7

Axioms for polar spaces

The axiomatisation of polar spaces was begun by Veldkamp, completed by Tits, and simplified by Buekenhout, Shult, Hanssens, and others. In this chapter, the analogue of Chapter 3, these results are discussed, and proofs given in some cases as illustrations. We begin with a discussion of generalised quadrangles, which play a similar rôle here to that of projective planes in the theory of projective spaces.

7.1 Generalised quadrangles

We saw the definition of a generalised quadrangle in Section 6.4: it is a rank 2 geometry satisfying the conditions

(Q1) two points lie on at most one line;

(Q2) if the point p is not on the line L, then p is collinear with exactly one point of L;

(Q3) no point is collinear with all others.

For later use, we represent generalised quadrangles by a diagram with a double arc, thus:

.

The axioms (Q1)–(Q3) are self-dual; so the dual of a generalised quadrangle is also a generalised quadrangle.

Two simple classes of examples are provided by the *complete bipartite graphs*, whose points fall into two disjoint sets (with at least two points in each, and whose lines consist of all pairs of points containing one from each set), and their duals, the *grids*, some of which we met in Section 6.4. Any generalised quadrangle in which lines have just two points is a complete bipartite graph, and dually (Exercise 2). We note that any line contains at least two points, and dually: if L were a singleton line $\{p\}$, then every other point would be collinear with p (by (Q2)), contradicting (Q3).

Apart from complete bipartite graphs and grids, all generalised quadrangles have orders:

Theorem 7.1 *Let G be a generalised quadrangle in which there is a line with at least three points and a point on at least three lines. Then the number of points on a line, and the number of lines through a point, are constants.*

Proof First observe that, if lines L_1 and L_2 are disjoint, then they have the same cardinality; for collinearity sets up a bijection between the points on L_1 and those on L_2.

Now suppose that L_1 and L_2 intersect. Let p be a point on neither of these lines. Then one line through p meets L_1, and one meets L_2, so there is a line L_3 disjoint from both L_1 and L_2. It follows that L_1 and L_2 both have the same cardinality as L_3.

The other assertion is proved dually. ■

This proof works in both the finite and the infinite case. If G is finite, we let s and t be the orders of G; that is, any line has $s+1$ points and any point lies on $t+1$ lines, so that the diagram is

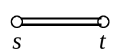

For the classical polar spaces over GF(q), we have $s = q$ and $t = q^{1+\varepsilon}$, where ε is given in Table 6.5.1.

From now on, "generalised quadrangle" will be abbreviated to GQ.

The next result summarises some properties of finite GQs.

Theorem 7.2 *Let G be a finite GQ with orders s,t.*

(a) G has $(s+1)(st+1)$ points and $(t+1)(st+1)$ lines.

(b) $s+t$ divides $st(s+1)(t+1)$;

(c) if $s>1$, then $t \le s^2$;

(d) if $t > 1$, then $s \le t^2$.

Proof (a) is proved by elementary counting, like that in Section 6.5. (b) is shown by an argument involving eigenvalues of matrices, in the spirit of the proof of the Friendship Theorem outlined in Exercise 2.2.4. (c) is proved by elementary counting (see Exercise 3), and (d) is dual to (c). ■

In particular, if $s = 2$, then $t \le 4$; and the case $t = 3$ is excluded by (b) above. So $t = 1, 2$ or 4. These three values are realised by the three orthogonal rank 2 polar spaces over $\mathrm{GF}(2)$. We will see, as a special case of a later result, that these are the only GQs with $s = 2$. However, this result is sufficiently interesting to be worth another proof which generalises it in a different direction.

Theorem 7.3 *Let G be a GQ with orders $s = 2$ and t. Then $t = 1, 2$ or 4; and there is a unique geometry for each value of t.*

Note the generalisation: t is not assumed to be finite!

Proof Take a point and call it ∞; let $\{L_i : i \in I\}$ be the set of lines containing ∞. Number the points other than ∞ on L_i as p_{i0} and p_{i1}. Now, for any point q not collinear with p, there is a function $f_q : I \to \{0, 1\}$ defined by the rule that the unique point of L_i collinear with q is $p_{i f_q(i)}$. We use the function f_q as a label for q. Let X be the set of points not collinear with ∞. We consider the possible relationships of points in X. Write $q \sim r$ if q and r are collinear.

1. If $q, r \in X$ satisfy $q \sim r$, then f_q and f_r agree in just one position, viz., the unique index i for which the line L_i through ∞ meets the line qr.

2. If q, r are not collinear but some point of X is collinear with both, then f_q and f_r agree in all but two positions; for all but two values are changed twice, the remaining two being changed just once.

3. Otherwise, $f_q = f_r$; for all the common neighbours of q and r are adjacent to ∞.

Note too that, for any $i \in I$ and $q \in X$, there is a point $r \sim q$ for which f_q and f_r agree only in i, viz., the last point of the line through q meeting L_i.

Now suppose (as we may) that $|I| > 2$, and choose distinct $i, j, k \in I$. Given $q \in X$, choose $r, s, t \in X$ such that f_q and f_r agree only in i, f_r and f_s only in j, and f_s and f_t only on k. Then clearly f_q and f_t agree in precisely the three points i, j, k, since these values are changed twice and all others three times. By the case analysis, it follows that $|I| = 3$ or $|I| - 2 = 3$, as required.

The uniqueness also follows from this analysis, with a little more work: we know enough about the structure of X that the entire geometry can be reconstructed. ■

Problem. Can there exist a GQ with s finite ($s > 1$) and t infinite?

The proof above shows that there is no such GQ with $s = 2$. It is also known that there is no GQ with $s = 3$ or $s = 4$ and t infinite, though the proofs are much harder. (This is due to Kantor and Brouwer for $s = 3$, and Cherlin for $s = 4$.) Beyond this, nothing is known, though Cherlin's argument could in principle be extended to larger values of s.

The GQs with $s = 2$ and $t = 1,2$ have simple descriptions. For $s = 1$ we have the 3×3 grid. For $t = 2$, take the points to be all the 2-element subsets of a set of cardinality 6, and the lines to be all partitions of the 6-set into three disjoint 2-subsets. The GQ with order $(2,4)$ is a little harder to describe. The implicit construction in Theorem 7.3 is one of the simplest — the functions f_q are all those which take the value 1 an even number of times, each such function representing a unique point. This GQ also arises in classical algebraic geometry, as the Schläfli configuration of 27 lines in a general cubic surface, lying three at a time in 45 planes.

In the classical polar spaces, the orders s and t are both powers of the same prime. There are examples where this is not the case — see Exercise 5.

Exercises

1. Prove that the dual of a GQ is a GQ.

2. Prove that a GQ with two points on any line is a complete bipartite graph.

3. Let G be a finite GQ with orders s,t, where $s > 1$. Let p_1 and p_2 be non-adjacent points, and let x_n be the number of points p_3 adjacent to neither p_1 nor p_2 for which there are exactly n common neighbours of p_1, p_2 and p_3. Show that

$$\begin{aligned}\sum x_n &= s^2t - st - s + t,\\ \sum n x_n &= s(t+1)(t-1),\\ \sum n(n-1)x_n &= (t+1)t(t-1).\end{aligned}$$

Hence prove that $t \leq s^2$, with equality if and only if any three pairwise non-collinear points have exactly $s+1$ common neighbours.

4. In this exercise, we use the terminology of coding theory (as in Section 3.2). Consider the space V of words of length 6 with even weight. This is a vector space of rank 5 over $\mathrm{GF}(2)$. The "standard inner product" on V is a bilinear form which is alternating (by the even weight condition); its radical $V^\perp$ is spanned by the unique word of weight 6. Thus, $V/V^\perp$ is a vector space of rank 4 carrying a non-degenerate alternating bilinear form. The 15 non-zero vectors of this space

are cosets of $V^{\perp}$ containing a word of weight 2 and the complementary word of weight 4, and so can be identified with the 2-subsets of a 6-set. Extend this identification to an isomorphism between the combinatorial description of the GQ with orders (2, 2) and the rank 2 symplectic polar space over $\mathrm{GF}(2)$.

5. Let q be an even prime power, and let C be a hyperoval in $\Pi = \mathrm{PG}(2,q)$, a set of $q+2$ points meeting every line in 0 or 2 points (see Section 4.3). Now take Π to be the hyperplane at infinity of $\mathrm{AG}(3,q)$. Let G be the geometry whose points are all the points of $\mathrm{AG}(3,q)$, and whose lines are all the lines of $\mathrm{AG}(3,q)$ which meet Π in a point of C. Prove that G is a GQ with orders $(q-1,q+1)$.

6. Construct "free" GQs.

7.2 Diagrams for polar spaces

The inductive properties of polar spaces are exactly what is needed to show that they are diagram geometries.

Proposition 7.4 *A classical polar space of rank n belongs to the diagram*

with n nodes.

Proof Given a variety U of rank d, the varieties contained in it form a projective space of dimension $d-1$, while the varieties containing it are those of the polar space $U^{\perp}/U$ of rank $n-d$; moreover, any variety contained in U is incident with any variety containing U. Since a rank 2 polar space is a generalised quadrangle, it follows by induction that residues of varieties are correctly described by the diagram. ∎

This diagram is commonly referred to as C_n.

By analogy with Section 5.2, it might be thought that any geometry with diagram C_n for $n \geq 3$ is a classical polar space. This is false for several reasons, which we will see at various points. But first, here is one example of a geometry with diagram C_3 which is nothing like a polar space, even though it is highly symmetrical. This geometry was discovered by Arnold Neumaier, and is referred to as *Neumaier's geometry* or the *A_7-geometry*.

Let X be a set of seven points. The structure of a projective plane $\mathrm{PG}(2,2)$ can be imposed on X in 30 different ways — this number is the index of $\mathrm{PGL}(3,2)$

in the symmetric group S_7. Since $\mathrm{PGL}(3,2)$ contains no odd permutations, it is contained in the alternating group A_7 with index 15, and so the 30 planes fall into two orbits of length 15 under A_7. Now we take the points, lines, and planes of the geometry to be respectively the elements of X, the 3-element subsets of X, and one orbit of A_7 on $\mathrm{PG}(2,2)$s. Incidence between points and lines, or between lines and planes, is defined by membership; and every point is incident with every plane.

It is clear that the residue of a plane is a projective plane $\mathrm{PG}(2,2)$, while the residue of a line is a digon. Consider the residue of a point x. The lines incident with x can be identified with the 2-element subsets of the 6-element set $Y = X \setminus \{x\}$. Given a plane, its three lines containing x partition Y into three 2-sets. It is easy to check that, given such a triple of lines, there are just two ways to draw the remaining four lines to complete $\mathrm{PG}(2,2)$, and that these two are related by an odd permutation of X. So our chosen orbit of planes has exactly one member inducing the given partition of Y, and the planes incident with x can be identified with all the partitions of Y into three 2-sets. As we saw in Section 7.1, this incidence structure is a generalised quadrangle with order $2,2$.

We conclude that the geometry has the diagram

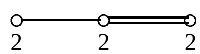

This example shows that, even in a geometry with such a simple diagram, a variety is not necessarily determined by its point-shadow (all planes have the same point-shadow!); the intersection of point-shadows of varieties need not be the point-shadow of a variety, and the points and lines need not form a partial linear space. So the special properties of linear diagrams with all strokes $\circ\overset{L}{\text{———}}\circ$ do not extend. However, classical polar spaces do have these nice properties.

A C_3-geometry in which every point and every plane are incident is called *flat*. Neumaier's geometry is the only known finite example of such a geometry. Some infinite examples were constructed by Sarah Rees; we now describe these. First, a re-interpretation of Neumaier's geometry.

Consider the rank 6 vector space V of all binary words of length 7 having even weight. On V, we can define a quadratic form by the rule

$$f(\mathbf{v}) = \tfrac{1}{2}\,\mathrm{wt}(\mathbf{v}) \pmod 2.$$

The bilinear form obtained by polarising f is just the usual dot product, since

$$\mathrm{wt}(\mathbf{v}+\mathbf{w}) = \mathrm{wt}(\mathbf{v}) + \mathrm{wt}(\mathbf{w}) - 2\mathbf{v}.\mathbf{w}\,.$$

It follows that f is non-singular: the only vector orthogonal to V is the all-1 word, which is not in V. Now the points of X, which index the coordinates, are in one-one correspondence with the seven words of weight 6, which are non-singular vectors. The lines correspond to the vectors of weight 4, which comprise all the singular vectors.

We saw in Section 6.4 that the planes on the quadric fall into two families, such that two planes of the same family meet in a subspace of even codimension (necessarily a point), while planes of different families meet in a subspace of odd codimension (the empty set or a line). Now a plane on the quadric contains seven non-zero singular vectors (of weight 4), any two of which are orthogonal, and so meet in an even number of points, necessarily 2. The complements of these 4-sets form seven 3-sets, any two meeting in one point, so forming a projective plane $\mathrm{PG}(2,2)$. It is readily checked that the two classes of planes correspond exactly to the two orbits of A_7 we described earlier. So the points, lines and planes of Neumaier's geometry can be identified with a special set of seven non-singular points, the singular points, and one family of planes on the quadric. Incidence between the non-singular and the singular points is defined by orthogonality.

Now we reverse the procedure. We start with a hyperbolic quadric Q in $\mathrm{PG}(5,F)$, that is, a quadric of rank 3 with germ zero. A set S of non-singular points is called an *exterior set* if it has the property that, given any line L of Q, a unique point of S is orthogonal to L. Now consider the geometry G whose POINTS, LINES and PLANES are the points of S, the points of Q, and one family of planes on Q; incidence between POINTS and LINES is defined by orthogonality, that between LINES and PLANES is incidence in the polar space, and every POINT is incident with every PLANE.

Such a geometry belongs to the diagram C_3. For the residue of a PLANE Π is a projective plane, naturally the dual of Π. (The correspondence between points of S and lines of Π is bijective; for, given $x \in S$, $x^\perp$ cannot contain Π, since a polar space in $\mathrm{PG}(4,q)$ cannot have rank 3, and so it meets Π in a line.) The residue of a POINT x is the polar space $x^\perp$, which as we've seen is rank 2, and so a GQ. And of course the POINTS and PLANES incident with a LINE form a digon.

Thas showed that no further finite examples can be constructed in this way:

Theorem 7.5 *There is no exterior set for the hyperbolic quadric in* $\mathrm{PG}(5,q)$ *for* $q > 2$.

However, Rees (who first described this construction) observed that there are infinite examples. Consider the case where $F = \mathbb{R}$; let the form be $x_1x_2 + x_3x_4 +$

x_5x_6. Now the space of rank 3 spanned by $(1,1,0,0,0,0)$, $(0,0,1,1,0,0)$ and $(0,0,0,0,1,1)$ is positive definite, and so is disjoint from the quadric; the points spanned by vectors in this space form an exterior set.

Now we turn to hyperbolic quadrics in general. As we saw in Section 6.4, the maximal t.s. subspaces on such a quadric Q of rank n can be partitioned into two families, so that a flat of dimension $n-2$ lies in a unique member of each family. We construct a new geometry by letting these flats be varieties of different types. Now there is no need to retain the flats of dimension $n-2$, since such a flat is the intersection of the two maximal flats containing it.

Theorem 7.6 *Let Q be a hyperbolic quadric of rank* $n \geq 3$. *Let G be the geometry whose flats are the t.s. subspaces of dimension different from* $n-2$, *where the two families of flats of dimension* $n-1$ *are assigned different types. Incidence between flats, at least one of which has dimension less than* $n-1$, *is as usual; while* $(n-1)$*-flats of different types are incident if they intersect in an* $(n-2)$*-flat. Then the geometry has diagram*

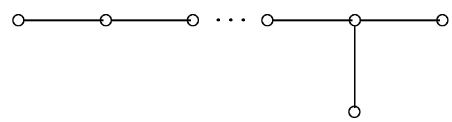

(n nodes).

Proof We need only check the residue of a flat of dimension $n-3$: the rest follows by induction, as in Proposition 7.4. Such a flat cannot be the intersection of two $(n-1)$-flats of different types; so any two such flats of different types containing it are incident. ■

This diagram is denoted by D_n. The result holds also for $n = 2$, provided that we interpret D_2 as two unconnected nodes — the quadric has two families of lines, each line of one family meeting each line of the other.

Exercises

1. Prove that the line joining two points of an exterior set to the quadric Q is disjoint from Q.

2. Prove that an exterior set to a quadric in $\mathrm{PG}(5,q)$ must have q^2+q+1 points.

3. Show that the plane constructed in Rees' example is an exterior set.

7.3 Tits and Buekenhout–Shult

We now begin working towards the axiomatisation of polar spaces. This major result of Tits (building on earlier work of Veldkamp) will not be proved completely here, but the next four sections should give some impression of how the proof works.

Tits' theorem characterises a class of spaces which almost coincides with the classical polar spaces of rank at least 3. There are a few additional examples of rank 3, some of which will be described later. I will use the term *abstract polar space* for a geometry satisfying the axioms. In fact, Tits' axioms describe all subspaces of arbitrary dimension; an alternative axiom system, proposed by Buekenhout and Shult, involves only points and lines (in the spirit of the Veblen–Young axioms for projective spaces). In this section, I show the equivalence of these axiom systems.

Temporarily, then, an *abstract polar space of type T* is a geometry satisfying the conditions (P1)–(P4) of Section 6.4, repeated here for convenience.

(P1) Any flat, together with the flats it contains, is a projective space of dimension at most $r-1$.

(P2) The intersection of any family of flats is a flat.

(P3) If U is a flat of dimension $r-1$ and p a point not in U, then the union of the lines joining p to points of U is a flat W of dimension $r-1$; and $U \cap W$ is a hyperplane in both U and W.

(P4) There exist two disjoint flats of dimension $r-1$.

An *abstract polar space of type BS* is a geometry of points and lines satisfying the following conditions. In these axioms, a *subspace* is a set S of points with the property that if a line L contains two points of S, then $L \subseteq S$; a *singular subspace* is a subspace, any two of whose points are collinear.

(BS1) Any line contains at least three points.

(BS2) No point is collinear with all others.

(BS3) Any chain of singular subspaces is of finite length.

(BS4) If the point p is not on the line L, then p is collinear with one or all points of L.

(Note that (BS4) is our earlier (BS), and is the key condition here.)

Theorem 7.7 *(a) The points and lines of an abstract polar space of type T form an abstract polar space of type BS.*

(b) The singular subspaces of an abstract polar space of type BS form an abstract polar space of type T.

Proof (a) It is an easy deduction from (P1)–(P4) that any subspace is contained in a subspace of dimension $n-1$. For let U be a subspace, and W a subspace of dimension $n-1$ for which $U \cap W$ has dimension as large as possible; if $p \in U \setminus W$, then (P3) gives a subspace of dimension $n-1$ containing p and $U \cap W$, contradicting maximality.

Now, if L is a line and p a point not on L, let W be a subspace of dimension $n-1$ containing L. If $p \in W$, then p is collinear with every point of L; otherwise, the neighbours of p in W form a hyperplane, meeting L in one or all of its points.

Thus, (BS4) holds. The other conditions are clear.

(b) Now let G be an abstract polar space of type BS. Call two points *adjacent* if they are collinear; this gives the point set a graph structure. Every maximal clique in the graph is a subspace. For let S be a maximal clique, and $p,q \in S$; let L be a line containing p and q. Any point of $S \setminus L$ is collinear with p and q, and so with every point of L; thus $S \cup L$ is a clique, and by maximality, $L \subseteq S$.

If $p \notin S$ (where S is a maximal clique), then the set of neighbours of p in S is a hyperplane. Every point $q \in S$ lies outside such a hyperplane; for, by (BS2), there is a point p not adjacent to q. As we saw in Section 3.1, if every line has size 3, then this implies that S is a projective space; but this deduction cannot be made in general. However, in the present situation, Buekenhout and Shult are able to show that S is indeed a projective space. (In particular, this implies that two points lie on at most one line. For the union of two lines through two common points is a clique by (BS4), and so would be contained in a maximal clique. However, Buekenhout and Shult have to show that two points lie on at most one line before they know that the subspaces are projective spaces; the proof is surprisingly tricky.)

Any singular subspace lies in some maximal clique, and so is itself a projective space. Thus (P1) holds; and the remaining axioms can now be verified. ■

We will now simplify the terminology by using the term "abstract polar space" equally for either type.

The induction principles we used in classical polar spaces work in almost the same way in abstract polar spaces.

Proposition 7.8 *Let U be a $(d-1)$-dimensional subspace of an abstract polar space of rank n. Then the subspaces containing U form an abstract polar space of rank $n-d$.* ■

Exercise

1. Show directly that, in an abstract polar space of type BS having three points on any line, any two points lie on at most one line, and singular subspaces are projective.

7.4 Recognising hyperbolic quadrics

There are two special cases where the proof of the characterisation of polar spaces is substantially easier, namely, hyperbolic quadrics and quadrics over $\mathrm{GF}(2)$; they will be treated in this section and the next.

In the case of a hyperbolic quadric, we bypass the need to reconstruct the quadric by simply showing that there is a unique example of each rank over any field. First, we observe that the partition of the maximal subspaces into two types follows directly from the axioms; properties of the actual model are not required. We begin with a general result on abstract polar spaces.

An abstract polar space G can be regarded as a point-line geometry, as we've seen. Sometimes it is useful to consider a "dual" situation, defining a geometry G^* whose POINTs are the maximal subspaces of G and whose LINEs are the next-to-maximal subspaces, incidence being reversed inclusion. We call this geometry a *dual polar space*. In a dual polar space, we define the distance between two POINTs to be the number of LINEs on a shortest path joining them.

Proposition 7.9 *Let G^* be a dual polar space.*

(a) The distance between two POINTs is the codimension of their intersection.

(b) Given a POINT p and a LINE L, there is a unique POINT of L nearest to p.

Proof Let U_1, U_2 be maximal subspaces. By the inductive principle (Proposition 7.8), we may assume that $U_1 \cap U_2 = \emptyset$. (It is clear that any path from U_1 to U_2, in which not all terms contain $U_1 \cap U_2$, must have length strictly greater than the codimension of $U_1 \cap U_2$; so, once the result is proved in the quotient, no such path can be minimal.)

Now each point of U_1 is collinear (in G) with a hyperplane in U_2, and *vice versa*; so, given any hyperplane H in U_2, there is a unique point of U_1 adjacent to H, and hence (by (P3)) a unique maximal subspace containing H and meeting U_1. The result follows.

(b) Let U be a maximal subspace and W a subspace of rank one less than maximal. As before, we may assume that $U \cap W = \emptyset$. Now there is a unique point $p \in U$ collinear with all points of W. Then $\langle W, p \rangle$ is the uniue POINT on the LINE W nearest to the POINT U. ■

Proposition 7.10 *Let G be an abstract polar space of rank n, in which any* $(n-2)$*-dimensional subspace is contained in exactly two maximal subspaces. Then the maximal subspaces fall into two families, the intersection of two subspaces having even codimension in each if and only if the subspaces belong to the same family.*

Proof The associated dual polar space is a graph. By Proposition 7.9(b), the graph is bipartite, since if an odd circuit exists, then there is one of minimal length, and both vertices on any edge are then equidistant from the opposite vertex in the cycle. ■

Now, in any abstract polar space of rank $n \geq 4$, in which lines contain at least three points, any maximal subspace is isomorphic to $\mathrm{PG}(n-1, F)$ for some skew field F. Now an easy connectedness argument shows that the same field F coordinatises every maximal subspace.

Theorem 7.11 *Let G be an abstract polar space of rank n* ≥ 4*, in which each next-to-maximal subspace is contained in exactly two maximal subspaces. Assume that some maximal subspace is isomorphic to* $\mathrm{PG}(n-1, F)$*. Then F is commutative, and G is isomorphic to the hyperbolic quadric of rank n over F.*

Proof It is enough to show that F is commutative and that n and F uniquely determine the geometry, since the hyperbolic quadric clearly has the required property.

Rather than prove F commutative, I will show merely that it is isomorphic to its opposite. It suffices to show this when $n = 4$. Take two maximal subspaces meeting in a plane Π, and a point $p \in \Pi$. By the FTPG, both maximal subspaces are isomorphic to $\mathrm{PG}(3, F)$. Now consider the residue of p. This is a projective space, in which there is a plane isomorphic to $\mathrm{PG}(2, F^\circ)$, and a point residue isomorphic to $\mathrm{PG}(2, F)$. Hence $F \cong F^\circ$. The stronger statement that F is commutative is shown by Tits. He observes that the quotient of p has a polarity interchanging a point and a plane incident with it, and fixing every line incident with both; and this can only happen in a projective 3-space over a commutative field.

Let U_1 and U_2 be disjoint maximal subspaces. Note that they have the same type if n is even, opposite types if n is odd. Let p be any point in neither subspace. Then for $i = 1,2$, there is a unique maximal subspace W_i containing p and meeting U_i in a hyperplane. Then W_i has the opposite type to U_i, so W_1 and W_2 have the same type if n is even, opposite types if n is odd. Thus, their intersection has codimension congruent to n mod 2. Since $p \in W_1 \cap W_2$, the intersection is at least a line. But their distance in the dual polar space is at least $n-2$, since U_1 and U_2 have distance n; so $W_1 \cap W_2$ is a line L. Clearly L meets both U_1 and U_2.

Each point of U_1 is adjacent to a hyperplane of U_2, and *vice versa*; so U_1 and U_2 are naturally duals. Now the lines joining points of U_1 and U_2 are easily described, and it is not hard to show that the whole geometry is determined. ■

7.5 Recognising quadrics over GF(2)

In this section, we determine the abstract polar spaces with three points on every line. Since we are given information only about points and lines, the BS approach is the natural one. The result here was first found by Shult (assuming a constant number of lines per point) and Seidel (in general), and was a crucial precursor of the Buekenhout–Shult Theorem (Theorem 7.7). Shult and Seidel proved the theorem by induction on the rank: a rank 2 polar space is a generalised quadrangle, and the classification in this case is Theorem 7.3. The elegant direct argument given here is due to Jonathan Hall.

Let G be an abstract polar space with three points per line. We have already seen that the facts that two points lie on at most one line, and that maximal singular subspaces are projective spaces, are proved more easily under this hypothesis than in general. But here is a direct proof of the first assertion. Suppose that the points a and b lie on two lines $\{a,b,x\}$ and $\{a,b,y\}$. Then y is collinear with a and b, and so also with x; so there is a line $\{x,y,z\}$ for some z, and both a and b are joined to z. Any further point is joined to both or neither x and y, and so is joined to z, contradicting (BS2).

Define a graph Γ whose vertices are the points, two vertices being adjacent if they are collinear. The graph has the following property:

(T) every edge $\{x,y\}$ lies in a triangle $\{x,y,z\}$ with the property that any further point is joined to one or all of $\{x,y,z\}$.

This is called the *triangle property*. Shult and Seidel phrased their result as the determination of finite graphs with the triangle property. (The argument just given shows that, in a graph with the triangle property in which no vertex is adjacent to all others, there is a unique triangle with the property specified by (T) containing any edge. Thus, the graph and the polar space determine each other.) The proof given below is not the original argument of Shult and Seidel, which used induction, but is a direct argument due to Jonathan Hall (having the added feature that it works equally well for infinite-dimensional spaces).

Theorem 7.12 *An abstract polar space in which each line contains three points is a quadric over* $\mathrm{GF}(2)$.

Proof As noted above, we may assume instead that we have a graph Γ with the triangle property (T), having at least one edge, and having no vertex adjacent to all others. Let X be the vertex set of the graph Γ, and let $F = \mathrm{GF}(2)$. We begin with the vector space $\hat{V}$ of all functions from X to F which are zero everywhere except on a finite set, with pointwise operations. (If X is finite, then V is just the space F^X of all functions from X to F.) Let $\hat{x} \in \hat{V}$ be the characteristic function of the singleton set $\{x\}$. The functions $\hat{x}$, for $x \in X$, form a basis for $\hat{V}$. We define a bilinear form $\hat{b}$ on $\hat{V}$ by setting

$$\hat{b}(\hat{x},\hat{y}) = \begin{cases} 0 & \text{if } x = y \text{ or } x \text{ is joined to } y, \\ 1 & \text{otherwise,} \end{cases}$$

and extending linearly, and a quadratic form $\hat{f}$ by setting $\hat{f}(\hat{x}) = 0$ for all $x \in X$ and extending to $\hat{V}$ by the rule

$$\hat{f}(\mathbf{v}+\mathbf{w}) = \hat{f}(\mathbf{v}) + \hat{f}(\mathbf{w}) + \hat{b}(\mathbf{v},\mathbf{w}).$$

Note that both $\hat{b}$ and $\hat{f}$ are well-defined.

Let R be the *radical* of $\hat{f}$; that is, R is the *subspace*

$$\{\mathbf{v} \in \hat{V} : \hat{f}(\mathbf{v}) = 0,\ \hat{b}(\mathbf{v},\mathbf{w}) = 0 \text{ for all } \mathbf{w} \in \hat{V}\},$$

and set $V = \hat{V}/R$. Then $\hat{b}$ and $\hat{f}$ induce bilinear and quadratic forms b, f on V: for example, we have $f(\mathbf{v}+R) = \hat{f}(\mathbf{v})$ (and this is well-defined, that is, independent of the choice of coset representative). Now let $\overline{x} = \hat{x} + R \in V$.

We claim that the embedding $x \mapsto \overline{x}$ has the required properties; in other words, it is one-to-one; its image is the quadric defined by f; and two vertices are adjacent if and only if the corresponding points of the quadric are orthogonal. We proceed in a series of steps.

Step 1 *Let* $\{x,y,z\}$ *be a special triangle, as in the statement of the triangle property (T). Then* $\bar{x}+\bar{y}+\bar{z}=0$.

It is required to show that $r=\hat{x}+\hat{y}+\hat{z}\in R$. We have

$$\hat{b}(r,\hat{v})=\hat{b}(\hat{x},\hat{v})+\hat{b}(\hat{y},\hat{v})+\hat{b}(\hat{z},\hat{v})=0$$

for all $v\in X$, by the triangle property; and

$$\hat{f}(r)=\hat{f}(\hat{x})+\hat{f}(\hat{y})+\hat{f}(\hat{z})+\hat{b}(\hat{x},\hat{y})+\hat{b}(\hat{y},\hat{z})+\hat{b}(\hat{z},\hat{x})=0$$

by definition.

Step 2 *The map* $x\mapsto\bar{x}$ *is one-to-one on* X.

Suppose that $\bar{x}=\bar{y}$. Then $r=\hat{x}+\hat{y}\in R$. Hence $\hat{b}(\hat{x},\hat{y})=0$, and so x is joined to y. Let z be the third vertex of the special triangle containing x and y. Then $\hat{z}=\hat{x}+\hat{y}\in R$ by Step 1, and so z is joined to all other points of X, contrary to assumption.

Step 3 *Any quadrangle is contained in a* 3×3 *grid.*

Let $\{x,y,z,w\}$ be a quadrangle. Letting $\overline{x+y}=\bar{x}+\bar{y}$, etc., we see that $x+y$ is not joined to z or w, and hence is joined to $z+w$. Similarly, $y+z$ is joined to $w+x$; and the third point in the special triangle through each of these pairs is $x+y+z+w$, completing the grid. (See Fig. 7.1.)

Step 4 *For any* $v\in V$, *write* $v=\sum_{i\in I}\bar{x_i}$, *where* $x_i\in X$, *and the number* $m=|I|$ *of summands is minimal (for the given* v*). Then*

(a) $m\leq 3$*;*

(b) the points x_i *are pairwise non-adjacent.*

This is the crucial step, and needs four sub-stages.

Substep 4.1 *Assertion (b) is true.*

If $x_i\sim x_j$, we could replace x_i+x_j by the third point x_k of the special triangle, and obtain a shorter expression.

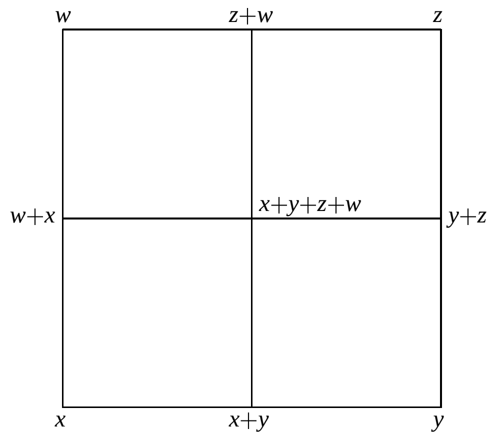

Figure 7.1: A grid

Substep 4.2 *If L is a line on x_1, and y a point of L which is adjacent to x_2, then $y \sim x_i$ for all $i \in I$.*

If not, let $L = \{x_1, y, z\}$, and suppose that $x_i \sim z$. Then x_i is joined to the third point w of the line $x_2 y$. Let u be the third point on $x_i w$. Then $\overline{z} + \overline{p_1} + \overline{u} + \overline{x_i} = \overline{x_2}$, and we can replace $\overline{x_1} + \overline{x_2} + \overline{x_i}$ by the shorter expression $\overline{z} + \overline{u}$.

Substep 4.3 *There are two points y, z joined to all x_i.*

Each line through x_1 contains a point with this property, by Substep 4.2. It is easily seen that if x_1 lies on a unique line, then one of the points on this line is adjacent to all others, contrary to assumption.

Substep 4.4 $m \leq 3$.

Suppose not. Considering the quadrangles $\{x_1, y, x_2, z\}$ and $\{x_3, y, x_4, z\}$, we find (by Step 3) points a and b with

$$\overline{x_1} + \overline{y} + \overline{x_2} + \overline{z} = \overline{a}, \qquad \overline{x_3} + \overline{y} + \overline{x_4} + \overline{z} = \overline{b}.$$

But then $\overline{x_1} + \overline{x_2} + \overline{x_3} + \overline{x_4} = \overline{a} + \overline{b}$, a shorter expression.

Step 5 *If $v \in V$, $v \neq 0$, and $f(v) = 0$, then $v = \overline{x}$ for some $x \in X$.*

If not then, by Step 4, either $v = \overline{x} + \overline{y}$, or $v = \overline{x} + \overline{y} + \overline{z}$, where points x, y (and z) are (pairwise) non-adjacent. In the second case,

$$f(v) = f(\overline{x}) + f(\overline{y}) + f(\overline{z}) + b(\overline{x},\overline{y}) + b(\overline{y},\overline{z}) + b(\overline{z},\overline{x}) = 0+0+0+1+1+1 = 1.$$

The other case is similar but easier.

Step 6 *$x \sim y$ if and only if $b(\overline{x},\overline{y}) = 0$.*
This is true by definition. ∎

7.6 The general case

A weak form of the general classification of polar spaces, by Veldkamp and Tits, can be stated as follows.

Theorem 7.13 *A polar space of type T having finite rank $n \geq 4$ is either classical, or defined by a pseudoquadratic form on a vector space over a division ring of characteristic 2.* ∎

I will not attempt to outline the proof of this theorem, but merely make some remarks, including a "definition" of a pseudoquadratic form.

Let V be a vector space over a skew field F of characteristic 2, and σ an anti-automorphism of F satisfying $\sigma^2 = 1$. Let K_0 be the additive subgroup $\{x + x^\sigma\}$ of F, and $K^* = K/K_0$. A function $f : V \to K^*$ is called a *pseudo-quadratic form* relative to σ if there is a σ-sesquilinear form g such that $f(\mathbf{v}) = g(\mathbf{v},\mathbf{v}) \bmod K_0$. Equivalently, f polarises to a σ-Hermitian form f satisfying $(\forall \mathbf{v} \in V)(\exists c \in F)(f(\mathbf{v},\mathbf{v}) = c + c^\sigma)$, that is, a trace-valued form. The function f defines a polar space, consisting of the subspaces of V on which f vanishes (mod K_0). If K_0 is equal to the fixed field of σ, then the same polar space is defined by the Hermitian form g; so we may assume that this is not the case in the second conclusion of Theorem 7.13. For further discussion, see Tits [S].

Tits' result is actually better than indicated: all polar spaces of rank $n \geq 3$ are classified. There are two types of polar spaces of rank 3 which are not covered by Theorem 7.13. The first exists over any non-commutative field, and will be described in the first section of Chapter 8. The other is remarkable in consisting of the only polar spaces whose planes are non-Desarguesian. This type is constructed by Tits from the algebraic groups of type E_6, and again I refer to Tits for the construction, which requires detailed knowledge of these algebraic groups. The planes actually satisfy a weakening of Desargues' theorem known as the *Moufang condition*, and can be "coordinatised" by certain *alternative division rings* which generalise the *Cayley numbers* or *octonions*.

Of course, the determination of polar spaces of rank 2 (GQs) is a hopeless task! Nevertheless, it is possible to formulate the Moufang condition for generalised quadrangles; and all GQs satisfying the Moufang condition have been determined (by Fong and Seitz in the finite case, Tits and Weiss in general.) This effectively completes the analogy with coordinatisation theorems for projective spaces.

The other geometric achievement of Tits in the 1974 lecture notes is the analogue of the Fundamental Theorem of Projective Geometry:

Theorem 7.14 *Any isomorphism between classical polar spaces of rank at least* 2*, which are not of symplectic or orthogonal type in characteristic* 2*, is induced by a semilinear transformation of the underlying vector spaces.* ∎

The reason for the exception will be seen in Section 8.4. As in Section 1.3, this result shows that the automorphism groups of classical polar spaces consist of semilinear transformations modulo scalars. These groups, with some exceptions of small rank, have "large" simple subgroups, just as happened for the automorphism groups of projective spaces in Section 4.6. These groups are the *classical groups*, and are named after their polar spaces: *symplectic*, *orthogonal* and *unitary* groups. For details, see the classic accounts: Dickson [K], Dieudonné [L], and Artin [B], or for more recent accounts Taylor [R], Cameron [10]. In the symplectic or unitary case, the classical group consists of all the linear transformations of determinant 1 preserving the form defining the geometry, modulo scalars. In the orthogonal case, it is sometimes necessary to pass to a subgroup of index 2. (For example, if the polar space is a hyperbolic quadric in characteristic 2, take the subgroup fixing the two families of maximal t.s. subspaces.)