Finite Geometry and Combinatorial Applications

Anurag Bishnoi

Preface

Incidence geometry is the study of abstract structures satisfying certain geometric axioms inspired from the incidence properties of points, lines, planes, etc. For example, the axiom that through any two points there is a unique line, or that any two lines intersect in at most one point. In this course we will introduce several finite incidence structures and explore how these structures interact with combinatorics. In particular, we will study finite projective and affine spaces, generalized polygons and polar spaces. On the combinatorial side we will study Latin squares, blocking sets, strongly regular graphs, finite field Kakeya and Nikodym problems, etc.

Recommended Reading

References for the course

- "Finite Geometry and Combinatorial Applications" by Simeon Ball.
- "An Introduction to Incidence Geometry" by Bart De Bruyn.
- "Incidence Geometry" by G. Eric Moorhouse, `http://math.ucr.edu/home/baez/qg-fall2016/incidence_geometry.pdf`.
- "Projective Geometry" by Rey Casse.
- "Algebraic Graph Theory" by Chris Godsil and Gordon Royle.

Further reading

- "Finite Geometries" by P. Dembowski.
- "Points and Lines" by E. Shult.
- "Projective Geometries over finite fields" by Hirschfeld and Thas.
- "Foundations of Incidence Geometry" by Johannes Ueberberg.
- "Projective Geometry: From Foundations to Applications" by Albrecht Beutelspacher and Ute Rosenbaum.
- "Combinatorics of finite geometries" by Lynn Margaret Batten.
- "Distance Regular Graphs" by Brouwer, Cohen and Neumaier.
- "Designs, Graphs, Codes and their Links" by Peter J. Cameron and J. H. van Lint.
- "Spectra of Graphs" by Brouwer and Haemers.

Background required

1 Linear Algebra

The reader should be familiar with the following notions from linear algebra: vector spaces over a field, linear independence, span, basis of a vector space, the lattice of subspaces, linear maps between vector spaces, linear forms, dual spaces, ranks of matrices, determinants, quotient spaces, eigenvalues and eigenspaces.

The following basic theorems will often be directly useful to us.

Theorem 0.1. *Let U_1, U_2 be two finite dimensional subspaces of a vector space V. Then*

$$\dim U_1 + \dim U_2 = \dim(U_1 \cap U_2) + \dim(U_1 + U_2).$$

Theorem 0.2 (Rank-Nullity Theorem)**.** *Let L be a linear map from a vector space V to a vector space W. Then*

$$\dim \ker(L) + \dim \operatorname{im}(L) = \dim V.$$

Theorem 0.3. *Let U be a subspace of an n dimensional vector space V, and let V/U be the quotient space with respect to U. Then $\dim V/U = n - \dim U$.*

Theorem 0.4. *Let A be a square matrix of order n. Then the eigenvalues of A are the solutions of the degree n polynomial equation $\det(A - \lambda I) = 0$, where I is the identity matrix.*

Corollary 0.5. *The sum of all eigenvalues of A is equal to the trace of A, that is, the sum of the diagonal entries of A, and the product of all eigenvalues is equal to the determinant of A.*

Theorem 0.6. *Let A be a real symmetric matrix of order n. Then A has n eigenvalues, $\lambda_1 \geq \cdots \geq \lambda_n$, counted with multiplicity. Moreover, there exists an orthonormal basis of $\mathbb{R}^n$, $x_1, \ldots, x_n$, such that x_i is an eigenvector of M with eigenvalue equal to λ_i.*

2 Finite Fields

The reader should be familiar with the notions of groups, rings, ideals, fields, quotient rings, isomorphism of rings/fields and field extensions.

The simplest example of a finite field is $\mathbb{F}_p = \mathbb{Z}/p\mathbb{Z}$, that is, the integers modulo p, for a prime number p. The main thing to check is that every element has an inverse, which follows from Euclid's theorem: for any integers a, b with $\mathrm{GCD}(a, b) = 1$, there exist integers x and y such that $ax + by = 1$. Let F be an arbitrary field which is finite. By the repeated addition of 1, we see that there exists a smallest integer m such that $1 + \cdots + 1$ (m times) is equal to 0 in F. This smallest integer must necessarily be a prime p, since there are no zero divisors in F, and this prime number is known as the characteristic of the finite field F. The set $\{0, s_1, \ldots, s_{p-1}\} \subseteq F$ where s_i is the i-fold sum of 1, is a subfield of F, which is known as the *prime subfield*, and it is isomorphic to $\mathbb{F}_p$.

Theorem 0.7. *For every finite field F, there exists a prime number p and a positive integer n such that $|F| = p^n$.*

Proof. Let p be the characteristic of F, and let n be the dimension of F when seen as a vector space over its prime subfield. The number of elements in an n-dimensional vector space over $\mathbb{F}_p$ is equal to p^n. □

Therefore, every finite field must have order equal to a prime power. In fact, we have the following converse.

Theorem 0.8. *For every prime power q, there exists a unique field of order q (up to isomorphism).*

Proof. Let $q = p^n$ where p is a prime. Any finite field of characteristic p is algebraic over the subfield $\mathbb{F}_p$, since it is finite, and hence contained in the algebraic closure F of $\mathbb{F}_p$. The unique field of order q is then just the set of zeros of the polynomial $x^q - x$ in F. □

This finite field of order $q = p^n$ is denoted by $\mathbb{F}_q$.[1] The proof above is not really constructive. The most concrete way to construct the finite field $\mathbb{F}_q$, and to do computations in it, is to take an irreducible polynomial $f(x) \in \mathbb{F}_p[x]$ of degree n (they always exist), to get the field $\mathbb{F}_q$ as $\mathbb{F}_p[x]/(f(x))$.

We now enumerate some basic properties of $\mathbb{F}_q$ that will be useful to us.

(a) For every $\alpha \in \mathbb{F}_q$ we have $\alpha^q = \alpha$. Moreover, the elements of the prime subfield are precisely the zeros of the polynomial $x^p - x \in \mathbb{F}_q[x]$.

(b) $\mathbb{F}_{p^m}$ is a subfield of $\mathbb{F}_{p^n}$ if and only if m divides n. Moreover, for every divisor m of n, there exists a unique copy of $\mathbb{F}_{p^m}$ in the subfield lattice of $\mathbb{F}_{p^n}$, given by the elements satisfying $\alpha^{p^m} = \alpha$ in $\mathbb{F}_{p^n}$.

(c) The additive group of $\mathbb{F}_q$ is isomorphic to $(\mathbb{Z}/p\mathbb{Z})^n$. The multiplicative group $\mathbb{F}_q^*$ of $\mathbb{F}_q$, is generated by a single element of order $q - 1$, and therefore it's isomorphic to $\mathbb{Z}/(q-1)\mathbb{Z}$. Any generator of this group is called a primitive element of $\mathbb{F}_q$.

(d) The map $x \mapsto x^p$ is an automorphism of the field $\mathbb{F}_q$. The full automorphism group of $\mathbb{F}_q$ is a cyclic group of order n consisting of the maps $x \mapsto x^{p^i}$, for $i = 1, \ldots, n$.

(e) For q odd, there are precisely $(q-1)/2$ squares (quadratic residues) in $\mathbb{F}_q^*$ and $(q-1)/2$ non-squares. Multiplying $\mathbb{F}_q^*$ by a square fixes these two sets while multiplying by a non-square permutes the two of them. For q even, every element of $\mathbb{F}_q$ is a square since $x \mapsto x^2$ is a field automorphism.

[1] In finite geometry literature it is sometimes denoted by GF(q), where GF stands for *Galois Field*.

Lemma 0.9. *For all $i \in \mathbb{N}$, $i \neq 0$, the sum*

$$\sum_{a \in \mathbb{F}_q} a^i$$

is -1 if i is a multiple of $q-1$ and 0 otherwise.

Proof. Let S be the sum and α be the generator of the group $\mathbb{F}_q^*$. Then

$$S = \sum_{a \in \mathbb{F}_q} a^i = \sum_{a \in \mathbb{F}_q} (\alpha a)^i = \alpha^i \sum_{a \in \mathbb{F}_q} a^i = \alpha^i S.$$

Now, α^i is equal to 1 if and only if i is a multiple of $q-1$, and hence for all $i \notin (q-1)\mathbb{N}$, we must have $S = 0$ for the equality above to hold. When $i \in (q-1)\mathbb{N}$, then since $a^{q-1} = 1$ for all $a \neq 0$, S is the $(q-1)$-fold sum of the identity 1, and hence it's equal to -1. □

Definition 0.10. Let $\mathbb{F}_{q^n}$ be a field extension of $\mathbb{F}_q$. Then the *trace* function, with respect to this extension, is defined as

$$\mathrm{Tr}(x) = x + x^q + x^{q^2} + \cdots + x^{q^{n-1}}.$$

The *norm* function with respect to this extension is defined as

$$\mathrm{Norm}(x) = x x^q \cdots x^{q^{n-1}} = x^{(q^n-1)/(q-1)}.$$

If $\mathbb{F}_q$ is the prime subfield, then these maps are known as *absolute trace* and *absolute norm*, respectively.

Lemma 0.11. Tr *is an additive surjective function from $\mathbb{F}_{q^n}$ to $\mathbb{F}_q$ with $|\mathrm{Tr}^{-1}(a)| = q^{n-1}$ for all $a \in \mathrm{F}_q$.*

Lemma 0.12. Norm *is a multiplicative function from $\mathbb{F}_{q^n}^*$ to $\mathbb{F}_q^*$, with $|\mathrm{Norm}^{-1}(a)| = (q^n-1)/(q-1)$ for all $a \in \mathbb{F}_q^*$.*

Contents

Preface iii

Background required v

1 Linear Algebra . . . v

2 Finite Fields . . . v

1 Affine and Projective Spaces 3

1.1 Introduction . . . 3

1.2 Affine planes . . . 4

1.3 Mutually Orthogonal Latin Squares . . . 6

1.4 Projective planes . . . 9

1.5 Duality . . . 12

1.6 Projetive spaces from vector spaces . . . 13

1.7 Desarguesian planes . . . 14

1.8 Exercises . . . 16

2 Conics and Ovals 19

2.1 Coordinates and Hypersurfaces . . . 19

2.2 Conics . . . 20

2.3 Ovals and Hyperovals . . . 23

2.4 Exercises . . . 26

3 Some Combinatorial Questions 29

3.1 Blocking Sets . . . 29

3.2 Kakeya and Nikodym Sets . . . 33

3.2.1 Nikodym sets and large minimal blocking sets . . . 36

3.3 Exercises . . . 38

4 Strongly Regular Graphs 41

4.1 Introdction . . . 41

4.2 Moore graphs . . . 42

4.3 Spectral Methods . . . 43

4.4 Two-intersection sets . . . 48

4.5 Exercises . . . 50

5 Generalized Quadrangles 53

5.1 Introduction . . . 53

5.2 Basic properties . . . 54

5.3 Some examples . . . 56

5.4 An application to Ball Packings . . . 58

5.5 Generalized polygons . . . 61

5.6 Exercises 66

6 Polar Spaces 69

6.1 Introduction 69
6.2 Collinearity Graph of a Polar Space 71
6.3 Quadrics and Symmetric Bilinear Forms 72
6.4 Sesquilinear Forms and Classical Polar Spaces 77
6.5 Exercises 78

1 Affine and Projective Spaces

1.1 Introduction

Incidence geometry is the study of mathematical structures that are formed by abstracting out the notion of incidence from Euclidean geometry. In this course we will be studying some finite incidence structures, and how they interact with combinatorics. Our main objects of study are the following.

Definition 1.1 (Point-Line Geometries)**.** A point-line geometry is a triple $\mathcal{S} = (\mathcal{P}, \mathcal{L}, \mathrm{I})$ where $\mathcal{P}$ is a non-empty set, $\mathcal{L}$ is a set disjoint from $\mathcal{P}$, and I is a subset of $\mathcal{P} \times \mathcal{L}$ such that for every $\ell \in \mathcal{L}$, there exists at least two $x \in \mathcal{P}$ such that $(x, \ell) \in \mathrm{I}$. The elements of $\mathcal{P}$ are referred to as *points* of $\mathcal{S}$, elements of $\mathcal{L}$ as *lines* of $\mathcal{S}$ and I is the *incidence relation* between points and lines.

Example 1.2. The Euclidean plane is one of the classical examples of a point-line geometry, where $\mathcal{P} = \mathbb{R}^2$ is the set of ordered pairs (x, y) of real numbers, $\mathcal{L}$ is the set of all straight lines, that is, subsets of $\mathcal{P}$ of the form $\{(x, mx+c) : x \in \mathbb{R}\}$ or of the form $\{(c, y) : y \in \mathbb{R}\}$ where $m, c \in \mathbb{R}$, and the incidence relation I is simply set containment.

Example 1.3. Here is a finite example of a point-line geometry. Let $\mathcal{P} = [4]$, $\mathcal{L} = \binom{[4]}{2}$ and I set containment (see Figure 1.1).

While these two examples might look quite different, they both belong to the family of *affine planes*, which will be discussed in the next section.

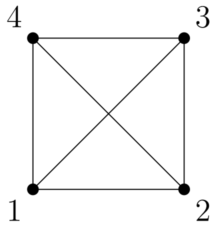

Figure 1.1: Affine plane of order 2

Example 1.4. Every multi-graph without a loop is a point-line geometry, where each line is incident with precisely two points. Every hypergraph[1] where each edge contains at least two vertices is a point-line geometry.

We will often use geometrical language and write the statement "$(x, \ell) \in \mathrm{I}$" as "the point x lies on the line ℓ", or "ℓ is a line through x", etc. Two points will be called *collinear* if there is a common line through both of them, and two lines will be said to *meet* each other if there is a common point incident to both of them.

[1] A hypergraph is a pair (V, E) where V is a non-empty set and E is a collection of subsets of V.

Definition 1.5. A point-line geometry is called a *partial linear space* if through every pair of distinct points there is at most one line. It is called a *linear space* if there is *exactly one* line through every pair of points.

Note that two distinct lines in a partial linear space intersect each other in at most one point, and hence these lines do behave like the lines that we are used to in Euclidean geometry. In a linear space, we will denote the unique line through two distinct points x, y by xy.

For partial linear spaces, we can uniquely identify each line with the subset of points it is incident with. Thus, we can think of these point-line geometries as a hypergraph, with the incidence relation as set containment. In fact, partial linear spaces are equivalent to the so-called *linear hypergraphs.*

1.2 Affine planes

Abstracting out the incidence properties of the Euclidean plane, we get the following structure.

Definition 1.6. An affine plane is a linear space with the following properties:

(A1) (*Playfair axiom*) For every point x and a line ℓ not through x, there exists a unique line ℓ' through x that does not meet ℓ.

(A2) There exist three non-collinear points.

It can be easily checked that the point-line geometries from Example 1.2 are both affine planes.

Definition 1.7. In a affine plane, we say that two lines ℓ and m are *parallel*, denoted by $\ell \parallel m$, if either $\ell = m$ or ℓ and m do not meet each other.

Proposition 1.8. *Parallelism is an equivalence relation on the lines of an affine plane, and each parallel class is a partition of the set of points.*

Proof. The fact that it is an equivalence relation follows easily from the definition. Now take a parallel class $\mathcal{S}$ and a point x of $\mathcal{S}$. Let ℓ be a line of $\mathcal{S}$. Then there exists a unique line through x which is parallel to ℓ, and hence a unique line of $\mathcal{S}$ through x. Since x was arbitrary, this shows that the point set is partitioned by $\mathcal{S}$. □

Theorem 1.9. *For every finite affine plane there exists an integer $n \geq 2$, called the order of the plane, such that:*

(*a*) *Each point is incident with exactly $n+1$ lines.*

(*b*) *Each line is incident with exactly n points.*

(*c*) *Each parallel class has n lines.*

(*d*) *There are $n+1$ parallel classes.*

(*e*) *There are n^2 points.*

(*f*) *There are* $n^2 + n$ *lines.*

Proof. (*a*) Since we have a finite affine plane, the number of lines through each point is finite. Let $x \neq y$ be two points of the plane. From (A1) it follows that for every line ℓ through x there is a unique line m through y which is parallel to ℓ. This gives a bijection between lines through x and lines through y. Define n to be such that $n+1$ is the common number of lines through every point. We now show that $n \geq 2$.

Pick three non-collinear points x, y, z, which exist by (A2). Besides the lines xy and xz through x, there is the unique line parallel to yz through x, giving us at least three lines through x. Therefore $n + 1 \geq 3$.

(*b*) Let ℓ be a line and x a point not in ℓ. Such a point exists since each line has at least two points on it, and through each point of ℓ there is at least one other line (recall that $n \geq 2$). There is a unique line through x which is parallel to ℓ, whereas the remaining n lines through x intersect ℓ in a unique point. These are all the points on ℓ since for every point y on ℓ the line joining x and y is a line through x that intersects ℓ.

(*c*) Let $\mathcal{S}$ be a parallel class and $\ell \in \mathcal{S}$. Take a point x on ℓ and a line $m \neq \ell$ through x. Through each of the $n - 1$ points on m other than x, there is a unique line parallel to ℓ, giving us $|\mathcal{S}| \geq n$. Moreover, every line in $\mathcal{S}$ must intersect m in a point since $m \notin \mathcal{S}$, giving us $|\mathcal{S}| \leq n$.

(*d*) Let x be a point. For every parallel class $\mathcal{S}$, there exists a unique line through x which is contained in $\mathcal{S}$. Moreover, there are $n + 1$ lines through x, giving us $n + 1$ different parallel classes.

(*e*) Since a parallel class partitions the set of points, and there are n lines in a parallel class with each line containing n points, we have n^2 points in total.

(*f*) The $n + 1$ parallel classes, each containing n lines, partition the set of lines, giving us $n(n + 1)$ lines in total.

□

Example 1.10. Let F be any field. The point-line geometry $\mathrm{AG}(2, F)$, with point set equal to F^2 and lines as sets of the form $\{(x, y) : x = c, y \in F\}$ and $\{(x, y) : y = mx + c, x \in F\}$ for $m, c \in F$, is an affine plane. A parallel class is either the set of "vertical" lines or the set of lines with a common slope m. If F is a finite field of order q, then we get a finite affine plane, which will be denoted by $\mathrm{AG}(2, q)$. See Figure 1.2 for a drawing of $\mathrm{AG}(2, 3)$, where for example, the lines AIE, CFH and BDG form a parallel class.

Remark 1.11. An alternate way of describing the affine plane $\mathrm{AG}(2, F)$ is as follows: the points are the elements of the vector space F^2 and the lines are the cosets (with respect to vector space addition) of 1-dimensional subspaces of F^2.

With this background on affine planes, we are ready to see our first connection between finite geometry and combinatorics.

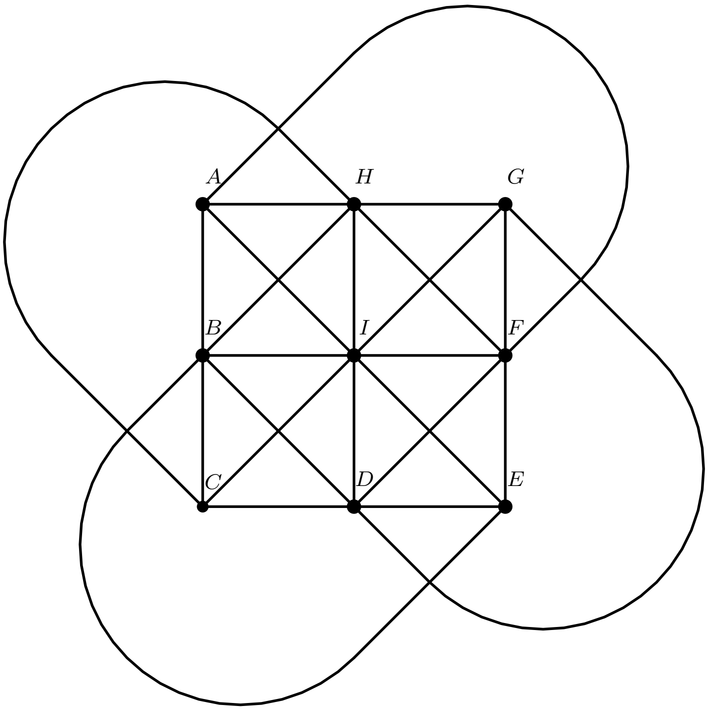

Figure 1.2: Affine plane over $\mathbb{F}_3$

1.3 Mutually Orthogonal Latin Squares

Given a set S of n elements, a Latin square L is a function $L : [n] \times [n] \to S$, i.e., an $n \times n$ array with elements in S, such that each element of S appears exactly once in each row and each column. For example,

$$\begin{array}{ccc} 1 & 2 & 3 \\ 3 & 1 & 2 \\ 2 & 3 & 1 \end{array}$$

is a 3×3 Latin square with $[3]$ as its set of elements. Sudoku puzzles are examples of 9×9 Latin squares, with the extra constraints that in each of the nine 3×3 subgrids all the symbols are distinct. The multiplication table of a finite group is also a Latin square.

Puzzle: From a standard deck of playing cards, take the aces, kings, queens and jacks. Can you arrange these 16 cards in a 4×4 grid such that in each row/column no two cards share the same suit or the same face value?

Solution: Here is one possible arrangement.

$$\begin{array}{cccc} J\heartsuit & Q\diamondsuit & K\spadesuit & A\clubsuit \\ Q\clubsuit & J\spadesuit & A\diamondsuit & K\heartsuit \\ K\diamondsuit & A\heartsuit & J\clubsuit & Q\spadesuit \\ A\spadesuit & K\clubsuit & Q\heartsuit & J\diamondsuit \end{array}$$

Definition 1.12. Let L_1 and L_2 be two Latin squares over the ground sets S_1, S_2, respectively. They are called orthogonal if for every $(x_1, x_2) \in S_1 \times S_2$ there exists a unique $(i, j) \in [n] \times [n]$ such that $L_1(i, j) = x_1$ and $L_2(i, j) = x_2$.

For example, the following are two orthogonal Latin squares of order 3.

$$\begin{array}{ccc} 1 & 2 & 3 \\ 3 & 1 & 2 \\ 2 & 3 & 1 \end{array} \qquad \begin{array}{ccc} 2 & 3 & 1 \\ 3 & 1 & 2 \\ 1 & 2 & 3 \end{array}$$

Can you find another Latin square of order 3 which is orthogonal to both of the Latin squares above? At most how many mutually orthogonal Latin squares (MOLS) of order n can there be, in terms of n? This is somewhat similar the question "how many pairwise orthogonal (non-zero) vectors can there by in $\mathbb{R}^n$?".

Lemma 1.13. *Let L be a Latin square with S as its ground set, and let $\pi : S \to S'$ be a bijection from S to another set S'. Then $\pi(L)$ defined by $\pi(L)(i,j) = \pi(L(i,j))$ is a Latin square on the ground set S'.*

Lemma 1.14. *Let L_1, L_2 be two orthogonal Latin squares on base sets S_1 and S_2. Let $\pi_1 : S_1 \to S_1'$, $\pi_2 : S_2 \to S_2'$ be two bijections. Then $\pi_1(L_1)$ and $\pi_1(L_2)$ are also orthogonal Latin squares.*

Proof. Let $(x,y) \in S_1' \times S_2'$. Let (i,j) be the unique index for which $L_1(i,j) = \pi_1^{-1}(x)$ and $L_2(i,j) = \pi_2^{-1}(y)$. Then (i,j) is also the unique index for which $\pi_1(L_1)(i,j) = x$ and $\pi_2(L_2)(i,j) = y$. □

Theorem 1.15. *For every n, there exist at most $n-1$ mutually orthogonal Latin squares of order n.*[2]

Proof. Let $L_1, \ldots, L_k$ be a set of MOLS of order n. Since changing the ground set does not affect orthogonality, we may assume that all of them have elements from $[n]$. Moreover, applying a permutation to the symbols of any of the Latin squares does not affect orthogonality of the system. So, we may assume that the first row of each L_i is $1\ 2 \cdots n$. Now say $k \geq n$. Since $L_i(2,1) \in \{2, \ldots, n\}$ for all i, by the pigeonhole principle there exist $i, j \in [k]$ such that $L_i(2,1) = L_j(2,1) = a$ for some $a \in \{2, \ldots, n\}$. But then (a,a) appears more than once in the superimposition of L_i and L_j, showing that they are not orthogonal to each other. □

We now have the natural extremal question, *is this bound sharp*? Can we find $n-1$ mutually orthogonal Latin squares for every n? The answer to the former is yes, while to the latter is no, in the following sense.

Theorem 1.16 (Bose 1938)**.** *There exist $n-1$ mutually orthogonal Latin squares of order n if and only if there exists an affine plane of order n.*

Proof. We will show the sharpness of the bound on MOLS whenever there exists an affine plane of order n. The other direction is left o the reader.

Let $\mathcal{A}$ be an affine plane of order n, and let $\mathcal{S}_1, \ldots, \mathcal{S}_{n+1}$ be its parallel classes of lines. Enumerate each $\mathcal{S}_i$ as $\{\ell_{i,1}, \ldots, \ell_{i,n}\}$. Denote the unique point in the intersection of $\ell_{n,i}$ and $\ell_{n+1,j}$ by (i,j). This gives us the domain $[n] \times [n]$ for the Latin squares we will construct. For $1 \leq k \leq n-1$, define the k'th Latin square L_k by $L_k(i,j) = m$ where m is the unique index for which a line $\ell_{k,m}$ of the k-th parallel class contains the point (i,j). We claim that this gives us $n-1$ MOLS of order n (all of them with the ground set equal to $[n]$).

[2] Compare this to the statement that in $\mathbb{R}^n$ there are at most n mutually orthogonal vectors.

Firstly, let L_k be one of the arrays. If there was a repetition in a row of L_k, then the line of $\mathcal{S}_k$ corresponding to this repeated entry intersects the line of $\mathcal{S}_n$ corresponding to the row in more than one points, which is a contradiction to the fact that an affine plane is a partial linear space. We get a similar contradiction for repetitions in columns, and hence L_k is a Latin square.

Now let L_k and $L_{k'}$ be two Latin squares, corresponding to the parallel classes $\mathcal{S}_k$ and $\mathcal{S}_{k'}$, respecitvely. Let $(a, b) \in [n] \times [n]$. Then the unique point of intersection of the lines $\ell_{k,a}$ and $\ell_{k',b}$ gives the unique coordinate (i, j) for which $L_k(i, j) = a$ and $L_{k'}(i, j) = b$. □

Example 1.10 gives us an affine plane of order q for each prime power q. This is equivalent to saying that we for all prime powers q, there exists $q - 1$ mutually orthogonal Latin squares of order q. Are there more orders for which we have an affine plane? This is one of the oldest and central open problems in finite geometry.

Conjecture 1.17 (Prime power conjecture)**.** *Every finite affine plane has order equal to a prime power.*

The only known restriction that is known on the order of an affine plane, that rules out infinitely many values, is as follows.

Theorem 1.18 (Bruck-Ryser 1949)**.** *If a finite affine plane of order n congruent to 1 or 2 modulo 4 exists, then n is the sum of two integral squares. In particular, any prime factor of n of the form $4k + 3$, for some $k \in \mathbb{N}$, has an even exponent in the prime factorisation of n.*

Corollary 1.19. *There exist no affine planes of order $6, 14, 21, 22, 30, 33 \ldots$.*

Beyond this, there is a celebrated result of Lam, Thiel and Swiercz from 1989 which showed that there is no affine plane of order 10. [3]

The non-existence of an affine plane of order 6 was already known before the Bruck-Ryser theorem. It was a famous problem of Euler to find two orthogonal Latin squares of order 6, who conjectured in 1782 that this can't be done.[4] Tarry proved in 1901 that Euler was correct, and his result in particular implies that there are no affine planes of order 6. Euler in-fact made a stronger conjecture, that for every $n \equiv 2 \pmod 4$, there does not exist a pair of orthogonal Latin squares. In the homework you'll see that for other values of n one can always construct orthogonal Latin squares of order n. This conjecture (whose only evidence was the trivial $n = 2$ case, and an incomplete proof of $n = 6$) was proven to be false in 1959 by Parker, Bose, and Shrikhande, a.k.a., the *Euler's Spoilers*[5].

Theorem 1.20 (Bose-Shrikhande-Parker 1960)**.** *For every $n > 6$, there exist two orthogonal Latin squares of order n.*

[3] See `http://web.thu.edu.tw/wang/www/emcc_Helly/Lam_finite_Proj_plane_order_10.pdf` for the story behind this result.

[4] It is known as the thirty-six officers puzzle.

[5] A phrase coined by the New York Times, who published the news of this result on their front page in a Sunday edition.

1.4 Projective planes

As we saw before, one of the main features of affine planes is the notion of parallelism. In some situations this feature can be a deficiency. This is “fixed” by the point-line geometries known as projective planes, which play a key role in Finite Geometry, Combinatorics, and Algebraic Geometry.

Definition 1.21. A projective plane is a linear space with the following properties.

(P1) Every two lines intersect in a unique point.

(P2) There exists a set of four points, no three of which are collinear.

Say we have an affine plane and let Π be set of parallel classes of the plane. For a line ℓ of the affine plane, let $\Pi(\ell)$ denote the unique parallel class it belongs to. For each parallel class $\mathcal{S} \in \Pi$, introduce a single new point $\sigma(\mathcal{S})$, with different parallel classes getting different points. Extend each of the original line ℓ by adding to it the new point $\sigma(\Pi(\ell))$. Introduce a new line $\ell_\infty = \{\sigma(\mathcal{S}) : \mathcal{S} \in \Pi\}$[6] consisting of all the new points that have been introduced. Then we claim that the point-line geometry that we get is a projective plane.

Proof. Let $(\mathcal{P}, \mathcal{L}, I)$ be the point line geometry that we get by the process above, so that the set of points of the affine planes is $\mathcal{P} \backslash \ell_\infty$ and the set of lines is $\{\ell \backslash \ell \cap \ell_\infty : \ell \in \mathcal{L} \backslash \{\ell_\infty\}\}$. Every two lines of $\mathcal{L}$ corresponding to affine lines in the same parallel class $\mathcal{S}$ intersect at the point $\sigma(\mathcal{S})$. Every two lines corresponding to two non-parallel affine lines intersect each other in their point of intersections in the affine plane and nowhere else, as the new points on them are not the same. Every line corresponding to a line of the affine plane intersects ℓ_∞ intersect in a single point, corresponding to the unique parallel class that the affine line belongs to. Therefore, $(\mathcal{P}, \mathcal{L}, I)$ is a linear space which satisfies (P1).

Now let x, y, z be three points of the affine plane which do not lie on a common line. Let ℓ be the line in $\mathcal{L}$ corresponding to the unique line through x parallel to the line yz, and let w be any other point on ℓ. Then no three points in the set x, y, z, w are incident to the same line. Therefore, (P2) is also satisfied. □

The projective plane that we get from the affine plane $\mathrm{AG}(2, F)$, is denoted by $\mathrm{PG}(2, F)$. The points of $\mathrm{PG}(2, F)$ are the points (x, y) of $\mathrm{AG}(2, F)$, the points corresponding to the lines of slope $m \in F$, which we denoted by (m), and the point (∞) corresponding to the parallel class of vertical lines. The lines are the lines of $\mathrm{AG}(2, F)$, extended with the point corresponding to the slope of the line, and a new line $\ell_\infty = \{(m) : m \in F \cup \{\infty\}\}$.

The finite projective plane of order q that we get when $F = \mathbb{F}_q$ is denoted by $\mathrm{PG}(2, q)$. The projective plane $\mathrm{PG}(2, 2)$ is known as the Fano plane, named after the Italian mathematician Gino Fano, and it turns out that it is the unique projective plane of order 2 (see Figure 1.4).

We can reverse the process of constructing a projective plane from an affine plane, as follows.

[6] We tend to denote this line as ℓ_∞ because of the case of real perspective geometry, where one can think of the parallel lines meeting each other at a far away point, that is “at infinity”.

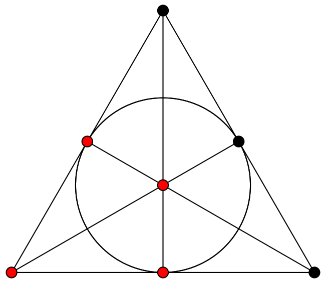

Figure 1.3: The Fano plane

Proposition 1.22. *For any projective plane, removing a line of the plane and all points on this line gives rise to an affine plane.*

Proof. Let ℓ be the line which is removed. We still have a linear space, so all that needs to be checked are the axioms (A1) and (A2). Let x be a point not on ℓ, and m a line distinct from ℓ such that $x \notin m$. Let $y = \ell \cap m$. Then every line through x intersects m in a point outside ℓ, except for the unique line xy, which is then the unique line through x parallel to m.

Let x, y, z, w be four points in the projective plane, such that no three of them are incident to the same line. If ℓ contains at most one of these four points, then (A2) is satisfied by the remaining three points. WLOG say $z, w \in \ell$. Let p be the point of intersection of the lines xz and yw. Then $p \notin \ell$, and $p \notin xy$. Therefore, p, x, y satisfy the condition in (A2), thus proving that we have an affine plane. □

We can use this relationship between affine and projective planes to obtain the following useful result on finite projective planes.

Proposition 1.23. *For every finite projective plane, there exists an integer $n \geq 2$, called the order of the plane, such that:*

(*a*) *Each line is incident with $n+1$ points.*

(*b*) *Each point is incident with $n+1$ lines.*

(*c*) *There are n^2+n+1 points.*

(*d*) *There are n^2+n+1 lines.*

Proof. Remove a line ℓ_∞ from the projective plane to get a finite affine plane. Let $n \geq 2$ be the order of the affine plane, from Theorem 1.9. Each affine line is incident with n points, and hence adding the point at infinity gives us $n+1$ points on each line in the projective plane. Through every affine point there are $n+1$ lines, and if we have a point at infinity then there are n affine lines through it, corresponding to a parallel class, and the line ℓ_∞. There are n^2 affine points and $n+1$ points at infinity. There are n^2+n affine lines and one ℓ_∞. □

By the discussion above the existence question of finite projective planes of a given order n is equivalent to the existence of an affine plane of order n, in particular, we again have the prime power conjecture which states that all finite projective plane must have their order equal to a prime power, and the Bruck-Ryser theorem.

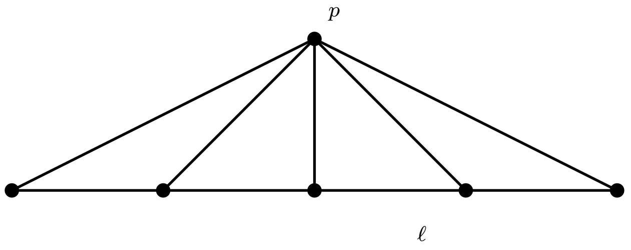

Figure 1.4: Degenerate projective plane

Projective planes play an important role finite geometry and combinatorics, as they often turn out to be extremal objects in many theorems. We will see one such examples below.

If we remove (P2) from the axioms of a projective plane, and add the condition that we have at least two lines, then it can be shown that the only new possibility we get is the following point line geometry, which we call the *degenerate* projective plane (see Figure 1.4). The points set contains a point p and the line set contains a line ℓ such that every point except p lies on ℓ and every line except ℓ passes through p.

Theorem 1.24 (De Bruijn-Erdős). *Let* $(\mathcal{P}, \mathcal{L}, \mathrm{I})$ *be a finite linear space, which has at least* 2 *lines. Then* $|\mathcal{P}| \leq |\mathcal{L}|$, *with equality if and only if the linear space is a (possibly degenerate) projective plane.*

Proof. (due to Conway/Motzkin) Note that in both degenerate projective planes and non-degenerate projective planes the number of points is equal to the number of lines. We will show that if $|\mathcal{L}| \leq |\mathcal{P}|$, then $|\mathcal{L}| = |\mathcal{P}|$ and we have a possibly degenerate projective plane. For this it suffices to show that every two lines intersect.

Let $b = |\mathcal{L}|$ and $v = |\mathcal{P}|$ and assume that $b \leq v$. For every $x \in \mathcal{P}$ let n_x denote the number of lines through x, and for every $\ell \in \mathcal{L}$ let m_ℓ denote the number of points on ℓ. Let $\mathcal{A} = (\mathcal{P} \times \mathcal{L}) \setminus I$ be the set of point-line pairs (x, ℓ) such that x is not incident with ℓ.[7]. Then from the fact that we have a linear space, we get $n_x \geq m_\ell$ for all $(x, \ell) \in \mathcal{A}$, as for every point y on the line ℓ we get a unique line xy through x. This implies

$$\frac{1}{b - n_x} \geq \frac{1}{b - m_\ell},$$

for all $(x, \ell) \in \mathcal{A}$. Note that since there are at least two lines, and each line contains at least two points, the set $\mathcal{A}$ is non-empty. By summing over all anti-flags, we get the inequality

$$\sum_{(x,\ell)\in\mathcal{A}} \frac{1}{b - n_x} \geq \sum_{(x,\ell)\in\mathcal{A}} \frac{1}{b - m_\ell},$$

which can be simplified to the following by noting that for each x there are exactly $b - n_x$ lines not incident with x and for each ℓ there are exactly $v - m_\ell$ points not incident with ℓ,

$$\sum_{x\in\mathcal{P}} \frac{b - n_x}{b - n_x} \geq \sum_{\ell\in\mathcal{L}} \frac{v - m_\ell}{b - m_\ell}.$$

[7] These are called anti-flags in the point-line geometry

The left hand side is equal to v whereas each term in the right hand side is at least v/b since we assumed $v \geq b$, and hence the right hand side is greater than or equal to v. We must have equality everywhere, which implies that $v = b$ and for all $(x, \ell) \in A$ we have $n_x = m_\ell$. In particular this means that any two pair of lines must intersect, as otherwise we'll get an $(x, \ell) \in A$ for which $n_x > m_\ell$. Therefore, the linear space must be a possibly degenerate finite projective plane. □

This result also has the following consequence on discrete sets of points in Euclidean spaces.

Corollary 1.25. *Let S be a set of n points in $\mathbb{R}^2$. Then S determines at least n lines, where a line is said to be determined by S if it contains at least two points of S.*

Proof. The set S along with the lines it determines forms a finite linear space. Moreover, it is known that no finite projective plane can be embedded in $\mathbb{R}^2$ using straight lines.[8] Therefore, we must have a degenerate projective plane in case of equality. □

1.5 Duality

Another reason to study projective planes is that there is a "duality" between the points and lines, in the following sense.

Definition 1.26. Let $\mathcal{S} = (\mathcal{P}, \mathcal{L}, I)$ be a point-line geometry such that every point is incident with at least two lines. Then the dual of of $\mathcal{S}$ is the point-line geometry $\mathcal{S}^D = (\mathcal{P}^D, \mathcal{L}^D, I^D)$ where $\mathcal{P}^D = \mathcal{L}$, $\mathcal{L}^D = \mathcal{P}$ and $(\ell, x) \in I^D$ if and only if $(x, \ell) \in I$.

Theorem 1.27. *Let π be a projective plane. Then π^D is also a projective plane.*

Proof. The first axiom is satisfied in π^D. For the second one, let x, y, z, w be four points in π such that no three of them are collinear. Then the lines xy, yz, zw and wx are four lines of π such that no three of them are concurrent, giving us the second axiom in π^D. □

Theorem 1.28 (Principle of duality). *If a theorem is valid for all projective planes, then the dual theorem obtained by interchanging the notions of point and line is also valid for all projective planes.*

Definition 1.29. An isomorphism between point-line geometries $\mathcal{S} = (\mathcal{P}, \mathcal{L}, I)$ and $\mathcal{S}' = (\mathcal{P}', 'L, I')$ is a pair of bijections $f_1 : \mathcal{P} \to \mathcal{P}'$ and $f_2 : \mathcal{L} \to \mathcal{L}'$ such that $(x, \ell) \in I$ if and only if $(f_1(x), f_2(\ell)) \in I'$. We say that $\mathcal{S}$ is isomorphic to $\mathcal{S}'$, denoted by $\mathcal{S} \cong \mathcal{S}'$ if there exists an isomorphism between $\mathcal{S}$ and $\mathcal{S}'$.

Definition 1.30. A point-line geometry $\mathcal{S}$ is called self-dual if $\mathcal{S} \cong \mathcal{S}^D$.

For example, the graph C_n gives a self-dual point-line geometry. We will soon see that $\mathrm{PG}(2, F)$ is self-dual for every field F.

[8] Try to prove it!

1.6 Projetive spaces from vector spaces

Projective geometry pre-dates the concept of a vector spaces, but since linear algebra has become ubiquitous in the world of mathematics it has become standard to define projective planes (and higher dimensional spaces) using vector spaces. However, this only defines a particular class of projective planes, whereas for higher dimensional spaces this definition is equivalent to the axiomatic definition of a projective space.

Theorem 1.31. *Let V be the 3-dimensional vector space over a field F. Let $\mathcal{P}$ be the set of 1-dimensional subspaces of V, and $\mathcal{L}$ the set of 2-dimensional spaces. Then $(\mathcal{P}, \mathcal{L}, I)$ with I as the subspace relation, is a projective plane.*

Proof. Any two distinct 1-dimensional subspaces are contained in a unique two dimensional subspace, which is the subspace spanned by them. Let U_1, U_2 be two 2-dimensional subspaces. Since $U_1 + U_2$ is a subspace of V, we have $\dim(U_1 + U_2) = \dim U_1 + \dim U_2 - \dim U_1 \cap U_2 \leq 3$, which implies that $\dim U_1 \cap U_2 \geq 1$ since $\dim U_1 = \dim U_2 = 2$. Moreover, since $U_1 \neq U_2$, we must have $\dim U_1 \cap U_2 < 2$, and hence $U_1 \cap U_2$ is the unique 1-dimensional subspace contained in both U_1 and U_2.

The 1-dimensional subspaces spanned by the vectors $(1, 0, 0)$, $(0, 1, 0)$, $(0, 0, 1)$ and $(1, 1, 1)$ satisfy (P2), since the matrix formed by taking these vectors as columns has rank 3 over any field F. □

Lemma 1.32. *The projective plane constructed above from the vector space F^3 is isomorphic to* $\mathrm{PG}(2, F)$.

Proof. We adopt a geometrical model of F^3 where the 1-dim subspaces are the lines through origin and the 2-dim subspaces are the planes through origin. Let π be the plane defined by the equation $z = 1$ in $AG(3, F)$. We know that π is isomorphic to $\mathrm{AG}(2, F)$. Every 1-dim subspace, except for those lying in the plane $z = 0$, intersect π in a unique point. Each 1-dim subspace contained in $z = 0$ naturally corresponds to a parallel class in π, and hence to the point at infinity in the projective completion of $\mathrm{PG}(2, F)$, as it gives a direction for the lines in π. Every 2-dim subspace of F^3, except for the space $z = 0$, intersects π in a line.

□

Corollary 1.33. $\mathrm{PG}(2, F)$ *is self-dual.*

Proof. Any isomorphic between the vector space F^3 and its dual gives us an isomorphism between $\mathrm{PG}(2, F)$ and $\mathrm{PG}(2, F)^D$. □

The projective planes arising from 3-dimensional vector spaces immediately suggest the following higher dimensional point-line geometry known as a projective space.

Definition 1.34. Let V be an $n + 1$ dimensional vector space over a field F. Then the n dimensional projective space over F, $\mathrm{PG}(n, F)$ is the point-line geometry where the points are the 1-dimensional subspaces of V and the lines are the 2-dimensional subspaces of V.[9]

We also have an axiomatic definition of projective spaces.

[9] Note that this n has nothing to do with the earlier n which was the order of a finite projective plane.

Definition 1.35. An abstract projective space is a linear space satisfying the following properties:

(PS1) Every line is incident with at least three points.

(PS2) Let ℓ_1, ℓ_2 be two lines through a point z, and for $i \in \{1, 2\}$ let x_i, y_i be two distinct points on $\ell_i \setminus \{z\}$. Then the lines x_1x_2 and y_1y_2 meet in a point.

Starting from this definition, one can build up the notion of subspaces, basis, and dimension, without relying on any linear algebra. See for example the book by Beutelspacher and Rosenbaum, "Projective Geometry: From Foundations to Applications". We will not do so here, and we justify that by referring to the following important result.

Theorem 1.36 (Veblen-Young, 1910)**.** *An abstract projective space is either a point, a line, a projective plane, or the n dimensional projective space over a (skew) field[10] F, for some $n \geq 3$.*

In particular, in an abstract projective space if we assume that there exist two disjoint lines then it must be isomorphic to $\mathrm{PG}(n, F)$ for some (skew) field F and $n \geq 3$. So, all we have to deal with are the projective planes, which we defined axiomatically, and projective spaces over vector spaces of dimension at least 4.

Even though the point-line incidences are sufficient to describe higher dimensional projective spaces, they have more structure than just points and lines. A k-dimensional subspace of $\mathrm{PG}(n, F)$ is the set of points, i.e., 1-dimensional subspaces of the underlying vector space F^{n+1}, contained in a $(k+1)$-dimensional subspace of F^{n+1}. The $0, 1, 2, 3, \ldots, n-1$ dimensional subspaces of $\mathrm{PG}(n, F)$ are known as *points*, *lines*, *planes*, *solids*, ..., *hyperplanes*, respectively, and the correspond to the $1, 2, 3, 4, \ldots, n$ dimensional subspace of the underlying vector space F^{n+1}.

If we remove a hyperplane from $\mathrm{PG}(n, F)$ and all points incident with it, then we get the n-dimensional affine space $\mathrm{AG}(n, F)$. The point-line geometry $\mathrm{AG}(n, F)$ can also be defined without any reference to the projective spaces as the point-line geometry whose points are the elements of the vector space F^n and lines are the cosets of the 1-dimensional subspaces. The k-dimensional subspaces are then the cosets of the k-dimensional subspaces of the vector space F^n. The equivalence relation of parallelism defined on the hyperplanes, i.e., the $n-1$ dimensional subspaces of $\mathrm{AG}(n, F)$, gives us the projective space $\mathrm{PG}(n, F)$ by adding a hyperplane at infinity, in the same was as it did for $n = 2$. When $F = \mathbb{F}_q$, we denote these geometries by $\mathrm{AG}(n, q)$ and $\mathrm{PG}(n, q)$, respectively.

1.7 Desarguesian planes

We can characterise the projective planes $\mathrm{PG}(2, F)$ among the class of abstract projective planes via the following theorem.

Theorem 1.37 (Desargues' Theorem)**.** *Let F be a field. Let $a_1b_1c_1$ and $a_2b_2c_2$ be two triangles in $\mathrm{PG}(2, F)$ such that the lines a_1a_2, b_1b_2 and c_1c_2 all pass through a common point. Then the points $x = b_1c_1 \cap b_2c_2$, $y = c_1a_1 \cap c_2a_2$ and $z = a_1b_1 \cap a_2b_2$ are collinear.*

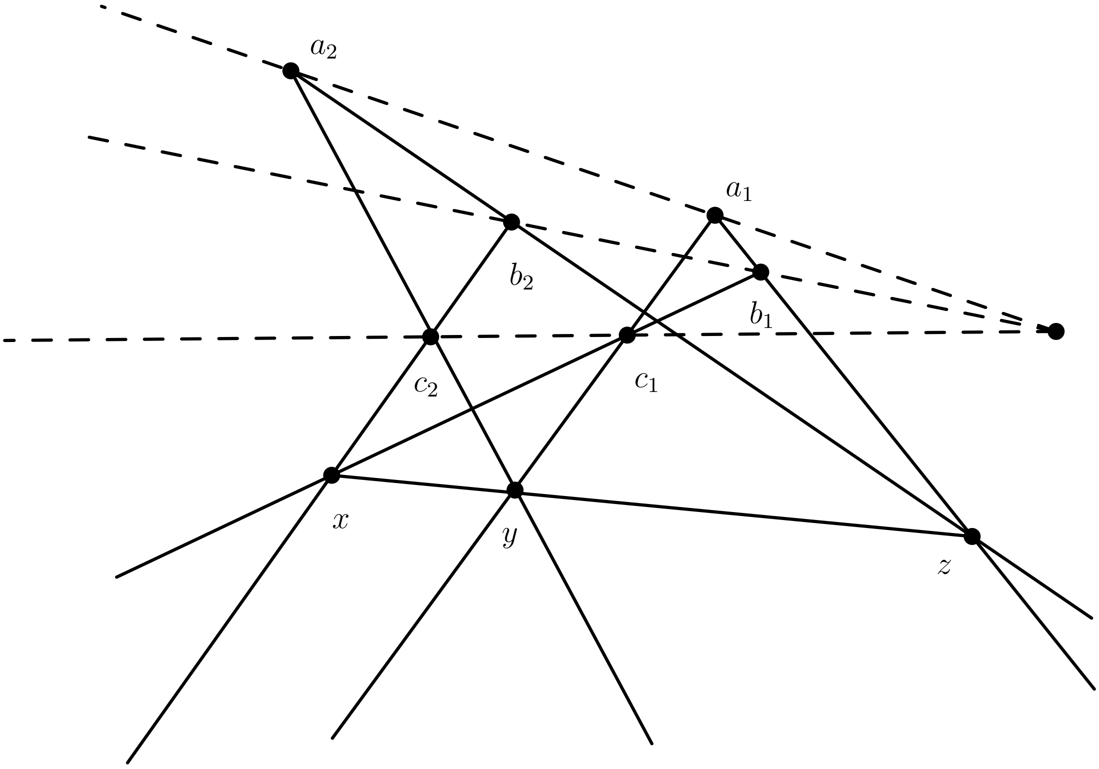

Figure 1.5: Desargues' theorem

Theorem 1.38. *Any projective plane in which the Desargues' theorem holds true is isomorphic to* $\mathrm{PG}(2, F)$ *for some division ring* F.

Remark 1.39. It is a nice exercise to prove that in any abstract projective space that contains two disjoint lines, Desargues' theorem is true. This then implies Theorem 1.36.

It is a classical result of Wedderburn, from 1907, which then states that any finite division ring is a finite field, giving us the following.

Corollary 1.40. *Any finite projective plane in which Desargues' theorem holds is isomorphic to* $\mathrm{PG}(2, q)$ *for some prime power* q.

Therefore, the planes $\mathrm{PG}(2, q)$ defined over the finite field $\mathbb{F}_q$ are known as the Desarguesian projective planes. There are many known families of finite projective planes which are non-Desarguesian. In fact there exists an infinite sequence of n for which the number of non-isomorphic projective planes of order n grows with at least a super polynomial function of n.

While the constructions of finite non-Desarguesian planes are fairly involved, and you will see some once we have developed more theory, it is fairly easy to construct a non-Desarguesian projective plane over $\mathbb{R}$: Let $\mathcal{P} = \mathbb{R}^2$, and let $\mathcal{L}$ consist of the same vertical and positively sloped lines of the Euclidean plane, and the negatively sloped lines replaced by the following, $\{(x, y) : y = mx + c \text{ if } x \geq 0 \text{ and } y = \frac{1}{2}mx + c \text{ if } x \leq 0\}$. Then it can be checked that $(\mathcal{P}, \mathcal{L})$ is a non-Desarguesian affine plane[11], and its projective completion

[10] Skew fields, a.k.a. division rings, satisfy all axioms of a field except possibly the commutativity of multiplication.

[11] What is the corresponding statement of Desargues' theorem for affine planes?

gives us a non-Desarguesian projective plane.

1.8 Exercises

1. Prove that up-to isomorphism there are unique projective plane of order 2 and 3.[12]

2. Prove that a degenerate projective plane is self-dual.

3. (a) Let $\mathcal{S}$ be a linear space on n^2 points where each line is incident with exactly n points, for some integer $n \geq 2$. Prove that $\mathcal{S}$ is an affine plane.

   (b) Prove that if a system of $n-1$ mutually orthogonal Latin squares of order n exists, then there exists an affine plane of order n.

4. (a) Given two Latin squares of orders m and n, construct a Latin square of order mn.

   (b) Given two orthogonal Latin squares of order m, and two orthogonal Latin squares of order n, construct two orthogonal Latin squares of order mn.

   (c) Prove that for every $n \not\equiv 2 \pmod 4$, there exist two orthogonal Latin squares of order n.

   (Hint: In part (a) try to construct the new Latin square over the ground set $S_1 \times S_2$, where S_1 is the ground set of the first Latin square and S_2 of the second Latin square.)

5. For $n \geq 2$ an integer, define a graph G_n whose vertex set is equal to the set of all $n \times n$ Latin squares on the ground set $[n]$, and two Latin squares are adjacent if they are orthogonal to each other. Prove that for all n, the chromatic number $\chi(G_n)$ is at most $n-1$. For what values of n is the chromatic number equal to $n-1$?

6. A subplane of a projective plane $(\mathcal{P}, \mathcal{L}, \mathrm{I})$ is a projective plane $(\mathcal{P}', \mathcal{L}', \mathrm{I}')$ such that $\mathcal{P}' \subsetneq \mathcal{P}$, $\mathcal{L}' \subsetneq \mathcal{L}$ and $\mathrm{I}' = \mathrm{I} \cap (\mathcal{P}' \times \mathcal{L}')$.

   (a) Prove that if a projective plane has order n and m is the order of a subplane of the projective plane, then $n \geq m^2$.

   (b) Prove that the bound above is tight by constructing a projective plane of order m^2 and a subplane in it that has order m, for some m (preferably an infinite family of m's).

   (c) Show that if $n > m^2$, then $n \geq m^2 + m$.

7. For a prime power q, and integers $0 \leq k \leq n$, the Gaussian binomial coefficient is defined as
$$\begin{bmatrix} n \\ k \end{bmatrix}_q = \frac{(q^n - 1)(q^{n-1} - 1) \cdots (q^{n-k+1} - 1)}{(q^k - 1)(q^{k-1} - 1) \cdots (q - 1)}.$$

   (a) Prove that the number of k-dimensional subspaces of $\mathrm{PG}(n, q)$ is equal to $\begin{bmatrix} n+1 \\ k+1 \end{bmatrix}_q$.

   (b) Give combinatorial proofs of the following identities:

[12] It has been conjecture that for every prime p, there is a unique projective plane of order p.

(i)

$$\begin{bmatrix} n \\ k \end{bmatrix}_q = \begin{bmatrix} n-1 \\ k-1 \end{bmatrix}_q + q^k \begin{bmatrix} n-1 \\ k \end{bmatrix}_q .$$

(ii)

$$\begin{bmatrix} n \\ k \end{bmatrix}_q = \begin{bmatrix} n \\ n-k \end{bmatrix}_q .$$

(iii)

$$\begin{bmatrix} n \\ k \end{bmatrix}_q = q^{n-k} \begin{bmatrix} n-1 \\ k-1 \end{bmatrix}_q + \begin{bmatrix} n-1 \\ k \end{bmatrix}_q .$$

8. For integers $1 \leq k \leq n$, prove the points of $\mathrm{PG}(n-1, q)$ can be partitioned into a collection of pairwise disjoint $k-1$ dimensional subspaces if and only if k divides n.

9. Determine the number of k-dimensional subspaces in $\mathrm{AG}(n, q)$, for $0 \leq k \leq n$.

10. Prove the following claims for $\mathrm{PG}(n, F)$, $n \geq 2$:

    (a) Let H be a hyperplane and S a k-dimensional subspace, for $k < n-1$. Then either S is contained in H or S intersects H in a $(k-1)$-dimensional subspace.

    (b) Any subspace that intersects every line in at least one point must be a hyperplane.

    (c) Let ℓ_1, ℓ_2 be two lines through a point z, and for $i \in \{1, 2\}$ let x_i, y_i be two distinct points on $\ell_i \setminus \{z\}$. Then the lines x_1x_2 and y_1y_2 meet in a point.[13]

[13] This shows that indeed these are abstract projective spaces.

2 Conics and Ovals

2.1 Coordinates and Hypersurfaces

Let $\mathrm{PG}(n, F)$ be the n-dimensional projective space with V as the underlying n-dimensional vector sapce over the field F. After picking a basis of V, we can give coordinates to each element of V. These coordinates can also be used for the points of $\mathrm{PG}(n, F)$, i.e., the 1-dimensional subspaces of V. A 1-dimensional subspace is determined by a non-zero vector, upto scalar multiplication by a non-zero element of F. Thus, a point P of $\mathrm{PG}(n, F)$ corresponding to $\langle v\rangle$ for some $v \in V \setminus \{0\}$ is given the coordinates $(x_0, \ldots, x_n)$ of v, where we assume that $(x_0, \ldots, x_n) = \lambda(x_0, \ldots, x_n)$ for all $\lambda \in F^\star$. In other words, we put an equivalence relation on the non-zero elements of F^{n+1} given by $(x_0, \ldots, x_n) \sim (x'_0, \ldots, x_n$ if $\exists\lambda \in F^\star$ such that $(x_0, \ldots, x_n) = \lambda(x'_0, \ldots, x'_n)$, and then the coordinates of the points of $PG(n, F)$ correspond to the equivalence classes of this relation.

If we *right-normalise*, that is, make the first non-zero coordinate from the right equal to 1 by dividing every coordinate by that element, then we get a unique representative of each of these equivalence classes, and hence a unique coordinate for every point of $\mathrm{PG}(n, F)$. For example, the right normalised coordinates of the points of $\mathrm{PG}(2, F)$ are

$$\{(\lambda, \mu, 1) : \lambda, \mu \in F\} \cup \{(\lambda, 1, 0) : \lambda \in F\} \cup \{(1, 0, 0)\}.$$

We can think of the points with $z = 0$ as the points lying at infinity, and $\{(\lambda, \mu, 1) : \lambda, \mu \in F\}$ as the points of the affine plane we get by removing the line $z = 0$. The point $(\lambda, 1, 0)$ denotes the point at infinity where the lines $x = \lambda y + c$ meet the line at infinity, whereas the point $(1, 0, 0)$ denotes the point where all horizontal lines of the affine plane meet. There are other ways to give unique coordinates to each point as well. For example, we could have right normalised the points with $z \neq 0$ and left normalised the points with $z = 0$, to get the points at infinity corresponding to the slopes of the affine lines.

Once we have coordinates, we can consider sets of points defined using algebraic equations.

Definition 2.1. Let d be a positive integer. A homogenous polynomial of degree d in variables $x_0, \ldots, x_n$ is a non-zero polynomial $\phi \in F[x_0, \ldots, x_n]$ whose each term is a monomial of the form $ax_0^{e_0}x_1^{e_1}\cdots x_n^{e_n}$ with $\sum e_i = d$ and $a \in F$.

For example, $x_0^2x_1+x_2x_3^2$ is a homogeneous polynomial of degree 3 in variables x_0, x_1, x_2, x_3 whereas $x_0x_1^2 + x_2^3 + x_3$ is not a homogeneous polynomial.

In $\mathrm{PG}(n, F)$, identifying the non-zero solutions to a polynomial equation $\phi(x) = 0$ with points of $\mathrm{PG}(n, F)$ make sense if ϕ is a homogeneous polynomial in $F[x_0, \ldots, x_n]$ since for such a polynomial we have $\phi(\lambda x) = \lambda^{\deg\phi}\phi(x)$, and hence the condition $\phi(u) = 0$ is well defined on the equivalence class of the non-zero scalar multiples of u.

Definition 2.2. A hypersurface of degree d in $\mathrm{PG}(n, F)$ is the set of points satisfying a degree d homogeneous polynomial equation. A hypersurface in $\mathrm{PG}(2, F)$ is known as a curve.

In particular, every line of $\mathrm{PG}(2, F)$ is a curve as it is given by the points satisfying an equation of the form $ax + by + cz = 0$, for some $(a, b, c) \in F^3 \setminus \{0\}$ (recall that a line is just a 2-dimensional subspace of F^3.)

Remark 2.3. A *projective algebraic variety* is defined as the set of common zeros of a system of homogeneous polynomial equations. In particular, every k-dimensional subspace of $\mathrm{PG}(n, F)$ is a variety. What are the homogeneous polynomial equations that define the set of points of such a subspace?

2.2 Conics

Definition 2.4. Let $\phi \in F[x, y, z]$ be a homogeneous polynomial of degree 2. Then the set $C = \{x \in \mathrm{PG}(2, F) : \phi(x) = 0\}$ is called a *conic*. A *conic* is called irreducible, or nondegenerate, if the polynomial ϕ cannot be written as a product of two degree 1 polynomials over F and every extension of F.

Example 2.5. 1. In $\mathrm{PG}(2, \mathbb{R})$, the conic C defined by the equation $x^2 + y^2 + z^2 + 2xy = 0$ is degenerate, since the polynomial is equal to $(x + y)^2 + z^2$, which splits into two linear factors $(x + y + iz)(x + y - iz)$ over the extension $\mathbb{C}$ of $\mathbb{R}$.

2. The conic C defined by $y^2 - zx = 0$ in $\mathrm{PG}(2, F)$ is non-degenerate and it can be parametrically written as $C = \{(\lambda^2, \lambda, 1) : \lambda \in F\} \cup (1, 0, 0)$. Therefore, if $F = \mathbb{F}_q$, then C has $q + 1$ points.

3. The parabolas, hyperbolas, and ellipses given by $y^2 = 4ax$, $\frac{x^2}{a^2} - \frac{y^2}{b^2} = 1$ and $\frac{x^2}{a^2} + \frac{y^2}{b^2} = 1$, in $\mathbb{R}^2$, can be *homogenized* to give us the following conics in $\mathrm{PG}(2, \mathbb{R})$: $y^2 = 4axz$, $\frac{x^2}{a^2} - \frac{y^2}{b^2} = z^2$ and $\frac{x^2}{a^2} + \frac{y^2}{b^2} = z^2$. How many points at infinity, that is, the points with $z = 0$, do each of these conics have?

Our aim in this section is to understand conics in Desarguesian projective planes and then look at their combinatorial generalizations to arbitrary finite projective planes.

Lemma 2.6. *Let C be a conic in* $\mathrm{PG}(2, F)$ *and ℓ a line. Then either all points of ℓ are contained in C or ℓ intersects C in at most two points.*

Proof. Let C be defined by $\phi(x, y, z) = ax^2 + by^2 + cz^2 + fyz + gxz + hxy = 0$. Say ℓ is not contained in C, and let $P = (x_1, y_1, z_1) \in \ell \setminus C$. Then $\phi(x_1, y_1, z_1) \neq 0$. Let $Q = (x_0, y_0, z_0)$ be any other point of ℓ. Then an arbitrary point of $\ell \setminus \{P\}$ is given by $Q + \lambda P = (x_1, y_1, z_1) + \lambda(x_0, y_0, z_0)$ for some $\lambda \in F$ (where $\lambda = 0$ corresponds to the point Q). This point in $\ell \setminus \{P\}$ is contained in C if and only if

$$\begin{aligned}\phi(Q + \lambda P) = \;& \phi(x_0, y_0, z_0) \\ & + ((2ax_0 + gz_0 + hy_0)x_1 + (2by_0 + fz_0 + hx_0)y_1 + (2cz_0 + fy_0 + gx_0)z_1)\,\lambda \\ & + \phi(x_1, y_1, z_1)\lambda^2 = 0.\end{aligned} \tag{2.1}$$

Since this is a quadratic equation in λ with a non-zero leading term, it has at most two zeros in F. □

Remark 2.7. Note that the quadratic equation in the proof above always has zeros in a degree 2 extension of F; in fact, it has exactly two zeros counted with multiplicity. Moreover, when $F = \mathbb{F}_q$, there is a unique quadratic extension, $\mathbb{F}_{q^2}$, and so every line not contained in C intersects the conic C' obtained by using the same equation as that of C but over $\mathbb{F}_q^2$, in two points of $\mathrm{PG}(2, q^2)$.

If E is a quadratic extension of a field F, obtained by adjoining the square root of a non-square k in F, then every element of E can be written uniquely as $\lambda + \mu\sqrt{k}$, with $\lambda, \mu \in F$.[1] The conjugate $\overline{a}$ of an element $a \in E$, with $a = \lambda + \mu\sqrt{k}$ is defined as $\overline{a} = \lambda - \mu\sqrt{k}$. Note that $a\overline{a} \in F$ and $a + \overline{a} \in F$ for all $a \in E$. For a line ℓ in $\mathrm{PG}(2, E)$ given by the equation $ax + by + cz = 0$, we define $\overline{\ell}$ to be the line given by the equation $\overline{a}x + \overline{b}y + \overline{c}z = 0$, which is equivalent to taking the $\overline{\ell}$ to be the set of conjugates of points on ℓ. Note that if $\ell \neq \overline{\ell}$, then the point $\ell \cap \overline{\ell}$ always lies in the subplane $\mathrm{PG}(2, F)$ of $\mathrm{PG}(2, E)$.

Lemma 2.8. *Let C be a conic in $\mathrm{PG}(2, F)$ defined by $\phi \in F[x, y, z]$. Then ϕ is reducible if and only if C contains all the points of a line in F, or in a quadratic extension E of F. Moreover, if it's reducible and one of the linear factors does not have all of its coefficients in F then the two linear factors define conjugate lines in $\mathrm{PG}(2, E)$.*

Proof. Left to the reader. □

If $Q = (x_0, y_0, z_0)$ is a point on the conic and $P = (x_1, y_1, z_1)$ is an arbitrary point outside the conic then Equation 2.1 is of the form $r\lambda + s\lambda^2 = 0$, where $s \neq 0$. Therefore it has roots $\lambda = 0$ and $\lambda = -r/s$. In particular, $\lambda = 0$ is a double root, i.e. the line PQ intersects C in only one point, if and only if $r = 0$, where

$$r = (2ax_0 + gz_0 + hy_0)x_1 + (2by_0 + fz_0 + hx_0)y_1 + (2cz_0 + fy_0 + gx_0)z_1 = 0.$$

It then follows that if the three equations

$$\frac{\partial\phi}{\partial x} := 2ax + hy + gz = 0$$

$$\frac{\partial\phi}{\partial y} := hx + 2by + fz = 0$$

$$\frac{\partial\phi}{\partial x} := gx + fy + 2cz = 0$$

are not simultaneously satisfied by the point (x_0, y_0, z_0) of the conic, then the line given by

$$(2ax_0 + gz_0 + hy_0)x + (2by_0 + fz_0 + hx_0)y + (2cz_0 + fy_0 + gx_0)z = 0$$

is the unique line through $Q = (x_0, y_0, z_0)$ that intersect C in only Q. It is called the **tangent** line at Q. [2]

Definition 2.9. Let C be a curve in $\mathrm{PG}(2, F)$ defined by the equation $\phi = 0$. Then a point P on C is called a **singular** point of C if

$$\frac{\partial\phi}{\partial x} = \frac{\partial\phi}{\partial y} = \frac{\partial\phi}{\partial z} = 0$$

[1] Think of how we get $\mathbb{C}$ from $\mathbb{R}$.

[2] The partial derivatives above are purely formal algebraic operations, where we define the derivative of x^r to be rx^{r-1} and extend this linearly to define the derivative of any univariate polynomial.

when evaluated at P. The curve is called **singular** if it has a singular point on it, and **non-singular** otherwise.

For conics we have the following nice characterisation of irreducibility, which we will use to count the total number of points on a conic in $\mathrm{PG}(2, q)$.

Proposition 2.10. *A conic in* $\mathrm{PG}(2, F)$ *is irreducible, if and only if it is non-singular.*

Proof. Let C be the conic defined by the polynomial $\phi \in F[x, y, z]$. Say ϕ is reducible and $\phi = (ax + by + cz)(a'x + b'y + c'z)$, possibly over a quadratic extension E of F. If $(a, b, c) \in E^3 \setminus F^3$, then $a' = \lambda\bar{a}, b' = \lambda\bar{b}$ and $c' = \lambda\bar{c}$, for some $\lambda \in F^\star$ (by Lemma 2.8). Let P be the point of intersection of the lines ℓ and m defined by the zero sets of these two linear factors of E (take P to be any point on the line if $\ell = m$). Then P is in $\mathrm{PG}(2, F)$ and it can be easily checked that all the partial derivatives vanish at P, making C singular.

Say $\phi = ax^2 + by^2 + cz^2 + fyz + gxz + hxy$ and say C is singular. After a change of basis, we can assume that the singular point has coordinates $(1, 0, 0)$. Since $\phi(1, 0, 0) = 0$ and all the three partial derivatives vanish at this point, we get $a = g = h = 0$, and hence $\phi = by^2 + cz^2 + fyz$. This has linear factors

$$(\sqrt{b})y + \left(\frac{f - \sqrt{f^2 - 4bc}}{2\sqrt{b}}\right) z \text{ and } (\sqrt{b})y + \left(\frac{f + \sqrt{f^2 - 4bc}}{2\sqrt{b}}\right) z,$$

over an extension of F where b and $f^2 - 4bc$ are perfect squares. Therefore, C is non-degenerate. □

Example 2.11. Using this proposition, we can easily check that the conic defined by $x^2 + y^2 + z^2 = 0$ is non-degenerate whenever the characteristic of the field is not equal to 2, as then the only common solutions that the three equations from the partial derivatives have is $(0, 0, 0)$. The polynomial $xy + yz + zx$ gives us an irreducible conic over fields of characteristic 2.

Lemma 2.12. *If an irreducible conic* C *of* $\mathrm{PG}(2, q)$ *contains a point, then it has exactly* $q + 1$ *points in it.*

Proof. Since it is irreducible, it is non-singular. Say P is a point of C, and let ℓ be the unique line through P which is tangent to C. Each of the q lines through P, other than the line ℓ, intersect C in exactly one other point, and conversely every point of the conic other than P defines one of these lines. Therefore, C has exactly $q + 1$ points in it. □

Remark 2.13. Note that in this proof we only used the fact that every irreducible conic is also non-singular; the other direction is not needed.

It will follow from the Chevalley-Warning theorem, which we will see in the next chapter, that every conic (whether irreducible or not) in $\mathrm{PG}(2, q)$ contains a point, i.e., there are no "point-less" conics over a finite field.[3]

[3] In $\mathrm{PG}(2, R)$ the conic given by $x^2 + y^2 + z^2 = 0$ is point-less. Note that it is also irreducible.

2.3 Ovals and Hyperovals

Definition 2.14. An *arc* in a projective plane is a set of points in which no three points are collinear. An arc of size k is called a k-arc.

Clearly, any subset of an irreducible conic is an arc, as otherwise the conic will contain a line. The conic itself gives us a $(q+1)$-arc in $\mathrm{PG}(2,q)$. Is this the largest possible arc we can have in a finite projective plane of order q?

Theorem 2.15. *An arc in a projective plane of order n has at most $n+2$ points in it. Moreover, if n is odd then it has at most $n+1$ points.*

Proof. Let S be an arc in a projective plane π of order n. Let $x \in S$. For each of the $n+1$ lines of π through x, we have at most one point of $S \setminus x$ contained in the line, and thee must be all points of S since every point $y \in S \setminus \{x\}$ defines the line yx through x. Therefore, $|S \setminus \{x\}| \leq n+1$, and hence $|S| \leq n+2$.

Let S be an arc of size $n+2$. Then every line must intersect S in 0 or 2 points, as if we have a line intersecting S in exactly one point, then looking at the $n+1$ lines through this point we will get $|S| \leq n+1$. Now consider the lines through a point $y \notin S$. They give rise to a partition of S into pairs, showing that $n+2$ is even, and hence n is even. Therefore, for n odd, every arc has size at most $n+1$. □

Definition 2.16. An arc of size $n+1$ in a projective plane of order n is called an *oval*, and an arc of size $n+2$ is called a *hyperoval*.

From the proof of Theorem 2.15 it follows that through every point of an oval there is a unique line that intersects the oval in a single point. This gives us an alternate proof of the fact that every point of a nondegenerate conic in $\mathrm{PG}(2,q)$ is contained in a unique tangent, if we can show that an nondegenerate conic has $q+1$ points in it, without using the fact that each point is in a unique tangent.

The nondegenerate conics in Desarguesian projective planes show that ovals exist. But what about hyperovals? Can you find a hyperoval in $\mathrm{PG}(2,q)$ for any even prime power q? It turns out that whenever there exists an oval in a an arbitrary projective plane of even order, there is also a hyperoval. In fact, something stronger holds true.

Theorem 2.17. *Let π be a projective plane of even order n, and $\mathcal{O}$ an oval in π. Then there is a unique hyperoval in π that contains $\mathcal{O}$.*

Proof. We are looking for a unique point $p \notin \mathcal{O}$ such that all lines through p are tangents to the oval, i.e., p must be the intersection of all $n+1$ tangents. For $0 \leq i \leq n+1$, let e_i denote the number of points in the plane, outside $\mathcal{O}$, that are incident with exactly i tangents. We want to show that $e_{n+1} = 1$.[4]

Noting that each of the n^2 points outside can be incident to at most $n+1$ tangents of $\mathcal{O}$, we have

$$\sum_{i=0}^{n+1} e_i = n^2.$$

[4] In fact, showing that $e_{n+1} > 0$ would suffice.

Double counting the pairs (x, ℓ) where x is a point of π outside $\mathcal{O}$ and ℓ is a tangent of $\mathcal{O}$ through x, we get

$$\sum_{i=1}^{n+1} ie_i = (n+1)n.$$

Counting the triples (x, ℓ, m) where ℓ, m are two distinct tangents of $\mathcal{O}$ and x is their point of intersection, we get

$$\sum_{i=2}^{n+1} i(i-1)e_i = (n+1)n.$$

Now observe that since $n+1$ is odd, each point outside must be incident to at least one tangent, since otherwise we will get a partition of $\mathcal{O}$ into pairs using the lines through this point. Thus we have $e_0 = 0$, and $\sum_{i=1}^{n+1} e_i = n^2$.

Consider the sum

$$S = \sum_{i=1}^{n+1} (i-1)(n+1-i)e_i,$$

which is always ≥ 0 with equality if and only if $e_i = 0$ for all $i \in \{2, \ldots, n\}$. We can simplify it to $S = \sum(n+1)ie_i - \sum(n+1)e_i - \sum i(i-1)e_i$. Using the equations above we get $S = (n+1)^2 n - n^2(n+1) - (n+1)n = 0$.

We can now compute e_{n+1} from the third equation where we get $(n+1)n \cdot e_{n+1} = (n+1)n$, and hence $e_{n+1} = 1$. □

Remark 2.18. The counting trick above is known as the *variance trick*, or the *second moment method*. Over the years it has seen many applications, especially when studying "regular" combinatorial structures.

Definition 2.19. Given a set S of points in a projective plane, a line ℓ is called a *tangent* of S if $|\ell \cap S| = 1$, a *external* if $|\ell \cap S| = 0$ and a secant if $|\ell \cap S| > 1$.

Proposition 2.20. *Let π be a projective plane of order n, $\mathcal{O}$ an oval in π and $\mathcal{H}$ a hyperoval in $\mathcal{P}$.*

(1) *$\mathcal{O}$ has $n+1$ tangents, $\binom{n+1}{2}$ secants and $\binom{n}{2}$ externals.*

(2) *$\mathcal{H}$ has $\binom{n+2}{2}$ secants and $\binom{n}{2}$ externals.*

Proof. In an oval there is a unique tangent through each of the $n+1$ points, a unique secant through every pair of points on the oval, and the rest are external. In a hyperoval there are no tangent lines. □

The fact that there are $n+1$ tangents, and $n+1$ points in an oval, might suggest some sort of duality between these objects. Indeed, for n odd we have the following result.

Theorem 2.21. *Let $\mathcal{O}$ be an oval in a projective plane of odd order, and $\mathcal{O}^*$ the set of tangents to $\mathcal{O}$. Then $\mathcal{O}^*$ forms an oval in the dual projective plane.*

Proof. There are $n+1$ tangents of $\mathcal{O}$, one through each point, giving us $n+1$ points in the dual projective plane. All we need to show is that no three tangents are concurrent.

As above, we get the following three equations, where e_i denotes the number of points outside $\mathcal{O}$ which are incident to exactly i tangents of $\mathcal{O}$.

$$\sum_{i=0}^{n+1} e_i = n^2.$$

$$\sum_{i=1}^{n+1} ie_i = (n+1)n.$$

$$\sum_{i=2}^{n+1} i(i-1)e_i = (n+1)n.$$

We want to show that $e_i = 0$ for all $i \geq 3$. Note that since n is odd, no point outside $\mathcal{O}$ can be incident to a unique tangent of $\mathcal{O}$, giving us $e_1 = 0$. Now consider the sum $S = \sum_{i=3}^{n+1} i(i-2)e_i = \sum_{i=0}^{n+1} i(i-2)e_i = \sum_{i=0}^{n+1} i(i-1)e_i - \sum_{i=0}^{n+1} ie_i = (n+1)n-(n+1)n = 0$. Therefore, $e_i = 0$ for all $i \geq 3$. □

Corollary 2.22. *Let $\mathcal{O}$ be an oval in a projective plane π of odd order. Then the points of π are partitioned into the following three types:*

(1) *On points: the point through which there is a unique tangent to $\mathcal{O}$, that is, the points lying on $\mathcal{O}$; there are $n+1$ one such points.*

(2) *Exterior points: the points through which there are exactly 2 tangents to $\mathcal{O}$; there are $\binom{n+1}{2}$ such points*

(3) *Interior points: the points through which there are no tangents of $\mathcal{O}$; there are $\binom{n}{2}$ such points.*

We saw that we can construct ovals in finite Desarguesian planes by taking an irreducible conic. In fact, we used the property of irreducible conics being a set of $q+1$ points no three of which collinear to motivate the definition of ovals. A deep result of Segre shows that when q is odd, the converse also holds true.

Theorem 2.23 (Segre 1955)**.** *Every oval in $\mathrm{PG}(2,q)$, for q odd, is a nondegenerate conic.*

This statement was conjectured by Jarnefelt and Kustaanheimo in 1949. In a review, Marshall Hall Jr. said that "The reviewer finds this conjecture implausible." He was again a reviewer of the paper of Segre, where he then said "The fact that this conjecture seemed implausible to the reviewer seems to have been at least a partial incentive to the author to undertake this work. It would be very gratifying if further expressions of doubt were as fruitful.".

For q even it is not the case that every oval/hyperoval comes from a conic. In the exercise you will prove that for every even prime power $q \geq 8$ there exists an oval in $\mathrm{PG}(2,q)$ which is not a conic.

Classifying hyperovals in $\mathrm{PG}(2,q)$, for q even, is a major open problem in finite geometry. This problem is further motivated by various connects that hyperovals have with other important objects like generalized quadrangles and certain permutation polynomials over finite fields. Complete classification of hyperovals is only known for $q < 64$, and there are various infinite families which do not come from extending a conic by a point.

In non-Desarguesian planes, even the existence of ovals and hyperovals is a widely open problem in general. It was long conjectured that they always exist, until 1996 when Penttila, Royle and Simpson showed that there are projective planes of order 16 that do not contain any ovals, using computer aided search.

2.4 Exercises

1. (a) Prove that the Fano plane is not a subplane of $\mathrm{PG}(2, \mathbb{R})$.

   (b) Give a characterisation of all fields F, in terms of the characteristic of F, for which the Fano plane is a subplane of $\mathrm{PG}(2, F)$.

2. Let $f \in \mathbb{F}_q[x, y, z]$ be an irreducible homogeneous polynomial of degree 2, defining the non-degenerate conic C in $\mathrm{PG}(2, q)$. Let q be odd.

   (a) Prove that the function $\beta(u, v) = f(u + v) - f(u) - f(v)$ is a symmetric bilinear form on the vector space $\mathbb{F}_q^3$ such that for any non-zero vector u, the set $u^\perp = \{v \in \mathbb{F}_q^3 : \beta(u, v) = 0\}$ is a 2-dimensional subspace of $\mathbb{F}_q^3$, and hence a line of $\mathrm{PG}(2, q)$.

   (b) Prove that the map $x \mapsto x^\perp$ gives rise to an isomorphism between $\mathrm{PG}(2, q)$ and its dual $\mathrm{PG}(2, q)^D$. To which lines do the points of C get mapped to by this map?

3. In this exercise we give an alternate proof of the fact any oval in a finite projective plane of even order can be uniquely extended to a hyperoval. Let $\mathcal{O}$ be an oval in a finite projective plane of even order.

   (a) Prove that through every point of the plane which lies on a secant of $\mathcal{O}$, i.e. a line intersecting $\mathcal{O}$ in exactly two points, there is a unique line which is tangent to $\mathcal{O}$.

   (b) Deduce that all tangents of $\mathcal{O}$ intersect in a common point, and thus show that $\mathcal{O}$ can be uniquely extended to a hyperoval.

4. Let K_6 denote the complete graph on the vertex set $[6] := \{1, 2, 3, 4, 5, 6\}$. Define a point line geometry $(\mathcal{P}, \mathcal{L}, I)$ as follows:

   - $\mathcal{P} = \mathcal{P}_1 \cup \mathcal{P}_2$ where $\mathcal{P}_1 = [6]$ and $\mathcal{P}_2$ is the set of all distinct 1-factors (perfect matchings) of K_6.
   - $\mathcal{L} = \mathcal{L}_1 \cup \mathcal{L}_2$ where $\mathcal{L}_1 = \binom{[6]}{2}$ and $\mathcal{L}_2$ is the set of all distinct 1-factorizations of K_6.[5]
   - a vertex $x \in \mathcal{P}_1$ is incident to an edge $e \in \mathcal{L}_1$ if $x \in e$ and it is never incident to an element of $\mathcal{L}_2$; a 1-factor $f \in \mathcal{P}_2$ is incident to an edge $e \in \mathcal{L}_1$ if $e \in f$ and it is incident to a 1-factorization $F \in \mathcal{L}_2$ if $f \in F$.

   (a) Prove that this point-line geometry is a projective plane of order 4.

[5] A 1-factorization is a partition of the edges of K_6 into pairwise disjoint 1-factors.

(b) Prove that this projective plane is isomorphic to $\mathrm{PG}(2,4)$. (Hint: use the existence of a hyperoval in $\mathrm{PG}(2,4)$.)

(c) Prove that every finite projective plane of order 4 is isoomrphic to $\mathrm{PG}(2,4)$, by showing that every such plane has a hyperoval.

5. (a) For an arbitrary field F, let S be a set of 5 points in $\mathrm{PG}(2,F)$ such that no three of them are collinear. Prove that there exists a unique conic containing S. Moreover, show that the conic is non-degenerate.

   (b) Find the number of distinct non-degenerate conics in $\mathrm{PG}(2,q)$.

   (c) For every even prime power $q \geq 8$, show that there exists an oval in $\mathrm{PG}(2,q)$ which is not a conic.

6. Every conic in $\mathrm{PG}(2,\mathbb{F}_q)$ is one of the following five types:

   a) Conics made up of two coincident lines of $\mathrm{PG}(2,q)$.

   b) Conics whose equations split into two distinct linear equations over $\mathbb{F}_q$.

   c) Conics whose equations split over two distinct conjugate linear equations over $\mathbb{F}_{q^2}$.

   d) Irreducible conics that have a point in $\mathrm{PG}(2,q)$.

   e) Irreducible conics that have no point in $\mathrm{PG}(2,q)$.

   Count the total number of conics in $\mathrm{PG}(2,q)$, the number of conics of each type, and deduce that every irreducible conic in $\mathrm{PG}(2,q)$ has a point on it.

7. Let q be an odd prime power and let S be a set of points in $\mathrm{PG}(3,q)$ such that no three of them are collinear.

   (a) Prove that $|S| \leq q^2+1$ with equality if and only if every plane of $\mathrm{PG}(3,q)$ intersects S in 0, 1 or $q+1$ points.

   (b) Prove that if $|S| = q^2+1$ then through every point x of S, there exists a unique plane π_x such that $S \cap \pi_x = \{x\}$ and moreover, for any line ℓ through x, we have $\ell \cap S = \{x\}$ if and only if $\ell \in \pi_x$.

   (c) Construct such a set S with $|S| = q^2+1$, for every odd prime power q.

8. Let F be a finite field of characteristic 2, and let C be a conic given by $ax^2+by^2+cz^2+fyz+gzx+hxy=0$.

   (a) Prove that C is equal to a repeated line if and only if $f=g=h=0$.

   (b) Say C is not equal to a repeated line, and let N be the point of $\mathrm{PG}(2,F)$ with coordinates (f,g,h). Prove that C is singular if and only if N lies on C.

   (c) Say C is irreducible. Then prove that every line tangent to C passes through the point N.

9. (a) Prove that the set $\{(1,t,\sqrt{t}) : t \in \mathbb{F}_q\} \cup \{(0,0,1),(0,1,0)\}$ is a hyperoval in $\mathrm{PG}(2,q)$, if q is an even prime power.[6]

[6] Note that for q even every element of $\mathbb{F}_q$ is a square.

(b) Does the hyperoval constructed in (a) contain a conic in it?

3 Some Combinatorial Questions

In this chapter we study some combinatorial questions arising in finite geometry. We will see how beyond simple counting arguments, algebraic methods also play an important role in answering some of these questions when we are working over Desarguesian spaces.

3.1 Blocking Sets

Definition 3.1. A subset B of points in a projective plane is called a *blocking set* if B meets all lines of the plane non-trivially, i.e., for every line ℓ we have $\ell \cap B \neq \emptyset$.

Remark 3.2. The concept of vertex covers in hypergraphs is a generalization of blocking sets.

Remark 3.3. The notion of blocking sets goes back to a 1955 paper of Richardson "On Finite Projective Games", where such a set was called a blocking coalition, if we have the extra condition that B contains no-lines.

If B is a blocking set, then clearly every set containing B is also a blocking set. Moreover, a line in a projective plain of order n always forms a blocking set of size $n+1$, since any two lines intersect each other. It turns out that these are the smallest possible blocking sets.

Theorem 3.4. *Every blocking set in a projective plane of order n has size at least $n+1$, with equality if and only if the blocking set is a line.*

Proof. The bound itself easily follows from the fact that each point blocks at most $n+1$ lines, and hence if we have $\leq n$ points, then they blocks at most $n(n+1) = n^2+n$ lines, which is less than the total number of lines. To characterise the sharp examples, we prove a stronger claim that for any $1 \leq k \leq n+1$, the number of lines meeting a set of k points is at most $kn+1$, with equality if and only if the k points are contained in a line.

This is clearly true for $k=1$ since a point meets exactly $n+1$ lines. Say the statement is true for some $k \leq n$. We will prove it for $k+1$.

Let S be a set of $k+1$ points and $x \in S$. Then $S \setminus \{x\}$ meets at most $kn+1$ lines with equality if and only if $S \setminus \{x\}$ is contained in a line ℓ. There are $n+1$ lines through x, and at least one of them also meets another point of $S \setminus \{x\}$, and hence it has already been counted in the lines meeting $S \setminus \{x\}$. Therefore, S meets at most $kn+1+n = (k+1)n+1$ lines. Equality can only occur if there is exactly one line ℓ through x which meets a point of $S \setminus \{x\}$, and all the points of $S \setminus \{x\}$ lie on ℓ. The line ℓ also contains x, and hence S is a subset of ℓ. Conversely, $k+1$ collinear points meet precisely $(k+1)n+1$ lines. □

Inspired by this result, one might ask what is the corresponding result for a finite affine plane. In any affine plane of order n, a pair of intersecting lines forms a blocking set, since every other line is parallel to at most one of these lines, and hence intersects the other. This gives us a blocking set of size $2n-1$. But is this the best that we can do? It turns out that if the affine plane is Desarguesian, then indeed this size is optimal (though there are many other configurations giving us the same bound), but proving this is highly non-trivial! Moreover, it is known that the following result is not true if the planes are non-Desarguesian.

Theorem 3.5 (Jamison/Brouwer-Schrijver)**.** *Every blocking set in the affine plane* $\mathrm{AG}(2,q)$ *has size at least* $2q-1$.

Proof. Let B be a blocking set in $\mathrm{AG}(2,q)$. Let B' be the set of points corresponding to B in the projective plane $\mathrm{PG}(2,q)$ we obtain by adding ℓ_∞. Then B' is a set of points in $\mathrm{PG}(2,q)$ that meets every line except one, ℓ_∞. In the dual projective plane, which is isomorphic to $\mathrm{PG}(2,q)$, we get a set L' of lines in $\mathrm{PG}(2,q)$ that cover all points except one. Remove a line $\ell \in L'$ to get a set L of lines in $\mathrm{AG}(2,q)$ that cover all points except one. We will show that $|L| \geq 2(q-1)$, which would imply that $|B| = |B'| = |L'| = |L|+1 \geq 2q-1$.

Every line of L is given by the zero set of a degree one polynomial in $\mathbb{F}_q[x,y]$. Multiply these polynomials together to get a polynomial f of degree $|L|$ with the property that f vanishes on all points of $\mathbb{F}_q^2$, except one. From the Lemma below we see that $|L| = \deg f \geq 2(q-1)$, and hence $|B| = |L|+1 \geq 2q-1$. □

Lemma 3.6. *Let* $f \in \mathbb{F}_q[x_1,\dots,x_n]$. *If there exists a point* $u \in \mathbb{F}_q^n$ *such that* $f(u) \neq 0$ *and* $f(a) = 0$ *for all* $a \neq u$, *then* $\deg f \geq n(q-1)$.

Proof. Say $\deg f < n(q-1)$. Then in every monomial m of f, there exists an exponent which is less than $q-1$. This implies that $\sum_{a\in\mathbb{F}_q^n} m(a) = 0$, since we can split the sum according to the n variables, and the variable whose power is less than $q-1$ and greater than 0 gives us the 0 factor by Lemma 0.9. If the exponent of a variable is 0 then this also gives us a 0 sum, as then we have q times the sum of the evaluation of the monomials Therefore, $\sum_{a\in\mathbb{F}_q^n} f(a) = 0$. But this is impossible since the sum is equal to $f(u) \neq 0$. □

As mentioned before, the lower bound of $2q-1$ on size of a blocking set does not hold true for non-Desarguesian affine planes. In fact, recently it has been proved by De Beuele, Héger, Szönyi and Van de Voorde, that for every prime power $q \geq 3$, there exist non-Desarguesian affine planes of order q^2 that contain a blocking set of size $\lfloor 4q^2/3 + 5q/3 \rfloor$.

Another interesting consequence of Lemma 3.6 is the following classical result in number theory, which in particular implies that every conic in $\mathrm{PG}(2,q)$ has at least one point in it.

Theorem 3.7 (Chevalley-Warning Theorem)**.** *Let* $f \in \mathbb{F}_q[x_1,\dots,x_n]$ *be a polynomial of degree* $d < n$. *If the set* $Z = \{a \in \mathbb{F}_q^n : f(a) = 0\}$ *is non-empty, then* $|Z| > 1$.

Proof. Let $g = 1 - f^{q-1}$. Then g is a polynomial of degree $d(q-1) < n(q-1)$, such that $g(a) \neq 0$ for all $a \in Z$ and $g(a) = 0$ for all $a \notin Z$. If $|Z| = 1$, then g vanishes on all points except one, and hence by Lemma 3.6, it must have degree at least $n(q-1)$, a contradiction. Therefore, either $Z = \emptyset$ or $|Z| > 1$. □

Corollary 3.8. *Every conic in* $\mathrm{PG}(2,\mathbb{F}_q)$ *contains a point.*

Proof. Let $\phi \in \mathbb{F}_q[x, y, z]$ be the homogeneous polynomial of degree 2 defining the conic. Since ϕ is homogeneous, $(0, 0, 0) \in Z := \{a \in \mathbb{F}_q^3 : \phi(a) = 0\}$. Therefore, $|Z| > 1$, and hence there exists a non-zero vector at which ϕ vanishes, giving us a point in C. □

While Theorem 3.4 might suggest that there is nothing that interesting about blocking sets in finite projective planes, since we can completely characterise the minimal ones, and in fact we can easily construct blocking sets of any larger size by just adding extra points. Things start getting more interesting if we disallow our blocking sets to contain lines.

Definition 3.9. A blocking set in a projective plane is called *non-trivial* if it does not contain any line.

What are some examples of non-trivial blocking sets that you can think of? What size does your set have, if the order of the projective plane is n? One of the simplest examples is the set of size $3n - 3$ consisting of three sides of a triangle except the vertices, the so-called *vertex-less triangle*, which gives a non-trivial blocking set for all $n > 2$. Is this the smallest possible non-trivial blocking set? We will construct better ones for some families of finite projective planes.

Consider $\mathrm{PG}(2, q^2)$. Restricting the coordinates to the subfield $\mathbb{F}_q$, we get a copy of $\mathrm{PG}(2, q)$ inside $\mathrm{PG}(2, q^2)$. This is known as a Baer subplane. In fact, in general if $n = m^2$, then a Baer subplane in a projective plane of order n is defined as a subplane of order m.

Theorem 3.10 (Bruen 1970)**.** *Baer subplanes are non-trivial blocking sets.*

Proof. Let B be a collection of m points in a projective plane π of order $n = m^2$ such that the unique line through any 2 points of B contains exactly $m + 1$ points of B. Moreover, any two lines which contain $m+1$ points of B meet inside B. For the sake of contradiction, let ℓ be a line of π that does not contain any point of B. Let $x \in \ell$. There can be at most one line through x that contains $m + 1$ points of B, since B forms a Baer subplane, and every other line through x contains at most one point of B. Since the lines through x must partition the points of B, we get $|B| \leq m + 1 + (n - 1) = m^2 + m$, which is a contradiction since B is a Baer subplane of order m and hence has $m^2 + m + 1$ points. □

Note that from the proof it follows that every line intersects the set of points of a Baer subplane in 1 or $\sqrt{n} + 1$ points, where n is the order of the plane. In fact, it can be shown that any set of $n + \sqrt{n} + 1$ points with these intersection properties with respect to lines must correspond to a Baer subplane.

Theorem 3.11 (Bruen 1970)**.** *Every non-trivial blocking set in a projective plane of order n has size at least $n + \sqrt{n} + 1$, with equality if and only if it is the set of points of a Baer subplane.*

Proof. Let B be a non-trivial blocking set of the projective plane. We will assume that B is of size $n + m$, with $m \leq \sqrt{n} + 1$ and show that it must be a Baer subplane. First we claim that every line intersects a such a set B in at most m points. For this, let ℓ be a line and let x be a point in $\ell \setminus B$. Every line through x contains at least one point of B, giving us $n + |\ell \cap B| \leq |B| = n + m$, and hence $|\ell \cap B| \leq m$.

Let the lines of the plane be $\ell_1, \ldots, \ell_{n^2+n+1}$, and define $k_i = |\ell_i \cap B|$. By double counting the pairs (x, ℓ) where x is a point of B and ℓ a line through x, we get

$$\sum k_i = |B|(n+1).$$

By double counting the triplets (x, y, ℓ) where $x \neq y \in B$ and ℓ is the unique line joining x and y, we get

$$\sum k_i(k_i - 1) = |B|(|B| - 1).$$

Consider the sum $S = \sum (k_i - 1)(\sqrt{n} + 1 - k_i)$. Since $1 \leq k_i \leq \sqrt{n} + 1$ for all i, we have $S \geq 0$, with equality if and only if every line intersects B in 1 or $\sqrt{n}+1$ points. By expanding the sum and using the previous identities we get

$$S = -|B|(|B|-1)+(\sqrt{n}+1)(n+1)|B|-(\sqrt{n}+1)(n^2+n+1) = (|B|-n-\sqrt{n}-1)(n\sqrt{n}+1-|B|).$$

Since $S \geq 0$, we must have $|B| \geq n + \sqrt{n} + 1$, and since we assumed $|B| = n + m \leq n + \sqrt{n} + 1$, we get $|B| = n + \sqrt{n} + 1$. Therefore, $S = 0$ and hence every line intersects B in 1 or $\sqrt{n} + 1$ points.

□

Remark 3.12. Instead of working with k_i's, we can define e_i to be the number of lines which intersect the blocking set in exactly i points, which will then give us the following equations:

$$\sum_{i=1}^{\sqrt{n}+1} e_i = n^2 + n + 1,$$

$$\sum_{i=1}^{\sqrt{n}+1} ie_i = |B|(n+1),$$

$$\sum_{i=1}^{\sqrt{n}+1} i(i-1)e_i = |B|(|B|-1).$$

This choice of variables is similar to our earlier proofs on ovals, but it basically gives the same proof as above when we try to show that $e_i = 0$ for all $i \notin \{1, \sqrt{n}+1\}$.

What happens when the order of the projective plane is not a square? What's the smallest size of a non-trivial blocking set? This is a much more difficult question. It is known that in $\mathrm{PG}(2,p)$ the smallest non-trivial blocking set has size $3(p+1)/2$. Proving this requires another surprising application of the polynomial method, due to Blokhuis.

Theorem 3.13 (Blokhuis 1994). *Let S be a non-trivial blocking set in* $\mathrm{PG}(2,p)$. *Then*

$$|S| \geq 3(p+1)/2.$$

The construction however can be described using directions determined by certain sets of points in $\mathrm{AG}(2,p)$, as you will see in the exercises. We now look at another direction problem, which was solved quite recently using polynomial method.

3.2 Kakeya and Nikodym Sets

The classical Kakeya's needle problem asks for the "smallest" subset S of $\mathbb{R}^2$ such that S contains a unit line segment (needle) in every direction. Clearly a semi-circle of radius 1 is such an example with area $\pi/2$. A circle of diameter 1 is a smaller example of such a set, as it has area $\pi/4$. An even smaller example is an equilateral triangle of height 1, which has area $1/\sqrt{3}$[1]. The reader is invited to come up with further examples that have smaller area. Perhaps surprisingly, Besicovitch proved (in 1919) there is no $\delta > 0$ such that every Kakeya set in $\mathbb{R}^2$ has area at least δ. Therefore if area was a measure of how small a Kakeya set be in $\mathbb{R}^2$, then the answer is 0. Over the years, mathematicians used other "measures" (like Minkowski dimension and Hausdorff dimension), and made conjectures regarding lower bounds. For example, it was conjectured that any Kakeya set in $\mathbb{R}^n$ has Hausdorff/Minkowski dimension at least n (any subset of $\mathbb{R}^n$ has Hausdorff dimension at most n, and hence Kakeya sets are conjectured to have the highest possible Hausdorff dimension). This is known for $n = 1, 2$ but it is still wide open for all $n \geq 3$.

To solve these conjectures, many connections between Kakeya conjecture and some problems from areas like Additive Combinatorics, and Fourier Analysis were found.[2] In 1999, Wolff suggested that perhaps looking at the same question over the finite field setting might provide us some tools to make progress in the original setting. He made the following conjecture regarding Kakeya sets in $\mathrm{AG}(n, q)$:

Conjecture 3.14. *Let S be a set of points in the affine space $\mathrm{AG}(n, q)$ with the property that for every direction, there exists a line in that direction that is completely contained in S. Then*

$$|S| \geq c_n q^n,$$

for some constant c_n that only depends on n.

Remark 3.15. In $\mathrm{AG}(n, q)$, the line joining points $x = (x_1, \dots, x_n)$ and $y = (y_1, \dots, y_n)$ is given by $\{x + \lambda(y - x) : \lambda \in \mathbb{F}_q\}$, where $y - x = (y_1 - x_1, \dots, y_n - x_n)$.

Note that the "direction" of a line ℓ in $\mathrm{AG}(n, q)$ corresponds to a point p in the hyperplane at infinity, H_∞, in the projective completion $\mathrm{PG}(n, q)$ of the affine space, such that in $\mathrm{PG}(n, q)$ the line ℓ contains the point p. Therefore, a Kakeya set can alternately be defined as a set S of points in $\mathrm{PG}(n, q)$ such that there exists a hyperplane H disjoint from S and for every point $x \in H$, there exists a line ℓ_x such that $\ell_x \setminus \{x\} \subseteq S$. This definition leads to the concept of Nikodym sets.

Definition 3.16. A *Nikodym set* in $\mathrm{PG}(n, q)$ (or $\mathrm{AG}(n, q)$), is a set S of points such that for all $x \notin S$, there exists a line ℓ_x for which $(\ell_x \setminus \{x\}) \subseteq S$.

A similar conjecture for the size of Nikodym sets was also proposed. So, let us explore how small Kakeya and Nikodym can be in the finite field spaces. We will start with Kakeya sets in the plane.

Example 3.17. Let $\mathcal{H}$ be a hyperoval in $\mathrm{PG}(2, q)$ for q even. In the dual plane, which is also isomorphic to $\mathrm{PG}(2, q)$, $\mathcal{H}$ corresponds to a set $\mathcal{L}$ of $q+2$ lines such that no-three lines are concurrent. Let ℓ be a line of $\mathcal{L}$ and define K to be the set of points contain on the lines in $\mathcal{L}$, except those lying on ℓ. Then K is a Kakeya set of size $\binom{q+2}{2} - (q+1) = \binom{q+1}{2}$.

[1] This is in fact the minimum possible area if we only allow convex sets.

[2] See for a survey.

Theorem 3.18. *A Kakeya set in* $\mathrm{AG}(2,q)$ *has size at least* $\binom{q+1}{2}$.

Proof. Let K be the Kakeya set. We have $q+1$ distinct lines, corresponding to each of the directions, completely contained in our set. The first line gives us q points of K, the second one gives rise to at least $q-1$ further points, the third line $q-2$ points, and so on, because any two of these lines intersect in at most (in fact, exactly) one point. Therefore, $|K| \geq q(q+1)/2$. Equality occurs if and only if K is the union of these $q+1$ lines, and no three lines are concurrent,i.e., the lines along with the line at infinity form a hyperoval since all these lines are in different direction. □

Example 3.19. Let $\mathcal{O}$ be an oval in $\mathrm{PG}(2,q)$, with q odd. Take a point x on $\mathcal{O}$ and let ℓ be the tangent line through x to $\mathcal{O}$. Through every point y on ℓ, except for x, there is a unique line $\ell_y \neq \ell$ that is a tangent to $\mathcal{O}$. Pick any line ℓ_x through x which is not equal to ℓ. Let

$$K = \left(\bigcup_{y \in \ell} \ell_y \right) \setminus \ell.$$

Then in the affine plane that one obtains by treating ℓ as the line at infinity, K is a Kakeya set in $\mathrm{AG}(2,q)$ of size $q(q+1)/2 + (q-1)/2$.

Theorem 3.20 (Blokhuis-Mazzocca 2008)**.** *A Kakeya set in* $\mathrm{AG}(2,q)$*, for q odd, has size at least* $q(q+1)/2 + (q-1)/2$*, with equality if and only if it is equivalent to the example given above.*

The proof is beyond the scope of our course, but it's interesting to note that its main ingredients are Theorems 2.23 and 3.5. Blokhuis and Mazzocca also managed to give a linear algebraic proof of the fact that in $\mathrm{AG}(n,q)$, the size of a Kakeya set is at least $q^{n-1}/(n-1)!$, which was quite close to the conjecture and substantially improved the previously known best bound of $c_n q^{(n+2)/2}$. But soon after they proved their results, Dvir settled the conjecture completely, using a different (and much simpler) linear algebraic argument. We will now see his proof, first for Nikodym sets (since it's simpler) and then for Kakeya sets. The basic idea behind the proof is as follows:

(1) Say the set S has smaller cardinality then what we want to prove.

(2) Show that for such a small S there must be a non-zero polynomial f that vanishes on all points of S.

(3) Show that such a polynomial f then vanishes on too many points, and hence must be the 0 polynomial, giving us a contradiction.

We start with the algebraic lemmas that formalise steps (2) and (3) above.

Lemma 3.21. *Let F be a field, d, n positive integers, and $E \subseteq F^n$ a set of cardinality strictly less than $\binom{d+n}{n}$. Then there exists a non-zero polynomial $f \in F[x_1, \ldots, x_n]$ of degree at most d such that $f(x) = 0$ for all $x \in E$.*

Proof. The vector space F^E of functions from E to F, is of dimension $|E|$. Consider the vector space V of all polynomials in $F[x_1, \ldots, x_n]$ with degree at most d. Then by a simple counting argument, $\dim V = \binom{d+n}{n}$. Consider the linear map ev from V to F^E, defined by taking $\mathrm{ev}(f)$ to be the function $x \mapsto f(x)$. If $|E| < \dim V$, then this map cannot be injective, and in particular the kernel of ev contains a non-zero polynomial f. □

Lemma 3.22 (Ore, DeMillo-Lipton-Schwartz-Zippel). *Let $f \in \mathbb{F}_q[x_1, \ldots, x_n]$ be a non-zero polynomial of degree $d < q$. Then f has at most dq^{n-1} zeros in $\mathbb{F}_q^n$.*

Proof. We prove this by induction on n. For $n = 1$, this follows from the fact that any univariate polynomial of degree d over an arbitrary field has at most d zeros. We now proceed by induction. Let $n > 1$, and assume that the statement is true for $n - 1$. Write f as

$$f(x_1, \ldots, x_n) = \sum_{i=1}^{d_n} g_i(x_1, \ldots, x_{n-1})x_n^i,$$

where $d_n = \deg_n f$ is the highest exponent of x_n appearing in f. Note that $\deg g_n \leq d - d_n$, since $d = \deg f$, and $g_n \neq 0$. Let $(a_1, \ldots, a_{n-1}) \in \mathbb{F}_q^{n-1}$ such that $g_n(a_1, \ldots, a_{n-1}) \neq 0$, then $f(a_1, \ldots, a_{n-1}, x_n)$ is a non-zero univariate polynomial of degree d_n, and hence has at most d_n zeros. We can partition the set Z of zeros of f in $\mathbb{F}_q^n$ into the set $Z_1 = \{(a_1, \ldots, a_{n-1}, a_n) : g_n(a_1, \ldots, a_{n-1}) = 0, f(a_1, \ldots, a_n) = 0\}$ and the set $Z_2 = \{(a_1, \ldots, a_{n-1}, a_n) : g_n(a_1, \ldots, a_{n-1}) \neq 0, f(a_1, \ldots, a_n) = 0\}$. By the discussion above, $|Z_2| \leq d_n q^{n-1}$ and by the induction hypothesis, $|Z_1| \leq q(d - d_n)q^{n-2}$ since $\deg g_n = d - d_n$. Therefore, $|Z| \leq d_n q^{n-1} + dq^{n-1} - d_n q^{n-1} = dq^{n-1}$. Let $S \subseteq \mathbb{F}_q$ be the set of values t for which $f(x_1, \ldots, x_{n-1}, t)$ is the 0 polynomial in $\mathbb{F}_q[x_1, \ldots, x_{n-1}]$. □

We can now easily prove the conjectured bound for Nikodym sets.

Theorem 3.23. *A Nikodym set in* AG(n, q) *has size at least* $\binom{q-2+n}{n}$.

Proof. Let N be a Nikodym set. Say $|N| < \binom{q+n-2}{n}$. Then by Lemma 3.21 there exist a non-zero polynomial f of degree at most $q - 2$ which vanishes on N. Let $x \in \mathbb{F}_q^n \setminus N$. Since N is a Nikodym set, there exists a line ℓ through x such that $\ell \setminus \{x\} \subseteq N$. The points on ℓ are given by $\{x + \lambda v : \lambda \in \mathbb{F}_q\}$ for some non-zero vector v (which is the direction of the line). Evaluating f at an arbitrary point of ℓ, we get $f(x + \lambda v)$ which is univariate polynomial in λ of degree at most $q - 2$. Since it vanishes on the $q - 1$ points $\{x + \lambda v : \lambda \neq 0\}$, it must be identically 0. In particular, this means that $f(x) = 0$. Since x was an arbitrary point outside N, and f vanishes on every point of N, we have f vanishing entirely on $\mathbb{F}_q^n$, which is a contradiction to Lemma 3.22. □

Note that $\binom{q-2+n}{n} = (q + n - 2) \cdot (q + n - 3) \cdots (q - 1)/n! \geq q^n/(2n!)$ for all prime powers q and $n \geq 2$. For Kakeya sets, the proof above needs to be modified a bit.

Theorem 3.24. *A Kakeya set in* AG(n, q) *has size at least*

$$\binom{q-1+n}{n} \geq q^n/n!.$$

Proof. As before, say there is a Kakeya set K with a smaller size. Let $f \in \mathbb{F}_q[x_1, \ldots, x_n]$ be a non-zero polynomial of degree $d \leq q-1$ which vanishes on K. Let $f^h \in \mathbb{F}_q[x_1, \ldots, x_n, x_{n+1}]$ be the homogenized form of f obtained by making each monomial in f of degree exactly d by multiplying it with a suitable power of x_{n+1}. Note that $f^h(x_1, \ldots, x_n, 0)$ is the homogeneous part of the polynomial f, that is, the degree d terms in f, and hence this polynomial is non-zero. Also note that $f^h(x_1, \ldots, x_n, 1) = f(x_1, \ldots, x_n)$. Embed the set K in PG(n, q) formed by adding the hyperplane H_∞ given by $x_{n+1} = 0$. The fact that K is a Kakeya set is equivalent to saying that for all points in H_∞, there is a line through this point whose affine points are all contained in K. Let $(a_1, \ldots, a_n, 0)$ be such a point, with $(a_1, \ldots, a_n) \neq (0, \ldots, 0)$, and ℓ the line through it. Let $(b_1, \ldots, b_n, 1)$ be an affine point on ℓ. Then the affine points of ℓ are given by $\{(b_1 + \lambda a_1, \dots, b_n + \lambda a_n, 1) : \lambda \in \mathbb{F}_q\}$. As before, f^h restricted to ℓ is a univariate polynomial in λ of degree at most $d \leq q-1$ which has at least q zeros since $\ell \subseteq K$, implying that it is the zero polynomial. The leading coefficient of $f^h(b_1 + \lambda a_1, \dots, b_n + \lambda a_n, 1)$, is equal to $f^h(a_1, \dots, a_n, 0)$, which must then be equal to 0. Therefore, $g(x_1, \dots, x_n) = f^h(x_1, \dots, x_n, 0)$ vanishes on all non-zero vectors of $\mathbb{F}_q^n$. Since g is homogeneous, it in fact vanishes on all elements of $\mathbb{F}_q^n$. By Lemma 3.22 g should be identically zero since $\deg g = \deg f^h = d \leq q - 1$. This is the contradiction since $g = f^h(x_1, \dots, x_n, 0)$ is non-zero. □

Therefore, Kakeya (and Nikodym sets) are of cardinality at least $c_n q^n$, where $c_n = \Theta(1/n!)$. Is this bound tight, in terms of the expression for c_n? Dvir, Kopparty, Saraf and Sudan later showed that we can improve c_n to $1/2^n$, and they also showed that for Kakeya sets this bound is tight, up-to a factor of 2 and some lower order terms.[3] Interestingly, for Nikodym sets even the bound of $q^n/2^n$ is far from being tight. We explore this in the next section.

3.2.1 Nikodym sets and large minimal blocking sets

We know that every Nikodym set N in $\mathrm{AG}(n, q)$ has size at least $q^n/(2n!)$. In this section we will see how this bound is far from being sharp for $n = 2$. Surprisingly, the improvement in the lower bound will follow from a combinatorial argument, instead of polynomial method.

If N is a Nikodym set in $\mathrm{AG}(n, q)$, then the set $N \cup H_\infty$ is clearly a Nikodym set in $\mathrm{PG}(n, q)$ which has size $|N| + (q^{n-1} + \cdots + q + 1) = |N| + \Theta(q^{n-1})$. Conversely, if N is a Nikodym set in $\mathrm{PG}(n, q)$, then $N \setminus H_\infty$ is a Nikodym set in $\mathrm{AG}(n, q)$ of size at least $|N| - \Theta(q^{n-1})$. Therefore, any lower bound of the form $c_n q^n$, in projective/affine setting, implies a lower bound of the form $(c_n + o(1))q^n$, in the other setting. So it doesn't make much difference in studying Nikodym sets over projective spaces instead of affine spaces (at least as far as the asymptotic are concerned).

Let's fix our n to be 2 and study Nikodym sets in $\mathrm{PG}(2, q)$. Recall that a set N in $\mathrm{PG}(2, q)$ is a Nikodym set iff for every $x \notin N$, there exists a line ℓ_x such that $\ell_x \setminus \{x\} \subseteq N$. Therefore, the complement N^c has the property that through every point x of N^c there exists a line that intersects N^c in precisely one point. Such sets have been studied in finite geometry at least since 1999, when they were introduced by Bruen and Drudge.[4] This is about the same time when Wolff introduced the finite field Kakeya sets.

Definition 3.25. A tangency set in a finite projective plane is a set T such that for all $x \in T$, there exists a line ℓ such that $\ell \cap T = \{x\}$.

For example, a line is a tangency set of size $q + 1$, and so is a conic, in $\mathrm{PG}(2, q)$. A Baer subplane in $\mathrm{PG}(2, q^2)$ is also a tangency set, that has size $q + \sqrt{q} + 1$. The motivation for the study of tangency sets in fact came from the study of minimal blocking sets.

[3] The used a more involved polynomial interpolation, that took into account the multiplicities of the zeros.

[4] An earlier concept of strong representative systems in projective planes due to Illés, Szönyi, Wettle from 1991 is equivalent to tangency sets.

Definition 3.26. A minimal blocking set in a finite projective plane is a blocking set B such that no proper subset of B is also a blocking set.

Note that through every point x of a minimal blocking set B, there must be a line ℓ such that $\ell \cap B = \{x\}$, as otherwise $B \setminus \{x\}$ will also be a blocking set. In particular, every minimal blocking set is a tangency set, and hence the complement of a Nikodym set. It is then possible to answer the question about how small a Nikodym can be via the max-min question: How large can a minimal blocking set be?

Theorem 3.27 (Bruen-Thas 77)**.** *Let B be a tangency set in a projective plane of order q. Then $|B| \leq q\sqrt{q}+1$, with equality if and only if every line intersects B in 1 or $\sqrt{q}+1$ points.*

Proof. Say $|B| = tq+1$. We will prove that $t \leq \sqrt{q}$ with equality if and only if every line intersects B in either 1 point or $\sqrt{q}+1$ points. Let ℓ_i be the i-th line, for $i = 1, \dots, q^2+q+1$, and $k_i = |\ell_i \cap B|$. By the same double counting as before, we have

$$\sum k_i = |B|(q+1),$$

$$\sum k_i(k_i - 1) = |B|(|B| - 1).$$

Since through each point of B there is a line that intersects B in exactly 1 point, we have at least $|B|$ lines that intersect B in exactly one point. Say these are $\ell_1, \dots, \ell_{|B|}$, and hence $k_i = 1$ for $1 \leq i \leq |B|$. Let

$$S = \sum_{i=1}^{q^2+q+1} (k_i - t - 1)^2.$$

Then we get $S \geq \sum_{i=1}^{|B|} t^2 = |B|t^2$, with equality if and only if $k_i = t+1$ for all $i > |B|$, that is, there is exactly one tangent line through each point of S and every other line of the plane intersects S in exactly $t+1$ points. From the equations above, we get

$$S = \sum (k_i^2 - k_i) - \sum (2t+1)k_i + (t-1)^2(q^2+q+1).$$

Putting everything together, we get the inequality

$$(t+1)q(q-t^2) \geq 0,$$

which proves the result. □

Note that in case of equality, we also get that through each point there is a unique tangent line. Can there be equality in the bound? Indeed, is sharp, as shown by the so-called Hermitian curve $x^{q+1} + y^{q+1} + z^{q+1} = 0$ in $\mathrm{PG}(2, q^2)$. Since $x \mapsto x^q$ is an automorphism of $\mathbb{F}_{q^2}$, and in fact the only non-trivial automorphism, we can think of x^q as the conjugate of x, just like $a - ib$ is the conjugate of $a + ib$ in $\mathbb{C}$, which is a quadratic extension of $\mathbb{R}$. The product xx^q is always an element of the subfield $\mathbb{F}_q$, and the analogue of the Hermitian curve over $\mathrm{PG}(2, \mathbb{C})$

$$x\overline{x} + y\overline{y} + z\overline{z} = 0,$$

which in fact has no solutions, just like the conic defined by $x^2 + y^2 + z^2 = 0$ had no points in $\mathrm{PG}(2, \mathbb{R})$. We will later see how Hermitian curves and conics, are both related to *polarities* of a projective plane.

Theorem 3.27 implies that Nikodym set in $\mathrm{PG}(2,q)$ has size at most $q^2+q+1-q\sqrt{q}-q = (1-o(1))q^2$, where $o(1)$ is a function of q that approaches 0 as q approaches ∞. This clearly improves the polynomial method bound of $q^2/4$. In fact, for every $n \geq 2$, it has been conjectured that a Nikodym set in $\mathrm{PG}(n,q)$ (or $\mathrm{AG}(n,q)$) has size at least $(1-o(1))q^n$. The conjecture is true for $n = 2$ but widely open for $n > 2$. The best known general lower bound of $(1-o(1))q^n/2^n$ comes from the polynomial method with multiplicities.

We have shown that a minimal blocking set in $\mathrm{PG}(2,q)$ has size at least $q+\sqrt{q}+1$ and at most $q\sqrt{q}+1$. Determining the possible sizes of a minimal blocking sets, is a challenging problem in finite geometry. Most of the progress here is based on algebraic methods, but sometimes even probabilistic methods gives us existence proofs of certain large minimal blocking sets.

3.3 Exercises

1. Prove that every blocking set of the Fano plane is trivial. For every projective plane of order $n \geq 3$, construct a non-trivial blocking set of size $2n$.

2. Let n be a square. Prove that if B is a collection of $n + \sqrt{n} + 1$ points in a projective plane of order n such that every line intersects B in either 1 point or $\sqrt{n}+1$ points, then B is the point set of a Baer subplane.

3. Let π be a projective plane of order n and $\mathcal{L}$ be a subcollection of lines of π. Then a set B of points of π is called a blocking set relative to $\mathcal{L}$ if it for every line ℓ in $\mathcal{L}$, we have $\ell \cap B \neq \emptyset$. Let $\mathcal{O}$ be an oval of π and assume that n is odd. For each of the following subcollections $\mathcal{L}$ of lines, find the smallest possible size of a blocking set relative to $\mathcal{L}$.

   a) $\mathcal{L}$ is the set of all external lines to $\mathcal{O}$.

   b) $\mathcal{L}$ is the set of all secants to $\mathcal{O}$.

   c) $\mathcal{L}$ is the union of all tangents and externals of $\mathcal{O}$.

   d) $\mathcal{L}$ is the union of all tangents and secants of $\mathcal{O}$.

   In each case, characterise the extremal examples, assuming that you have a Desarguesian projective plane.

4. For a subset S of points in $\mathrm{AG}(2,q)$ we say that S determines a direction m if there exist two points in S such that the line joining them has slope m, where $m \in \mathbb{F}_q \cup \{\infty\}$ (vertical lines given by the equation $x = c$ are assumed to have slope ∞).

   (a) Prove that if $|S| > q$, then S determines all directions.

(b) Let f be a function from $\mathbb{F}_q$ to $\mathbb{F}_q$. Consider the set $S = \{(x, f(x)) : x \in \mathbb{F}_q\}$ of points in $\mathrm{AG}(2, q)$ and let D be the set of directions determined by S. Prove that if $|D| > 1$, then the set $B_f = S \cup D$ is a non-trivial blocking set in the projective completion $\mathrm{PG}(2, q)$ of $\mathrm{AG}(2, q)$.

(c) Determine the size of the blocking set B_f when $f(x) = x^{(q+1)/2}$, assuming q to be odd.

(d) For every even q, construct a minimal blocking set of size $(3q + 2)/2$.

5. Let S be a set of points in $\mathrm{AG}(2, q)$ such that S blocks every line in each direction except at most k directions where at most m lines remain unblocked. Prove that $|S| \geq 2q - k - m$.

6. A blocking set in a projective space is a set of points such that every hyperplane meets the set non-trivially. For all $n \geq 2$, and all prime powers q, find the minimum possible size of a blocking set in $\mathrm{PG}(n, q)$, and classify all minimal examples.

7. Let $s(q)$ be the smallest possible size of a set S of points in $\mathrm{AG}(3, q)$ such that every line intersects S non-trivially.

(a) Determine $s(2)$ and $s(3)$.

(b) Prove that
$$2q^2 - q \leq s(q) \leq 3q^2 - 3q + 1$$
for all q.

(c) Prove that $s(q) \geq 2q^2 - 1$ for all q, and $s(q) \leq 3q^2 - 3q$ for all $q \geq 3$.

(d) For q square, prove that $s(q) \leq 2q^2 + o(q^2)$, by using the fact that for such a q there are two disjoint Baer subplanes in $\mathrm{PG}(2, q)$.

(e) (**Open Problem**) Determine $s(q)$ for all q, at least asymptotically. In particular, show that $\limsup s(q)/q^2$ exists and determine its value.

8. Let S be a set of points in $\mathrm{PG}(n, q)$ with the property that through each point x of S there exists a hyperplane H_x such that $H_x \cap S = \{x\}$, a.k.a., a tangent hyperplane.

(a) Prove that $|S| \leq q^{(n+2)/2+1}$, and equality in the bound implies that through each point there is a unique hyperplane.

(b) Prove that $|S| < q^{(n+2)/2} + 1$ for all $n \geq 4$.

9. Let $\mathcal{A}$ be an affine plane of order q^2, and let π be the projective plane of order q^2 obtained from $\mathcal{A}$ by adding a line ℓ_∞. Let D be a set of $q + 1$ points on ℓ_∞ with the following property: for any two distinct points x, y of $\mathcal{A}$, if the line xy meets ℓ_∞ in a point of D, then there exists a Baer subplane B of π containing x and y such that $B \cap \ell_\infty = D$. Define the following point-line geometry $\mathcal{A}_D = (\mathcal{P}, \mathcal{L}, I)$:

(a) $\mathcal{P}$ is equal to the the point set of $\mathcal{A}$.

(b) $\mathcal{L} = \mathcal{L}_1 \cup \mathcal{L}_2$, where $\mathcal{L}_1$ is the set of lines of $\mathcal{A}$ that meet ℓ_∞ in a point outside D, and $\mathcal{L}_2$ is the set of Baer subplanes B of π that satisfy $B \cap \ell_\infty = D$.

(c) I is the natural incidence of containment.

Prove that $\mathcal{A}_D$ is an affine plane of order q^2. [5]

10. (a) Let S be be a subset of points in a projective plane of order q, and let $k_1, \ldots, k_{q^2+q+1}$ be the intersection sizes of the q^2+q+1 lines with the set S. Prove that

$$\sum_{i=1}^{q^2+q+1} \left(k_i - \frac{(q+1)|S|}{q^2+q+1} \right)^2 \leq q|S|.$$

(b) Let S be a set of points and T a set of lines in a projective plane of order q. Define $i(S,T)$ to be the number of incidences between the points of S and the lines of T, that is, $i(S,T) = |I \cap (S \times T)|$, where I is the point-line incidence relation of the projective plane. Prove that

$$\left| i(S,T) - \frac{q+1}{q^2+q+1} |S||T| \right| \leq \sqrt{q|S||T|}.$$

11. Let $\mathcal{H} = \{(x,y,z) \in \mathrm{PG}(2,q^2) : x^{q+1} + y^{q+1} + z^{q+1} = 0\}$.

a) Determine $|\mathcal{H}|$ by considering the following three types of points in the plane: $\{(x,1,0) : x \in \mathbb{F}_{q^2}\}$, $\{(0,y,1) : y \in \mathbb{F}_{q^2}\}$ and $\{(x,y,1) : x,y \in \mathbb{F}_{q^2}, x \neq 0\}$.

(Hint: The function $x \mapsto x^{q+1}$ is the norm function from $\mathbb{F}_{q^2}$ to $\mathbb{F}_q$.

b) Let $P = (x_0, y_0, z_0) \in \mathcal{H}$. Assuming that no line is completely contained in $\mathcal{H}$, show that the line ℓ_P defined by the equation $x_0^q x + y_0^q y + z_0^q z = 0$, satisfies $\ell_P \cap \mathcal{H} = \{P\}$. Deduce that every line meets $\mathcal{H}$ in exactly 1 or $q+1$ points.

c) (Bonus) Show that $\mathcal{H}$ does not contain any line completely.

12. (a) Let S be be a subset of points in a projective plane of order q, and let $k_1, \ldots, k_{q^2+q+1}$ be the intersection sizes of the q^2+q+1 lines with the set S. Prove that

$$\sum_{i=1}^{q^2+q+1} \left(k_i - \frac{(q+1)|S|}{q^2+q+1} \right)^2 \leq q|S|.$$

(b) Let S be a set of points and T a set of lines in a projective plane of order q. Define $i(S,T)$ to be the number of incidences between the points of S and the lines of T, that is, $i(S,T) = |I \cap (S \times T)|$, where I is the point-line incidence relation of the projective plane. Prove that

$$\left| i(S,T) - \frac{q+1}{q^2+q+1} |S||T| \right| \leq \sqrt{q|S||T|}.$$

[5] By choosing such sets D in $\mathrm{PG}(2,q^2)$ one can construct non-Desarguesian planes of order q^2. You can try to find such a set and show that the plane you construct is non-Desarguesian, for all $q \geq 3$.

4 Strongly Regular Graphs

4.1 Introdction

In finite geometry, the kind of graphs that we usually encounter are often highly symmetric. In particular, they are regular graphs, i.e., the degree of each vertex is the same. In fact, often the graphs have an even higher form of regularity, captured by the following notion due to R. C. Bose.

Definition 4.1. A non-complete and non-empty graph G is called a *strongly regular graph* with parameters n, k, λ, μ, or more concisely an $\mathrm{srg}(n, k, \lambda, \mu)$, if it is a k-regular graph on n vertices such that

- every pair of adjacent vertices have exactly λ common neighbours, and
- every pair of non-adjacent vertices have exactly μ common neighbours.

We will develop some theory of these graphs, which will be applied to finite geometry in later chapters. Using this theory, we will also see a proof of the famous friendship theorem of Erdős, Rényi and Sós, which says that in any finite society where every two individuals have a *unique* common friend, there must be a person who is everybody's friend. This result also has a surprising connection to finite projective planes.

Some "trivial" examples of strongly regular graphs are the disjoint union of m K_r's, which is an $\mathrm{srg}(mr, r-1, r-2, 0)$, and its complement, the m-partite complete graph with each part of size r, which is an $\mathrm{srg}(mr, (m-1)r, (m-2)r, (m-1)r)$. These are trivial because of the following reason:

Proposition 4.2. *Let G be an $\mathrm{srg}(n, k, \lambda, \mu)$. Then $0 \leq \mu \leq k$, with $\mu = 0$ if and only if G is a disjoint union of equally sized cliques and $\mu = k$ if and only if G is a complete multipartite graph with equal sized parts.*

Proof. Exercise. □

For non-trivial examples, consider C_5, which is an $\mathrm{srg}(5, 2, 0, 1)$ and the Petersen graph, an $\mathrm{srg}(10, 3, 0, 1)$. For infinite families of non-trivial examples, take the line graphs $L(K_n)$ which is an $\mathrm{srg}(n(n-1)/2, 2(n-2), n-2, 4)$, and $L(K_{n,n})$ which is an $\mathrm{srg}(n^2, 2(n-1), n-2, 2)$.[1] Many more examples of strongly regular graphs are given in the exercises, and in later chapters.

Our first results on strongly regular graphs says that the parameters are not entirely independent of each other. In fact, it gives us the first condition for the feasibility of the parameters.

[1] As a graph theory exercises, find all graphs G for which $L(G)$ is a strongly regular graph.

Proposition 4.3. *Let G be an* $\mathrm{srg}(n, k, \lambda, \mu)$*. Then*

$$k(k - \lambda - 1) = (n - k - 1)\mu.$$

Proof. Fix a vertex x of G. Double count the set of pairs of vertices (y, z) such that x is adjacent to y, y is adjacent to z and x, z are non-adjacent. There are k choices for y, and then since x, y have λ common neighbours, there are $k - \lambda - 1$ remaining choices for z. There are $n - k - 1$ choices for z, since x has exactly k neighbours, and then for every such z there are μ choices for y. □

If try to generalise the C_5 and the Petersen graph, both of which had $\lambda, \mu = 0, 1$ by asking for an $\mathrm{srg}(n, k, 0, 1)$, then from the proposition above we see that n must be equal to $k^2 + 1$.

Proposition 4.4. *The complement of an* $\mathrm{srg}(n, k, \lambda, \mu)$ *is an* $\mathrm{srg}(n, n - k - 1, n - 2k + \mu - 2, n - 2k + \lambda)$*.*

Proof. Let G be an $\mathrm{srg}(n, k, \lambda, \mu)$. The complement of G also has n vertices. The number of non-neighbours in G of a vertex is $n-k-1$. Let x, y be two non-adjacent vertices in G. They have μ common neighbours, and therefore the number of vertices that are adjacent to at least one of x or y is equal to $2k - \mu$. This implies that the number of common non-neighbours is equal to $n - 2 - (2k - \mu)$. Let x, y be two adjacent vertices in G. They have λ common neighbours, and hence the set of vertices that is adjacent to at least one of x or y, has cardinality $2(k - 1 - \lambda)$. Therefore, among the $n - 2$ vertices other than x, y there are $n - 2 - 2k + 2 - \lambda$ vertices which are non-adjacent to both x and y. □

4.2 Moore graphs

Strongly regular graphs are sometimes extremal examples of natural graph theoretical questions. One of these is as follows.

In a simple graph G, girth of G is the smallest length of a cycle contained in G. If G is acyclic, then we say that it's girth is ∞. The diameter of G is the largest distance between two vertices of G, where the distance is measured by the length of a shortest path. As one might expect, the graphs with large (but finite) girth, must also have a large diameter. This is captured by the following relation.

Lemma 4.5. *For any connected graph G of diameter d and girth $g < \infty$, we have:*

$$g \leq 2d + 1.$$

Proof. Say $g \geq 2d+2$, and let C be a cycle of length g. Let x, y be two opposite vertices on C. Since the two paths along C from x to y both have length at least $d + 1$, and d is the diameter of the graph, there must be a path P of length at most d from x to y that does not share all of its edges with C. By combining P and any path from x to y along C, we get a closed walk of length at most $d + g/2 < g$, which must contain a cycle. This is a contradiction to the fact that g is the length of the smallest cycle in G. □

A trivial example of graph in which the bound above is tight, is the complete graph on $n \geq 3$ vertices for which $g = 3$ and $d = 1$.

Definition 4.6. A Moore graph is a graph of diameter d and girth $2d+1$, for some $d > 1$.

For example, a C_{2d+1} is a Moore graph for all $d \geq 2$. Can you think of any other examples? The Petersen graph should be a good guess, and indeed it is a Moore graph of diameter 2. Are there more examples?

Lemma 4.7. *Every Moore graph is k-regular, for some $k \geq 2$, and has $1+k+k(k-1)+\cdots+k(k-1)^{d-1}$ vertices, where d is the diameter of the graph.*

Proof. Let d be the diameter of the Moore graph G. The girth of G is then $2d+1$. We first show that any two vertices x, y which are at distance d from each other have equal degrees. Let P the unique path of length d joining x and y. There is a neighbour x_1 of x on P and a neighbour y_1 of y on P. Let x_i be any other neighbour of x. Then $\mathrm{d}(x_i, y) = d$, and the unique path of length d between them gives rise to a neighbour y_i of y. All y_i's must be distinct, as otherwise we will get a cycle of length $\leq 2d$. This shows that $\deg(y) \geq \deg(x)$, and a similar argument shows that $\deg(x) \geq \deg(y)$.

Now let C be a cycle of length $2d+1$. Let x, y be two adjacent vertices on C. Then there exists a vertex z on C at distance d from both x and y. This implies that $\deg(x) = \deg(y) = \deg(z)$, and hence every vertex on C has the same degree. Let k be this degree. Let w be a vertex not on C, at distance i from C. Then there exists a vertex on C which has distance d from w,obtained by taking $d-i$ steps from the nearest vertex in C to w, implying that $\deg(w) = k$.

The number of vertices follows from an easy count that shows that there are exactly $k(k-1)^{i-1}$ vertices at distance i from a fixed vertex. □

Lemma 4.8. *Let G be a Moore graph of diameter 2, and k its regularity. Then G is an $\mathrm{srg}(k^2+1, k, 0, 1)$.*

Proof. We know that G is a k-regular graph on k^2+1 vertices. Let x, y be two adjacent vertices. If they have any common neighbour, then G will contain a triangle, which is impossible since the girth of any Moore graph is at least 5. Now let x, y be two non-adjacent vertices. There exists a path of length 2 between them, since the diameter of the Moore graph is 2. This should be unique path of length 2 as otherwise we will get a cycle of length 4. □

One can prove that Moore graphs of diameter greater than 2 are also "highly regular" in some sense (look up the notion of distance regular graphs), but we will not delve into that and just use the theorem of Bannai and Ito from 1971 states that there are no Moore graphs of diameter $d > 2$, except for the odd cycles C_{2d+1}. Therefore, we are only left with understanding the diameter 2 case, for which the fact that it is a strongly regular graph with the parameters above, turns out to be an extremely strong condition.

4.3 Spectral Methods

Let A be the adjacency matrix of a (simple undirected) graph G on n vertices, i.e., the $n \times n$ real matrix obtained by ordering the vertex set of G and then defining $a_{ij} = 1$ if the i-th vertex is adjacent to the j-th vertex, and 0 otherwise. Since A is a real symmetric matrix of order n, it has n real eigenvalues $\lambda_1 \geq \cdots \geq \lambda_n$, counted with multiplicity, where λ is an eigenvalue of A if there exists a non-zero vector x such that $Ax = \lambda x$. These are precisely the zeros of the degree n polynomial $\phi(t) = \det(tI - A) = 0$, known as the characteristic polynomial of A. So, we have $\phi(t) = \prod_{i=1}^{n}(t - \lambda_i)$. Note that this implies $\sum \lambda_i = \text{Tr } A = 0$ and $\prod \lambda_i = \det A$.

Two graphs G_1, G_2 on the same vertex set $V = \{v_1, \ldots, v_n\}$ are isomorphic if and only if there exist a permutation matrix P such that $P^T A_1 P = A_2$, where A_1 and A_2 are the adjacency matrices of G_1 and G_2. Since $P^T = P^{-1}$ for permutation matrices, we get $\det(tI - A_1) = \det(P^{-1}(tI - A_1)P) = \det(tI - A_2)$. Therefore, isomorphic graphs have the same characteristic polynomial, and hence the same eigenvalues. Moreover, this shows that the eigenvalues of a graph do not depend on the ordering of the vertices we choose to write its adjacency matrix.

Definition 4.9. The spectrum of a graph G is the set of distinct eigenvalues of G, along with their corresponding multiplicities. If the eigenvalues of G are $\lambda_1 > \lambda_2 > \cdots > \lambda_s$, and their multiplicities are $m_1, m_2, \ldots, m_s$, then we write

$$\text{Spec}(G) = \begin{pmatrix} \lambda_1 & \lambda_2 & \ldots & \lambda_s \\ m_1 & m_2 & \ldots & m_s \end{pmatrix}.$$

By direct computation, or by using some of the results that follow, one can easily compute the following spectra:

$$\text{Spec}(K_n) = \begin{pmatrix} n-1 & -1 \\ 1 & n-1 \end{pmatrix}.$$

$$\text{Spec}(K_{m,n}) = \begin{pmatrix} \sqrt{mn} & 0 & -\sqrt{mn} \\ 1 & m+n-2 & 1 \end{pmatrix}.$$

The multiplicity of an eigenvalue λ is also equal to the dimension of the eigenspace $V_\lambda = \{x \in \mathbb{R}^n : Ax = \lambda x\}$. Moreover, the real spectral theorem asserts that there exists an orthonormal basis of $\mathbb{R}^n$ consisting entirely of eigenvectors, or in other words, the matrix A can be diagonalised via a change of basis. In particular, it says that for two eigenvalues $\lambda \neq \mu$, any vector in V_λ is orthogonal to any vector in V_μ, a fact that can be shown easily.

Lemma 4.10. *Let x and y be two eigenvectors of a real symmetric matrix A with $Ax = \lambda x$, $Ay = \mu y$ and $\lambda \neq \mu$. Then $x \cdot y = 0$.*

Proof. From the definition of the dot product and the fact that A is symmetric, we get that $(Ax) \cdot y = y^T Ax = x^T A^T y = x^T Ay = x \cdot (Ay)$. Since $Ax = \lambda x$ and $Ay = \mu y$, we get $\lambda(x \cdot y) = \mu(x \cdot y)$ which implies that $x \cdot y = 0$ since $\lambda \neq \mu$. □

Therefore, if

$$\text{Spec}(G) = \begin{pmatrix} \lambda_1 & \lambda_2 & \ldots & \lambda_s \\ m_1 & m_2 & \ldots & m_s \end{pmatrix},$$

then

$$\mathbb{R}^n = V_{\lambda_1} \perp V_{\lambda_2} \perp \cdots \perp V_{\lambda_s}.$$

While the definition of eigenvalues might suggest that to find the eigenvalues of a graph on n vertices we will be finding the zeros of a degree n polynomial, in practice, this is something we never do. The way we will find eigenvalues is by looking at the eigenvectors graph. One can interpret the equation $Ax = \lambda x$, with $x = (x_1, \ldots, x_n)$ as labelling the n

vertices of G by real numbers $x_1, \dots, x_n$ such that for any vertex of the graph the sum of the labels on its neighbours is λ times the label of the vertex. Therefore, finding an eigenvector is the same as finding such a labelling. This immediately tells us that for any k-regular graph, the all-one vector $(1, \dots, 1)$ is an eigenvector with eigenvalue k. We can prove something stronger.

Lemma 4.11. *For any k-regular graph G, k is the largest eigenvalue and it has multiplicity 1 if and only if G is connected.*
Proof. Let $V(G) = [n]$, A the $n \times n$ adjacency matrix and let $(x_1, \dots, x_n)$ be an eigenvector of A with eigenvalue λ. We first show that $|\lambda| \leq k$. Let m be the vertex for which $|x_i|$ is takes the maximum value. Then looking at the m'th coordinate of the equation $Ax = \lambda x$, we get

$$\sum_{i \in N(m)} x_i = \lambda x_m,$$

where $N(m)$ is the neighbourhood of the vertex m, which by k-regularity has size k. Taking absolute value on both sides, and using the triangle inequality we get

$$|\lambda||x_m| \leq \sum_{i \in N(m)} |x_i| \leq k|x_m|,$$

which implies $|\lambda| \leq k$. Therefore, every eigenvalue of G lies in the range $[-k, k]$, and in particular, since k is an eigenvalue of the graph, it is the largest eigenvalue.

Let G be a disconnected graph, with $C \subset [n]$ as one of its components. Then the vector $(x_1, \dots, x_n)$ with $x_i = 1$ for $i \in C$ and $x_i = 0$ for $i \notin C$ is an eigenvector with eigenvalue k, but it is not a scalar multiple of the all-one vector which implies that the multiplicity of k is at least 2.

Now let G be a connected graph, and let $(x_1, \dots, x_n)$ be any eigenvector with eigenvalue k. We will show that $x_1 = x_2 = \cdots = x_n$. Let m be the vertex for which x_m is the largest. Then since $kx_m = \sum_{i \in N(m)} x_i$, and $x_i \leq x_m$ for all m, we must have $x_i = x_m$ for all $i \in N(m)$. By repeating this we get that $x_i = x_m$ for all i, because the graph is connected. □

Therefore, for connected k-regular graphs on n vertices the eigenvalues are $k = \lambda_1 > \lambda_2 \geq \lambda_3 \geq \cdots \geq \lambda_n \geq -k$. The eigenvalues of a graph can be used to characterise several properties of graphs. You will see some of these in the exercises. For example, you will show that for connected k-regular graphs, $\lambda_n = -k$ if and only if the graph is bipartite. But now let's get back to strongly regular graphs, where we can computer the eigenvalues, and their multiplicities precisely in terms of the parameters of the graphs.

Proposition 4.12. *Let G be an $\mathrm{srg}(n, k, \lambda, \mu)$, with $k \geq \mu > 0$. Then G has three distinct eigenvalues k, θ_1 and θ_2 with multiplicities 1, m_1 and m_2, respectively, where*

$$\theta_1, \theta_2 = \frac{1}{2}\left(\lambda - \mu \pm \sqrt{\Delta}\right)$$

and

$$m_1, m_2 = \frac{1}{2}\left(n - 1 \mp \frac{2k + (n-1)(\lambda - \mu)}{\sqrt{\Delta}}\right)$$

with $\Delta = (\lambda - \mu)^2 + 4(k - \mu)$.

Proof. Since $\mu > 0$, G must be connected, and hence k is an eigenvalue of multiplicity 1. Let θ be an eigenvalue of $A = A(G)$ with $\theta \neq k$, and let x be an eigenvector with $Ax = \theta x$. The fact that G is an $\mathrm{srg}(n, k, \lambda, \mu)$ implies the following:

$$A^2 = kI + \lambda A + \mu(J - A - I),$$

where I is the $n \times n$ identity matrix and J is the $n \times n$ all-one matrix. Evaluating at x on both sides, and using the fact that $Jx = 0$ (since x is orthogonal to the eigenspace corresponding to the eigenvalue k), we get that

$$\theta^2 x = kx + \lambda\theta x - \mu(\theta x + x).$$

Since x is non-zero, this implies that

$$\theta^2 - (\lambda - \mu)\theta - (k - \mu) = 0.$$

The discriminant Δ of this quadratic equation is $(\lambda - \mu)^2 + 4(k - \mu)$, and thus we get the solutions θ_1 and θ_2 as mentioned above.

Let m_1, m_2 be the multiplicities. We know that

$$1 + m_1 + m_2 = n,$$

since there are n eigenvalues counted with multiplicities. Since the sum of all eigenvalues is 0, we get

$$k + \theta_1 m_1 + \theta_2 m_2 = 0.$$

Solving these two equations we get the values of m_1 and m_2. □

From this we see that the spectrum of the Petersen graph, which is an $\mathrm{srg}(10, 3, 0, 1)$, is

$$\begin{pmatrix} 3 & 1 & -2 \\ 1 & 5 & 4 \end{pmatrix}.$$

One of the ways in which we can apply this result to give restrictions on parameters of a strongly regular graph is by observing that the multiplicities m_1 and m_2 must be integers.

Corollary 4.13. *Let G be a Moore graph of diameter 2, with regularity k. Then $k \in \{2, 3, 7, 57\}$.*

Proof. We know that G is an $\mathrm{srg}(k^2 + 1, k, 0, 1)$. The multiplicities of the eigenvalues of G are

$$\frac{1}{2}\left(k^2 \pm \frac{k^2 - 2k}{\sqrt{4k - 3}}\right).$$

One necessary condition for these numbers to be an integer is that $(k^2 - 2k)/\sqrt{4k - 3}$ is an integer. One possibility is that $k^2 - 2k = 0$, i.e, $k = 2$, which is feasible as shown by C_5. Say $k^2 - 2k \neq 0$, then we must have $4k - 3 = m^2$ for some integer m and $k^2 - 2k = (m^2 + 3)^2/16 - 2(m^2 + 3)/4 \equiv 0 \pmod{m}$. Multiplying both sides by 16, we get that $(m^2 + 3)^2 - 8(m^2 + 3) \equiv -15 \equiv 0 \pmod{m}$. Therefore $m = 1, 3, 5$ or 15, which gives us $k = 1, 3, 7$ or 57. Since $k > 1$, as we assume that the graph has finite girth, the only possibilities we are left with are $k = 2, 3, 7, 57$. □

Theorem 4.14 (Friendship Theorem)**.** *Let G be a graph in which any two vertices have a unique common neighbour, then G has a vertex which is adjacent to all other vertices, and as a consequence it is the windmill graph W_n on $2n+1$ vertices for some $n \geq 1$.*

Proof. Say there is no vertex in G that is adjacent to every other vertex. We claim that G is regular. Note that G is C_4-free. Let u, v be two non-adjacent vertices, which exist by our assumption. We know that u, v have a unique common neighbour, call it w. For every other neighbour of u, we get a unique neighbour of v, and these neighbours must be distinct since otherwise we will have C_4 in the graph. Therefore, $\deg(v) \geq \deg(u)$ and by symmetry we get $\deg(u) \geq \deg(v)$. Now let x be any vertex in $V(G) \setminus \{u, v, w\}$. Then x is non-adjacent to at least one of u or v, since otherwise u and v will have more than one common neighbours. Hence $\deg(x) = \deg(u) = \deg(v)$. Finally, since w can't be adjacent to all vertices, there exists a $y \in V(G) \setminus \{x, u, v\}$ with $\deg(w) = \deg(y)$. This shows that G is regular.

We now derive a contradiction. Let n be the number of vertices in G and k its regularity. Then G is an $\mathrm{srg}(n, k, 1, 1)$, which implies that $n = k^2 - k + 1$. The multiplicities of the eigenvalues of G are

$$\frac{1}{2}\left(k^2 - k \pm \frac{k}{\sqrt{k-1}}\right).$$

Therefore, $k - 1 = m^2$ for some integer m and m divides k. This is only possible if $m = 1$, that is $k = 2$. But then G is a triangle, which does have a vertex adjacent to all other vertices, a contradiction.

Therefore, G has a vertex adjacent to all other vertices. We leave it to the reader to now deduce that now G must be the windmill graph. □

An interesting corollary of the Friendship theorem is a classical result of Baer on finite projective planes.

Definition 4.15. Let $\pi = (\mathcal{P}, \mathcal{L}, I)$ be a projective plane. A *polarity* of π is a bijective map $\sigma : \mathcal{P} \cup \mathcal{L} \to \mathcal{P} \cup \mathcal{L}$ such that, $\sigma(\mathcal{P}) = \mathcal{L}$, $\sigma(\mathcal{L}) = \mathcal{P}$, $\sigma^2 = \mathrm{id}$ and $(x, \ell) \in I$ if and only if $(\sigma(\ell), \sigma(x)) \in I$. A point x of π is called an *absolute point* of σ if $x \in \sigma(x)$.

Remark 4.16. Clearly, a polarity of a projective plane π gives us an isomorphism between π and π^D. An arbitrary isomorphism between π and π^D is known as a duality of π, and then a polarity is a duality which is inverse of itself.

Theorem 4.17 (Baer 1946)**.** *Every polarity of a finite projective plane has an absolute point.*

Proof. Let $n \geq 2$ be the order of the plane, and assume that there are no absolute points of a polarity σ. Let G be the graph defined on the n^2+n+1 points by making x adjacent to y if $y \in \sigma(x)$. Since there are no absolute points, this defines a simple graph which is $(n+1)$-regular, as the neighbourhood of a point x is equal to the set of points on $\sigma(x)$. Moreover, since any two distinct lines in the plane intersect in a unique point, G has the property that any two distinct vertices have a unique common neighbour. By the Friendship theorem, G must be the triangle, which is a contradiction to $n \geq 2$. □

Remark 4.18. Baer also proved that in a finite projective plane of order n every polarity has at least $n+1$ absolute points. The map $(a, b, c) \mapsto \{(x, y, z) : ax + by + cz = 0\}$ is a polarity of the Desarguesian projective plane $\mathrm{PG}(2, q)$, that has exactly $q+1$ absolute points, which form a conic if q is odd and a line if q is even. The theorem of Baer can also be used to show that for odd q every irreducible conic in $\mathrm{PG}(2,q)$ has a point.

Remark 4.19. The absolute points form a tangency set because if $y \in \sigma(x)$ for some absolute point x, and $y \neq x$, then since $\sigma(y)$ must contain x it can't contain y, that is, y is not an absolute point. Therefore, the number of absolute points in a projective plane of order n is at most $n\sqrt{n}+1$, with the polarity $(a,b,c) \mapsto \{a^q x + b^q y + c^q z : x,y,z \in \mathbb{F}_{q^2}\}$ in $\mathrm{PG}(2,q^2)$ showing that equality is achieved for infinitely many values of n.

We end this section with a classic result in spectral graph theory that gives a non-trivial upper bound on the independence number of a graph, and hence a lower bound on the chromatic number of the graph.

Theorem 4.20 (Delsarte-Hoffman bound). *Let G be a k-regular graph on n vertices and $k = \lambda_1 \geq \cdots \geq \lambda_n$ its eigenvalues. Then*

$$\alpha(G) \leq \frac{-\lambda_n}{k-\lambda_n} \cdot n.$$

Proof. Let $v_1, \ldots, v_n$ be the vertices of G, and A the adjacency matrix. Let S be an independent set of G, and χ the characteristic vector of S, that is, $\chi(i) = 1$ if $v_i \in S$ and 0 otherwise. Since S is an independent set, we must have $\chi^T A \chi = 0$.[2] Let $u_1, \ldots, u_n$ be an orthonormal basis consisting eigenvectors of A, with $\lambda_1, \ldots, \lambda_n$ as the corresponding eigenvalues. We take u_1 to be the normalised all 1 vector, $(1/\sqrt{n}, \ldots, 1/\sqrt{n})$. Let $\chi = \sum c_i u_i$, for some $c_1, \ldots, c_n \in \mathbb{R}$. Then $c_i = \chi^T u_i$, and hence $c_1 = |S|/\sqrt{n}$. Moreover, $\sum c_i^2 = ||\chi|| = \chi^T \chi = |S|$. Since $A\chi = \sum c_i \lambda_i u_i$, using $\chi^T A \chi = 0$ we get

$$0 = |S|^2 \lambda_1/n + \sum_{i=2}^{n} c_i^2 \lambda_i \geq \lambda_1 |S|^2/n + \lambda_n \sum_{i=2}^{n} c_i^2 = k|S|^2/n + \lambda_n(|S| - |S|^2/n),$$

which implies the inequality

$$|S| \leq \frac{-\lambda_n}{k-\lambda_n} \cdot n.$$

Since S was an arbitrary independent set, this proves the result. □

4.4 Two-intersection sets

In this section we give a finite geometric construction of strongly regular graphs.

Definition 4.21. A set S of points in $\mathrm{PG}(n-1,q)$ is called a two-intersection set[3] if there exist two non-negative integers $k_1 < k_2$ such that for every hyperplane π, we have $|S \cap \pi| \in \{k_1, k_2\}$.

Example 4.22. A hyperoval in $\mathrm{PG}(2,q)$ is a two-intersection set with $k_1 = 0$ and $k_2 = 2$.

Example 4.23. A Baer subplane in $\mathrm{PG}(2,q)$ is a two-intersection set with $k_1 = 1$ and $k_2 = \sqrt{q}+1$.

[2] In general, if S, T are subsets of vertices, and χ_S, χ_T are their characteristic vectors, then $\chi_S^T A \chi_T$ counts the total number of edges between S and T.

[3] These are also referred to has two-character sets in the literature.

More examples in higher dimensions will be given later.

Theorem 4.24 (Delsarte 1972, Calderbank-Kantor 1986)**.** *Let H be a hyperplane in* $\mathrm{PG}(n,q)$*,* $n \geq 3$*, and S a two-intersection set in H. Then the graph $G(S)$, with vertex set equal to* $\mathrm{PG}(n,q) \setminus H$*, and a point x adjacent to a point y if the line joining x and y meets H in a point of S, is a strongly regular graph.*

Proof. Let S be a two-intersection set in H, with parameters k_1 and k_2. The graph is clearly $|S|(q-1)$-regular since there are $|S|$ lines through a point x that meet H in a point of S, and each of them contains $q-1$ points other than x and the point of S the line contains. Let u, v be two non-adjacent vertices of $G(S)$. Then the line uv intersects H in a point $t \notin S$. Let w be a common neighbour of u and v. Then the lines uw and vw meet H in points x and y of S, respectively. The planes spanned by the lines uw and vw meets H in the line xy, and hence t must be collinear with both x and y. Conversely, if x and y are two vertices in S contained on a line through t, then the intersection of ux and vy gives a common neighbour of u and v. This gives a bijectionbetween ordered pairs (x,y) in S for which the line xy contains t, and the common neighbours of S. We will show that the number μ of such ordered pairs is a constant not depending on our choice of u and v.

We count the triplets (x,y,π) where x,y are distinct points of S and π is a hyperplane of H containing t,x and y. Let θ_1 be the number of hyperplanes through t that meet S in k_1 points and θ_2 the number of hyperplanes through t that meet S in k_2 points. Then this number is equal to

$$k_1(k_1-1)\theta_1 + k_2(k_2-1)\theta_2.$$

For the μ pairs (x,y) where the line xy contains t the number of choices for π is $\begin{bmatrix} n-2 \\ n-3 \end{bmatrix}_q = \begin{bmatrix} n-2 \\ 1 \end{bmatrix}_q$ since we are counting the number of hyperplanes through a line and for the remaining $|S|(|S|-1)-\mu$ pairs there are $\begin{bmatrix} n-3 \\ 1 \end{bmatrix}_q$ such choices since we are counting the number of hyperplanes through a plane. Therefore we get

$$\mu \begin{bmatrix} n-2 \\ 1 \end{bmatrix}_q + (|S|(|S|-1)-\mu) \begin{bmatrix} n-3 \\ 1 \end{bmatrix}_q = k_1(k_1-1)\theta_1 + k_2(k_2-1)\theta_2.$$

It now suffices to show that θ_1 and θ_2 do not depend on the choice of t.

Counting the total number of hyperplanes through the point t in $H \cong \mathrm{PG}(n-1,q)$ we get

$$\theta_1 + \theta_2 = \begin{bmatrix} n-1 \\ 1 \end{bmatrix}_q.$$

Counting pairs (x,π) where $x \in S$ and π is a hyperplane of H containing both x and t, we get

$$k_1\theta_1 + k_2\theta_2 = \begin{bmatrix} n-2 \\ 1 \end{bmatrix}_q |S|.$$

Solving these linearly independent equations we get values of θ_1 and θ_2 in terms of the constants $|S|, k_1, k_2, n$ and q, and hence the value of μ.

Now let u, v be any two adjacent vertices, and let $t = H \cap \ell$, where ℓ is the line joining u and v. All the points in $\ell \setminus \{t,u,v\}$ are common neighbours of u and v giving rise to $q-2$ such points. The other common neighbours correspond to pairs (x,y), with $x \neq y \neq t$

such that t lies on the line joining x and y, and $x, y \in S$. If $\max\{k_1, k_2\} \leq 2$, then there are no such neighbours, and we get $\lambda = q - 2$. Otherwise, we get

$$(\lambda-(q-2))\begin{bmatrix} n-2 \\ 1 \end{bmatrix}_q+((|S|-1)(|S|-2)-\lambda+q-2)\begin{bmatrix} n-3 \\ 1 \end{bmatrix}_q = (k_1-1)(k_1-2)\lambda_1+(k_2-1)(k_2-2)\lambda_2,$$

where λ is the number of such neighbours, and λ_1, λ_2 are the solutions of

$$\lambda_1 + \lambda_2 = \begin{bmatrix} n-1 \\ 1 \end{bmatrix}_q$$

and

$$(k_1 - 1)\lambda_1 + (k_2 - 1)\lambda_2 = \begin{bmatrix} n-2 \\ 1 \end{bmatrix}_q (|S| - 1).$$

□

If we take S to be a hyperoval of $\mathrm{PG}(2, q)$, then $G(S)$ is an $\mathrm{srg}(q^3, (q-1)(q+2), q-2, q+2)$. In the next chapter we will see how this graph is in fact a special case of another general construction of strongly regular graphs coming from the collinearity graph of some partial linear spaces.

4.5 Exercises

1. Let G be an $\mathrm{srg}(v, k, \lambda, \mu)$.

   (a) Prove that $\mu \leq k$, with equality if and only if G is a complete multipartite graph.

   (b) Prove that the following are equivalent (i) $\mu = 0$, (ii) $k = \lambda + 1$, (iii) G is a disjoint union of complete graphs, (iv) G is disconnected.

2. Determine the spectrum of the graphs C_n and $K_{m,n}$ for all positive integer m, n.

3. Prove that the following graphs are strongly regular and determine their parameters.

   (a) Given a set of $m - 2$ MOLS of order n, $m \geq 2$, the graph $L_m(n)$ with vertex set $[n]^2$ and two vertices adjacent if they either share a coordinate, or have the same symbol appearing in the coordinates in one of the Latin squares (if $m > 2$).

   (b) The graph on the lines of $\mathrm{PG}(n, q)$, $n \geq 3$, with two lines adjacent if they meet in a point.

   (c) Let $P_1, \ldots, P_5, Q_1, \ldots, Q_5$ be disjoint copies of the cycle graph C_5. Label the vertices of P_i as $\{v_{i1}, v_{i2}, v_{i3}, v_{i4}, v_{i5}\}$ such that v_{ik} is adjacent to $v_{i(k+1)}$. Label the vertices of Q_j as $\{u_{j1}, u_{j2}, u_{j3}, u_{j4}, u_{j5}\}$ such that u_{jk} is adjacent to $u_{j(k+2)}$.[4] For all $1 \leq i, j, k \leq 5$ connect v_{ik} to $u_{jk'}$ with an edge where $k' = ij + k$ (the arithmetic on all indices is done modulo 5).

4. Let q be an odd prime power, and define a directed graph on $\mathbb{F}_q$ by taking (a, b) as an edge if $a - b$ is a square in $\mathbb{F}_q$.

[4] With this labelling P_i's look like pentagons whereas Q_j's look like pentagrams.

(a) Prove that this graph is an undirected graph if and only if $q \equiv 1 \pmod 4$.

(b) Prove that for $q \equiv 1 \pmod 4$, this graph is a strongly regular graph and compute its parameters.

(c) Find the independence number of the graph when $q \equiv 1 \pmod 4$ and q is an even power of a prime.

5. We give further characterisations of the Moore graphs.

   a) Prove that any k-regular strongly regular graph with $\mu \neq 0$ has at most $k^2 + 1$ vertices, with equality if and only if it is a Moore graph of diameter 2.

   b) Prove that a graph on n vertices with girth ≥ 5 has at most $n\sqrt{n-1}$ edges, with equality if and only if it's a Moore graph.

   c) Prove that a k-regular graph of diameter d has at most $1+k\sum_{i=0}^{d-1}(k-1)^i$ vertices, with equality if and only if it's a Moore graph.

6. (a) Let $\mathcal{O}$ be an oval in a finite projective plane of order 5. Prove that the graph defined on the interior points of the plane, that is, the points through which there are no tangents to $\mathcal{O}$, with two points adjacent if the line joining them is a secant of $\mathcal{O}$, is isomorphic to the Petersen Graph.

   (b) Use this to identify all points and lines of the projective plane with certain substructures of the Petersen graph (vertices, edges, 1-factors, and independent sets), and thus prove the uniqueness of the projective plane of order 5 assuming it contains an oval.

   (c) Show that every projective plane of order 5 has an oval.

7. Let G be a graph with eigenvalues $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_n$.

   (a) Prove that

$$\lambda_1 = \max_{||x||=1} x^T Ax = \max_{x\neq 0} \frac{x^T Ax}{||x||^2},$$

   and deduce that $\lambda_1 \geq \overline{d}(G)$, that is, the average degree of G.

   (b) Suppose G is k-regular. Prove that G is bipartite if and only if $\lambda_n = -k$.

8. Let G be a graph with $[n]$ as its vertex set, and let A be its adjacency matrix. Let $C = [2m]$ for some $1 \leq m \leq n/2$, and $D = [n] \setminus C$, such that the following two properties hold: (i) the induced subgraph on the vertex set C is regular and (ii) every vertex in D is adjacent to 0, m or $2m$ vertices in C. Let

$$Q = \begin{pmatrix} \frac{1}{m}J_{2m} - I_{2m} & 0 \\ 0 & I_{n-2m} \end{pmatrix}$$

   (a) Prove that QAQ^T is an adjacency matrix of a simple graph G' such that $\mathrm{Spec}(G') = \mathrm{Spec(G)}$.

   (b) Use this idea to construct a pair of non-isomorphic graphs on 9 vertices that have the same spectrum.

9. Let S be a two-intersection set in $\mathrm{PG}(2, q)$, with every line meeting S in either k_1 or k_2 points, $k_1 < k_2$.

   (a) Show that $k_2 - k_1$ divides q.

   (b) Prove that $|S|$ is one of the roots of the quadratic equation

$$x^2 - (q(k_1 + k_2 - 1) + k_1 + k_2)x + k_1k_2(q^2 + q + 1) = 0.$$

5 Generalized Quadrangles

5.1 Introduction

Recall that a partial linear space is a point-line geometry where through every pair of distinct points there is at most one line. So far we have only seen partial spaces that are linear spaces, that is, through every pair of points there is a *unique* line. We now introduce our first family of proper partial linear spaces, the generalized quadrangles, which are intimately linked with strongly regular graphs, bipartite Moore graphs, and many interesting combinatorial questions.

Definition 5.1. A generalized quadrangle is a partial linear space satisfying the following axioms:

(GQ1) for every non-incident point-line pair x, ℓ there exists a unique point x' incident to ℓ such that x and x' are collinear,

(GQ2) every point is incident with at least two lines.

For example, the quadrangle, with $\mathcal{P} = \{1, 2, 3, 4\}$ and $\mathcal{L} = \{12, 23, 34, 41\}$, is a generalized quadrangle. An easy generalisation of this is the $m \times n$ grid for integers $m, n \geq 1$, where the lines are either the vertical lines (which have n points on them) or the horizontal lines (which have m points on them).

It is easy to see that the point-line dual of a generalized quadrangle is also a generalized quadrangle. The dual of an $m \times n$ grid is simply the complete graph $K_{m,n}$.

We will often abbreviate "generalized quadrangle" to GQ. We already notice that unlike projective planes, the number of points on a line in GQ can be different from the number of lines through a point. Even though we will show that the grid and the dual grid are in some sense degenerate examples, we will also see non-degenerate GQ's where the number of points on a line is not the same as the number of lines on a point. Even so, in the non-degenerate case all lines do still have the same number of points, and all points do have the same number of lines through them, that is, there is some *order*.

Definition 5.2. A generalized quadrangle is said to have an order (s, t), if there exists constants $s, t \in \mathbb{N}$ such that every line is incident with $s + 1$ points and every point is incident with $t + 1$ lines.

If a GQ has order (s, t), then clearly its dual has order (t, s). For example, the symmetrical $n \times n$ grid has order $(n - 1, 1)$ while its dual, $K_{n,n}$ has order $(1, n - 1)$. Are there any GQ's where all the lines have at least three points on them *and* all the points have at least three lines through them? Here is our first such example.

Example 5.3. Let H be a hyperoval in $\mathrm{PG}(2,4)$. Consider the point-line structure where the points are the points of $\mathrm{PG}(2,4)$ not contained in H, and the lines are the secants of $\mathcal{H}$.

The above example can also be described as follows: the points are all the 1-factors of K_6, the lines are the edges of K_6, and a point is incident to a line if the 1-factor contains the edge.[1] This example even has a pretty drawing, which gives it the name "Doily" (see Figure 5.1).

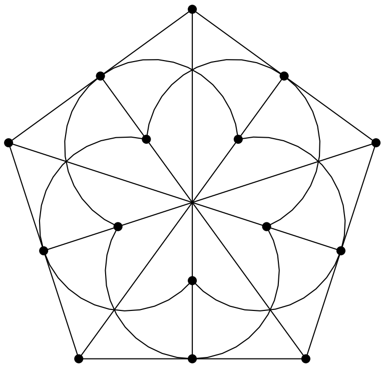

Figure 5.1: Doily, a GQ of order $(2,2)$

Before we proceed with more interesting examples of GQ's, let's prove some basic structural properties.

5.2 Basic properties

Lemma 5.4. *(a) If $\mathcal{S}$ is a generalized quadrangle in which every point is incident with exactly* 2 *lines, then $\mathcal{S}$ is a grid.*

(b) If $\mathcal{S}$ is a generalized quadrangle in which every line is incident with exactly 2 *points, then $\mathcal{S}$ is a dual grid.*

Proof. It suffices to prove the second statement from which the first follows by duality. So let $\mathcal{S}$ be a generalized quadrangle in which every line is incident with exactly two points. We can think of $\mathcal{S}$ as a simple graph on the vertex set equal to the point set $\mathcal{P}$ where two vertices are adjacent if they are collinear. This graph has no triangles. Let x be a point of $\mathcal{S}$. The set $N(x)$ of neighbours of x is an independent set since there are no triangles. Now let $z \neq x$ be a point which is not in $N(x)$. For every $y \in N(x)$, zy must be a line by (GQ2) applied to the point z and the line xy. No two distinct points z_1, z_2 in $\mathcal{P} \setminus (\{x\} \cup N(x))$ can then be collinear as that would give rise to a triangle. Therefore, $\mathcal{S}$ is a complete bipartite graph with $N(x)$ and $\mathcal{P} \setminus N(x)$ as its parts. □

We now show that almost every generalized quadrangle has an order.

Proposition 5.5. *If $\mathcal{S}$ is a generalized quadrangle, then one of the following is true:*

[1] This incidence structure goes back to the work of Sylvester from 1861, and it can also be used to prove that the group S_6 is the only symmetric group S_n which has an outer automorphism.

(a) $\mathcal{S}$ is a non-symmetrical grid.

(b) $\mathcal{S}$ is a non-symmetrical dual grid.

(c) $\mathcal{S}$ has an order (s,t).

Proof. It suffices to show that if $\mathcal{S}$ is not a grid or a dual-grid then it has an order. For a point x, let t_x+1 be the number of lines through x. Let x,y be two non-collinear points of $\mathcal{S}$. For every line through x there exists a unique point on the line collinear with y, giving us a unique line through y. This is a bijection proving $t_x = t_y$.

Now, since $\mathcal{S}$ is not a dual grid, the previous lemma shows that there exists a line ℓ with at least three points, say u,v,w. Let z be any point on a second line through w. Then z is non-collinear with u and v proving that $t_u = t_v = t_z$. Repeatedly applying the same argument we get $t_u = t_v = t_w = t$, and that every point on ℓ is incident with exactly $t+1$ lines. Every point not on ℓ has at least one point in ℓ it is non-collinear to, proving that each point of the generalized quadrangle is incident with $t+1$ lines.

Using the same argument in the dual generalized quadrangle, and using the fact that $\mathcal{S}$ is not a grid, we can show that every line is incident to $s+1$ points for some constant s. □

Proposition 5.6. *Let $\mathcal{S}$ be a generalized quadrangle of order (s,t). The collinearity graph G, defined on the points of $\mathcal{S}$ by making two points adjacent if they are incident to a common line, is an*

$$\mathrm{srg}\left((s+1)(st+1), s(t+1), s-1, t+1\right).$$

Proof. Let x be a point of $\mathcal{S}$. It is incident with $t+1$ lines, each of which has s points besides x, giving us $s(t+1)$ neighbours of x. Let x,y be two collinear points. Then the only points collinear with both of them are the points on the line joining x and y, giving us $s-1$ of them. Let x,y be two non-collinear points. On each of the $t+1$ lines through x, there is a unique point collinear with y, and vice versa, giving us $t+1$ common neighbours between x and y. Therefore, the graph is an $\mathrm{srg}(n, s(t+1), s-1, t+1)$, where now n can be computed using Proposition 4.3. Alternately, let ℓ be a line of the GQ, then the set of points not incident to ℓ is the disjoint union of $N(x)\setminus\{\ell\}$, for $x \in \ell$, giving us $s+1+(s+1)st$ points in total. □

Corollary 5.7. *The eigenvalues of the collinearity graph of a generalized quadrangle of order (s,t) are $s(t+1)$, $s-1$ and $-t-1$, with respective multiplicities*

$$1,\ \frac{st(s+1)(t+1)}{s+t},\ \frac{s^2(st+1)}{s+t}.$$

Corollary 5.8. *If a generalized quadrangle has order (s,t), then $s+t$ divides $st(s+1)(t+1)$.*

Theorem 5.9 (Higman 1971)**.** *If a generalized quadrangle has order (s,t), with $s>1$ and $t>1$, then $s \le t^2$, and dually $t \le s^2$.*

Proof. (Cameron 1975) Let x,y be two non-collinear points. Let V be the set of points which are not collinear to either x or y, and let T be the set of points collinear with both x and y. Then $|T| = t+1$ and $|V| = (s+1)(st+1)-2-2s(t+1)+t+1 = s^2t-st-s+t$.

For the i-th point z of V, let t_i be the number of points of T that z is collinear with. We then have

$$\sum t_i = (t+1)(t-1)s$$

by double counting $(z, u) \in V \times T$, z collinear with u, and

$$\sum t_i(t_i - 1) = (t+1)t(t-1)$$

by double counting $(z, u, u') \in V \times T \times T$, $u \neq u'$ and z collinear with both u and u'. By Cauch-Schwarz inequality we have

$$|V| \sum t_i^2 \geq \left(\sum t_i\right)^2 ,$$

which simplifies to $t(s-1)(s^2-t) \geq 0$. Since $s > 1$, we get $t \leq s^2$. Arguing the same way in the dual GQ we get $s \leq t^2$. □

As an application of the restrictions to the parameters above, let us obtain all values of t for which there exists a GQ of order $(2, t)$. Firstly, we must have $2+t$ dividing $6t(t+1) = 6(t+2-2)(t+2-1)$, which implies that $t+2$ is a divisor of 12. This gives $t \in \{1, 2, 4, 10\}$, but since $t \leq s^2 = 4$, we know that $t = 10$ is not possible. It turns out that for all of these 3 values of t, there is a unique generalized quadrangle of order $(2, t)$.

5.3 Some examples

Example 5.10 (R. W. Ahrens and G. Szekeres 1969). [2] Let H be a hyperoval in $\pi \cong \mathrm{PG}(2, q)$ embedded inside a $\mathrm{PG}(3, q)$, for q a power of 2. The following point-line geometry is a generalized quadrangle of order $(q-1, q+1)$:

- Points are all the points of $\mathrm{PG}(3, q)$ that lie outside π.
- Lines are all lines of $\mathrm{PG}(3, q)$ that are not contained in π and intersect π in a point of H.
- Incidence is containment.

The proof that the above example is a GQ is pretty straightforward. The fact that it's a partial linear space follows from the fact that the line set is a subset of lines of $\mathrm{PG}(3, q)$, which is a linear space. Also note that each line contains $q \geq 2$ points. If we take a line ℓ and a point p not in ℓ, then the span of p and ℓ is a plane that intersects π in a line ℓ'. This ℓ' contains the point x where ℓ meets H. Since H is a hyperoval, ℓ' must contain another unique point y of H. The line py of $\mathrm{PG}(3, q)$ also lies in the plane spanned by p and ℓ, and hence intersects ℓ in a unique point p'. This point p' is the unique point on ℓ which is collinear to p, and hence (GQ2) is satisfied. For (GQ1), we can see that through each point there are $q+2 > 2$ lines.

[2] They came up with this example, and another one of order $(q-1, q+1)$ which works for all prime powers q, as a way to generalized the incidence structure of "27 lines on a cubic surface", which in fact corresponds to the $q = 3$ case.

Remark 5.11. For $q = 2$ this examples gives us the complete bipartite graph $K_{4,4}$, and for all even $q > 2$, we get a non-degenerate generalized quadrangle. Ahrens and Szekeres in fact gave another construction of a GQ of order $(q-1, q+1)$ which works for all prime powers q.

The next example is a GQ of order (q, q) where q is an arbitrary prime power, and the construction is slightly more involved. The proof is left to the reader, as it is merely a careful case analysis.

Example 5.12. Let O be an oval in $\pi \cong \mathrm{PG}(2, q)$ embedded inside a $\mathrm{PG}(3, q)$. The points are one of the following three types:

(i) the points of $\mathrm{PG}(3, q)$ not contained in π;

(ii) the planes of $\mathrm{PG}(3, q)$ which meet O in a unique point;

(iii) a new symbol (∞).

The lines are of two types:

(a) the lines of $\mathrm{PG}(3, q)$ not contained in π that meet O in a point;

(b) the points of O.

A point of type (i) is incident with no line of type (b), and with all those lines of type (a) on which it lies in $\mathrm{PG}(3, q)$. A point of type (ii) is incident with those lines of type (a) and (b) which are contained in it in $\mathrm{PG}(3, q)$. A point of type (iii) is incident with no lines of type (a) and with all lines of type (b).

Example 5.13. Let $\beta : \mathbb{F}_q^4 \times \mathbb{F}_q^4 \to \mathbb{F}_q$ be the bilinear form defined by $\beta(x, y) := x_0y_0 - x_1y_1 + x_2y_2 - x_3y_3$. For a set S of points of $\mathrm{PG}(3, q)$, define $S^\perp = \{y \in \mathrm{PG}(3, q) : \beta(x, y) = 0,\ \forall x \in S\}$. Observe that $\beta(x, x) = 0$ for all x, and $\beta(x, y) = -\beta(x, y)$, which implies that $x \in x^\perp$ for all x, and $x \in y^\perp$ implies $y \in x^\perp$ for all x, y.

Let $\mathcal{P}$ be the set of all points of $\mathrm{PG}(3, q)$ and $\mathcal{L}$ be the set of those lines ℓ which satisfy $\ell \subseteq \ell^\perp$, i.e., $\beta(x, y) = 0$ for all $x, y \in \ell$. Then $(\mathcal{P}, \mathcal{L})$ is a generalized quadrangle of order (q, q).

To see this first observe that a line ℓ joining points x and y satisfies $\ell \subseteq \ell^\perp$ if and only if $\beta(x, y) = 0$. One direction is clear from the definition of $\ell^\perp$, and the other direction follows from the fact that $\beta(\lambda x + \mu y, \lambda' x + \mu' y) = \lambda\lambda'\beta(x, x) + \mu\mu'\beta(y, y) + \lambda\mu'\beta(x, y) + \mu\lambda'\beta(y, x)$ for any $\lambda, \lambda', \mu, \mu' \in \mathbb{F}_q$. Therefore, for any point x, the set of points collinear with it is given by $x^\perp$, which is a plane π in $PG(3, q)$, that contains x. Let ℓ be a line not containing x. Then $x^\perp$, being a plane in $\mathrm{PG}(3, q)$, meets ℓ in a unique point unless $\ell \subseteq x^\perp$. We show that the latter cannot happen.

Say $x \in \mathcal{P}$ and $\ell \in \mathcal{L}$, such that $x \notin \ell$ and $\ell \subseteq x^\perp$. Let y, z be two distinct points on ℓ. Let A be the non-singular 4×4 matrix over $\mathbb{F}_q$ that defines β, that is,

$$A = \begin{pmatrix} 0 & 1 & 0 & 0 \\ -1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & -1 & 0 \end{pmatrix}$$

and $\beta(x,y) = x^T A y$. The vectors $x^T A$, $y^T A$ and $z^T A$ are linearly independent, since x, y and z are linearly independent. Therefore, the null space of the 3×4 matrix B that has $x^T A$, $y^T A$ and $z^T A$ as its rows, is 1-dimensional. Note that this null space is equal to the set of points w for which $\beta(x,w) = \beta(y,w) = \beta(z,w) = 0$, which implies that there is a unique such point. But by our assumption x, y, z all satisfy these three equations, a contradiction.

Remark 5.14. The last two examples are due to Jacques Tits, from his paper of 1959 where he introduced the notion of generalized polygons (of which generalized quadrangles is a special case). He gave many other classical examples, some of which we will study later.

In the homework you will construct a GQ of order (q, q^2) for all prime powers q. Despite 60 years of research, the only known values of (s,t), $s \leq t$, for which we know the existence of a non-degenerate GQ of order (s,t) are: $(q-1, q+1)$, (q,q), (q,q^2) and (q^2, q^3), where q is an arbitrary prime power. Showing the existence of any new GQ of a different order would be a big breakthrough.

Another interesting open problem regarding generalized quadrangles is showing the existence or non-existence of GQ's in which every line as a finite number of points on it, while through each point there are infinitely many lines, the so-called *semi-finite generalized quadrangles*. It is known that $GQ(s, \infty)$ doesn't exist for $s = 2, 3$ and 4, where while the first two cases (due to Cameron and Brouwer) are proved using direct combinatorial methods whereas the last one (due to Cherlin) employs tools from model theory, which is an area of log

5.4 An application to Ball Packings

The celebrated four colour theorem says that the chromatic number of any planar graph is at most 4. How could one generalise this question to higher dimensions? Perhaps, one could ask the chromatic number of a graph that can be embedded in $\mathbb{R}^d$, for $d > 2$, using straight lines. The answer to this question is that the chromatic number of such graphs cannot be bounded from above for any $d \geq 3$, since any graph can be embedded in $\mathbb{R}^3$!

Consider the following question instead. Let $\mathcal{B}$ be a finite collection of disks in $\mathbb{R}^2$ such that the interiors of any two disks are disjoint. Define a graph G, known as the *contact graph* of $\mathcal{B}$, with vertices as the centres of the disks in $\mathcal{B}$ and and two vertices adjacent if the corresponding disks touch each other at the boundary. Then clearly G is planar, and hence $\chi(G) \leq 4$ by the four colour theorem. Determining the chromatic number of such graphs then seems like a weaker version of the four colour theorem. Surprisingly, it is equivalent. A famous result of Koebe, Andreev, and Thurston, known as the *circle packing theorem* says that for any planar graph we can find a collection $\mathcal{B}$ of disks whose contact graph is isomorphic to the planar graph.

Therefore, we can ask the following question for higher dimensions:

What is maximum chromatic number $\chi(d)$ of the contact graph of a finite collection $\mathcal{B}$ of balls in $\mathbb{R}^d$ such that any two balls in $\mathcal{B}$ have disjoint interior?

Such a collection $\mathcal{B}$ is known as a *ball packing*. From four colour theorem we know that $\chi(2) = 4$. In general it's not hard to show that $\chi(d) \geq d + 2$ by picking balls centred at the vertices of a regular simplex, and its barycenter. One can prove an upper bound on $\chi(d)$ in terms of the so-called kissing number $\kappa(d)$, which is the maximum number of unit spheres that can touch a given sphere. Greedily colouring the balls starting from the one with the smallest radius and its neighbours shows that $\chi(d) \leq \kappa(d) + 1$[3], and an easy crude estimate on the kissing number is $\kappa(d) \leq 3^d - 1$.[4] But what about the lower bounds? Can we do better than $d + 2$?

It was shown by Maehara in 2007 that $\chi(d) \geq d+3$, for all $d \geq 3$, and it was observed by Hao Chen in 2014 that some existing results imply that $\chi(d) \geq d + 4$ for all $d = 2^k - 2$, $k \geq 3$. One might suspect that $\chi(d)$ is always $d + c$ for some constant c.[5] Chen showed that this is not true by giving an infinite sequence of d's for which $\chi(d)-d \geq f(d)$ for some function $f(d) = \Theta(d^{2/3})$. The more interesting part of his result for us is that he used finite geometries, and in particular generalized quadrangles, to obtain this result! The method was inspired from the counterexamples to Borsuk's conjecture found by Bondarenko using strongly regular graphs.

The ball packings that we will construct use geometric representations of strongly regular graphs. Let G be an $\mathrm{srg}(n, k, \lambda, \mu)$ with adjacency matrix A and

$$\mathrm{Spec}(G) = \begin{pmatrix} k & \theta_1 & \theta_2 \\ 1 & m_1 & m_2 \end{pmatrix}$$

where $\theta_1 = \frac{1}{2}(\lambda-\mu+\sqrt{\Delta}) > 0$ and $\theta_2 = \frac{1}{2}(\lambda-\mu-\sqrt{\Delta}) < 0$, with $\Delta = (\lambda-\mu)^2+4(k-\mu)$. We will construct a ball packing in $\mathbb{R}^{m_1}$ whose contact graph is G.

Since G is strongly regular, we have

$$A^2 - (\lambda - \mu)A - (k - \mu)I = \mu J. \tag{5.1}$$

We can in fact factorise the left and side to write

$$(A - \theta_1 I)(A - \theta_2 I) = \mu J. \tag{5.2}$$

Consider the matrix

$$E = (A - \theta_2 I)(I - J/n)$$

Lemma 5.15. *Columns of E are eigenvectors of A with eigenvalue θ_1.*

Proof. It suffices to show that $(A - \theta_1 I)E = 0$. Plugging in the value of E, we get

$$(A - \theta_1 I)E = (A - \theta_1 I)(A - \theta_2 I)(I - J/n)$$

which we can simplify using Equation 5.2, to get

$$\mu J(I - J/n) = \mu(J - J^2/n) = \mu(J - J) = 0,$$

[3] Note that at any step the neighbours of a an uncoloured ball of smallest radius r will have radius at least r.

[4] Better exponential bounds are known for $\kappa(d)$, and its exact value is known only for $d = 1, 2, 3, 4, 8, 24$. There is also a lower bound on $\kappa(d)$ which is exponential.

[5] Though looking at the known upper bounds one might suspect this is far away from the truth.

where we have used that $J^2 = nJ$. □

Therefore, we can see that the column vectors live in an m_1 dimensional subspace of $\mathbb{R}^n$. These vectors, or points in $\mathbb{R}^{m_1}$, have the useful property that they form a 2-distance set.

Lemma 5.16. *Let $u_1, \dots, u_n$ be the columns of E. Then*

$$u_i^T u_j = \begin{cases} a + c, & \text{if } i = j \\ b + c, & \text{if } A_{ij} = 1 \\ c, & \text{if } A_{ij} = 0 \text{ and } i \neq j \end{cases}$$

for some numbers a, b, c such that $a, b \geq 0$.

Proof. To check the dot products we look at the matrix $E^T E$, whose ij'th entry is equal to $u_i^T u_j$. Since $E = (A - \theta_2 I)(I - J/n) = A - \theta_2 I - kJ/n + \theta_2 J/n = A - \theta_2 I - (k - \theta_2)J/n$, we have

$$E^T E = E^2 = A^2 + \theta_2^2 I + (k - \theta_2)^2 J/n - 2\theta_2 A - 2(k - \theta_2)kJ/n + 2\theta_2(k - \theta_2)J/n$$

Simplifying and using Equation 5.1 for the value of A^2, we get

$$E^T E = kI + \lambda A + \mu(J - I - A) - 2\theta_2 A + \theta_2^2 I - (k - \theta_2)^2 J/n,$$

where the right hand side can be written as $aI + bA + cJ$, with $a = k - \mu + \theta_2^2$, $b = \lambda - \mu - 2\theta_2$, and $c = \mu - (k - \theta_2)^2/n$. □

This lemma implies that the set $u_1, \dots, u_n$ of points define a two-distance set since $\mathrm{d}(u_i, u_j) = ||u_i - u_j|| = \sqrt{||u_i||^2 + ||u_j||^2 - 2u_i \cdot u_j}$. Moreover, since $b + c \geq c$, the distance is smaller when the i-th vertex is adjacent to the j-th vertex. Therefore, we can define a ball packing in $\mathbb{R}^{m_1}$ with balls of radius equal to half of the smaller distance, centred around u_i's, such that the contact graph is isomorphic to G. So all we need now is the chromatic number of G to be much larger than m_1. To estimate this we use Hoffman's bound, which says that $\chi(G) \geq 1 - k/\theta_2$. There are several examples of strongly regular graphs for which $1 - k/\theta_2 > m_1 + 3$, as one can check on the online database of Brouwer[6]. We look at the *compliment* of the collinearity graph of a generalized quadrangle of order (q, q^2), which is a strongly regular graph with parameters

$$\big((q+1)(q^3+1), q^4, q(q-1)(q^2+1), (q-1)q^3\big),$$

and spectrum

$$\begin{pmatrix} q^4 & q^2 & -q \\ 1 & q^3 - q^2 + q & q^4 + q^2 \end{pmatrix}$$

The lower bound on the chromatic number turns out to be $q^3 + 1$, whereas the dimension m_1 is equal to $q^3 - q^2 + q$. Therefore, $\chi(d) - d \geq q^2 - q + 1$, for all $d = q^3 - q^2 + q$, where q is a prime power.

Question: Can this bound be improved further using strongly regular graphs, perhaps by finding a family where the chromatic number is much larger than the one given by the Delsarte-Hoffman bound?

[6] `https://www.win.tue.nl/~aeb/graphs/srg/srgtab.html`

The following result puts a limit on the best improvement on the lower bound that one can hope for by using strongly regular graphs, while also providing a useful restriction on the parameters.

Theorem 5.17. *Let S be a two-distance set in $\mathbb{R}^d$. Then $|S| \leq (d+2)(d+1)/2$.*

This bound can be improved to $(d+2)(d+1)/2 - 1$ if we assume that all the points lie on a sphere, which happens to be the case with the two-distance set we have constructed from a strongly regular graph.

Corollary 5.18. *For any strongly regular graph on n vertices, with*

$$\mathrm{Spec}(G) = \begin{pmatrix} k & \theta_1 & \theta_2 \\ 1 & m_1 & m_2 \end{pmatrix},$$

we have $n \leq m_1(m_1+3)/2$.

This in particular implies that the number of vertices, and hence the chromatic number, in the contact graph that we construct using a strongly regular graph is at most a quadratic function of the dimension, and hence strongly regular graphs can only be used to prove $\chi(d) \geq \Theta(d^2)$. It is not clear what the actual asymptotic of $\chi(d)$ should be.

5.5 Generalized polygons

We saw earlier that a graph of diameter d has girth at most $2d+1$, and defined the graphs meeting this bound as Moore graphs. If we add the condition that the graph under consideration is bipartite, that is, it has no odd cycles, then the girth g cannot be $2d+1$, and hence $g \leq 2d$. What can we say about bipartite graphs of diameter d and girth $2d$?

For $d = 2$, complete bipartite graphs are examples of such graphs, and in fact the only ones. For any $d \geq 2$, the cycle C_{2d} is another such example. Are there some "non-trivial" examples? We will see that many such bipartite graphs come from finite geometry.

Definition 5.19. The incidence graph of a point-line geometry $(\mathcal{P}, \mathcal{L}, I)$ is a bipartite graph with parts $\mathcal{P}$ and $\mathcal{L}$ with an edge between $x \in \mathcal{P}$ and $\ell \in \mathcal{L}$ if $(x, \ell) \in I$.

Lemma 5.20. *The incidence graph of a partial linear space has girth at least* 6.

Proof. Say there exists a cycle of length 4, given by x, ℓ, y, m where x, y are points, and ℓ, m are lines. Then x and y are contained in two lines, a contradiction to the fact that we have a partial linear space. □

Proposition 5.21. *The incidence graph of a projective plane has diameter* 3 *and girth* 6.

Proof. Any two points have a line through them, and any two lines intersect in a point. Therefore, two vertices in the same part of the incidence graph have distance 2 between each other. Let x be a point and ℓ a line. If $x \in \ell$, then the distance between them in the incidence graph is 1, and if $x \notin \ell$, then since x is at distance 2 from any point on ℓ, the distance between x and ℓ is 3, which shows that the diameter is 3. To get a cycle of length 6, consider a triangle $x_1, \ell_1, x_2, \ell_2, x_3, \ell_3$ in the projective plane. □

Proposition 5.22. *The incidence graph of a generalized quadrangle has diameter* 4 *and girth* 8*.*

Proof. Left to the reader. □

We will now show that being a bipartite Moore graphs of the corresponding diameter characterises projective panes and generalized quadrangles. But first we prove some structural results on bipartite Moore graphs.

Lemma 5.23. *Let G be a bipartite graph of diameter d and girth $2d$. Then the following hold:*

(a) *For any two vertices u and v at distance $< d$, there is a unique shortest path joining them.*

(b) *For any two vertices u and v at distance d from each other,* $\deg(u) = \deg(v)$*.*

(c) *The minimum degree of G is at least* 2*.*

Proof. For (a) observe that two distinct paths of length $< d$ would then give rise to a cycle that has length strictly less than $2d$.

Now let u, v be two vertices at distance d from each other. Let u' be a neighbour of v. If $\mathrm{d}(u', v) = d$, then we will get an odd cycle of length $2d + 1$ since the girth is $2d$. This contradicts the fact that we have a bipartite graph, and hence $\mathrm{d}(u', v) = d - 1$. By (a) there exists a unique path of length $d - 1$ between u_i and v, which contains a neighbour v' of v as its last vertex. These neigbours of v have to be distinct as otherwise we will get a cycle of length $< 2d$. Therefore, $\deg(v) \geq \deg(u)$. Similarly we get $\deg(u) \geq \deg(v)$.

Let C be a cycle of length $2d$. Every vertex on C has degree at least 2. Let x be a vertex outside C, and let P be the shortest path from x to C, that has length i. By going $d - i$ steps on C we get a vertex y that has distance d from x, and hence $\deg(y) = \deg(x)$ by (b). This shows that $\deg(x) \geq 2$. □

Theorem 5.24. *A bipartite graph has diameter* 3 *and girth* 6 *if and only if it is the incidence graph of a possibly degenerate projective plane.*

Proof. We have already shown one side of this claim. Now let G be such a graph. By Lemma 5.23, every vertex in G is of degree at least 2, and hence we can identify the parts of G as points and lines to get a point-line geometry[7], with G as its incidence graph. Let u, v be two vertices in one of the parts, that is, two points or two lines. The distance between them must be even since the graph is bipartite, and it can't be ≥ 4 since the diameter is 3. Therefore, $\mathrm{d}(u, v) = 2$. Moreover, they must have a unique common neighbour as otherwise we will get a cycle of length 4, which is not possible in a graph of girth 6. Thus, the axioms of a possibly degenerate projective plane are satisfied. □

Theorem 5.25. *A bipartite graph has diameter* 4 *and girth* 8 *if and only if it is the incidence graph of a generalized quadrangle.*

Proof. Let G be a bipartite graph of diameter 4 and girth 8. Form lemma 5.23, we know that every vertex has degree 2. The fact that there are no cycles of length 4 implies that we have an incidence graph of a partial linear space. We just need to check (GQ1). Let u be a vertex in one part and v in the other part, such that v is not adjacent to u. This corresponds to a non-incidenct point-line pair of the point-line geometry. As the diameter

[7] recall that in our definition we required each line to be incidence to at least 2 points

is 4, we have $\mathrm{d}(u, v) \leq 4$, but the distance must be odd and hence $\mathrm{d}(u, v) \geq 3$. There must be equality as u is not adjacent to v. This gives a unique neighbour of v which is at distance 2 from u (by Lemma 5.23), proving (GQ1). □

These results motivate the following definition.

Definition 5.26. A generalized d-gon, $d \geq 2$, is a point-line geometry whose incidence graph has diameter d and girth $2d$.

These point-line geometries were introduces by Jacques Tits in 1959, whose motivation came from a geometric study of Lie Groups.

It turns out that generalized d-gons, for $d > 4$, also have the property that either they are "degenerate", and can be easily described, or have an order. We prove this after a sequence of structural lemmas. In all of them, assume that we have a generalized d-gon, and we refer to the points/lines as vertices of the incidence graph. All distances between points, or lines, will be measured in this incidence graph.

Lemma 5.27. *Any two vertices lie in a cycle of length* $2d$*.*

Proof. Let x, y be two vertices, and let P be the shortest path connecting them. Extend P, using Lemma 5.23(c) and the fact the girth is at least $2d$, to get to a path of length d with z as the other end point. Let z' be a neighbour of z which is not in this path. It has distance $d-1$ from x, and hence a path of length $d-1$ between x and z' gives a cycle of length $2d$ containing x and y. □

Definition 5.28. A vertex u is called thick if it has degree at least 3. A generalized polygon is called thick, every vertex has degree at least 3.

Lemma 5.29. *Let C be a cycle of length $2d$, then any two vertices on C that are at the same distance from a thick vertex of C have the same degree.*

Proof. Let x be the thick vertex, and y, z vertices on the cycle that are at equal distance i from x. Let x' be the point at distance d from x on the cycle. Since x is thick, there is a neighbour u of x which is not on C. The distance between u and x' is $d-1$ and hence there exists a path P between them which is disjoint from C except for sharing the vertex x'. This gives us three internally disjoint paths between x and x'. On P take a vertex w at distance $d-i$ from x. Then $\mathrm{d}(w, y) = \mathrm{d}(w, z) = d$, and hence by Lemma 5.23(b), $\deg(y) = \deg(z)$. □

Lemma 5.30. *The minimum distance k between any two thick vertices is a divisor of d. If d/k is even, then thick vertices have at most two degrees. If d/k is odd, then every thick vertex has the same degree.*

Proof. If $k = d$, then this is a consequence of Lemma 5.23(b). So assume that $k < d$. Let u, v be two thick vertices at distance k, from each other, and let C be a cycle of length $2d$ containing them. Applying the previous lemma starting at v, we get that every k-th vertex on C starting from u is thick. Let $x_0 = u$, $x_1 = v$, …, x_i, …, be the thick vertices on C obtained this way, with x_i being the vertex on C at distance ki from u. The vertex u' at distance d from u on C has the same degree as u and hence is thick. Write $d = mk + r$, for some $0 \leq r < k$. If $r > 0$, then x_m is at distance $0 < r < k$ from u', which contradicts the minimality of k. Therefore, $d = mk$, that is, $x_m = u'$. The same argument yields that $x_0 = u$, $x_1 = v$, …, $x_m = u'$, …, x_{2m-1} are all the thick vertices on this cycle. Again from the previous lemma, we see that $\deg(x_i) = \deg(x_{i+2})$ for all i, with arithmetic modulo $2m$. This implies that any thick vertex of C has degree equal to $\deg(u)$ or $\deg(v)$. Moreover, if $m = d/k$ is odd, then $\deg(x_m) = \deg(x_0)$ which implies that $\deg(u) = \deg(v)$ and hence every thick vertex has the same degree.

We can now finish the proof by showing that for any arbitrary thick vertex x, there is a cycle of length $2d$ containing x, u and v. WLOG say u is closer to x than v, and let P be a shortest path between x and u. Extend P to a path P' of length d by moving around C in the direction from u to v. Let y be the end point of P', and z the point on C after y. Then z is at distance $d-1$ from x, and hence there is a path from x to z which is internally disjoint from P', giving us a cycle C' of length $2d$ containing P'. The cycle C' which contains $z \neq u$ can only meet C in a path between two thick vertices, and since there are no thick vertices between u and v, the vertex v must belong to C'. □

Corollary 5.31. *Every thick generalized d-gon has an order (s,t), $s > 1$, $t > 1$. Moreover, if d is odd, then $s = t$.*

The main result for thick generalized polygons is the following theorem of Feit and Higman, that can be proved using the theory of distance regular graphs, which we haven't introduced in this course, and hence we will skip the proof.

Theorem 5.32 (Feit-Higman 1964)**.** *If G is a thick generalized d-gon, then $d \in \{3,4,6,8\}$.*

We will prove a special case instead, using our knowledge of strongly regular graphs.

Proposition 5.33. *There are no thick generalized 5-gons.*

Proof. Say such a generalized 5-gon exists, and let (s,s) be its order. We will show that s must be equal to 1. Let G be the collinearity graph on points. Then it is an srg of parameters $(s^4+s^3+s^2+s+1, s(s+1), s-1, 1)$. Therefore, the eigenvalues are

$$\theta_1, \theta_2 = \frac{1}{2}(s - 2 \pm s\sqrt{5}).$$

Since the sum of eigenvalues with multiplicities is equal to 0, we must have $m_1 = m_2$ as otherwise we will have an irrational number equal to a rational number. Therefore, we have $m_1 = m_2 = (n-1)/2$. By taking the sum of the eigenvalues, we now get

$$s(s+1) + \frac{1}{2}s(s+1)(s^2+1)(s-2) = 0.$$

If we take $s \geq 2$, then the left hand side is > 0.

Alternately, the multiplicity of θ_1 is

$$\frac{1}{2}\left((s^2+1)(s+1) - \frac{2s(s+1) + (s^2+1)(s+1)(s-2)}{s\sqrt{5}}\right),$$

which is an integer if and only if $s = 1$. □

A k-fold subdivision of a graph G is the graph obtained by replacing each edge of G with a path of length k. It turns out that the non-thick generalized polygons can be described via subdivisions of thick generalized polygons.

Theorem 5.34. *The incidence graph of a generalized polygon that is not thick is either an even cycle, the k-fold subdivision of a multiple edge, or the k-fold subdivision of the incidence graph of a thick generalized polygon for some $k \geq 2$.*

Proof. Say G is the incidence graph of a generalized d-gon. Say there are no thick vertices, then G is C_{2d}. If there is one thick vertex, then there are at least 2 since there is a cycle C_{2d} containing this vertex in which the opposite point has the same degree. Let k be the minimum distance between two thick vertices of G.

Say $k = d$, and let u, v be two thick vertices at distance d from each other. As in the proof of the Lemma 5.30, any thick vertex $x \neq u, v$ must be contained in a cycle of length $2d$ with u and v, which will contradict the minimality of $k = d$. Hence, there are no other thick vertices. Both u and v must be contained in every C_{2d} of G, and hence G looks like a collection of (at least three) internally disjoint paths of length $k = d$ from u to v, that is, it is a k-fold subdivision of a multiple edge.

Now assume that $k < d$. Then there exist at least 3 thick vertices. Define a graph G' on thick vertices by making two of them adjacent if they are at distance k from each other. We claim that G' is the incidence graph of a generalized d'-gon, with $d' = d/k$, and G is a k-fold subdivision of G'.

Two distinct paths can only have thick vertices in common, and hence G is indeed a k-fold subdivision of G'. We first show that every cycle of length $2d$ in G is a k-fold subdivision of a cycle of length $2d'$ in G'. In particular, this will show that every vertex in G' has degree at least 3. Let C be a cycle of length $2d$. All we need to show is that there are two thick vertices at distance k from each other on C. There is a vertex on C that is at distance d from a thick vertex, and hence C contains at least one thick vertex x. The vertex y opposite to x must also be thick, and hence C is the union of two internally disjoint paths P_1 and P_2, between x and y. Now let u, v be two thick vertices at distance k from each other, and C' a cycle of length $2d$ containing x, u and v. Any neighbour of x on C' is a vertex at distance $d-1$ from y. Let P_3 be a shortest path from y to this vertex. Then P_3 must contain the thick vertex of C' next to x. $P_3 \cup P_2$ contains two thick vertices at distance k from each other, and hence $P_1 \cup P_2 = C$ also does.

The diameter of G' is at least d' since on any two thick vertices at distance d from each other on a cycle C_{2d} of G are at distance d' from each other in G'. Now let x, y be two thick vertices of G that are at distance $> d'$ from each other in G'. Let z be a vertex on a shortest path from x to y which is at distance d' from x. The thin neighbour z' of z, closer to y, must be at distance $d-1$ from x in G giving rise to a cycle C_{2d} containing x, z, z'. Let w be the thick vertex on the nearest thick vertex on this cycle for which z' is between w and z. Then z, w, y show that G' is not a subdivision of G, a contradiction. We know that G' has a cycle of length $2d'$, and any cycle of smaller length will give rise to a cycle of length smaller than $2d$ in G, a contradiction.

All that is left to show now is that G' is bipartite. Let C be a minimum length odd cycle in G', necessarily of length at least $2d'+1$. Consider two vertices x, y at distance d' from each other in G' on this cycle. Looking at the two thin neighbours of y in G, which are at distance $d-1$ from G, gives rise to a cycle of length $2d$ in G, which must then be a subdivision of a cycle of length $2d'$ in G', and that gives an odd cycle of smaller length in G', a contradiction. □

The following are further restrictions known on the order of thick generalized polygons.

Theorem 5.35. *Let $\mathcal{S}$ be a finite generalized d-gon of order (s,t), with $s,t \geq 2$.*

(a) If $d = 4$, then $s \leq t^2$ and $t \leq s^2$. (Higman)

(b) If $d = 6$, then st is a square, $s \leq t^3$ and $t \leq s^3$. (Haemers and Roos)

(c) If $d = 8$, then $2st$ is a square, $s \leq t^2$ and $t \leq s^2$. (Higman)

The known generalized hexagons have orders (q,q), (q,q^3) and (q^3,q) where q is a prime power. The known generalized octagons have orders (q,q^2) and (q^2,q) where $q = 2^{2k+1}$, for some $k \in \mathbb{N}$.

5.6 Exercises

1. Let H be a hyperoval in $\pi \cong \mathrm{PG}(2,q)$ embedded in $\mathrm{PG}(3,q)$, for q even. Define the following point-line geometry $\mathcal{S}(H,\{x,y\})$, where x,y are distinct points on H. The points are the points of $\mathrm{PG}(3,q)$ not contained in π, the planes through x not containing y, and the planes through y not containing x. The lines are the lines of $\mathrm{PG}(3,q)$ not contained in π and meeting π in $H \setminus \{x,y\}$. A point and a line are incident if they are incident as the objects in $\mathrm{PG}(3,q)$. Prove that $\mathcal{S}(H,\{x,y\})$ is a generalized quadrangle, and determine its order.

2. Find all possible values of t for which there exists a generalized quadrangle of order $(2,t)$. Prove that there is a unique generalized quadrangle of order $(2,2)$.

3. An ovoid in $\mathrm{PG}(3,q)$ is a set of q^2+1 points, no three of which are collinear. Let O be an ovoid in a hyperplane of $\mathrm{PG}(4,q)$. Give a construction of a generalized quadrangle of order (q,q^2) using O, and prove the correctness of your construction.

4. Let $\beta : \mathbb{F}_{q^2}^4 \times \mathbb{F}_{q^2}^4 \to \mathbb{F}_{q^2}^4$ be defined as follows,

$$\beta((x_0,x_1,x_2,x_3),(y_0,y_1,y_2,y_3)) = x_0y_0^q + x_1y_1^q + x_2y_2^q + x_3y_3^q.$$

   A subspace S of $\mathrm{PG}(3,q^2)$ is called totally isotropic with respect to β if $\beta(x,y) = 0$ for all $x,y \in S$. Prove that the point-line geometry of the totally isotropic points, and the totally isotropic lines in $\mathrm{PG}(3,q)$ is a generalized quadrangle. Also determine the order of this generalized quadrangle.

5. Let G be a non-empty finite simple graph with the following property:

   Every edge xy of G is contained in a triangle x,y,z with the property that any vertex $u \notin \{x,y,z\}$ is adjacent to exactly one of x,y,z.

   (a) Prove that G is either the Windmill graph, or a k-regular graph with $k \in \{4,6,10\}$.

   (b) For each of these values of k, construct such a graph G.

   (Bonus) Prove (a) without assuming that the graph is finite.

6. Borsuk's conjecture[8] states that any bounded subset S of $\mathbb{R}^d$, consisting of at least 2 points, can be partitioned into $d+1$ subsets of smaller diameter, for all $d \geq 2$.[9]

   (a) Let G be a strongly regular graph on n vertices with
   $$\mathrm{Spec}(G) = \begin{pmatrix} k & \theta_1 & \theta_2 \\ 1 & m_1 & m_2 \end{pmatrix}.$$
   Construct a finite set S of n points in $\mathbb{R}^{m_1}$ that cannot be partitioned into less than $n/\omega(G)$ parts of smaller diameter, where $\omega(G)$ is the clique number of G.

   (b) Use the existence of an $\mathrm{srg}(416, 100, 36, 20)$ with clique number equal to 5 to disprove Borsuk's conjecture.[10]

7. A connected graph G of diameter d is called distance regular if there exist constants a_i, b_i, c_i, with $i \in \{0, 1, \ldots, d\}$ such that the following holds for any two vertices x and y of G at distance i from each other:

   - there are a_i neighbours of y at distance i from x;
   - there are b_i neighbours of y at distance $i+1$ from x;
   - there are c_i neighbours of y at distance $i-1$ from x.

   Clearly $a_0 = c_0 = b_d = 0$ and G is k-regular with $k = a_0 + b_0 + c_0 = \cdots = a_d + b_d + c_d$.

   (a) Prove that the collinearity graph of a generalized n-gon of order (s,t) is a distance regular graph of diameter $\lfloor \frac{n}{2} \rfloor$.

   (b) Determine the number of points and lines in a generalized n-gon of order (s,t).

8. (a) Prove that a k-regular graph of diameter d has at most
   $$1 + k \sum_{i=0}^{d-1} (k-1)^i$$
   vertices, with equality if and only if the graph is a Moore graph.

   (b) Prove that a k-regular bipartite graph of diameter d has at most
   $$2 \sum_{i=0}^{d-1} (k-1)^i$$
   vertices, with equality if and only if the graph is the incidence graph of a generalized d-gon of order $(k-1, k-1)$.

9. A *near* $2d$*-gon*, for $d \geq 2$, is a partial linear space whose collinearity graph has diameter d, and for every point x and every line ℓ, there exists a unique point $x' \in \ell$ that is closest to x in the collinearity graph, among all points of ℓ.

   (a) Prove that every generalized $2d$-gon is a near $2d$-gon.

[8] now known to be false in general

[9] You should try to see why d such subsets would clearly not suffice.

[10] This strongly regular graph is related to many interesting finite simple groups and finite geometries, but it's difficult to give a well motivated construction here.

(b) Prove that every near $2d$-gon that satisfies the following two properties is a generalized $2d$-gon: (i) every point is incident with at least two lines, and (ii) for every pair of points x and y at distance $i \geq 2$ from each other in the collinearity graph there is a unique neighbour of y at distance $i-1$ from x.

10. Let π be the Fano plane $\mathrm{PG}(2,2)$. Construct the following point-line geometry $\mathcal{S} = (\mathcal{P}, \mathcal{L})$.

- $\mathcal{P}$ is equal to the set of all points, lines, and all point-line pairs (p, ℓ) of π.
- $\mathcal{L}$ is equal to the set of all the following 3-element subsets of $\mathcal{P}$: for every incident point-line pair (p, ℓ) in π, the sets $\{p, \ell, (p, \ell)\}$, $\{(p, \ell), (p_1, \ell_1), (p_2, \ell_2)\}$ and $\{(p, \ell), (p_1, \ell_2), (p_2, \ell_1)\}$, where p, p_1, p_2 are the points on ℓ in π and ℓ, ℓ_1, ℓ_2 are the lines through p in π.

(a) Determine the number of points and the number of lines in $\mathcal{S}$ and prove that it is a partial-linear space of order $(2,2)$.

(b) Prove that $\mathcal{S}$ is a generalized hexagon.

11. (a) Prove that a bipartite k-regular graph of girth $g \geq 4$ has at least

$$2 \sum_{i=0}^{(g-2)/2} (k-1)^i,$$

vertices, and for every $g \in \{4, 6, 8, 12\}$ the bound is sharp for infinitely many values of k.

(b) For every prime power q, construct a bipartite q-regular graph of girth 6 with $2q^2$ vertices, and of girth 8 with $2q^3$ vertices.

(c) (Bonus) Improve the constructions from (b) to get as close as you can to the lower bounds of $2(q^2 - q + 1)$ and $2(q^3 - 2q^2 + 2q)$, respectively.

6 Polar Spaces

6.1 Introduction

Definition 6.1. A polar space is a partial linear space satisfying the following axioms:

(PS1) for every point x and every line ℓ not incident with x, either exactly one point of ℓ is collinear with x or all points of ℓ are collinear with x.

(PS2) there is no point which is collinear with all points.

A partial linear space satisfying (PS1) but not (PS2) is called a *degenerate* polar space. In particular, any linear space is a degenerate polar space. Generalized quadrangles are examples of non-degenerate polar spaces. We will see that there are several other examples, and as the name suggests, many (though not all) of them come from polarities.

Definition 6.2. A duality of $\mathrm{PG}(n, F)$ is a bijection δ on the collection of all subspaces that maps every k-space to an $(n-k-1)$-space, such that $S \subseteq T$ if and only if $T^\delta \subseteq S^\delta$.[1] A duality δ is called a polarity if $\delta = \delta^{-1}$.

We denote polarities by $\perp$, a notation which will become clearer when we talk about orthogonal polar spaces. Note that if $\perp$ is a polarity then $S \subseteq T^\perp$ implies that $T \subseteq S^\perp$, and in particular for two points x and y, x lies in the hyperplane $y^\perp$ if and only if y lies in the hyperplane $x^\perp$.

Lemma 6.3. *If S and T are two subspaces of $\mathrm{PG}(n, F)$ and $\perp$ is a polarity then $(S+T)^\perp = S^\perp \cap T^\perp$, and $(S \cap T)^\perp = S^\perp + T^\perp$.*

Proof. Since S, T are subsets of $S+T$, we get $S^\perp \supseteq (S+T)^\perp$ and $T^\perp \supseteq (S+T)^\perp$. Therefore, $S^\perp \cap T^\perp \supseteq (S+T)^\perp$. Since $S^\perp \cap T^\perp$ is a subset of both $S^\perp$, and $T^\perp$, we get $(S^\perp \cap T^\perp)^\perp \supseteq S+T$, and hence $(S+T)^\perp \subseteq S^\perp \cap T^\perp$.

The proof of the other claim is obtained by replacing S with $S^\perp$ And T with $T^\perp$. □

Corollary 6.4. *A polarity is completely determined by its action on the points.*

Definition 6.5. Let $\perp$ be a polarity of $\mathrm{PG}(n, F)$. A projective subspace S is called isotropic if $S \cap S^\perp \neq \emptyset$, totally isotropic if $S \subseteq S^\perp$ and non-isotropic if $S \cap S^\perp = \emptyset$.

Note that for a point the notion of being isotropic is the same as being totally isotropic. Such points are also referred to as absolute points, as we did earlier when we defined polarities of projective planes.

[1] In other words, it's an incidence reversing bijection on the subspace lattice.

Lemma 6.6. *If S is a totally isotropic subspace, then every subspace contained in S is also totally isotropic.*

Proof. Let $T \subseteq S$, be a subspace. Then $S^{\perp} \subseteq T^{\perp}$. And since $S \subseteq S^{\perp}$, we get $T \subseteq S^{\perp}$. □

Lemma 6.7. *Let x, y be two totally isotropic points with respect to a polarity $\perp$. Then the line $\ell = xy$ is totally isotropic if and only if $x \in y^{\perp}$.*

Proof. We have $\ell^{\perp} = x^{\perp} \cap y^{\perp}$. If $x \in y^{\perp}$, then $y \in x^{\perp}$ and hence $\{x, y\} \subset \ell^{\perp}$. But since $\ell^{\perp}$ is a subspace, the span of x and y must then also be in $\ell^{\perp}$. Conversely, if $\ell \in \ell^{\perp}$, then $x \in \ell^{\perp} = x^{\perp} \cap y^{\perp}$, and hence $x \in y^{\perp}$. □

Proposition 6.8. *Let $\perp$ be a polarity of* $\mathrm{PG}(n, F)$. *Then the totally isotropic points and lines with respect to $\perp$ form a possibly degenerate polar space. The polar space is non-degenerate if and only if there is no isotropic point x such that $x^{\perp}$ containing all the isotropic points.*

Proof. Let x be a totally isotropic point. From the previous lemma we know that for any other isotropic point y, the line joining them in $\mathrm{PG}(n, F)$ is a line of this point-line geometry if and only if $x \in y^{\perp}$. Therefore, the points collinear with x are precisely the set of all isotropic points contained in $x^{\perp}$. In particular, the polar space is non-degenerate, that is, satisfies (PS2), if and only if there is no isotropic point x such that all isotropic points are contained in $x^{\perp}$.

Now let ℓ be any totally isotropic line that does not contain x. Note that every point on ℓ in $\mathrm{PG}(n, F)$ is an isotropic point, and hence a point of the point-line geometry. Since $x^{\perp}$ is a hyperplane, it either contains the line ℓ or intersects ℓ in exactly one point, thus proving (PS1).

□

When defining polar spaces using a polarity of $\mathrm{PG}(n, F)$ the following is a natural parameter associated with it.

Definition 6.9. The rank of the polar space formed by a polarity $\perp$ of $\mathrm{PG}(n, F)$ is the maximum vector space dimension of a totally isotropic subspace with respect to $\perp$.

We will see later that the dimension of any *maximal* totally isotropic subspace, is also equal to the rank of the polar space.

Lemma 6.10. *The rank of a polar space formed by a polarity of $\perp$ of* $\mathrm{PG}(n, F)$ *is at most* $(n + 1)/2$.

Proof. Let k be the vector space dimension of a totally isotropic space S. Then the vector space dimension of $S^{\perp}$ is $n + 1 - k$. Since $S \subseteq S^{\perp}$, we get $k \leq n + 1 - k$, that is, $k \leq (n + 1)/2$. □

Recall that the rank of the polar space associated to a polarity of $\mathrm{PG}(2, q)$ is 1, as it cannot have any totally isotropic lines. The examples of these polarities that we have seen are $(a, b, c) \mapsto \ell : ax + by + cz = 0$ which has $q + 1$ isotropic points, and $(a, b, c) \mapsto \ell : a^{\sqrt{q}}x + b^{\sqrt{q}}y + c^{\sqrt{q}}z = 0$ which has $q\sqrt{q} + 1$ isotropic points (when q is a square). We will see more examples of rank 1 polar spaces later. First we consider rank 2 polar spaces, and show that these give rise to generalized quadrangles.

Proposition 6.11. *Any non-degenerate rank* 2 *polar space over* $\mathrm{PG}(n, F)$ *is a generalized quadrangle.*

Proof. Let x be a totally isotropic point, and ℓ a totally isotropic line not containing x. Since $x^\perp$ is a hyperplane, it either meets ℓ in one point, and thus (GQ1) is satisfied for this pair, or it contains ℓ. We show that the latter cannot happen. Say $\ell \subseteq x^\perp$. Let y, z be two points on ℓ, and let $\pi = x \oplus y \oplus z$ be the plane spanned by x and ℓ. Then $\pi^\perp = x^\perp \cap y^\perp \cap z^\perp = x^\perp \cap (y^\perp \cap z^\perp) = x^\perp \cap \ell^\perp$. Since $\ell \in \ell^\perp$ and $x \in x^\perp$, we have $\pi \in x^\perp \cap \ell^\perp$, and hence $\pi \in \pi^\perp$. But this contradicts the fact that the rank of the polar space is 2.

Now we show that (GQ2) is satisfied. Let x be a totally isotropic point. Since the rank of the polar space is 2, there exists a totally isotropic line ℓ. Say $x \notin \ell$. Then by (GQ1) there exists a line m through x meeting ℓ. Say $x \in \ell$, then pick a point not collinear with x and then a line m through that point that meets ℓ. In both cases we have two lines meeting each other and x lies on one of these lines. Let y be the point of intersection, and z a point non-collinear with y. Let ℓ' be the line through z that meets the line m if $x \in \ell$, or the line ℓ if $x \in m$. By (GQ1) there exists a line through x meeting ℓ', giving us the second line through x and thus proving (GQ2). □

Many of the examples of generalized quadrangles that we saw were in fact rank 2 polar spaces coming from polarities of $\mathrm{PG}(n, \mathbb{F}_q)$. We will see more soon.

6.2 Collinearity Graph of a Polar Space

We know that the collinearity graph of a generalized quadrangle with an order (s, t) is always strongly regular. It turns out we can say the same for polar spaces that have an order.

Theorem 6.12. *If $\mathcal{S}$ be a polar space of order (s, t) on n points, $s, t \geq 1$, which is not a linear space, then the collinearity graph of $\mathcal{S}$ is an srg(n, k, λ, μ) with $k = s(t + 1)$, $\lambda = k - 1 - (n - k - 1)/s$ and $\mu = t + 1$.*

Proof. Since the order is (s, t), the graph is clearly $s(t + 1)$-regular. Moreover, if x and y are two non-collinear points, then every line through x contains a unique point collinear with y by (PS1). Therefore, $\mu = t + 1$.

Let x, y be two points on a line ℓ. All the $s - 1$ points in $\ell \setminus \{x, y\}$ are common neighbours of x and y. If z is any other common neighbour of x and y, then z must be collinear with all points of ℓ by (PS1). Therefore, z is also a common neighbour of every other pair of points on ℓ. This implies that there exists a constant $\lambda_\ell \geq s - 1$ such that any two points on ℓ have the same number of common neighbours.

Now fix a point p on ℓ and double count the pairs (q, r) where $q \in \ell \setminus \{p\}$ and r is a point collinear with q but non-collinear with p. For each of the s choices of q, there are $k - \lambda_\ell - 1$ choices of r. For each of the $n - k - 1$ choices of r, there is a unique choice of q by (PS1) and the fact that p, q are non-collinear. Therefore, $s(k - \lambda_\ell - 1) = n - k - 1$ which gives us $\lambda_\ell = k - 1 - (n - k - 1)/s$. In particular, λ_ℓ does not depend on the choice of ℓ, and the graph is strongly regular with parameters $(n, s(t + 1), s(t + 1) - 1 - (n - s(t + 1) - 1)/s, t + 1)$. □

6.3 Quadrics and Symmetric Bilinear Forms

We now look at some of the standard examples of polar spaces, coming from polarities of finite projective spaces.

Definition 6.13. Let F be a field and V a vector space over $F>$ A function $\beta : V \times V \to F$ is called a *bilinear form* if $\beta(\lambda u + \mu v, w) = \lambda\beta(u, w) + \mu\beta(v, w)$ and $\beta(u, \lambda v + \mu w)$ for all $\lambda, \mu \in F$ and $u, v, w \in V$. The form is called symmetric if $\beta(u, v) = \beta(v, u)$ for all $u, v \in V$. It is called *non-degenerate* if for every $u \in V \setminus \{0\}$, there exists a v such that $\beta(u, v) \neq 0$.

Lemma 6.14. *Let β be a non-degenerate symmetric bilinear form on F^{n+1}. Then $S \mapsto S^\perp = \{x \in \mathrm{PG}(n, F) : \beta(x, y) = 0\ \forall y \in S\}$, is a polarity of* $\mathrm{PG}(n, F)$.

Proof. Say $S \subseteq T$, then $T^\perp$ is clearly contained in $S^\perp$, since for every $x \in V$, $\beta(x, y) = 0$ for all $y \in T$ implies that $\beta(x, y) = 0$ for all $y \in S$. Therefore, we have an inclusion reversing map.

Let $e_1, \dots, e_{n+1}$ be a basis, and let $b_{ij} = \beta(e_i, e_j)$. Then by bi-linearity, we get that if $(u_1, \dots, u_{n+1})$ and $(v_1, \dots, v_{n+1})$ are the coordinates of u and v, then

$$\beta(u, v) = \sum u_i v_j b_{ij} = u^T B v,$$

where $B = [b_{ij}]_{(n+1)\times(n+1)}$ and we write u, v for the column vectors consisting of the coordinates. The non-degeneracy implies that the matrix is non-singular, since if there exists a non-zero v such that $Bv = 0$, then $\beta(u, v) = 0$ for all u. Now let S be a k-dimensional projective subspace, and let A be a $(k+1) \times (n+1)$ matrix whose rows form a basis of S. Then the coordinates of vectors in $S^\perp$ is equal to the set of solutions of $(AB)x = 0$. Since B is non-singular, the matrix AB is also a $(k+1) \times (n+1)$ matrix of rank $k+1$, and $S^\perp$ is its null space, which by the rank-nullity theorem has vector space dimension $n-k$, and hence projective dimension $n-1-k$. Therefore, $S \mapsto S^\perp$ is a map from k-dimensional subspaces to $n-1-k$ dimensional subspaces of $\mathrm{PG}(n, F)$.

Now $(S^\perp)^\perp = \{x \in \mathrm{PG}(n, F) : \beta(x, y) = 0 \text{ for all } y \in S^\perp\}\}$. Since $\beta(x, y) = \beta(y, x)$, we know from the definition of $S^\perp$ we know that $\beta(x, y) = 0$ for all $x \in S$, that is, $S \subseteq (S^\perp)^\perp$. But the dimension on the right hand side is $n - 1 - (n - 1 - \dim S) = \dim S$. This also shows that $\perp$ is a bijection, as it is its own inverse. □

Therefore, non-degenerate symmetric bilinear forms give rise to possibly degenerate polar spaces. In fact, degeneracy of the polar space can only occur in characteristic 2 fields.

Theorem 6.15. *Let F be a field of characteristic not equal to 2, and β a non-degenerate symmetric bilinear form on F^{n+1}. Then the polarity of* $\mathrm{PG}(n, F)$ *arising from β gives a non-degenerate polar space.*

Proof. We only need to check (PS2). Let x be an isotropic point and let $y \notin x^\perp$. Then the line ℓ joining them meets $x^\perp$ in a unique point, x. An arbitrary point of $\ell \setminus \{x\}$ can be written as $y + \lambda x$ for some $\lambda \in F$. We have $\beta(y + \lambda x, y + \lambda x) = \beta(y, y) + 2\lambda\beta(x, y)$. We know that $\beta(x, y) \neq 0$ since $y \notin x^\perp$. Pick $\lambda = -\frac{1}{2}\beta(y, y)/\beta(x, y)$, to get a point $z = x + \lambda y$ for which $\beta(z, z) = 0$, but $\beta(x, z) \neq 0$. □

For obtaining finite polar spaces, we let F to be the finite field $\mathbb{F}_q$, with q odd. The symmetric bilinear forms in this case are completely classified, and thus we know what the polar spaces look like. It's convenient to write this classification in terms of a related notion of quadratic forms, which also has the benefit that for q even quadratic forms also give us a non-degenerate polar spaces.

Definition 6.16. A quadratic form on a vector space V, over a field F, is a function $Q : V \to F$ such that

(i) $Q(\lambda v) = \lambda^2 Q(v)$ for all $v \in V$, and

(ii) $\beta(u, v) = Q(u + v) - Q(u) - Q(v)$ is a symmetric bilinear form.

We say that the bilinear form β above is obtained by *polarising* the quadratic form Q. The standard examples of quadratic forms, which are in fact the only ones after choosing a basis, come from evaluations of degree 2 homogeneous polynomials in $F[x_1, \dots, x_n]$.

Definition 6.17. Let Q be a quadratic form over a vector space V. A non-zero vector u is called singular if $Q(u) = 0$. A subspace S is called totally singular if $Q(u) = 0$ for all $u \in S$. Q is called non-degenerate (or non-singular) if there is no non-zero vector u for which both $Q(u) = 0$ and $\beta(u, v) = 0$ for all vectors v.

Note that the notion of non-degeneracy of a quadratic form is weaker than just saying the its corresponding bilinear form β is non-degenerate.

Lemma 6.18. *Let Q be a quadratic form and β its corresponding bilinear form. If a non-zero vector is singular with respect to Q, then it is isotropic with respect to β.*

Proof. Let u be a singular vector. Then $Q(u + u) = Q(u) + Q(u) + \beta(u, u)$. Since $Q(u + u) = 2^2 Q(u) = 0$, we have $\beta(u, u) = 0$. □

Proposition 6.19. *Let Q be a non-degenerate quadratic form over F^{n+1}. Then the totally singular points and lines in* $\mathrm{PG}(n, F)$ *form a non-degenerate polar space.*

Proof. Let x be a singular point. Let y be any other singular point. Since $Q(x) = Q(y) = 0$, we have $\beta(x, y) = Q(x + y)$. Therefore, if y is collinear with x, then the line joining them is totally singular, and in particular $Q(x + y) = 0$, that is, $\beta(x, y) = 0$. Conversely, if $\beta(x, y) = 0$, then for any λ, μ we have $Q(\lambda x + \mu y) = \lambda^2 Q(x) + \mu^2 Q(y) + \lambda\mu\beta(x, y) = 0$, and hence x and y are collinear. This shows that the collinearity is the same as it would be with respect to the symmetric bilinear form β. Moreover, for any singular point x, the space $x^\perp = \{y \in \mathrm{PG}(n, F) : \beta(x, y) = 0\}$ is a hyperplane, and not the whole space, by non-degeneracy of Q. Hence, (PS1) is satisfied.

For (PS2), let x be a point of Q and let $y \notin x^\perp$ be an arbitrary point outside the hyperplane $x^\perp$ in $\mathrm{PG}(n, q)$. Let $x + \lambda y$ be an arbitrary point of the line joining them (excluding x). We have $Q(x + \lambda y) = Q(x) + \lambda^2 Q(y) + \lambda\beta(x, y) = \lambda^2 Q(y) + \lambda\beta(x, y)$. If $Q(y) = 0$, then we have found singular point which is not collinear with x. If not, then let $\lambda = -\beta(x, y)/Q(y)$, which gives us a singular point not collinear with x. Therefore (PS2) is satisfied. □

The polar spaces associated to a quadratic form are known as *orthogonal polar spaces.* The set of singular points of a quadratic form, in the projective space, is known as a *quadric.* The rank of a polar space associated to a non-degenerate quadratic form Q is the maximum vector space dimension of a totally singular subspace.

Lemma 6.20. *Let F be a field of characteristic not equal to 2. For every symmetric bilinear form β on a vector space over F, the function $Q(u) = \beta(u,u)$ is a quadratic form. Moreover, a subspace S is totally singular with respect to Q if and only if it is totally isotropic with respect to β.*

Proof. We have $Q(\lambda u) = \beta(\lambda u, \lambda u) = \lambda^2 \beta(u,u) = \lambda^2 Q(u)$. We have $Q(u+v) - Q(u) - Q(v) = \beta(u+v, u+v) - \beta(u,u) - \beta(v,u) = 2\beta(u,v)$, which is a symmetric bilinear form since $2 \neq 0$. All we need to show is that every totally isotropic space is also totally singular. Let S be such that $\beta(x,y) = 0$ for all $x, y \in S$. Then $Q(x) = \beta(x,x) = 0$, for any $x \in S$. □

Therefore, as far as fields of characteristic not equal to 2 are concerned, symmetric bilinear forms are equivalent to quadratic forms, and we don't gain anything by studying either of the two. On the other hand, in characteristic 2 fields we gain something since the symmetric bilinear forms do not necessarily give a non-degenerate polar space, whereas quadratic forms always do.

Lemma 6.21. *Let Q be a quadratic form and ℓ a line in* $\mathrm{PG}(n,F)$*. Then either ℓ is totally singular, or it has at most 2 singular points on it.*

Proof. Say ℓ is not totally singular, and let x be a non-singular point on it. Let $y \in \ell \backslash \{x\}$, then an arbitrary point of $\ell \setminus \{x\}$ can be written as $y + \lambda x$, with $\lambda \in F$. This point is singular if $Q(y + \lambda x) = Q(y) + \lambda^2 Q(x) + \lambda \beta(y,x) = 0$, which is a non-trivial degree 2 equation in λ since $Q(x) \neq 0$ and hence has at most two solutions. □

The lines with exactly 0, 1 and 2 singular points on them are sometimes called *external*, *tangent* and *hyperbolic lines.*

Lemma 6.22. *Let f be an irreducible degree 2 homogeneous polynomial in $F[x,y]$. Then f has no zeros in* $\mathrm{PG}(1,F)$*.*

Proof. Let $f = ax^2 + bxy + cy^2$ and let (x_0, y_0) be a zero of f with $(x_0, y_0) \neq (0,0)$. WLOG say $x_0 \neq 0$.[2] The $ax_0^2 + bx_0y_0 + cy_0^2 = x_0^2(a + by_0/x_0 + cy_0^2/x_0^2) = 0$ and hence $a + by_0/x_0 + cy_0^2/x_0^2 = 0$. Then y_0/x_0 is a zero of the polynomial $g(z) = a + bz + cz^2$, and hence $g(z) = c(z - z_1)(z - z_2)$ for some $z_1, z_2 \in F$. Since $f(x,y) = x^2 g(y/x)$, This implies that $f(x,y) = cx^2(y/x - z_1)(y/x - z_2) = c(y - z_1 x)(y - z_2 x)$, contradicting that f is irreducible. □

The following classification of quadratic forms over $\mathbb{F}_q$ follows from a classical result of Witt, and we will not give a complete proof here.

Theorem 6.23. *Let Q be a non-degenerate quadratic form on an $n+1$ dimensional vector space over $\mathbb{F}_q$. Let r be the maximum vector space dimension of a totally singular subspace with respect to Q. Then there exists a basis in which Q has one of the following forms, with $n+1 = 2r$, $2r+1$ and $2r+2$, respectively.*

- $Q(x_1, \ldots, x_{n+1}) = x_1x_2 + x_3x_4 + \cdots + x_{2r-1}x_{2r}$.
- $Q(x_1, \ldots, x_{n+1}) = x_1x_2 + x_3x_4 + \cdots + x_{2r-1}x_{2r} + \xi x_{2r+1}^2$, *where $\xi = 1$ if q is even, and $\xi = 1$ or a chosen non-square if q is odd.*

[2] In fact, both x_0 and y_0 should be non-zero.

- *$Q(x_1, \dots, x_{n+1}) = x_1x_2 + x_3x_4 + \cdots + x_{2r-1}x_{2r} + x_{2r+1}^2 + \xi x_{2r+1}x_{2r+2} + \eta x_{2r+2}^2$, where $\eta = 1$ and $Tr(\xi^{-1}) = 1$ if q is even, and $\xi = 0$ and $-\eta$ is a chosen non-square if q is odd.*

Proof. (Sketch) A projective line spanned by two linearly independent vectors u, v is called a hyperbolic line if $Q(u) = Q(v) = 0$ and $\beta(u, v) \neq 0$. By Chevalley-Warning theorem (or in fact a direct argument), there exists a non-zero vector u such that $Q(u) = 0$, if $n \geq 2$. Pick such a vector u, and then pick a v such that u, v span a hyperbolic line W (this can always be done because of (PS2)). Write V as $W \oplus W^\perp$, where $\perp$ is defined using β. The restriction of Q on $W^\perp$ is still non-degenerate, and then we can use induction. In this basis consisting of hyperbolic lines the quadratic form looks as above, except until the last step where we have a subspace X of dimension at most 2 for which Q has no singular vectors. We then classify this case to get the above where the basis consists of r hyperbolic lines, and then a basis for X. Note that Lemma 6.22 gives the irreducible quadratic form on a 2-dimensional subspace, which can be taken to be the particular forms mentioned in the third case. It can also be shown that this r does not depend on which hyperbolic lines we use. □

The three different quadratic forms, and the polar spaces associated with them, are called *hyperbolic*, *parabolic* and *elliptic*, respectively.

As a convention, we also associated a parameter ϵ with these spaces, which is equal to -1, 0 and $+1$, respectively. The role of ϵ will be clear from the following proposition. Note that the dimension of the underlying projective space of a quadric of rank r and parameter ϵ is $2r + \epsilon$. These three quadrics, or more precisely the polar spaces associated with them, are also (somewhat confusingly) denoted as $Q^+(2r-1, q)$, $Q(2r, q)$ and $Q^-(2r+1, q)$. You can think of the symbol here as the amount that needs to be added to the dimension n so that we get twice the rank.

Thanks to this classification, we can compute various properties of quadrics/orthogonal polar spaces over $\mathbb{F}_q$ by looking at these special forms. In particular, we can count the total number of singular points.

Proposition 6.24. *Let Q be a quadric of rank r and parameter ϵ. Then the number of points in Q is*

$$(q^r - 1)(q^{r+\epsilon} + 1)/(q - 1).$$

Proof. We can write $Q(x_1, \dots, x_{n+1}) = x_1x_2 + x_3x_4 + \cdots + x_{2r-1}x_{2r} + \varphi$ where, $\varphi = 0$, ξx_{2r+1}^2 or an irreducible degree 2 homogeneous polynomial in x_{2r+1} and x_{2r+2}, depending on $\epsilon = -1, 0$ or 1, respectively, where $n + 1 = 2r + \epsilon$. Let $f(r)$ be the number of singular points in with respect to Q in $\mathrm{PG}(n, q)$. Each such point corresponds to $q-1$ solutions of the equation corresponding to Q in $\mathbb{F}_q^{n+1}$. Moreover, $(0, \dots, 0)$ is also a solution in $\mathbb{F}_q^{n+1}$. Partition the set of solutions in $\mathbb{F}_q^{n+1}$ by whether $Q'(x_3, \dots, x_{n+1}) = x_3x_4 + \cdots + \varphi$ is 0 or not. If it is 0 then we have $2q - 1$ choices for x_1 and x_2, whereas if it is non-zero then we have $q - 1$ choices. This gives us the following recurrence relation.

$$(q-1)f(r) + 1 = (2q-1)((q-1)f(r-1) + 1) + (q-1)(q^{2r-1+\epsilon} - (q-1)f(r-1) - 1).$$

The recurrence relation simplifies to

$$f(r) = qf(r-1) + q^{2r-1+\epsilon} + 1.$$

Now $f(0) = 0$, since a rank 0 polar space has no points in it, and we can easily verify that looking at these forms as well. This gives us the required value of $f(r)$. □

Corollary 6.25. *There exists an ovoid in* $\mathrm{PG}(3, q)$.

Proof. The rank 1 polar space of type +1 with equation $Q(x_1, x_2, x_3, x_4) = x_1x_2 + \varphi(x_3, x_4)$, where $\varphi \in \mathbb{F}_q[x_3, x_4]$ is an irreducible polynomial, has $\mathrm{PG}(3, q)$ as its underlying space and it has $q^2 + 1$ points on it. No three points here can be collinear by Lemma 6.21. □

To compute the number of lines in an orthogonal polar space, we compute the number of lines through each point, which gives us the order (q, t) of these spaces. The number of lines is then equal to the $N(t + 1)/(q + 1)$ where N is the number of points. We get t immediately from looking at *quotient polar spaces.*

Let U be a $(k + 1)$-dimensional subspace of an $(n + 1)$-dimensional vector space V. Then the quotient V/U, defined naturally as the vector space on the equivalence classes of the relation $x \sim y$ if $x - y \in U$, is an $n - k$ dimensional vector space. Therefore, the projective space we obtain from this quotient V/U is an $n - k - 1$ dimensional projective space. For example, in $\mathrm{PG}(n, F)$, the point-line geometry where points are all the lines passing through a fixed point x and the lines are all the planes passing through x is in itself a projective space isomorphism to $\mathrm{PG}(n-1, F)$. A similar notion of quotient spaces can be defined for polar spaces as follows.

Let Q be a quadratic form giving rise to an orthogonal polar space of rank r in $\mathrm{PG}(n, F)$. Let U be a totally singular subspace. Define a quadratic form Q_U on $U^\perp/U$ by $Q_U(x+U) = Q(x)$. Then this is well defined since for any $u \in U$, we have $Q_U(x+u+U) = Q(x+u) = Q(x) + Q(u) + \beta(x, u) = Q(x)$, as $Q(u) = \beta(x, u) = 0$. The *quotient polar space* at U is defined as the orthogonal polar space on $U^\perp/U$ with quadratic form Q_U.

Proposition 6.26. *Let $\mathcal{S}$ be an orthogonal polar space of rank r. For a point x of $\mathcal{S}$, the point-line geometry $\mathcal{S}/x$ where the points are all the totally singular lines through x and the lines are all the totally singular planes through x is a polar space of rank $r - 1$ of the same type as $\mathcal{S}$ with $x^\perp/x$ as its underlying projective space.*

Remark 6.27. The recurrence relation $f(r) = qf(r-1)+q^{2r-1+\epsilon}+1$ can now be interpreted as x being collinear to $qf(r - 1)$ points and non-collinear to $q^{2r-1+\epsilon}$ points.

Corollary 6.28. *The orthogonal polar space of rank r and parameter ϵ is of order*

$$\left(q, \frac{(q^{r-1} - 1)(q^{r+\epsilon-1} + 1)}{q - 1}\right).$$

For example, the rank 2 orthogonal polar spaces give the generalized quadrangles of order $(q, 1)$, (q, q) and (q, q^2) (embedded inside $\mathrm{PG}(3, q)$, $\mathrm{PG}(4, q)$ and $\mathrm{PG}(5, q)$, respectively).

Corollary 6.29. *The collinearity graph graph of an orthogonal polar space of rank r and type ϵ, is strongly regular with parameters*

$$n = \frac{q^{2r+\epsilon} - q^{r+\epsilon} + q^r - 1}{q - 1},$$

$$k = \frac{q^{2r+\epsilon-1} - q^{r+\epsilon} + q^r + q(q-2)}{q-1},$$

$$\lambda = \frac{q^{2r+\epsilon-2} - q^{r+\epsilon} + q^r + q(q-2)}{q-1},$$

$$\mu = \frac{q^{2r+\epsilon-2} - q^{r+\epsilon-1} + q^{r-1} + (q-2)}{q-1}.$$

One interesting thing to observe here is that because $\lambda - \mu = -q^{r+\epsilon-1} + q^{r-1} + q - 2$, the second largest eigenvalue in absolute value is $O(\sqrt{k})$ for $\epsilon = 0$ (and pretty close to $\sqrt{k}$ for the other two values of ϵ), which is in fact (asymptotically) the lowest possible second largest eigenvalue value any k-regular graph on $n > 2k$ vertices can have. Thus we get something known as optimally pseudorandom graphs from the collinearity graphs of parabolic polar spaces.

Another way of getting strongly regular graphs from quadrics is as follows.

Proposition 6.30. *Let Q be a quadric of rank r and type $\epsilon \in \{-1, +1\}$ in $\mathrm{PG}(2r + \epsilon, q)$. Then Q forms a two-intersection set.*

Proof. Let H be a hyperplane. Then either H is equal to $x^\perp$ for some $x \in Q$ or H intersects Q in a parabolic quadric of $\mathrm{PG}(2r - 1 + \epsilon, q)$. □

Another application of orthogonal polar spaces is in the construction of generalized hexagons. One of the constructions of the only known family of generalized hexagons of order (q, q), known as the *split Cayley hexagons*, takes the set of points as the points of $Q(6, q)$, and through each point x we pick a special set of $q+1$ totally singular lines through x forming a totally singular plane. The way these lines are chosen can be described geometrically, via something known as a triality, or algebraically using the standard equation of the quadric. Quadrics, and polar spaces in general are therefore fundamental objects in finite geometry. We now look at the classification of all polar spaces coming from polarities.

6.4 Sesquilinear Forms and Classical Polar Spaces

Definition 6.31. Let V be a vector space over F and $\sigma \in \mathrm{Aut}(F)$. A σ-sesquilinear form on V is a map $\beta : V \times V \to F$ such that

$$\beta(u_1 + u_2, v) = \beta(u_1, v) + \beta(u_2, v),$$

$$\beta(u, v_1 + v_2) = \beta(u, v_1) + \beta(u, v_2),$$

$$\beta(\lambda u, \mu v) = \lambda\sigma(\mu)\beta(u, v),$$

for all $u, u_1, u_2, v, v_1, v_2 \in V$ and all $\lambda, \mu \in F$. If σ is the identity then this form is called *bilinear*. The form β is called non-degenerate if $\beta(u, v) = 0$ for all v implies $u = 0$, or equivalently, $\beta(u, v) = 0$ for all u implies $v = 0$.

Proposition 6.32. *Every duality δ of $\mathrm{PG}(n, F)$ is induced by a σ-sesquilinear form β on F^{n+1}, with $S \mapsto S^\delta = \{x : \beta(x, y) = 0$ for all $y \in S\}$*

Definition 6.33. A sesquilinear form β is called reflexive if $\beta(u,v) = 0$ if and only if $\beta(v,u) = 0$ for all u, v.

Lemma 6.34. *A duality is a polarity if and only if the sesquilinear form corresponding to it is reflexive.*

Theorem 6.35 (Birkhoff-von Neumann)**.** *Let V be a vector space of dimension at least 3 and $\perp$ a polarity of $\mathrm{PG}(V)$ Then the reflexive σ-sesquilinear form corresponding to $\perp$ is one of the following:*

(a) Alternating: $\sigma = 1$ and $\beta(u,u) = 0$ for all $u \in V$.

(b) Symmetric: $\sigma = 1$ and $\beta(u,v) = \beta(v,u)$ for all $u, v \in V$.

(c) Hermitian: $\sigma^2 = 1$, $\sigma \neq 1$ and $\beta(u,v) = \sigma(\beta(v,u))$ for all $u, v \in V$.

The polar spaces corresponding to quadratic forms or σ-sesquilinear forms are known as classical polar spaces. In general F need not be a field, and a division ring suffices, but for our purposes of finite polar spaces we look at $\mathbb{F}_q$. In any finite abstract polar space one can define a notion of “rank” similar to the one that we have in these classical polar spaces, and a deep result of Tits shows that every finite abstract polar space of rank ≥ 3 must be one of the classical polar spaces. Therefore, the only polar spaces that are not completely understood are the rank 2 ones, a.k.a., generalized quadrangles.

6.5 Exercises

1. An ovoid $\mathcal{O}$ in $\mathrm{PG}(n,q)$ is a set of points with the following two properties: (i) no three distinct points of $\mathcal{O}$ are collinear, and (ii) for all points $x \in \mathcal{O}$ there exists a hyperplane H_x such that the set of lines through x that are tangent to $\mathcal{O}$ is equal to the set of lines through x in H_x.

   (a) Prove that an ovoid in $\mathrm{PG}(n,q)$ has $q^{n-1} + 1$ points.

   (b) Prove that for every $n \geq 4$, $\mathrm{PG}(n,q)$ does not contain any ovoids.

2. A bilinear form $\beta : V \times V \to F$, for some vector space V over F, is called *alternating* if $\beta(u,u) = 0$ for all $u \in V$. It is called non-degenerate if $\beta(u,v) = 0$ for all $v \in V$ implies that $u = 0$.

   (a) Show that if β is an alternating bilinear form then $\beta(u,v) = -\beta(v,u)$ for all u, v.

   (b) Prove that the maximum vector space dimension of a totally isotropic subspace with respect to a non-degenerate alternating bilinear form over a vector space of dimension $n+1$ is equal to $(n+1)/2$, and hence deduce that n must be odd.

   (c) Prove that the totally isotropic points and lines of $\mathrm{PG}(n,F)$ with respect to a non-degenerate alternating bilinear form over the underlying vector space F^{n+1}, form a non-degenerate polar space. Also determine the order (s,t) of this polar space if $F = \mathbb{F}_q$.

3. Let $\mathcal{S}$ be an orthogonal polar space of rank r and type ϵ, over $\mathbb{F}_q$. Prove that the number of totally singular subspaces of vector space dimension r is equal to

$$\prod_{i=1}^{r}(q^{i+\epsilon}+1).$$

4. (a) Give an example of a non-degenerate symmetric bilinear form that contains totally isotropic points and lines, such that these point and lines give rise to a degenerate polar space.

   (b) Prove that a quadratic form over a field of characteristic not equal to 2 is non-degenerate if and only if the bilinear form associated with it is non-degenerate.

   (c) Give an example of a non-degenerate quadratic form Q such that the bilinear form β associated with it is degenerate.

5. Let β be a non-degenerate symmetric bilinear form over $\mathbb{F}_q^5$, that gives a polarity $\perp$ of $\mathrm{PG}(4,q)$ by mapping a point x to the hyperplane $x^\perp = \{y \in \mathrm{PG}(4,q) : \beta(x,y)=0\}$. Let z be a point of $\mathrm{PG}(4,q)$ such that $z \notin z^\perp$, and let $\mathcal{O}$ be an ovoid in $z^\perp \cong \mathrm{PG}(3,q)$.

   Define a graph G with vertex set equal to the set of points x in $\mathrm{PG}(4,q)\setminus(z^\perp \cup \{z\})$ such that x lies on a line joining z and a point of $\mathcal{O}$, and making two vertices x, y adjacent if $x \in y^\perp$.

   (a) Determine the number of vertices n in G and show that the number of edges is at least $Cn^{5/3}$, for some constant C and large enough n.

   (b) Prove that G does not contain any copies of the graph $K_{3,3}$.[3]

[3] In fact, the graph G is asymptotically the densest possible graph on n vertices that does not contain a $K_{3,3}$.