigma\tau}$ may be formed without prior knowledge of V .

CHAPTER 6

SHEAVES ON THE BUILDING OF $G_2(k)$

Having introduced the group $G_2(k)$ and its associated building (the generalised hexagon) in Chapter 3, and sheaves on buildings in Chapter 5, we can now turn our attention to the specific problem at hand. The fixed-point sheaf $\mathcal{F}_{\overline{C}}$ on the building of $G_2(k)$, where \overline{C} is the Cayley module, is investigated by Ronan and Smith in [23, Example 3.3] and by Segev and Smith in [26]. Their results and discussions are summarised below.

Using computer calculation, Ronan and Smith provide dimension data for $H_0(\mathcal{F}_{\overline{C}})$ over fields of order $q = 2, 3$. Using the computer algebra system MAGMA we have computed data for finite fields of order up to 11.

6.1 Data for some small fields

Uniquely in characteristic 2, the module \overline{C} is not irreducible but contains a 1-dimensional irreducible submodule X . Taking instead the irreducible 6-dimensional quotient \overline{C}/X , the dimension of $H_0(\mathcal{F}_{\overline{C}/X})$ for the field $k = \mathbb{F}_2$ was calculated by Ronan and Smith to be 14 [23, p.335, Section (3.3)]. For $k = \mathbb{F}_3$, the 7-dimensional module \overline{C} is irreducible. In this case, Ronan and Smith also obtained $\dim H_0(\mathcal{F}_{\overline{C}}) = 14$.

Using the computer algebra system MAGMA [5], results for some larger fields were obtained which are displayed in Table 6.1.

$q = k $	V	$\dim V$	$\dim H_0(\mathcal{F}_V)$
2	\overline{C}/X	6	14
3	\overline{C}	7	14
9	\overline{C}	7	14
5	\overline{C}	7	7
7	\overline{C}	7	7
11	\overline{C}	7	7

Table 6.1: Dimensions of $H_0(\mathcal{F}_V)$ for some small fields, all calculated using MAGMA except for $q = 2$ which is from Ronan and Smith [23, p.335, Section (3.3)].

Segev and Smith note that, in the calculated examples available to them (*i.e.* for $q = 2, 3$), we have $\dim H_0(\mathcal{F}_{\overline{C}}) > q$ — but that ‘if [this] holds in general, the situation would run counter to the intuition of many geometers’ [26, p.497]. Indeed, the calculations for $q = 7$ and $q = 11$ have demonstrated that $\dim H_0(\mathcal{F}_{\overline{C}}) > q$ does not hold in general. Table 6.1 hints at a possible conjecture: that $\dim H_0(\mathcal{F}_{\overline{C}}) = 7$ if $\pi > 3$ (where $q = \pi^a$). With this in mind, we introduce the following standing assumption:

Assumption 6.1.1 *The order of k is $q = \pi^a$ for some $\pi > 3$.*

6.2 Setting up notation

Denote $\underline{C}_0 := C_0(\mathcal{F}_{\overline{C}}) = Z_0(\mathcal{F}_{\overline{C}})$ and $H := H_0(\mathcal{F}_{\overline{C}})$. The module H is defined as a quotient of \underline{C}_0 , and by Theorem 5.3.2 we know that \overline{C} is a quotient of H . Hence we have a sequence of kG -modules with natural maps between them given by

$$\underline{C}_0 \xrightarrow{\alpha} H \xrightarrow{\theta} \overline{C}.$$

We will use the following notational convention. Subspaces of \underline{C}_0 will be written with an underline $\underline{\quad}$, and subspaces of the Cayley module \overline{C} with an overline $\overline{\quad}$. Thus we will

write, for example, $(\underline{X})\alpha = X$ and $(X)\theta = \overline{X}$.

6.2.1 Subspaces of \underline{C}_0 , H and \overline{C}

Recall that \underline{C}_0 is the direct sum of the subspaces corresponding to points and lines of the geometry. For a point p in the geometry, denote by \underline{C}_p the subspace of \underline{C}_0 corresponding to p , so $\underline{C}_p := (\mathcal{F}_{\overline{C}})_p$. Now, the sheaf term corresponding to p is given by

$$(\mathcal{F}_{\overline{C}})_p = C_{\overline{C}}(U_p),$$

where $G_p = \langle U_p, L_p \rangle$ is a Levi decomposition of the point stabiliser of p .

Let us compute the dimension of \underline{C}_p , which is independent of the choice of p . Without loss of generality, therefore, we may take $p = \mathbf{p} = \langle \overline{1} \rangle$. As calculated in Chapter 3, we have $U_{\mathbf{p}} = \langle A(\lambda), \dots, D(\lambda), F(\lambda) \mid \lambda \in k \rangle$. Inspecting the matrices of these generators, we see that they fix pointwise only the 1-space $\langle \overline{1} \rangle$ and so $\dim \underline{C}_p = \dim(\mathcal{F}_{\overline{C}})_p = 1$.

Similarly, in $\mathcal{F}_{\overline{C}}$ we have a sheaf term associated to the line ℓ , given by

$$(\mathcal{F}_{\overline{C}})_\ell = C_{\overline{C}}(U_\ell),$$

where $G_\ell = \langle U_\ell, L_\ell \rangle$. Without loss of generality, take $\ell = \mathbf{\ell} = \langle \overline{1}, \overline{2} \rangle$. Again we have already calculated $U_{\mathbf{\ell}} = \langle A(\lambda), \dots, E(\lambda) \mid \lambda \in k \rangle$, which fixes the 2-space $\langle \overline{1}, \overline{2} \rangle$ pointwise and so $\dim(\mathcal{F}_{\overline{C}})_\ell = 2$. We define \underline{C}_ℓ as the subspace of \underline{C}_0 spanned by not only the sheaf term at ℓ , but also the sheaf terms corresponding to all $q + 1$ points $p \in \ell$. Thus

$$\underline{C}_\ell := \langle (\mathcal{F}_{\overline{C}})_\ell, (\mathcal{F}_{\overline{C}})_p \mid p \in \ell \rangle$$

and $\dim \underline{C}_\ell = q + 3$.

Recall the subspaces $\overline{C}_p, \overline{C}_\ell, \overline{D}_p, \overline{E}_\ell, \overline{F}_p \leq \overline{C}$ we introduced in Definition 3.3.30, and

the associated balls in the graph Γ centred at p or ℓ . Let us define equivalent subspaces of \underline{C}_0 and H :

Definition 6.2.1 *Let p be a point and ℓ be a line. Define*

$$\begin{aligned}\underline{D}_p &:= \langle \underline{C}_\ell \mid \ell \ni p \rangle, \\ \underline{E}_\ell &:= \langle \underline{D}_p \mid p \in \ell \rangle, \text{ and} \\ \underline{F}_p &:= \langle \underline{E}_\ell \mid \ell \ni p \rangle.\end{aligned}$$

The images under the natural map $\alpha: \underline{C}_0 \rightarrow H$ of \underline{C}_p , \underline{C}_ℓ , and the spaces defined above are denoted C_p , C_ℓ , D_p , E_ℓ and F_p respectively.

Because of the transitivity of G on the points and lines of Γ , the dimensions of these subspaces are independent of the choice of p or ℓ . Indeed, for any $p \in P$ or $\ell \in L$, and any $g \in G$, we have $X_p g = X_{pg}$ and $Y_\ell g = Y_{\ell g}$ for $X \in \{C, D, F\}$ and $Y \in \{C, E\}$.

Recall that we have the map $\theta: H \rightarrow \overline{C}$ from Theorem 5.3.2. It follows from the definition of θ that $\theta(C_p) = \overline{C}_p$, so we have a sequence of surjective maps

$$\underline{C}_p \xrightarrow{\alpha} C_p \xrightarrow{\theta} \overline{C}_p.$$

We showed above that $\dim \underline{C}_p = \dim \overline{C}_p = 1$ by definition, which therefore forces $\dim C_p = 1$. Thus $(\alpha \circ \theta)|_{\underline{C}_p}$ is an isomorphism.

The situation for C_ℓ is a little more complex. Again from the definition of θ we have that $\theta(C_\ell) = \overline{C}_\ell$, so we have a sequence of surjections

$$\underline{C}_\ell \xrightarrow{\alpha} C_\ell \xrightarrow{\theta} \overline{C}_\ell.$$

We have shown that $\dim \underline{C}_\ell = q + 3$ and $\dim \overline{C}_\ell = 2$, so we have $2 \leq \dim C_\ell \leq q + 3$. For

all $p \in \ell$ we have a 1-simplex (p, ℓ) with boundary map $\delta_{(p,\ell)} = \pm (\varphi_{(p,\ell),p} - \varphi_{(p,\ell),\ell})$ from $(\mathcal{F}_{\overline{C}})_{(p,\ell)}$ into $C_0(\mathcal{F}_{\overline{C}})$. Lemma 5.1.11 tells us that the term $(\mathcal{F}_{\overline{C}})_{(p,\ell)}$ is 1-dimensional, and is therefore isomorphic to the term $(\mathcal{F}_{\overline{C}})_p$ with the map $\varphi_{(p,\ell),p}$ being the identity; the map $\varphi_{(p,\ell),\ell}$ is the embedding of this 1-space into the 2-space $(\mathcal{F}_{\overline{C}})_\ell$. Therefore when we take the quotient $H = C_0(\mathcal{F}_{\overline{C}})/B_0(\mathcal{F}_{\overline{C}})$, each 1-space C_p corresponding to a point $p \in \ell$ is identified with part of the 2-space C_ℓ , and so $\dim C_\ell = 2$.

This leads us to the following lemmas:

Lemma 6.2.2 $H = \langle C_\ell \mid \ell \in L \rangle$.

Proof. By the argument above, each point space C_p is identified with a 1-dimensional subspace of C_ℓ for each $\ell \ni p$. So $H = \langle C_p, C_\ell \mid p \in P, \ell \in L \rangle = \langle C_\ell \mid \ell \in L \rangle$. \square

Lemma 6.2.3 $H = \langle C_p \mid p \in P \rangle$.

Proof. For any $\ell \in L$, we have $C_\ell = \langle C_p \mid p \in \ell \rangle$. So $\langle C_p \mid p \in P \rangle = \langle C_\ell \mid \ell \in L \rangle = H$ by Lemma 6.2.2. \square

The spaces \overline{D}_p , \overline{E}_ℓ and \overline{F}_p are spanned by copies of \overline{C}_x corresponding to points and lines in a ball centred at p or ℓ , so the above arguments also imply that $\theta(D_p) = \overline{D}_p$, $\theta(E_\ell) = \overline{E}_\ell$ and $\theta(F_p) = \overline{F}_p$.

Calculating the dimension of D_p is rather more difficult. We have $\dim \overline{D}_p = 3$, so this is a lower bound on $\dim D_p$. By the argument above, we can ignore the dimension contribution of the point spaces C_p , as they are identified with parts of the line spaces C_ℓ . Furthermore, all $q + 1$ of the 2-spaces $\{C_\ell \mid \ell \ni p\}$ intersect in the 1-space C_p , so $\dim D_p \leq 1 + (q + 1) = q + 2$.

Segev and Smith comment without proof that, over a prime field, the dimension of D_p is either 3 or $q + 2$ [26, p.497]. We prove this here.

Lemma 6.2.4 *Suppose that k is a prime field, so $q = \pi$. Then $\dim D_p \in \{3, \pi + 2\}$. Furthermore, if $\dim D_p = \pi + 2$ then D_p/C_p is a non-split extension of the natural 2-dimensional module for $\mathrm{SL}_2(k)$ by a $\pi - 1$ -dimensional module isomorphic to $V((\pi - 2)\lambda)$, where λ is the fundamental weight for $\mathrm{SL}_2(k)$.*

Proof. Since G is transitive on points, it is sufficient to consider the case $p = \mathfrak{p}$. From Lemma 3.3.37 we have that $G_{\mathfrak{p}} = \langle U_{\mathfrak{p}}, L_{\mathfrak{p}} \rangle$ acts on the 2-space $\overline{D_{\mathfrak{p}}}/\overline{C_{\mathfrak{p}}}$ as $\mathrm{GL}_2(k)$, with $U_{\mathfrak{p}}$ fixing every vector in the 2-space and $L_{\mathfrak{p}} = \langle T, r_{\mathfrak{p}}, E(\lambda) \mid \lambda \in k \rangle \cong \mathrm{GL}_2(k)$.

Set $H := \langle E(\lambda), E(\lambda)^{r_{\mathfrak{p}}} \mid \lambda \in k \rangle$. Notice that $S := \langle E(\lambda) \mid \lambda \in k \rangle$ is a Sylow π -subgroup of $L_{\mathfrak{p}}$, and any two distinct Sylow π -subgroups of $\mathrm{GL}_2(k)$ generate $\mathrm{SL}_2(k)$. Therefore $H \cong \mathrm{SL}_2(k)$. Now S is also a Sylow π -subgroup of H , and $N_H(S) = \langle T_2(\mu), E(\lambda) \mid \mu \in k^\times, \lambda \in k \rangle$ which has order $\pi(\pi - 1)$. Thus we may regard $D_{\mathfrak{p}}/C_{\mathfrak{p}}$ as a $k \mathrm{SL}_2(k)$ -module.

Notice that $N_H(S)$ leaves invariant the line $\ell = \langle \overline{1}, \overline{2} \rangle$, so on restriction to $N_H(S)$, the module $D_{\mathfrak{p}}/C_{\mathfrak{p}}$ certainly contains the 1-dimensional $kN_H(S)$ -submodule $W := C_{\ell}/C_{\mathfrak{p}}$. Hence, as a kH -module, $D_{\mathfrak{p}}/C_{\mathfrak{p}}$ is a quotient of the induced module $V := \mathrm{Ind}_{N_H(S)}^H(W)$. This induced module has a 2-dimensional irreducible quotient $\overline{D_{\mathfrak{p}}}/\overline{C_{\mathfrak{p}}}$, so we can apply Lemma 4.2.1 to conclude that V is a non-split extension of M by $\overline{D_{\mathfrak{p}}}/\overline{C_{\mathfrak{p}}}$, where M is the unique irreducible kH -module of dimension $\pi - 1 = q - 1$.

This leaves us with two possibilities: either $\dim D_{\mathfrak{p}}/C_{\mathfrak{p}} = 2$, or $\dim D_{\mathfrak{p}}/C_{\mathfrak{p}} = 2 + (q - 1) = q + 1$. Since $C_{\mathfrak{p}}$ is 1-dimensional, we therefore have $\dim D_{\mathfrak{p}} \in \{3, q + 2\}$ as required. \square

6.3 Segev and Smith's Theorem

The main result of Segev and Smith is as follows [26, Theorem on p.495].

Theorem 6.3.1 *Suppose $\pi > 2$. Let M be a submodule of H , and V the quotient $V = H/M$. If $\dim(D_p + M)/M = 3$, then $\dim V = 7$, and therefore $V \cong \overline{C}$.*

We may take $M = 0$, and so $\dim D_p = 3$ implies that $H \cong \overline{C}$.

Corollary 6.3.2 *Suppose $M \leq H$ such that $\dim(D_p + M)/M = 3$. Then $M = \ker \theta$.*

Proof. Since $\dim(D_p + M)/M = 3$, then by Theorem 6.3.1 we have $H/M \cong \overline{C}$. But by Theorem 5.3.2, H contains a unique maximal submodule K and $H/K \cong \overline{C}$. Thus $M = K = \ker \theta$. \square

For the remainder of the thesis we make the following standing assumption, and attempt to arrive at a contradiction:

Assumption 6.3.3 $H \not\cong \overline{C}$.

This immediately gives us some further information. Let K be the maximal submodule of H .

Corollary 6.3.4 *K is non-zero.*

Proof. The module H maps onto \overline{C} via the map θ , and since $H \not\cong \overline{C}$ by Assumption 6.3.3 there must be a non-zero kernel $\ker \theta \leq H$. Furthermore, $\overline{C} \cong H/\ker \theta$ is irreducible since $\pi > 3$ by Assumption 6.1.1, so $K = \ker \theta$. \square

Corollary 6.3.5 *We have $\dim H = \dim K + 7$.*

Proof. This follows immediately from the fact that $K = \ker \theta$, where $\theta : H \rightarrow \overline{C}$ is surjective, and $\dim \overline{C} = 7$. \square

Corollary 6.3.6 *Suppose k is a prime field, so that $q = \pi$. Then the dimension of D_p is $q + 2$.*

Proof. By Lemma 6.2.4 we have $\dim D_p \in \{3, q + 2\}$ if $q = \pi$. But by Theorem 6.3.1, $\dim D_p = 3$ implies $H \cong \overline{C}$, contradicting Assumption 6.3.3. Therefore $\dim D_p = q + 2$. \square

CHAPTER 7

RESTRICTION TO THE PRIME SUBFIELD

Since Lemma 6.2.4, and therefore Corollary 6.3.6, only hold for prime fields, we will need an alternative method to deal with other finite fields k . In this chapter we develop some theory to achieve this, working towards a proof that if $H \cong \overline{C}$ over a prime field k_0 of order π , then in fact $H \cong \overline{C}$ over any finite field k of order $q = \pi^a$.

Lemma 7.0.1 *Suppose that ℓ_1, ℓ_2, ℓ_3 are three distinct lines of Δ all containing a point p . If $C_{\ell_3} \leq C_{\ell_1} + C_{\ell_2}$, then $\dim D_p = 3$ and $\dim H = 7$.*

Proof. We have $C_{\ell_1} + C_{\ell_2} \leq D_p$ and $\dim(C_{\ell_1} + C_{\ell_2}) = 3$. Hence $C_{\ell_1} + C_{\ell_2} = C_{\ell_1} + C_{\ell_3} = C_{\ell_2} + C_{\ell_3}$. Let ℓ_4 be a line containing p . Then, by Lemma 3.3.37, there exists $x \in G_p$ such that $\ell_1 x = \ell_1$, $\ell_2 x = \ell_2$ and $\ell_3 x = \ell_4$. But then $C_{\ell_1} + C_{\ell_2}$ is x -invariant, and $C_{\ell_3} x = C_{\ell_4}$ from which we deduce that $C_{\ell_4} \leq C_{\ell_1} + C_{\ell_2}$. Therefore $D_p = C_{\ell_1} + C_{\ell_2}$ as claimed. That $\dim H = 7$ now follows from Theorem 6.3.1. \square

If k_0 is a subfield of k , then $G_2(k_0)$ is a subgroup of $G_2(k)$. Furthermore, considered as a simplicial complex, the building Δ^0 of $G_2(k_0)$ is a subcomplex of the building Δ of $G_2(k)$ [1, p.194]. This motivates the following definitions:

Definition 7.0.2 (Subfield building) *Suppose G is a group of Lie type defined over a finite field k with building Δ , and k_0 is a subfield of k . Let $G^0 \leq G$ be the subgroup of*

G defined over k_0 instead of k . The **subfield building** Δ^0 of G^0 is the building of G^0 , considered as a subcomplex of Δ .

Definition 7.0.3 (Restricted sheaf \mathcal{F}^0) Let G be a group of Lie type defined over a finite field k , let Δ be the building of G , and suppose that \mathcal{F} is a sheaf on Δ defined over k . Let k_0 be a subfield of k with corresponding subfield building $\Delta^0 \subseteq \Delta$. We define the **restricted sheaf** \mathcal{F}^0 to be the sheaf on the subfield building Δ^0 for which $\mathcal{F}_\sigma^0 := \mathcal{F}_\sigma$ for all $\sigma \in \Delta^0$, with connecting maps $\varphi_{\sigma\tau}$ inherited from \mathcal{F} .

Notice that, whilst the building Δ^0 is a geometry for a group defined over k_0 , the restricted sheaf \mathcal{F}^0 is defined over k .

Lemma 7.0.4 Assume that Δ^0 is a subfield building of Δ , and let \mathcal{F}^0 be the restricted sheaf on Δ^0 of the Δ -sheaf $\mathcal{F}_{\overline{C}}$. Denote $H = H_0(\Delta, \mathcal{F}_{\overline{C}})$. Then $J := \langle C_p, C_\ell \mid p, \ell \in \Delta^0 \rangle \leq H$ is a quotient of $H_0(\Delta^0, \mathcal{F}^0)$.

Proof. We regard $C_0(\Delta^0, \mathcal{F}^0)$ as a subspace of the chain space $C_0(\Delta, \mathcal{F}_{\overline{C}})$, and $C_1(\Delta^0, \mathcal{F}^0)$ as a subspace of $C_1(\Delta, \mathcal{F}_{\overline{C}})$. Then the boundary map $\delta_1^0 : C_1(\Delta^0, \mathcal{F}^0) \rightarrow C_0(\Delta^0, \mathcal{F}^0)$ is just the restriction to Δ^0 of the boundary map $\delta_1 : C_1(\Delta, \mathcal{F}_{\overline{C}}) \rightarrow C_0(\Delta, \mathcal{F}_{\overline{C}})$, since the connecting maps $\varphi_{\sigma\tau}$ are inherited, and the cotypes of the simplices (corresponding to their designation as a point or a line) are unchanged in the restricted sheaf. Hence $B_0(\Delta^0, \mathcal{F}^0) = \text{Im}(\delta_1^0) \leq \text{Im}(\delta_1) = B_0(\Delta, \mathcal{F}_{\overline{C}})$. Therefore, $B_0(\Delta^0, \mathcal{F}^0) \leq C_0(\Delta^0, \mathcal{F}^0) \cap B_0(\Delta, \mathcal{F}_{\overline{C}})$ and consequently

$$\begin{aligned} J &= (C_0(\Delta^0, \mathcal{F}^0) + B_0(\Delta, \mathcal{F}_{\overline{C}})) / B_0(\Delta, \mathcal{F}_{\overline{C}}) \\ &\cong C_0(\Delta^0, \mathcal{F}^0) / (C_0(\Delta^0, \mathcal{F}^0) \cap B_0(\Delta, \mathcal{F}_{\overline{C}})) \end{aligned}$$

is a quotient of $C_0(\Delta^0, \mathcal{F}^0) / B_0(\Delta^0, \mathcal{F}^0) = H_0(\Delta^0, \mathcal{F}^0)$. □

We can now state and prove the main result of this chapter.

Proposition 7.0.5 *Let Δ^0 be the subfield building in Δ corresponding to the subfield $k_0 = \mathbb{F}_\pi$ of k . Let \mathcal{G}^0 be the fixed-point sheaf of Δ^0 corresponding to the Cayley module \overline{C} defined over k_0 . If $H_0(\Delta^0, \mathcal{G}^0)$ is irreducible of dimension 7, then $H = H_0(\Delta, \mathcal{F}_{\overline{C}}) \cong_k \overline{C}$.*

Proof. Assume that $H_0(\Delta^0, \mathcal{G}^0)$ is irreducible of dimension 7.

Let \mathcal{F}^0 be the restriction of the sheaf $\mathcal{F}_{\overline{C}}$ to the subfield building Δ^0 . By [24, p.152], \mathcal{F}^0 is isomorphic to the sheaf constructed as follows. For each $\sigma \in \Delta^0$ set $\mathcal{F}_\sigma^0 := \mathcal{G}_\sigma^0 \otimes k$. The connecting maps are tensored similarly (see the discussion in [24, p.152]). Then by [24, p.152, Theorem A2 (ii)], we have $H_0(\Delta^0, \mathcal{F}^0) \cong_k H_0(\Delta^0, \mathcal{G}^0) \otimes k$. Hence $\dim H_0(\Delta^0, \mathcal{F}^0) = 7$ and so $H_0(\Delta^0, \mathcal{F}^0)$ is irreducible as $\pi > 2$.

By Lemma 7.0.4, setting $J = \langle C_p, C_\ell \mid p, \ell \in \Delta^0 \rangle \leq H$, we have that J is a quotient of $H_0(\Delta^0, \mathcal{F}^0)$. Thus $\dim J \leq 7$, and since $J \neq 0$ and $H_0(\Delta^0, \mathcal{F}^0)$ is irreducible, we have $\dim J = 7$.

Let $p \in \Delta^0$ and let $\mathcal{L} = \{\ell \mid \ell \in \Delta^0 \text{ with } \ell \ni p\}$ be the set of lines in Δ^0 which are neighbours in Δ^0 of p . Then $|\mathcal{L}| = \pi + 1 \geq 3$ and $\mathcal{D} := \langle C_\ell \mid \ell \in \mathcal{L} \rangle \leq J \cap D_p$ with $\dim \mathcal{D} = 3$ as $\dim J = 7$. In particular, for $\ell_1, \ell_2, \ell_3 \in \mathcal{L}$ distinct lines we have

$$\mathcal{D} = C_{\ell_1} + C_{\ell_2} \geq C_{\ell_3}.$$

Applying Lemma 7.0.1 yields $\dim H = 7$ and concludes the proof. \square

The major consequence of Proposition 7.0.5 is that to prove our main theorem and to prove the Segev-Smith speculation in general, it suffices to prove it for buildings defined over \mathbb{F}_π and for \mathbb{F}_π -sheaves.

We can apply this immediately with some of the results which we have computed:

Theorem 7.0.6 *Suppose that q is a finite field of characteristic 5, 7 or 11. Then $H_0(\mathcal{F}_{\overline{C}}) \cong \overline{C}$.*

Proof. We have computed $\dim H_0(\mathcal{F}_{\overline{C}}) = 7$ in the cases $q = 5, 7, 11$ (see Table 6.1). Then by Proposition 7.0.5, this suffices to show that $H_0(\mathcal{F}_{\overline{C}}) \cong \overline{C}$ for any finite field $k = \mathbb{F}_q$, where $q = \pi^a$ and $\pi \in \{5, 7, 11\}$. \square

CHAPTER 8

A LOWER BOUND ON THE DIMENSION OF H

The main consequence of Proposition 7.0.5 is that, for the proof of our main theorems, it suffices to consider only the prime field case. This assumption allows us to apply certain results from the theory of weights, as any weight we encounter will automatically be π -restricted. Throughout this chapter, therefore, we assume that $q = \pi$ and $k = \mathbb{F}_\pi$.

In this chapter, we intend to recognise a certain submodule of H using the theory of weights. For $G_2(k)$, following the standard notation established in Section 4.3, we have two fundamental weights, λ_1 and λ_2 , and these are fixed respectively by the reflections σ_1 and σ_2 in the Weyl group. For $\pi > 3$ we know that the Cayley module \overline{C} is irreducible of dimension 7. Therefore we can use Table 2 from [15, p.124], which shows the weight multiplicities of each highest weight $kG_2(k)$ -module $V(\lambda)$, in order to deduce that $\overline{C} \cong V(\lambda_1)$.²

In our construction of G , the Weyl group is given by $W = N/T$ where $N = \langle r_{\mathbf{p}}, r_{\mathbf{\ell}}, T \rangle$.

²The table row corresponding to the weight λ_1 is labelled '10'. It shows that $V(\lambda_1)$ contains weights 0 (orbit length 1) with multiplicity 1, and λ_1 (orbit length 6) with multiplicity 1, and thus has dimension 7. Here we have also used Premet's result (Theorem 4.1.10) which allows us to deduce $\dim L(\lambda_1) = 7$. By contrast, the table row corresponding to the weight λ_2 shows that $V(\lambda_2)$ contains weights 0 (with multiplicity 2), λ_1 and λ_2 , and thus has dimension at least 14. Then Premet's result shows that $13 \leq \dim L(\lambda_2) \leq 14$.

The two generators σ_1 and σ_2 thus correspond to the involutions $r_p T$ and $r_\ell T$; the only problem is that we don't yet know which way round they go. That is, we need to orient our choice of 'points' and 'lines' with our simple reflections σ_1 and σ_2 .

This is relatively straightforward. Since $\overline{C} \cong V(\lambda_1) = V(1 \cdot \lambda_1 + 0 \cdot \lambda_2)$, by Smith's Lemma (Lemma 5.1.9) the coefficient of $\mathcal{F}_{\overline{C}}$ at a σ_1 -panel is isomorphic to the $\mathrm{SL}_2(q)$ -module $V(\lambda)$ of dimension 2, and the coefficient at a σ_2 -panel is isomorphic to the $\mathrm{SL}_2(q)$ -module $V(0)$ of dimension 1. (Here, λ is the fundamental dominant weight of $\mathrm{SL}_2(q)$.) Since G_p leaves the 1-space C_p invariant, and G_ℓ acts irreducibly on C_ℓ which has dimension 2, we see that $\langle r_\ell, T \rangle / T$ corresponds to σ_1 , and $\langle r_p, T \rangle / T$ corresponds to σ_2 .

Recall that $\theta : H \rightarrow \overline{C}$ is the natural map from H onto the Cayley module \overline{C} , and $K := \ker \theta$. Let Y be a maximal proper kG -submodule of K , so that K/Y is an irreducible kG -module. We will attempt to find a lower bound on the dimension of H by calculating (partially) the highest weight of the module K/Y . To achieve this we shall identify $C_{K/Y}(U_p)$ as a $kO^{\pi'}(L_p)$ -module, thus finding the contribution of λ_2 in the weight of K/Y . Finally, we find a lower bound on the dimension of such a module.

8.1 Calculating the highest weight of K/Y

Fix a line $\ell \in L$ and a point $p \in \ell$. Consider the restriction of θ to the subspace D_p and denote $K_p := K \cap D_p$.

Lemma 8.1.1 $\dim K_p = \pi - 1$.

Proof. By Lemma 3.3.31, the image $\overline{D_p}$ has dimension 3. But by Corollary 6.3.6 (due to Assumption 6.3.3), we have $\dim D_p = \pi + 2$. Thus $\dim K_p = \pi - 1$. \square

Lemma 8.1.2 *Let $r \ni p$, and $U = \langle U_p, U_r \rangle$ be the unipotent radical of the minimal parabolic fixing the chamber (p, r) . Then $[C_r, U] = [C_r, U_p] = C_p$.*

Proof. Transitivity of G on chambers (p, r) follows as an immediate consequence of Lemma 3.3.12 (arc transitivity), so we may take $p = \mathbf{p}$ and $r = \boldsymbol{\ell}$ here without loss of generality. We can then consider the action of $U_{\boldsymbol{\ell}}$ and $U_{\mathbf{p}}$ on the 2-space $C_{\boldsymbol{\ell}}$. We note that $U_{\boldsymbol{\ell}}$ fixes $C_{\boldsymbol{\ell}}$ pointwise, whereas $U_{\mathbf{p}}$ acts as $\langle (\frac{1}{\alpha} \ 0) \mid \alpha \in k \rangle$. Hence $[C_{\boldsymbol{\ell}}, U] = [C_{\boldsymbol{\ell}}, U_{\mathbf{p}}] = \langle (1, 0) \rangle = C_{\mathbf{p}}$.

Therefore in general we have $[C_r, U] = [C_r, U_p] = C_p$, as required. \square

Lemma 8.1.3 $K_p \leq C_H(U_p)$.

Proof. By Lemma 8.1.2 and Lemma 5.1.8 we have

$$[D_p, U_p] = \left[\sum_{r \ni p} C_r, U_p \right] = \sum_{r \ni p} [C_r, U_p] = \sum_{r \ni p} C_p = C_p.$$

Now, $K_p \leq D_p$ so we have

$$[K_p, U_p] = K_p \cup [D_p, U_p] \leq [D_p, U_p] = C_p.$$

Also we have that D_p is a kG_p -module, so $\theta|_{D_p}$ is a kG_p -module homomorphism. Now, K_p is the kernel of $\theta|_{D_p}$, and so is a kG_p -submodule of D_p . Thus by Lemma 5.1.6 we have $[K_p, U_p] \leq K_p$. Putting these two together, we obtain

$$[K_p, U_p] \leq K_p \cap C_p. \quad (*)$$

Now, both C_p and $\overline{C_p}$ are 1-spaces, so $\theta|_{C_p}$ is a kG_p -module isomorphism. So, if C_p were contained in K_p , then $\dim \overline{K_p} \cap \overline{C_p} \geq 1$ — an impossibility since, by definition, $\dim \overline{K_p} = 0$. Therefore we must have $K_p \cap C_p = \{0\}$ and so by (*) we see that $[K_p, U_p] = \{0\}$. Hence $K_p \leq C_H(U_p)$ as required. \square

Let us investigate the module K_p . Set

$$G_p^* := O^{\pi'}(G_p) = \langle \text{Syl}_{\pi}(G_p) \rangle.$$

To investigate the structure of this group it suffices to consider the case $p = \mathfrak{p}$. Note that one Sylow π -subgroup of $G_{\mathfrak{p}}$ is given by $\langle A(\lambda), \dots, F(\lambda) \mid \lambda \in k \rangle$. Conjugation by $r_{\mathfrak{p}}$ swaps the pair of generators $\{A(\lambda), B(\lambda)\}$ as well as the pair $\{D(\lambda), F(\lambda)\}$, maps the generator $C(\lambda)$ to $C(-\lambda)$ and maps the generator $E(\lambda)$ to its transpose. So $G_{\mathfrak{p}}^* = \langle A(\lambda), \dots, F(\lambda), E(\lambda)^{r_{\mathfrak{p}}} \mid \lambda \in k \rangle$.

Then $O_{\pi}(G_{\mathfrak{p}}^*) = U_{\mathfrak{p}} = \langle A(\lambda), \dots, D(\lambda), F(\lambda) \mid \lambda \in k \rangle$, and so

$$G_{\mathfrak{p}}^*/O_{\pi}(G_{\mathfrak{p}}^*) \cong \langle E(\lambda), E(\lambda)^{r_{\mathfrak{p}}} \mid \lambda \in k \rangle \cong \text{SL}_2(k).$$

Hence $G_p^*/O_{\pi}(G_p^*) \cong \text{SL}_2(k)$ for any point p .

Lemma 8.1.4 *K_p is irreducible as a kG_p -module, and is isomorphic to $V((\pi - 2)\lambda_2)$ as a $kG_p^*/O_{\pi}(G_p^*)$ -module.*

Proof. This follows immediately from Lemma 6.2.4. □

Lemma 8.1.5 $K_p \cap Y = \{0\}$.

Proof. By Lemma 8.1.4, K_p is irreducible as a kG_p -module. Now, $K_p \cap Y$ is G_p -invariant, so either $K_p \cap Y = \{0\}$ (in which case we are done), or $K_p \cap Y = K_p$.

For the sake of a contradiction, let us assume that $K_p \cap Y = K_p$. Then $\dim(D_p + Y)/Y = 3$, since D_p has dimension $q + 2$ and the entire $(q - 1)$ -dimensional kernel K_p is contained in Y . Therefore by Corollary 6.3.2 we have $Y = K$; a contradiction since Y is a maximal proper submodule of K . Thus the assumption that $K_p \cap Y = K_p$ cannot hold, and so $K_p \cap Y = \{0\}$. □

With these lemmas in hand, we arrive at the following result.

Proposition 8.1.6 $(K_p + Y)/Y = C_{K/Y}(U_p)$.

Proof. K_p is fixed pointwise by U_p by Lemma 8.1.3, so $(K_p + Y)/Y \leq C_{K/Y}(U_p)$. Due to our choice of Y as a maximal proper submodule of K , we have that K/Y is an irreducible kG -module. Therefore by Smith's Lemma (Lemma 5.1.9), $C_{K/Y}(U_p)$ is an irreducible kG_p -module.

We cannot have $(K_p + Y)/Y = \{0\}$ because K_p is non-zero by Lemma 8.1.1 and $K_p \cap Y = \{0\}$ by Lemma 8.1.5. Therefore

$$(K_p + Y)/Y = C_{K/Y}(U_p),$$

as required. □

This allows us to partially calculate the weight of K/Y :

Proposition 8.1.7 K/Y is a kG -module with highest weight

$$\psi = a\lambda_1 + (\pi - 2)\lambda_2,$$

for some $a \in \mathbb{Z}^+$.

Proof. Suppose the irreducible kG -module K/Y has highest weight $\psi = a\lambda_1 + b\lambda_2$. In the fixed-point sheaf for K/Y , the sheaf term at a point p is given by $C_{K/Y}(U_p)$. By Smith's Lemma, this is an irreducible kG_p -module with highest weight $\psi|_{T \cap O^{\pi'}(L_p)} = b\lambda_2$.

By Lemma 8.1.6, we have $C_{K/Y}(U_p) = (K_p + Y)/Y$. Then by Lemma 8.1.5 we have $K_p \cap Y = \{0\}$, and so $(K_p + Y)/Y \cong K_p$, which in turn is isomorphic to $V((\pi - 2)\lambda_2)$ as a $kG_p^*/O_\pi(G_p^*) \cong k\mathrm{SL}_2$ -module by Lemma 8.1.4. Hence $b = \pi - 2$, and $\psi = a\lambda_1 + (\pi - 2)\lambda_2$ for some $a \in \mathbb{Z}^+$. □

In the next section, we obtain a lower bound for the dimension of a module with this highest weight.

8.2 A lower bound on the dimension of K/Y

Having found that the highest weight of K/Y is of the form $\psi = a\lambda_1 + (\pi - 2)\lambda_2$, we can now calculate a lower bound on its dimension.

Proposition 8.2.1 $\dim K/Y \geq 9\pi^2 - 36\pi + 40$.

Proof. K/Y is a kG -module with highest weight ψ , and so $\dim K/Y \geq \dim L(\psi)$. Let $\psi_0 := (\pi - 2)\lambda_2$. Then $\psi - \psi_0 = a\lambda_1 = a(2\alpha_1 + \alpha_2)$, a sum of simple roots with non-negative integer coefficients, and so $\psi_0 \prec \psi$. Thus by Lemma 4.1.9, we have that

$$\Pi(V(\psi_0)) \subseteq \Pi(V(\psi)).$$

Since $(\pi - 2)\lambda_2$ is π -restricted, we can use Lemma 4.1.9 and Theorem 4.1.10 to obtain

$$\dim L(\psi) \geq |\Pi(V((\pi - 2)\lambda_2))|.$$

At this point we refer back to Example 4.3.2, in which we calculated that

$$|\Pi(V((\pi - 2)\lambda_2))| = 1 + 6 \left\lfloor \frac{5(\pi - 2)}{2} \right\rfloor + 12 \sum_{i=1}^{\pi-3} \left\lfloor \frac{3i}{2} \right\rfloor = B_{\pi-2}^{(1)}.$$

We can approximate the bound with a polynomial $B_{\pi-2}^{(2)} \leq B_{\pi-2}^{(1)}$. We define

$$\begin{aligned}
B_{\pi-2}^{(2)} &:= 1 + 6 \left(\frac{5(\pi-2) - 1}{2} \right) + 12 \sum_{i=1}^{\pi-3} \left(\frac{3i-1}{2} \right) \\
&= 1 + 6 \left(\frac{5(\pi-2) - 1}{2} \right) + 12 \left(\frac{3}{2} \frac{(\pi-3)^2 + (\pi-3)}{2} - \frac{\pi-3}{2} \right) \\
&= 9\pi^2 - 36\pi + 40 \\
&\leq B_{\pi-2}^{(1)}.
\end{aligned}$$

Thus we obtain

$$\dim K/Y \geq \dim L(\psi) \geq 9\pi^2 - 36\pi + 40,$$

which completes the proof. □

The main importance of this section is the following bound.

Theorem 8.2.2 *Suppose that $q = \pi$ and $K \neq \{0\}$. Then $\dim H \geq 9\pi^2 - 36\pi + 47$.*

Proof. We have

$$\dim H \geq \dim H/K + \dim K/Y = 7 + \dim K/Y = 9\pi^2 - 36\pi + 47.$$

CHAPTER 9

A CUBIC UPPER BOUND ON THE DIMENSION OF H

With Proposition 7.0.5 in mind, in this chapter we continue to restrict our attention to the case $q = \pi$. Assumption 6.1.1 still holds, so we have $\pi > 3$. The aim of this chapter is to prove the following cubic bound on the dimension of H :

$$\dim H \leq q^3 + 2q^2 + q + 2.$$

First we require bounds on the dimensions of some subspaces of H .

9.1 Preliminaries

Recall that $\dim C_p = 1$ and $\dim C_\ell = 2$. We can bound the dimensions of our other defined subspaces as follows. (Note that due to Assumption 6.3.6 we have $\dim D_p = q + 2$ already fixed, but this lemma holds in more generality so here we let $m := \dim D_p$.)

Lemma 9.1.1 *Let p be a point and ℓ be a line, and let C_p , C_ℓ , D_p , E_ℓ and F_p be as defined in Definition 6.2.1. Set $m := \dim D_p$. Then $m \leq q + 2$ and*

(i) $\dim E_\ell \leq m(q + 1) - 2q \leq q^2 + q + 2$.

$$(ii) \dim F_p \leq q^2(m-2) + q(m-2) + m \leq q^3 + q^2 + q + 2.$$

Proof. We have already observed that the dimensions of these spaces are independent of the choice of p or ℓ .

To see that $m \leq q + 2$, we simply note that $D_p = \langle C_r \mid r \ni p \rangle$ and each line $r \ni p$ contributes at most 1 to the dimension of D_p , except for the first one which contributes 2.

(i) The line ℓ is incident to $(q + 1)$ points, which we denote p_1, \dots, p_{q+1} . Firstly note that

$$\dim(D_{p_1} + D_{p_2}) \leq 2m - 2,$$

since each D_p has dimension m and the 2-space C_ℓ is contained in their intersection. Hence each additional D_{p_i} contributes at most an additional $m - 2$ to the total dimension. Thus

$$\begin{aligned} \dim E_\ell &= \dim(D_{p_1} + \dots + D_{p_{q+1}}) \leq m + q(m - 2) \\ &= m(q + 1) - 2q. \end{aligned}$$

Taking $m = q + 2$ yields the bound $\dim E_\ell \leq q^2 + q + 2$.

(ii) Let the $(q + 1)$ lines incident to p be $\ell_1, \dots, \ell_{q+1}$. Firstly we have

$$\dim(E_{\ell_1} + E_{\ell_2}) \leq 2 \cdot (m(q + 1) - 2q) - m,$$

since each copy of E_ℓ has dimension at most $m(q + 1) - 2q$ and their intersection contains D_p which has dimension m . Adding in each additional E_{ℓ_i} increases the

dimension by another $m(q+1) - 2q - m = q(m-2)$, and so we obtain

$$\begin{aligned} \dim(E_{\ell_1} + \cdots + E_{\ell_{q+1}}) &\leq (m(q+1) - 2q) + q(q(m-2)) \\ &= q^2(m-2) + q(m-2) + m. \end{aligned}$$

Taking $m = q + 2$ yields the bound $\dim E_\ell \leq q^3 + q^2 + q + 2$. \square

9.2 Forming a cubic bound

We begin with some lemmas.

Lemma 9.2.1 *Let p and p^\dagger be a pair of opposite points. Then $D_{p^\dagger} \leq \langle F_p, C_{p^\dagger} \rangle$.*

Proof. By Lemma 3.3.18, for every line h incident to p^\dagger we can find a point $r = r(p, h) \in h$ which has distance 4 from p . Letting $p, \ell_1, p_1, \ell_2, r$ be a shortest path from p to r , we see that

$$C_r \leq C_{\ell_2} \leq D_{p_1} \leq E_{\ell_1} \leq F_p.$$

None of the points $r(p, h)$ are equal to p^\dagger since they all have distance 4 from p (whereas $d(p, p^\dagger) = 6$). Therefore, $C_h = \langle C_r, C_{p^\dagger} \rangle \leq \langle F_p, C_{p^\dagger} \rangle$. Since this is true for all choices of $h \ni p^\dagger$, we have $D_{p^\dagger} \leq \langle F_p, C_{p^\dagger} \rangle$ as claimed. \square

Note that $\{r(p, h) \mid h \ni p^\dagger\}$ is the ideal line $I_{p^\dagger}(p)$.

Definition 9.2.2 *For p, p^\dagger opposite points, define*

$$U_{p, p^\dagger} = \langle F_p, D_r \mid r \in I_{p^\dagger}(p) \rangle.$$

Lemma 9.2.3 $\langle F_p, C_{p^\dagger} \rangle \leq U_{p, p^\dagger}$.

Proof. For any $r \in I_{p^\dagger}(p)$ we have $C_{p^\dagger} \leq D_r$, so $\langle F_p, C_{p^\dagger} \rangle \leq U_{p, p^\dagger}$. \square

Lemma 9.2.4 *The dimension of U_{p,p^\dagger} is at most $(q+1)^2(m-2) + m - q \leq q^3 + 2q^2 + q + 2$, where $m = \dim D_p$.*

Proof. A naïve upper bound for $\dim U_{p,p^\dagger}$ would be $\dim F_p + (q+1)m$; one copy of F_p , plus a copy of D_p for each of the $(q+1)$ lines on p^\dagger . We can refine this by considering the intersection $F_p \cap D_r$, where $r = r(p, h)$ for some $h \ni p^\dagger$. Note that $F_p \cap D_r$ contains the 2-space C_f , where f is the unique line incident to r which lies on the geodesic from r to p . Thus the dimension contribution of each D_r is at most $m-2$, and so we can form the bound

$$\dim U_{p,p^\dagger} \leq \dim F_p + (m-2)(q+1).$$

This can be refined slightly further; we have covered the 1-space C_{p^\dagger} a total of $q+1$ times, since $C_{p^\dagger} \leq D_{r_1} \cap D_{r_2}$ for any $r_1, r_2 \in I_{p^\dagger}(p)$. Subtracting the extra q dimensions gives us

$$\dim U_{p,p^\dagger} \leq \dim F_p + (m-2)(q+1) - q.$$

Applying Lemma 9.1.1, we obtain

$$\begin{aligned} \dim U_{p,p^\dagger} &\leq q^2(m-2) + q(m-2) + m + (m-2)(q+1) - q \\ &= (q+1)^2(m-2) + m - q \\ &\leq q^3 + 2q^2 + q + 2, \end{aligned}$$

as required. □

This leads us to the main result of this section.

Proposition 9.2.5 *Let p be a point and p^\dagger be a point opposite to p . Then $H = U_{p,p^\dagger}$. In particular, $\dim H \leq (q+1)^2(m-2) + m - q \leq q^3 + 2q^2 + q + 2$, where $m = \dim D_p$.*

We will first demonstrate the following.

Lemma 9.2.6 $U_{p,p^\dagger} = U_{p^\dagger,p}$.

Proof. We will begin by showing that $U_{p^\dagger,p} \leq U_{p,p^\dagger}$. By definition, $U_{p^\dagger,p} = \langle F_{p^\dagger}, D_r \mid r \in I_p(p^\dagger) \rangle$, so we can achieve this by demonstrating that $F_{p^\dagger} \leq U_{p,p^\dagger}$ and that $\langle D_r \mid r \in I_p(p^\dagger) \rangle \leq U_{p,p^\dagger}$.

Let h be a line through p^\dagger ; it contains a point $r = r(p, h)$ such that $d(p, r) = 4$ due to Lemma 3.3.18. Now, let u be a point on h . There are three cases to consider.

- **Case 1:** $u = p^\dagger$. We have $D_{p^\dagger} \leq U_{p,p^\dagger}$ by Lemma 9.2.1 and Lemma 9.2.3.
- **Case 2:** $u = r$. We have $D_r \leq U_{p,p^\dagger}$ by the definition of U_{p,p^\dagger} .
- **Case 3:** $u \notin \{p^\dagger, r\}$. Then $d(p, u) = 6$ so we have $D_u \leq \langle F_p, C_u \rangle$ by Lemma 9.2.1, applied with u opposite p . But $F_p \leq U_{p,p^\dagger}$ by definition, and $C_u \leq C_h \leq U_{p,p^\dagger}$, so $D_u \leq \langle F_p, C_u \rangle \leq U_{p,p^\dagger}$.

Thus $D_u \leq U_{p,p^\dagger}$ for all $u \in h$, so $E_h \leq U_{p,p^\dagger}$, and since the initial choice of h was arbitrary, $F_{p^\dagger} \leq U_{p,p^\dagger}$.

Now, $U_{p^\dagger,p} = \langle F_{p^\dagger}, D_r \mid r \in I_p(p^\dagger) \rangle$, so it remains to show that $D_r \leq U_{p,p^\dagger}$ for $r \in I_p(p^\dagger)$. Choose such an r . Then r is collinear with p , so $D_r \leq F_p \leq U_{p,p^\dagger}$ as required. Hence $U_{p^\dagger,p} \leq U_{p,p^\dagger}$.

Applying the same argument with the roles of p and p^\dagger reversed then gives us $U_{p^\dagger,p} = U_{p,p^\dagger}$. □

9.2.1 Points opposite both p and p^\dagger : Brouwer's theorem and rogues

By definition, U_{p,p^\dagger} includes all point spaces C_s for points with $d(p, s) \leq 4$, and therefore by Lemma 9.2.6 it also includes all point spaces C_t for points with $d(p^\dagger, t) \leq 4$. The remaining points to be checked are those which are opposite both p and p^\dagger . In order to show that these points are included, we will need the following theorem:

Theorem 9.2.7 (Brouwer) *Let Γ be a generalised hexagon which is not 3-regular. (Note: In our case, where Γ is the generalised hexagon for $G_2(k)$, this forces $q \neq 2$ which is true by Assumption 6.1.1.)*

Let $x \in \Gamma$ be a point. Then the subgraph of Γ induced by the points and lines of distance 6 and 5 respectively from x is connected.

The above theorem is due to Brouwer and was proved in [7, p.54, Theorem 1.1 (i)]. With this in hand, we are able to show the following:

Lemma 9.2.8 *Suppose p^\dagger is opposite p , and that p^\ddagger is another point which is opposite to both p and p^\dagger . Then $C_{p^\ddagger} \leq U_{p,p^\dagger}$.*

Proof. Let $\ell_1, \dots, \ell_{q+1}$ be the lines going through p^\ddagger . For each $1 \leq i \leq q+1$, let $r_i = r(p, \ell_i)$ and $s_i = r(p^\dagger, \ell_i)$. By Lemma 3.3.18, each $C_{r_i} \leq F_p \leq U_{p,p^\dagger}$ and each $C_{s_i} \leq F_{p^\dagger} \leq U_{p,p^\dagger}$. Note that the set $\{r_i \mid 1 \leq i \leq q+1\}$ forms the ideal line $I_{p^\ddagger}(p)$ and the set $\{s_i \mid 1 \leq i \leq q+1\}$ forms the ideal line $I_{p^\ddagger}(p^\dagger)$.

Suppose that $r_i \neq s_i$ for some $1 \leq i \leq q+1$, so that the two ideal lines are different. Then $\langle C_{r_i}, C_{s_i} \rangle = C_{\ell_i}$, which contains C_{p^\ddagger} , so $C_{p^\ddagger} \leq U_{p,p^\dagger}$.

Therefore, let us now suppose that $r_i = s_i$ for all $1 \leq i \leq q+1$. That is, the ideal lines $I_{p^\ddagger}(p)$ and $I_{p^\ddagger}(p^\dagger)$ are equal. If this is true, and $C_{p^\ddagger} \not\leq U_{p,p^\dagger}$, then we call p^\ddagger a *rogue*. We will show that it is impossible for there to be any rogues, and therefore that $C_{p^\ddagger} \leq U_{p,p^\dagger}$ for all p^\ddagger opposite both p and p^\dagger , as required.

This setup is displayed in Figure 9.1, which again displays a portion of the geometry Δ as a collinearity hypergraph (and not the graph Γ). That is, vertices in the figure represent points, and edges represent lines. We will adopt a new convention which we will use going forward, where the end vertices of a line in the figure are coloured black and others are coloured white. Of course, in the geometry (and in the graph Γ) there is no distinction between ‘end’ points and ‘middle’ points of a line, but drawing the figure




Figure 9.1: A pair of opposite points p and p^\dagger , with p^\ddagger a rogue opposite to both such that the ideal line $I_{p^\ddagger}(p)$ is equal to the ideal line $I_{p^\ddagger}(p^\dagger)$ (here shown by the circle marked I).

in this manner clearly differentiates the lines from one another in cases where the figure might otherwise be ambiguous. (Note that Figure 9.1 displays at most two points per line, and so the use of white dots is not required until Figure 9.2.)

Let \mathcal{O}_p denote the subgraph of Γ induced by all the points opposite p , as well as all the lines at distance 5 from p . By Theorem 9.2.7, \mathcal{O}_p is a connected graph. Assume for a contradiction that there are rogues. Since \mathcal{O}_p contains all points opposite p , it contains all the rogues as well as the point p^\dagger . Let p_0 be a rogue which is closest to p^\dagger in the graph \mathcal{O}_p , and let γ be a shortest path from p_0 to p^\dagger in \mathcal{O}_p . (Note that since p_0 is a rogue, we have $d(p_0, p^\dagger) = 6$, so $|\gamma| \geq 6$). Let p_1 be the point immediately after p_0 along the path γ .

This situation is shown in Figure 9.2; here p_0 is a rogue, with $I := I_{p_0}(p) = I_{p_0}(p^\dagger)$. The point p_1 is collinear with p_0 but cannot be contained in the ideal line I , since $p_1 \in \mathcal{O}_p$ and all such points are opposite p , whereas points in I are at distance 4 from p .

We reach a contradiction by showing that p_1 is another rogue which is closer to p^\dagger in \mathcal{O}_p . Firstly, notice that p_1 is opposite p and p^\dagger , because the only points collinear with p_0 at distance 4 from p or p^\dagger are those contained in I . Let y be the unique point of I which is collinear with p_1 . We cannot have $C_{p_1} \leq U_{p,p^\dagger}$, because we know that $C_y \leq U_{p,p^\dagger}$ and




Figure 9.2: We take p_0 to be the closest rogue to p^\dagger in \mathcal{O}_p — but show that p_1 is another rogue which is closer.

$C_{p_0} \leq \langle C_{p_1}, C_y \rangle$ but $C_{p_0} \not\leq U_{p,p^\dagger}$ by assumption. Since $C_{p_1} \not\leq U_{p,p^\dagger}$, the ideal lines $I_{p_1}(p)$ and $I_{p_1}(p^\dagger)$ are equal, and so p_1 is another rogue.

However, this contradicts the choice of p_0 as a rogue of minimal distance from p^\dagger in \mathcal{O}_p . Therefore there cannot exist any rogues and the lemma is proven. \square

So finally we have:

Proof (Proof of Proposition 9.2.5). By definition U_{p,p^\dagger} includes all point spaces for points of distance ≤ 4 from p , and by symmetry (Lemma 9.2.6) all point spaces for points of distance ≤ 4 from p^\dagger . Also, by Lemma 9.2.8, U_{p,p^\dagger} also contains all the point spaces for points opposite both p and p^\dagger . This covers all the point spaces; by Lemma 6.2.3 this covers all of H . Hence $H = U_{p,p^\dagger}$ and so $\dim H \leq (q+1)^2(m-2) + m - q \leq q^3 + 2q^2 + q + 2$ by Lemma 9.1.1. \square

CHAPTER 10

A QUADRATIC UPPER BOUND ON THE DIMENSION OF H

Here we refine our cubic bound to a quadratic one. We will show that, in most cases, we have

$$\dim H \leq 3q^2 + q + 2.$$

In this chapter we again restrict our attention to the case $q = \pi$, which is sufficient to prove the main theorem due to Proposition 7.0.5. Recall also that $q > 3$ by Assumption 6.1.1.

We begin with some notation. Fix a point p , a line $\ell \ni p$ and a line ℓ^* opposite ℓ . By Corollary 3.3.20, there is a unique path of length 5 from p to ℓ^* . Denote the point on this path incident to ℓ^* as p^* , so that $d(p, p^*) = 4$. Indeed, for any point $r \in \ell$, there is a corresponding point $r^* \in \ell^*$ at distance 4 from r . This gives us a bijection $*$ from the $q + 1$ points on ℓ to the $q + 1$ points on ℓ^* , where each pair of points is connected by a path of length 4.

Recall that for points x and y with $d(x, y) = 4$, the point $x * y$ is defined as the unique point in $\Gamma_2(x) \cap \Gamma_2(y)$.

Definition 10.0.1 (Midpoint m_r) *Let $r \in \ell$ and r^* be the point on ℓ^* at distance 4 from r . Define the midpoint $m_r := r * r^* =: m_{r^*}$.*




Figure 10.1: A pair of opposite edges ℓ and ℓ^* with specified points $p \in \ell$ and $p^* \in \ell^*$, and corresponding midpoint m_p .

The set of midpoints $\{m_r \mid r \in \ell\}$ is given by $\Gamma_3(\ell) \cap \Gamma_3(\ell^*)$.

This setup is displayed in Figure 10.1, which again displays a portion of the geometry Δ displayed as a collinearity hypergraph. We continue to use the convention that a line is represented on the figure by a straight line segment containing white vertices sandwiched between a black vertex at each end; all of the vertices represent points and this distinction is purely to remove any ambiguity from the figure.

In this chapter we will prove the following:

Proposition 10.0.2 *Suppose that $q = \pi$, and $3 \nmid q - 1$. Let $p \in \ell$ and ℓ^* be opposite ℓ .*

Define

$$X_0 := \langle E_\ell, E_{\ell^*} \rangle$$

and

$$X := \langle X_0, D_{m_r} \mid r \in \ell \rangle.$$

Then $H = X$. In particular, we have $\dim H \leq 3q^2 + q + 2$.

In the previous chapter, we showed that for any pair of opposite points $\{p, p^\dagger\}$ we have $H = U_{p,p^\dagger} = \langle F_p, D_r \mid r \in I_{p^\dagger}(p) \rangle$. Therefore we may prove Proposition 10.0.2 by showing

that the following two lemmas hold:

Lemma 10.0.3 *Suppose $3 \nmid q - 1$. Then $F_p \leq X$.*

Lemma 10.0.4 *Suppose $3 \nmid q - 1$ and let $p^\dagger \in \ell^*$ be a point opposite p (that is, p^\dagger is any point on ℓ^* other than p^*). Then $\langle D_r \mid r \in I_{p^\dagger}(p) \rangle \leq X$.*

The following sets will help us to keep track of those points and lines for which we have demonstrated that the corresponding subspace C_p or C_ℓ lies in X .

Definition 10.0.5 (P_X and L_X) *Define $P_X := \{p \in P \mid C_p \leq X\}$ and $L_X := \{\ell \in L \mid C_\ell \leq X\}$.*

10.1 Spines, pages and books

We begin with some new notation.

Definition 10.1.1 (Spines, pages and books) *Take a pair of opposite lines ℓ and ℓ^* , as shown in Figure 10.1. We will refer to one of the paths of length 4 from ℓ to ℓ^* as a **spine**. For any pair of opposite lines there are $q + 1$ spines between them; let these be denoted $\{s_0, \dots, s_q\}$. We use $[r, r^*]$ as shorthand for the spine through r and r^* . The midpoint $m_r = r * r^* = r^* * r = m_{r^*}$ lies on the spine $[r, r^*]$.*

*A choice of any two distinct spines between a pair of opposite lines defines a **page** $\{s_i, s_j\}$, which we interpret as containing not only the points and lines of the two spines, but also the points and lines of any (diagonal) paths of length 6 from the point in one corner to the point in the opposite corner — in a sense, these diagonal paths are drawn onto the page.*

*Fix one spine s_i in the diagram. The set of pages $\{\{s_i, s_j\} \mid j \neq i\}$ is called a **book with spine s_i** .*




Figure 10.2: A diagonal path on the page $\{[r, r^*], [w, w^*]\}$.

Note that the two spines bordering a page correspond to a 12-circuit in Γ and therefore form an apartment. Furthermore, the points at opposite corners of a page are also opposite in Γ , so the paths of length 6 from one to the other are in fact shortest paths between the points.

Figure 10.2 shows two spines, $[r, r^*]$ and $[w, w^*]$, and a diagonal path on the page $\{[r, r^*], [w, w^*]\}$ running from r to w^* . We will denote the lines contained in a spine $[x, x^*]$ by h_x (incident to x) and h_{x^*} (incident to x^*).

By definition $F_p = \langle E_f \mid f \ni p \rangle$, so in order to prove Lemma 10.0.3, we must show that X contains E_f for every line $f \ni p$. The case $f = \ell$ is trivial as $E_\ell \leq X$ by definition, so it remains to check the other q lines incident to p . We will deal with the $q - 1$ cases where $f \neq h_p$ first, and then finally we will show that $E_{h_p} \leq X$.

10.2 Diagonals

Choose two points $r, w \in \ell$, with corresponding points $r^*, w^* \in \ell^*$ such that $[r, r^*]$ and $[w, w^*]$ are two spines forming a page. Suppose a is a line incident to r which is not ℓ or h_r . Note that $d(r, w^*) = 6$ so by Corollary 3.3.20 there is a unique shortest path of length 5 from a to w^* . This path may not involve either ℓ^* or h_{w^*} , because this would force a

cycle of length at most 10 (contradicting the girth of Γ being 12), and therefore the final line on the path must be one of the other $q - 1$ lines incident to w^* . See Figure 10.2 for an illustration of such a path, which we refer to as a **diagonal** on the page $\{[r, r^*], [w, w^*]\}$.

Definition 10.2.1 (Diagonal) *A path of length 6 from one corner of a page to an opposite corner is called a **diagonal**.*

Definition 10.2.2 (Central line of a diagonal) *For d a diagonal, denote by d' be the central line of d (the line equidistant from the endpoints of d).*

Lemma 10.2.3 *A page determines $q - 1$ diagonals from each corner to its opposite corner, none of which share any points or lines other than the endpoints.*

Proof. Without loss of generality, assume that the diagonals begin at $r \in \ell$, and a is a line incident to r other than h_r or ℓ . If we make a different choice for the line a , we obtain a new diagonal which does not share any of the points or lines of the first one. To see this, draw a new line $b \ni r$. By the argument above, there is a point $y \in b$ and a line $b^* \ni w^*$ such that $d(y, b^*) = 3$ and either b^* is a new line, or $b^* = a^*$. But if $b^* = a^*$ then we form a cycle of length at most 10, a contradiction since the girth of Γ is 12. Similarly, if we denote the line at distance 2 from both b and b^* by b' , then we cannot have $a' = b'$ or else we would obtain a 6-cycle. This proves the lemma. \square

Excluding the point r itself, there are q other points on the line a in Figure 10.2. We will now demonstrate that there is a bijection between these q points, and the q points of ℓ^* excluding r^* , given by those pairs which are distance 4 apart. (For example, the bijection maps v to w^* in Figure 10.2.)

Lemma 10.2.4 *For any point $r \in \ell$, any line $a \ni r$ (with $a \notin \{\ell, h_r\}$) and any point $v \in a$ (with $v \neq r$), there is a unique spine $[z, z^*]$ such that there is a diagonal on the page $\{[r, r^*], [z, z^*]\}$, beginning with the path r, a, v .*




Figure 10.3: A diagonal on the page $\{[r, r^*], [w, w^*]\}$.

Proof. By Lemma 10.2.3, any page contains $q - 1$ diagonals between each pair of opposite corners, and none of them share any lines. Therefore, *every* line incident to r , other than ℓ and h_r , lies on one of these diagonals.

Consider the book with spine $[r, r^*]$, and fix a line $a \ni r$. By the above argument, there is a diagonal from r starting with the line a on every page of the book. Thus for any spine $[w_1, w_1^*] \neq [r, r^*]$ there exists precisely one point v on a which has distance 4 from w_1^* . It remains to show that no such point v can correspond in this way to two different spines $[w_1, w_1^*]$ and $[w_2, w_2^*]$. But if that were the case then there would be a path γ of length 4 from v to w_1^* , and a path ω of length 4 from v to w_2^* , so the concatenation $\gamma \circ (w_1^*, w_2^*) \circ \omega^{-1}$ forms a closed walk containing a cycle of length at most 10, a contradiction to the girth of Γ being 12. Therefore the pair (a, v) determines a unique $w^* \in \ell^*$ and so also a unique spine $[w, w^*]$. \square

The previous lemma shows that the point v can be considered to be ‘in general position’, in the sense that (with ℓ and ℓ^* already determined), for any choice of $r \in \ell, a \ni r$ with $a \notin \{\ell, h_r\}$ and $v \in a$ with $v \neq r$, we can locate a spine $[w, w^*]$ and draw Figure 10.2.

The significance of adding these diagonals is that they allow us to show that parts of E_a lie in X .

Lemma 10.2.5 *Fix a point $r \in \ell$, a line $a \ni r$ with $a \notin \{\ell, h_r\}$, and a point $v \in a$ such that $v \neq r$. The pair (a, v) determines a unique diagonal from r to some $w^* \in \ell^*$ beginning with r, a, v . Denote the next line in this path as b , and the final line as c . Then $b \in L_X$.*

Proof. Label the vertices along the diagonal as v and u , as depicted in Figure 10.3. Note that $C_a \leq E_\ell \leq X$ and $C_c \leq E_{\ell^*} \leq X$ by definition. This implies that $C_v \leq X$ and $C_u \leq X$, and these together span the 2-space C_b , so $C_b \leq X$ and $b \in L_X$. \square

Since r, a and v were chosen arbitrarily, we obtain the following:

Corollary 10.2.6 *For any diagonal d , the 2-space $C_{d'}$ (corresponding to the central line d' of d) lies in X . In other words, we have $d' \in L_X$.*

We have now shown that $a, b \in L_X$, so the subspaces corresponding to two lines incident to v lie in X . To show that $D_v \leq X$ there are $q - 1$ further lines which we must demonstrate are in L_X . We do this by identifying more connections in Γ .

10.3 Crossbraces

In this section we will find lines in Δ which connect pairs of diagonals. Many of the arguments used involve locating illegal cycles of length shorter than 12. To make this easier, we will use the following notation for a closed walk: $(a_1 \sim_{i_1} a_2 \sim_{i_2} \cdots \sim_{i_{n-1}} a_n)$, with $a_1 = a_n$ and $1 \leq i_j \leq 6$. The walk has length $\sum_{j=1}^{n-1} i_j$, so if some edge appears in the walk only once, then it necessarily contains a cycle of positive length $d \leq \sum_{j=1}^{n-1} i_j$.

Definition 10.3.1 ($\mathcal{D}(x, y^*)$ and $\mathcal{D}_{\mathcal{P}}(x)$) *Let $\mathcal{P} = \{[x, x^*], [y, y^*]\}$ be a page. We define $\mathcal{D}(x, y^*)$ to be the complete set of $q - 1$ diagonals from x to y^* on \mathcal{P} . Similarly, $\mathcal{D}(y, x^*)$ is the complete set of $q - 1$ diagonals from y to x^* on \mathcal{P} .*

An alternative notation is useful when only one of the spines is important: we let

$$\mathcal{D}_{\mathcal{P}}(x) = \mathcal{D}_{\mathcal{P}}(y^*) = \mathcal{D}(x, y^*)$$

and

$$\mathcal{D}_{\mathcal{P}}(x^*) = \mathcal{D}_{\mathcal{P}}(y) = \mathcal{D}(x^*, y).$$

Definition 10.3.2 (Bridged pair of sets of diagonals) *Suppose \mathcal{P} and \mathcal{Q} are two different pages sharing a common spine $[s, s^*]$. Then we refer to the pair $(\mathcal{D}_{\mathcal{P}}(s), \mathcal{D}_{\mathcal{Q}}(s^*))$ as a **bridged pair of sets of diagonals**, with the spine $[s, s^*]$ being the **bridge**.*

We remark that a bridged pair of sets of diagonals share a spine, but not a point. In the next proposition we consider the positions of the central lines of these diagonals relative to one another:

Proposition 10.3.3 *Let $(\mathcal{D}_1, \mathcal{D}_2)$ be a bridged pair of sets of diagonals. For any diagonal $d \in \mathcal{D}_2$, there is precisely one diagonal $c \in \mathcal{D}_1$ such that $d(d', c') = 4$. If b is any of the remaining $q - 2$ diagonals in \mathcal{D}_1 then d' and b' are opposite.*

We will prove this proposition via three lemmas.

Lemma 10.3.4 *Let $[y, y^*]$ and $[z, z^*]$ be two spines forming a page, and d be a diagonal on the page from y^* to z . Let the point on d at distance 2 from y^* be denoted f (see Figure 10.4). Then y is opposite both z^* and f , and the ideal lines $I_y(z^*)$ and $I_y(f)$ are equal.*

Proof. It is clear from Figure 10.4 that $d(y, f) = d(y, z^*) = 6$. Note that both ideal lines $I_y(z^*)$ and $I_y(f)$ are spanned by the points z and m_y . Hence, by Theorem 3.3.41 they are equal. \square

We now begin to look for new lines connecting the diagonals so we can ascertain $d(c', d')$.

Lemma 10.3.5 *Let $(\mathcal{D}_1, \mathcal{D}_2)$ be a bridged pair of sets of diagonals. For any $d \in \mathcal{D}_2$, there must be at least one $c \in \mathcal{D}_1$ such that d' is not opposite c' . In particular, we have $d(c', d') = 4$.*




Figure 10.4: A diagonal in $\mathcal{D}(z, y^*)$.




Figure 10.5: Diagonals on two pages sharing a spine.

Proof. Let $\mathcal{P}_1 = \{[x, x^*], [y, y^*]\}$ and $\mathcal{P}_2 = \{[y, y^*], [z, z^*]\}$. Let $\mathcal{D}_1 = \mathcal{D}_{\mathcal{P}_1}(y)$ and $\mathcal{D}_2 = \mathcal{D}_{\mathcal{P}_2}(y^*)$. Choose one diagonal d from \mathcal{D}_2 , going from y^* to z . Label the points and lines of this diagonal, and of the $q - 1$ diagonals in \mathcal{D}_1 (going from x^* to y) as shown in Figure 10.5. The central line of d , denoted d' , is incident to points e and f . Suppose for a contradiction that the line d' is opposite the central lines of all diagonals in \mathcal{D}_1 .

Note that $I_y(f) = I_y(z^*)$ by Lemma 10.3.4. Next, let $i \in \{1, \dots, q - 1\}$ and note that g_i is opposite z^* (a path of length 6 passes through points h_i and x^*), so g_i cannot be in the ideal line $I_y(z^*)$, and therefore $g_i \notin I_y(f)$.




Figure 10.6: The ideal line $I_y(f)$ is given by $\{s_i \mid i = 1, \dots, q - 1\} \cup \{z, m_y\}$.

There must, however, be a point of $I_y(f)$ on the line $\langle y, g_i \rangle$. Denote this point s_i . It has distance 4 from f , so must be collinear with a point on either $\langle f, y^* \rangle$, $\langle f, e \rangle$, or one of the $q - 1$ other lines incident to f . However, s_i cannot be collinear with a point on either $\langle f, y^* \rangle$ or $\langle f, e \rangle$ as this would form a 10-cycle in either case: $(s_i \sim_4 y^* \sim_2 m_y \sim_2 y \sim_2 s_i)$, or $(s_i \sim_4 e \sim_2 z \sim_2 y \sim_2 s_i)$. Thus we may draw on these new points and lines as in Figure 10.6.

Next, note that the lines $\langle g_i, y \rangle$ and $\langle f, e \rangle$ are opposite (a path of length 6 goes through points y, z and e). By Corollary 3.3.20, there is a unique path of length 5 from g_i to $\langle f, e \rangle$. Let the final point on this path (incident to the line $\langle f, e \rangle$) be denoted t_i . We will show that the line $\langle f, e \rangle$ consists precisely of the points $\{t_i \mid i = 1, \dots, q - 1\} \cup \{f, e\}$.

Firstly suppose that $t_i = t_j$ for some $i \neq j$. Then $I_y(t_i) = \langle g_i, g_j \rangle = I_y(x^*)$. But $z \in I_y(t_i)$ whereas $z \notin I_y(x^*)$, a contradiction. Now suppose that $t_i = f$ for some i . But then $s_i = g_i$, which we have already shown to be impossible. Finally, if $t_i = e$ then this would form a 10-cycle given by $(e \sim_2 z \sim_2 y \sim_2 g_i \sim_4 e)$, a contradiction. Therefore all the t_i are distinct, and together with f and e form the entire line $\langle f, e \rangle$.

Before we can add the paths from each g_j to t_j to our figure, it remains for us to determine whether the line out of each g_j is $\langle h_j, g_j \rangle$ or a different line. However, we have




Figure 10.7: Shortest paths from each g_j to the line $\langle f, e \rangle$.

assumed that $\langle f, e \rangle$ is opposite $\langle h_j, g_j \rangle$, so the line out of each g_j must be a new one otherwise the distance between central lines $\langle f, e \rangle$ and $\langle h_j, g_j \rangle$ would be 4, not 6. Thus we may annotate our diagram further, as in Figure 10.7.

Fix some $1 \leq i \leq q - 1$. Since $\langle h_i, g_i \rangle$ is opposite $\langle f, e \rangle$, there is some point on $\langle f, e \rangle$ at distance 4 from h_i . Firstly note that $h_i \sim_4 f$ is impossible as it creates the 10-cycle $(h_i \sim_2 x^* \sim_2 y^* \sim_2 f \sim_4 h_i)$. Next, suppose that $h_i \sim_4 e$. This means that $I_{x^*}(e) = \langle y^*, h_i \rangle = I_{x^*}(y)$, and so $m_x \in I_{x^*}(e)$. Then we have $e \sim_4 m_x$ which gives a contradiction since $d(e, m_x) = 6$.

The only remaining possibility is that $h_i \sim_4 t_j$ for some j . The intermediate point on this path cannot be on the line $\langle h_i, x^* \rangle$ as this would create a 10-cycle $(x^* \sim_4 t_j \sim_2 f \sim_2 y^* \sim_2 x^*)$. It also cannot be on the line $\langle h_i, g_i \rangle$ as we have assumed that $\langle h_i, g_i \rangle$ is opposite $\langle f, e \rangle$. Therefore it must be on a new line out of h_i .

We have $I_{x^*}(t_j) = \langle h_i, y^* \rangle = I_{x^*}(y)$. But $I_{x^*}(y)$ also contains h_j , so therefore $h_j \in I_{x^*}(t_j)$ and $t_j \sim_4 h_j$ (see Figure 10.8). This forms the illegal 10-cycle $(t_j \sim_4 g_j \sim_2 h_j \sim_4 t_j)$, a contradiction. Hence the line d' cannot be opposite the central lines of all diagonals in \mathcal{D}_1 ; there must be at least one $c \in \mathcal{D}_1$ such that $d(c', d') \leq 4$.

Note finally that $d(c', d') = 2$ is impossible; if c' and d' were collinear (meeting in a




Figure 10.8: A shortest path from h_i to $\langle f, e \rangle$, and the illegal 10-cycle it induces involving h_j, g_j and t_j .

point s , say) then an illegal 10-cycle would be given by $(x^* \sim_2 y^* \sim_2 f \sim_2 s \sim_2 h_i)$, for some $1 \leq i \leq q - 1$. Thus we have $d(c', d') = 4$. \square

Lemma 10.3.6 *Let $(\mathcal{D}_1, \mathcal{D}_2)$ be a bridged pair of sets of diagonals. For any $d \in \mathcal{D}_2$, there is at most one $c \in \mathcal{D}_1$ such that d' is not opposite c' .*

Proof. Let the setup be as in Figure 10.5, and suppose that there are two diagonals in \mathcal{D}_1 such that $\langle f, e \rangle$ is not opposite either of their central lines — without loss of generality, suppose these are $\langle h_1, g_1 \rangle$ and $\langle h_2, g_2 \rangle$. Then there is a line containing a point on both $\langle f, e \rangle$ and $\langle h_1, g_1 \rangle$, and another line containing a point on both $\langle f, e \rangle$ and $\langle h_2, g_2 \rangle$.

These points on $\langle h_i, g_i \rangle$ cannot be h_i or g_i , else we would create a 10-cycle given by $(h_i \sim_2 x^* \sim_2 y^* \sim_2 f \sim_4 h_i)$ or $(g_i \sim_2 y \sim_2 z \sim_2 e \sim_4 h_i)$. For the same reason, the connecting lines may not be incident to e or f . Thus we may add new points to our diagram, labelling them $t_1 \in \langle h_1, g_1 \rangle, t_2 \in \langle h_2, g_2 \rangle$ and $u_1, u_2 \in \langle f, e \rangle$ as per Figure 10.9 where we assume for a contradiction that $u_1 \neq u_2$.

We have $I_{x^*}(u_1) = \langle y^*, h_1 \rangle$ and $I_{x^*}(u_2) = \langle y^*, h_2 \rangle$. But note that $I_{x^*}(y)$ contains h_1, h_2 and y^* , so $I_{x^*}(y) = I_{x^*}(u_1) = I_{x^*}(u_2)$. Therefore $h_2 \in I_{x^*}(u_1)$ so $h_2 \sim_4 u_1$. Now, u_1 cannot be collinear with a point w on $\langle h_2, g_2 \rangle$ otherwise we obtain the 8-cycle




Figure 10.9: For a contradiction, we assume here that $u_1 \neq u_2$.

$(u_1 \sim_2 w \sim_2 t_2 \sim_2 u_2 \sim_2 u_1)$ (an even smaller cycle is formed if $w = t_2$). But then since we have shown that $h_2 \sim_4 u_1$, we obtain the 10-cycle $(u_1 \sim_4 h_2 \sim_2 t_2 \sim_2 u_2 \sim_2 u_1)$, a contradiction. Thus we must have $u_1 = u_2$ (see Figure 10.10).

Now, we have $I_y(u_1) = \langle g_1, g_2 \rangle = I_y(x^*)$. But $z \in I_y(u_1)$ whereas $z \notin I_y(x^*)$, a contradiction. So $\langle f, e \rangle$ cannot be opposite more than one of the lines $\langle h_i, g_i \rangle$. \square

Proof (Proof of Proposition 10.3.3). This is proved by a combination of Lemma 10.3.5 and Lemma 10.3.6. \square

This motivates a definition:

Definition 10.3.7 (Crossbrace) *Let $(\mathcal{D}_1, \mathcal{D}_2)$ be a bridged pair of sets of diagonals. A line connecting a point on the central line d'_1 of some $d_1 \in \mathcal{D}_1$ and a point on the central line d'_2 of some $d_2 \in \mathcal{D}_2$ is called a crossbrace.*

We can deduce a little more information about how d'_1 and d'_2 meet the crossbrace c .

Lemma 10.3.8 *A crossbrace c cannot meet the central lines d'_1 and d'_2 at their ‘endpoints’, by which we mean the points on the paths d_1 and d_2 . Furthermore, every crossbrace c is opposite both ℓ and ℓ^* .*




Figure 10.10: We conclude that t_1 and t_2 both must be collinear with u_1 ; but this induces an illegal 10-cycle $(u_1 \sim_4 h_2 \sim_2 t_2 \sim_2 u_2 \sim_2 u_1)$. Therefore, the line $\langle f, e \rangle$ cannot be at distance 4 from more than one line $\langle h_i, g_i \rangle$.

Proof. Let the setup be as above, and let c be a crossbrace which goes through the points $t_1 \in d'_1$ and $t_2 \in d'_1$. If either t_1 or t_2 is an ‘endpoint’, an illegal 10-cycle is induced (see Figure 10.11 for an example) — or even an illegal 8-cycle if both t_1 and t_2 are endpoints near ℓ or both near ℓ^* .

Now, if $d(c, \ell^*) = 2$ then the crossbrace c would share a point with the line ℓ^* , which would create an illegal 8-cycle. Similarly, if $d(c, \ell^*) = 4$ then there is a point on c which is collinear with a point on ℓ^* ; but this creates an illegal 10-cycle as displayed in Figure 10.12. Therefore, we have $d(c, \ell^*) = 6$. A similar argument shows that $d(c, \ell) = 6$. \square

Our interest in crossbraces stems from the following proposition:

Proposition 10.3.9 *Suppose that c is a crossbrace. Then $C_c \leq X$.*

Proof. By definition, a crossbrace c is spanned by two points a, b which lie on diagonals. By Corollary 10.2.6 the diagonals are in L_X , so $a, b \in P_X$. Hence $C_c \leq X$. \square




Figure 10.11: The points t_1 and t_2 cannot be endpoints, otherwise an illegal 10-cycle is induced given by $(t_1 \sim_2 h_2 \sim_2 y^* \sim_2 x^* \sim_2 h_1 \sim_2 t_1)$.




Figure 10.12: We cannot have $d(c, \ell^*) = 4$, as this induces an illegal 10-cycle given by $(t_2 \sim_2 r \sim_2 w^* \sim_2 y^* \sim_2 h_2 \sim_2 t_2)$.

10.4 Locating an interior point with three incident lines contained in L_X

Thus far we have proven a series of results which hold for any pair of opposite lines ℓ and ℓ^* . Since G acts transitively on pairs of opposite lines by Corollary 3.3.13, our choice of ℓ and ℓ^* does not matter. At this point, therefore, it will simplify calculations to fix $\ell = \ell = \langle \bar{1}, \bar{2} \rangle$ and $\ell^* = \langle \bar{7}, \bar{8} \rangle$, and calculate explicitly in \bar{C} . Note that these lines ℓ and ℓ^* belong to the apartment depicted in Figure 3.1.

Each point $p \in \ell$ has a unique corresponding point $p^* \in \ell^*$ at distance 4, creating a spine $[p, p^*]$. We can see from Figure 3.1 that $\langle \bar{1} \rangle$ and $\langle \bar{7} \rangle$ form a spine with midpoint $\langle \bar{3} \rangle$, and $\langle \bar{2} \rangle$ and $\langle \bar{8} \rangle$ form a spine with midpoint $\langle \bar{6} \rangle$. In fact, the remaining spines are given by the pairs

$$(\langle \bar{1} + \lambda \bar{2} \rangle, \langle \bar{7} - \lambda \bar{8} \rangle)$$

for $1 \leq \lambda < k$. (We will prove this in Lemma 10.4.2.)

Consider the pages $\mathcal{P} = \{[\langle \bar{1} \rangle, \langle \bar{7} \rangle], [\langle \bar{1} + \bar{2} \rangle, \langle \bar{7} - \bar{8} \rangle]\}$ and $\mathcal{Q} = \{[\langle \bar{1} + \bar{2} \rangle, \langle \bar{7} - \bar{8} \rangle], [\langle \bar{2} \rangle, \langle \bar{8} \rangle]\}$ which share a common spine. The sets $\mathcal{D}_1 := \mathcal{D}(\langle \bar{1} \rangle, \langle \bar{7} - \bar{8} \rangle)$ and $\mathcal{D}_2 := \mathcal{D}(\langle \bar{1} + \bar{2} \rangle, \langle \bar{8} \rangle)$ are a bridged pair of sets of diagonals. By Proposition 10.3.3, for each choice of diagonal $d \in \mathcal{D}_1$, there is precisely one diagonal $e \in \mathcal{D}_2$ such that $d(d', e') = 4$.

Let us choose the points $\langle \bar{1} + \bar{2} + \bar{3} \rangle$ and $\langle \bar{3} - \bar{4} + \bar{5} + \bar{6} - \bar{7} \rangle$, which form a diagonal $d_1 \in \mathcal{D}_1$ on \mathcal{P} . We verify that the central line of d_1 is at distance 4 from the central line of the diagonal $d_2 \in \mathcal{D}_2$ formed by $\langle \bar{2} - \bar{3} + \bar{4} - \bar{5} - \bar{6} \rangle$ and $\langle \bar{6} + \bar{7} - \bar{8} \rangle$ on \mathcal{Q} . The setup so far is shown in Figure 10.13, with crossbrace c connecting points $x_1 \in d_1'$ points $x_2 \in d_2'$. The point $x_1 \in d_1'$ is given by $\langle \bar{1} + \bar{2} + \bar{4} - \bar{5} - \bar{6} + \bar{7} \rangle$, and the point $x_2 \in d_2'$ is given by $\langle \bar{2} - \bar{3} + \bar{4} - \bar{5} + \bar{7} - \bar{8} \rangle$.

Definition 10.4.1 (Interior point) *Let s be a point at distance 3 from either ℓ or ℓ^* .*




Figure 10.13: A crossbrace c realised in geometry. Here we have $x_1 = \langle \bar{1} + \bar{2} + \bar{4} - \bar{5} - \bar{6} + \bar{7} \rangle$, and $x_2 = \langle \bar{2} - \bar{3} + \bar{4} - \bar{5} + \bar{7} - \bar{8} \rangle$.

We say that s is an interior point (with respect to ℓ and ℓ^*) if does not lie on any line contained in a spine from ℓ to ℓ^* .

The points $\langle \bar{1} + \bar{2} + \bar{3} \rangle$ and $\langle \bar{3} - \bar{4} + \bar{5} + \bar{6} - \bar{7} \rangle$ are examples of interior points in Figure 10.13.

We begin by showing that there exist interior points s_1 and s_2 such that three of the lines incident to each s_i are in L_X . Then we will use a group theoretic argument to show that this is in fact true for *all* interior points, not just s_1 and s_2 .

Firstly, an easy lemma about the spines between ℓ and ℓ^* .

Lemma 10.4.2 *Let $w_1 = \langle \bar{1} + \alpha_1 \bar{2} \rangle$ and $w_2^* = \langle \bar{7} + \alpha_2 \bar{8} \rangle$. Then $[w_1, w_2^*]$ is a spine if and only if $\alpha_1 + \alpha_2 = \pi$ or $\alpha_1 = \alpha_2 = 0$.*

Proof. Since $w_1 \in \ell$ and $w_2^* \in \ell^*$, they form a spine if and only if $d(w_1, w_2^*) = 4$; otherwise, w_1 and w_2^* are opposite. So by Lemma 3.3.26, w_1 and w_2^* form a spine if and only if $b(\bar{1} + \alpha_1 \bar{2}, \bar{7} + \alpha_2 \bar{8}) = 0$. Now,

$$(\bar{1} + \alpha_1 \bar{2}) \cdot (\bar{7} + \alpha_2 \bar{8}) = -\bar{3} - \alpha_2 \bar{4} - \alpha_1 \bar{5} + \alpha_1 \alpha_2 \bar{6}$$

and so, by looking at the coefficients of $\bar{4}$ and $\bar{5}$ in the product, we see that $b(\bar{1} + \alpha_1\bar{2}, \bar{7} + \alpha_2\bar{8}) = 0$ if and only if $\alpha_1 + \alpha_2 = 0 \pmod{\pi}$. Since $0 \leq \alpha_1, \alpha_2 < \pi$, this is equivalent to $\alpha_1 + \alpha_2 = \pi$ or $\alpha_1 = \alpha_2 = 0$. \square

Note that this gives q of the spines between ℓ and ℓ^* . The remaining spine is given by $[\langle\bar{2}\rangle, \langle\bar{8}\rangle]$.

Proposition 10.4.3 *There exist interior points s_1 at distance 3 from the line ℓ and s_2 at distance 3 from ℓ^* such that three of the lines incident to each s_i are in L_X . In particular, L_X contains:*

- a line connecting s_1 with ℓ and a line connecting s_2 with ℓ^* ,
- a line incident to each s_i which is the central line of a diagonal, and
- one further line incident to each s_i .

Proof. To prove this, we will build up a list of points contained in P_X in stages. We begin with the set $P_X^{(1)}$ of points contained in some D_x for $x \in \ell$ or ℓ^* (these points are in P_X by the definition of X). Next, by Corollary 10.2.6 the central line d' of any diagonal d is in L_X , because two of its points are in $P_X^{(1)}$. Thus the next set of points included in P_X are those which lie on the central line of a diagonal but are not contained in $P_X^{(1)}$; we refer the set of these points as $P_X^{(2)}$.

There are more points in P_X to be found on crossbraces. Let $P_X^{(3)}$ be the set of points which lie on a crossbrace but are not in $P_X^{(1)}$ or $P_X^{(2)}$; we will demonstrate that $P_X^{(3)}$ is non-empty. In particular, we consider the points on our crossbrace c (Figure 10.13). Denote $v_1 := \bar{1} + \bar{2} + \bar{4} - \bar{5} - \bar{6} + \bar{7}$ and $v_2 := \bar{2} - \bar{3} + \bar{4} - \bar{5} + \bar{7} - \bar{8}$, so that $x_1 = \langle v_1 \rangle$ and $x_2 = \langle v_2 \rangle$. We know that c is spanned by the two 1-spaces x_1 and x_2 , and since we have stipulated that $q > 3$ (Assumption 6.1.1), c contains at least 5 points.

Let $x_3 := \langle v_1 + v_2 \rangle = \langle \bar{1} + 2 \cdot \bar{2} - \bar{3} + 2 \cdot \bar{4} - 2 \cdot \bar{5} - \bar{6} + 2 \cdot \bar{7} - \bar{8} \rangle \in \langle x_1, x_2 \rangle = c$. It is easy to verify that $(v_1 + v_2) \cdot (v_1 + v_2) = 0$, and so x_3 is a point on c . We will demonstrate that x_3 is a ‘new’ point, *i.e.* it is not in $P_X^{(1)}$ or $P_X^{(2)}$. Note firstly that c is opposite both ℓ and ℓ^* by Lemma 10.3.8, so $x_3 \notin P_X^{(1)}$ since it cannot be in D_x for any $x \in \ell$ or ℓ^* . It remains to show that $x_3 \notin P_X^{(2)}$.

Now, since c is opposite both ℓ and ℓ^* by Lemma 10.3.8, there are unique points $w_1 \in \ell$ and $w_2^* \in \ell^*$ such that $d(x_3, w_1) = 4$ and $d(x_3, w_2^*) = 4$. Since $d(x_3, w_1) = 4$ we have $b(x_3, w_1) = 0$. It is easily verified that $b(x_3, \langle \bar{2} \rangle) \neq 0$, so we can write $w_1 = \langle \bar{1} + \alpha_1 \bar{2} \rangle$. We calculate

$$\begin{aligned} x_3 \cdot w_1 &= -2 \cdot \bar{1} - \bar{2} - 2 \cdot \bar{3} + \bar{4} + \alpha_1 \cdot \bar{1} + 2\alpha_1 \cdot \bar{2} - 2\alpha_1 \cdot \bar{5} - \alpha_1 \cdot \bar{6} \\ &= (\alpha_1 - 2) \cdot \bar{1} + (2\alpha_1 - 1) \cdot \bar{2} - 2 \cdot \bar{3} + \bar{4} - 2\alpha_1 \cdot \bar{5} - \alpha_1 \cdot \bar{6}, \end{aligned}$$

and so $1 - 2\alpha_1 \equiv 0 \pmod{\pi}$. Since $0 \leq \alpha_1 < \pi$, we must have $\alpha_1 = (\pi + 1)/2$ and

$$w_1 = \langle \bar{1} + \frac{\pi + 1}{2} \bar{2} \rangle.$$

Similarly, it is easily verified that $b(x_3, \langle \bar{8} \rangle) \neq 0$, so let $w_2^* = \langle \bar{7} + \alpha_2 \bar{8} \rangle$. Then using the same technique, we have

$$\begin{aligned} x_3 \cdot w_2^* &= \bar{3} - 2 \cdot \bar{4} - 2 \cdot \bar{7} + \bar{8} - \alpha_2 \cdot \bar{5} - 2\alpha_2 \cdot \bar{6} - \alpha_2 \cdot \bar{7} + 2\alpha_2 \cdot \bar{8} \\ &= \bar{3} - 2 \cdot \bar{4} - \alpha_2 \cdot \bar{5} - 2\alpha_2 \cdot \bar{6} - (\alpha_2 + 2) \cdot \bar{7} + (2\alpha_2 + 1) \cdot \bar{8}, \end{aligned}$$

so $\alpha_2 + 2 = 0 \pmod{\pi}$. Therefore $\alpha_2 = \pi - 2$ and $w_2^* = \langle \bar{7} + (\pi - 2)\bar{8} \rangle$.

Now, by Lemma 10.4.2 the points w_1 and w_2^* form a spine if and only if $\alpha_1 + \alpha_2 = \pi$;




Figure 10.14: When $\pi \neq 3$, the points $w_1 \in \ell$ and $w_2^* \in \ell^*$ do not form a spine.

that is, if and only if

$$\frac{\pi + 1}{2} + \pi - 2 = \pi \iff \pi = 3.$$

Therefore, since $\pi > 3$ by Assumption 6.1.1, w_1 and w_2^* do not form a spine.

The situation is shown in Figure 10.14. Suppose for a contradiction that $x_3 \in P_X^{(2)}$, so x_3 is a point on the central line e' of some diagonal e . Then e is a diagonal from w_1 to w_2^* , as these are the unique points on ℓ and ℓ^* at distance 4 from x_3 . This implies that the lines incident to x_3 on the shortest paths from x_3 to w_1 and w_2^* respectively (the lines $\langle x_3, s_1 \rangle$ and $\langle x_3, s_2 \rangle$ on the figure) are in fact both the line e' , and thus the points $\{x_3, s_1, s_2\}$ are all collinear.

We can calculate the points s_1 and s_2 . Set $s_1 = \beta_1 \cdot \bar{1} + \dots + \beta_8 \cdot \bar{8}$. Using that $s_1 \cdot w_1 = 0$, $w_1 \cdot s_1 = 0$, $s_1 \cdot x_3 = 0$ and $x_3 \cdot s_1 = 0$, we can form a set of simultaneous equations to find the β_i . The solution, which is easily verified by multiplication with w_1 and with x_3 , is

$$s_1 = \langle 3 \cdot \bar{1} + 4 \cdot \bar{3} - 2 \cdot \bar{4} + 2 \cdot \bar{5} + \bar{6} \rangle.$$

Similarly, setting $s_2 = \gamma_1 \cdot \bar{1} + \dots + \gamma_8 \cdot \bar{8}$ and multiplying with x_3 and w_2^* to find the γ_i ,




Figure 10.15: We have shown that three of the lines incident to s_2 are in L_X .

we obtain

$$s_2 = \langle \bar{3} - 2 \cdot \bar{4} + 2 \cdot \bar{5} + 4 \cdot \bar{6} - 3 \cdot \bar{8} \rangle.$$

By inspection, we can see that (when $\pi \neq 3$) the vectors s_1, s_2 and x_3 span a 3-space and not a 2-space. Therefore they are not all collinear, a contradiction. Hence $x_3 \notin P_X^{(2)}$ and thus it is a ‘new’ point, not contained in $P_X^{(1)}$ or $P_X^{(2)}$. Thus $P_X^{(3)}$ is non-empty.

Now, by Lemma 10.2.4, the point s_2 determines a unique spine $[z, z^*]$ such that there exists a diagonal d_1 on the page $\{[z, z^*], [w_2, w_2^*]\}$ going from w_2^* via s_2 to z . In fact, $[z, z^*] = [\langle \bar{2} \rangle, \langle \bar{8} \rangle]$ which we can verify by confirming that $b(s_2, \langle \bar{2} \rangle) = 0$. But this means we have shown that three distinct lines incident to the point s_2 are contained in L_X ; the central line d'_1 of the diagonal d_1 , the line $\langle s_2, w_2^* \rangle$ and the line $\langle s_2, x_3 \rangle$ (see Figure 10.15). A similar argument for s_1 , which lies on a diagonal from w_1 to $\langle \bar{7} \rangle$, completes the proof. \square

10.5 Transitivity arguments

Having shown that there exist interior points such that three of their incident lines lie in L_X , we will use some transitivity arguments to prove the following proposition. We continue with $\ell = \ell = \langle \bar{1}, \bar{2} \rangle$ and $\ell^* = \langle \bar{7}, \bar{8} \rangle$.

Proposition 10.5.1 *Suppose that $3 \nmid q - 1$. Then if x is any interior point with respect to lines ℓ and ℓ^* , all of the lines incident to x are in L_X . Hence $D_x \leq X$.*

10.5.1 All interior points have three of their incident lines contained in L_X

We prove Proposition 10.5.1 in two stages; firstly we will prove a transitivity result on the set of all interior points (in relation to ℓ and ℓ^*).

Lemma 10.5.2 *Let x be an interior point and suppose $D_x \leq X$. Then $D_y \leq X$ for any interior point y .*

We begin by constructing the stabiliser $G_{\ell, \ell^*} = G_\ell \cap G_{\ell^*}$.

Lemma 10.5.3 *We have $G_{\ell, \ell^*} = \langle T, r_\ell, F(\lambda) : \lambda \in k \rangle$. G_{ℓ, ℓ^*} is isomorphic to $\mathrm{GL}_2(k)$ and has order $q(q^2 - 1)(q - 1)$. Furthermore,*

- (i) T fixes the points $\langle \bar{1} \rangle, \langle \bar{2} \rangle, \langle \bar{3} \rangle, \langle \bar{6} \rangle, \langle \bar{7} \rangle$ and $\langle \bar{8} \rangle$. $T_1(\lambda)$ maps $\langle \bar{2} + \gamma\bar{3} \rangle$ to $\langle \bar{2} + \gamma\lambda\bar{3} \rangle$.
- (ii) r_ℓ swaps the points $\langle \bar{1} \rangle$ and $\langle \bar{2} \rangle$, the points $\langle \bar{3} \rangle$ and $\langle \bar{6} \rangle$, and the points $\langle \bar{7} \rangle$ and $\langle \bar{8} \rangle$.
- (iii) $F(\lambda)$ maps $\langle \bar{2} \rangle$ to $\langle \bar{2} - \lambda\bar{1} \rangle$.

Proof. We showed that $G_\ell \cap G_{\ell^*} = \langle T, r_\ell, F(\lambda) : \lambda \in k \rangle$ in Chapter 3; in Lemma 3.3.36 we showed that $L_\ell = G_\ell \cap G_{\ell^*} \cong \mathrm{GL}_2(k)$. Finally, (i), (ii) and (iii) are clear by inspection. \square

To this, we add one further generator; the element $w_0 = (r_{\mathbf{p}r_\ell})^3$, which rotates our standard apartment from Figure 3.1 by a half-turn, mapping $\langle \bar{1} \rangle$ to $\langle \bar{8} \rangle$, $\langle \bar{2} \rangle$ to $\langle \bar{7} \rangle$ and so on. It therefore swaps ℓ and ℓ^* . We define the group

$$G_{\{\ell, \ell^*\}} := \langle G_{\ell, \ell^*}, w_0 \rangle,$$

where the set notation in the subscript means that the set of lines $\{\ell, \ell^*\}$ is left invariant, rather than each line being fixed individually. $G_{\{\ell, \ell^*\}}$ has order $2q(q^2 - 1)(q - 1)$.

Lemma 10.5.4 *The group $G_{\{\ell, \ell^*\}}$ leaves X invariant.*

Proof. It is clear that $G_{\{\ell, \ell^*\}}$ leaves $X_0 = \langle E_\ell, E_{\ell^*} \rangle$ invariant, and it also leaves invariant set of midpoints $\{m_r \mid r \in \ell\}$ and hence the subspace $\langle D_{m_r} \mid r \in \ell \rangle$. This proves the lemma. \square

Hence, if $Y \leq X$ is some subspace then $Yg \leq X$ for any $g \in G_{\{\ell, \ell^*\}}$. We will use this fact to our advantage during this section.

Proof (Proof of Lemma 10.5.2). We will prove the lemma by showing the transitivity of $G_{\{\ell, \ell^*\}}$ on the interior points. We achieve this via a calculation of the stabiliser $G_{\{\ell, \ell^*\}, \langle \bar{1} \rangle, \langle \bar{1}, \bar{2} + \bar{3} \rangle, \langle \bar{2} + \bar{3} \rangle}$ (see Figure 10.16).

Firstly, we show that $G_{\{\ell, \ell^*\}}$ is transitive on the set of points of ℓ and ℓ^* . By Lemma 10.5.3 (ii), the element r_ℓ swaps (and negates) the coefficients of $\bar{1}$ and $\bar{2}$, so if $x \in \ell$ is arbitrary, then at least one of $\{x, x \cdot r_\ell\}$ is a point on ℓ for which the coefficient of $\bar{2}$ is non-zero. Write this point as $y = \langle \alpha \bar{1} + \bar{2} \rangle$. Then due to Lemma 10.5.3 (iii) we have $y \cdot F(\lambda) = \langle (\alpha - \lambda) \bar{1} + \bar{2} \rangle$, and so $y \cdot F(\lambda) r_\ell = \langle \bar{1} + (\alpha - \lambda) \bar{2} \rangle$. The set $\{\langle (\alpha - \lambda) \bar{1} + \bar{2} \rangle, \langle \bar{1} + (\alpha - \lambda) \bar{2} \rangle \mid \lambda \in k\}$ contains all points on ℓ , so $G_{\{\ell, \ell^*\}}$ is transitive on these points. Then the element $w_0 \in G_{\{\ell, \ell^*\}}$ provides a bijection between the points of ℓ and ℓ^* , so we have transitivity on the complete set of points on ℓ and ℓ^* .

Now let us calculate the stabiliser $G_{\{\ell, \ell^*\}, \langle \bar{1} \rangle}$. (Note that fixing a point on the line $\langle \bar{1}, \bar{2} \rangle$ also fixes the corresponding spine and therefore a point on the line $\langle \bar{7}, \bar{8} \rangle$ as well.) T fixes the point $\langle \bar{1} \rangle$, as do the elements $F(\lambda)$ for $\lambda \in k$, but the elements r_ℓ and w_0 do not. So $G_{\{\ell, \ell^*\}, \langle \bar{1} \rangle} = \langle T, F(\lambda) : \lambda \in k \rangle$ and has order $q(q - 1)^2$.

Next we show the transitivity of this group on the set of $q - 1$ lines incident to $\langle \bar{1} \rangle$, excluding the line $\ell = \langle \bar{1}, \bar{2} \rangle$ and the line $h_{\langle \bar{1} \rangle} = \langle \bar{1}, \bar{3} \rangle$ on the spine. These lines are




Figure 10.16: We calculate the stabiliser $G_{\{\ell, \ell^*, \langle \bar{1} \rangle, \langle \bar{1}, \bar{2} + \bar{3} \rangle, \langle \bar{2} + \bar{3} \rangle\}}$.

parametrised by $\langle \bar{1}, \gamma \bar{2} + \mu \bar{3} \rangle$ where $\gamma, \mu > 0$. (Note that $\gamma = 0$ gives us the line $h_{\langle \bar{1} \rangle}$ on the spine, and $\mu = 0$ gives us the line ℓ .) We have

$$\langle \bar{1}, \gamma \bar{2} + \mu \bar{3} \rangle \cdot T_1(\gamma \mu^{-1}) = \langle \gamma \mu^{-1} \bar{1}, \gamma \bar{2} + \gamma \bar{3} \rangle = \langle \bar{1}, \bar{2} + \bar{3} \rangle.$$

Hence all of these lines are in a single $G_{\{\ell, \ell^*, \langle \bar{1} \rangle\}}$ orbit as required.

Now we will calculate the stabiliser inside $G_{\{\ell, \ell^*, \langle \bar{1} \rangle\}}$ of a line $f = \langle \bar{1}, \bar{2} + \bar{3} \rangle$ incident to $\langle \bar{1} \rangle$, with $f \neq \ell$ and $f \neq h_{\langle \bar{1} \rangle}$. The element $F(\lambda)$ fixes the vector $\bar{1}$ and maps $\bar{2} + \bar{3} \mapsto \lambda \bar{1} + \bar{2} + \bar{3}$, and therefore fixes f . A general element of T is of the form

$$t = T_1(\lambda)T_2(\mu) = \text{diag}(\lambda, \mu, \lambda \mu^{-1}, 1, 1, \lambda^{-1} \mu, \mu^{-1}, \lambda^{-1}).$$

If t fixes f then it preserves the ratio of the coefficients of $\bar{2}$ and $\bar{3}$, forcing $\mu = \lambda \mu^{-1}$ which is equivalent to $\lambda = \mu^2$. Thus we have

$$G_{\{\ell, \ell^*, \langle \bar{1} \rangle, \langle \bar{1}, \bar{2} + \bar{3} \rangle\}} = \langle T_{\dagger}(\mu), F(\lambda) : \lambda \in k, \mu \in k^{\times} \rangle,$$

where $T_{\dagger}(\mu) = \text{diag}(\mu^2, \mu, \mu, 1, 1, \mu^{-1}, \mu^{-1}, \mu^{-2})$. Then $G_{\{\ell, \ell^*, \langle \bar{1} \rangle, \langle \bar{1}, \bar{2} + \bar{3} \rangle\}}$ has order $q(q-1)$.

The next step is to show transitivity of $G_{\{\ell, \ell^*\}, \langle \bar{1} \rangle, \langle \bar{1}, \bar{2} + \bar{3} \rangle}$ on the q points of the line $\langle \bar{1}, \bar{2} + \bar{3} \rangle$ excluding the point $\langle \bar{1} \rangle$. Here, the elements parametrised by $F(\lambda)$ do the trick; by Lemma 10.5.3 (iii) we have, for $\alpha \neq 0$,

$$\langle \bar{1} + \alpha(\bar{2} + \bar{3}) \rangle \cdot F(\lambda) = \langle (1 - \alpha\lambda)\bar{1} + \alpha(\bar{2} + \bar{3}) \rangle,$$

so only $F(0)$ fixes the point $\langle \bar{1} + \alpha(\bar{2} + \bar{3}) \rangle$. This gives transitivity on these points.

In all, we have shown that the group $G_{\{\ell, \ell^*\}}$ acts transitively on the set of interior points (with respect to the lines ℓ and ℓ^*). So if $D_x \leq X$ for some interior point x , then $D_y \leq X$ for any interior point y . \square

For completeness, the stabiliser $G_{\{\ell, \ell^*\}, \langle \bar{1} \rangle, \langle \bar{1}, \bar{2} + \bar{3} \rangle, \langle \bar{2} + \bar{3} \rangle}$ is given by $\langle T_{\dagger}(\mu) : \mu \in k^\times \rangle$ and has order $q - 1$.

By Proposition 10.4.3, the points s_1 and s_2 both have three of their incident lines in L_X , including a line from s_i to ℓ (or ℓ^*), a central line of a diagonal, and one other line. Thus by the this transitivity result, *every* interior point has three such incident lines in L_X .

10.5.2 All interior points have all their incident lines contained in L_X

Next we will prove a transitivity result on the lines incident to an interior point. For this we need another condition on q ; specifically that $3 \nmid q - 1$.

Lemma 10.5.5 *Let x be an interior point at distance 3 from ℓ (respectively, from ℓ^*). Of the $q + 1$ lines incident to x , one of them contains a point on ℓ (or ℓ^* respectively), and another of them is the central line of some diagonal. If $3 \nmid q - 1$ then the stabiliser $G_{\{\ell, \ell^*\}, x}$ acts transitively on the remaining $q - 1$ lines.*




Figure 10.17: Two of the lines incident to x are distinguished, being the line connecting x to ℓ (labelled y), and the central line of a diagonal (labelled d) respectively. We want to show transitivity on the remaining $q - 1$ lines incident to x .

Proof. Using the transitivity of $G_{\{\ell, \ell^*\}}$ on the interior points, we may set $x = \langle \bar{2} + \bar{3} \rangle$ without loss of generality. Then we have $d(x, \ell) = 3$, with the line $y = \langle \bar{1}, \bar{2} + \bar{3} \rangle$ connecting x with $\langle \bar{1} \rangle \in \ell$.

Due to Lemma 10.2.4, another one of the lines incident to an interior point is distinguished, being a central line of a diagonal. In the case of the point $x = \langle \bar{2} + \bar{3} \rangle$, this is the line $d = \langle \bar{2} + \bar{3}, \bar{6} - \bar{7} \rangle$, which lies on a diagonal from $\langle \bar{1} \rangle$ to $\langle \bar{8} \rangle$ (see Figure 10.17).

We would now like to show transitivity on the remaining $q - 1$ lines incident to x . These lines take the form

$$\langle \bar{2} + \bar{3}, \bar{1} + \gamma \bar{6} - \gamma \bar{7} \rangle$$

for $\gamma \neq 0$. Without loss of generality, take $\gamma = 1$ and consider the line $f = \langle \bar{2} + \bar{3}, \bar{1} + \bar{6} - \bar{7} \rangle$. If there is a non-trivial stabiliser of this line within the group $G_{\{\ell, \ell^*\}, \langle \bar{2} + \bar{3} \rangle} = G_{\{\ell, \ell^*\}, \langle \bar{1} \rangle, \langle \bar{1}, \bar{2} + \bar{3} \rangle, \langle \bar{2} + \bar{3} \rangle} = \langle T_{\dagger}(\mu) : \mu \in k \rangle$, then the action cannot be transitive, since the group has order $q - 1$ and there are $q - 1$ lines to permute.

We note that if $3 \mid q - 1$ then there are 3 solutions in k to the equation $\mu^3 = 1$, and for each of these we get an element $T_{\dagger}(\mu_0) = \text{diag}(\mu_0^2, \mu_0, \mu_0, 1, 1, \mu_0^2, \mu_0^2, \mu_0)$ which fixes f ,

giving a stabiliser of order 3. However, if $3 \nmid q - 1$ there is only one solution in k to the equation $\mu^3 = 1$, giving a stabiliser of order 1. Thus $G_{\{\ell, \ell^*\}, \langle \bar{1} \rangle, \langle \bar{1}, \bar{2} + \bar{3} \rangle, \langle \bar{2} + \bar{3} \rangle, \langle \bar{2} + \bar{3}, \bar{1} + \bar{6} - \bar{7} \rangle}$ has order $\gcd(3, q - 1)$.

Hence if $3 \nmid q - 1$, then $G_{\{\ell, \ell^*\}, \langle \bar{1} \rangle, \langle \bar{1}, \bar{2} + \bar{3} \rangle, \langle \bar{2} + \bar{3} \rangle} = G_{\{\ell, \ell^*\}, \langle \bar{2} + \bar{3} \rangle}$ is transitive on the $q - 1$ lines. This proves the lemma. \square

This leads us to:

Proof (Proof of Proposition 10.5.1). By Proposition 10.4.3, there exists an interior point s_1 such that three of the lines incident to s_1 are in L_X . Furthermore, these lines include the line y connecting s_1 to a point on ℓ , the central line d of a diagonal, and one other line f . By Lemma 10.5.5, $f \in L_X$ implies that $\{j \mid j \ni s_1\} \setminus \{y, d\} \subseteq L_X$. Thus all of the lines incident to s_1 are in L_X , so $D_{s_1} \leq X$. Then by Lemma 10.5.1, $D_x \leq X$ for any interior point x . This proves the proposition. \square

10.6 Points on spines

Recall that our goal is to demonstrate that

$$\begin{aligned} H &= U_{p, p^\dagger} = \langle F_p, D_r \mid r \in I_{p^\dagger}(p) \rangle \\ &\leq \langle E_\ell, E_{\ell^*}, D_{m_r} \mid r \in \ell \rangle \\ &= X. \end{aligned}$$

Keep $\ell = \langle \bar{1}, \bar{2} \rangle$ and $\ell^* = \langle \bar{7}, \bar{8} \rangle$, and set $p = \langle \bar{1} \rangle$ so that $p^* = \langle \bar{7} \rangle$. Recall that the space F_p is defined as $\langle E_f \mid f \ni p \rangle$; we need to show, therefore, that $E_f \leq X$ for all lines $f \ni p$.

By the definition of X we have $E_\ell \leq X$. In the previous section we showed that $D_x \leq X$ for all interior points x , which implies that $E_f \leq X$ whenever $f \ni p$ with $f \notin \{\ell, \langle \bar{1}, \bar{3} \rangle\}$ (see Figure 10.18). It remains, therefore, to show that $E_{\langle \bar{1}, \bar{3} \rangle} \leq X$. To do this, we must show that $D_r \leq X$ for all points $r \in \langle \bar{1}, \bar{3} \rangle$.




Figure 10.18: Interior points include x and y , along with the white dots on the lines $\langle \bar{1}, x \rangle$ and $\langle \bar{1}, y \rangle$. We now need to show that $D_r \leq X$ for all points r on spines, such as the point z on the figure.

Note firstly that $D_{\langle \bar{1} \rangle} \leq E_l \leq X$, and $D_{\langle \bar{3} \rangle} \leq X$ by definition since $\langle \bar{3} \rangle$ is a midpoint. There are $q - 1$ remaining points on $\langle \bar{1}, \bar{3} \rangle$; one such point is marked z on Figure 10.18.

Choose a line h connecting to $\langle \bar{8} \rangle$ to an interior point. The line $\langle \bar{1}, \bar{3} \rangle$ is opposite to h , so there is a bijection between the points on the two lines mapping points at distance 4 to each other. (One point, $w \in h$, is on a diagonal from $\langle \bar{8} \rangle$ to $\langle \bar{1} \rangle$, and therefore is mapped to $\langle \bar{1} \rangle$ under this bijection. The point $\langle \bar{8} \rangle \in h$ is mapped to $\langle \bar{3} \rangle$; the path of length 4 goes via the point $\langle \bar{7} \rangle$.) Let $s \in h$ be such that $d(s, z) = 4$ and let $t = s * z$. (See Figure 10.19.)

Now, since s is an interior point, provided $3 \nmid q - 1$ we have $D_s \leq X$ by Proposition 10.5.1. Hence $t \in P_X$. Also, $C_z \leq D_{\langle \bar{1} \rangle} \leq E_l \leq X$, so $z \in P_X$. Thus the line $\langle t, z \rangle$ is in L_X .

10.6.1 Transitivity on the set of points like z

Having fixed $p = \mathbf{p}$, we only need to show that $D_x \leq X$ for points x on the spine $h_{\mathbf{p}}$. However, we find that it is easy to show that $D_x \leq X$ for *all* spine points x , so we prove this slightly stronger result instead. We follow a similar approach as before, by establishing that $G_{\{\ell, \ell^*\}}$ acts transitively on the points under consideration, such as the




Figure 10.19: The line h is opposite $\langle \bar{1}, \bar{3} \rangle$, and $d(s, z) = 4$.

point z from Figure 10.19.

Lemma 10.6.1 *The group*

$$G_{\{\ell, \ell^*\}} = \langle w_0, T, r_\ell, F(\lambda) \mid \lambda \in k \rangle$$

acts transitively on the set of points on spines which are not midpoints and do not lie on ℓ or ℓ^ .*

Proof. Recall that w_0 rotates the apartment containing ℓ and ℓ^* , so in fact it is enough to show that

$$G_{\ell, \ell^*} = \langle T, r_\ell, F(\lambda) \mid \lambda \in k \rangle$$

acts transitively on the subset of these points lying on a spine line h_r for some $r \in \ell$.

We know that G_{ℓ, ℓ^*} stabilises the set $\{h_r \mid r \in \ell\}$, since it leaves invariant both ℓ and the set of all midpoints between ℓ and ℓ^* . Let $m_p = \langle \bar{3} \rangle$ and $h_p = \langle \bar{1}, \bar{3} \rangle$. We begin by calculating the stabiliser G_{ℓ, ℓ^*, h_p} . This contains both $\langle F(\lambda) \mid \lambda \in k \rangle$ and T , since both fix

the points $\langle \bar{1} \rangle$ and $\langle \bar{3} \rangle$. However, r_ℓ maps $\langle \bar{1}, \bar{3} \rangle$ to $\langle \bar{2}, \bar{6} \rangle$ so $r_\ell \notin G_{\ell, \ell^*, h_p}$. Hence we have

$$G_{\ell, \ell^*, h_p} = \langle T, F(\lambda) \mid \lambda \in k \rangle,$$

which has order $q(q-1)^2$, and so $[G_{\ell, \ell^*} : G_{\ell, \ell^*, h_p}] = q+1$. This is equal to $|\{h_r \mid r \in \ell\}|$, and so G_{ℓ, ℓ^*} is transitive these lines.

Suppose that $z = \langle \bar{1} - \bar{3} \rangle \in h_p$. We calculate the stabiliser $G_{\ell, \ell^*, h_p, z}$. Again this must contain $\langle F(\lambda) \mid \lambda \in k \rangle$ which fixes the vectors $\bar{1}$ and $\bar{3}$, and it also contains $\langle T_1(\lambda) \mid \lambda \in k^\times \rangle$. However, $z \cdot T_2(\lambda) = \langle \bar{1} - \lambda^{-1}\bar{3} \rangle$, so $\langle T_2(\lambda) \mid \lambda \in k^\times \rangle$ intersects $G_{\ell, \ell^*, h_p, z}$ trivially. Hence we have

$$G_{\ell, \ell^*, h_p, z} = \langle T_1(\lambda), F(\mu) \mid \lambda \in k^\times, \mu \in k \rangle$$

with order $q(q-1)$. The index $[G_{\ell, \ell^*, h_p} : G_{\ell, \ell^*, h_p, z}]$ is $q-1$, and so G_{ℓ, ℓ^*, h_p} is transitive on the $q-1$ points $\{t \mid t \in h_p\} \setminus \{p, m_p\}$.

Therefore, $G_{\{\ell, \ell^*\}}$ acts transitively on the spine points as desired. \square

10.6.2 Transitivity on the q non-spine lines through z

With transitivity of $G_{\{\ell, \ell^*\}}$ on the points like z in hand, we may fix $z = \langle \bar{1} - \bar{3} \rangle$ without loss of generality.

Lemma 10.6.2 *Let $z = \langle \bar{1} - \bar{3} \rangle \in h_p$. Then $D_z \leq X$.*

Proof. Since we know that $C_{h_p} \in X$, and (provided $3 \nmid q-1$) we have located one non-spine line ($f = \langle t, z \rangle$ in Figure 10.19) incident to z such that $C_f \leq X$, it is enough to prove that the group $G_{\{\ell, \ell^*\}, z}$ acts transitively on the q non-spine lines through z .

Firstly we note that these lines are given by $\{f_\psi = \langle \bar{1} - \bar{3}, \bar{2} + \psi\bar{3} - \bar{4} + \bar{5} + \bar{7} \rangle \mid \psi \in k\}$. (It is easy to verify by multiplication in C^+ that $\langle \bar{2} + \psi\bar{3} - \bar{4} + \bar{5} + \bar{7} \rangle$ is a point and that it forms a line with $\langle \bar{1} - \bar{3} \rangle$.)

Now, we have

$$G_{\{\ell, \ell^*\}, z} = G_{\ell, \ell^*, h_p, z} = \langle T_1(\lambda), F(\mu) \mid \lambda \in k^\times, \mu \in k \rangle.$$

We calculate

$$\begin{aligned} \langle \bar{2} + \psi \bar{3} - \bar{4} + \bar{5} + \bar{7} \rangle \cdot F(\mu) &= \langle -\lambda \bar{1} + \bar{2} + (\psi - 2\lambda) \bar{3} - \bar{4} + \bar{5} + \bar{7} \rangle \\ &= \langle \bar{2} + (\psi - 3\lambda) \bar{3} - \bar{4} + \bar{5} + \bar{7} \rangle - \lambda \langle \bar{1} - \bar{3} \rangle, \end{aligned}$$

and so $f_\psi \cdot F(\lambda) = f_{\psi-3\lambda}$. Hence to map f_ψ to any f_φ , we just need to multiply by $F(\lambda)$ where $\lambda = 3^{-1}(\psi - \varphi)$. This proves the transitivity result. \square

Bringing this all together, we have the following proposition:

Proposition 10.6.3 *Suppose $3 \nmid q - 1$ and let x be a point on a spine between ℓ and ℓ^* . Then $D_x \leq X$.*

Proof. If $x \in \ell$, $x \in \ell^*$ or $x = m_r$ for some $r \in \ell$ then $D_x \leq X$ by definition. Lemma 10.6.2 and Lemma 10.6.1 prove the proposition in all other cases. \square

10.7 Proof that $H = X$

We can now prove the two lemmas which together give us Proposition 10.0.2.

Proof (Proof of Lemma 10.0.3). By definition we have $F_p = \langle E_f \mid f \ni p \rangle$, so we must show that $E_f \leq X$ for every line $f \ni p$. For $f = \ell$ this follows immediately from the definition of X . For $f = h_p$, we must show that $\langle D_x \mid x \in h_p \rangle = E_{h_p} \leq X$; this follows from Proposition 10.6.3. Finally, suppose f is any of the remaining $q - 1$ lines. Then f contains the point $p \in \ell$ and q other points, all of which are interior points. By Proposition 10.5.1, if x is one of these interior points then $D_x \leq X$, and $D_p \leq X$ by definition. Hence $E_f \leq X$. Therefore $F_p \leq X$. \square

Proof (Proof of Lemma 10.0.4). We must show that $D_x \leq X$ for all $x \in I_{p^\dagger}(p)$, where $p^\dagger \in \ell^*$ with $d(p, p^\dagger) = 6$ (in other words, with $p^\dagger \neq p^*$). The ideal line $I_{p^\dagger}(p)$ contains the points p^* and m_{p^\dagger} , as well as $q - 1$ interior points which lie on the diagonals from p^\dagger to p . We have $D_{p^*} \leq E_{\ell^*} \leq X$ and $D_{m_{p^\dagger}} \leq X$ by definition, and Proposition 10.5.1 tells us that $D_x \leq X$ for any interior point x , so we are done. \square

Proof (Proof of Proposition 10.0.2). Lemma 10.0.3 and Lemma 10.0.4 together prove that $H = X$.

We obtain the dimension bound as follows. We have $\dim E_\ell \leq q^2 + q + 2$ from Lemma 9.1.1, and we know that $\dim D_p \leq q + 2$. Now, E_ℓ and E_{ℓ^*} intersect in a 1-space at each of the $q + 1$ midpoints, so

$$\dim X_0 = \dim \langle E_\ell, E_{\ell^*} \rangle \leq 2(q^2 + q + 2) - (q + 1) = 2q^2 + q + 3.$$

Each D_{m_r} overlaps X_0 in the lines h_r and h_{r^*} , and therefore contributes at most $q - 1$ to the dimension of X . Thus we have

$$\dim X \leq 2q^2 + q + 3 + (q + 1)(q - 1) = 3q^2 + q + 2,$$

as required. \square

10.7.1 A note about the condition $3 \nmid q - 1$

By computer calculation (see Table 6.1) we have already demonstrated that $\dim H = 7$ over the field \mathbb{F}_7 . This is generalised further by Proposition 7.0.5 to fields of order $q = 7^a$. This provides some evidence that the condition $3 \nmid q - 1$ is not actually required, since $3 \mid 7^a - 1$ for all $a \geq 1$.

In the same way that identifying the presence of ‘crossbraces’ allowed us to demonstrate that new points and lines were contained in P_X and L_X , it seems likely that finding

more connections in Γ might be enough to demonstrate that $D_x \leq X$ for all interior points x , even in the case that $3 \mid q - 1$. For example, in the case $q = 7$, the $q + 1$ lines incident to an interior point x are composed of the following $G_{\{\ell, \ell^*\}, x}$ orbits:

- 1 line connecting x to either ℓ or ℓ^* ,
- 1 line which is the central line of a diagonal,
- 2 lines which each contain a point lying on a crossbrace,
- 2 more lines which each contain a point lying on a crossbrace, and
- 2 lines which do not contain any points lying on a crossbrace.

(This contrasts with the case $3 \nmid q - 1$, where Lemma 10.5.5 shows that all lines incident to x , except the first two orbits in the list above, lie in one single orbit of length $q - 1$; and furthermore that all of these lines contain a point lying on a crossbrace.)

A line containing a point z lying on a crossbrace is automatically in L_X , as it is spanned by x and z , both of which are in P_X . Therefore it is only the last orbit in the list above which poses a potential problem. However, for $q = 7$ at least, the lines in the last orbit do each contain a point z on a line f , which in turn contains two points z_1, z_2 lying on crossbraces. Thus $z_1, z_2 \in P_X$, so $f \in L_X$ and $z \in P_X$, and so the lines in the last orbit are also in L_X .

It seems likely than an approach like this will generalise to any finite field, but this has not yet been shown.

CHAPTER 11

PROOF OF THE MAIN THEOREM

We are now able to bring all of the results together in our main theorem, which we restate here.

Theorem A *Let $k = \mathbb{F}_q$, where $q = \pi^a$ (π a prime) with $\pi > 3$ and $3 \nmid \pi - 1$. Set $G = G_2(k)$, and let Δ be the building of G . Denote by \overline{C} the Cayley module for G , and let $\mathcal{F}_{\overline{C}}$ be the corresponding fixed-point sheaf on Δ . Then $H_0(\Delta, \mathcal{F}_{\overline{C}}) \cong \overline{C}$ and so $\dim H_0(\Delta, \mathcal{F}_{\overline{C}}) = 7$.*

Proof. By Proposition 7.0.5, it suffices to consider the case where $q = \pi$. Let $H = H_0(\mathcal{F}_{\overline{C}})$. If K , the kernel of the natural map $\theta: H \rightarrow \overline{C}$, is non-zero, then by Theorem 8.2.2 we have $\dim H \geq 9q^2 - 36q + 47$. But by Proposition 10.0.2 we have $\dim H \leq 3q^2 + q + 2$. This is only possible if $9q^2 - 36q + 47 \leq 3q^2 + q + 2$, which is true if and only if $5/3 \leq q \leq 9/2$. The possible values for q are therefore 2, 3 or 4, all of which are ruled out by our assumption that $\pi > 3$.

Hence $K = 0$, and so $H \cong \overline{C}$. Thus $\dim H = 7$ as required. □

As per the discussion in Section 10.7.1, we suspect that in fact the result holds without the need for the condition $3 \nmid \pi - 1$; the case $q = 7$ has been computed using MAGMA, and here we get $\dim H = 7$ as well, which provides some evidence for this speculation.

The case $q = 13$, which would provide further evidence, is currently just out of our reach computationally. We plan to optimise our MAGMA programs in the hope that the case $q = 13$ becomes viable.

APPENDIX A

MAGMA PROGRAMS

The MAGMA programs written during the completion of this thesis are available at the following GitHub repository:

<https://github.com/mw-butler/sheaves-on-buildings>

LIST OF REFERENCES

- [1] Peter Abramenko and Kenneth S. Brown. *Buildings: Theory and Applications*. Graduate Texts in Mathematics. Springer New York, 2008.
- [2] Jonathan L. Alperin. *Local Representation Theory: Modular Representations as an Introduction to the Local Representation Theory of Finite Groups*. Cambridge Studies in Advanced Mathematics. Cambridge University Press, 1986.
- [3] David J. Benson. *Representations and Cohomology: Volume 1, Basic Representation Theory of Finite Groups and Associative Algebras*. Cambridge Studies in Advanced Mathematics. Cambridge University Press, 1998.
- [4] Norman L. Biggs and Arthur T. White. *Permutation Groups and Combinatorial Structures*. Cambridge University Press London ; New York ; Melbourne, 1979.
- [5] Wieb Bosma, John Cannon, and Catherine Playoust. The MAGMA algebra system. *J. Symbolic Comput.*, 24(3-4):235–265, 1997. Computational algebra and number theory (London, 1993).
- [6] Nicolas Bourbaki. *Lie Groups and Lie Algebras: Chapters 4-6*. Number 4-6 in Elements de mathematique. Springer Berlin Heidelberg, 2008.
- [7] Andries E. Brouwer. *The complement of a geometric hyperplane in a generalized polygon is usually connected*, page 53–58. London Mathematical Society Lecture Note Series. Cambridge University Press, 1993.
- [8] Kenneth S. Brown. What is... a building? *Notices of the AMS*, pages 1244–1245, 2002.

- [9] Steve M. Cohen and Stephen D. Smith. A geometric presentation of the 26-dimensional module for $F_4(q)$. *Journal of Algebra*, 132(2):474–487, 1990.
- [10] David Craven. *Representation Theory of Finite Groups: a Guidebook*. Universitext. Springer Nature, United States, September 2019.
- [11] Walter Feit and Graham Higman. The nonexistence of certain generalized polygons. *Journal of Algebra*, 1(2):114–131, 1964.
- [12] Tony J Fisher. A bound for the dimension of $H_0(\mathcal{F}_V)$ of modules $V(\lambda_i + \lambda_j)$ for $A_\ell(q)$. *Journal of Algebra*, 143(1):119–129, 1991.
- [13] Paul B. Garrett. *Buildings and Classical Groups*. Taylor & Francis, 1997.
- [14] James E. Humphreys. *Linear Algebraic Groups*. Graduate Texts in Mathematics. Springer, 1975.
- [15] James E. Humphreys. *Introduction to Lie Algebras and Representation Theory*. Graduate Texts in Mathematics. Springer New York, 1994.
- [16] James E. Humphreys. *Modular Representations of Finite Groups of Lie Type*. Cambridge University Press, 2006.
- [17] Frank Lübeck. Small degree representations of finite chevalley groups in defining characteristic. *LMS Journal of Computation and Mathematics*, 4:135–169, 2001.
- [18] Ulrich Meierfrankenfeld and Bernd Stellmacher. The general FF-module theorem. *Journal of Algebra*, 351(1):1–63, 2012.
- [19] Alexander A. Premet. Weights of infinitesimally irreducible representations of Chevalley groups over a field of prime characteristic. *Sbornik: Mathematics*, 61(1):167–183, February 1988.
- [20] Mark A. Ronan. Embeddings and hyperplanes of discrete geometries. *European Journal of Combinatorics*, 8(2):179–185, 1987.

- [21] Mark A. Ronan. Duality for presheaves on chamber systems, and a related chain complex. *Journal of Algebra*, 121(2):263–274, 1989.
- [22] Mark A. Ronan. *Lectures on Buildings*. Academic Press, 1989.
- [23] Mark A. Ronan and Stephen D. Smith. Sheaves on buildings and modular representations of Chevalley groups. *Journal of Algebra*, 96(2):319–346, 1985.
- [24] Mark A. Ronan and Stephen D. Smith. Universal presheaves on group geometries, and modular representations. *Journal of Algebra*, 102:135–154, 1986.
- [25] Mark A. Ronan and Stephen D. Smith. Computation of 2-modular sheaves and representations for $L_4(2)$, A_7 , $3S_6$, and M_{24} . *Communications in Algebra*, 17(5):1199–1237, 1989.
- [26] Yoav Segev and Stephen D. Smith. Apartments and the Cayley-algebra module for $G_2(k)$. *Communications in Algebra*, 14(3):493–506, 1986.
- [27] Ernest E. Shult. *Points and Lines: Characterizing the Classical Geometries*. Universitext. Springer Berlin Heidelberg, 2010.
- [28] Stephen D. Smith. Irreducible modules and parabolic subgroups. *Journal of Algebra*, 75(1):286–289, 1982.
- [29] Stephen D. Smith and Helmut Völklein. A geometric presentation for the adjoint module of $SL_3(k)$. *Journal of Algebra*, 127:127–138, 1987.
- [30] Robert Steinberg. *Lectures on Chevalley groups*. University Lecture Series. American Mathematical Society, 2016.
- [31] Franz G. Timmesfeld. A remark on irreducible modules for finite Lie-type groups. *Archiv der Mathematik*, 46:499–500, 1986.
- [32] Jacques Tits. *Buildings of Spherical Type and Finite BN-pairs*. Lecture notes in mathematics. Springer-Verlag, 1974.

- [33] Jacques Tits and Richard M. Weiss. *Moufang Polygons*, pages 19–22. Springer Berlin Heidelberg, Berlin, Heidelberg, 2002.

- [34] David Ward. *Topics in Finite Groups: Homology Groups, π -Product Graphs, Wreath Products and Cuspidal Characters*. PhD thesis, School of Mathematics, The University of Manchester, 2015.

- [35] Richard M. Weiss. *The Structure of Spherical Buildings*. Princeton University Press, 2003.

- [36] Robert A. Wilson. *The Finite Simple Groups*. Graduate Texts in Mathematics. Springer London, 2009.