RECTILINEAR EMBEDDINGS OF K_n IN THE PLANE

JOHN WILTSHIRE-GORDON

Abstract. One way to draw a graph is to position its vertices in the plane, connecting adjacent vertices with line segments that correspond to edges. This rectilinear embedding raises many questions: "Which graphs have embeddings that have no crossings?", "How many essentially different embeddings does a certain graph have?", or even "Can I find an embedding that has a certain set of desired crossings?". This paper will concern itself primarily with the third question, restricting its attention to rectilinear embeddings of complete graphs. We will recast these geometrical problems in combinatorial terms, and end with a conjectured "good characterization" of valid combinatorial configurations.

Contents

1. Introduction 1
2. Definitions and Preliminary Observations 3
3. Theorems relating geometrical and combinatorial configurations 4
4. A conjecture 7
References 8

1. Introduction

Some branches of mathematics are known for the ease with which a novice can stumble across difficult problems. Although number theory is the most celebrated of these branches, the relatively young field of combinatorial geometry is certainly not far behind. The question of rectilinear crossing number—a quintessential problem of combinatorial geometry—is a beautiful example of a simple question that remains quite open.

The question first came up in railway design, of all places. If a rail yard has several loading docks or depots that all need to be connected by track, the stationmaster would do well to place the depots in such a way that the tracks do not cross, so as to lower the probability of collisions, derailments, etc.. For instance, if there were four depots, only a foolish stationmaster would place them at the corners of a convex quadrilateral: putting one depot in the convex hull of the other three avoids crossings altogether. Of course, a rail yard with five depots cannot avoid a crossing: K_5 is non-planar, and remains so with the extra restriction of line-segment edges. Nevertheless, a wise stationmaster can limit the configuration to a single crossing by putting two depots in the convex hull of the other three.

Date: July 22, 2008.

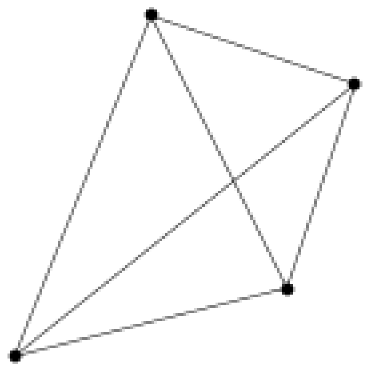

A dumb stationmaster

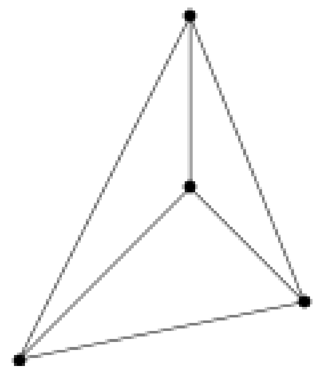

A smart stationmaster

With as few as twenty depots, however, the minimum number of crossings is already unknown. Asymptotics have been investigated, and bounds (usually derived from smaller cases) have been discovered, but a formula for the exact number remains well out of reach. Part of the difficulty of tackling the problem is the infinite size of the space of configurations: if each of n depots can move freely in two dimensions, the space of configurations will be $2n$-dimensional, with the same cardinality as the continuum. An exhaustive listing of these configurations would be not only infeasible, but impossible: the continuum is uncountable.

Still, the situation is not hopeless. If we consider two configurations to be different when their crossings are different, then we have a finite space of options: with only $m = \binom{n}{2}$ edges, there are at most $2^{\binom{m}{2}}$ options because for every pair of edges, we must decide whether they cross. Many of these options will be impossible, however. For instance, a combinatorial configuration may prescribe exactly 1000 crosses for K_{18}, even though the minimum number of crosses is known to be 1029. Such a configuration would therefore be invalid.

If there were an easy way to tell the valid configurations from the invalid ones, perhaps a sticky problem of geometry could be reduced to standard combinatorics: each configuration could be represented as a graph that takes the $\binom{n}{2}$ line segments as vertices, with two segments sharing an edge iff they cross. Theorems might be proved about graphs that correspond to valid configurations, and conclusions could be drawn about geometry using techniques from combinatorics.

This paper will conjecture an “easy test”—which is also known as a “good characterization”—for the geometrical validity of a combinatorial configuration.

2. Definitions and Preliminary Observations

Before we begin, we will make the problem precise. The above description relies on points "in the plane." As it turns out, we do not want this plane to have an origin or any other distinguished points. We want the "affine plane," a plane of points that behaves just like $\mathbb{R}^2$ translated some unknown amount from the origin. It will make sense to speak of the difference of two points—it is simply a vector—but their sum will remain undefined because the origin could be anywhere.

Definition 2.1. Let $\mathbb{E}^2$, the **affine plane**, be $\{\vec{p} - \vec{o} | \vec{p} \in \mathbb{R}^2\}$, where $\vec{o}$ is some unknown vector in $\mathbb{R}^2$ that represents the location of the origin. For convenience, we will write $\vec{p} - \vec{o}$ as simply P. Note that the difference of two elements of the affine plane is a vector in $\mathbb{R}^2$.

Definition 2.2. A **configuration** is a finite set of points in the affine plane.

Definition 2.3. A **linear combination** of P and Q in the affine plane is an element of the set

$$\{\alpha P + (1 - \alpha) Q | \alpha \in \mathbb{R}\}$$

These points are weighted averages of P and Q, and make up the straight line containing P and Q.

Observation 2.4. $\mathbb{E}^2$ *is closed under linear combinations.*

Proof.

$$\begin{aligned} Z &= \alpha P + (1 - \alpha) Q \\ Z &= \alpha (\vec{p} - \vec{o}) + (1 - \alpha) (\vec{q} - \vec{o}) \\ Z &= (\alpha (\vec{p} - \vec{q}) + \vec{q}) - \vec{o} \end{aligned}$$

So $Z \in \mathbb{E}^2$, as required. □

Definition 2.5. Given a configuration $\mathcal{A} = \{A_1, A_2, A_3, ..., A_k\}$ the **convex hull** of $\mathcal{A}$, which we denote $\overline{\mathcal{A}} = \overline{A_1 A_2 A_3 ... A_k}$, is the set of convex combinations of the A_i. That is, B is an element of $\overline{A_1 A_2 A_3 ... A_k}$ exactly when there exist β_i with $\sum_{i=1}^{k} \beta_i = 1$ and $0 \leq \beta_i$, such that

$$B = \sum_{i=1}^{k} \beta_i A_i$$

The previous proof can be extended to show that $\overline{\mathcal{A}} \subseteq \mathbb{E}^2$, that is, every element of a convex hull of points in the affine plane will also be in the affine plane.

Note that the conventional notation for a line segment between two points can be considered a special case of this notation: $\overline{PQ}$ means the set of points between P and Q, just as it should.

Definition 2.6. We say a point P is **inside** a configuration $\mathcal{C}$ when $P \in \overline{\mathcal{C}}$, where $\overline{\mathcal{C}}$ stands for the convex hull of the points in the configuration $\mathcal{C}$.

Definition 2.7. We say a configuration $\mathcal{C}$ is **convex** if removing a point makes its convex hull smaller. That is, $\mathcal{C}$ is convex if for all P in $\mathcal{C}$, $\overline{\mathcal{C} \setminus \{P\}} \subsetneq \overline{\mathcal{C}}$.

Definition 2.8. We say that $\overline{PQ}$ **crosses** $\overline{RS}$—usually written $\overline{PQ} \times \overline{RS}$—if the two line segments intersect at some point other than an endpoint. More precisely, $\left(\overline{PQ} \setminus \{P, Q\}\right) \cap \left(\overline{RS} \setminus \{R, S\}\right) \neq \emptyset$.

Definition 2.9. We say that $\overline{PQ}$ **avoids** $\overline{RS}$—usually written $\overline{PQ} \parallel \overline{RS}$—if the two line segments do not intersect: $\overline{PQ} \cap \overline{RS} = \emptyset$. Similarly, we say a segment **completely avoids** a configuration $\mathcal{C}$ if it avoids every segment with endpoints in $\mathcal{C}$.

Axiom 2.10. Two line segments never intersect at more than one point:

$$\|\overline{PQ} \cap \overline{RS}\| > 1 \Rightarrow \overline{PQ} = \overline{RS}$$

.

Definition 2.11. The **cross-count** of a configuration $\mathcal{C}$ is the number of crosses in $\mathcal{C}$. In other words, it is the cardinality of the following set of unordered quadruples:

$$\{\{P, Q, R, S\} | P, Q, R, S \in \mathcal{C} \wedge \left(\overline{PQ} \times \overline{RS} \vee \overline{PR} \times \overline{QS} \vee \overline{PS} \times \overline{QR}\right)\}$$

3. Theorems relating geometrical and combinatorial configurations

The following theorems show a correspondence between geometrical configurations and their combinatorial counterparts. These theorems are biconditional: they relate the two spaces in a powerful way. They allow us to understand the geometry in terms of the combinatorics and vice-versa.

Theorem 3.1. *A configuration $\mathcal{G}$ is convex $\Leftrightarrow$ $\mathcal{G}$ has a cross-count of $\binom{\|\mathcal{G}\|}{4}$*

Proof. Let $\mathcal{Q}$ be the set of unordered quadruples taken from the configuration $\mathcal{G}$:

$$\mathcal{Q} = \{\{P, Q, R, S\} | P, Q, R, S \in \mathcal{G}\}$$

$\Rightarrow$: Assuming $\mathcal{G}$ is convex, we must show that there are exactly $\binom{n}{4}$ crosses. Since the cardinality of $\mathcal{Q}$ is already $\binom{n}{4}$, it will suffice to show that every quadruple of points leads to a crossing. Let $\{P, Q, R, S\}$ be a quadruple in $\mathcal{G}$. If one of these three points were to lie in the convex hull of the other three, say $P = \alpha_1 Q + \alpha_2 R + \alpha_3 S$ then the set $\mathcal{G}$ would not be convex: removing the point P would have no effect on the convex hull because you could enlist the above mixture of Q, R, and S to take its place. We may assume, therefore, that there is no way to write any of these four points as a convex combination of the other three.

Let's take a look at three vectors emerging from P, $\vec{u} = Q - P$, $\vec{v} = R - P$, $\vec{w} = S - P$. The plane is two-dimensional, so these three vectors must be linearly dependent:

$$\mu\vec{u} + \nu\vec{v} + \omega\vec{w} = \vec{0}$$

Now consider the signs of these three coefficients, μ, ν, and ω. If they all had the same sign, then the point P would be in the convex hull of Q, R, and S: simply divide both sides by $\zeta = \mu + \nu + \omega$ and add P to find the right combination:

$$\begin{aligned}
\frac{\mu\vec{u}+\nu\vec{v}+\omega\vec{w}}{\zeta}+P &= P\\
\frac{\mu\vec{u}+\nu\vec{v}+\omega\vec{w}}{\zeta}+\frac{\zeta}{\zeta}P &= P\\
\frac{\mu\vec{u}+\nu\vec{v}+\omega\vec{w}+\zeta P}{\zeta} &= P\\
\frac{\mu\left(\vec{u}+P\right)+\nu\left(\vec{v}+P\right)+\omega\left(\vec{w}+P\right)}{\zeta} &= P\\
\frac{\mu Q+\nu R+\omega S}{\zeta} &= P\\
\frac{\mu}{\mu+\nu+\omega}Q+\frac{\nu}{\mu+\nu+\omega}R+\frac{\omega}{\mu+\nu+\omega}S &= P
\end{aligned}$$

This final combination is a convex combination not only because $\mu/(\mu+\nu+\omega)+\nu/(\mu+\nu+\omega)+\omega/(\mu+\nu+\omega)=1$, but also because these three values are positive—a result of μ, ν, and ω having the same sign.

Since we have already shown that P does not lie in the convex hull of the other three points, we may assume that μ, ν, and ω have differing signs. This means we will be able to rewrite the original statement of linear dependency. We will use only positive coefficients by shifting terms that would have negative coefficients to the other side. Our equation will now have two terms on one side and one on the other since the three signs are not all the same. If we let μ', ν', and ω' be the positive versions of the coefficients, our equation may look something like this:

$$\nu'\vec{v}=\mu'\vec{u}+\omega'\vec{w}$$

We now claim that—at least in this case—segment $\overline{PR}$ crosses segment $\overline{QS}$. In general, the side of the equation with a single term will correspond to the point opposite P, and the other two points will form the other diagonal. In order to see this, observe that vector addition will put $\vec{v}$ pointing somewhere between the directions of $\vec{u}$ and $\vec{w}$. This makes $\vec{v}$ the middle vector, and makes $\overline{PR}$ a diagonal. It makes sense that the two diagonals should cross, and indeed they do: $\alpha=\mu'\omega'/(\mu'+\omega')$ of the way from P to R and $\beta=\omega'/(\mu'+\omega')$ of the way from Q to S.

In order to complete this half of the proof, it will help if we pin down μ, ν, and ω. An astute reader will note that they are not uniquely defined—in fact, any multiple of their values would work. Therefore, let us assume that the product of the two coefficients that were on the same side equals the third. In this case, $\mu\omega=\nu$. We can do this safely because the left side gives quadratic growth over the right side, which grows linearly.

$$\begin{aligned}
\mu\omega &= \nu \\
\mu\omega\vec{v} &= \nu\vec{v} \\
\mu P + \omega P + \mu\omega\vec{v} &= \mu P + \omega P + (\mu\vec{u} + \omega\vec{w}) \\
(\mu + \omega - \mu\omega) P + \mu\omega(\vec{v} + P) &= \mu(\vec{u} + P) + \omega(\vec{w} + P) \\
(\mu + \omega - \mu\omega) P + \mu\omega R &= \mu Q + \omega S \\
\frac{\mu + \omega - \mu\omega}{\mu + \omega} P + \frac{\mu\omega}{\mu + \omega} R &= \frac{\mu}{\mu + \omega} Q + \frac{\omega}{\mu + \omega} S \\
(1 - \alpha) P + \alpha R &= (1 - \beta) Q + \beta S
\end{aligned}$$

This proves that $\overline{PR}$ and $\overline{QS}$ cross. An assiduous reader will want to check where exactly we used Axiom 2.10.

$\Leftarrow$: Assuming there are exactly $\binom{n}{4}$ crosses, we must now show that $\mathcal{G}$ is convex.

The cardinality of $\mathcal{Q}$ is of course $\binom{n}{4}$, where n is the size of $\mathcal{G}$. This means that every quadruple of points in $\mathcal{G}$ must be convex if we are to have enough crosses in total. We will prove the contrapositive. Assume $\mathcal{G}$ is not convex. This means that there is a point in $\mathcal{G}$ that can be removed without changing the convex hull of $\mathcal{G}$. Call this point P. We will show that P is part of a quadruple in $\mathcal{G}$ that has no cross, thus proving that there are not as many as $\binom{n}{4}$ crosses.

Let $\{G_1, G_2, G_3, ...\}$ be the other points in $\mathcal{G}$, that is, $\mathcal{G} \setminus \{P\}$. Now, by the definition of convex hull, we have nonnegative β_i such that:

$$P = \sum_{i=1}^{n-1} \beta_i G_i$$

Because we are in two dimensions, we know that any set of three points has $\vec{0}$ as one of its linear combinations (this follows from linear dependence). Therefore, as long as we can find three non-zero coefficients, we can shift them around until one of them becomes 0. We have to be careful, though: the sum of the β_i must always be 1, which means the sum of the remaining coefficients must be non-zero in order to pick up the slack.

Therefore: we can reduce all but three of the β_i to 0, leaving P inside the convex hull of three other points. This set of four points will not have a crossing, and we are done. □

Now we may safely consider convexity to be a combinatorial property in addition to a geometrical one.

Theorem 3.2. *A point P in the affine plane is inside a convex configuration $\mathcal{G}$ with more than three points iff*

(1) $\mathcal{G} \cup \{P\}$ is non-convex and
(2) there are at most two segments connecting P to a point in $\mathcal{G}$ that completely avoid $\mathcal{G}$.

Proof. This proof will be less formal than the last, and correspondingly clearer.

$\Rightarrow$: Imagine a convex polygon with every diagonal drawn in: lines are crossing every which way, giving the impression of an enormously complicated spider web. We must place a point inside this polygon, connecting it to each vertex. If we place our point (red) very near to a side of the polygon, then the two segments to the two nearest vertices will completely avoid the configuration. This, however, is the best we can do. There is no way to have three segments avoid crossings because a fourth point—$\mathcal{G}$ has more than three—will be separated from the first by segments to the second and third.

Furthermore, the resulting configuration will not be convex because removing the point you just added will not alter the convex hull.

$\Leftarrow$: This time, let us try to place a point outside the polygon. We have two options. The first is to place our point (blue) close to one of the sides of the polygon. However, this fails because the new configuration is still convex. We may try to place a point (green) so that some other point is swallowed, but ours remains on the edge. This fails as well, because such a point will have at least three segments coming out of it that completely avoid the original configuration.

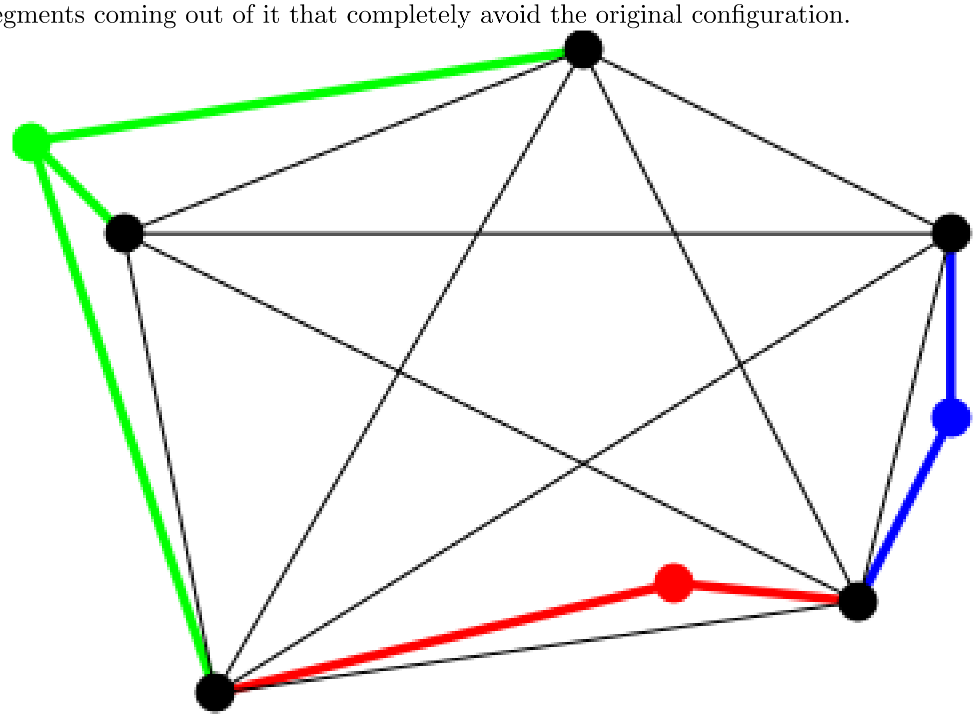

□

4. A conjecture

Conjecture 4.1. *A combinatorial configuration of a complete graph can be realized geometrically iff*

(1) any segment that begins outside a convex subconfiguration and ends inside that configuration must not completely avoid that configuration and
(2) any subconfiguration with exactly five points must have at least one crossing.

Although the above conjecture is probably false, the process of finding a "good characterization" usually begins with several conditions that must be met. Later, when a counterexample is found, the conditions can be strengthened until a "good characterization" is reached.

References

[1] J. P. May. A Concise Course in Algebraic Topology. University of Chicago Press. 1999.