Conference Athletic Schedules: An Application of Projective Geometry, Finite Fields, and Graph Theory

A THESIS SUBMITTED TO THE FACULTY OF THE
UNIVERSITY OF MINNESOTA BY

Alexandra Nelson

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF MASTER OF SCIENCE

Adviser: Bert Fristedt

May 2016

Abstract

This paper solves a problem faced by the Suburban East Conference of the Minnesota State High School League in 2009, of designing a consistent wrestling schedule to accommodate a new school. We adapt the application of projective geometry of finite fields to general scheduling problems, and develop an algorithm for determining a schedule. We prove that this algorithm can be completed, will yield the desired schedule, and can yield all possible schedules in the desired format. We also model the schedule with bipartite graphs, and use edge colorings to complete the schedule with home and away assignments.

Contents

List of Figures

1 The Scheduling Problem

The problem to be solved in this paper was posed by Daniel Willaert, former graduate student at the University of Minnesota and assistant wrestling coach at Hastings High School in Hastings, Minnesota, and current head coach at Cretin-Derham Hall High School in St. Paul, Minnesota. The problem was one which arose in 2009 in response to changes in the Suburban East Conference of the Minnesota State High School League, and went unsolved.

Until spring 2009, the Suburban East Conference consisted of nine high schools: Cretin-Derham Hall, Forest Lake, Hastings, Mounds View, Park-Cottage Grove, Roseville, Stillwater, White Bear Lake, and Woodbury. The conference schedule for the nine wrestling teams of these schools was designed in Round Robin style, in which each team wrestled in eight dual meets, consisting of two teams matched up, through the season. However, coaches would often shuffle schedules and trade assignments to create triangular meets, consisting of three teams matched up, so as to eliminate the total number of conference meets in which their team competed. This would open up their schedules to either additional bye weeks to rest, or the scheduling of nonconference meets.

In the fall of 2009, East Ridge High School in Woodbury opened and joined the Suburban East Conference. By necessity, the Round Robin schedule had to be extended to a nine-week schedule from its previous eight-week schedule, which was even less desirable to coaches trying to work in byes or nonconference meets. In response, Willaert and the rest of the coaching staff at Hastings High School attempted to design a schedule for the now ten teams in the conference which would incorporate triangular meets from the beginning and thus condense the conference season. They proposed a design which would take six weeks, and in which each team would compete in a total of three triangular meets and three dual meets, each of the first five weeks would consist of two triangular meets and two dual meets, and the final sixth week would consist of five dual meets.

There were three huge advantages to this design. First, the final sixth week could be "powerized" in that the teams could purposely be assigned so that the previous year's first place and second place teams would compete head to head, the third and fourth place teams would compete head to head, and so on. This would create more exciting end-of-season meets for all schools. Second, not only was the number of dual and triangular meets in which each team competed consistent, but because the entire conference schedule was only six weeks, the teams were all on bye from each other at the same time, eliminating that advantage. Third, with each team competing in three triangular meets and two dual meets in the bulk of the season, plus the one dual meet in the final week, the assignments of home and away teams could be attempted in an even manner, so that each team hosted one triangular meet and one dual meet, and in the sixth week, the higher seed team could have the home advantage.

Though this proposed design would work, the task of arranging the ten teams into such a design so that each team played each of the other nine exactly once turned out to be a formidable task. The Hastings staff was unable to design such a schedule, and the project was shelved. Willaert designed instead a seven-week schedule that incorporated a bye week for each team but still kept the number of dual and triangular meets in which each team competed consistent. However, even this schedule was ultimately rejected by the conference, as some coaches preferred the autonomy of tweaking their own schedules over consistency.

The problem remained shelved until 2014, when Willaert and I were introduced in a mathematics class to the application of projective geometry and finite fields to scheduling problems. Willaert conjectured that this application could be modified to solve his wrestling problem. In fact, he did produce his desired schedule, using a much more structured trial-and-error process than his staff had managed in 2009. This paper, however, will focus on not a single solution to Willaert's problem, but the development of an algorithm to solve it. The algorithm presented will not only yield the desired schedule, but can yield a variety of possible schedules, depending on choices made as part of the process. The algorithm is also invariant under permutations of the ten teams, so that after the algorithm is run, the ten teams may be assigned so that the sixth week of the schedule is powerized as desired.

In fall 2014, Hastings High School left the Suburban East Conference to join the new Metro East Conference, leaving Suburban East with only nine schools. This solution to the original scheduling problem is therefore no longer relevant to the current situation. However, this algorithm could potentially be generalized to different numbers of teams, schedule sizes, and meet choices (i.e. working in bye weeks, or four-team meets, and so on).

2 Introduction to Projective Geometry

2.1 What Is Projective Geometry?

Projective geometry is a particular type of geometry determined by four axioms. These axioms, taken from the definition of nongenerative projective geometry of Beutelspacher and Rosenbaum [2], are the following:

Axiom 1. For any two distinct points P and Q, there is exactly one line incident with both P and Q; a name for that line is $\overleftrightarrow{PQ}$.

Axiom 2. For any distinct points A, B, C, and D: if $\overleftrightarrow{AB}$ meets $\overleftrightarrow{CD}$, then $\overleftrightarrow{AC}$ meets $\overleftrightarrow{BD}$.

Axiom 3. Every line is incident with at least three points.

Axiom 4. There are at least two lines.

Axiom 1 will be of especial use in the scheduling application presented here. This axiom implies that two different projective lines may intersect at no more than one projective point.

2.2 Projective Geometry with Field Elements

A field $\mathbb{F}$ is a set of elements which is closed under two binary operations, addition $(+)$ and multiplication $(\times)$. Both addition and multiplication must be associative and commutative in $\mathbb{F}$. $\mathbb{F}$ must also include the additive and multiplicative inverses of its elements, and the additive and multiplicative identities. By convention, the additive identity is designated as 0, and the multiplicative identity is designated as 1. Furthermore, the operations of $\mathbb{F}$ must obey the distributive property.

We can create an n-dimensional projective geometry based on a field $\mathbb{F}_m$ having m members by identifying as points the equivalence classes of $(n+1)$-tuples $(x_0, ..., x_n)$, where each x_i is an element of

$\mathbb{F}$ and not every x_i is the additive identity element 0, under the equivalence $(x_0, ..., x_n) \sim r(x_0, ..., x_n)$ for any $r \in \mathbb{F}_m \setminus \{0\}$. We call this geometry $\mathbb{FP}^n$. We may also deal with fields having infinitely many members, for which we use a common name for the field rather than $\mathbb{F}_\infty$. For example, if we let each x_i be any element of the field of real numbers $\mathbb{R}$, then the set of $(n+1)$-tuples $(x_0, ..., x_n)$, under the restriction $(x_0, ..., x_n) \sim r(x_0, ..., x_n)$ for any $r \in \mathbb{R} \setminus \{0\}$, is the real projective geometry $\mathbb{RP}^n$. We may add projective points coordinatewise as we add vectors in $\mathbb{R}^n$.

Furthermore, we can classify points in $\mathbb{FP}^n$ according to whether the first element x_0 is zero or nonzero:

Definition 1. An *ordinary point* in a projective geometry $\mathbb{FP}^n$ is a point $(x_0, ..., x_n)$ such that $x_0 \neq 0$.

Because $x_0 \neq 0$, the point $(x_0, ..., x_n)$ may be uniquely written in the form $(1, x_1 x_0^{-1}, ..., x_n x_0^{-1})$.

Definition 2. An *ideal point* in a projective geometry $\mathbb{FP}^n$ is a point $(x_0, ..., x_n)$ such that $x_0 = 0$.

Each ideal point $(0, x_1, ..., x_n)$ may also be written uniquely, with the first nonzero coordinate as a 1. For example, the projective point $(0, 0, 3, -12)$ in $\mathbb{RP}^3$ may be written as $(0, 0, 1, -4)$. We can summarize these unique representations of ordinary and ideal points using the following definition:

Definition 3. The *standard representation* of a point $X = (x_0, ..., x_n)$ in $\mathbb{FP}^n$, is the unique $(n+1)$-tuple $X^\circ = (y_0, ..., y_n)^\circ$ in the equivalence class of $(x_0, ..., x_n)$ whose first nonzero coordinate is the multiplicative identity element 1.

It is an exercise for the reader to show that the standard representation of a point $(x_0, ..., x_n)$ is unique as claimed.

We can define a line in $\mathbb{FP}^n$ as follows:

Definition 4. Let X and Y be two distinct projective points. Then the line $\overleftrightarrow{XY}$ is the following set of projective points:

$$\overleftrightarrow{XY} = \{X\} \cup \{Z \mid Z = rX + Y \text{ for some } r \in \mathbb{F}\}.$$

Note that if X and Y are both ideal points, then for any representation $(x_0, ..., x_n)$ of X and $(y_0, ..., y_n)$ of Y, $x_0 = y_0 = 0$. Then any representation $(z_0, ..., z_n)$ of Z satisfies $z_0 = rx_0 + y + 0 = 0$. This implies that if a projective line contains two ideal points, every point on the line must be ideal. We will call such a line an *ideal line*. Conversely, if there is a single point on a projective line which is ordinary, the line contains at most one ideal point. In fact, suppose that the line contains only ordinary points. Then for $r = -y_0 x_0^{-1}$, $z_0 = rx_0 + y_0 = -y_0 + y_0 = 0$, implying that Z is an ideal point - a contradiction. This implies that any projective line which has at least one ordinary point has exactly one ideal point. We will call such a line an *ordinary line*.

For our scheduling application, we will consider 2-dimensional projective geometries based on a finite field $\mathbb{F}_m$ having m members. We characterize these fields according to two theorems stated in Chapter 15 of Artin [1]:

Theorem 1. *The order of any finite field is a power of a prime.*

Theorem 2. *All fields of order m are isomorphic.*

In particular, we may consider $\mathbb{F}_p$, with p prime, which may be represented as the following:

Corollary 1. *Any field $\mathbb{F}_p$ with p a prime number is isomorphic to the field $\mathbb{Z}/p$, consisting of the set of integers $\{0, 1, ..., p-1\}$ under the operation of arithmetic modulo p.*

Each ordinary point in $\mathbb{F}_p\mathbb{P}^n$ has the standard representation $(1, x_1, ..., x_n)^\circ$, where each $x_i \in \mathbb{F}_p$. As there are p possibilities for each x_i, there are p^n ordinary points in $\mathbb{F}_p\mathbb{P}^n$. We can take this one step further to determine how many points are in each ordinary line, and how many ordinary lines are incident with each point.

Proposition 1. *Each line in $\mathbb{F}_p\mathbb{P}^n$ contains exactly $p+1$ points. Each ordinary line contains exactly p ordinary points and one ideal point, and each ideal line contains exactly $p+1$ ideal points.*

Proof. By Definition 4, the set of points in a projective line incident with two distinct points Y and X is $\{Y + 0X, Y + X, ..., Y + (p-1)X, X\}$. These $p+1$ points will not represent unique projective points if for some distinct $r, s \in \mathbb{F}_p$, $Y + rX \sim Y + sX$, or if $Y + rX \sim X$. This is equivalent to, for $t \in \mathbb{F}_p$:

$$\begin{aligned} Y + rX &= t(Y + sX) \qquad \text{or} \qquad & Y + rX &= tX \\ Y(1-t) &= X(s-r) & Y &= (t-r)X \\ Y(1-t)(s-r)^{-1} &= X & & \end{aligned}$$

Because $r \neq s$, $s - r \neq 0$ and so $(s-r)$ has the given multiplicative inverse. The above implies that Y and X are, in either case, in fact equivalent, which contradicts the assertion that Y and X are distinct. Therefore, the $p+1$ points on the projective line are distinct.

Because each ordinary line contains exactly one ideal point, this leaves p ordinary points on the ordinary line. Because an ideal line has no ordinary points, it must have $p+1$ ideal points. This completes the proof. □

Corollary 2. *There is only one ideal line in $\mathbb{F}_p\mathbb{P}^2$.*

Proof. By Proposition 1, each ideal line in $\mathbb{F}_p\mathbb{P}^2$ contains exactly $p+1$ ideal points. An ideal point in $\mathbb{F}_p\mathbb{P}^2$ has standard representation $(0, 1, x_2)$ or $(0, 0, 1)$; hence there are $p+1$ ideal points in $\mathbb{F}_p\mathbb{P}^2$. Therefore, each ideal line contains all of the ideal points in $\mathbb{F}_p\mathbb{P}^2$. By Axiom 1, the line between any two ideal points is unique. Thus there is only one ideal line, incident with all $p+1$ ideal points. □

Proposition 2. *Each point in $\mathbb{F}_p\mathbb{P}^2$ is incident with exactly $p+1$ lines. Each ordinary point is incident with $p+1$ ordinary lines, and each ideal point is incident with p ordinary lines, and one ideal line.*

Proof. There are $p^2 + p + 1$ points in $\mathbb{F}_p\mathbb{P}^2$, and so the number of pairs of points in this geometry is

$$\binom{p^2+p+1}{2} = \frac{(p^2+p+1)(p^2+p)}{2}.$$

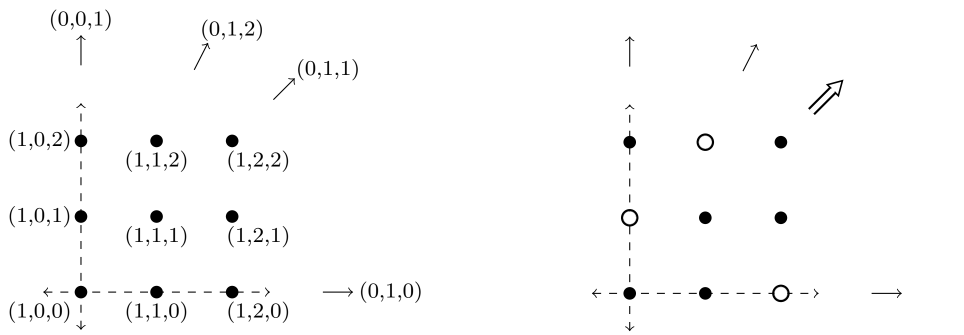

(a) The ordinary and ideal points of $\mathbb{F}_3\mathbb{P}^2$.

(b) The line through $(1,0,1)^\circ$ and $(1,1,2)^\circ$.

Figure 1: A visual representation of $\mathbb{F}_3\mathbb{P}^2$.

Each of these pairs of points defines a projective line in $\mathbb{F}_p\mathbb{P}^2$. Each line contains $p+1$ points, and so can be defined by using any of $\binom{p+1}{2} = \frac{p(p+1)}{2}$ pairs of points on the line; this means that each line has been counted $\frac{p(p+1)}{2}$ times. There are therefore

$$\frac{(p^2+p+1)(p^2+p)}{2} \div \frac{p(p+1)}{2} = p^2+p+1$$

lines in $\mathbb{F}_p\mathbb{P}^2$. Since each line contains $p+1$ points, there are $(p^2+p+1)(p+1)$ points; however, each point has been counted multiple times. As there are in fact p^2+p+1 points, we can infer that each point has been counted $p+1$ times. This implies that each point is incident with $p+1$ lines.

Because no ordinary point may be incident with the ideal line, each ordinary point must be incident with exactly $p+1$ ordinary lines. Because there is only one ideal line, each ideal point is incident with this one ideal line, and therefore is incident with p ordinary lines. $\square$

We can visually represent $\mathbb{F}_p\mathbb{P}^2$ in a two-dimensional grid. Each ordinary point $(1,x_1,x_2)^\circ$ is represented in the grid as the ordered pair (x_1,x_2), and the ideal points are visualized as "points at infinity." The ideal points $(0,1,x_2)^\circ$ are represented as points at infinity if one continued to travel along the line through the origin with slope x_2, and the ideal point $(0,0,1)^\circ$ is represented at a point at infinity along the x_2-axis. For example, a visual representation of $\mathbb{F}_3\mathbb{P}^2$ is shown in Figure 1a.

An example of an ordinary line in $\mathbb{F}_3\mathbb{P}^2$, the line between the ordinary points $(1,0,1)^\circ$ and $(1,1,2)^\circ$, is shown in Figure 1b. As we treat $\mathbb{F}_p$ with arithmetic mod p, the point $(1,2,0)^\circ$ may be placed on the grid at $(2,0)$ as shown, or at $(-1,0)$ or $(2,3)$. By Proposition 1, there are four points on this line: the ordinary points $(1,0,1)^\circ$, $(1,1,2)^\circ$, and $(1,2,0)^\circ$, and the ideal point $(0,1,1)^\circ$. That this ideal point is on this line is illustrated by the fact that the line, in this visual representation, has a slope of 1, and so is in the direction of the two-dimensional point $(1,1)$.

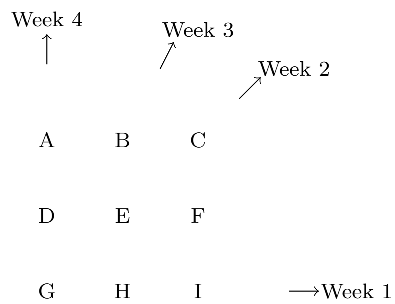

Figure 2: The assignment of teams and dates to the ordinary and ideal points of $\mathbb{F}_3\mathbb{P}^2$.

2.3 A Simple Schedule Obtained via Projective Geometry

The two-dimensional representation of $\mathbb{F}_p\mathbb{P}^2$ can be used as a model for scheduling problems. Given a certain number of teams, and a certain number of weeks in the desired schedule, one can use this representation to assign the teams to multi-team meets so that each team plays each other team exactly once. The model is designed as follows:

- Each ordinary point represents a team.
- Each ideal point represents a date in the schedule.
- Each ordinary line represents a meet.
- The ideal line represents the collection of dates (the calendar).

The axioms and propositions of projective geometry that are described earlier in this paper ensure that designing a schedule in this way meets the requirements. By Axiom 1, there is only one line incident with both of two distinct projective points. If both points are ordinary points, then this axiom implies that two teams compete together in only one meet. If one point is ordinary and the second point is ideal, this axiom implies that the team competes in only one meet on a given date. Proposition 1 states that there are p ordinary points and one ideal point on any ordinary line. This implies that every meet consists of p teams competing, and occurs on exactly one date. Proposition 1 also states that the ideal line contains all $p+1$ ideal points, which implies simply that each of the $p+1$ dates chosen is included in the schedule. By Proposition 2, there are $p+1$ ordinary lines through any ordinary point. This implies that each team competes in $p+1$ meets over the $p+1$ dates in the schedule.

For example, if we have nine teams to match up in a four-week schedule, we can use the visual representation of $\mathbb{F}_3\mathbb{P}^2$ to design this schedule. This perfectly fits the setup, since there are nine (3^2) ordinary points in $\mathbb{F}_3\mathbb{P}^2$ and four $(3+1)$ ideal points. Suppose we label these teams Teams A, B, C, D, E, F, G, H, and I, who need to compete against each other across Weeks 1, 2, 3, and 4. Then we can assign these nine teams and four dates in the visual representation as shown in Figure 2.

Consider Team D. There are four lines through the point assigned to Team D; each of these four lines is shown in Figure 3. In Week 1, Team D competes with Teams E and F; in Week 2, Team D competes with Teams B and I; in Week 3, Team D competes with Teams C and H; and in Week 4, Team D competes with Teams A and G.

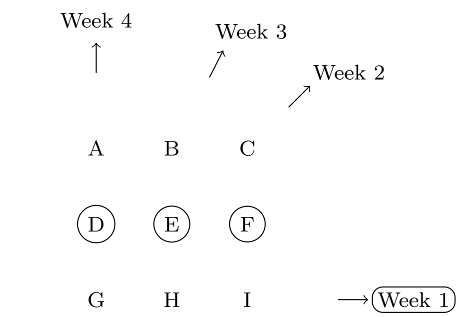

(a) The teams competing with Team D in Week 1.

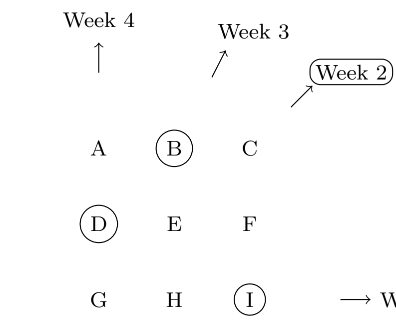

(b) The teams competing with Team D in Week 2.

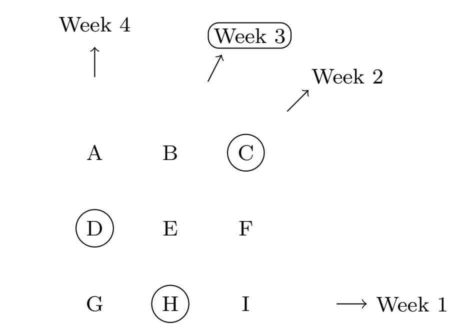

(c) The teams competing with Team D in Week 3.

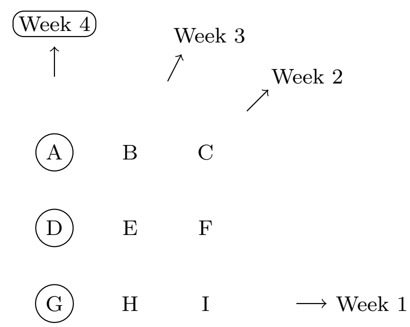

(d) The teams competing with Team D in Week 4.

Figure 3: The teams competing with Team D each of the four weeks.

Week 1:	$(---)(---)(--)(--)$
Week 2:	$(---)(---)(--)(--)$
Week 3:	$(---)(---)(--)(--)$
Week 4:	$(---)(---)(--)(--)$
Week 5:	$(---)(---)(--)(--)$
Week 6:	$(--)(--)(--)(--)(--)$

Figure 4: The basic setup for the desired schedule. Each $-$ represents one team.

3 The Algorithm

In the example given above, the number of teams and the number of weeks in the schedule were chosen so that each point in $\mathbb{F}_3\mathbb{P}^2$ corresponded to a team or to a date in the schedule. In practical use, the situation often arises in which the number of teams and number of weeks in the schedule do not match up so nicely with a representation of $\mathbb{F}_p\mathbb{P}^2$.

However, we can tweak the design. We can let the number of weeks in the schedule still match the number of ideal points in the representation, but we can allow fewer teams than ordinary points. In such cases, every team is assigned to an ordinary point in $\mathbb{F}_p\mathbb{P}^2$, but not every ordinary point is assigned to a team. There are then a number of blank spaces in the model, where a meet that contains three ordinary points, for example, may contain only one or two teams matched up, or none at all.

We are faced with the task of assigning ten teams in a six-week schedule, using a combination of meets which consist of three teams and meets which consist of two teams. Because the schedule takes six weeks, we can model the schedule using $\mathbb{F}_5\mathbb{P}^2$, which has six ideal points. There are 25 ordinary points in $\mathbb{F}_5\mathbb{P}^2$, but we can choose ten of these 25 points to represent our ten teams. In this way, an ordinary line which contains five ordinary points will represent a meet of only two or three teams instead of five.

The difficulty lies in choosing these ten points. If the ten points are chosen incorrectly, then the schedule designed may contain meets with only one team, or meets with four or five teams. The schedule could also contain only meets with two or three teams, but in uneven numbers.

The algorithm presented below will choose the ten ordinary points indirectly - by choosing the fifteen ordinary points that will not represent any of the ten teams in the schedule. (It is this scheme together with a proof that it works that is the focus of this paper.) The basic setup of the schedule is shown in Figure 4.

In the first five weeks, there are only four meets. However, each ideal point in the projective model of $\mathbb{F}_5\mathbb{P}^2$ is incident with five ordinary lines, implying that five meets take place in each of these weeks. Therefore, we can eliminate our fifteen points by eliminating five ordinary lines, one incident with each of five ideal points. The sixth ideal point will correspond to the last week of the schedule, which still has five meets. In this way, we will have five weeks of four meets each (corresponding to five ideal points incident with four ordinary lines each) and a sixth week of five meets (corresponding to the sixth ideal point incident with five ordinary lines). Because we deal only with ordinary lines in this model to represent the meets, it may be assumed that for the remainder of this paper, the lines mentioned are ordinary lines. We will not involve the ideal line in the algorithm.

Each line eliminated contains five ordinary points, and no two of these five lines are incident with the same ideal point. In Section 3.3, we will prove that in order to eliminate fifteen points from 25 using five lines, it is necessary that each of the five lines intersects each of the other four lines at four distinct points. This is equivalent to claiming that no three of the lines intersect in the same point. Furthermore, this choice of five lines will always yield a schedule of the pattern shown in Figure 4.

This, then, is the algorithm: choosing these five lines. Suppose that $\alpha, \beta, \gamma, \delta, \epsilon, \zeta$ are the six ideal points in $\mathbb{F}_5\mathbb{P}^2$, in no particular order. We may label ordinary lines in $\mathbb{F}_5\mathbb{P}^2$ according to the single ideal point with which it is incident. An ordinary line incident with the ideal point α will be called an α-line, an ordinary line incident with the ideal point β will be called a β-line, and so on. Then the algorithm proceeds as follows:

1. Eliminate all five points on an α-line in the model.
2. Eliminate the four points on a β-line that have not been previously eliminated. (The fifth point on the β-line was also on the α-line and has, thus, already been eliminated.)
3. Eliminate the three points on a γ-line in the model that intersects each of the previous lines at distinct points, that have not been previously eliminated. (The fourth and fifth points on the γ-line were on the α-line or β-line and have already been eliminated.)
4. Eliminate the two points on a δ-line in the model that intersects each of the previous lines at distinct points, that have not been previously eliminated.
5. Eliminate the single point on an ϵ-line in the model that intersects each of the previous lines at distinct points, that have not been previously eliminated.
6. By this step, $5 + 4 + 3 + 2 + 1 = 15$ points will have been eliminated. Assign the ten teams to the remaining points in the model, yielding the desired schedule.

This final step may be done using any permutation of the ten teams assigned to the ten remaining points. Therefore, if we want the sixth and final week of the schedule to be powerized as described in Section 1, we may assign the teams so that the first and second place teams are on the same ζ-line, the third and fourth place teams are on the same ζ-line, and so on. The algorithm is not difficult to perform, but the fact that the algorithm does generate the desired schedule is not trivial. In order to completely solve this scheduling problem, we must prove four assertions about our algorithm:

- That the pattern in Figure 4 is the only one that meets the requirements that over six weeks, each of the ten teams plays each of the other nine exactly once, and that each team competes in the same number of dual and triangular meets. See Section 3.1.
- That the algorithm can always be performed to completion. In other words, regardless of the choices of lines to eliminate in previous steps, a line may be found in each step which is incident with the desired ideal point and which intersects the other four elimination lines at four distinct points. See Section 3.2.

- That the algorithm works: that after running the algorithm, assigning the ten teams to the remaining ten ordinary points of the model does yield a schedule of the pattern shown in Figure 4. See Section 3.3.
- That the method of choosing eliminated lines as detailed in the algorithm is in fact the only way to eliminate the necessary 15 points in the model. This implies that not only does the algorithm yield some solution, but that it can yield any solution obtained by any other method. See Section 3.4.

3.1 The Combinatorial Design

The following theorem and its proof address the first of the four bullet points listed at the end of the last section; the proof relies on an elementary counting scheme.

Theorem 3. *Suppose that ten teams are to be scheduled in competition across six weeks, such that each team plays each other team exactly once, that each team plays in either a dual meet or triangular meet each week, and that each team plays the same number of dual and triangular meets as the other teams. Then five weeks must consist of two dual meets and two triangular meets, and one week must consist of five dual meets.*

In other words, the setup designed by the Hastings athletic staff in 2009, as shown in Figure 4, is the only possible design for this schedule. That the week consisting of five dual meets is chosen to be the sixth week is an arbitrary choice, as any of the ideal points in the model may represent any of the weeks.

Proof. Each team must play nine teams in six weeks. If the team is not allowed a bye, then it must compete against at least one team every week. For each team, then, we can assign six of the remaining nine teams to compete against it in each week of the schedule. This leaves three teams remaining. Because no team can compete against three or more other teams, we assign each of the remaining three teams to compete against the original team in one week of the schedule. This yields three weeks in which the team plays only one other team, and three weeks in which the team plays two other teams. This confirms that each team must play in exactly three dual meets and three triangular meets.

Because there are ten teams, each competing in three dual meets and three triangular meets, there are a total of 30 teams competing in dual meets (including multiplicity) and 30 teams competing in triangular meets (including multiplicity), where each team is counted a total of six times, once per meet competed in. As each dual meet consists of two teams, and each triangular meet consists of three teams, there are 15 dual meets and 10 triangular meets over the six-week schedule.

Because there are ten teams, it is not possible for a given week to consist of only triangular meets. There are therefore two possibilities: in any given week, either 5 dual meets occur, or (if three of those dual meets are replaced by two triangular meets) 2 dual meets and 2 triangular meets occur.

Because there are ten triangular meets that must occur, five weeks out of the six must consist of the latter format, with two dual meets and two triangular meets. This leaves the sixth week to consist of five dual meets, which does yield 15 dual meets in total. □

3.2 Completion of the Algorithm

It is necessary to prove that the algorithm may always be run to completion. How do we know, for example, that after choosing the α-line, β-line, and γ-line, that a δ-line with the required number of intersections does, in fact, still exist? As we will prove, each elimination line in the algorithm may be chosen without regard to the choices that need to be made later. Regardless of previous choices, there will always be a line that may be chosen as an elimination line in each step. The last step in the algorithm then may be completed if the first five steps of the algorithm yield ten points which have not been eliminated.

First, we may define *single elimination points* and *double elimination points* in this model. A single elimination point will denote any point which is incident with exactly one elimination line, and a double elimination point will denote any point which is incident with exactly two elimination lines. We may extend this definition to triple or quadruple elimination points, but because our model permits no three elimination lines to intersect in a single point, we will not encounter these.

Furthermore, we will denote the five elimination lines by ℓ_α, ℓ_β, ℓ_γ, ℓ_δ, and ℓ_ϵ, the elimination lines incident with ideal points α, β, γ, δ, and ϵ, respectively. We may also label the double elimination points by the elimination lines with which each is incident; for example, we may let $p_{\alpha,\beta}$ be the point incident with ℓ_α and ℓ_β, and so on.

Lemma 1. *There is an α-line which may be chosen as the first elimination line ℓ_α.*

Proof. This lemma is trivial. Regardless of our choice of which ideal point is α, there are five α-lines, and we may choose any of these five. □

Lemma 2. *There is a β-line which may be chosen as the second elimination line ℓ_β.*

Proof. Regardless of our choice of which ideal point is β, there are five β-lines. Each of these five lines intersects ℓ_α exactly once, by Axiom 1. Therefore, we may choose any of these lines to be ℓ_β. □

Lemma 3. *There is a γ-line which may be chosen as the third elimination line ℓ_γ.*

Proof. Regardless of our choice of which ideal point is γ, there are five γ-lines. Each of these lines intersects ℓ_α and ℓ_β exactly once. However, our choice of ℓ_γ may not be incident with $p_{\alpha,\beta}$. Because all five γ-lines intersect at γ, no two of them may intersect at $p_{\alpha,\beta}$ by Axiom 1. Therefore, only one γ-line is incident with $p_{\alpha,\beta}$. We may choose any of the other four γ-lines to be ℓ_γ. □

Lemma 4. *There is a δ-line which may be chosen as the fourth elimination line ℓ_δ.*

Proof. Regardless of our choice of which ideal point is δ, there are five δ-lines, each intersecting ℓ_α, ℓ_β, and ℓ_γ exactly once. As when choosing ℓ_γ, we have double elimination points to avoid: our choice of ℓ_δ may not be incident with $p_{\alpha,\beta}$, $p_{\alpha,\gamma}$, and $p_{\beta,\gamma}$. By Axiom 1, there are at most three δ-lines incident with one of these double elimination points, leaving at least two δ-lines that are incident with none of them. We may choose either of these two δ-lines to be ℓ_δ. $\square$

The construction implicit in Lemmas 1-4 yields no triple elimination points, as we allow no new line to include any double elimination points from previously chosen lines. These lemmas imply that the six double elimination points $p_{\alpha,\beta}$, $p_{\alpha,\gamma}$, $p_{\alpha,\delta}$, $p_{\beta,\gamma}$, $p_{\beta,\delta}$, and $p_{\gamma,\delta}$ generated by the first four elimination lines are distinct.

In choosing the first four elimination lines, we may choose any of the six ideal points to be α, β, γ, and δ. However, the choice of ϵ is not arbitrary, as shown in the proof of the following lemma:

Lemma 5. *There is exactly one line which may be chosen as the fifth and final elimination line ℓ_ϵ.*

Proof. After the choices implied by Lemmas 1-4, there are two ideal points not incident with the elimination lines. Incident with each are five lines, each intersecting ℓ_α, ℓ_β, ℓ_γ, and ℓ_δ exactly once. Our choice for ℓ_ϵ may not be incident with the six double elimination points. By counting alone, it is possible that for a given ideal point, each of the five lines incident with it is also incident with one of the six double elimination points. Therefore we instead consider the ten lines incident with one of the two remaining ideal points. Of these ten lines, we will show that exactly nine are disqualified from being an elimination line by virtue of passing through one of the above six double elimination points. Then we choose the remaining line to be ℓ_ϵ, and the ideal point with which it is incident to be ϵ. The remaining ideal point we choose to be ζ.

There is, by Axiom 1, exactly one line incident with any pair of these points. For any pair of points incident with the same elimination line, this elimination line is that unique line. For example, the unique line between points $p_{\alpha,\gamma}$ and $p_{\beta,\gamma}$ is ℓ_γ. For any pair of points which are not incident with any of the same elimination lines, the unique line between them is either an ϵ-line or a ζ-line. For example, consider the unique line between points $p_{\alpha,\gamma}$ and $p_{\beta,\delta}$. This line cannot be an α-line or γ-line, since the only α-line and γ-line incident with $p_{\alpha,\gamma}$ are ℓ_α and ℓ_γ, and $p_{\beta,\delta}$ is incident with neither. Similarly, the line cannot be a β-line or δ-line, and so it must be either an ϵ-line or a ζ-line. Therefore three of the ten ϵ-lines and ζ-lines are incident with two double elimination points: one through $p_{\alpha,\beta}$ and $p_{\gamma,\delta}$, one through $p_{\alpha,\gamma}$ and $p_{\beta,\delta}$, and one through $p_{\alpha,\delta}$ and $p_{\beta,\gamma}$.

Each of the six double elimination points is incident with one ϵ-line and one ζ-line, yielding twelve ϵ-lines and ζ-lines in total. By the above argument, three of those lines are incident with two double elimination points, and have thus been counted twice. Therefore, nine of the ϵ-lines and ζ-lines are incident with the collection of double elimination points. By interchanging the labels ϵ and ζ if necessary, we may assume that the tenth line, which is not incident with any double elimination points, is an ϵ-line, which we choose to be ℓ_ϵ. $\square$

Thus we have eliminated a total of fifteen points, leaving ten for assignment to the ten teams of the schedule. Ten of the elimination points are double elimination points, leaving five single elimination points. This completes the proof of the following theorem:

Theorem 4. *The elimination lines in the first five steps of the algorithm may be chosen without consideration of later steps, and the algorithm may be run to completion.*

3.3 Generation of a Correct Schedule

In order to prove that the schedule generated by the algorithm is indeed the schedule shown in Figure 4, we need to prove two items. First, we must prove that after the algorithm is run, the five ζ-lines in the model each yield a dual meet, corresponding to the sixth week of the schedule shown in Figure 4. Second, we must prove that each of the sets of five α-lines, β-lines, γ-lines, δ-lines, and ϵ-lines yields two dual meets and two triangular meets (each of four lines corresponds to a meet, and the fifth line, the elimination line, corresponds to no meet).

The schedule generated by the algorithm depends on the order in which we chose the ideal points α, β, γ, δ, and ϵ and their corresponding elimination lines, but only in terms of the order of the first five weeks of the schedule. We may consider two schedules equivalent if one may be obtained by permuting the weeks or teams of the other; under this equivalence, then, the schedule does not depend on the order in which we chose the elimination lines.

Theorem 5. *After running the algorithm, each of the five ζ-lines corresponds to a dual meet which takes place on the date corresponding to ideal point ζ.*

Proof. Each ζ-line intersects each of the five elimination lines exactly once. Because each elimination point on a ζ-line is either a single or double elimination point, there must be at least three elimination points on each ζ-line (two double elimination points, and one single elimination point). Any fewer would require triple or quadruple elimination points, which is not possible given the choices of lines in the algorithm.

There are fifteen elimination points in total, five ζ-lines, and no elimination points which are on multiple ζ-lines. Thus each of five lines contains at least three elimination points, and there are fifteen elimination points total; hence none of the lines may contain more than three. Therefore, there are exactly three elimination points on each ζ-line, leaving two ordinary points on each ζ-line which are assigned to teams. Each ζ-line, representing a meet, therefore is a dual meet of the two teams assigned to the two points remaining. $\square$

Theorem 6. *After running the algorithm, each of the sets of five α-lines, five β-lines, five γ-lines, five δ-lines, and five ϵ-lines consists of two lines which correspond to a dual meet, two lines with correspond to a triangular meet, and one line, the elimination line, which corresponds to no meet.*

Proof. The proof of this theorem is threefold. First, we prove that none of these lines correspond to a meet of four or more teams. Second, we prove that if one of these sets of five lines corresponds to four meets which are either dual or triangular, then two lines must correspond to dual meets and two lines must correspond to triangular meets. Third and finally, we prove that none of these lines corresponds to a meet of one team.

Claim 1: No α-line, β-line, γ-line, δ-line, or ϵ-line may correspond to a meet of four or more teams.

This is equivalent to claiming that no such line contains one or fewer elimination points, as one or zero elimination points would leave four or five points in the meet to be assigned to teams.

Consider any α-line, β-line, γ-line, δ-line, or ϵ-line remaining on the model which is not one of the five elimination lines (each elimination line trivially contains five elimination points). Let this line be m. The line m is incident with one ideal point, and so intersects one of the elimination lines at this ideal point. It must intersect the other four elimination lines at not necessarily distinct ordinary points; it is therefore impossible for m to be incident with no elimination points.

Suppose that m contains exactly one elimination point. Then this elimination point must be incident with each of those four elimination lines. This is impossible, because the choice of lines in the algorithm does not allow for quadruple elimination points. Therefore, m must contain more than one elimination point. This completes the proof of Claim 1.

Claim 2: If any set of five α-lines, β-lines, γ-lines, δ-lines, or ϵ-lines contains one elimination line corresponding to no meet, and four lines corresponding to dual or triangular meets, then two lines in the set must correspond to dual meets, and two lines in the set must correspond to triangular meets.

Suppose that each set of five α-lines, β-lines, γ-lines, δ-lines, and ϵ-lines contains the elimination line ℓ_α, ℓ_β, ℓ_γ, ℓ_δ, or ℓ_ϵ, respectively, and four other ordinary lines which correspond to meets of two or three teams. Then each line is incident at least two and at most three elimination points.

Consider one set of five lines; without loss of generality, consider the five α-lines. Aside from ℓ_α, there are four α-lines, each consisting of either two or three elimination points. The line ℓ_α contains five of the fifteen elimination points, leaving ten to be incident with the four remaining α-lines. By the pigeonhole principle, eight of these ten remaining elimination points must be incident with these four lines, two per α-line. This leaves two elimination points to be assigned to some two of the four α-lines (because none of these α-lines, save ℓ_α, may be incident with four elimination points, the two final elimination points must be incident with different lines). Therefore, two of the remaining four α-lines are incident with two elimination points, and two are incident with three elimination points. This yields two α-lines which correspond to triangular meets, and two α-lines which correspond to dual meets, respectively. This completes the proof of Claim 2.

Claim 3: No α-line, β-line, γ-line, δ-line, or ϵ-line may correspond to a meet of one team.

We prove that no such line may be incident with four elimination points. Consider the collection of four α-lines that are not ℓ_α. Each α-line is incident with ℓ_β, ℓ_γ, ℓ_δ, and ℓ_ϵ, but not ℓ_α; hence if an α-line is incident with four elimination points, these must be single elimination points.

Suppose that ℓ_ϵ has not yet been eliminated (though it is determined once we have eliminated ℓ_α, ℓ_β, ℓ_γ, and ℓ_δ). Then there are three double elimination points $p_{\beta,\gamma}$, $p_{\beta,\delta}$, and $p_{\gamma,\delta}$ which are incident with the α-lines (excluding ℓ_α). By the proof of Lemma 5 none of these three double elimination points are incident with the same α-line. Hence three α-lines are incident with a distinct point in $\{p_{\beta,\gamma},\ p_{\beta,\delta},\ p_{\gamma,\delta}\}$. Each of these α-lines thus contains two elimination points at this point in the algorithm, and after eliminating ℓ_ϵ will contain no more than three. The fourth α-line is incident with none of them, but instead with three single elimination points. Let this line be m_α.

Similar proofs show that there are also one β-line, one γ-line, and one δ-line each incident with three elimination points. Let these lines be denoted by m_β, m_γ, and m_δ.

After eliminating ℓ_α, ℓ_β, ℓ_γ, and ℓ_δ, there are $14 = 5+4+3+2$ elimination points. Six of these are double elimination points, leaving eight single elimination points. Eliminating ℓ_ϵ yields ten double elimination points and five single elimination points. Hence ℓ_ϵ must change four single elimination points to double elimination points by intersection, and adds one single elimination point.

Next, we prove that all eight single elimination points belong to the union of m_α, m_β, m_γ, and m_δ. Notice that each single elimination point belongs to the union of ℓ_α, ℓ_β, ℓ_γ, and ℓ_δ. Suppose, for instance, that a particular single elimination point p belongs to ℓ_α; then it is not incident with m_α. There are three distinct double elimination points on ℓ_α. Once ℓ_ϵ is eliminated, Theorem 5 implies that the ζ-line through p must contain two double elimination points and one single elimination point. Thus, it contains at least one double elimination point prior to eliminating ℓ_ϵ, since ℓ_ϵ cannot be incident with more than one point on this ζ-line. p is thus connected to three of the six double elimination points by ℓ_α, and to at least one additional double elimination point by a ζ-line. Of the β-line, γ-line, and δ-line incident with p, at most two connect p to the remaining two double elimination points; hence one of these lines must not be incident with any double elimination points. This line is m_β, m_γ, or m_δ.

As our choice of ℓ_α in the preceding argument is arbitrary, we may conclude that all of the eight single elimination points belongs to the union of m_α, m_β, m_γ, or m_δ. ℓ_ϵ is incident with exactly one point on each of these four lines. Hence if ℓ_ϵ changes four single elimination points into double elimination points by intersection, each of the four points must be on distinct lines m_α, m_β, m_γ, and m_δ, since ℓ_ϵ is incident with only one α-line, one β-line, one γ-line, and one δ-line. Then each line after eliminating ℓ_ϵ still contains three total elimination points, not four.

Since the order of using α, β, γ, δ, and ϵ is irrelevant to the argument, this completes the proof that no α-line, β-line, γ-line, δ-line, or ϵ-line contains more than three elimination points.

Now, from Claims 1 and 3, we have proven that aside from the five elimination lines, every α-line, β-line, γ-line, δ-line, and ϵ-line corresponds to either a dual or triangular meet. By Claim 2, the number of dual and triangular meets is in accordance with the first five weeks of the schedule shown in Figure 4. This completes the proof of the theorem. □

Theorems 5 and 6 together imply that the six-week schedule yielded by the algorithm does, in fact, have the desired format.

3.4 Generation of All Correct Schedules

We have shown by now that the algorithm may always be run to completion, and that when run to completion, it produces a schedule in the desired format. Now, we prove that any schedule of the desired format has a corresponding set of five elimination lines in $\mathbb{F}_5\mathbb{P}^2$ that may be chosen from our algorithm. Because the first five weeks of the schedule consist of four meets, and there are five ordinary lines incident with each ideal point in our model, it is necessary to eliminate an entire line incident with each of five ideal points, to limit those first five weeks to four meets. However, it is not trivial that we must choose the lines so that each pair of lines intersects in distinct points. That this is necessary, not just sufficient, may be shown by a combinatorial argument.

Theorem 7. *If a collection of five elimination lines is incident with fifteen elimination points, it is necessary that each pair of elimination lines intersects at distinct points.*

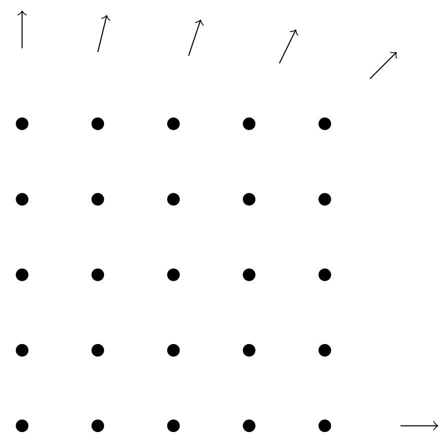

Figure 5: The geometry $\mathbb{F}_5\mathbb{P}^2$.

Proof. We must eliminate 15 out of 25 ordinary points in the model, to yield ten points to assign to ten teams. Each line eliminated contains five points. We should consider the number of points eliminated with each choice of elimination line. Recall that because each elimation line is necessarily incident with a different ideal point, each pair of elimination lines must intersect at some ordinary point.

The first elimination line chosen eliminates five points from the model.

The second elimination line chosen eliminates four points from the model, because it intersects the first line at an ordinary point which has already been eliminated.

The third elimination line must eliminate at least three points from the model. It will eliminate three points if it intersects each of the first two at distinct points, but will eliminate four if it intersects the first two at the same intersection point.

In similar fashion, the fourth elimination line must eliminate at least two points from the model: two points if it intersects each of the first three at distinct points, but more than two if it intersects the first three at any existing intersection points. The fifth line must eliminate at least one point from the model, by the same reasoning.

Therefore, the minimum number of points eliminated is $5+4+3+2+1 = 15$, which is attained only in the case where each line chosen does not share any existing intersection points. This completes the proof.

$\square$

3.5 Running the Algorithm

In this section we will work through the algorithm to obtain a desired schedule. Figure 5 shows the geometry $\mathbb{F}_5\mathbb{P}^2$ in which we will design our schedule.

As discussed, the choices of our first and second lines, ℓ_α and ℓ_β, are arbitrary. For this example, these choices are shown in Figures 6a and 6b. In these figures, a single slash indicates a single elimination point, and a double slash indicates a double elimination point. For our choice of ℓ_γ, we have only the restriction that it not include the double elimination point in the lower left corner.

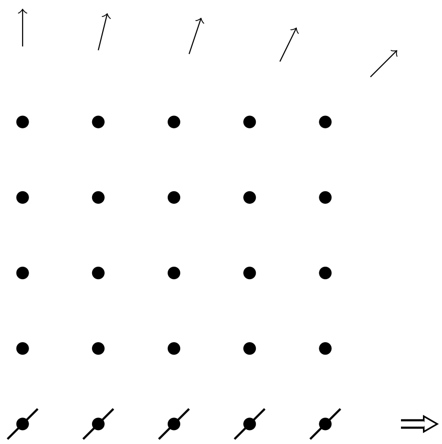

(a) The α-line, highlighted in black.

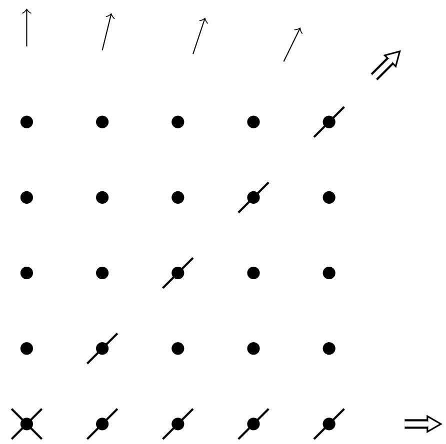

(b) The α- and β-lines, highlighted in black.

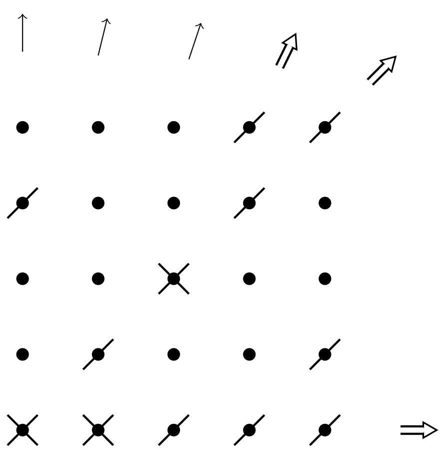

(c) The α-, β-, and γ-lines.

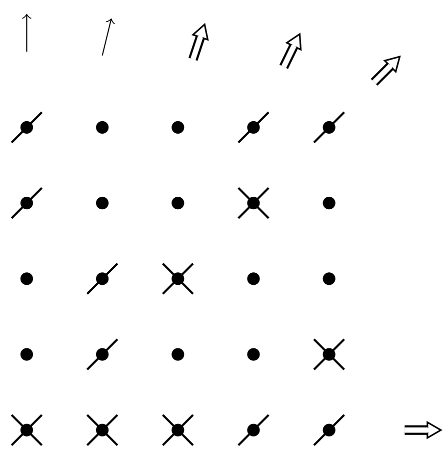

(d) The α-, β-, γ-, and δ-lines.

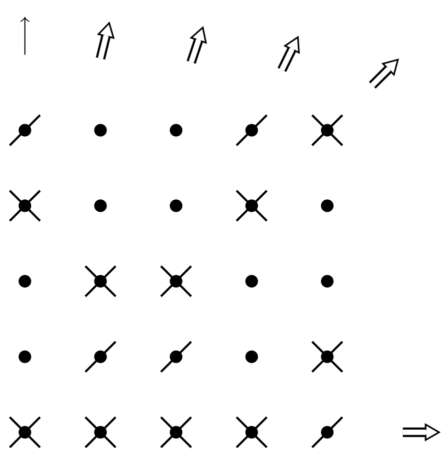

(e) All five lines chosen in the algorithm.

Figure 6: The five lines chosen in the algorithm.

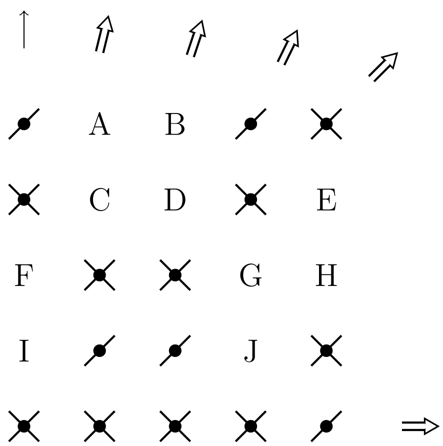

(a) The assignment of the ten teams.

Week 1:	(CDE) (FGH) (AB) (IJ)
Week 2:	(AHJ) (BCF) (DI) (EG)
Week 3:	(BEJ) (CGI) (AF) (DH)
Week 4:	(AEI) (DFJ) (BG) (CH)
Week 5:	(ADG) (BHI) (CJ) (EF)
Week 6:	(AC) (BD) (EH) (FI) (GJ)

(b) A sample schedule obtained by our algorithm.

Figure 7: The schedule generated by our algorithm. Letters A through J represent the ten teams.

There are still quite a few possibilities; one option is shown in Figure 6c. After choosing our ℓ_γ, we now have three double elimination points to avoid. If we choose the ideal point with which ℓ_δ is incident, there are five total δ-lines; we may choose either of the two which are not incident with the three double elimination points. One option is shown in Figure 6d.

We now only have one choice for ℓ_ϵ. Either of the two unhighlighted ideal points from Figure 6d could be ϵ. Each of the five vertical lines, however, is incident with at least one double elimination point; hence ϵ must be the ideal point $(0,1,4)^\circ$, not $(0,0,1)^\circ$. ℓ_ϵ must intersect the bottom row but not at any of the three double elimination points; hence it intersects either $(1,3,0)^\circ$ or $(1,4,0)^\circ$. The ϵ-line through $(1,4,0)^\circ$ intersects the double elimination point $(1,2,2)^\circ$, hence the ϵ-line through $(1,3,0)^\circ$ is ℓ_ϵ, shown in Figure 6e.

The single remaining unhighlighted ideal point corresponds to the sixth and final week of the schedule, consisting of five dual meets (each of the five lines incident with this ideal point contains exactly two ordinary points). For every other ideal point, the collection of five lines incident with this idea point consists of two lines containing exactly two ordinary points, two lines containing exactly three ordinary points, and the elimination line, indicating a week of two dual meets and two triangular meets as desired.

What remains is to assign the ten teams to the ten uneliminated points. This is shown in Figure 7a, and the schedule generated is shown in Figure 7b.

4 Home and Away Assignments: An Exercise in Graph Theory

Thus far we have determined who plays whom in each week of our schedule, but we have yet to determine where. Given the significant home advantage in nearly any sporting event, it is important to consider whether the ten teams are able to host the same number of meets. Furthermore, for complete fairness we may attempt to not only ensure that each team hosts the same number of meets, but given the disparity in nature of a dual meet compared to a triangular meet, also attempt to ensure that each team hosts the same number of each type of meet.

From a combinatorial perspective, accomplishing this with the triangular meets is an easier task. There are a total of ten triangular meets, and ten teams, hence we may be most fair by assigning each team one triangular meet to host. However, there are fifteen dual meets: ten in the first five weeks, but an additional five in the sixth week. We may instead assign each team one dual meet to host during the first five weeks, but in the sixth week, assign home advantages based on last year's results. We designed our schedule so that in the final week, last season's first and second place teams would compete head to head, the third and fourth place teams would complete head to head, and so on. As is typical in playoff brackets of most athletic events, we may simply assign home advantage to the higher placed team in each game.

However, determining home and away assignments for the first five weeks of the schedule is not arbitrary. To determine these assignments, we model the schedule of the first five weeks with two bipartite graphs and construct edge matchings.

4.1 Bipartite Graphs and Perfect Matchings

A *graph* is a collection of vertices connected by edges. We define a bipartite graph as follows:

Definition 5. A *bipartite graph* is a graph whose vertices may be partitioned into two sets such that no two vertices in the same partition are connected by an edge.

In addition to partitioning the vertices of a graph, we may consider partitioning the edges. A *matching* in a graph is a collection of edges without common vertices: no two edges share the same vertex. Clearly any graph may be partitioned into some collection of matchings (we may decide that, trivially, each edge is its own partition). However, it is more useful to attempt to partition a graph into as few matchings as possible, or to partition the graph into as large of matchings as possible. Because each vertex of the graph may be incident with at most one edge in the matching, we may never have a matching with more edges than the number of vertices. We define a perfect matching as follows:

Definition 6. Given a graph G, a *perfect matching* is a matching whose collection of edges is incident with every vertex in G.

Whether or not a graph allows a perfect matching, or whether its edges may be partitioned into only perfect matchings, is a significant open question in graph theory. Restricting our discussion to bipartite graphs, however, yields answers. Which finite bipartite graphs contain a matching that spans one partition of vertices was determined by Philip Hall [3]. In the following theorem, for a set of vertices A, $|A|$ denotes the number of vertices in the set:

Theorem 8 (Hall's Marriage Theorem)**.** *Let G be a bipartite graphs with vertex partitions X and Y. For a vertex subset $A \subset X$, let $f(A)$ denote the collection of vertices in partition Y which share an edge with a vertex in A. Then there exists a matching which is incident with every vertex in X if and only if for every subset A, $|A| \leq |f(A)|$.*

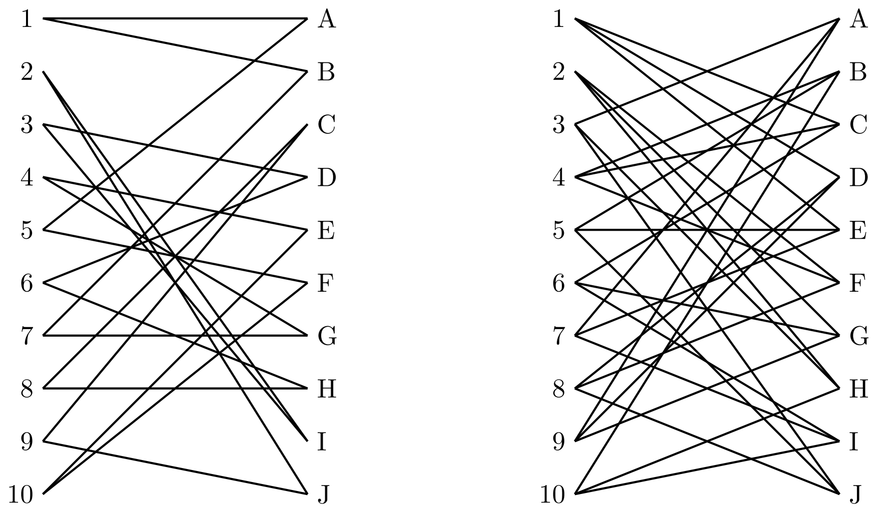

(a) The bipartite graph representing the dual meets.

(b) The bipartite graph representing the triangular meets.

Figure 8: Bipartite graphs representing the dual and triangular meets.

Note that if the two partitions X and Y contain the same number of vertices, this matching is a perfect matching. For such graphs, a perfect matching may be found by applying the Hopcroft-Karp algorithm [4], which produces a maximum cardinality matching given any bipartite graph (if a perfect matching exists, it is by definition maximum in size). This algorithm is illustrated in the latter portion of Section 4.2 below.

4.2 Home and Away Assignments

We may model the dual and triangular meets of the first five weeks of our schedule using two bipartite graphs. We will refer to the example schedule shown in Figure 7b.

For the collection of dual meets, we construct a bipartite graph with one partition of ten vertices corresponding to the ten dual meets in the first five weeks of the schedule, and a second partition of ten vertices corresponding to the ten teams. We label the ten dual meets 1-10 in the order in which they appear in the schedule written in Figure 7b. We then add edges connecting each team in the second partition to each dual meet in the first partition in which the team competes. Each vertex corresponding to a team is therefore incident with two edges, corresponding to the two dual meets in which the team competes, and each vertex corresponding to a meet is incident with two edges, corresponding to the two teams which compete in it. This graph is shown in Figure 8a.

We similarly construct a bipartite graph corresponding to the triangular meets. One partition contains ten vertices corresponding to the ten triangular meets, and the second partition contains ten vertices corresponding to the teams. We then add edges connecting each team in the second partition to each triangular meet in the first partition in which the team competes. Each vertex corresponding to a team is therefore incident with two edges, corresponding to the three triangular meets in which the team competes, and each vertex corresponding to a meet is incident with three

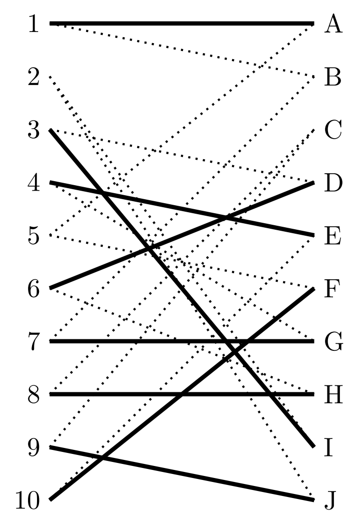

Figure 9: A partial matching M on the graph of dual meets.

edges, corresponding to the three teams which compete in it. This graph is shown in Figure 8b.

For any schedule our algorithm generates, consider a collection of k dual meets. Each meet consists of two teams, yielding $2k$ teams competing in these k meets. Since each team participates in two dual meets, these $2k$ teams may include teams counted at most twice, so there are at least k different teams competing in the k dual meets. This satisfies the hypothesis of Hall's theorem, so there is a perfect matching in the bipartite graph representing the dual meets. This implies that it is possible to determine home and away assignments so that each team hosts one dual meet and is away for one dual meet. Similarly, at least k different teams must compete in any k triangular meets, so the theorem implies that there is also a perfect matching in the bipartite graph representing the triangular meets. This implies that it is possible to determine home and away assignments so that each team hosts one triangular meet and is away for two triangular meets.

These perfect matchings can be found using the Hopcroft-Karp algorithm. To run the algorithm, we begin with any partial matching of edges, not necessarily incident with every vertex (this partial matching may be as small as a single edge, but we prefer to start with as large a matching as we can manage by arbitrary choices). Such a partial matching is shown in Figure 9, on the graph of dual meets; we will call this matching M. M is maximal, as no edge may be added to it to increase the matching. Vertices 2 and 5 on the left, and vertices B and C on the right, are not incident with M. We choose any of these remaining vertices, and choose an edge connecting it to a vertex on the other side; say we choose vertex 5 and the edge connecting it to vertex F, as shown in Figure 10a. Since M is maximal, we end up on a vertex which is incident with M. We then construct a path of edges, alternating between edges in M and edges not in M. Edges in M continue the path from right to left, and edges not in M continue the path from left to right. Since no vertex is incident with only edges in M, we may continue until we reach a vertex which is incident with no edges in M. Such a path is shown in Figure 10b, which ends on vertex B.

We now have an alternating path whose first and last edges are not in M and are not adjacent

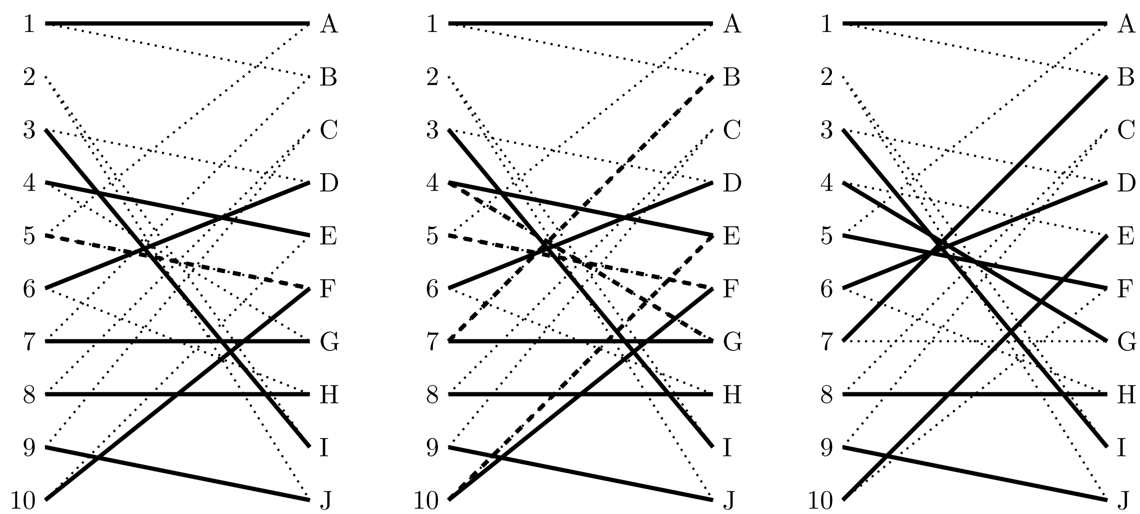

(a) An augmenting edge of M. (b) An augmenting path of M. (c) A larger partial matching M'.

Figure 10: Extending a partial matching M

to edges in M. The number of edges in the path which are not in M is one greater than the number of edges in the path in M. We then switch the edges in the path, so that the edges originally in M now are not, and the edges originally not in M now are. This adds one edge to M to form a new matching M', shown in Figure 10c. We may continue this with any vertex which is not in M' to form a new matching M'', and so on, until we run out of vertices, yielding a perfect matching.

Perfect matchings for the graphs representing dual meets and triangular meets are shown in Figure 11. The solid lines that form the perfect matchings indicate which team on the right hosts the meet on the left. Combined with the home assignments already determined for the final sixth week of the schedule, the task of determining home and away assignments is now complete.

5 Further Considerations

This paper focuses on a specific scenario regarding the number of teams involved, the number of weeks of the schedule, and allowing both dual and triangular meets to occur. It remains to be seen whether the algorithm presented may be adapted to a general case.

The number of teams in the schedule determines the number of elimination points that must be incident with the elimination lines in the algorithm. If, instead of a ten-team conference, the six-week schedule must accommodate eleven teams, we may only eliminate fourteen points from $\mathbb{F}_5\mathbb{P}^2$. As discussed in proving Theorem 7, if the elimination lines intersect at distinct points, at least fifteen points are necessarily eliminated by five elimination lines. Therefore if eleven teams compete, the algorithm must be adapted to either use fewer elimination lines, or allow for three lines to intersect at a single point. Whether these adaptations still allow the algorithm to work, according to assertions proven in Section 3, is an open question.

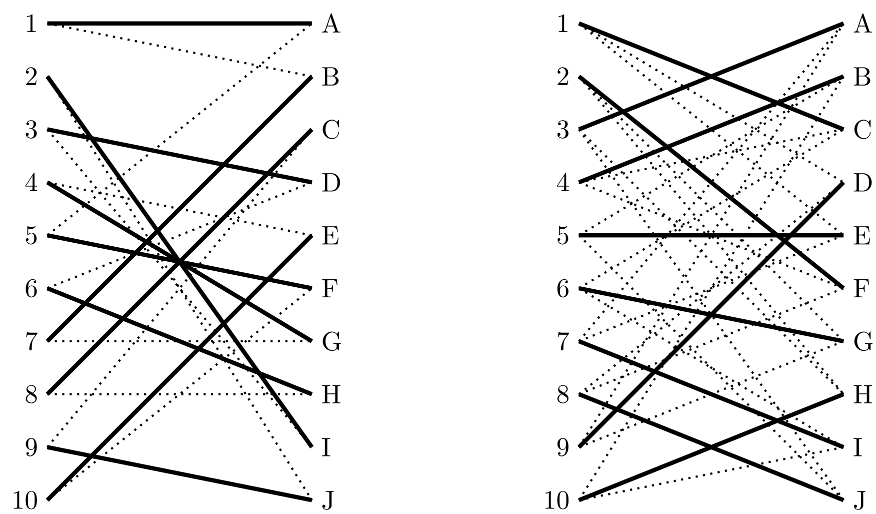

(a) The home team assignments for the ten dual meets of the first five weeks.

(b) The home team assignments for the ten triangular meets.

Figure 11: Home team assignments for the first five weeks of the schedule.

The number of weeks of the schedule determines the geometry used in the construction. Because the order of any finite field is a power of a prime number (Theorem 1), this model may only be directly adapted to construct a schedule of a number of weeks one greater than a power of a prime. For example, the algorithm presented in this paper could not be used directly to construct a seven-week schedule.

The algorithm presented allows for various final schedules, depending on the choices for elimination lines. Furthermore, given any schedule generated by the algorithm, there are 3840 different permutations of the ten teams that result in the same sixth week matchups but different matchups in the first five weeks. It is an open question whether the schedules generated by two different collections of five elimination lines are truly different, or if one may be obtained from the other by permuting the ten teams or the first five weeks.

References

[1] Michael Artin. *Algebra*. Prentice Hall, second edition, 2011.

[2] Albrecht Beutelspacher and Ute Rosenbaum. *Projective geometry: from foundations to applications*. Cambridge University Press, Cambridge, 1998.

[3] Philip Hall. On representatives of subsets. *Journal of the London Mathematical Society*, s1-10(1):26–30, 1935.

[4] John E. Hopcroft and Richard M. Karp. An $n^{5/2}$ algorithm for maximum matchings in bipartite graphs. *SIAM J. Comput.*, 2:225–231, 1973.