2022

Counting The Moduli Space Of Pentagons On Finite Projective Planes

Maxwell Hosler
The College of Wooster, mhosler22@wooster.edu

Counting the Moduli Space of Pentagons on Finite Projective Planes

Independent Study Thesis

Presented in Partial Fulfillment of the Requirements for
the Degree Degree Bachelor of Arts in the
Department of Mathematics at The College of Wooster

by
Maxwell Hosler
The College of Wooster
2022

Advised by:

Dr. Robert Kelvey (Mathematics)

THE COLLEGE OF

WOOSTER

Abstract

Finite projective planes are finite incidence structures which generalize the concept of the real projective plane. In this paper, we consider structures of points embedded in these planes. In particular, we investigate pentagons in general position, meaning no three vertices are colinear. We are interested in properties of these pentagons that are preserved by collineation of the plane, and so can be conceived as properties of the equivalence class of polygons up to collineation as a whole. Amongst these are the symmetries of a pentagon and the periodicity of the pentagon under the pentagram map, and a generalization of the concepts of rotational and reflective symmetry. We are also interested in counting exactly how many such equivalence classes of pentagons exist on a given projective plane.

This work is dedicated to my parents, Lisa and Jay Hosler.

Acknowledgments

Thank you to Dr. Rob Kelvey for his guidance. He helped me to develop a handful of ideas that struck me when trawling Wikipedia into a complete project.

Thanks and love to my parents and brother, Lisa, Jay and Jack Hosler, for supporting me in this and other pursuits. They have both inspired me and pushed me to explore what I am passionate about. They have patiently listened to me as I try and fail to explain whatever I'm interested in in that moment.

I extend a similar gratitude to my friends, both at The College of Wooster and at home, who have listened to me try my best to explain what I'm exploring even when its far beyond their personal areas of expertise.

Finally, I would like to thank the developers and maintainers of Python, Geogebra, and Desmos for providing the free tools which I have and will continue to use to explore mathematics.

Vita

Fields of Study Major field: Mathematics

Minor field: Computer Science

Specialization: Finite geometry, abstract algebra

Contents

List of Figures

List of Listings

CHAPTER 1

Introduction

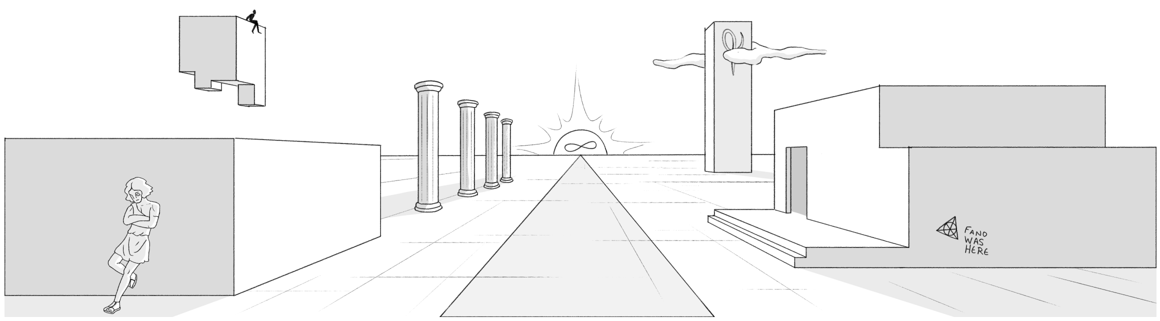

Figure 1.1: Parallel lines seem to meet at the horizon in linear perspective art.

In art that uses the technique of one-point perspective, parallel lines seem to meet at infinity (see fig. 1.1). In the realm of the real projective plane, $\mathbb{RP}^2$, this appearance is reality. Parallel lines are defined to meet at their own *point at infinity*, and these *points at infinity* lie on a *line at infinity*. Beyond being a realization of a human perceptual intuition, this change corrects a fundamental asymmetry of the regular, affine plane $\mathbb{R}^2$. Just as every pair of points has a line between them, now every pair of lines has a point between them.

Of course, we must be more rigorous in our definition of $\mathbb{RP}^2$ in order for it to be a well-defined mathematical object. For every pair of lines ℓ_1, ℓ_2, we define the point they share to be their intersection if they aren't parallel, and a new point m if they are parallel, where m is their shared slope. Note that this means that every line in each equivalence class under the relation of being parallel meets at a single point. These 'slope points' are our points at infinity, and we add a single line to our plane which contains exactly these points, which is our line at infinity.

An important difference from our perspective intuition is that parallel lines only meet at a single point. A person standing on a pair of train tracks might think that they meet at the horizon in front of them, and then turn around to see the tracks meeting at the horizon again behind them. While this hypothetical person might think of these two meetings as two points, they are, in $\mathbb{RP}^2$ at least, the same point. In a very loose sense, the lines 'wrap around,' meeting at the same point in either direction.

The real projective plane is a very important object in its own right, but it opens up an idea that I find much more interesting. The fundamental structure of $\mathbb{RP}^2$ defines a plane, not in the language of vector spaces, topologies, or curvature, but as an interrelated structure of points and lines called an *incidence structure*. Lines are returned to their pre-Cartesian status as prime objects in of themselves, instead of just a set of points. What this allows us to do is explore a new form of geometry, abstracted from the concept of spacial position, or even any sense of continuous space at all, built entirely around the concept of incidence. With this we can explore the world of finite projective planes.

Within these new spaces, we will examine the geometry of polygons. More specifically, we will look at pentagons, the smallest polygons that are 'nontrivial' with respect to the concepts we'll be applying. Some of it will take familiar language and form; for example, much of our time will be spent focused on the symmetries of these pentagons. In this realm, much of what we deomonstrate reflects our intuitions about the subject, or at the very least has a direct analogue on the real plane. We will even encounter many planes' own version of the golden ratio, that which famously appears in the study of pentagons on the real plane. However, the nature of finite planes also gives rise to very alien realities. Much of this paper will be focused on counting the similarity classes of polygons on finite planes, a notion that makes no sense in the continuous, infinite world of the real plane. We'll also study an operation on pentagons whose iteration on real pentagons causes them to shrink to a point, but is periodic on pentagons of the finite plane. The goal of this paper is to bring these elements of familiarity and unfamiliarity together while building up our knowledge of the structure of these spaces.

CHAPTER 2

Finite Projective Planes

2.1 Definition

Projective planes are defined using axioms derived from $\mathbb{RP}^2$. Specifically, a projective plane π is a tuple $(\mathcal{P}, \mathcal{L}, \mathcal{I})$, where $\mathcal{P}$ is defined as a set of points, $\mathcal{L}$ is a set of lines, and $\mathcal{I} \subset (\mathcal{P} \times \mathcal{L}) \cup (\mathcal{L} \times \mathcal{P})$ is a relation between points and lines, which satisfy the following axioms [1].

1. For all distinct $p_1, p_2 \in \mathcal{P}$, there is a unique $\ell \in \mathcal{L}$ such that $p_1 \mathcal{I} \ell$ and $p_2 \mathcal{I} \ell$.

2. For all distinct $\ell_1, \ell_2 \in \mathcal{L}$, there is a unique $p \in \mathcal{P}$ such that $\ell_1 \mathcal{I} p$ and $\ell_2 \mathcal{I} p$.

3. There exists p_1, p_2, p_3, p_4 such that, for any three of them p_a, p_b, p_c, there is no $\ell \in \mathcal{L}$ such that $p_a \mathcal{I} \ell$, $p_b \mathcal{I} \ell$, and $p_c \mathcal{I} \ell$.

Note that $\mathcal{P}$ and $\mathcal{L}$ need not correspond to any 'real' points or lines in any plane. They are defined to be these things, and even constructing such an equivalence is impossible in general.

Of course, these axioms are quite opaque, but they can be clarified by borrowing more terminology from regular geometry. If $p \mathcal{I} \ell$, we say that p lies on ℓ and ℓ contains p. Naturally, if a collection of points lies on a line, they are colinear. This gives us a much more readable set of axioms:

1. Any two points in π both lie on precisely one shared line in π.
2. Any two lines in π have precisely one shared point that lies on both of them in π.
3. There exists four points in π such that no three are colinear.

The third axiom is also sometimes written as 'There exists a quadrangle' to clarify its intuition further. While this is clearly true for the real projective plane, it may seem out of place. It serves to prevent the formation of *degenerate planes*, guaranteeing our structure has certain symmetries. As we will see later, it also allows us to introduce a coordinate system to arbitary projective planes.

A *finite projective plane* is exactly what its name implies. It is a projective plane wherein its point and line sets have a finite number of elements. In fact, as we will see, they must have the same number of elements, in addition to other numerical symmetries.

For this paper, we will be using the following notation to discuss unique incidence. For any two distinct points, A and B, we will call the unique line they both lie on $[A \cdot B]$, and for any two distinct lines M abd N we will call the unique point that lies on both of them $M \cap N$.

2.2 A Basic Construction

To understand how we construct a finite projective plane, let's start with a 3×3 grid of points [10]. The size of our starting grid will be what we call the *order* of our plane. We need to construct a set of lines on this grid, with some being parallel. We can construct some as you would expect, horizontally and vertically.

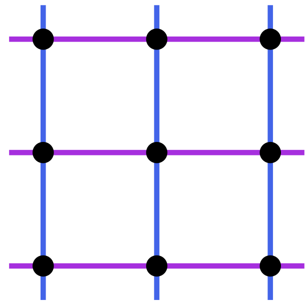

Figure 2.1: Horizontal and vertical lines in our construction.

However, simply choosing all the diagonal lines in this diagram would give us lines that do not satisfy our axioms. In particular, some would not intersect while also not being parallel, and thus meeting at infinity. Instead, we must define lines in terms of a starting point, and a 'direction' of the form $(1, k)$, with k being greater than zero and less than the order. Start at a point and add the direction to its position, wrapping around when necessary, until you return to the original point. This will give you a line, and two lines with the same direction but different starting points will either be the same line or parallel.

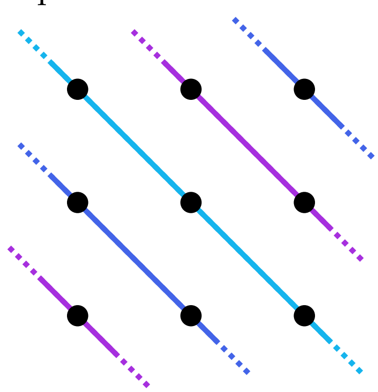

Figure 2.2: A class of three parallel lines. Sections of the same color are part of the same line.

Finally, just as we did with the real projective plane, we add a new point connecting all classes of parallel lines, and a new line through all of those. You will note that this procedure works whenever the order we choose is a prime number.

Otherwise, if it is not, then our manner of generating 'parallel' lines will sometimes create parallel lines that intersect, when the direction is not coprime to the order. The end result of this process is figure 2.3.

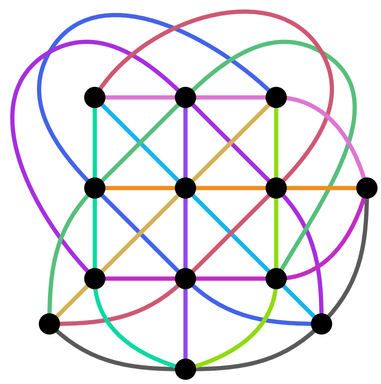

Figure 2.3: The finite projective plane of order 3.

The smallest of these planes is the Fano plane, having order 2 [1]:

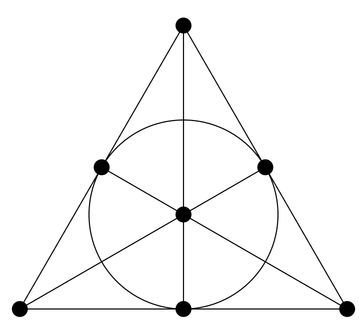

Figure 2.4: The classical way of representing the Fano plane.

In this image, each of the diagonals, sides, and the circle all represent lines. The Fano plane serves as a useful model object for demonstrating some, although not all, properties of projective planes.

2.3 Properties

Fundamental to the structure of finite projective planes are the concepts of order and duality. Order we have already encountered, but to develop a better understanding of it we must first consider duality.

Duality for finite projective planes is much like duality for regular polyhedra; just as faces can be changed to vertices and vice versa in regular polyhedra to form other regular polyhedra, lines and points in projective planes can be inverted. In fact, this duality is even stronger here, as any theorem about points and lines holds true if you swap the words 'lines' and 'points' [1]. This is fundamentally because the projective plane axioms treat points and lines symmetrically, such that the categories of 'lines' and 'points' are arbitrary namings. This is clear for axioms 1 and 2, which are transparently duals of each other. However, axiom 3 is also dualistic, as the existence of a quadrangle also implies the existence of four lines, no three of which go through a single point, those lines being the sides of the quadrangle.

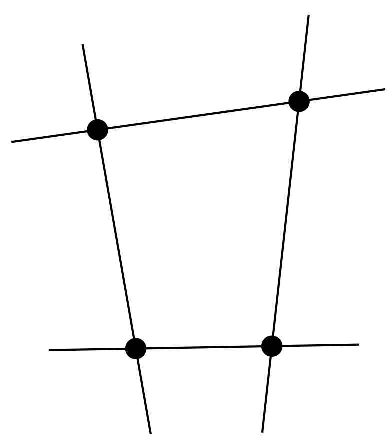

Figure 2.5: A quadrangle with its 4 distinct sides.

Duality also appears in the quantitative structure of finite projective planes, and this is where the concept relates to order. In general, we denote finite planes of order p with π_p, although as we will see this notation is ambiguous for some values of p, without a few qualifications. Nonetheless, we can say with certainty that π_p

has precisely $p^2 + p + 1$ points, and dually, $p^2 + p + 1$ lines. Furthermore, every line passes through $p + 1$ points and every point lies on $p + 1$ lines [1].

While this won't demonstrate this principle for all planes, let us consider why this is true for our construction. Consider that there are p^2 points in our grid. Beyond that, there p directions of the form $(1, n)$, and 1 additional direction, that being $(0, 1)$. As such, there are $p + 1$ points at infinity. So, there are $p^2 + p + 1$ points.

2.4 Representations of Finite Projective Planes

2.4.1 Over Finite Fields

The construction given in the previous section has a very direct connection to how we constructed $\mathbb{RP}^2$ intuitively. However, because of this, it inherits one of the misleading aspects of that construction. It seems to differentiate between 'normal' points and 'points at infinity.' Mathematically, there is no such distinction. There exist incidence-preserving automorphisms of any plane which take any line to any other line. In other words, we can rearrange the plane so that anything is the 'line at infinity,' meaning that this concept is a matter of perspective. However, now that we have the motivation for why we call these incidence structures 'planes,' we can move to a new way of constructing our incidence structures with both a strong visual intuition and a connection to abstract algebra.

Once again, we will consider the case of $\mathbb{RP}^2$, however, in this case we will be building it starting with $\mathbb{R}^3$, considered as a vector space over the field of reals. First, we remove the zero vector from consideration. Then, for any $\vec{v}, \vec{u} \in \mathbb{R}^3$, we identify these vectors with each other if there is some $c \in \mathbb{R} \setminus \{0\}$ such that $\vec{v} = c\vec{u}$. This system is a called homogenous coordinates, and it is common in both projective geometry and computer graphics. By adding an additional Euclidean dimension to $\mathbb{R}^2$, then removing a dimension by reducing the space modulo scaling, we get a new two dimensional space with a different structure. To motivate this, imagine a plane in $\mathbb{R}^3$ at $z = 1$. If the z coordinate of a point (a, b, c) in our space is nonzero, we can scale it onto that plane by multiplying by the inverse of c, to get $(\frac{a}{c}, \frac{b}{c}, 1)$. So, the points $(x, y, 1)$ act like 'normal' points of the plane.

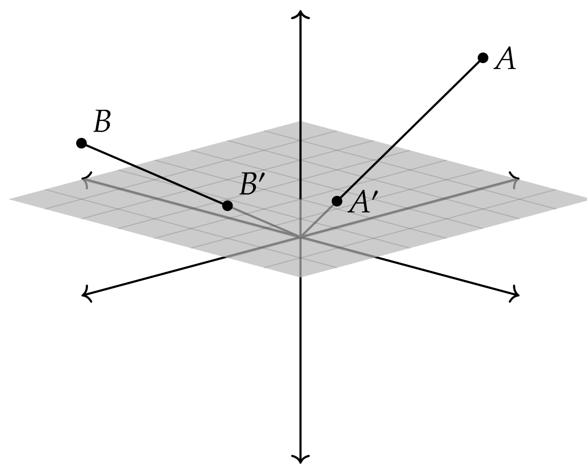

Figure 2.6: Projection onto $z = 1$

Now, consider (a, b, c) as c goes to 0. We can see that $(\frac{a}{c}, \frac{b}{c}, 1)$ has its first two coordinates rush off to infinity. Indeed, for this reason, $(x, y, 0)$ act as our points at infinity.

In addition, if we consider only the unit vectors as representatives, we get the unit sphere with antipodal points identified. This is a standard representation of $\mathbb{RP}^2$ in topology. However, for our purposes, it will be best to think of these classes of points as the one dimensional subspaces of $\mathbb{R}^3$. Ironically, these subspaces, which look like lines, will be the points of $\mathbb{RP}^2$.

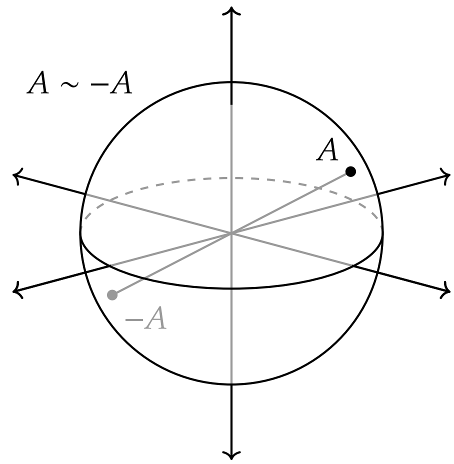

Figure 2.7: The real projective plane as a manifold $S^2/(x \sim -x)$. Like the Klein bottle, it can't be represented in 3-dimensional space without holes, self-intersection, or redundancy. We see the latter approach here.

Our lines, then, will be planes. In specific, for any pair of one dimensional subspaces $A, B \subset \mathbb{R}^3$ with $A \neq B$, the 'line' connecting them will be their shared plane[1]. It is intuitively clear this plane will always exist, as any two distinct lines that share a point (in this case the origin) will define a plane, and that that plane is unique. Thus, the first axiom, that any two points sit on one shared line, is satified. It is also intuitively clear that any two planes that pass through the origin will share a line, meaning that the second axiom, that every pair of lines intersects at precisely one point, is satisfied.

[1]Interestingly, projecting these planes onto our unit sphere representation will give you a great circle of that sphere connecting our two points. This is one way of thinking about how that representation works.

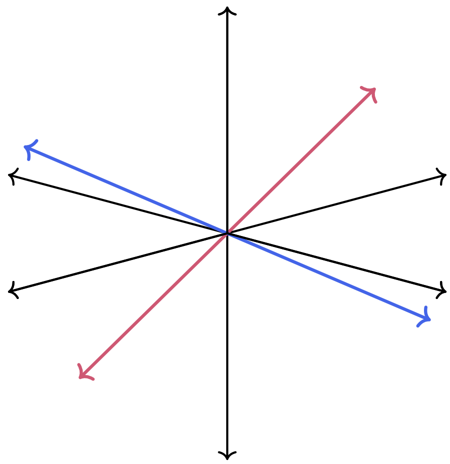

Figure 2.8: 1-dimensional subspaces of $\mathbb{R}^3$.

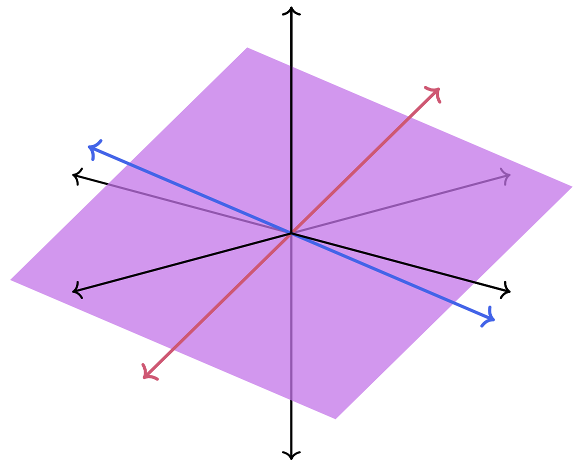

Figure 2.9: Two distinct 1-dimensional subspaces define a plane.

Intuition aside, it will be useful for us to justify this with linear algebra. Note that, by definition, A and B are one-dimensional subspaces of $\mathbb{R}^3$. This means that each is generated by some single vector. Let those vectors be $a, b \in \mathbb{R}$, respectively. We will denote this $A = \langle \vec{a} \rangle$ and $B = \langle \vec{b} \rangle$. Furthermore, since A and B are distinct, $\vec{a}$ and $\vec{b}$ must be linearly independent. So, $L = \langle \vec{a}, \vec{b} \rangle$ must be a two dimensional linear subspace of $\mathbb{R}^3$, and because of the properties of bases, it must be unique. Thus, axiom 1 is satisfied.

Now, consider two distinct two-dimensional subspaces $\langle \vec{a}, \vec{b} \rangle, \langle \vec{c}, \vec{d} \rangle$. Clearly, their intersection cannot be two-dimensional, because they are distinct spaces. However, if it were zero dimensional, the intersection would be the set containing only the zero vector, which would imply there was only one solution to $\alpha\vec{a} + \beta\vec{b} = \gamma\vec{c} + \delta\vec{d}$, that being $\alpha = \beta = \gamma = \delta = 0$. In other words, there would only be the trivial solution to $\alpha\vec{a} + \beta\vec{b} - \gamma\vec{c} - \delta\vec{d} = \vec{0}$, meaning $\vec{a}, \vec{b}, \vec{c}, \vec{d}$ are linearly independent. However, this cannot be true because $\mathbb{R}^3$ is three dimensional, so at most three vectors can be linearly independent in it. So, the intersection must be one-dimensional. So, axiom 2 is satisfied.

Finally, consider any three basis vectors $\vec{a}, \vec{b}, \vec{c}$ of $\mathbb{R}^3$ and also the vector $\vec{a} + \vec{b} + \vec{c}$.

By definition, $\vec{a}, \vec{b}, \vec{c}$ are linearly independent, and therefore not coplanar. Without loss of generality, we can consider $\vec{a}, \vec{b}, \vec{a}+\vec{b}+\vec{c}$ as a representative of any other triple. Consider the equation $\vec{0} = \alpha\vec{a}+\beta\vec{b}+\gamma(\vec{a}+\vec{b}+\vec{c}) = (\alpha+\gamma)\vec{a}+(\beta+\gamma)\vec{b}+\gamma\vec{c}$. If $\vec{a}, \vec{b}, \vec{a}+\vec{b}+\vec{c}$ weren't linearly independent, this would have a nonzero solution, but if that were true, then that would imply $(\alpha+\gamma)\vec{a}+(\beta+\gamma)\vec{b}+\gamma\vec{c} = \vec{0}$ had a nonzero solution, which would mean $\vec{a}, \vec{b}, \vec{c}$ were linearly dependent. This is false by definition, so any triple in our set is linearly independent. Thus, axiom 3 is satisfied.

Note that the process we used to generate $\mathbb{RP}^2$ from $\mathbb{R}^3$ can be entirely justified using its properties as a vector space and the real numbers' status as a field. So, even though we have a visual intuition for why this process works, we can safely extend it to more abstract vector spaces and fields. It comes naturally, then, that one can build finite projective planes from the *finite fields*, $\mathbb{F}_p$. Simply repeat the same process on the vector space $\mathbb{F}_p^3$, with the $\mathbb{F}_p$ being the scalar field. It is known that $\mathbb{F}_p$ exists precisely when p is a prime or a power of a prime [9].

In fact, this method of construction seems to be deeply connected to the nature of finite projective planes. Not only can any finite field be used to construct a plane, but every known finite projective plane can be constructed from some finite field. It has not been proven that this holds in general, however, it is highly suspected [4].

Now, let's consider how many points are in the plane constructed from $\mathbb{F}_p^3$. There are p^3 objects in $\mathbb{F}_p^3$. Since we removed $(0,0,0)$, we have p^3-1. Now, consider some $\vec{v} \in \mathbb{F}_p^3 \setminus \{\vec{0}\}$. Clearly, it has some nonzero coordinate. Without loss of generality, assume its in the first position, v_0. For any $c \in \mathbb{F}_p \setminus \{0\}$, consider $c\vec{v}$. Note that, since $\mathbb{F}_p \setminus \{0\}$ is a group under multiplication, cv_0 must take on a different value for every value of c, of which there are $p-1$ possibilities. Since all these multiples are equivalent in our system, this means that every unique point is represented $p-1$ times in our collection of p^3-1 objects. So, the number of points is $\frac{p^3-1}{p-1} = \frac{(p-1)(p^2+p+1)}{p-1} = p^2+p+1$. Thus, there are p^2+p+1 points. But note what this means! It shows that the order of our finite field, p, is also *precisely* the order of our projective plane! Thus, the bond between finite fields and finite projective planes is strengthened even further. What's more, this demonstrates that, for any p which is a prime or a power of a prime, some finite plane π_p exists. If we add the additional condition that a theorem called *Desargues's theorem* holds, π_p refers unambiguously to a single plane. There are sometimes other *non-Desarguesian* planes of order p, however, we will not be dealing with those. This is in part because of their additional complexity, and in part because a plane is a field plane if and only if it is Desarguesian [14]. Since we will leveraging the algebraic structure of field planes to investigate them, this is a necessity.

However, we can push the usefulness and elegance of this representation even further. Consider that, in three dimensions, every 2-dimensional subspace has a 1-dimensional subspace of orthogonal vectors. In other words, every plane through the origin has a line through the origin orthogonal to it. This means we can uniquely represent the lines of our projective planes as vectors instead of planes. Now, consider finding the line ℓ between two points, A and B. We need to find some vector which is orthogonal to the plane spanned by $\vec{a}$ and $\vec{b}$, in other words, a vector orthogonal to both of these vectors. However, we know such an operation exists in the form of the cross product! So, $[A \cdot B] = \ell = A \times B$. Furthermore, consider the 1-dimensional subspace shared by two 2-dimensional subspaces must be orthogonal to both the 2-dimensional subspaces' orthogonal space. So, given two lines ℓ, m, the cross product once again defines the incidence relation, with $\ell \cap m = A = \ell \times m$. With this, we not only get to represent incidence with a very familiar operation, but we also get a perfect representation of the complete duality of finite projective planes. Our set of lines and our set of points are precisely the same set of objects under different names, with two vectors being incident if they are orthogonal, meaning our incidence relation is simply $\mathcal{I} = \{(x, y) : x \cdot y = 0\}$, with the shared incidence of two points or lines being the same single operation, the cross product.

This construction of finite projective planes as the linear subspaces of $\mathbb{F}_p^3$ will be the main way we approach them. In particular, it gives us very powerful tools to study the symmetries of these spaces using matrices, and from there derive combinatorial results. We will also find its characterization of non-colinearity in terms of linear independence very useful in finding out what kinds of configurations are 'valid' in some sense.

2.4.2 Levi Graphs

For planes of orders greater than 2, it quickly becomes hard to visualize the incidence structure. For planes larger than the Fano plane, direct depictions become rare as the resulting image is an incomprehensible net of lines and points. However, while it will always be hard to visualize projective planes of high enough order, we can at least get depictions of lower order graphs where some of the structure is at least vaguely apparent. We achieve this using Levi graphs.

A Levi graph of any incidence structure is a graph where every point and line is assigned a node, and there is an edge between two nodes if they respectively represent a point and a line which are incident [11].

Using only this definition and what we know about projective planes, we can start to make claims about the structure of Levi(π_p). For example, it is clear that it must be $(p + 1)$-regular, since every point is incident with $p + 1$ lines and every line is incident with $p + 1$ points. It also must be bipartite, since there can't be an edge between a line and a line or a point and a point.

The cycles of Levi(π_p) are of particular interest. First, there can be no odd cycles, since Levi(π_p) is bipartite. There also cannot be a 4-cycle, as the two lines represented in the cycle would both be incident with the two points on the cycle. This would

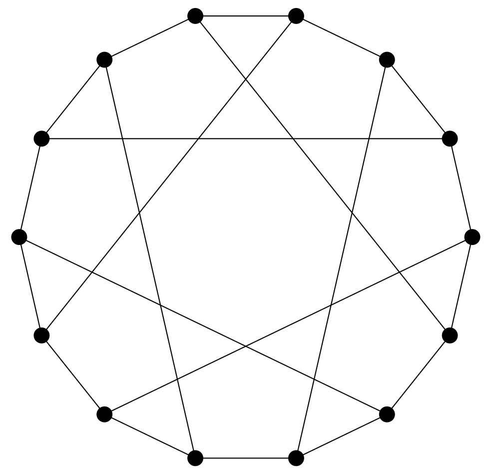

Figure 2.10: The Levi graph of the Fano plane, Levi(π_2), also called the Heawood graph. This graph representation is an excellent tool for studying projective planes.

mean two lines crossed each other at two points, violating both axiom 1 and axiom 2. However, since axiom 3 guarantees the existence of three non-colinear points, there must be a 6-cycle in Levi(π_p). Thus, the smallest cycle in Levi(π_p) will always have length 6. This is called the *girth* of the graph.

Finally, consider any two vertices a, c of Levi(π_p). If they are of the same type (as in point or line), there must be some point or line they share in common, c. So, there is a path of length 2 between them, $a \to b \to c$. If they are of different type, let a be the line. Choose any point on the line, b_1, and let the line shared by b_1 and c be called b_2. So, there is a path of length 3 between a and c, $a \to b_1 \to b_2 \to c$. Thus, the furthest distance between any two vertices of Levi(π_p) is 3, meaning the diameter of Levi(π_p) is 3 [11].

CHAPTER 3

Polygons on Finite Projective Planes

3.1 Automorphisms

Like all mathematical structures, projective planes have a naturally defined set of automorphisms which preserve their inherent structure. For a plane $\pi = (\mathcal{P}, \mathcal{L}, \mathcal{I})$, an automorphism is a bijective function $\phi : \mathcal{P} \to \mathcal{P}$ with an induced function $\overline{\phi} : \mathcal{L} \to \mathcal{L}$ such that, for $p \in \mathcal{P}$ and $\ell \in \mathcal{L}$, it is true that $p\mathcal{I}\ell \iff \phi(p)\mathcal{I}\overline{\phi}(\ell)$. In practice, it is usually clear enough to let a single symbol ϕ to refer to both ϕ and $\overline{\phi}$. Since it maps lines to lines, ϕ is also sometimes called a *collineation* [1].

Thankfully, our representation of π_p using $\mathbb{F}_p^3$ gives us a natural way to think about the automorphisms of π_p. Any invertible linear transformation of the vector space $\mathbb{F}_p^3$ defines an automorphism, the vectors representing points being taken to new points, and remaining on the planes that represent lines. However, $\mathrm{GL}(3, p)$, the group of invertible 3×3 matrices on $\mathbb{F}_p$, isn't isomorphic to the collineation group of π_p. Notice that any pure scaling operation, represented by some scalar multiple of the identity, is always equivalent to the same automorphism, the identity automorphism. This is because we consider vectors the same up to scaling, so any operation that doesn't knock at least one vector off its own span will be the identity automorphism. In other words, our group is 'too large' insofar as every collineation has multiple representations in $\mathrm{GL}(3,p)$. Thankfully, when p is prime, this is the only hangup. So, the automorphism group of π_p is $\mathrm{GL}(3,p)/Z(\mathrm{GL}(3,p))$, with $Z(G)$ being the center of a group G, the set of all elements that commute with every other element. In this case, $Z(\mathrm{GL}(3,p))$ is in this case the group of all scalar transformations. This group is called the projective linear group, $\mathrm{PGL}(3,p)$. It represents all invertible linear transformations, with transformations that are scalar multiples of each other being considered the same [5].

For non-prime orders, there are additional complications which make the automorphism group $\mathrm{P\Gamma L}(3,p)$, the projective semilinear group on $\mathbb{F}_p^3$. Without going into too much detail, this is because fields of non-prime order have nontrivial automorphisms. These automorphisms induce a type of collineation called an *automorphic collineation* which is not representable by matrices. This is as opposed to the collineations we will use, which can be, and are called *homographies*. Thankfully, due to the Fundemental Theorem of Projective Geometry, we know that every collineation is a product of a homography and an automorphic collineation, called a *semilinear map* [2]. This means that, in the case of fields of prime order, all collineations are just homographies. Due to its additional complexity, we will be considering only planes of prime order from now on.

Unlike transformations in the general linear group, transformations in the projective linear group are not uniquely defined by where three linearly independent vectors are taken. Consider three such vectors $\vec{a}_1, \vec{a}_2, \vec{a}_3$ which are taken to vectors $\vec{b}_1, \vec{b}_2, \vec{b}_3$ by transformation T_1. Now, consider a second transformation T_2 which is different in that it takes $\vec{a}_1$ to $2\vec{b}_1$. It is clear that T_1 and T_2 aren't scalar multiples of each other, and thus aren't the same transformation, but they are nonetheless the same with respect to how it moves $\vec{a}_1, \vec{a}_2, \vec{a}_3$, since we consider $\vec{b}_1$ and $2\vec{b}_1$ to be the same. A vector it does change by knocking off its span, is $\vec{a}_1 + \vec{a}_2 + \vec{a}_3$ which T_1 maps to $\vec{b}_1 + \vec{b}_2 + \vec{b}_3$ and T_2 maps to $\vec{b}_1 + \vec{b}_2 + 2\vec{b}_3$, which are not scalar multiples of each other. What this shows us is that we need at least 4 one-dimensional subspaces to define a transformation in the projective linear group. In more geometric language, a collineation is defined entirely by where it takes a given quadrangle [6].

3.2 Polygons

It is relatively natural to transfer our idea of polygons from the real plane to finite planes:

Definition 3.1 *An **n-gon** is a collection of n points, arranged in a cycle such that every point is adjacent to two others.*

This abstracts our euclidean idea of polygons: a polygon is a cycle of points with a sense of adjacency. This means that when we represent a polygon as a tuple, it is considered the same polygon when you reverse the order or shift the position. For example, the polygon P could equally well be represented by (p_1, p_2, p_3, p_4), (p_4, p_3, p_2, p_1) or (p_4, p_1, p_2, p_3). When we use such ordered pairs, we will call them *representations*, and will have to be careful of the fact that any representation is just one of many[1].

There are various possible constraints one can place on these polygons, two of which are relevant here. The weaker condition, that no three *adjacent* points are colinear, has been studied by Lazebnik et al [11]. However, when we speak of polygons from now on, we will assume a stronger condition, that they are in general position. Our use of this concept originates with Shwartz's work on the pentagram map [12].

Definition 3.2 *An n-gon is in **general position** iff no three points are colinear.*

This condition will serve us in particular when we consider the pentagram map ourselves in Section 4.

[1]Specifically, it is clear to see that any n-gon with $n \geq 3$ has $2n$ representations.

Finally, we will consider how automorphisms relate to polygons. For this, we will borrow language from Euclidean geometry. When we compare shapes going back to elementary school, we considered them *similar* if one could be taken to another by the angle-preserving automorphisms of the plane: scaling, rotation, and translation. These are, in a sense, the fundamental automorphisms of Euclidean geometry, as they preserve the relations of angles and relative length which it usually studies. Analogously, collineations are the fundamental automorphisms of projective geometry, as they preserve incidence. As such, we will define the similarity of polygons on a projective plane as such:

Definition 3.3 *Two n-gons A, B on a projective plane π are* ***similar*** *iff there exists an collineation $\phi : \pi \to \pi$ such that $(a_1, a_2, \ldots, a_n)$ is a representation of A and $(\phi a_1, \phi a_2, \ldots, \phi a_n)$ is a representation of B. We write this as $\phi A = B$.*

In other words, ϕ maps A to B in a way that preserves the adjacency of the vertices of B.

By the nature of collineation, on any plane π, all triangles in general position are similar to each other, and all quadrangles in general position are similar to each other. Thus, the smallest polygons for which similarity is nontrivial are pentagons. As such, these will be the focus of our consideration.

Theorem 3.1.

If A is an n-gon in general position and ϕ is an collineation, ϕA is in general position.

Proof. Consider some pentagon A in general position and some collineation ϕ.

Let a, b, c be three vertices of ϕA. Note that $\phi^{-1}a, \phi^{-1}b, \phi^{-1}c$ are vertices of A. Since A is in general position, $\phi^{-1}a, \phi^{-1}b, \phi^{-1}c$ are noncolinear. So, $[\phi^{-1}a \cdot \phi^{-1}b]$ is not incident with $\phi^{-1}c$. Since ϕ is a collineation, $\phi[\phi^{-1}a \cdot \phi^{-1}b] = \phi\phi^{-1}[a \cdot b] = [a \cdot b]$ is not incident with $\phi(\phi^{-1}(c)) = c$. Since c doesn't lie on $[a, b]$, we know a, b, c are not colinear.

So, any three vertices of ϕA are noncolinear. Thus, ϕA is in general position. □

So, being in general position is a property of an equivalence class of pentagons.

3.3 Classes of Similar Polygons

Let $\pi_p = (\mathcal{P}, \mathcal{L}, \mathcal{I})$ be any finite projective plane of order p.

Using our finite field representation, consider the following quadrangle:

$$Q = (\langle 1, 0, 0 \rangle, \langle 0, 1, 0 \rangle, \langle 0, 0, 1 \rangle, \langle 1, 1, 1 \rangle)$$

This is the quadrangle we proved must exist for any finite projective plane. Let us denote its elements u_1, u_2, u_3, u_4, respectively. We will call this the unit quadrangle in whatever plane π_p in which we encounter it.

Finally, for any pentagon P, let $[P]$ denote the equivalence class of pentagons similar to P. Most of the rest of this chapter will be moving towards the end of counting exactly how many of these equivalence classes exist on a given plane π_p. In other words, we will be attempting to count the size of the *moduli space* of pentagons on π_p. However, in service of this combinatorial end, we will derive other insights, particularly about the symmetries and general form of these pentagons.

Now, let's consider a subset of the pentagons in general position, which we'll call *unitary pentagons.*

Definition 3.4 *A pentagon is* ***unitary*** *if and only if it is in general position and has a representation* (u_1, u_2, u_3, u_4, a)*, with* $a \in \mathcal{P}$*.*

Theorem 3.2.

In π_p, there are $(p-2)(p-3)$ unitary pentagons.

Proof. Consider some unitary pentagon P on π_p with representation (u_1, u_2, u_3, u_4, x). Since P is in general position, x must not be a linear combination of any two of u_1, u_2, u_3, u_4.

Thus, these (and only these) are the forms that x cannot take:

$$au_1 + bu_2 = \langle a, b, 0 \rangle$$

$$au_1 + bu_3 = \langle a, 0, b \rangle$$
$$au_2 + bu_3 = \langle 0, a, b \rangle$$
$$au_1 + bu_4 = \langle a+b, b, b \rangle$$
$$au_2 + bu_4 = \langle b, a+b, b \rangle$$
$$au_3 + bu_4 = \langle b, b, a+b \rangle$$

From this, we can see that x cannot have a zero in any of its three positions, and no two positions can be equal. Having these properties, however, guarantees x is not in any of the above forms, and so is not a linear combination of any two points of our unit quadrangle.

We can set $x_0 = 1$ since we consider vectors the same up to scaling. So, let $x = \langle 1, a, b \rangle$. Since there are p elements of $\mathbb{F}_p$, and we're not allow to choose 0 or 1, we can choose $p-2$ possible values for a. Since $a \neq b$, this leaves us with $p-3$ possibilities for b, giving us $(p-2)(p-3)$ possibilities.

So, since everything else is fixed and there are $(p-2)(p-3)$ possible values of x, there are $(p-2)(p-3)$ unitary pentagons on π_p. □

Theorem 3.3.

Any pentagon in general position is similar to some unitary pentagon.

Proof. Consider any pentagon in general position, T on π_p. Let $(t_1, t_2, t_3, t_4, t_5)$ by a representation of T. Since T is in general position, $Q = (t_1, t_2, t_3, t_4)$ is a quadrangle. Let ϕ be the unique transformation which takes that quadrangle Q to the unit quadrangle U.

So, $\phi T = (u_1, u_2, u_3, u_4, \phi t_5)$. By definition, T is similar to ϕT, and ϕT is unitary. So, for any pentagon is similar to some unitary pentagon. □

For us, unitary pentagons will serve an important purpose as representatives of their similarity equivalence classes. In particular, as these classes get larger and larger as the order p of π_p grows, we will demonstrate that the number of unitary pentagons in each class is nonetheless very strictly constrained. In fact, we will show the only possible numbers of unitary pentagons in an equivalence class $[P]$ are 10 and its divisors. Establishing this requires accounting for the symmetries of pentagons. We can leverage these symmetries to give an exact account of how many pentagon classes exist with nontrivial symmetries.

Definition 3.5 *A **symmetry** of a pentagon P is an collineation ϕ such that $\phi P = P$.*

Theorem 3.4.

For any pentagon P on π_p, there is some s such that $s|10$ and each $Q \in [P]$ has s symmetries. Furthermore, the symmetries of each $Q \in [P]$ form a group, and that group is isomorphic with the symmetry group of every other pentagon in $[P]$.

Proof. Consider some pentagon P on π_p. Consider the set of symmetries $\mathcal{S}$ of P. Note that for $\phi_1, \phi_2 \in \mathcal{S}$, $\phi_1\phi_2 P = \phi_1 P = P$, so $\phi_1\phi_2 \in \mathcal{S}$. Furthermore, since $\phi_1 P = P$, we see that $P = \phi_1^{-1}P$, so $\phi_1^{-1} \in \mathcal{S}$. So, since $\mathcal{S}$ is a closed subset of $\mathrm{PGL}(3, p)$, $\mathcal{S}$ is a group.

Note that each automorphism in $\mathcal{S}$ permutes the 5 verticies of P in some way that preserves their adjacency. Furthermore, this permutation is unique for each unique symmetry, as each permutation tells us exactly where 4 of the vertices go, which uniquely defines a collineation. The group of all such permutations is of course the 5th dihedral group, D_5. So, there is some injective homomorphism $h : \mathcal{S} \to D_5$ which maps each automorphism in $\mathcal{S}$ uniquely to the element of D_5 which represents how it permutes the vertices of P. Since h is injective, $\mathcal{S} \cong \text{im}(h)$. However, since $\text{im}(h)$ is a subgroup of D_5, its size must divide $|D_5| = 10$. So, $|\text{im}(h)| = |\mathcal{S}|$ divides 10. Let $s = |\mathcal{S}|$.

Now, consider some symmetry ϕ of $Q \in [P]$, which permutes the vertices of Q some way. For any other pentagon $R \in [P]$, we can construct a symmetry that permutes the vertices of R in the same way. By definition, there exists some γ which maps Q to R. Thus, the map $\gamma\phi\gamma^{-1}$ maps R to Q, permutes the vertices, and then reverses the mapping, taking Q back to R. Thus, any symmetries Q has, R does as well. Since this relationship is symmetrical, this means that all members of $[P]$ have the same number of symmetries. In particular, they have s symmetries.

For any pentagon P, there is some s such that $s|10$ and each $Q \in [P]$ has s symmetries, and the symmetry group of each $Q \in [P]$ is isomorphic to the symmetry group of every other pentagon in $[P]$. □

Remark 3.1 *Due to this result, we can speak of a given equivalence class as having some number of symmetries s and some symmetry group (up to isomorphism) $\mathcal{S}$. In other words, it the symmetries of a polygon are invariant under collineation.*

Now, to prove the next result, we will define a new operation on pentagons. However, its motivation might be unclear if the thought process if left unexplained. To begin, consider a much easier to imagine space: $\mathbb{R}^2$ along with the angle-preserving transformations: those which are a combination of translations, rotations, reflections, and scaling. If we imagine an equilateral triangle in this space, we can see that any permuation σ of its vertices which preserves adjacency (which in the case of triangles, is all of them) has a corresponding angle-preserving transformation τ_σ which permutes the vertices in that way.

However, what about irregular triangles? You can't transform $\mathbb{R}^2$ in such a way that it takes one vertex with angle θ to another with angle $\neq \theta$, since by definition that would require a transformation that does not preserve angles. For this reason, angle-preserving transformations cannot in general take any triangle to any other triangle. It's this lack of transitivity that makes the concept of similarity in geometry nontrivial. However, these transformations can take any line segment to any other line segment, meaning we can permute at least two of the three vertices correctly, and let the last one go where it may. We will call this *forcing*.

Take, for example, the triangle ABC in figure 3.1. We will try to force the permutation $ABC \to CAB$ on it. We've chosen A, B can be permuted correctly, so since the permutation does not correspond to a symmetry of ABC, C is taken to an entirely new point, C'.

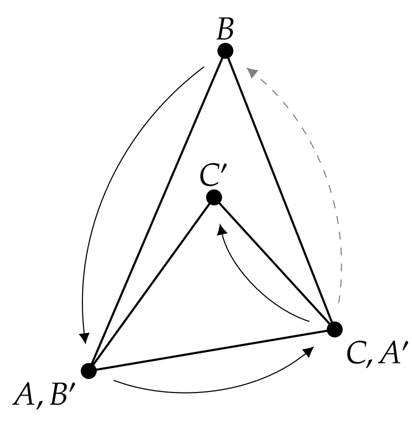

Figure 3.1: Forcing a triangle with the permutation $ABC \to CBA$.

If we were to fix some standard line segment, like for example the one between $(0, 0)$ and $(1, 0)$, we can consider only the triangles that contain that segment. From there, we can always make sure that whatever segment gets taken to that side of the triangle by a permutation is the one that is mapped correctly, giving us a comparison of all possible permutation forcings, as seen in figure 3.2.

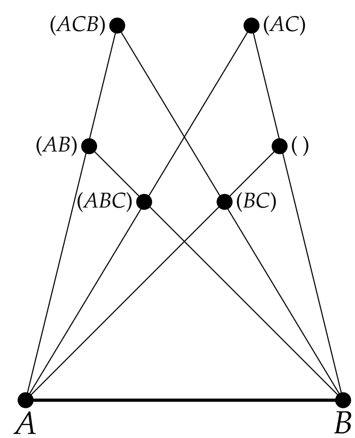

Figure 3.2: All possible forcings, with whatever lands on AB always being the correct part. The resulting third points are labeled with the permutation that created them.

There is a slight issue with this strategy in that two points don't contain enough information about orientation, and so our choice of reflection is arbitrary. I've chosen the orientation which makes whatever is permuted to C as close to C as possible. However, this is not an issue with more free types of transformations, which are transitive on larger sets of points. For example, Something analogous can be done with quadrilaterals and affine transformations, which can take any triangle to any other triangle. See figure 3.3. Note that in the example chosen, pairs of permutations land in the same place; we will see that things like this correspond to there being nontrivial symmetries of our shape. In this case, the symmetry is the one that permutes the vertices like $(AB)(CD)$.

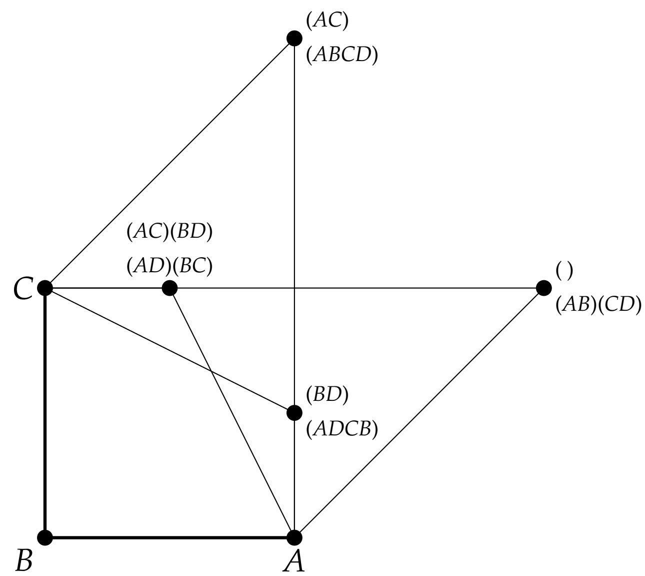

Figure 3.3: Forcings of affine transformations on a quadrilateral. Note that pairs of points land in the same place.

Moving up one more level of freedom with our transformations, that's what we're going to do with unitary pentagons. Given some unitary pentagon U and a permutation of its vertices σ, we want to see which vertices are taken to the unit quadrangle, 'reverse engineer' a collineation which does that, and apply it to the whole pentagon.

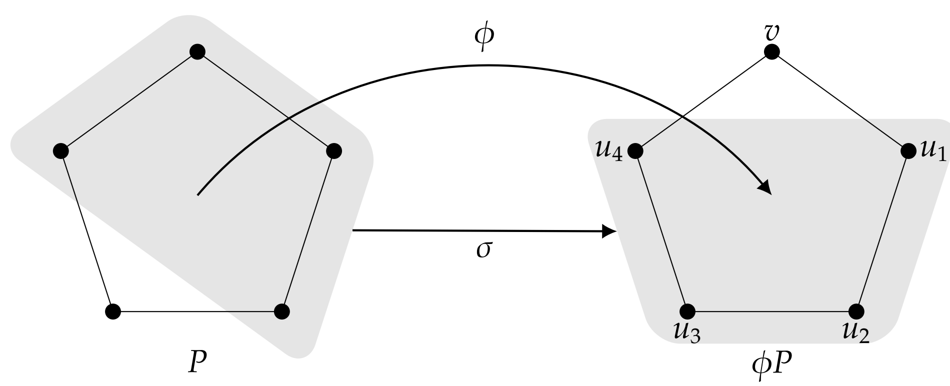

Figure 3.4: Relationships inherent to the forcing operation of collineations on pentagons.

With that motivation in mind, here is the formal definition:

Definition 3.6 *Let Σ be the set of set of permutations of a cycle of 5 points that preserves adjacency*[2]*, and $\mathcal{U}_p$ be the set of unitary pentagons on π_p. Then, let the* ***forcing operation*** *$\otimes : \Sigma \times \mathcal{U}_p \to \mathcal{U}_p$ be an operation defined as the following:*

For some $\sigma \in \Sigma$ and $U \in \mathcal{U}_p$, let Q be the quadrangle in U which σ maps to the unit quadrangle. Let ϕ be the collineation which takes Q to the unit quadrangle. Then, $\sigma \otimes U = \phi U$.

Remark 3.2 *It should be clear that if there is a symmetry of $U \in \mathcal{U}_p$ that permutes the vertices the same way σ does, then $\sigma \otimes U = U$. Otherwise, $\sigma \otimes U$ will be some other unitary pentagon similar to U by definition.*

Theorem 3.5.

$\otimes$ is a group action of Σ on $\mathcal{U}_p$.

Proof. Consider any $U \in \mathcal{U}_p$. Let ϕ_σ denote the collineation that is generated as part of the definition of $\sigma \otimes U$.

The identity permutation e takes the unit quadrangle in U to itself, so the collineation ϕ_e is the identity, so $e \otimes U = \phi_e U = U$. Thus, the identity requirement of a group action is satisfied.

Now, consider any two permutations $\alpha, \beta \in \Sigma$. Consider only the unit quadrangle in U. By construction, ϕ_α, ϕ_β permute this quadrangle the same way as α, β do, respectively. So, $\phi_{\alpha\beta}$ permutes it the same way $\alpha\beta$ does. Since a collineation is defined by where it takes a given quadrangle,

$$\alpha \otimes (\beta \otimes U) = \phi_\alpha \phi_\beta U = \phi_{\alpha\beta} U = (\alpha\beta) \otimes U$$

Thus, the compatibility requirement of group actions is satisfied.

As both requirements are satisfied, $\otimes$ is a group action of Σ on $\mathcal{U}_p$. □

[2]This is clearly isomorphic to D_5, however, it is better to think of them as permutations.

Theorem 3.6.

For any pentagon P on some π_p, $[P]$ contains $\frac{10}{s}$ unitary pentagons, where s is the number of symmetries of $[P]$.

Proof. Consider any pentagon P on some finite projective plane. Consider some unitary pentagon $U \in [P]$, which must exist by Theorem 3.3.

Let Σ be the 10 permutations of the vertices of U which preserve adjacency. Now, consider the subgroup Σ' of Σ of permutations which correspond to symmetries of U. We know this is a subgroup because it must be isomorphic to the symmetry group of U. Obviously, $|\Sigma'| = s$.

Now, consider the left cosets of the form $\alpha\Sigma'$, where $\alpha \in \Sigma$. There must be $\frac{|\Sigma|}{|\Sigma'|} = \frac{10}{s}$ of these. Now, consider any two $\sigma_1, \sigma_2 \in \alpha\Sigma'$. By definition, $\sigma_i = \alpha\sigma_i'$, with $\sigma_i' \in \Sigma'$. However, since both σ_1', σ_2' are symmetries, they don't change U, meaning $\sigma_1 \otimes U = (\alpha\sigma_1') \otimes U = \alpha \otimes (\sigma_1' \otimes U) = \alpha \otimes U$, and $\sigma_2 \otimes U = (\alpha\sigma_2') \otimes U = \alpha \otimes (\sigma_2' \otimes U) = \alpha \otimes U$. Therefore, all elements of the same right coset map U to the same pentagon under $\otimes$.

Now, consider some α, β such that $\alpha \otimes U = \beta \otimes U$. So, $(\alpha^{-1}\beta) \otimes U = U$, meaning $\alpha^{-1}\beta \in \Sigma'$. Thus, $\beta = \alpha(\alpha^{-1}\beta) \in \alpha\Sigma'$. Thus, two permutations take U to the same pentagon under $\otimes$ only if they are in the same right coset of Σ'.

So, the number of distinct pentagons that U is be taken to by Σ under $\otimes$ is the number of right cosets of Σ', which is $\frac{10}{s}$.

Now, consider $\mathcal{U} = \{\sigma \otimes U \mid \sigma \in \Sigma\}$. This must be the set of all the unitary pentagons in $[P]$, since if a unitary pentagon is similar to U, then four of U's vertices must map to its unitary quadrangle, which would correspond to some permutation. We've seen that there must be $\frac{10}{s}$ elements in $\mathcal{U}$.

For any pentagon P, $[P]$ contains $\frac{10}{s}$ unitary pentagons, where s is the number of symmetries of $[P]$. □

With this relationship, we gain more leverage over the problem of the number of equivalence classes. We will proceed by trying to count the number of classes that display 2, 5, and 10 symmetries. While these form a very small porportion of the total number of pentagons, since we know that there are $(p-2)(p-3)$ unitary pentagons, and know exactly how many unitary pentagons should be in each class as a function of its symmetries, we can solve for the number of pentagons with only the trivial symmetry. The reason this is attractive is that the presence of nontrivial symmetries allows us to explore what form those pentagons take in terms of the solutions to linear equations. We can do this by solving for the collineation matrix and fifth vertex of a general unitary pentagon. The number of different solutions to these equations will also tell us how many such pentagons exist.

Definition 3.7 *A polygon P has* ***n-fold symmetry*** *if it has a symmetry of order n.*

Theorem 3.7.

The unitary pentagons with 2-fold symmetry with a symmetry that fixes u_1 are precisely those of the form $(u_1, u_2, u_3, u_4, \langle 1, \frac{1}{2-b}, b\rangle)$ for $b \notin \{0, 1, 2\}$.

Proof. Let p be prime. Consider π_p.

Let us consider the most general possible unitary pentagon with 2-fold symmetry. Let us call the pentagon P and the symmetry σ. Every P has some point r that is fixed by σ, since it has an odd number of sides. Since we want a unique representative of P up to automorphism, we can fix four of P's 5 points such that P is a certain unique unitary pentagon. Thus, consider the following pentagon, with $r = \langle 1, 0, 0\rangle$:

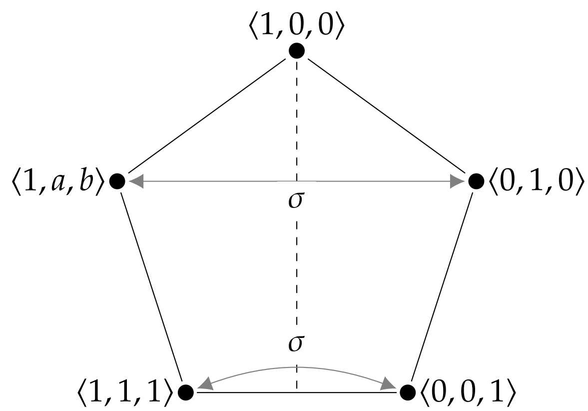

Figure 3.5: Our representative pentagon

We know that the last vertex must be of the form $\langle 1, a, b \rangle$, with $a \neq b$ and $a, b \notin \{0, 1\}$, thanks to how Theorem 3.2 was proven. So, we have the following conditions on σ, up to scaling:

$$\sigma\langle 1,0,0\rangle = \langle 1,0,0\rangle$$

$$\sigma\langle 0,1,0\rangle = \langle 1,a,b\rangle$$

$$\sigma\langle 0,0,1\rangle = \langle 1,1,1\rangle$$

$$\sigma\langle 1,1,1\rangle = \langle 0,0,1\rangle$$

$$\sigma\langle 1,a,b\rangle = \langle 0,1,0\rangle$$

The first three conditions are satisfied in the usual way, by setting the columns of the matrix to be the desired output:

$$\begin{pmatrix} 1 & 1 & 1 \\ 0 & a & 1 \\ 0 & b & 1 \end{pmatrix}$$

We can scale each column by s_1, s_2, s_3 respectively, and look for a solution in terms of a, b that satisfies the third condition.

$$\begin{pmatrix} s_1 & s_2 & s_3 \\ 0 & s_2 a & s_3 \\ 0 & s_2 b & s_3 \end{pmatrix} \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}$$

Multiplying this out gives $s_1 + s_2 + s_3 = 0$, $s_2 a + s_3 = 0$, and $s_2 b + s_3 = 1$. The solution to this linear system of equations is $s_1 = \frac{1-a}{a-b}$, $s_2 = \frac{1}{b-a}$, and $s_3 = \frac{a}{a-b}$. So, we have σ in terms of a, b:

$$\sigma = \begin{pmatrix} \frac{1-a}{a-b} & \frac{1}{b-a} & \frac{a}{a-b} \\ 0 & \frac{a}{b-a} & \frac{a}{a-b} \\ 0 & \frac{b}{b-a} & \frac{a}{a-b} \end{pmatrix}$$

Since $a \neq b$, we can scale this matrix by $a - b$ to simplify it:

$$\sigma = \begin{pmatrix} 1-a & -1 & a \\ 0 & -a & a \\ 0 & -b & a \end{pmatrix}$$

So, consider the third condition:

$$\begin{pmatrix} 1-a & -1 & a \\ 0 & -a & a \\ 0 & -b & a \end{pmatrix} \begin{pmatrix} 1 \\ a \\ b \end{pmatrix} = \begin{pmatrix} 1 - 2a + ab \\ -a^2 + ab \\ 0 \end{pmatrix} = \begin{pmatrix} 0 \\ k \\ 0 \end{pmatrix}$$

Thus, $2a - ab - 1 = 0$. Manipulating this gives $a = \frac{1}{2-b}$.

Since a is expressed in terms of b, we've parametrized σ in terms of b. Furthermore, this is valid by construction whenever a exists, so $b \neq 2$, meaning $b \notin \{0, 1, 2\}$ when we include our previous requirements for b.

Now, we will demonstrate every pentagon of the form $P = (u_1, u_2, u_3, u_4, \langle 1, \frac{1}{2-b}, b\rangle)$ is a unitary pentagon with 2-fold symmetry and a symmetry that fixes u_1.

From our previous computation we get the following matrix. We scale it by $2-b$ to simplify it.

$$\sigma = \begin{pmatrix} 1-\frac{1}{2-b} & -1 & \frac{1}{2-b} \\ 0 & -\frac{1}{2-b} & \frac{1}{2-b} \\ 0 & -b & \frac{1}{2-b} \end{pmatrix} \cong \begin{pmatrix} 1-b & b-2 & 1 \\ 0 & -1 & 1 \\ 0 & b(b-2) & 1 \end{pmatrix}$$

We can show σ reflects P by fixing u_1 and exchanging the other two pairs of vertices by direct computation:

$$\begin{aligned}
\sigma\langle 1,0,0\rangle &= \langle 1-b,0,0\rangle \cong \langle 1,0,0\rangle \\
\sigma\langle 0,1,0\rangle &= \langle b-2,-1,b(b-2)\rangle \cong \langle 1,\tfrac{1}{2-b},b\rangle \\
\sigma\langle 0,0,1\rangle &= \langle 1,1,1\rangle \\
\sigma\langle 1,1,1\rangle &= \langle 0,0,b(b-2)+1\rangle \cong \langle 0,0,1\rangle \\
\sigma\langle 1,\tfrac{1}{2-b},b\rangle &= \langle 0,b+\tfrac{1}{b-2},0\rangle \cong \langle 0,1,0\rangle
\end{aligned}$$

So, P is a pentagon for which σ is a symmetry of order 2 which fixes u_1, meaning P has 2-fold symmetry, and whose symmetry fixes u_1.

So, the unitary pentagons with 2-fold symmetry with a symmetry that fixes u_1 are precisely those of the form $(u_1, u_2, u_3, u_4, \langle 1, \frac{1}{2-b}, b\rangle)$ for $b \notin \{0,1,2\}$. □

Theorem 3.8.

The unitary pentagons with 5-fold symmetry are precisely those of the form $(u_1, u_2, u_3, u_4, \langle 1, -g, g+1\rangle)$ *where* $g^2 - g - 1 - 0$.

Proof. Let p be prime. Consider π_p.

Consider any unitary pentagon R with 5-fold symmetry. Just as in the previous theorem, it must be of the form $R = (u_1, u_2, u_3, u_4, \langle 1, a, b\rangle)$ with $a \neq b$. By definition, there is some ρ' of R that is of order 5. This means that there is some ρ such that:

$$\rho\langle 1,0,0\rangle = \langle 0,1,0\rangle$$
$$\rho\langle 0,1,0\rangle = \langle 0,0,1\rangle$$
$$\rho\langle 0,0,1\rangle = \langle 1,1,1\rangle$$
$$\rho\langle 1,1,1\rangle = \langle 1,a,b\rangle$$
$$\rho\langle 1,a,b\rangle = \langle 1,0,0\rangle$$

This is because $(\rho')^n$ would have to be ρ for some n, because there are 5 distinct ways to permute the vertices of R which preserve orientation (of which ρ does one), and 5 distinct powers of ρ'. ρ' must preserve orientation, as otherwise its order world be even.

To satisfy the first three constraints, ρ must be in the following form, with $s_1, s_2, s_3 \in \mathbb{F}_p$:

$$\rho = \begin{pmatrix} 0 & 0 & s_3 \\ s_1 & 0 & s_3 \\ 0 & s_2 & s_3 \end{pmatrix}$$

From the third constraint, we get:

$$\rho = \begin{pmatrix} 0 & 0 & s_3 \\ s_1 & 0 & s_3 \\ 0 & s_2 & s_3 \end{pmatrix}\begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} = \begin{pmatrix} s_3 \\ s_1 + s_3 \\ s_2 + s_3 \end{pmatrix} = \begin{pmatrix} 1 \\ a \\ b \end{pmatrix}$$

This means that $s_3 = 1$, $s_1 = a - 1$, and $s_2 = b - 1$. Finally, let's consider the final constraint, with k as some arbitrary constant in $\mathbb{F}_p$:

$$\begin{pmatrix} 0 & 0 & 1 \\ a-1 & 0 & 1 \\ 0 & b-1 & 1 \end{pmatrix} \begin{pmatrix} 1 \\ a \\ b \end{pmatrix} = \begin{pmatrix} b \\ a+b-1 \\ a(b-1)+b \end{pmatrix} = \begin{pmatrix} k \\ 0 \\ 0 \end{pmatrix}$$

So, $a + b - 1 = 0$, meaning $a = 1 - b$. So:

$$a(b-1) + b = (1-b)(b-1) + b = -b^2 + 3b - 1 = 0$$

Thus, $b^2 - 3b + 1 = 0$. Now, let $g = b - 1$. So, $(g-1)^2 - 3(g-1) + 1 = g^2 - t - 1 = 0$. Note that $a = 1 - b = 1 - (g+1) = -g$. So, we can parametrize all unitary pentagons with 5-fold symmetry on π_p as the unit quadrangle followed by $\langle 1, -g, g+1 \rangle$, for all g such that $g^2 - g - 1 = 0$.

We will now demonstrate every such pentagon is a unitary pentagon with 5-fold symmetry. So, let $P = (u_1, u_2, u_3, u_4, \langle 1, -g, g+1 \rangle)$ be a pentagon with $g^2 - g - 1 = 0$.

From our previous computation we get the following matrix.

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

We can show ρ rotates P by taking u_1 to u_2, u_2 to u_3, and so on, by direct computation:

$$\rho\langle 1,0,0\rangle = \langle 0,-g-1,0\rangle \cong \langle 0,1,0\rangle$$
$$\rho\langle 0,1,0\rangle = \langle 0,0,g\rangle \cong \langle 0,0,1\rangle$$
$$\rho\langle 0,0,1\rangle = \langle 1,1,1\rangle$$
$$\rho\langle 1,1,1\rangle = \langle 1,-g,g+1\rangle$$
$$\rho\langle 1,-g,g+1\rangle = \langle g+1,0,-g^2+g+1\rangle = \langle g+1,0,0\rangle \cong \langle 1,0,0\rangle$$

So, P is a pentagon for which ρ is a symmetry of order 5, meaning P has 5-fold symmetry.

So, the unitary pentagons with 5-fold symmetry are precisely those of the form $(u_1, u_2, u_3, u_4, \langle 1, -g, g+1\rangle)$ where $g^2 - g - 1 - 0$. □

Definition 3.8 *A* ***regular pentagon*** *P is one for which the symmetry group $\mathcal{S}$ of $[P]$ is isomorphic to D_5.*

Theorem 3.9.

A pentagon in π_p with 5-fold symmetry must be a regular pentagon.

Proof. Consider any pentagon with 5-fold symmetry R. It must be similar to some unitary pentagon U, and by Theorem 3.8, we know that U is of the form $(u_1, u_2, u_3, u_4, \langle 1, -g, g+1\rangle)$, for some g such that $g^2 - g - 1 = 0$.

Now, let $b = g + 1$. For $\frac{1}{2-b} = \frac{1}{1-g} = \frac{1+g}{1-g^2}$. Note that, since $g^2 - g - 1 = 0$, $1 - g^2 = -g$. So, $\frac{1+g}{1-g^2} = \frac{1+g}{-g} = -1 - \frac{1}{g}$. Since $g^2 - g - 1 = 0$, $g^2 = g + 1$, meaning $g = 1 + \frac{1}{g}$. So, $-1 - \frac{1}{g} = -g$. Thus, $\frac{1}{2-b} = -g$. This means that U is of the form $(u_1, u_2, u_3, u_4, \langle 1, \frac{1}{2-b}, b\rangle)$.

Note that $g \notin \{-1, 0, 1\}$ because none of those are roots of $x^2 - x - 1$ modulo any prime, since $(-1)^2 - (-1) - 1 = 1$, $0^2 - 0 - 1 = -1$, and $1^2 - 1 - 1 = -1$. So, $b \notin \{0, 1, 2\}$. These two facts mean that, by Theorem 3.7, U has a 2-fold symmetry.

Since R and U have the same symmetry group by Theorem 3.4, R has a 2-fold symmetry. However, since R has 2-fold symmetry and 5-fold symmetry, that means it has a symmetry of order 5 and a symmetry of order 2. Since its symmetry group is a subgroup of D_5, this means its symmetry group must be D_5. So, R is a regular pentagon.

So, a pentagon on π_p with 5-fold symmetry must be a regular pentagon. □

What we've proven here is a fact that we know is intuitively true in the real plane; you can't have a pentagon with rotational symmetry that doesn't have reflective symmetry. However, unlike the proofs one might have seen in other types of geometry, we've had to prove this fact without the use of a metric and the derived concept of angles. This implies that this is a property that is more general than simply being a property of metric spaces.

Theorem 3.10.

For prime p, π_p has a regular pentagon only if $p \bmod 5$ is congruent to -1, 0, or 1.

Proof. Since p is prime, the collineation group of π_p is $\mathrm{PGL}(3, p)$. Since π_p is prime, then there is some regular pentagon R, for which the symmetry group $\mathcal{S}$ of R is isomorphic to D_5. Thus, $D_5 \leq \mathrm{PGL}(3, p)$. So, $|D_5|$ divides $|\mathrm{PGL}(3, p)|$, meaning that 10 divides $p^3(p^3 - 1)(p^2 - 1)$, since that is the order of $\mathrm{PGL}(3, p)$ [5].

Since $p^3(p^3 - 1)(p^2 - 1) \equiv 0 \mod 10$, it must also be true that:

$$p^3(p^3 - 1)(p^2 - 1) \equiv 0 \mod 2$$
$$p^3(p^3 - 1)(p^2 - 1) \equiv 0 \mod 5$$

The first is always true, as when p^3 is odd, $p^3 - 1$ is even and vice versa, so the whole expression must always be even. Since 5 is prime, the second implies that one of the following are true:

$$p^3 \equiv 0 \mod 5$$
$$p^3 - 1 \equiv 0 \mod 5$$
$$p^2 - 1 \equiv 0 \mod 5$$

Manipulating, we get:

$$p^3 \equiv 0 \mod 5$$
$$p^3 \equiv 1 \mod 5$$
$$p^2 \equiv 1 \mod 5$$

Finally, taking roots gives us:

$$p \equiv 0 \mod 5$$
$$p \equiv 1 \mod 5$$
$$p \equiv \pm 1 \mod 5$$

So, $p \mod 5$ is congruent to -1, 0, or 1.

Thus, for prime p, if π_p has a regular pentagon, then $p \mod 5$ is congruent to -1, 0, or 1. □

We now finally can begin to count the equivalence classes of similar pentagons. We'll approach this by counting the number of pentagons with 2, 5, and 10 symmetries, and then use that fact, and what we know about unitary pentagons, to count the classes of entirely irregular pentagons. After that, counting the total number is as simple as summing these values up.

Lemma 3.1.

If p is a prime with $p \equiv \pm 1 \mod 5$, then there exist two $g \in \mathbb{F}_p$ such that $g^2 - g - 1 = 0$.

Proof. Let p be a prime with $p \equiv \pm 1 \mod 5$. By the law of quadratic reciprocity, we know that:

$$\left(\frac{p}{5}\right)\left(\frac{5}{p}\right) = (-1)^{\frac{5-1}{2}\frac{p-1}{2}} = ((-1)^2)^{\frac{p-1}{2}} = 1$$

$\left(\frac{r}{q}\right)$ is the Legendre symbol, which is 1 when there is some $n^2 \equiv r \mod q$ and -1 otherwise. In the first case, we call r a *quadratic residue* $\mod q$. Since the only possible values of $\left(\frac{p}{5}\right)$ and $\left(\frac{5}{p}\right)$ are ± 1 and their product is 1, then we can see that:

$$\left(\frac{p}{5}\right) = \left(\frac{5}{p}\right)$$

Now, consider that, since $1^2 \equiv 1 \mod 5$ and $2^2 = -1 \mod 5$, and $p \equiv \pm 1 \mod 5$, this means that p is a quadratic residue $\mod 5$. As such, $\left(\frac{p}{5}\right) = 1$. However, this means that $\left(\frac{5}{p}\right) = 1$, meaning there is some $n \in \mathbb{Z}$ such that $n^2 = 5 \mod p$.

So, there is some $G \in \mathbb{F}_p$ such that $G^2 = 5$. Consider $g = \frac{1}{2}(G+1)$. Since $G = 2g-1$, $(2g-1)^2 = 4g^2 - 4g + 1 = 5$. Subtracting 5 from both sides and dividing by 4 gives us $g^2 - g - 1 = 0$.

Now, consider that:

$$\begin{aligned}(1-g)^2 - (1-g) - 1 &= 1 - 2g + g^2 - 1 + g - 1 \\ &= g^2 - g - 1 \\ &= 0\end{aligned}$$

So, $1 - g$ is also a root of $x^2 - x - 1 = 0$.

So, if p is a prime with $p \equiv \pm 1 \mod 5$, then exist two $g \in \mathbb{F}_p$ such that $g^2 - g - 1 = 0$.

□

Definition 3.9 *Let* $\mathbf{C_{n,p}}$ *be the number of equivalence classes with exactly n symmetries on* π_p.

Definition 3.10 *Let the indexed constants* $\mathbf{r_p}$ *be the following:*

$$r_p = \begin{cases} 1 & p \equiv 0 \mod 5 \\ 2 & p \equiv \pm 1 \mod 5 \\ 0 & p \equiv \pm 2 \mod 5 \end{cases}$$

Theorem 3.11.

For any prime number p, $C_{10,p} = r_p$.

Proof. Let p be prime.

From Theorem 3.8 and Theorem 3.9, we know that the unitary pentagons are precisely those of the form $(u_1, u_2, u_3, u_4, \langle 1, -g, g+1)$ with $g^2 - g - 1 = 0$. Since, by Theorem 3.6, each similarity class with 10 symmetries has only $\frac{10}{10} = 1$ unitary pentagons, the number of similarity classes is the same as the number of unitary pentagons, which is the same as the number of roots of $x^2 - x - 1$ in $\mathbb{F}_p$.

When $p \equiv 0 \mod 5$, we can see $p = 5$, since p is prime. There is only one root over $\mathbb{F}_5$, which is 3. This can be checked directly. So, in this case, $C_{10,p} = 1 = r_p$.

When $p \equiv \pm 1 \mod 5$, by Lemma 3.1, $x^2 - x - 1$ has two roots over $\mathbb{F}_p$. So, $C_{10,p} = 2 = r_p$.

Finally, when $p \equiv \pm 2 \mod 5$, it follows from the contrapositive of Theorem 3.10 that there are no regular pentagons[3] on π_p. So, $C_{10,p} = 0 = r_p$.

So, $C_{10,p} = r_p$ in all cases. □

[3]This also demonstrates obliquely that there are no roots of $x^2 - x - 1$ on $\mathbb{F}_p$ when $p \equiv \pm 2 \mod 5$.

This result can also be seen in terms of the quadratic formula. Since $g = \frac{1 \pm \sqrt{5}}{2}$, there are two roots when $\sqrt{5}$ exists. There is an exception when $p = 5$, since then $\sqrt{5} = \sqrt{0} = 0$, meaning there is only one $g = \frac{1}{2}$. Finally, $\sqrt{5}$ doesn't exist for $\mathbb{F}_p$ when $p \equiv \pm 2 \bmod 5$, so there are no roots.

It might seem strange that there are almost always two classes of regular pentagons whenever they exist. We'd usually think of all regular pentagons as being similar; in this case, only π_5 matches our intuition. However, the fact that π_5 has only one similarity class of regular pentagons is actually the strange thing here. We usually think of all regular pentagons being similar because it's common in euclidean geometry to exclude star polygons, due to their self-intersection. However, on the real projective plane, the idea of the line segment depends on your particular perspective; on finite planes it is entirely incoherent, especially without a clear idea of the 'inside' of a polygon. However, if (a, b, c, d, e) is regular then (a, c, e, b, d) has the same symmetries and so is also regular. It's pretty clear to see that the two shapes aren't similar to each other in $\mathbb{RP}^2$; this is also true in most finite planes. It's actually a surprising and interesting property of π_5 that a regular pentagon is similar to its own star polygon! One caveat to this language is that, between the two classes of regular pentagons on π_p, there's no nonarbitrary way two call one class the 'normal' regular pentagons, and the other the 'star' regular pentagons. Being a star polygon is a dual relation within which, without any auxillary concept of distance, it is impossible to privilege one side over the other.

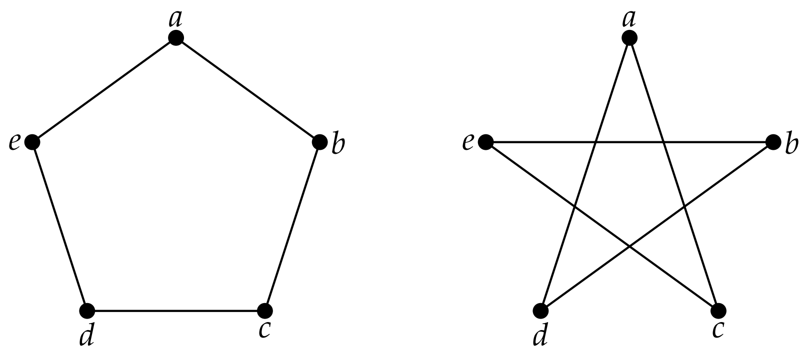

Figure 3.6: A pentagon and its star pentagon.

Theorem 3.12.

For any prime number p, $C_{5,p} = 0$.

Proof. Let p be prime. This follows directly from Theorem 3.9, as any pentagon with 5 symmetries must be regular and thus have 10 symmetries. So, there are no pentagons with exactly 5 symmetries, meaning $C_{5,p} = 0$. □

For the next theorem, we'll proceed in much the same way as Theorem 3.11, attempting to parametrize the unitary pentagons of a plane in such a way that we generate exactly 1 from each equivalence class, and then counting the number of pentagons generated.

Theorem 3.13.

For any prime number p, $C_{2,p} = p - 3 - r_p$.

Proof. Let p be prime. Consider the number of pentagons with 2-fold symmetry, which we will call C. Since the only such pentagons are those with 2 symmetries or with 10, $C = C_{2,p} + C_{10,p}$. So, $C_{2,p} = C - r_p$, by Theorem 3.11.

By Theorem 3.7, the pentagons that fix u_1 under a symmetry of order 2 are precisely those of the form $P_b = (u_1, u_2, u_3, u_4, \langle 1, \frac{1}{2-b}, b\rangle)$ with $b \notin \{0, 1, 2\}$. Since there are p elements of $\mathbb{F}_p$, there are $p - 3$ unit pentagons P_b parametrized by b. Furthermore, by construction, the fixed point of any P_b under reflection is u_1. If there are 4 other unitary pentagons similar to some P_b, there must be one for each possible vertex to fix (of u_2, u_3, u_4 and v), as we can transform the fixed point to be mapped onto any of them. So, it can't be similar to any other P_b since all P_b fix u_1. Otherwise, the P_b is regular. In either case, each P_b is in a different equivalence class, since there can be only one per class. Furthermore, any class must have a pentagon P_b, because we can transform any pentagon so it has the unitary quadrangle where we had it relative to the fixed point. So, we have parametrized one pentagon for each equivalence class. Thus, there are $p-3$ equivalence classes of pentagons with 2-fold symmetry.

Thus, $C = p - 3$. So, $C_{2,p} = C - r_p = p - 3 - r_p$.

So, for any prime number p, $C_{2,p} = p - 3 - r_p$. □

Theorem 3.14.
For any prime number p, $C_{1,p} = \frac{1}{10}((p-3)(p-7) + 4r_p)$.

Proof. Let p be prime. Since each equivalence class of unitary pentagons with s symmetries contains $\frac{10}{s}$ unitary pentagons (Theorem 3.6), and there are $(p-2)(p-3)$ unitary pentagons on π_p (Theorem 3.2), and because the only symmetries of possible of a pentagon are 1, 2, 5, and 10 (Theorem 3.4), $(p-2)(p-3) = \frac{10}{1}C_{1,p} + \frac{10}{2}C_{2,p} + \frac{10}{5}C_{5,p} + \frac{10}{10}C_{10,p}$. We can plug in the values gotten in the previous 3 theorems, (Theorem 3.11, Theorem 3.12, Theorem 3.13), to get:

$$\begin{aligned}\frac{10}{1}C_{1,p} + \frac{10}{2}C_{2,p} + \frac{10}{5}C_{5,p} + \frac{10}{10}C_{10,p} &= 10C_{1,p} + 5(p-3-r_p) + 2(0) + r_p \\ &= 10C_{1,p} + 5p - 15 - 4r_p\end{aligned}$$

So, $10C_{1,p} = (p-2)(p-3) - (5p - 15 - 4r_p)$. We can manipulate this expression like so:

$$\begin{aligned}(p-2)(p-3) - (5p-15-4r_p) &= p^2 - 5p + 6 - 5p + 15 + 4r_p \\ &= p^2 - 10p + 21 + 4r_p \\ &= (p-3)(p-7) + 4r_p\end{aligned}$$

Dividing this expression by 10, we get $C_{1,p} = \frac{1}{10}((p-3)(p-7) + 4r_p)$. □

Theorem 3.15.

For any prime p, the total number of equivalence classes of similar pentagons on π_p *is* $\frac{1}{10}((p+3)(p-3) + 4r_p)$.

Proof. Let p be prime. Because the only symmetries of possible of a pentagon are 1, 2, 5, 10 (Theorem 3.4), this can be verified simply by computing the sum $C_{1,p} + C_{2,p} + C_{5,p} + C_{10,p}$, the terms of which are given in Theorem 3.11, Theorem 3.12, Theorem 3.13, and Theorem 3.14, and rearranging. □

With this, we've proven all we want to for now about similarity classes of pentagons on π_p. Although our endpoint was an exact count of the number of these classes, this is not the only understanding we've gained. We've also characterized and classified these pentagon classes in terms of their symmetries, and applied the same combinatorial scrutiny to them as we did similarity classes as a whole. This in particular will give us the grounds to examine the pentagram map in the next chapter.

3.4 Other n-gons

Proving generalizations of the theorems in the previous section for n-gons with $n > 5$ is beyond the scope of this paper. However, it seems appropriate to provide an analysis of how such a generalization should be approached, what I believe is likely, and what I believe will pose difficulties.

Counting the number of unitary n-gons (and by counting extension equivalence classes as seen in Theorem 3.15) as we did for pentagons in Theorem 3.2 already poses certain problems. For one, we can only ever fix 4 points to be our unit quadrangle. This leaves $n - 4$ remaining points that can vary. Even for 2 or 3 such points, finding out which points allow for a polygon in general position is difficult because now the corresponding vectors need not only to be linearly independent of the known points that we fixed, but also of each other. Avoiding this kind of interdependence was precisely why unitary pentagons were employed. Still, as seen in papers like Lazebnik's [11], this sort of problem can be approached for fixed n by splitting the possibilities into cases and counting from there. However, what is less clear is how to approach the general case, where both the order of the plane p and the number of vertices n of the polygon are both unknown.

Theorem 3.10, where we proved which planes could contain regular pentagons, is different. At the very least, it seems that there is the outline of a generalized proof, at least for n-gons with n prime. The forward direction of the if and only if is generalizes for any prime; having regular n-gons implies a subgroup of $\mathrm{PGL}(3, p)$ isomorphic to D_n, which in turn implies that $2n$ divides the order of $\mathrm{PGL}(3, p)$. The rest follows naturally.

The strategy for the other direction is not so quickly derived from the pentagon case, but there is a clear suggestion of a path forward. Central to the pentagon proof is the construction of a matrix of order 5 over $\mathbb{F}_p$. This was made possible by the existence of a root to $g^2 - g - 1$, which was in turn guaranteed by the fact that 5 has a square root in $\mathbb{F}_p$. At first, it seemed to me that is was possible that a matrix of order n can be constructed using some relative of a square root of n (or possibly $-n$). However, after further investigation, I think that's unlikely. I currently suspect that what would be needed to construct a rotation matrix is the solution to a kind of polynomial. In $\mathbb{RP}^2$, the short diagonal[4] of a pentagon with side-length 1 has length $\varphi = \frac{1+\sqrt{5}}{2}$, which is a root of $x^2 - x - 1$. In fact, this is the *minimal polynomial* of φ, meaning it is the polynomial with rational coefficients of lowest degree for which φ is a root. We can extend this process further, seeing that the short diagonal of a hexagon has length $\sqrt{3}$, which has minimal polynomial $x^2 - 3$. Unfortunately, after this we aren't guaranteed that the length can be written without the use of transcendental functions, sin in particular. Luckily, however, we can still compute the minimal polynomials. In general, the short-diagonal-length of an n-gon with side length 1 is

$$d_n = 2\sin\left(\frac{(n-2)\pi}{2n}\right)$$

which can be derived from elementary geometry and trigonometry. I used WolframAlpha to compute the minimal polynomials up to $n = 18$, which can be seen in figure 3.7. While this minimal-polynomial approach, if it is the correct way, does mean that we aren't limited to prime n, it also makes it much harder to characterize when a plane will have an n-gon with n-fold symmetry. Above 6, the relevant roots either can't be expressed as radicals or are expressed with nested radicals. In either event, quadratic reciprocity can't be applied so directly. However, if some root is assumed, for fixed n, I suspect symmetry matrices can be produced in the way seen in Theorem 3.11 and Theorem 3.13, by solving a system of vector-matrix multiplications. However, this tactic probably does not effectively generalize to a polygon with an unknown number of sides.

[4]The short diagonal is the diagonal that connects two vertices that are one vertex apart. Pentagons only have short diagonals, but other n-gons have more types.

n	d_n if it can be expressed in terms of arithmetic and radicals.	The minimal polynomial of d_n
4	$\sqrt{2}$	$x^2 - 2$
5	$\frac{1}{2}(1+\sqrt{5})$	$x^2 - x - 1$
6	$\sqrt{3}$	$x^2 - 3$
7		$x^3 - x^2 - 2x + 1$
8	$\sqrt{2+\sqrt{2}}$	$x^4 - 4x^2 + 2$
9		$x^3 - 3x - 1$
10	$\sqrt{\frac{1}{2}(5+\sqrt{5})}$	$x^4 - 5x^2 + 5$
11		$x^5 - x^4 - 4x^3 + 3x^2 + 3x - 1$
12	$\sqrt{2+\sqrt{3}}$	$x^4 - 4x^2 + 1$
13		$x^6 - x^5 - 5x^4 + 4x^3 + 6x^2 - 3x - 1$
14		$x^6 - 7x^4 + 14x^2 - 7$
15	$\frac{1}{4}\left(-1+\sqrt{5}+\sqrt{6(5+\sqrt{5})}\right)$	$x^4 + x^3 - 4x^2 - 4x + 1$
16	$\sqrt{2+\sqrt{2+\sqrt{2}}}$	$x^8 - 8x^6 + 20x^4 - 16x^2 + 2$
17		$x^8 - x^7 - 7x^6 + 6x^5 + 15x^4 - 10x^3 - 10x^2 + 4x + 1$
18		$x^6 - 6x^4 + 9x^2 - 3$

Figure 3.7: Minimal polynomials of the length of short diagonals of real n-gons with side length 1.

This sort of problem for composite n is much more unclear. When we can no longer guarantee that certain terms in our equations are coprime, things get more difficult to prove. The fact that if something has a 5-cycle it can't possibly have any smaller cycles besides a 1-cycle is used so often that it became implicit. D_n also has more than the 4 possible types of subgroups we are able to limit our scope to when n is prime. However, none of these issues are irresovable, they simply require more number-theoretic care to solve.

CHAPTER 4

The Pentagram Map

4.1 Introduction to the Pentagram Map

The pentagram map was brought into prominence by Richard Shwartz [12] and was studied by him across a series of papers. The visual intuition is very clear; it maps any pentagon to the smaller pentagon at the center of the pentagram formed by the pentagon's diagonals (see figure 4.1). The pentagram map is usually represented with T.

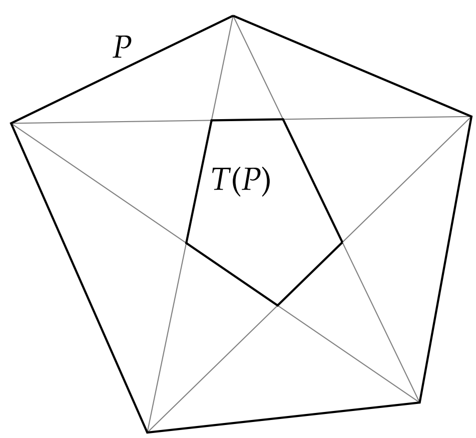

Figure 4.1: The relationship between a pentagon P and its image under the pentagram map $T(P)$.

Noting that each corner of the inner pentagon is the intersection of two consecutive diagonals, we can use this definition:

Definition 4.1 *Consider a pentagon P in general position represented by* $(p_0, p_1, p_2, p_3, p_4)$*. Let* D_n *be the diagonal of P given by* $[p_n \cdot p_{n+2}]$*, with subscript addition done mod 5. Then the* ***pentagram map*** *T is the map which takes P to* $(D_0 \cap D_1, D_1 \cap D_2, D_2 \cap D_3, D_3 \cap D_4, D_4 \cap D_0)$*.*

This definition can be extended to more general n-gons by letting the D_i be the shortest diagonals of that polygon. This map has several interesting properties. One highly relevant aspect to this discussion is that T is the identity map on the space of equivalence classes of pentagons, called the *moduli space*. In other words, $[T(P)] = [P]$. More generally, T^2, the pentagram map iterated twice, is the identity on the moduli space of hexagons. No similar identity property holds for n-gons with $n > 6$ [12]. This fact makes the pentagram map interesting to us in our study of pentagon equivalence classes.

Theorem 4.1.

(Conway, via Shwartz) For any pentagon P on π_p*, P is similar to* $T(P)$*. [12]*

The main reason for Shwartz's study of the pentagram map concern its properties as a dynamical system. Much of this work is beyond the scope of this paper and frankly beyond my current level of understanding, but will still be summarized to the best of my ability. The work of Shwartz, Serge Tabachnikov, and Fedor Soloviev eventually showed that the pentagram map is a completely integrable system on the moduli space of polygons. This means, roughly, that iteration of the pentagram map preserves enough invariants for the system to be in some sense 'well behaved.' For reference, the three body problem is *not* an integrable system. In this case, it means that a given polygon P stays on a certain higher-dimensional torus no matter how many times the pentagram map is applied to it, and that these torii fill the moduli space in layers, or *foilate* the space [13].

4.2 Periodicity of the Pentagram Map

In the world of discrete and finite geometry, the pentagram map takes on different properties. One such property is that, since the number of possible pentagons on any π_p is finite, and T is invertible, it is always periodic:

Theorem 4.2.

For any pentagon in general position P on π_p, there is some $n \in \mathbb{N}$ such that $T^n(P) = P$.

Proof. First we must show that T is invertible. Let the sides of a pentagon $P = (p_0, p_1, p_2, p_3, p_4)$ be $S_n = [p_n \cdot p_{n+1}]$, with subscript addition done mod 5. Let E be the map with takes P to $(S_0 \cap S_2, S_1 \cap S_3, S_2 \cap S_4, S_3 \cap S_0, S_4 \cap S_2)$. So long as P is in general position, we know all S_n are distinct, since otherwise that would mean there would be three colinear points in P. Thus, this map is well defined. Now, consider the geometric relationship between P and $E(P)$, shown in figure 4.2.

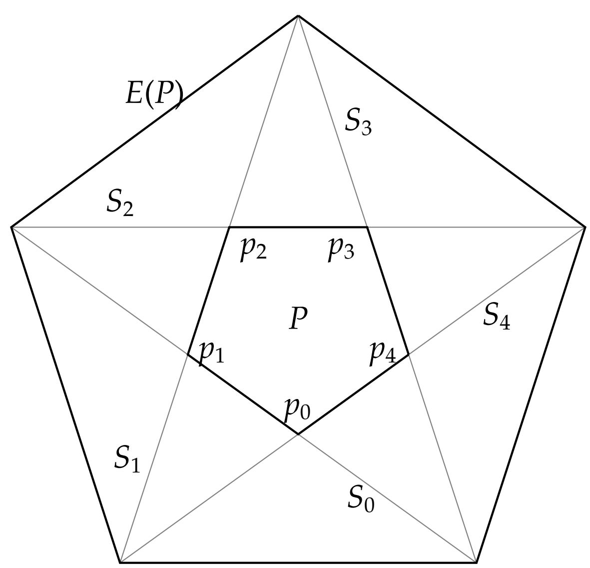

Figure 4.2: The relationship between a pentagon P and $E(P)$.

Because all S_n are distinct, these relationships are guaranteed to hold, even if diagrammed differently. As such, we can see that the sides of P are the diagonals of

$E(P)$; this means that $T(E(P)) = P$. Furthermore, if we let $W = E(P)$, we can see that $E(T(W)) = E(P) = W$. So, E is T^{-1}. Thus, T is invertible.

Now, choose some pentagon P and define the sequence $(x_n) = T^n(P)$. Now, consider that, since T is a function over a finite set of pentagons, there are only finitely many values x_i can take. Since (x_n) is an infinite sequence, this means there must be some $k \neq j$ such that $x_k = x_j$. Without loss of generality, let $k > j$. So, $T^k(P) = T^j(P)$. Since T is invertible, we can take $T^{-j}(T^k(P)) = T^{-j}(T^j(P))$, so $T^{k-j}(P) = P$. Since $k > j$, $k - j \in \mathbb{N}$.

For any pentagon P on π_p, there is some $n \in \mathbb{N}$ such that $T^n(P) = P$. □

This is in stark contrast to T's behavior on $\mathbb{RP}^2$. T does display quasiperiodicity[1] in the *moduli space* of polygons, it is absolutely not periodic on the space of polygons themselves; it is proven (and relatively intuitive) that repeated applications of the pentagram map cause the resulting sequence of pentagons to collapse exponentially in size [12]. However, in finite projective planes, on which the concept of scale is incoherent, this is not an issue. It is for this reason that our discussion our exploration will go in a very different direction that Shwartz's. While he is concerned with T's effect on moduli spaces, T's behavior is trivial on the moduli space of pentagons, meaning that we will instead concern T's behavior on individual pentagons *within* a given equivalence class. This will still give us some account of the differences between equivalence classes, however, because T acts on similar pentagons in analogus ways for the following reason:

Theorem 4.3.

(Shwartz) For any collineation α and polygon P, $\alpha T(P) = T(\alpha P)$. [12]

Proof. Since T is defined entirely in terms of incidence relations, and collineations

[1] Quasiperiodicity is the property that a system is the sum of two or more frequencies which are in irrational ratio to each other. This means that, while it isn't periodic, it is as close to periodic as you want it to be if you choose a large enough 'period.'

preserve incidence relations, for any collineation α, $\alpha T(P) = T(\alpha P)$. □

This theorem implies that classes of similar pentagons all have the same period under the iterated pentagram map. In this chapter our main focus will be on this periodic behavior. However, this is surprisingly complicated. As seen in Appendix B, figure B.1, different classes of similar pentagons can have very different periods.

Theorem 4.1 shows that a pentagon and its pentagram map image are similar, so there is some collineation ϕ such that $\phi P = T(P)$. However, we can prove a stronger result, that it's the same collineation for that pentagon's entire orbit when iterating the pentagram map. This means the period of $T^n(P)$ is just the order of ϕ in $\mathrm{PGL}(3, p)$. A proof of this result was also given by Shwartz, but we will demonstrate it here [12].

Theorem 4.4.

For any pentagon P, there exists some collineation ϕ such that $T^n(P) = \phi^n P$ for $n \in \mathbb{Z}^+$.

Proof. Consider a pentagon P on some π_p. Let ϕ be the collineation that takes P to $T(P)$, which must exist by Theorem 4.1. We will prodceed by induction.

Let S_k be the proposition that $T^k(P) = \phi^k P$. Clearly S_0 and S_1 are true, as $T^0(P) = P = \phi^0 P$ trivially and $\phi P = T(P)$ by construction.

Now, assume S_i is true for some $i > 1$. So, $T^i(P) = \phi^i P$. Multiply both sides by ϕ to get $\phi T^i(P) = \phi^{i+1} P$. By theorem Theorem 4.4, we know ϕ and T commute, so $\phi T^i(P) = T^i(\phi P) = \phi^{i+1} P$. Since $\phi P = T(P)$, $T^i(\phi P) = T^i(T(P)) = T^{i+1}(P) = \phi^{i+1} P$. So, S_{i+1} is true. Thus, $S_i \implies S_{i+1}$.

So, by induction, for all $n \in \mathbb{Z}^+$, $T^n(P) = \phi^n P$.

For any pentagon P, there exists some collineation ϕ such that $T^n(P) = \phi^n P$ for $n \in \mathbb{Z}^+$. □

One might think that this implies the period of $T^n(P)$ must be the order of

$\phi \in \mathrm{PGL}(3,p)$. Unfortunately, this is not true. It is possible that there is some k less than the order of ϕ such that ϕ^k is a symmetry of P, meaning $T^k(P) = \phi^k P = P$. However, this does imply that the period of $T^n(P)$ divides the order of ϕ.

As such, this still gives us important information. For example, we know that the period of $T^n(P)$ divides $p^3(p^3-1)(p^2-1)$, the order of $\mathrm{PGL}(3,p)$. What's more, if we can give ϕ as a matrix for some pentagon, then we can algebraically work out the period of the pentagram map on that pentagon. Thankfully, we can work out this matrix for most unitary regular pentagons. To help simplify discussion, we will give this collineation a name:

Definition 4.2 *A* ***pentagram collineation*** τ_P *of a pentagon P is a collineation such that $T(P) = \tau_P P$.*

The reason we say 'a' rather than 'the' is because pentagon symmetries mean that pentagram collineations are not unique.

Lemma 4.1.

Let $p > 5$ be prime, and U a regular unitary pentagon on π_p. Then, a pentagram collineation of U is given by:

$$\begin{pmatrix} g+1 & -1 & 1-g \\ 0 & g & 1 \\ g+1 & -1 & 1 \end{pmatrix}$$

where $g \in \mathbb{F}_p$ such that $g^2 - g - 1 = 0$.

Proof. Let $p > 5$ be prime, and U a regular unitary pentagon on π_p.

By the construction we used as part of Theorem 3.9, U must be of the form $(u_1, u_2, u_3, u_4, \langle 1, -g, g+1 \rangle)$, for some $g \in \mathbb{F}_p$ such that $g^2 - g - 1 = 0$. We will let $v = \langle 1, -g, g+1 \rangle$ for brevity.

As we discussed in subsection 2.4.1, intersection and finding colinear points are simply the cross product[2]. As such, we can compute the diagonals of U using the cross product:

$$D_0 = u_2 \times v = \langle -g-1, 0, 1\rangle$$
$$D_1 = u_1 \times u_3 = \langle 0, 1, 0\rangle$$
$$D_2 = u_2 \times u_4 = \langle 1, 0, -1\rangle$$
$$D_3 = u_1 \times v = \langle -g, -1, 0\rangle$$
$$D_4 = u_1 \times v = \langle 0, 1, -1\rangle$$

And then we compute their intersections using the cross product again:

$$D_4 \times D_0 = \langle 1, g+1, g+1\rangle$$
$$D_0 \times D_1 = \langle 1, 0, g+1\rangle$$
$$D_1 \times D_2 = \langle 1, 0, 1\rangle$$
$$D_2 \times D_3 = \langle 1, -g, 1\rangle$$
$$D_3 \times D_4 = \langle 1, -g, -g\rangle$$

So, $T(U) = (\langle 1, g+1, g+1\rangle, \langle 1, 0, g+1\rangle, \langle 1, 0, 1\rangle, \langle 1, -g, 1\rangle, \langle 1, -g, -g\rangle)$.

Now, let τ be the matrix given by:

$$\tau = \begin{pmatrix} g+1 & -1 & 1-g \\ 0 & g & 1 \\ g+1 & -1 & 1 \end{pmatrix}$$

It's a matter of direct computation to show that $\tau U = T(U)$. Thus, τ is a pentagram

[2]More specifically, we are using the cross product $\vec{w} \times \vec{z} = \langle w_2 z_3 - w_3 z_2, w_3 z_1 - w_1 z_3, w_1 z_2 - w_2 z_1\rangle$. This gives us the unique vector (up to scaling) that is orthogonal to $\vec{w}$ and $\vec{z}$, which is what we need.

collineation of U. □

We can use this to relate the period of regular pentagons P to the order of a much simpler object: an element from a finite field.

Lemma 4.2.

For some p, all $a \in \mathbb{F}_p$ such that $a^2 + 3a + 1 = 0$ have the same order in $\mathbb{F}_p$'s multiplicative group.

Proof. Let $a_1, a_2 \in \mathbb{F}_p$ be roots of $a^2 + 3a + 1$. If $a_1 = a_2$, then clearly they have the same order. So, consider the case where $a_1 \neq a_2$.

Clearly, $(x - a_1)(x - a_2) = x^2 + 3x + 1$. This means $x^2 - (a_1 + a_2)x + a_1a_2 = x^2 + 3x + 1$. In particular, $a_1a_2 = 1$. So, $a_1 = a_2^{-1}$. Since they are multiplicative inverses, they must have the same order in $\mathbb{F}_p$'s multiplicative group.

Thus, all roots of $a^2 + 3a + 1$ have the same order in $\mathbb{F}_p$'s multiplicative group. □

We have shown it is well-defined to talk about $a \in \mathbb{F}_p$ having a unique order when $a^2 + 3a + 1 = 0$. With this, we can make a conjecture.

Conjecture 4.1.

The period of the iterated pentagram map on a regular pentagon on π_p with $p > 5$ is the same as the order of a in $\mathbb{F}_p$'s multiplicative group, where $a^2 + 3a + 1 = 0$.

I have not been able to prove this statement as of yet. I can, however, discuss why I believe it to be true, and the difficulties I have had in proving it.

First of all, in Lemma 4.1, we showed that the pentagram collineation τ_R of any regular pentagon R could be written in terms of g, where $g^2 - g - 1 = 0$. Now, if we let $a = -1 - g$, we see that $(-1 - a)^2 - (-1 - a) - 1 = a^2 + 3a + 1 = 0$. So, τ_R can be written in terms of a as well. This algebraic relationship makes the conjecture seem more plausible; if τ_R could not be written in terms of a, it would be hard to see the connection between the two objects.

Given this, the main reason is computational. I used a script to check that this conjecture holds on π_p for all primes $p < 10,000$. If these quantites are independent, this initial of correspondence is particularly surprising because the potential period of $T^n(R)$ is significantly larger than a. As seen in figure B.1, the maximum period of the iterated pentagram map on some pentagon appears to be on the order of $O(p^2)$. Meanwhile, a's order must be less than $p-1$, as it is a member of $\mathbb{F}_p^\times$. So, the 'probability,' loosely speaking, of continued equality between these values drops as p gets large.

As we discussed, computing the order of τ_R is not sufficient to computing the period of the iterated pentagram map; however, even this first step proves difficult. The matrix we use to represent τ_R seems to resist analysis. It is not, in general, diagonalizable. 1 is always a root of its characteristic polynomial, however, experimentation seems to suggest it only has two other roots when $p \equiv 1 \bmod 5$. In fact, brute force attempts with SymPy to solve for $\tau_R = BTB^{-1}$, with T being an upper triangular matrix with powers of a on its diagonal, yielded no results. Solving this problem seems to either require more advanced algebraic concepts or a different approach entirely that bypasses the matrix.

There is a final issue of a slightly different type. While I believe this to be true by virtue of the computed data, I have no clear idea why it might be true. This conjecture is based on a pattern that I discovered wholly by accident when experimenting with scripts. As such, my intuitive understanding, and thus my ability to approach the problem in novel ways, is stymied. This intellectual problem is in part because this is a case where the behavior on finite planes is qualitatively different from on the real plane, meaning geometric analogies can only extend so far.

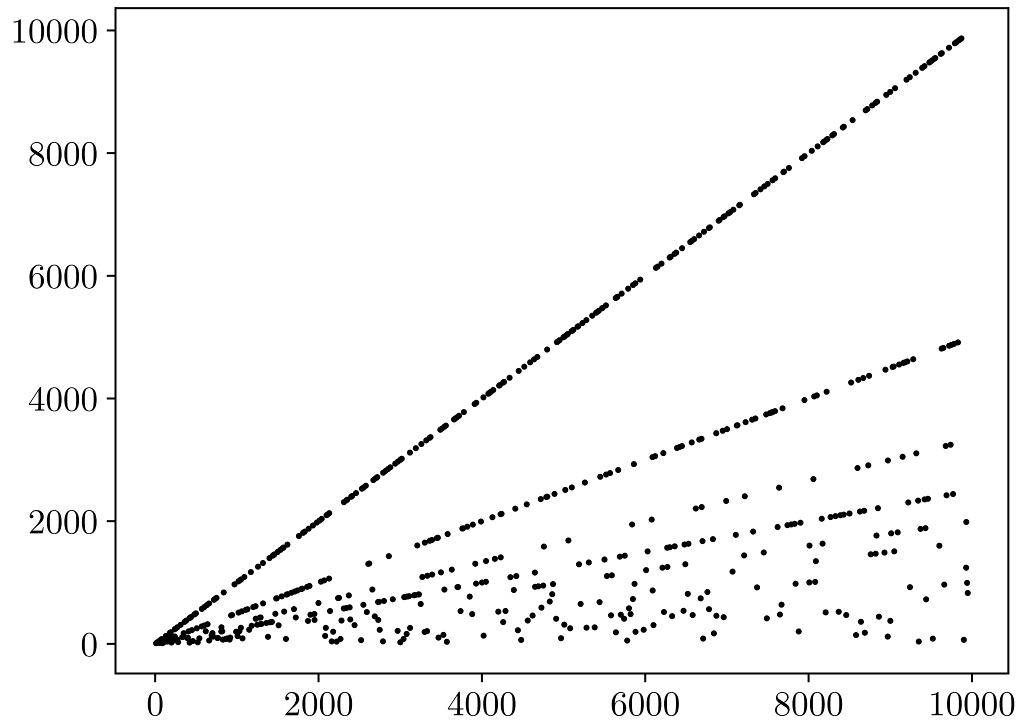

Figure 4.3: The primes p versus the order of a in $\mathbb{F}_p^\times$.

Actually determining the order of a for any field is en even more difficult task, and was not something I ever expected to be able to solve for this project. Of course, we know for sure that $|a|$ divides $p-1$, since that is the order of $\mathbb{F}_p$'s multiplicative group. However, beyond that the question becomes difficult to answer in general. Because $a^2 = -3a - 1$, we can show that $a^n = b_n a - b_{n-1}$, with $b_0 = 0$, $b_1 = 1$, and $b_{n+1} = -3b_n - b_{n-1}$. However, this means that finding the order is roughly equivalent to computing the period modulo p of $(b_n)_{n\in\mathbb{N}}$. However, this is the generalized Pisano period of $(b_n)_{n\in\mathbb{N}}$. A closed form for this value is not even known for the original Pisano periods of the Fibbonacci numbers [8], after which it is named[3]. Because of this, I consider computing this period beyond the scope of this paper.

[3]The man often called Fibonacci was known as Leonardo Bigollo Pisano during his lifetime.

Barring some number-theoretic technique which works only in this case, this is the best I can hope to achieve from this approach.

CHAPTER 5

Computation

A large portion of the results in this project were first encountered or evidenced by computer exploration. In fact, the very first proof of concept that this line of examination had nontrivial results was a Python script designed to check the periodicity of the pentagram map on the pentagons of a various finite projective planes. This section will explore the various scripts used, including their development and optimization as my theoretical understanding increased.

5.1 Representing Finite Planes

5.1.1 The Original Implementation

My first implementation of finite projective planes was taken from the internet [10]; this script generated the plane with the same process used in Section 2.2. Points were single objects, either tuples or singles, and lines were sets of points. While this sufficed for a preliminary exploration, there were obvious inefficiencies. Computing the line shared by two points involved iterating over the set of lines and checking which one both points sat on. This meant that this operation on a field of order p involved iterating over a list of $O(p^2)$ elements and iterating over a list of $O(p)$ elements for each of those, meaning it was a $O(p^3)$ operation relative to the order of the plane. Finding the intersection of two lines was better, but still involved finding the shared elements of two unsorted lists of $O(p)$ size. There may be a better runtime, but my niave implementation was naturally $O(p^2)$. Eventually, I mitigated these issues by wrapping that previous implementation in a `FiniteProjectivePlane` class which generated the original representation, used it to precompute line intersections and shared lines, and stored them in a matrix. This made the operation of finding these things constant time.

What could not so easily be done in this representation was collineation. Given two quadrangles A and B, I needed to compute a collineation that takes A to B pairwise. In this representation, the best I could do was start with a dictionary containing the points of A as keys and the points of B as values. From here, I would let the computer iteratively 'deduce' more and more points. This was inefficient, especially since I was at this point aware that I could represent the plane with $\mathbb{F}_p^3$ and use matrices to represent collineation. Because collineations were becoming more and more important for me to study, I eventually had to reimplement everything using finite fields.

5.1.2 Finite Fields

After an admittedly cursory search, I could find no Python package which implemented finite fields in a way that suited my needs. As such, I decided to implement my own, which I called `FFelem`. Since I only planned to explore fields of prime order, I did not have to worry about reducing modulo a polynomial as with fields of higher prime power order. As such, each FFelem need only be instantiated with a field order p and a numerical value a, the second of which is reduced modulo p before being stored. This step in the constructor meant I didn't have to worry about reducing modulo p in each of the operations. The `FFelem` class also has an internal variable for its inverse, so this didn't need to be computed more than once[1]. The inverse is computed using Fermat's Little Theorem, specifically by computing $a^{p-2} \bmod p$. Positive powers are computed by repeated multiplication, and negative powers are computed using the positive power of the inverse. Division is simply computed as multiplication by the inverse. All of these operations are overrides of the normal Python operators `+,-,*,/,**`, so `FFelem`s can be treated like regular integers. In addition, `FFelem`s can be used in operations with normal integers, where those integers will be treated like `FFelem`s of the same order. The implementation of this class can be seen in section C.1 of Appendix C. Finally, there was a simple function `get_field(n)` to get all elements of the field of order n. This had the advantage that is you stored the field in a list, with `f = get_field(n)`, then `f[5]`, for example, would get you the element 5 from $\mathbb{F}_n$.

Points and lines were represented by a single object called `ProjObj`. `ProjObj` essentially wraps a NumPy 3-vector of `FFelems` and normalizes it so that the first nonzero term of the vector is 1. This wrapper also overrides the bitwise xor operator, `^`, so that it returns the shared point between lines or the shared line between points, using the cross product. The bitwise or operator, `|`, is used to test for incidence using the dot product. There is also a transform method which returns the product of the vector and a 3×3 matrix of `FFelems`; however, this should only be used on representations of points. In our representation, for a collineation whose action on the points is given by a matrix A, its action on the lines is given by the inverse of the transpose of A, $(A^T)^{-1}$. This is due to the fact that we swapped out the 2-dimensional subspaces that represented our lines for their 1-dimensional tangent space. This implementation is given in section C.2. Finally, there was a function `get_plane(n)`

[1]However, given how `FFelem` was implemented, this probably was of minimal importance much of the time, as the addition, multiplication, and exponentiation operations create new instances of `FFelem` objects.

which returned all the points on π_n by simply iterating over all points of the three possible forms given how they are normalized: $(0,0,1)$, $(0,1,a)$, and $(1,a,b)$.

The advantages of finite field representation were reaped with the collineation representation. Unfortunately, while `numpy` could multiply matrices of FFelems, it was not general enough to solve linear systems or do matrix inversion. As such, I had to reimplement Gauss-Jordan elimination myself. Admittedly, all that was required was a relatively 'lazy' implementation which only needed to work on $3 \times n$ matrices; if it was given a 3×4 matrix, it returned the last column as a solution, and if it was a 3×6 matrix, it returned the rightmost 3×3 matrix, as I only used this to find inverses. The `base_to_quad( quad )` function generated an automorphism which took the unit quadrangle (see section 3.3) to the input quadrange, $Q = [\vec{a}, \vec{b}, \vec{c}, \vec{d}]$. First it generates a matrix m that takes the first three vertices of the unit quadrangle to the first three vertices of Q. To get the fourth vertices to correspond, we need to see how to scale each of the columns relative to each other to get the desired vector d. This turns out to be the solution $\vec{w}$ to the equation $m\vec{w} = \vec{d}$. Using the Gauss-Jordan function we wrote, we can solve this equation and multiply each column of m with its corresponding value from $\vec{v}$ to get the solution, M. From here, the `gen_automorphism( quad1, quad2 )` function, which takes an arbitary quad to another, was relatively easy to implement. It simply calls `base_to_quad` on both of our inputs Q_1 and Q_2, to get two matrices M_1 and M_2. Since M_1 takes the unit quadrangle to Q_1, M_1^{-1} takes Q_1 to the unit quadrangle. Thus, $M_2M_1^{-1}$ takes Q_1 to the unit quadrangle to Q_2, and is our solution. The implementation of these functions is found in Appendix C, section C.3.

5.2 Iterating Over the Pentagons

When computing periods periods of the pentagram map and equivalence classes of pentagons under collineation, one needs a way to iterate over the pentagons in general position of a given plane π_p. At first, I took the naive approach, iterating over every permutation of 5 points on the plane, and testing those in general position. Obviously, this had many inefficiencies; first of all, it the same pentagon is processed 10 times, once for every possible orientation. This could be fixed by requiring the first element in the permutation be the first in lexical order, locking it into a certain rotation, and that the second element comes before the last in lexical order, fixing it into a certain orientation. However, since most permutations of 5 will be in general position, this still meant checking $O({}_{p^2+p+1}C_5) = O(p^{10})$ possibilities.

Once I began to understand collineations, particularly the fact that they are defined by where they take a quadrangle, I was able to significantly increase the efficiency of iteration. Since I only ever needed one representative from each similarity class, I could fix 4 of the 5 vertices, $\langle 1, 0, 0\rangle$, $\langle 0, 1, 0\rangle$, $\langle 0, 0, 1\rangle$, $\langle 1, 1, 1\rangle$, and only ever iterate over the 5th. This is the origin of my deployment of unitary pentagons (see section 3.3) This made searching an entire plane $O(p^2)$, as that is the order of the number of points.

However, there were still multiple unitary pentagons from each equivalence class; for a similarity class with no non-identity symmetries, there would be 10, for example. The elimination of these redundancy wasn't an optimization strategy, however, as there were still $O(p^2)$ classes. This was for the sake of getting accurate counts. It was done by iterating over all possibilities, but applying every possible mapping of the 4 unitary points of the pentagon to another 4 points, and collecting them in a list. If a pentagon was already in one of these lists, it was skipped, and at the end, one representative from each list was taken.

5.3 Other Scripts

This is not a full list of every script I wrote while working on this project. However, the tools and methods discussed formed a basis of many smaller scripts, of which I've written over 20. They primarily serve one of two purposes: the gathering of examples, and checking intuitions. For an example of the former, I computed the rotational symmetry collineation of each each regular unitary pentagon on each plane, and used those, along with some algebra, to work out the general form used in Theorem 3.10. For the latter, it is harder to show specific fruits. It was more a 'sanity check' than anything else, making sure things that I was more or less certain were true held up for real examples. Most of the results in this paper were tested this way long before I ever proved them.

CHAPTER 6

Conclusion

The pentagram map was what originally inspired this project, even though my focus on it in the end was much smaller than on pentagon symmetries and similarity-class counting. Having chanced upon its Wikipedia page one day on a family outing, I was eventually inspired to apply the process to finite projective planes, another concept that I had discovered when trawling Wikipedia. The code I wrote hastily showed me something that surprised me at the time: different starting pentagons yielded wildly different periods. This was surprising because, with my limited understanding of projective planes, I saw them as spaces so symmetrical that I found it strange any pentagon was differentiable from any other. The resulting desire to gain a handle on *how* pentagons could be different in these spaces resulted in an effort to classify and count them which is what inspired much of chapter 3.

Over the course of this project I have at least started this process. I have found ways of counting classes on pentagons by their symmetries. In doing so, I've given parametrized forms for representatives of each of these classes, shedding some light on the structure of individual pentagons themselves. In doing so, I've also shown a fundemental connection between a finite field's equivalents of the golden ratio and pentagon regularity. This demonstrates that the appearance of ϕ on the short edges of regular pentagons in $\mathbb{R}^2$ is not a fluke within the world of field planes, it has a fundemental connection to pentagon regularity that transcends the need for a metric; regular pentagons can only exist if $x^2 - x - 1$ has a root.

This line of exploration is far from complete, of course. Obviously, I've only dealt with pentagons, and have almost entirely ignored n-gons of larger size. But even for pentagons it is not as if the classification of their properties is complete. I demonstrated in Theorem 3.15 that there were $\frac{1}{10}((p+3)(p-3)+4r_p)$ different classes of similar pentagons on π_p, however, I've only been able to categorize these classes by a single property with three possibilities: 1 symmetry, 2 symmetries, or 10 symmetries. The pentagram map demonstrates clearly that more classification is possible; when classified by pentagram-map period, the number of classes that are qualitatively different from each other increases as the size of the plane increases, as seen in figure B.1.

Of course, I was not expecting to complete this incredibly broad line of inquiry in this IS, and I am satisfied with what I have accomplished. The pentagram map inspired this process, and even though what followed was often very different, I'm glad I was able to explore what it inspired me to do, and then connect it back to the map in the end.

APPENDIX A

Proof Addenda

A.1 A regular pentagon on π_5

The pentagon on π_5 given by $R = (\langle 1,2,4\rangle, u_1, u_2, u_3, u_4)$ is regular, with rotation symmetry given by:

$$\rho = \begin{pmatrix} 4 & 1 & 0 \\ 3 & 0 & 2 \\ 1 & 0 & 0 \end{pmatrix}$$

It can be checked by direct computation that $\rho R_0 = R_1$, $\rho R_1 = R_2$, $\rho R_2 = R_3$, etc.

Because of the contraints put on the fifth vertex of a unitary pentagon in Theorem 3.2, we know the only other unitary pentagons are those with $\langle 1,2,3\rangle, \langle 1,3,2\rangle, \langle 1,4,2\rangle, \langle 1,4,3\rangle$ or $\langle 1,3,4\rangle$ as their fifth vertex. Call them P_1, P_2, P_3, P_4, P_5 respectively.

You can check directly that the following are true, meaning the remaining 5 pentagons are all similar to each other:

$$\begin{pmatrix} 0 & 0 & 1 \\ 2 & 0 & 1 \\ 0 & 1 & 1 \end{pmatrix} P_1 = P_2$$

$$\begin{pmatrix} 0 & 0 & 1 \\ 3 & 0 & 1 \\ 0 & 1 & 1 \end{pmatrix} P_2 = P_3$$

$$\begin{pmatrix} 0 & 0 & 1 \\ 3 & 0 & 1 \\ 0 & 2 & 1 \end{pmatrix} P_3 = P_4$$

$$\begin{pmatrix} 0 & 0 & 1 \\ 2 & 0 & 1 \\ 0 & 3 & 1 \end{pmatrix} P_4 = P_5$$

$$\begin{pmatrix} 0 & 0 & 1 \\ 1 & 0 & 1 \\ 0 & 2 & 1 \end{pmatrix} P_5 = P_1$$

Since they are all similar, that means that there are 5 unitary pentagons in their equivalence class, so by Theorem 3.6, they have $\frac{10}{5} = 2 < 10$ symmetries, meaning they aren't regular. So, R is the only regular unitary pentagon on π_5. Thus, there is only one equivalence class of regular pentagons on π_5.

A.2 Alternative argument that regular pentagons exist on every plane where $p \equiv 0, \pm 1 \mod 5$

The following is an argument cut from Theorem 3.10 due to its redundancy. It existed when the thereom needed to be if and only if, rather than just only if. It is preserved here because it was the theorem that in part originated the matrix techniques that I would use regularly elsewhere in the paper.

Now, consider any projective plane π_p with prime order $p \equiv 0 \mod 5$ or $p \equiv \pm 1 \mod 5$. Note that, since p is prime, its collineation group is $\mathrm{PGL}(3, p)$.

In the first case, since p is prime, then the only possibility is that $p = 5$. You can check directly that π_5 has a regular pentagon. An example is given in Appendix A, section A.1.

Otherwise, $p \equiv \pm 1 \mod 5$. So, by Lemma 3.1, we have $g \in \mathbb{F}_p$ such that $g^2 - g - 1 = 0$. Note that this is the polynomial associated with the golden ratio[1]. It is trivial to see that $g^2 = g + 1$ and $g^2 - g = 1$. We will use these in simplifications.

Consider the matrix:

$$\rho = \begin{pmatrix} 1 & g-2 & 1-g \\ 1 & g-1 & 0 \\ 1 & -1 & 0 \end{pmatrix}$$

We can compute ρ^2 as the following, and simplify using the relations we deduced:

$$\rho^2 = \begin{pmatrix} 0 & g^2-g-1 & 1-g \\ g & g^2-g-1 & 1-g \\ 0 & -1 & 1-g \end{pmatrix} = \begin{pmatrix} 0 & 0 & 1-g \\ g & 0 & 1-g \\ 0 & -1 & 1-g \end{pmatrix}$$

Then, we can compute ρ^4 as $(\rho^2)^2$.

$$\rho^4 = \begin{pmatrix} 0 & g-1 & g^2-2g+1 \\ 0 & g-1 & 1-g \\ -g & g-1 & g^2-g \end{pmatrix} =$$

[1] It is interesting to consider the fact that regular pentagons in the real plane are also associated with the golden ratio; thus, it seems fitting (and is probably mathematically meaningful) that regular pentagons only exist in π_p when its underlying field $\mathbb{F}_p$ admits a golden ratio equivalent.

Finally, we compute ρ^5 as $\rho^4\rho$:

$$\rho^5 = \begin{pmatrix} 1 & g^2-g-1 & 0 \\ 0 & g^2-g & 0 \\ 0 & 0 & g^2-g \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

So, $\rho^5 = e$, the identity matrix, and thus represents a collineation of order 5 in the collineation group of π_p. Now, with $P_0 = [1,0,0]$, one can compute $P_1 = \rho P_0 = [1,1,1]$ and $P_2 = \rho^2 P_0 = [0,g,0] = [0,1,0]$. Note that:

$$[P_0 \cdot P_1] = [0,1,-1] \neq [1,0,-1] = [P_1 \cdot P_2]$$

So, $[P_0 \cdot \rho P_0] \neq [\rho P_0 \cdot \rho^2 P_0] = \rho[P_0 \cdot \rho P_0]$. This means that there is some point P for which ρ does not fix the line $[P \cdot \rho P]$.

So, there is some point P on π_p and some collineation ρ of π_p such that ρ has order 5 and $[P \cdot \rho P]$ is not fixed by ρ.

Now, consider the pentagon $R = (P, \rho P, \rho^2 P, \rho^3 P, \rho^4 P)$. Since ρ doesn't fix $[P \cdot \rho P]$, this means that the sides of this pentagon are all distinct lines. Thus, they are in general position. Note that it has 5 rotational symmetries, as:

$$\rho R = (\rho P, \rho^2 P, \rho^3 P, \rho^4 P, \rho^5 P) = (\rho P, \rho^2 P, \rho^3 P, \rho^4 P, P) = R$$

This means that the set generated by ρ, $\langle\rho\rangle$, are all symmetries of R. However, since $|\langle\rho\rangle| = 5$, by Theorem 3.9, this means that R is regular. Thus, π_p contains a regular pentagon.

So, for prime p, π_p has a regular pentagon if and only if $p \mod 5$ is congruent to -1, 0, or 1.

A.3 Alternate argument that ρ doesn't fix $[P \cdot \rho P]$ when $p \equiv -1 \mod 5$

The following is an argument cut from theorem above due to its redundancy. It shows that ρ, a collineation of order 5 on π_p, cannot fix every line of the form $[P \cdot \rho P]$ when $p \equiv -1 \mod 5$. It is preserved here because it is a mostly geometric argument that the author believes provides an interesting intuitive insight into the questions at hand.

Case: $p \equiv -1 \mod 10$

Since $p \equiv -1 \mod 10$, thanks to a result by Bloom [3], we know that PSL(2, 5) is isomorphic to a subgroup of PSL(3, p). Furthermore, it is known that A_5 is isomorphic to PSL(2, 5), and $D_5 \leq D_{10}$, and D_{10} is isomorphic to a subgroup of A_5. Thus, D_5 is isomorphic to a subgroup of PSL(3, p) $\leq$ PGL(3, p). So, D_5 is isomorphic to a subgroup of PGL(3, p).

So, in any case, D_5 is isomorphic to a subgroup of the collineation group of π_p. Call one such group $\mathcal{S}$ and consider generators $\rho, \sigma \in \mathcal{S}$ with $\rho^5 = \sigma^2 = e$, the identity.

Assume that, for all points $P \in \pi_p$, the line $[P \cdot \rho P]$ is fixed by ρ. Note that the fixed points and lines must conform to the first two axioms of projective planes; thus, those fixed points and lines of ρ must form either a true subplane of π_p or a degenerate plane. It cannot be a true subplane, as p is prime, so π_p is a Desarguesian plane, meaning subplanes correspond to proper subfields of $\mathbb{F}_p$, which do not exist since p is prime [7].

Since the fixed points of ρ form a degenerate plane, they will be in one of two configurations [1].

The first configuration is lines $\mathbb{L}, L_1, L_2, \cdots L_n$ and points $\mathbb{P}, P_1, P_2, \cdots P_m$ for some nonnegative integers m, n. All points lie on $\mathbb{L}$ and all lines pass through $\mathbb{P}$.

Note that, by our assumption, every point in $P \in \pi_p$ lies on a fixed line, since

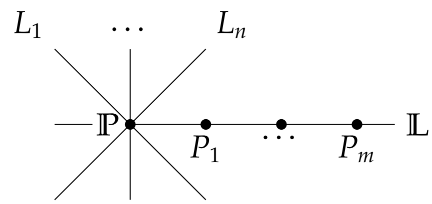

Figure A.1: The first possible configuration.

$[P \cdot \rho P]$ is fixed. Since there are $p^2 + p + 1$ points and there can be at most $p + 1$ points on $\mathbb{L}$, that implies there is at least one other fixed line, which we will say is L_1. Note that L_1 has only one fixed point $\mathbb{P}$, and so the rest of the points on L_1 must be permuted by ρ. However, this is impossible. Since L_1 has $(p + 1) - 1 = p$ unfixed points, and ρ has order 5, that must mean that 5 divides p. However, $p \equiv -1 \mod 5$, so this is impossible. Thus, if our assumption is true, that means that the fixed points of ρ must be in the other configuration.

The second configuration is lines $\mathbb{L}, L_1, L_2, \cdots L_n$ and points $\mathbb{P}, P_1, P_2, \cdots P_n$ for some nonnegative integer n. All points except $\mathbb{P}$ lie on $\mathbb{L}$ and all $L_n = [P_n \cdot \mathbb{P}]$.

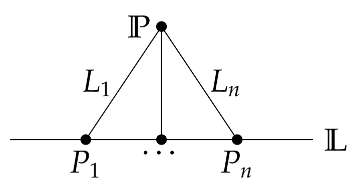

Figure A.2: The second possible configuration.

Once again, since every point lies on a fixed line and there are $p^2 + p + 1$ points and each line has $p + 1$ points, there must be at least one other fixed line, L_1, besides $\mathbb{L}$. There are two fixed points (P_1 and $\mathbb{P}$) on L_1, so ρ must permute the remaining $(p+1)-2 = p-1$ remaining points. Since none of these are fixed, and $\rho^5 = e$, $5|(p-1)$. So, $p \equiv 1 \mod 5$. However, we are considering the case where $p \equiv -1 \mod 5$, so this is a contradiction. Thus, our assumption is false, and there is some point P such that $[P \cdot \rho P]$ is not fixed by ρ.

APPENDIX B

Computed Data

The following is a table which contains the period of pentagon equivalence classes under the pentagram map. The first column contains the order of the plane, and the second contains each periodicity paired with the number of period with that period.

Figure B.1: Computed periods under the pentagram map.

p	Period	# of classes
5	3	1
	5	1
7	8	2
	12	2
11	10	2
	15	2
	55	2
	60	2
	133	4
13	6	1
	7	1
	12	2
	28	2
	84	4
	168	2
	183	4
17	8	3
	9	1
	16	6
	18	2
	36	2
	48	2
	96	2
	307	10
19	9	6
	18	2
	60	4
	120	4
	127	2
	171	2
	180	2
	342	2
	360	2
	381	10
23	11	2
	22	4
	44	2
	66	2
	88	4
	132	6
	176	4
	264	12
	528	6
	553	10

p	Period	# of classes
29	7	1
	14	5
	21	6
	28	4
	67	4
	70	4
	84	2
	140	10
	168	6
	203	2
	210	4
	280	4
	406	4
	420	8
	840	4
	871	16
31	15	12
	30	10
	32	2
	60	2
	64	4
	96	2
	120	6
	160	4
	192	6
	240	4
	320	10
	331	10
	465	2
	930	4
	993	18
37	9	5
	18	4
	19	1
	36	16
	38	4
	57	2
	76	4
	114	6
	152	12
	171	4
	342	14
	469	12
	684	14
	1368	8
	1407	30

p	Period	# of classes
41	14	2
	20	6
	21	10
	28	6
	30	2
	40	16
	56	2
	60	6
	112	8
	168	2
	280	4
	336	4
	410	2
	560	16
	820	6
	840	8
	1680	24
	1723	44
43	7	2
	14	8
	21	14
	42	16
	66	2
	77	2
	132	6
	154	8
	231	10
	308	12
	616	16
	631	10
	924	4
	1848	12
	1893	62
47	23	16
	46	28
	92	2
	184	10
	276	2
	368	4
	552	14
	736	16
	1104	30
	2208	22
	2257	76

APPENDIX C

Code

C.1 Finite field element class implementation.

```
class FFelem:

    def __init__(self, num, order):
        self.order = order
        self.num = num%order

        self._inv = None

    def __add__(self, other):

        if isinstance(other, FFelem):
            if self.order != other.order:
                raise ValueError("Field orders don't match")
            return FFelem( self.num + other.num , self.order)
        elif isinstance(other, int):
            return FFelem( self.num + other     , self.order)
        else:
            raise ValueError("Cannot add objects of this type")

    def __radd__(self, other):
        return self+other

    def __sub__(self, other):

        if isinstance(other, FFelem):
            if self.order != other.order:
                raise ValueError("Field orders don't match")
            return FFelem( self.num - other.num , self.order)
        elif isinstance(other, int):
            return FFelem( self.num - other     , self.order)
        else:
            raise ValueError("Cannot subtract objects of this type")
```

```
    def __mul__(self, other):

        if isinstance(other, FFelem):
            if self.order != other.order:
                raise ValueError("Field orders don't match")
            return FFelem( self.num * other.num , self.order)
        elif isinstance(other, int):
            return FFelem( self.num * other     , self.order)
        else:
            raise ValueError("Cannot multiply by objects of this type")

    def __rmul__(self, other):
        return self*other

    def __truediv__(self, other):

        if isinstance(other, FFelem):
            if self.order != other.order:
                raise ValueError("Field orders don't match")
            return self*other.inverse()
        elif isinstance(other, int):
            return self*FFelem(other, self.order).inverse()
        else:
            raise ValueError("Cannot divide by objects of this type")

    def __neg__(self):
        return FFelem( -self.num , self.order)

    def __pow__(self,power):
        if not isinstance(power, int):
            raise ValueError("Power not an integer")

        if power < 0:
            return (self.inverse())**(-power)

        return FFelem( self.num ** power, self.order )

    def inverse(self):
        if self.num == 0:
            raise ValueError("Cannot take inverse of zero element")

        if self._inv != None:
            return self._inv

        i = self**(self.order-2)
        self._inv = i
        return i

    def __eq__(self,other):
        if isinstance(other, FFelem):
            if self.order != other.order:
                return False
            return self.num == other.num
```

```
        if isinstance(other, int):
            return self.num == other

        return False

    def __repr__(self):
        return str(self)

    def __str__(self):
        return str(self.num)
```

C.2 Point/line element class implementation, and its helper function, homogenize.

```
from finitefield import get_field,FFelem
import numpy as np

def homogenize(pt):

    for i in range(3):
        if pt[i] != 0:
            return pt * pt[i]**-1

    raise ValueError("Zero vector is not homogenous")

class ProjObj:

    def __init__(self,a,b,c):
        self.vec = homogenize(np.array([a,b,c]))

    def __xor__(self,other):
        #xor (^) gives the intersection of two
        #lines or the line connecting two points
        if not isinstance(other,ProjObj):
            raise ValueError("Must be two ProjObj elements")
        return ProjObj( *np.cross(self.vec,other.vec) )

    def __or__(self,other):
        #or (|) tests incidence of point and line,
        #a, b (or line and point)
        if not isinstance(other,ProjObj):
            raise ValueError("Must be two ProjObj elements")
        return np.dot(self.vec,other.vec) == 0

    def transform(self,matrix): #matrix multiplication
        v = [matrix[:,i]*self.vec[i] for i in range(3)]
        v = v[0]+v[1]+v[2]
        return ProjObj(*v)

    def __eq__(self,other):
        return all(self.vec[i] == other.vec[i] for i in range(3))

    def __repr__(self):
        return str(self)

    def __str__(self):
        return str(self.vec)

    def __hash__(self): #returns number which allows lex sort
        o = self.vec[0].order
        a,b,c = [self.vec[k].num for k in range(3)]
```

```
        return a*o**2 + b*o + c
```

C.3 Implementation of collineation-related functions.

```
import numpy as np
from proj_plane import ProjObj,get_unit_quadrangle,get_plane
from finitefield import FFelem


#maps quad1 = [a b c d] to quad2 = [e f g h]
def gen_automorphism( quad1, quad2 ):
    t1 = base_to_quad(quad1)  #unit_quad -> quad1
    t2 = base_to_quad(quad2)  #unit_quad -> quad2
    return t2@mat_inverse(t1) #quad1 -> unit_quad -> quad2

def mat_inverse(mat):
    o = mat[0,0].order
    z = FFelem(0,o)
    n = FFelem(1,o)

    I = np.array([[n,z,z],[z,n,z],[z,z,n]])

    return gauss_jordan(np.concatenate((np.copy(mat),I),axis=1))

#maps unit_quad = [001 010 100 111] to quad = [a b c d]
def base_to_quad( quad ):

    a,b,c,d = [k.vec for k in quad]
    m = np.empty(shape=(3,4)).astype(FFelem)
    M = np.empty(shape=(3,3)).astype(FFelem)

    for i,v in enumerate([a,b,c,d]):
        for j in range(3):
            m[j,i] = v[j]
            if i < 3:
                M[j,i] = v[j]

    w = gauss_jordan(m)

    for i in range(3):
        M[:,i] *= w[i]

    return M

def gauss_jordan( mat ): #hard coded for 3x4 matrices

    if mat[0,0] == 0:
        if mat[1,0] != 0:
            mat[[0,1]] = mat[[1,0]]
        elif mat[2,0] != 0:
            mat[[0,2]] = mat[[2,0]]
```

```
        else:
            raise ValueError("Matrix not invertible")

    mat[0] *= mat[0,0]**-1
    mat[1] -= mat[0]*mat[1,0]
    mat[2] -= mat[0]*mat[2,0]

    if mat[1,1] == 0:
        if mat[2,1] != 0:
            mat[[1,2]] = mat[[2,1]]
        else:
            raise ValueError("Matrix not invertible")

    mat[1] *= mat[1,1]**-1
    mat[2] -= mat[1]*mat[2,1]

    mat[2] *= mat[2,2]**-1

    #mat is now upper triangular

    mat[0] -= mat[1]*mat[0,1]
    mat[0] -= mat[2]*mat[0,2]

    mat[1] -= mat[2]*mat[1,2]

    if mat.shape == (3,4):
        return np.array([mat[0,3],mat[1,3],mat[2,3]])
    if mat.shape == (3,6):
        return mat[:,3:]
```

References

1. A. Adrian Albert and Reuben Sandler. *An introduction to finite projective planes*. Dover Publications, Incorporated, 2015. 5, 8, 9, 10, 19, 75

2. Arnaud Beauville. Fundamental theorem of projective geometry, 2022. 20

3. David M. Bloom. The subgroups of psl(3, q) for odd q. *Transactions of the American Mathematical Society*, 127(1):150, 1967. doi: 10.2307/1994638. 75

4. R. H. Bruck and H. J. Ryser. The nonexistence of certain finite projective planes. *Canadian Journal of Mathematics*, 1(1):88–93, 1949. doi: 10.4153/CJM-1949-009-2. 14

5. John Horton Conway. *Atlas of finite groups: maximal subgroups and ordinary characters for simple groups*. Clarendon Press; Oxford University Press, 1985. ISBN 9780198531999,0-19-853199-0. 20, 39

6. H.S.M. Coxeter. *Projective Geometry*. Springer New York, 2003. ISBN 9780387406237. URL `https://books.google.com/books?id=LfTPdnYLQ-AC`. 21

7. D.S. Dummit and R.M. Foote. *Abstract Algebra*. Wiley, 2003. ISBN 9780471433347. URL `https://books.google.com/books?id=KJDBQgAACAAJ`. 75

8. Sergio Falcon and Ángel Plaza. k-fibonacci sequences modulo m. *Chaos, Solitons Fractals*, 41(1):497–504, 2009. ISSN 0960-0779. doi: https://doi.org/10.1016/j.chaos.2008.02.014. URL `https://www.sciencedirect.com/science/article/pii/S0960077908000726`. 60

9. Saunders Mac Lane Garrett Birkhoff. *A Survey of Modern Algebra*. Akp Classics. A K Peters/CRC Press, 5th edition, 2010. ISBN 1-56881-068-7,9781568810683,9781568814544,1568814542. 14

10. Joel Grus. Spot-it code. `https://github.com/joelgrus/spot-it`, 2013. 6, 63

11. Felix Lazebnik, Keith Mellinger, and Oscar Vega. On the number of k-gons in flnite projective planes. *Note di Matematica*, 29, 01 2009. doi: 10.1285/i15900932v29n1supplp135. 16, 17, 21, 47

12. Richard Schwartz. The pentagram map. *Experimental Mathematics*, 1(1):71 – 81, 1992. doi: em/1048709118. URL https://doi.org/. 21, 51, 52, 54, 55

13. Fedor Soloviev. Integrability of the pentagram map. *Duke Mathematical Journal*, 162(15), Dec 2013. ISSN 0012-7094. doi: 10.1215/00127094-2382228. URL http://dx.doi.org/10.1215/00127094-2382228. 52

14. Charles A. Weibel. Survey of non-desarguesian planes. 2007. 15