Levi's Lemma, pseudolinear drawings of K_n, and empty triangles

Alan Arroyo[*+], Dan McQuillan[±],
R. Bruce Richter[†+], and Gelasio Salazar[*×]

LaTeX-ed: November 20, 2015

Abstract

There are three main thrusts to this article: a new proof of Levi's Enlargement Lemma for pseudoline arrangements in the real projective plane; a new characterization of pseudolinear drawings of the complete graph; and proofs that pseudolinear and convex drawings of K_n have $n^2 + \mathrm{O}(n \log n)$ and $\mathrm{O}(n^2)$, respectively, empty triangles. All the arguments are elementary, algorithmic, and self-contained.

AMS Subject Classification Primary 52C30; Secondary 05C10, 68R10

1 Introduction

The Harary-Hill Conjecture asserts that the crossing number of the complete graph K_n is equal to

$$H(n) := \frac{1}{4}\left\lfloor\frac{n}{2}\right\rfloor\left\lfloor\frac{n-1}{2}\right\rfloor\left\lfloor\frac{n-2}{2}\right\rfloor\left\lfloor\frac{n-3}{2}\right\rfloor .$$

The work of Ábrego et al [2] verifies this conjecture for "shellable" drawings of K_n; this is one of the first works that identifies a topological, as opposed to geometric, criterion for a drawing to have at least $H(n)$ crossings.

Throughout this work, all drawings of graphs are *good drawings*: no two edges incident with a common vertex cross; no three edges cross at a common point; and no two edges cross each other more than once.

It is well-known that the *rectilinear* crossing number (all edges are required to be straight-line segments) of K_n is, for $n \geq 10$, strictly larger than $H(n)$. In fact, this applies to the more general *pseudolinear* crossing number.

[*]Supported by CONACYT.
[†]Supported by NSERC.
[+]University of Waterloo, [±]Norwich University, and [×]UASLP

An *arrangement of pseudolines* Σ is a finite set of simple open arcs in the plane $\mathbb{R}^2$ such that: for each $\sigma \in \Sigma$, $\mathbb{R}^2 \setminus \sigma$ is not connected; and for distinct σ and σ' in Σ, $\sigma \cap \sigma'$ consists of a single point, which is a crossing.

A drawing of K_n is *pseudolinear* if there is an arrangement of Σ of pseudolines such that each edge of K_n is contained in one of the pseudolines and each pseudoline contains just one edge. It is clear that a rectilinear drawing (chosen so no two lines are parallel) is pseudolinear.

The arguments (originally due to Lovász et al [13] and, independently, Ábrego and Fernández-Merchant [1]) that show every rectilinear drawing of K_n has at least $H(n)$ crossings apply equally well to pseudolinear drawings.

The proof that every optimal pseudolinear drawing of K_n has its outer face bounded by a triangle [5] uses the "allowable sequence" characterization of pseudoline arrangements of Goodman and Pollack [8]. Our principal result is that there is another, topological, characterization of pseudolinear drawings of K_n.

Let D be a drawing of K_n in the sphere. For any three distinct vertices u, v, w of K_n, the triangle T induced by u, v, w is such that $D[T]$ (the subdrawing of D induced by the subgraph T) is a simple closed curve in the sphere.

This simple observation leads to the natural ideas of a convex drawing of K_n and a face-convex drawing of K_n, which capture at different levels of generality the notion of a convex set in Euclidean space.

Definition 1.1 *Let D be a drawing of K_n in the sphere.*

1. *Let T be a 3-cycle in K_n. Then a closed disc Δ bounded by $D[T]$ is* convex *if, for any distinct vertices u and v of K_n such that both $D[u]$ and $D[v]$ are in Δ, then $D[uv] \subseteq \Delta$.*

2. *The drawing D is* convex *if, for every 3-cycle T in K_n, at least one of the closed discs bounded by $D[T]$ is convex.*

3. *A* face *of D is a component of $\mathbb{R}^2 \setminus D[K_n]$.*

4. *The drawing D is* face-convex *if there is a face F of D such that, for every triangle T of D, the closed disc bounded by $D[T]$ and not containing F is convex. The face F is the* outer face *of D.*

There seem to be interesting connections between convexity and Knuth's CC systems [11], but we have not yet formalized this.

In Definition 1.1 4, there is necessarily at least one outer face that shows the drawing to be face-convex. The unique drawing of K_6 with three crossings has two such faces. (See Figure 1.2.)

It is convenient for the definition of convexity to use drawings in the sphere: every simple closed curve is the boundary of two closed discs. Every drawing in the plane is converted by the standard 1-point compactification into a spherical drawing. Keeping track of the infinite face F in a pseudolinear drawing in the plane results in a face-convex drawing in the sphere with outer face F. The interesting point is the converse: if we convert the face F in the definition of face-convex to be the unbounded face, then the resulting drawing in the plane is pseudolinear.

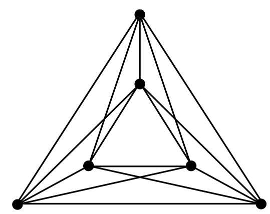

Figure 1.2: The two faces bounded by 3-cycles can each be the outer face.

Theorem 1.3 *A drawing of K_n in the plane is face-convex if and only if it is pseudolinear.*

This theorem is proved in Section 3. An independent recent proof has been found by Aichholzer et al [3]; their proof uses Knuth's CC systems [11] (reinforcing the interest in the connection with convexity), the duals of which are realizable as pseudolinear arrangements of lines. Moreover, their statement is in terms of a forbidden configuration. Properly speaking, their result is of the form, "there exists a face relative to which the forbidden configuration does not occur". Their face and our face are the same. However, our proof is completely different, yielding directly a polynomial time algorithm for finding the pseudolines.

Aichholzer et al show that the there is a pseudolinear drawing of K_n having the same crossing pairs of edges as the given drawing of K_n. Gioan's Theorem [7] that any two drawings of K_n with the same crossing pairs of edges are equivalent up to Reidemeister III moves is then invoked to show that the original drawing is also pseudolinear. Our proof is completely self-contained; in particular, it does not involve CC-systems and does not invoke Gioan's Theorem.

The ideas we use are elementary and derive from a simple, direct proof of Levi's Enlargement Lemma given in Section 2. In a separate paper [4], we give a proof of Gioan's Theorem in the same spirit.

In Section 4, we extend the Bárány and Füredi [6] theorem that a rectilinear drawing of K_n has at least $n^2+O(n\log n)$ empty triangles to pseudolinear drawings of K_n. Moreover, we show that a convex drawing of K_n has at least $n^2/3+O(n)$ empty triangles.

2 Proof of Levi's Enlargement Lemma

In this section, we prove Levi's Enlargement Lemma [12]. This important fact seems to have only one proof by direct geometric methods in English [9]. The proof in [9] includes a simple step that Grünbaum admits seems clumsy. Our proof avoids this technicality. (There is another proof by Sturmfels and Ziegler via oriented matroids [14].)

One fact we do use is that there is an alternative definition of an arrangement of pseudolines. Equivalent to the definition given in the introduction, a *pseudoline* is a non-contractible simple closed curve in the real projective plane, and an arrangement of pseudolines is a set of pseudolines, any two intersecting in exactly one point; the intersection is necessarily a crossing point. This perspective will be used in the proof of Levi's Enlargement Lemma.

Theorem 2.1 (Levi's Enlargement Lemma) *Let Σ be an arrangement of pseudolines and let a,b be any two points in the plane not both in the same pseudoline in Σ. Then there is a pseudoline σ that contains both a and b and such that $\Sigma\cup\{\sigma\}$ is an arrangement of pseudolines.*

The principal ingredients in all our arguments are two considerations of the facial structure of an arrangement of pseudolines. In fact, we need something slightly more general. An *arrangement of arcs* is a finite set Σ of open arcs in the plane $\mathbb{R}^2$ such that, for every $\sigma\in\Sigma$, $\mathbb{R}^2\setminus\sigma$ is not connected and any two elements of Σ have at most one point in common, which must be a crossing. Thus, two arcs in an arrangement of arcs may have no intersection and so be "parallel".

Let Σ be an arrangement of arcs. Set $\mathcal{P}(\Sigma)$ to be the set $\bigcup_{\sigma\in\Sigma}\sigma$ of points in the plane. A *face* of Σ is a component of $\mathbb{R}^2\setminus\mathcal{P}(\Sigma)$. Since Σ is finite, there are only finitely many faces of Σ.

The *dual* Σ^* of Σ is the finite graph whose vertices are the faces of Σ and there is one edge for each segment α of each $\sigma\in\Sigma$ such that α is one of the components of $\sigma\setminus\mathcal{P}(\Sigma\setminus\{\sigma\})$. The dual edge corresponding to α joins the faces of Σ on either side of α.

Levi's Lemma is a consequence of our first consideration of the facial structure of an arrangement of arcs.

Lemma 2.2 (Existence of dual paths) *Let Σ be an arrangement of arcs and let a, b be points of the plane not in any line in Σ. Then there is an ab-path in Σ^* crossing each arc in Σ at most once.*

Proof. We proceed by induction on the number of curves in Σ that separate a from b, the result being trivial if there are none. Otherwise, for $x \in \{a, b\}$, let F_x be the face of Σ containing x and let $\sigma \in \Sigma$ be incident with F_a and separate a from b. Then Σ^* has an edge F_aF that crosses σ.

Let R be the region of $\mathbb{R}^2 \setminus \sigma$ that contains F_b and let Σ' be the set $\{\sigma' \cap R \mid \sigma' \in \Sigma,\ \sigma' \cap R \neq \varnothing\}$. The induction implies there is an FF_b-path in Σ'^*. Together with F_aF, we have an F_aF_b-path in Σ^*, as required. ■

We now turn to the proof of Levi's Lemma.

Proof of Theorem 2.1. In this proof, we view the pseudoline arrangement Σ as non-contractible simple closed curves in the real projective plane, any two intersecting exactly once.

If a is not in any arc in Σ, then let F be the face of Σ containing a; replace a with any point in the boundary of F and not in the intersection of two arcs in Σ. Likewise, for b. In all cases, the points representing a and b are chosen to be in different arcs in Σ.

If we find the required ab-arc σ to extend Σ using one or two replacement points, then σ goes through the face(s) of Σ containing the original point(s), and so we may reroute σ to go through the original points, as required. Thus, we may assume a and b are both in arcs in Σ.

Let Σ_a consist of the arcs in Σ containing a and let $F_b^{(a)}$ be the face of $\mathcal{P}(\Sigma_a)$ containing b. Up to spherical homeomorphisms, there is a unique small arc α through a that has one end in $F_b^{(a)}$ and crosses all the arcs in Σ_a at a. The ends of this arc are the two points a', a''. In a similar way, we get the small arc β through b joining the two points b', b''.

The choices for α and β show that we may label a', a'' and b', b'' so that a' and b' are in the same face F' of $\Sigma_a \cup \Sigma_b$. We apply Lemma 2.2 to this component to obtain an $a'b'$-arc γ' contained in F'.

The arc composed of γ' together with the little arcs α and β crosses every arc in $\Sigma_a \cup \Sigma_b$ exactly once. This shows that a'' and b'' are in the same face F'' of $\Sigma_a \cup \Sigma_b$.

Let Σ'' be the set $\{\sigma \cap F'' \mid \sigma \in \Sigma\, \sigma \cap F'' \neq \varnothing\}$. Lemma 2.2 implies there is an $a''b''$-arc γ'' in F'' crossing each element of Σ'' at most once.

Let γ be the closed curve $\gamma' \cup \alpha \cup \gamma'' \cup \beta$, adjusted as necessary near a', a'', b', and b'' so that γ is a simple closed curve. It is clear that γ crosses each arc in $\Sigma_a \cup \Sigma_b$ exactly once and, therefore, is non-contractible. By construction, γ crosses any arc in Σ at most twice; both being non-contractible implies this is in fact at most once. Therefore, $\Sigma \cup \{\gamma\}$ is the desired arrangement of pseudolines. ■

3 Proof of Theorem 1.3

In this section we prove Theorem 1.3: a face-convex drawing of K_n in the sphere with outer face F is a pseudolinear drawing in the plane by making F the infinite face.

It is evident that face-convexity is inherited in the sense that if D is a face-convex drawing of K_n and v is any vertex of K_n, then $D[K_n - v]$ is a face-convex drawing of $K_n - v$. We begin with a simple observation.

Lemma 3.1 *Let D be a face-convex drawing of K_n with outer face F. If J is any K_4 in K_n such that $D[J]$ has a crossing, then F is in the face of $D[J]$ bounded by a 4-cycle of J. In particular, no crossing of D is incident with F, so F is bounded by a cycle in K_n.*

Proof. Let v, w, x, y be the four vertices of J labelled so that vw crosses xy in D. Consider, for example, the triangle $T = (v, w, x, v)$. The vertex y is in a closed face F_y of $D[T]$. Since xy crosses vw, $D[xy]$ is not contained in F_y, so F_y is not convex. Since D is face-convex, it follows that $F \subseteq F_y$.

Thus, none of v, w, x, y is on the convex side of the triangle containing the other three vertices. It follows that F is contained in the face of $D[J]$ bounded by the 4-cycle (v, x, w, y, v), as required.

Inserting a vertex at every crossing point of D produces a 2-connected planar embedding of the resulting graph having F as a face. This face is bounded by a cycle; since no inserted vertex is incident with F, this cycle is a cycle of K_n. ■

We remark that Lemma 3.1 shows that a face-convex drawing does not have the forbidden configuration of Aichholzer et al [3]. The converse is no harder.

For a face-convex drawing D of K_n with outer face F, let C_F denote the cycle of K_n bounding F and let Δ_F denote the closed disc bounded by C_F and disjoint from F. For any subset W of vertices of K_n, let $D[W]$ denote the subdrawing of D induced by the complete subgraph having precisely the vertices in W. Since $D[W]$ is a face-convex drawing, if $|W| \geq 3$, then its face F_W containing F is bounded by a cycle C_W. The closed disc Δ_W bounded by $D[C_W]$ and disjoint from F is the *convex hull of W*.

For each edge uv of G, $D[uv]$ is a simple arc in the sphere. Arbitrarily giving $D[uv]$ a direction distinguishes a left and right side to the arc $D[uv]$. We prefer not to use the labels 'left' and 'right', as we find them somewhat confusing. For now, we shall arbitrarily label them as *side 1* and *side 2* of uv.

For each vertex w different from u and v, *w is on side i of uv* if the face of $D[\{u,v,w\}]$ disjoint from F is on side i of uv. We set Σ^i_{uv} to be the set of vertices on side i of uv; for convenience, we include u and v in Σ^i_{uv}.

It is clear that $\Sigma^1_{uv} \cap \Sigma^2_{uv} = \{u,v\}$. What is less clear is that $D[\Sigma^1_{uv}] \cap D[\Sigma^2_{uv}]$ consists just of u, v, and uv. The next lemma is a useful step in proving this.

Lemma 3.2 *Let D be a face-convex drawing of K_n with outer face F and let u,v,x,y be distinct vertices of K_n.*

- **(3.2.1)** *Then x and y are on the same side of uv if and only if uv is incident with $F_{\{u,v,x,y\}}$.*
- **(3.2.2)** *In particular, if x and y are on different sides of uv, then $D[\{u,v,x,y\}] - xy$ has no crossing.*
- **(3.2.3)** *If z is any vertex such that u is in the interior of $\Delta_{\{x,y,z\}}$, then some two of x, y, and z are on different sides of uv.*

Proof. Ultimately, the easiest way to understand (3.2.1) is to draw the two possible drawings of K_4 and, in both cases, check the two possibilities: uv is incident with $F_{\{u,v,x,y\}}$ and uv is not incident with $F_{\{u,v,x,y\}}$. In the case the K_4 has a crossing, $F_{\{u,v,x,y\}}$ is the face bounded by the 4-cycle.

For (3.2.2), let J be the K_4 induced by u,v,x,y. Since x and y are on different sides of uv, the preceding conclusion shows that either $D[J]$ has no crossing, in which case we are done, or uv is crossed in $D[J]$ and it crosses xy. As this is the only crossing in $D[J]$, $D[J] - xy$ has no crossing.

Finally, we consider (3.2.3). If $D[v] \notin \Delta_{\{x,y,z\}}$, then $D[uv]$ crosses the 3-cycle xyz. Now (3.2.1) shows that the ends of the edge crossing $D[uv]$ are on different sides of uv. Thus, we may assume $D[v] \in \Delta_{\{x,y,z\}}$.

Since $D[u] \subseteq \Delta_{\{x,y,z\}}$, $D[\{ux,uy,uz\}] \subseteq \Delta_{\{x,y,z\}}$. If $v = z$, then $D[uv]$ is not incident with $F_{\{u,v,x,y\}}$. Therefore, (3.2.1) shows x and y are on different sides of uv and consequently, we may assume $v \neq z$.

By definition, $\Delta_{\{u,x,y\}} \subseteq \Delta_{\{x,y,z\}}$, and likewise for $\Delta_{\{u,x,z\}}$ and $\Delta_{\{u,y,z\}}$. We may choose the labelling of x, y, and z so that $D[v] \in \Delta_{\{u,x,y\}}$. But now uv is not in the boundary of $D[\{u,v,x,y\}]$. Again, (3.2.1) shows x and y are on different sides of uv. ■

We are now ready for the first significant step, which is Item (3.3.4) in our next result.

Lemma 3.3 *Let D be a face-convex drawing of K_n with outer face F, let $W \subseteq V(K_n)$, and let uv be any edge of K_n.*

- **(3.3.1)** *Both $(W \cap \Sigma^1_{uv}) \setminus \{u,v\}$ and $(W \cap \Sigma^2_{uv}) \setminus \{u,v\}$ are not empty if and only if uv is not incident with $F_{W\cup\{u,v\}}$.*
- **(3.3.2)** *For $\{i,j\} = \{1,2\}$, no vertex of $\Sigma^i_{uv} \setminus \{u,v\}$ is in $\Delta_{(W\cup\{u,v\})\cap\Sigma^j_{uv}}$.*
- **(3.3.3)** *If, for $i = 1,2$, $x_i, y_i \in \Sigma^i_{uv}$, then x_1y_1 does not cross x_2y_2 in D.*
- **(3.3.4)** $\Delta_{\Sigma^1_{uv}} \cap \Delta_{\Sigma^2_{uv}} = D[\{u,v\}]$.

Proof. Suppose uv is incident with $F_{W\cup\{u,v\}}$. For any $x, y \in W \setminus \{u,v\}$, it follows that uv is incident with $F_{\{u,v,x,y\}}$, so Lemma (3.2.1) shows x and y are on the same side of uv.

Conversely, suppose all vertices in $W \setminus \{u,v\}$ are on the same side of uv. The closed disc $\Delta_{W\cup\{u,v\}}$ is the union of all the convex sides $\Delta_{\{x,y,z\}}$, for $x, y, z \in W \cup \{u,v\}$. If u is in the interior of some $\Delta_{\{x,y,z\}}$, then Lemma (3.2.3) shows some two of x and y are on different sides of uv. Thus, both u and v are in $C_{W\cup\{u,v\}}$. If $uv \notin E(C_{W\cup\{u,v\}})$, then $C_{W\cup\{u,v\}} - \{u,v\}$ is not connected; let x and y be in different components of $C_{W\cup\{u,v\}} - \{u,v\}$. Then $D[xy]$ crosses $D[uv]$, showing x and y are on different sides of uv. This contradiction completes the proof of (3.3.1).

For $i = 1,2$, let $W_i = (W \cup \{u,v\}) \cap \Sigma^i_{uv}$.

For (3.3.2), suppose $x \in \Sigma^i_{uv} \setminus \{u,v\}$ is in Δ_{W_j}. Since $W_j \subseteq W_j \cup \{x\}$, $\Delta_{W_j} \subseteq \Delta_{W_j\cup\{x\}}$. Since $C_{W_j\cup\{x\}}$ either contains x, in which case $D[x] \notin \Delta_{W_j}$, or is C_{W_j}, in which case $D[x]$ is in the interior of Δ_{W_j}.

Assume by way of contradiction that it is the latter case. Then $C_{W_j\cup\{x\}} = C_{W_j}$. Therefore, $\Delta_{W_j\cup\{x\}} = \Delta_{W_j}$. Since uv is in C_{W_j} by (3.3.1), the other direction of (3.3.1) implies the contradiction that $x \in \Sigma^j_{uv}$.

For (3.3.3), we suppose x_1y_1 and x_2y_2 cross in D. From (3.2.2), not both $\{x_1,y_1\}$ and $\{x_2,y_2\}$ can contain an element of $\{u,v\}$. We may choose the labelling so that $\{x_1,y_1\} \cap \{u,v\} = \varnothing$ and let J_1 be the K_4 induced by u, v, x_1, y_1. Since $\{x_2,y_2\} \neq \{u,v\}$, we may assume $x_2 \notin \{u,v\}$.

Claim 1 $y_2 \notin \{u,v\}$.

Proof. Suppose $y_2 \in \{u,v\}$. Apply (3.2.2) to each of the K_4's induced by u, v, x_2, x_1 and u, v, x_2, y_1. The conclusion is that x_2y_2 does not cross $D[J_1] - x_1y_1$. Thus, as we follow $D[x_2y_2]$ from $D[x_2]$, its first and only intersection with $D[J_1]$ is with $D[x_1y_1]$, showing x_1y_1 is incident with the face F_{J_1}. Since uv is also incident with F_{J_1}, we deduce that $D[J_1]$ is a crossing K_4.

However, continuing on to the end $y_2 \in \{u, v\}$, $D[x_2y_2]$ must cross C_{J_1} without crossing any edge of J_1, which is impossible, as required. $\square$

Let J_2 be the K_4 induced by u, v, x_2, y_2. Lemma (3.2.2) and Claim 1 show that the only possible crossing between $D[J_1]$ and $D[J_2]$ is the crossing of x_1y_1 with x_2y_2. However, both x_2 and y_2 are in F_{J_1}, showing that x_2y_2 must cross C_{J_1} an even number of times. As there is at least one crossing and all the crossings are with x_1y_1, we violate the requirement that, in a drawing, no two edges cross more than once.

Now for (3.3.4). From (3.3.1), no vertex of one side is inside the convex hull of the other side. Going one step further, suppose $x, y \in (W \cup \{u, v\}) \cap \Sigma^2_{uv}$ is such that $D[xy]$ has a point that is in $\Delta_{(W\cup\{u,v\})\cap\Sigma^1_{uv}}$. Then xy crosses some edge of $C_{\Sigma^1_{uv}}$, contradicting (3.3.3).

Finally, we show that $\Delta_{\Sigma^1_{uv}} \cap \Delta_{\Sigma^2_{uv}} = D[\{u, v\}]$. The cycles $C_{\Sigma^1_{uv}}$ and $C_{\Sigma^2_{uv}}$ are disjoint except for uv. If there is some point a of the sphere in $\Delta_{\Sigma^1_{uv}} \cap \Delta_{\Sigma^2_{uv}}$ that is not in uv, then a is in the convex hull of both $C_{\Sigma^1_{uv}}$ and $C_{\Sigma^2_{uv}}$. This implies that either $\Delta_{\Sigma^1_{uv}} \subseteq \Delta_{\Sigma^2_{uv}}$ or $\Delta_{\Sigma^2_{uv}} \subseteq \Delta_{\Sigma^1_{uv}}$, contradicting (3.3.1). $\blacksquare$

It follows from the above that, for every edge uv, $\Delta_{\Sigma^1_{uv}} \cup \Delta_{\Sigma^2_{uv}}$ includes all the vertices of K_n and all edges that have both ends in the same one of Σ^1_{uv} and Σ^2_{uv}. We obtain a more refined understanding of the relationship of this subdrawing with the entire drawing in the following.

Lemma 3.4 *Let D be a face-convex drawing of K_n with outer face F and let uv be any edge of K_n. Let W be any subset of $V(K_n)$. Then there are not four distinct vertices x_1, x_2, y_1, y_2 of C_W appearing in this cyclic order in C_W such that, for $i = 1, 2$, $x_i, y_i \in \Sigma^i_{uv}$.*

Proof. If such vertices exist, then the edges x_1y_1 and x_2y_2 are both in Δ_W and they cross, contradicting Lemma (3.3.3). $\blacksquare$

It follows from Lemma 3.4 that, for a face-convex drawing of K_n with outer face F, C_F has, for $i = 1, 2$, a path (possibly with no vertices; this happens only when uv is in C_F) contained in $\Sigma^i_{uv} \setminus \{u, v\}$. The ends of these paths are connected in C_F either by an edge or by a path of length 2, the middle vertex being one of u and v.

Henceforth, we assume $uv \notin E(C_F)$; that is, we assume both $\Sigma^1_{uv} \setminus \{u, v\}$ and $\Sigma^2_{uv} \setminus \{u, v\}$ are both non-empty. In this case, $C_F \cup C_{\Sigma^1_{uv}} \cup C_{\Sigma^2_{uv}}$ is a planar embedding of a 2-connected graph. Three of its faces are F, $\Delta_{\Sigma^1_{uv}}$, and $\Delta_{\Sigma^2_{uv}}$. The other faces, if any, are determined by whether or not u or v is in C_F.

Let A_{uv} consist of the ones of u and v not in C_F. For each $a \in A_{uv}$, there is a face F^a_{uv} of $C_F \cup C_{\Sigma^1_{uv}} \cup C_{\Sigma^2_{uv}}$ incident with a and an edge f^a_{uv} of C_F; the edge f^a_{uv} is incident with a vertex in Σ^1_{uv} and a vertex in Σ^2_{uv}. Let Q^a_{uv} be the cycle bounding F^a_{uv}. See Figure 3.5.

For $j = 1, 2$, $C_{\Sigma^j_{uv}}$ is the union of two internally disjoint paths, namely $C_F \cap C_{\Sigma^j_{uv}}$ and the path P^j_{uv} in $C_{\Sigma^j_{uv}}$ having its ends in C_F but otherwise disjoint from C_F. If a is in $\{u, v\} \setminus A_{uv}$, then a is in C_F and, therefore, is an end of both P^1_{uv} and P^2_{uv}. If $a \in A_{uv}$, then one end of f^a_{uv} is an end of P^1_{uv} and the other end of f^a_{uv} is an end of P^2_{uv}.

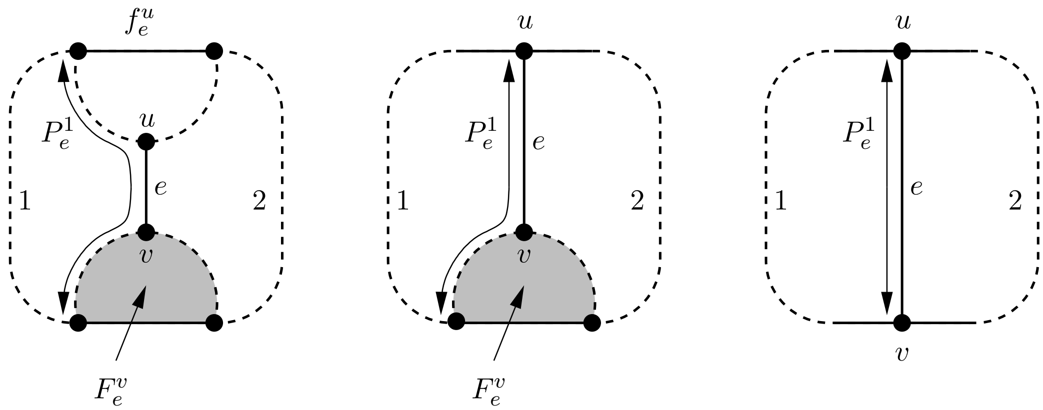

Figure 3.5: In the left-hand figure, $A_e = \{u, v\}$, in the middle $A_e = \{v\}$, and in the right $A_e = \varnothing$.

We are now prepared to prove our characterization of pseudolinear drawings. Recall that Δ_F is the closed disc bounded by C_F that is disjoint from F.

Proof of Theorem 1.3. We begin by finding, for each edge e that is not in C_F, an arc α_e such that:

(I) α_e consists of three parts, namely $D[e]$, and, for each vertex u incident with e, a subarc α^u_e, which is either just u, if $u \in V(C_F)$, or an arc in F^u_e joining u to a point in f^a_e and otherwise disjoint from Q^u_e;

(II) if α_e crosses an edge e' (including possibly $e' \in E(C_F)$), then e' has an incident vertex in each of Σ^1_e and Σ^2_e; and

(III) for any other edge e' not in C_F, α_e and $\alpha_{e'}$ intersect at most once, and if they intersect, the intersection is a crossing point.

Arbitrarily order the edges of K_n not in C_F as $e_1, \dots, e_r$. We suppose $i \geq 1$ and we have $\alpha_{e_1}, \dots, \alpha_{e_{i-1}}$ satisfying Items (I) – (III). We show there is an arc α_{e_i} such that $\alpha_{e_1}, \dots, \alpha_{e_i}$ also satisfy Items (I) – (III). Let $e_i = uv$.

Since e_i is not in C_F, $D[e_i]$ is inside Δ_F. In C_F there are vertices on each side of e_i.

Useful Fact: *Let $j \in \{1, 2, \dots, i-1\}$. By (II) and Lemma 3.4, α_{e_j} crosses each of $P^1_{e_i}$ and $P^2_{e_i}$ at most twice.*

Since part of the extension of $D[e_i]$ to α_{e_i} is trivial if either u or v is in C_F, we will generally proceed below as though neither u nor v is in C_F. When there is a subtlety in the event u or v is in C_F, we will specifically mention it.

We can apply Lemma 2.2 in the interior of $F^u_{e_i}$ and $F^v_{e_i}$ to extend e_i in both directions to points on (actually very near) $f^u_{e_i}$ and $f^v_{e_i}$ to create a possible α_{e_i}. These are all equivalent up to Reidemeister moves and any one is a potential solution. We let Λ_i denote the set of these dual path solutions.

For $j = 1, 2, \dots, i-1$, the segment α_{e_j} is *unavoidable for e_i* if α_{e_j} crosses both the paths $P^1_{e_i}$ and $P^2_{e_i}$. In particular, α_{e_j} is unavoidable if it crosses e_i.

It may be that e_j is incident with one of u and v, for example. As this forces a crossing of α_{e_j} with α_{e_i}, we take this as a crossing of both $P^1_{e_i}$ or $P^2_{e_i}$. On the other hand, if e_j is incident with an end w of $f^u_{e_i}$, then this constitutes a crossing of α_{e_j} with the one of $P^1_{e_i}$ and $P^2_{e_i}$ that contains w.

Claim 1 *For $j \in \{1, 2, \dots, i-1\}$, α_{e_j} is unavoidable for e_i if and only if every arc in Λ_i crosses α_{e_j}.*

Proof. Suppose first that α_{e_j} is unavoidable for e_i. If α_{e_j} has a point in $D[e_i]$, then evidently α_{e_j} crosses every solution in Λ_i.

In the case α_{e_j} is disjoint from the closed arc $D[e_i]$, there is some subarc of α_{e_j} with an end in each of $P^1_{e_i}$ and $P^2_{e_i}$, but otherwise disjoint from $P^1_{e_i} \cup P^2_{e_i}$. This arc must join two points in either $Q^u_{e_i}$ or $Q^v_{e_i}$. It is clear that every solution in Λ_i must cross this arc, as required.

Conversely, if α_{e_j} is not unavoidable, then it does not cross, say, $P^1_{e_i}$. In this case, there is a solution in Λ_i whose extensions of e_i go just inside $F^u_{e_i}$ and $F^v_{e_i}$, in both cases very close to $P^1_{e_i}$. This solution does not cross α_{e_j}. $\square$

Suppose that α_{e_j} is unavoidable for e_i and suppose there is an end a_j of α_{e_j} in $f^u_{e_i}$. Following α_{e_j} from a_j, we come to a crossing of, say $P^1_{e_i} \cap Q^u_{e_i}$. The segment of $f^u_{e_i}$ from its end $u^1_{e_i}$ in $P^1_{e_i}$ to a_j is *restricted for α_{e_i}*. We do not want α_{e_i} to cross α_{e_j} on this end segment of α_{e_j}, since they must cross elsewhere.

It may be that the portion of α_{e_j} from a_j to its first intersection in $P^1_{e_i} \cup P^2_{e_i}$ meets $P^1_{e_i} \cup P^2_{e_i}$ at u. In particular, u is an end of e_j. In this case, it is not immediately clear what the restriction should be. The other end of e_j is either in $\Sigma^1_{e_i}$ or $\Sigma^2_{e_i}$, so, correspondingly,

$D[e_j] \subseteq \Delta_{\Sigma^1_{e_i}}$ or $D[e_j] \subseteq \Delta_{\Sigma^2_{e_i}}$. As α_{e_i} must be made to cross α_{e_j} at their intersection u, only in the case $D[e_j] \subseteq \Delta_{\Sigma^2_{e_i}}$ do we get a restriction between $u^1_{e_i}$ and a_j. (In the other case, as in the next paragraph, the restriction is between $u^2_{e_i}$ and a_j.)

There is a completely analogous restriction from a_j to the other end $u^2_{e_i}$ of $f^u_{e_i}$ in $P^2_{e_i}$ if, traversing α_{e_j} from a_j, α_{e_j} first meets $P^2_{e_i}$.

Let R^1_u be the union of all the e_j-restricted portions, $j = 1, 2, \ldots, i-1$, of $f^u_{e_i}$ that contain the end $u^1_{e_i}$ of $f^u_{e_i}$ and let R^2_u be the union of all the restricted portions of $f^u_{e_i}$ that contain the other end $u^2_{e_i}$ of $f^u_{e_i}$.

If u is in C_F, then the u portion of α_{e_i} is just u and no extension at this end is required. The restrictions are required in the case u is not in C_F, the subject of the next claim.

Claim 2 *If u is not in C_F, then $R^1_u \cap R^2_u = \varnothing$.*

Proof. If the intersection is not empty, then there exist $j, j' \in \{1, 2, \ldots, i-1\}$ such that:

1. α_{e_j} proceeds from a_j in $f^u_{e_i}$ to $P^1_{e_i}$;
2. $\alpha_{e_{j'}}$ proceeds from $a_{j'}$ in $F^u_{e_i}$ to $P^2_{e_i}$; and,
3. in $f^u_{e_i}$, $a_{j'}$ is not further from $u^1_{e_i}$ than a_j is.

In particular, $\sigma^1_{u,j}$ and $\sigma^2_{u,j'}$ must cross in $F^u_{e_i}$, so they never cross again. (This is true even if the crossing is $a_j = a_{j'}$. It turns out that $a_j = a_{j'}$ does not occur in our construction, but we do not need this fact, so we do not use it.)

As we traverse α_{e_j} beginning at a_j, we first cross $P^1_{e_i}$ at $\times^1_{j,1}$ in $P^1_{e_i} \cap Q^u_{e_i}$. Since α_j is unavoidable, it must cross $P^2_{e_i}$ for the first time at the point $\times^2_{j,1}$. Between $\times^1_{j,1}$ and $\times^2_{j,1}$, there is a second crossing $\times^1_{j,2}$ with $P^1_{e_i}$; possibly $\times^1_{j,2} = \times^2_{j,1}$. The Useful Fact implies these are no other crossings of α_{e_j} with $P^1_{e_i}$. (In $\times^k_{(r,s)}$, the exponent k refers to which $P^k_{e_i}$ is being crossed; the subscripts (r, s) are indicating which arc α_{e_r} is under consideration and, for $s \in \{1, 2\}$, it is the s^{th} crossing of α_{e_r} with $P^k_{e_i}$ as we traverse α_{e_r} from a_r.)

We claim that the second crossing $\times^1_{j,2}$ of α_j with $P^1_{e_i}$ cannot be in the segment of $P^1_{e_i}$ between $u^1_{e_i}$ and $\times^1_{j,1}$. To see this, suppose $\times^1_{j,2}$ is in this segment; let σ_j be the segment of α_{e_j} from $\times^1_{j,2}$ to the other end a'_j. The Useful Fact and the non-self-crossing of α_{e_j} imply that σ_j is trapped inside the subregion of $F^u_{e_i}$ incident with $u^1_{e_i}$ and the segment of α_{e_j} from a_j to $\times^1_{j,1}$. The only place a'_j can be is in $f^u_{e_i}$, contradicting (I).

A very similar argument shows that $\alpha_{e_{j'}}$ cannot cross that same segment of $P^2_{e_i}$. (Such a crossing would be the second of $\alpha_{e_{j'}}$ with $P^2_{e_i}$. Thus, the other end $a'_{j'}$ of $\alpha_{e_{j'}}$ would also be in $f^u_{e_i}$.)

We conclude that $\times^1_{j,2}$ is in $P^1_{e_i}$ between $\times^1_{j,1}$ and the other end $v^1_{e_i}$ of $P^1_{e_i}$.

Since $\alpha_{e_{j'}}$ is unavoidable, as we traverse it from $a_{j'}$ in $f^u_{e_i}$, there is a first crossing $\times^1_{j',1}$ of $\alpha_{e_{j'}}$ with $P^1_{e_i}$. Between $a_{j'}$ and $\times^1_{j',1}$, there are the two crossings $\times^2_{j',1}$ and $\times^2_{j',2}$ of $\alpha_{e_{j'}}$ with $P^2_{e_i}$; possibly $\times^2_{j',2} = \times^1_{j',1}$. Note that the Useful Fact implies $\alpha_{e_{j'}}$ is disjoint from $C_F \cap C_{\Sigma^2_{e_i}}$.

Let γ be the simple closed curve consisting of the portion of α_{e_j} from a_j to $\times^2_{j,1}$, and then the portion of $C_{\Sigma^2_{uv}}$ from $\times^2_{j,1}$ to $v^2_{e_i}$ and along $C_F \cap C_{\Sigma^2_{uv}}$ to $u^2_{e_i}$, and then the portion of $f^u_{e_i}$ from $u^2_{e_i}$ back to a_j.

From $\times^2_{j',1}$ to the other end $a_{j'}$, $\alpha_{e_{j'}}$ must cross γ. The only segment it can cross is the portion of $P^2_{e_i}$ between $\times^2_{j,1}$ and $v^2_{e_i}$. This implies that $\times^2_{j',2}$ is between $\times^2_{j,1}$ and $v^2_{e_i}$ in $P^2_{e_i}$.

Reversing the roles of j and j' and of sides 1 and 2, we conclude that the preceding argument shows that $\times^1_{j,2}$ is between $\times^1_{j',1}$ and $v^1_{e_i}$ in $P^1_{e_i}$.

The simple closed curve γ above is crossed by $\alpha_{e_{j'}}$ at the point $\times^2_{j',2}$ in the segment of $P^2_{e_i}$ between $\times^2_{j,1}$ and $v^2_{e_i}$. See Figure 3.6. On the other hand, $\times^1_{j',1}$ is on the segment of $P^1_{e_i}$ between the two points $\times^1_{j,1}$ and $\times^1_{j,2}$ and so is on the other side of γ. This shows that $\alpha_{e_{j'}}$ must cross γ again and this is impossible. $\square$

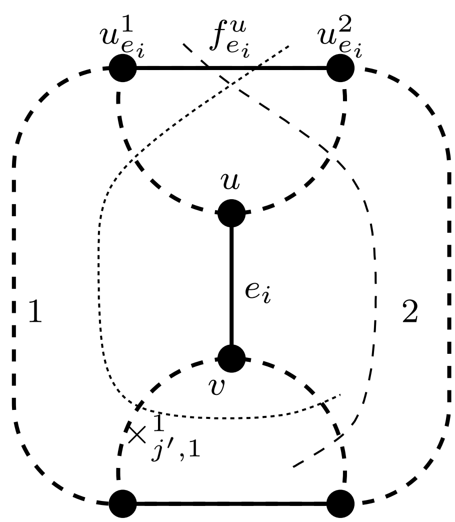

Figure 3.6: One instance of overlapping restrictions. There is no way for $\alpha_{e_{j'}}$ to get to $\times^1_{j',1}$.

If $f^u_{e_i}$ does not exist, then set $\eta_u = u$. Otherwise, let ρ_u be either $u^1_{e_i}$ or the point of R^1_u furthest from $u^1_{e_i}$. Then η_u is any point between ρ_u and the next point between ρ_u and $u^2_{e_i}$ that is an end of some α_j, for $j \in \{1, 2, \dots, i-1\}$. Likewise, η_v is any point of $f^v_{e_i}$ between the last point ρ_v of R^2_v and the next point between ρ_u and $v^1_{e_1}$ that is an end of some α_j, for $j \in \{1, 2, \dots, i-1\}$. (Notice that we use the $P^1_{e_i}$-side restrictions at the

"u-end" and the $P^2_{e_i}$-side restrictions at the v-end. We could have equally well used the $P^2_{e_i}$-side restrictions at the u-end and the $P^1_{e_i}$-restrictions at the v-end.)

We now apply Lemma 2.2 to the region $F^u_{e_i}$ (if it exists) using u and η_u as ends to be connected. We do the same thing in $F^v_{e_i}$ joining v and η_v. These arcs together with $D[e_i]$ give us α_{e_i}, as described in (I). (See Figure 3.7.)

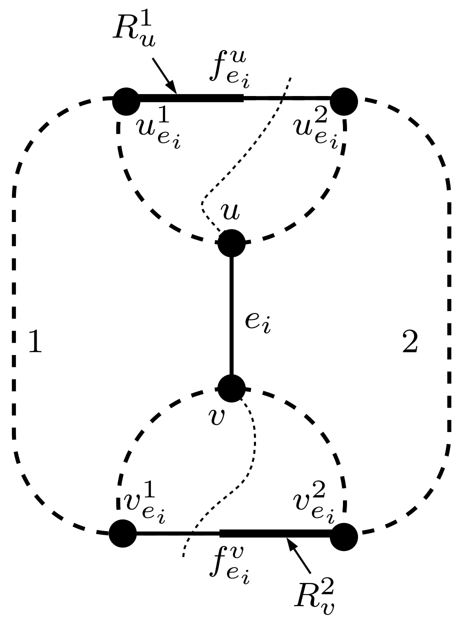

Figure 3.7: The arc α_{e_i}.

The construction of α_{e_i} makes it clear that α_{e_i} meets each of $\Delta_{\Sigma^1_{e_i}}$ and $\Delta_{\Sigma^2_{e_i}}$ in e_i. Therefore, α_{e_i} satisfies (II).

Claim 3 *For any $j \in \{1, 2, \ldots, i-1\}$ for which α_{e_j} is unavoidable, α_{e_i} crosses α_{e_j} exactly once.*

Proof. Because of the restrictions, α_{e_i} does not cross the portions (if either or both of these exist) of α_{e_j} from $f^u_{e_i}$ to its first intersection with $P^1_{e_i} \cup P^2_{e_i}$ and the analogous segment from $f^v_{e_i}$. It is enough to show that α_{e_j} does not have two completely disjoint segments that have one end in $P^1_{e_i}$ and one end in $P^2_{e_i}$, but otherwise disjoint from $P^1_{e_i} \cup P^2_{e_i}$; this includes the possibility of one segment consisting of just one point in e_i.

Suppose τ_1 and τ_2 are two such segments of α_j. Since each involves a crossing of each of $P^1_{e_i}$ and $P^2_{e_i}$, the Useful Fact implies that these are the only crossings of α_j with $P^1_{e_i} \cup P^2_{e_i}$. We may assume that, in traversing α_j from one end to the other, we first traverse τ_1 from its end in $P^1_{e_i}$ to its end $\times_{2,1}$ in $P^2_{e_i}$. As we continue along α_j from $\times_{2,1}$, we are inside $\Delta_{\Sigma^2_{e_i}}$ until we meet the second crossing $\times_{2,2}$ of α_j and $P^2_{e_i}$.

The earlier "Useful Fact" asserts that $\times_{2,1}$ and $\times_{2,2}$ are the only crossings of α_j with $P^2_{e_i}$. It follows that $\times_{2,2}$ is an end of τ_2 and, continuing along α_{e_j} from $\times_{2,2}$ we are traversing τ_2 up to its other end, which is in $P^1_{e_i}$.

In summary, α_{e_j} crosses $P^1_{e_i}$ at one end of τ_1, crosses $P^2_{e_i}$ at the other end of τ_1, then goes through $\Delta_{\Sigma^2_{e_i}}$ until it crosses $P^2_{e_i}$ a second (and final) time, beginning its traversal of τ_2 up to the second (and final) crossing of $P^1_{e_i}$. The rest of α_{e_j} is inside $\Delta_{\Sigma^1_{e_i}}$ and so its terminus must be in $C_{\Sigma^1_{e_i}}$.

However, Lemma 3.4 and (II) imply α_j crosses $C_{\Sigma^1_{e_i}}$ only twice, and we have three crossings: $\tau_1 \cap P^1_{e_i}$, $\tau_2 \cap P^1_{e_i}$, and the terminus of α_{e_j}, a contradiction. $\square$

The verification that $\alpha_{e_1}, \ldots, \alpha_{e_i}$ satisfy Conditions (I) – (III) is completed by showing that α_{e_i} does not cross any avoidable α_{e_j} more than once. By way of contradiction, we assume that α_{e_i} crosses the avoidable α_{e_j} more than once. Since α_{e_j} is avoidable, either it does not cross $P^1_{e_i}$ or it does not cross $P^2_{e_i}$; for the sake of definiteness, we assume the latter. In particular, α_{e_j} does not cross e_i and, therefore, must cross each of the subarcs $\alpha^u_{e_i}$ and $\alpha^v_{e_i}$ (see Figure 3.8).

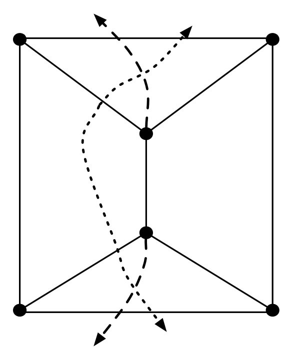

Figure 3.8: One instance of α_{e_i} crossing an avoidable α_{e_j} twice.

By the choice of η_v, there is an unavoidable $\alpha_{e_{j'}}$ that crosses $f^v_{e_i}$ between the intersections of α_{e_j} and α_{e_i} on $f^v_{e_i}$. Moreover, from its intersection with $f^v_{e_i}$, $\alpha_{e_{j'}}$ crosses $P^2_{e_i} \cap Q^v_{e_i}$ and, in going to that crossing, it must also cross the segment of α_{e_j} inside $Q^v_{e_i}$. Thus, (III) implies that $\alpha_{e_{j'}}$ and α_{e_j} cannot cross again.

As we follow $\alpha_{e_{j'}}$ from its end in $f^v_{e_i}$, we come first to the crossing with α_{e_j}, and then to a crossing with $P^2_{e_i}$. Continuing from this point, we cross $P^2_{e_i}$ again at $\times_{2,2}$ followed by the first crossing $\times_{1,1}$ with $P^1_{e_i}$. Some point $\times_{e_i}$ of $\alpha_{e_{j'}}$ in the closed subarc between $\times_{2,2}$ and $\times_{1,1}$ is in α_{e_i}. Claim 3 asserts that $\times_{e_i}$ is the unique crossing of $\alpha_{e_{j'}}$ with α_{e_i}.

The point $\times_{e_i}$ must lie on the segment of α_{e_i} between the two crossings of α_{e_i} with α_{e_j}, as otherwise $\alpha_{e_{j'}}$ must cross α_{e_j} a second time. It follows that, as we continue a short distance along $\alpha_{e_{j'}}$ beyond $\times_{e_i}$, there is a point of $\alpha_{e_{j'}}$ that is inside the simple closed curve bounded by the segments of each of α_{e_i} and α_{e_j} between their two intersection points. But $\alpha_{e_{j'}}$ must get to C_F from here without crossing $\alpha_{e_i} \cup \alpha_{e_j}$, which is impossible, showing that $\alpha_{e_1}, \ldots, \alpha_{e_i}$ satisfy all of (I)–(III).

What remains is to deal with the portions of the pseudolines that are in F. We begin by letting γ be a circle so that $D[K_n]$ is contained in the interior of γ. We label C_F as $(v_0, f_1, v_1, \ldots, f_k, v_0)$. Our first step is to extend one at a time each $D[f_i]$ to an arc β_{f_i} in $F \cup D[f_i]$ that, except for its endpoints, is contained in the open, bounded side of γ joining antipodal points a_i and b_i on γ. Pick arbitrarily two antipodal points a_1 and b_1 on γ and extend $D[f_1]$ in F to an arc β_{f_1} joining a_1 and b_1.

Suppose we have $\beta_{f_1}, \ldots, \beta_{f_{i-1}}$. The arc β_{f_i} will have to cross $\beta_{f_{i-1}}$ at v_{i-1}. (If $i = k$, then β_{f_i} will also have to cross β_{f_1} at v_0.) Extend $D[f_i]$ slightly so that it actually crosses $\beta_{f_{i-1}}$ at v_{i-1} (and β_{f_1} at v_0 when $i = k$). If $i < k$, then pick arbitrarily antipodal points a_i and b_i on γ distinct from $a_1, \ldots, a_{i-1}, b_1, \ldots, b_{i-1}$ and join the endpoints of f_i to these points, making sure to cross $\beta_{f_{i-1}}$ at v_{i-1}.

If $i = k$, then β_{f_1} and $\beta_{f_{k-1}}$ both constrain β_{f_k}. In this case, f_k is in precisely one of the four regions inside γ created by β_{f_1} and $\beta_{f_{k-1}}$. The arc in γ contained in the boundary of this region and its antipodal mate are to be avoided. The endpoints a_k and b_k of β_{f_k} are in the other antipodal pair of arcs in γ. We extend $D[f_k]$ slightly into each of these two regions.

In every case, we apply Lemma 2.2 to the two regions of $F \setminus \beta_{f_{i-1}}$. In particular, both slight extensions of $D[f_i]$ are constrained not to cross $\beta_{f_{i-1}}$ except at v_{i-1}. (For f_k, we restrict to the two of the four regions of $F \setminus (\beta_{f_{k-1}} \cup \beta_{f_1})$ that contain the slight extensions of $D[f_k]$.)

These restrictions guarantee that β_{f_i} does not cross $\beta_{f_{i-1}}$ more than once. Furthermore, β_{f_i} does not cross any of $\beta_{f_1}, \ldots, \beta_{f_{i-2}}$ more than once in each of the subregions of F. For $j = 1, 2, \ldots, i-2$, β_{f_i} and β_{f_j} have interlaced ends in γ and, therefore, they cross an odd number of times. It follows that they cross at most once.

We use a similar process to extend the arcs α_{e_i} that join points in C_F. Again, extend each one slightly into F in such a way that, every time two α_{e_i}'s meet at a common vertex in C_F, they cross at that vertex. The slightly extended α_{e_i} crosses some of the β_{f_j}'s: at least 2 (with equality if both endpoints of α_{e_i} are in the interiors of f_j's) and at most 4 (with equality if both endpoints of e_i are in C_F).

We will be proceeding with the α_{e_i} one by one, so that, when extending α_{e_i}, we already have extended $\alpha_{e_1}, \ldots, \alpha_{e_{i-1}}$, to arcs $\alpha^*_{e_1}, \ldots, \alpha^*_{e_{i-1}}$. Let Λ be the set consisting of those β_{f_j} and $\alpha^*_{e_k}$ that cross the slightly extended α_{e_i}. Since α_{e_i} crosses each arc in Λ exactly once, the two ends of α_{e_i} are in two regions of $F \setminus (\bigcup_{\lambda\in\Lambda} \lambda)$ that are incident with antipodal segments of γ. Choose arbitrarily an antipodal pair, one from each of these antipodal segments. Apply Lemma 2.2 to these two regions to extend the ends of α_{e_i} to these chosen antipodal points, yielding the arc $\alpha^*_{e_i}$.

Evidently, $\alpha^*_{e_i}$ crosses every arc in Λ exactly once. None of the other β_{f_j} and $\alpha^*_{e_k}$ crosses α_{e_i}. Therefore, $\alpha^*_{e_i}$ crosses each of these at most twice, once in each of the two regions in which $\alpha^*_{e_i}$ completes α_{e_i}. Since $\alpha^*_{e_i}$ and any β_{f_j} or $\alpha^*_{e_k}$ have interlaced ends in γ, they cross an odd number of times; thus, they cross exactly once. ■

4 Empty triangles in face-convex drawings

Let D be a drawing of K_n, let xyz be a 3-cycle in K_n and let Δ be an open disc bounded by $D[xyz]$. Then Δ is an *empty triangle* if $D[V(K_n)]\cap\Delta = \varnothing$. The classic theorem of Bárány and Füredi [6] asserts that, in any rectilinear drawing of K_n, there are $n^2 + \mathrm{O}(n \log n)$ empty triangles.

In Corollary 4.6, we extend the Bárány and Füredi theorem to pseudolinear drawings by proving the same theorem as theirs for face-convex drawings. Their proof adapts perfectly, as long as one has an appropriate "intermediate value" property. The other main result of this section is that any convex (not necessarily face-convex) drawing of K_n has at least $n^2/3 + \mathrm{O}(n)$ empty triangles.

Let D be a convex drawing of K_n and suppose that T is a transitive orientation of K_n with the additional property that, for each convex region Δ bounded by a triangle (u, v, w, u), if x is inside Δ, then x is neither a source nor a sink in the inherited orientation of the K_4 induced by u, v, w, x. A convex drawing of K_n with such a transitive orientation is *a convex intermediate value drawing*.

Convexity is not quite enough for our proof. A convex drawing D of K_n is *hereditarily convex* if, for each triangle T there is a specified side Δ_T of $D[T]$ that is convex and, moreover, for every triangle $T' \subseteq \Delta_T$, $\Delta_{T'} \subseteq \Delta_T$. Every pseudolinear drawing is trivially hereditarily convex, but so also is the "tin can" drawing of K_n that has $H(n)$ crossings. More generally, any drawing of K_n in which arcs are drawn as geodesics in the sphere is hereditarily convex.

Theorem 4.1 *A hereditarily convex intermediate value drawing of K_n has $n^2+\mathrm{O}(n\log n)$ empty triangles.*

Proof. Label $V(K_n)$ with $v_1, v_2, \ldots, v_n$ to match the intermediate value orientation, so $\overrightarrow{v_iv_j}$ is the orientation precisely when $i < j$. We henceforth ignore the arrow and use v_iv_j

for an edge only when $i < j$.

Bárány and Füredi [6, Lemma 8.1] prove the following fact.

Lemma 4.2 *Let G be a graph with vertex set $\{1, 2, \ldots, n\}$. Suppose that there are no four vertices $i < j < k < \ell$ such that ik, $i\ell$, and $j\ell$ are all edges of G. Then $|E(G)| \leq 3n\lceil \log n \rceil$.* ■

We will apply this lemma exactly as is done in [6]. For distinct i, j, k with $i < j$, say that v_k is *to the left of* v_iv_j if the face of the triangle $v_iv_jv_k$ containing the left side (as we traverse v_iv_j) of v_iv_j is convex.

Claim 1 *If v_k is to the left of v_iv_j, then the face of the triangle $v_iv_jv_k$ containing the left side of v_iv_j contains an empty triangle incident with v_iv_j.*

Proof. If v_ℓ is inside the convex triangle Δ, then v_ℓ is joined within Δ to the 3 corners of Δ and the triangle incident with v_ℓ contained in Δ, and incident with v_iv_j is convex by heredity. Since Δ contains a minimal convex triangle incident with v_iv_j, this one is necessarily empty. □

It follows that, as long as v_iv_j has a vertex to the left and a vertex to the right, then v_iv_j is in two empty triangles. The main point is to show that there are not many pairs (i, j) such that $i < j$ and v_iv_j is in only one empty triangle $v_iv_jv_k$ with $i < k < j$. We count those that have all intermediate vertices on the left.

Claim 2 *There do not exist $i < j < k < \ell$ such that: v_j and v_k are on the left of v_iv_ℓ; v_j is on the left of v_iv_k; v_k is on the left of v_jv_ℓ; and each of v_iv_ℓ, v_iv_k, and v_jv_ℓ is in at most one empty triangle.*

Proof. Suppose by way of contradiction that: $i < j < k < \ell$; v_j and v_k are on the left of v_iv_ℓ; v_j is on the left of v_iv_k; v_k is on the left of v_jv_ℓ; and each of v_iv_ℓ, v_iv_k, and v_jv_ℓ is in at most one empty triangle.

Let J be the complete subgraph induced by v_i, v_j, v_k, v_ℓ. Suppose first that $D[J]$ were a planar K_4. In order for both v_j and v_k to be on the left of v_iv_ℓ, one of v_j and v_k is inside the convex triangle Δ bounded by v_i, v_ℓ, and the other of v_j and v_k; let $v_{j'}$ be the one inside. Thus, the edges $v_iv_{j'}$ and $v_{j'}v_\ell$ are both inside Δ. (See Figure 4.3.)

The nested condition implies that the three smaller triangular regions inside Δ and incident with $v_{j'}$ are all convex. If $j' = j$, then v_j is to the right of vv_k, a contradiction. If $j' = k$, then v_k is to the right of v_jv_ℓ, a contradiction. Therefore, $D[J]$ is not a planar K_4.

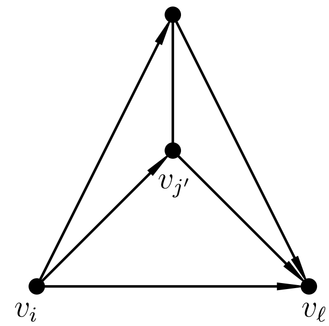

Figure 4.3: The case v_i, v_j, v_k, v_ℓ make a planar K_4.

In a crossing K_4, the convex sides of any of the triangles in the K_4 are the unions of two regions incident with the crossing. For each 3-cycle T in this K_4, the fourth vertex shows that it is on the non-convex side of T, so it is the bounded regions in the figure that are convex. With the assumption that v_j and v_k are both to the left of v_iv_ℓ, we see that v_iv_ℓ is in the 4-cycle that bounds a face of $D[J]$. Let $v_{j'}$ be the one of v_j and v_k that is the neighbour of v_i in this 4-cycle. (See Figure 4.4.)

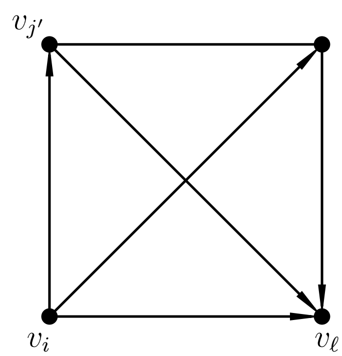

Figure 4.4: The case v, v_j, v_k, v_ℓ make a nonplanar K_4.

If $v_{j'} = v_k$, then v_j is to the right of v_iv_k, a contradiction. Therefore, $v_{j'} = v_j$. Consider the quarter of the inside of the 4-cycle that is incident with v_i, v_ℓ, and the crossing. We claim that there is no vertex of K_n in the interior of this region.

If v_r were in this region, then $r < i$ or $r > \ell$ violates the intermediate value property. If $i < r < k$, then heredity implies that we have the contradiction that v_r is to the right of v_iv_k. If $j < r < \ell$, then v_r is to the right of v_jv_ℓ, an analogous contradiction. It follows that the convex triangles bounded by both (v_i, v_ℓ, v_j, v_i) and (v_i, v_ℓ, v_k, v_i) contain empty triangles and these empty triangles are different, the final contradiction. □

It follows from Claim 2 and Lemma 4.2 that only $O(n \log n)$ pairs $i < j$ have the property that there is at most one k such that $i < k < j$ and $v_i v_k v_j$ is an empty triangle.

Since there are $\binom{n}{2}$ pairs $i < j$, there are $\binom{n}{2} - \mathrm{O}(n \log n)$ pairs with two empty triangles $v_i v_k v_j$ and $v_i v_{k'} v_j$, with $i < k, k' < j$. Thus, there are $n^2 - \mathrm{O}(n \log n)$ empty triangles, as required. ■

The proof that the Bárány-Füredi result holds for pseudolinear drawings comes from showing that every face-convex drawing has a transitive ordering with the intermediate value property.

Theorem 4.5 *If D is a face-convex drawing of K_n, then there is a transitive ordering with the intermediate value property.*

Proof. Let v_1 be any vertex incident with a face F witnessing face-convexity. Let $v_2, v_3, \dots, v_n$ be the cyclic rotation at v_1 induced by D, labelled so that $v_1 v_2$ and $v_1 v_n$ are incident with F.

We claim that this transitive ordering has the intermediate value property. Let J be an isomorph of K_4 such that $D[J]$ is planar; let $i < j < k < \ell$ be such that the four vertices of J are v_i, v_j, v_k, v_ℓ. Deleting $v_1, \dots, v_{i-1}$ and $v_{\ell+1}, \dots, v_n$ shows that v_i and v_ℓ are incident with the face of $D[J]$ that contains F. The convex side of the triangle of J bounding this face is, therefore, the side containing one of v_j and v_k. Evidently, this one is neither a sink nor a source of J. ■

Theorems 4.1 and 4.5 immediately imply the generalization of Bárány and Füredi to face-convex drawings.

Corollary 4.6 *Let D be a face-convex drawing of K_n. Then D has at least $n^2 + \mathrm{O}(n \log n)$ empty triangles.*

We conclude this work by showing that convexity is enough to guarantee $\mathrm{O}(n^2)$ empty triangles. This is somewhat surprising, since it is known that general drawings of K_n can have as few as $2n - 4$ empty triangles [10].

Theorem 4.7 *Let D be a convex drawing of K_n. Then D has $n^2/3 - \mathrm{O}(n)$ empty triangles.*

Proof. Let U be the subgraph of K_n induced by the edges not crossed in D. Then $D[U]$ is a planar embedding of the simple graph U, so U has at most $3n - 6$ edges. Thus, there are $n^2/2 - \mathrm{O}(n)$ edges that are crossed in D.

For each edge e that is crossed in D, let f be one of the edges that crosses e. Let $J_{e,f}$ be the K_4 induced by the vertices of K_n incident with e and f. Because $D[J_{e,f}]$ is a crossing K_4, each of the four triangles in $J_{e,f}$ has a unique convex side; it is the side not containing the fourth vertex. Let $\Delta_{e,f}$ be the closed disc that is the union of these four convex triangles.

It follows that, for any other vertex x that is on the convex side of any of these four triangles, x is joined inside $\Delta_{e,f}$ to the four corners of $\Delta_{e,f}$, implying one of the edges f' incident with x crosses e. Now $\Delta_{e,f'}$ is strictly contained in $\Delta_{e,f}$, proving that there is a minimal $\Delta_{e,f}$ that does not contain any vertex of K_n in its interior.

The two convex triangles incident with e and contained in such a minimal $\Delta_{e,f}$ are both empty. Thus, every crossed edge is in at least two empty triangles. Since every empty triangle contains precisely three edges, there are at least

$$\frac{1}{3}2\left(\frac{n^2}{2} - \mathrm{O}(n)\right) = \frac{n^2}{3} - \mathrm{O}(n)$$

empty triangles, as required. ■

References

[1] B.M. Ábrego and S. Fernández-Merchant, A lower bound for the rectilinear crossing number, Graphs Combin. **21** (2005), 293–300.

[2] B.M. Ábrego, O. Aichholzer, S. Fernández-Merchant, P. Ramos, and G. Salazar, Shellable drawings and the cylindrical crossing number of K_n, Discrete Comput. Geom. **52** (2014), no. 4, 743–753.

[3] O. Aichholzer, T. Hackl, A. Pilz, G. Salazar, and B. Vogtenhuber, Deciding monotonicity of good drawings of the complete graph, preprint, April 2015.

[4] A. Arroyo, D. McQuillan, G. Salazar, and R.B. Richter, A proof of Gioan's Theorem, preprint.

[5] J. Balogh, J. Leaños, S. Pan, R.B. Richter, and G. Salazar, The convex hull of every optimal pseudolinear drawing of K_n is a triangle, Australas. J. Combin. **38** (2007), 155–162.

[6] I. Bárány and Z. Füredi, Empty simplices in Euclidean space, Canad. Math. Bull. **30** (1987), no. 4, 436–445.

[7] E. Gioan, Complete graph drawings up to triangle mutations, D. Kratsch (Ed.): WG 2005, LNCS 3787, pp. 139–150, 2005, Springer-Verlag, Berlin-Heidelberg.

[8] J.E. Goodman and R. Pollack, On the combinatorial classification of non-degenerate configurations in the plane, J. Combin. Theory Ser. A. **29** (1980), 220–235.

[9] B. Grünbaum, Arrangements and Spreads, Conference Board of the Mathematical Sciences Regional Conference Series in Mathematics, No. 10. Amer. Math. Soc., Providence, R.I., 1972.

[10] H. Harborth, Empty triangles in drawings of the complete graph, Discrete Math. **191** (1998), no. 1-3, 109–111.

[11] D. Knuth, Axioms and Hulls, Vol. 606, Lecture Notes Comput. Sci., Springer, Berlin, 1992.

[12] F. Levi, Die Teilung der projektiven Ebene durch Gerade oder Pseudogerade, Ber. Math-Phys. Kl. Sächs. Akad. Wiss. **78** (1926), 256–267.

[13] L. Lovász, K. Vesztergombi, U. Wagner, and E. Welzl, Convex quadrilaterals and k-sets, in *Towards a theory of geometric graphs*, 139–148, Contemp. Math. **342**, Amer. Math. Soc., Providence, RI, 2004.

[14] B. Sturmfels and G.M. Ziegler, Extension spaces of oriented matroids, Discrete Comput. Geom. **10** (1993), no. 1, 23–45.