Extending drawings of complete graphs into arrangements of pseudocircles

Alan Arroyo,* R. Bruce Richter,† and Matthew Sunohara‡

alanarroyoguevara@gmail.com, brichter@uwaterloo.ca, matthew.sunohara@mail.utoronto.ca

LaTeX-ed: April 16, 2022

Abstract

Motivated by the successful application of geometry to proving the Harary-Hill Conjecture for "pseudolinear" drawings of K_n, we introduce "pseudospherical" drawings of graphs. A *spherical drawing* of a graph G is a drawing in the unit sphere $\mathbb{S}^2$ in which the vertices of G are represented as points—no three on a great circle—and the edges of G are shortest-arcs in $\mathbb{S}^2$ connecting pairs of vertices. Such a drawing has three properties: (1) every edge e is contained in a simple closed curve γ_e such that the only vertices in γ_e are the ends of e; (2) if $e \neq f$, then $\gamma_e \cap \gamma_f$ has precisely two crossings; and (3) if $e \neq f$, then e intersects γ_f at most once, either in a crossing or an end of e. We use Properties (1)–(3) to define a *pseudospherical drawing* of G. Our main result is that, for the complete graph, Properties (1)–(3) are equivalent to the same three properties but with "precisely two crossings" in (2) replaced by "at most two crossings".

The proof requires a result in the geometric transversal theory of arrangements of pseudocircles. This is proved using the surprising result that the absence of special arcs (*coherent spirals*) in an arrangement of simple closed curves characterizes the fact that any two curves in the arrangement have at most two crossings.

Our studies provide the necessary ideas for exhibiting a drawing of K_{10} that has no extension to an arrangement of pseudocircles and a drawing of K_9 that does extend to an arrangement of pseudocircles, but no such extension has all pairs of pseudocircles crossing twice.

1   Introduction

The Harary-Hill Conjecture states that the crossing number of the complete graph K_n is given by the formula

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

Some of the known families of drawings of K_n achieving $H(n)$ crossings have the geometric character of being spherical: a *spherical drawing* of a graph G is a drawing in the unit sphere $\mathbb{S}^2$ in which the vertices of G are represented as points—no three on a great circle—and the edges of G are shortest-arcs in $\mathbb{S}^2$ connecting pairs of vertices.

Examples of families having $H(n)$ crossings are: Hill's Tin Can Drawings [18]; Kynčl's general spherical drawing in his posting [20]; the family of Ábrego et al. [5] in which every edge is crossed at least once; and the crossing-minimal 2-page drawings in Ábrego et al. [6, Sec. 4.3, 4.4], where they show that the 2-page crossing number of K_n is $H(n)$. The first three of these are known to be spherical.

A surprising result by Moon [23] states that the number of crossings in a random spherical drawing of K_n has, as n goes to infinity, $\frac{3}{8}\binom{n}{4}$ crossings. Thus, spherical drawings are linked with the asymptotic version of the Harary-Hill Conjecture:

$$\lim_{n \to \infty} \frac{\mathrm{cr}(K_n)}{\binom{n}{4}} = \frac{3}{8}.$$

No drawing of K_n having $H(n)$ crossings is known to be non-spherical (at least up to Reidemeister-type moves; see below). (For n even, the 2-page drawings in [6] are "pseudospherical"; see the discussion in our Section 4.)

2

However, it is still unknown whether every spherical drawing of K_n has at least $H(n)$ crossings.

The analogue in the plane of spherical drawings is rectilinear drawings. A *rectilinear* drawing of a graph G is a drawing of G in the plane so that its edges are straight-line segments. One of the most important recent accomplishments in the study of crossing numbers is a result of Ábrego and Fernández-Merchant [1] and, simultaneously and independently, Lovász et al. [21], showing that rectilinear drawings of K_n have at least $H(n)$ crossings. (It follows from [7], especially Theorem 11 there, that, for $n \geq 10$, rectilinear drawings have strictly more than $H(n)$ crossings.) There is a quite direct line from this early work, via shellability [3, 4], to Ábrego et al [6] determining the 2-page crossing number of K_n.

The proofs that rectilinear drawings of K_n have at least of $H(n)$ crossings use machinery for studying arrangements of pseudolines, and only require the property that the edges in a rectilinear drawing can be extended to such an arrangement. Drawings whose edges can be extended into an arrangement of pseudolines are called *pseudolinear*. An analogous property is satisfied by spherical drawings: each edge can be extended to a great circle. This implies that the edges can be extended into an *arrangement of pseudocircles*, defined as a set of simple closed curves in $\mathbb{S}^2$ such that every two intersect at most twice, and every intersection is a crossing between two curves.

The success of the geometric approach for the rectilinear crossing number of K_n suggests trying an analogous approach for spherical drawings, replacing pseudolinear drawings with a suitable generalization of spherical drawings, which we call "pseudospherical" drawings and define below.

There are nice characterizations of pseudolinear drawings of K_n. Aichholzer et al. [8] prove that a drawing of K_n in the plane is pseudolinear if and only if every crossing K_4 has the facial 4-cycle bounding the infinite face of the K_4. Arroyo et al. [10] have an equivalent characterization: a drawing of K_n is pseudolinear if and only if it is f-convex. Arroyo et al. [9] characterize when a drawing of a general graph in the plane is pseudolinear.

The notion of an f-convex drawing is introduced by Arroyo et al. in [12] as part of the *convexity hierarchy*:

$$\{\text{convex drawings}\} \supset \{\text{h-convex drawings}\} \supset \{\text{f-convex drawings}\}.$$

A drawing D of K_n is *convex* if, for each 3-cycle T of K_n, there is a closed side Δ of the drawing $D[T]$ of T such that, if vertices x, y are drawn in Δ,

3

then $D[xy] \subseteq \Delta$. A principal theorem in [11] is a partial result suggesting that all crossing-minimal drawings of K_n might be convex.

A convex drawing D of K_n is *h-convex* if the convex side Δ_T of each 3-cycle T may be chosen so that, if a 3-cycle T' is, for another 3-cycle T, drawn in Δ_T, then $\Delta_{T'} \subseteq \Delta_T$. In unpublished work, the method for the principal theorem in [11] mentioned above shows that a convex, but not h-convex, drawing of K_n is not crossing-minimal. Putting these two ideas together, it is conceivable that every crossing-minimal drawing of K_n is h-convex.

It is quite easy to show that every spherical drawing is h-convex (see Section 4). However, it is not true that every h-convex drawing is spherical. The property of h-convexity is preserved under (the natural analogue of) Reidemeister III moves, but sphericity is not. Reidemeister III moves preserve three properties of spherical drawings which we take as the definition of a pseudospherical drawing.
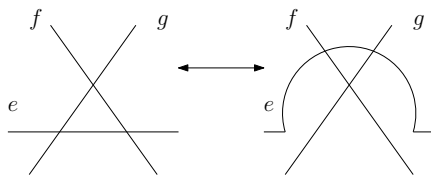

Figure 1: Reidemeister III move.

A drawing of a graph G in the sphere is *pseudospherical* if, for any distinct edges e and f:

(PS1) e is contained in a simple closed curve γ_e such that no vertex other than an end of e is contained in γ_e;

(PS2) $|\gamma_e \cap \gamma_f| = 2$ and all intersections are crossings; and

(PS3) $e \cap \gamma_f$ has at most one point.

The least obvious part of the definition is perhaps (PS3), which can be thought of as a combinatorial analogue of the property that an edge in a spherical drawing is not just a geodesic arc between its endpoints, but a shortest-arc between its endpoints.

Since the number of crossings of a drawing is also preserved under Reidemeister III moves, one obtains non-spherical drawings of K_n with $H(n)$ crossings from the spherical examples given above. It is more natural to

consider pseudospherical drawings than spherical drawings of K_n in connection with crossing-minimality because pseudosphericity is preserved under Reidemeister III moves.

The principal part of our main theorem is to show that a drawing of K_n is pseudospherical if and only if it is h-convex, a property which can be checked in polynomial time (see Section 4). This can be seen as a parallel to the characterization by Arroyo et al. [10] of pseudolinear drawings of K_n as f-convex drawings of K_n. The equivalence also enables the introduction of geometric methods to the study of h-convex drawings and bolsters the idea that all crossing-optimal drawings of K_n are h-convex/pseudospherical: evidence that crossing-optimal drawings of K_n are h-convex is evidence that they are pseudospherical and vice versa.

In fact, we will do quite a bit more than characterize pseudospherical drawings of K_n. The reader may have wondered about the choice of equality in (PS2). A natural variation of the notion of a pseudospherical drawing results from weakening (PS2). A drawing D is *weakly pseudospherical* if it satisfies (PS1), (PS3), while (PS2) is relaxed to

(PS2w) $|\gamma_e \cap \gamma_f| \leq 2$ and all intersections are crossings.

Our main result is the following.

Theorem 1.1. *For a drawing D of K_n, the following are equivalent:*

(1.1.1) *D is pseudospherical;*

(1.1.2) *D is weakly pseudospherical; and*

(1.1.3) *D is h-convex.*

The implication $(1.1.1) \Rightarrow (1.1.2)$ is trivial. Although the reader may not see it now, the implication $(1.1.2) \Rightarrow (1.1.3)$ is quite easy. The hard part is $(1.1.3) \Rightarrow (1.1.1)$. We do not see how to prove $(1.1.2) \Rightarrow (1.1.1)$ directly.

The proof of $(1.1.3) \Rightarrow (1.1.1)$ proceeds by iteratively finding a curve $\gamma_{e'}$ for the next edge e' to extend by one the current set Γ of γ_es satisfying the conditions (PS1)–(PS3). There are two principal steps involved. The first step is to find two initial approximations for $\gamma_{e'}$ (when added to Γ, either of these will satisfy (PS2w)), while the second is to repeatedly shift one of the initial approximations, gradually increasing the number of curves in Γ that it intersects until it intersects them all, at which point it is a Γ-transversal.

Each of these steps has its challenges. To find the initial approximations, we require an extensive study of h-convex drawings; this is done in Section 4. Crucially, each edge of an h-convex drawing partitions the vertex set of K_n into two pseudolinear drawings of (typically smaller) complete subgraphs; the initial approximations are near the outer boundary of each of these pseudolinear subdrawings. (Motivated by this study, we tried to show this partitioning holds for pseudospherical drawings of general graphs. In a personal communication, Xinyu Lily Wang has shown that it is false in the more general context.)

Producing the Γ-transversal requires shifting one of a pair of initial approximations from the preceding paragraph towards the other using analogues of Reidemeister II and III moves. The core of the shifting turns out to require a characterization of an arrangement of pseudocircles: a set of simple closed curves in the sphere such that every two intersect at most twice, and every intersection is a crossing between the two curves.

The characterization of arrangement of pseudocircles requires an even more general notion. An *arrangement of simple closed curves* in the sphere is a set of simple closed curves, any two of which have finitely many intersections, all of which are crossings. If Σ is an arrangement of simple closed curves, then a *spiral* of Σ is an arc (that is, a homeomorph of a compact interval) in the union $P(\Sigma)$ of the curves in Σ that always makes the same—either all left or all right—turn in changing from one curve to another (see Figure 2). In Section 2, we give a more precise definition of spiral and define *coherent spirals*. The auxiliary result that we need is the following.

Theorem 1.2. *Let Σ be an arrangement of simple closed curves. Then Σ is an arrangement of pseudocircles if and only if Σ has no coherent spirals.*

Our study of spirals led us to two drawings, one for each of K_{10} and K_9. The former has no extension of its edges to an arrangement of pseudocircles. The latter has such an extension, but no extension has all pseudocircles crossing exactly twice. These examples are exhibited in Section 7.

Independently, Aichholzer at the 2015 Crossing Number Workshop in Rio de Janeiro and Pilz at the Geometric Graph Week in Berlin (2015) asked if every drawing of K_n has an extension to an arrangement of pseudocircles. The drawing of K_{10} answers this question in the negative. The drawing of K_9 answers negatively the related question of improving a pseudocircular extension to a pseudospherical extension.

6

In the next section, we introduce spirals and prove Theorem 1.2. Section 3 contains the proof that an initial pair of approximations for Γ implies the existence of a Γ-transversal. Section 4 introduces h-convex drawings and proves the easy implication $(1.1.2) \Rightarrow (1.1.3)$. Section 5 contains the necessary discussion of h-convex drawings to obtain the initial approximations, which is done in Section 6, thereby completing the proof of Theorem 1.1. The interesting drawings of K_9 and K_{10} are in Section 7, while Section 8 has concluding remarks.

2   Spirals and Coherence

A spiral is a special arc, illustrated in Figure 2 and defined precisely below, in the union of an arrangement of simple closed curves that has all its crossings facing the same side of the arc. (Alternatively, a spiral always makes the same—either all left or all right—turn when it changes from one curve to another; it is permitted to continue straight through a crossing.) The main result in this section is a characterization of arrangements of pseudocircles as either (a) not having any "coherent" spirals or, equivalently, (b) every spiral has an "external segment". That an arrangement of pseudocircles has no coherent spiral is the point required for the proof of Theorem 1.1.
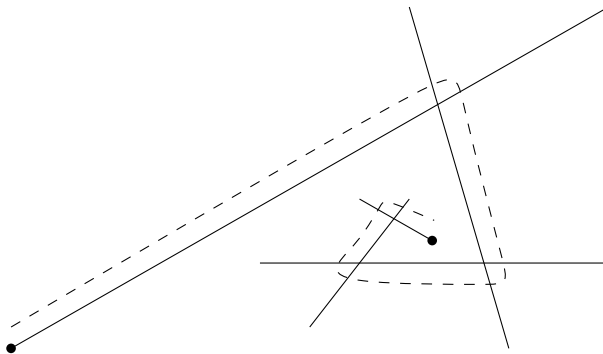

Figure 2: An example of a spiral.

In an arrangement of simple closed curves Γ (defined just before Theorem 1.2), three or more of the curves in Γ may cross at the same point. This is necessary for extensions of drawings of K_n, where $n-1$ curves pairwise cross at each vertex.

A non-trivial arc (that is, not just a single point) A in the union $P(\Gamma)$ of the curves in Γ with ends s and t has a unique decomposition sequence $\alpha_0 \alpha_1 \ldots \alpha_m$ of subarcs of A, as depicted in Figure 3, such that:

(i) s is an end of α_0, t is an end of α_m;

(ii) for each $i = 0, 1, \ldots, m$, there is a $\gamma_i \in \Gamma$ such that $\alpha_i \subseteq \gamma_i$; and

(iii) for $i = 1, 2, \ldots, m$, the curves γ_{i-1} and γ_i in Γ are distinct and $\alpha_{i-1} \cap \alpha_i$ is a crossing of γ_{i-1} and γ_i.

The number m is the *weight of A*. Figure 3 shows an arrangement of simple closed curves with an arc of weight 3. For $i = 1, 2, \ldots, m$, the crossing of α_{i-1} with α_i is denoted $\times_i$. For convenience, we set $\times_0 = s$ and $\times_{m+1} = t$.
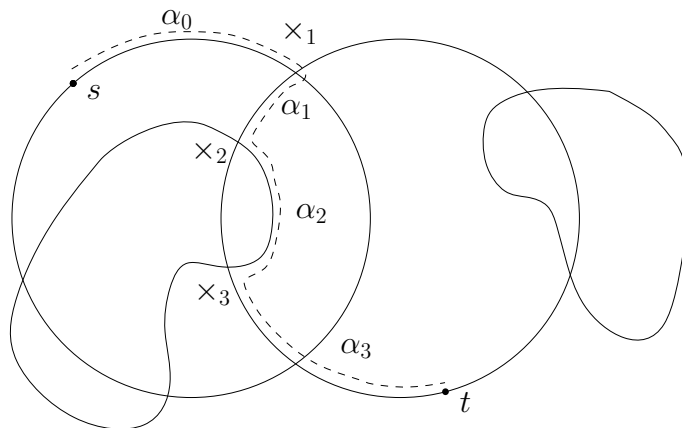

Figure 3: Illustrated is an arrangement of four simple closed curves. The st-arc indicated by the dotted curve has decomposition sequence $\alpha_0 \alpha_1 \alpha_2 \alpha_3$. Notice that the $\times_3 t$-arc α_3 makes no turn at the crossing in its middle and the right-hand circle is both γ_1 and γ_3.

For $i = 0, 1, \ldots, m$, α_i^- and α_i^+ are the closures of the components of $\gamma_i \setminus A$ incident with $\times_i$ and $\times_{i+1}$, respectively. Four such α_j^ε are illustrated in Figure 4. Evidently, $\alpha_i^\pm$ consists of the continuation of α_i through either $\times_i$ ($-$) or $\times_{i+1}$ ($+$) up to the next meeting with A.

The continuations α_i^+ and α_{i+1}^- both leave $\times_{i+1}$ on the same side of A. This is the *side of A that $\times_{i+1}$ faces*. In Figure 3, $\times_1$ and $\times_2$ face different

sides of the dotted arc. We are now prepared for the definitions of spiral, external segment, and coherence.

Definition 2.1 (Spiral, External Segment, Coherence). Let Γ be an arrangement of simple closed curves and let A be an arc in $P(\Gamma)$ with decomposition sequence $\alpha_0\alpha_1\ldots\alpha_m$. For each $i = 0, 1, 2\ldots, m$, let $\gamma_i \in \Gamma$ be such that $\alpha_i \subseteq \gamma_i$. (Only consecutive γ_i are required to be distinct; if $j > i + 1$, then γ_j could be the same as γ_i.)

(2.1.1) The arc A is a *spiral* if all of $\times_1, \ldots, \times_m$ face the same side of A.

(2.1.2) For $i \in \{0, 1, 2, \ldots, m\}$, the segment α_i is

$$\begin{cases} external \text{ for } A, & \gamma_i \cap A = \alpha_i \\ internal \text{ for } A, & \text{otherwise} \end{cases} .$$

(2.1.3) For $i \in \{0, 1, 2, \ldots, m\}$ and $\varepsilon \in \{+, -\}$, the arc α_i^ε is a *coherent extension* if: (a) α_i is internal for A and (b) α_i^ε has both ends on the same side (see discussion of "side" below) of the interior of A. (In Figure 4, α_1^- is not a coherent extension; the other three are.)

(2.1.4) For $i \in \{0, 1, 2, \ldots, m\}$, the segment α_i is *coherent* if at least one of α_i^- and α_i^+ is a coherent extension. In Figure 4, both α_0 and α_3 are coherent.

(2.1.5) The arc A is *coherent* if, for each $i = 0, 1, 2, \ldots, m$, α_i is coherent.

As we traverse A from one end to the other, there are naturally left and right sides. The two ends of the arc $\alpha_i^\pm$ are in A. The issue in the definition of coherence is: are the points near each end of $\alpha_i^\pm$ on the same side of A or not. "Left" and "right" depend on an orientation of A and are irrelevant to us.

To be on the same side, the ends of $\alpha_i^\pm$ must be in the interior of A. As the points $\times_0$ and $\times_{m+1}$ are not in the interior of A, we have the following.

Remark 2.2. *Neither α_0^- nor α_m^+ is coherent.*

The following characterization of an arrangement of pseudocircles in terms of its spirals seems to be quite interesting in its own right. We only need $(2.3.1) \Rightarrow (2.3.2)$ for the proof of Theorem 1.1.
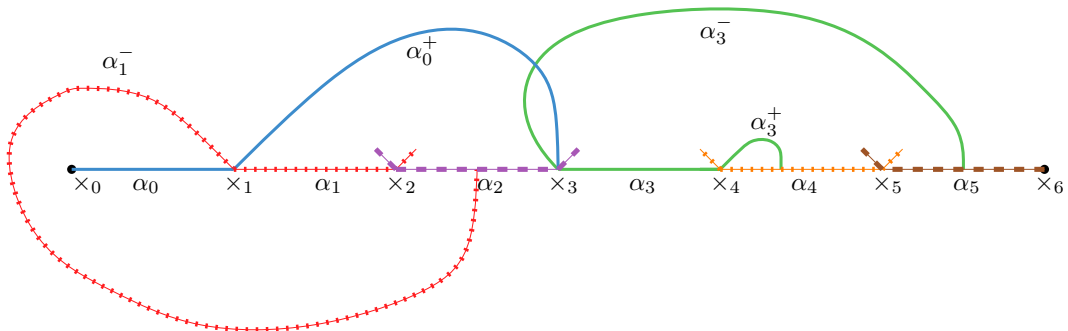
Figure 4: Illustrating several points in Definition 2.1. The arc A is homeomorphic to a straight line segment and is so represented in the diagram. (Color online.)

Theorem 2.3. *Let Γ be an arrangement of simple closed curves in the sphere. Then the following are equivalent:*

(2.3.1) Γ *is an arrangement of pseudocircles;*

(2.3.2) $P(\Gamma)$ *has no coherent spirals;*

(2.3.3) $P(\Gamma)$ *has no coherent spirals with weight 1; and*

(2.3.4) *every spiral A in $P(\Gamma)$ has a segment external for A.*

This characterization fits in well with much recent work on arrangements of pseudocircles. The paper of Felsner and Scheucher [15] is one example; their references [2], [7], [10], [15], [19], and [26] are others. Felsner and Scheucher have a web page devoted to pseudocircles [16]. Another recent work about unavoidable configurations in the sense of Ramsey's Theorem is Medina et al. [22].

The example to the left in Figure 5 has a spiral A in an arrangement of pseudocircles. This spiral is not coherent because α_1 is external for A.

In addition to its use in the proof of Theorem 1.1, Theorem 2.3 is useful for constructing drawings of graphs that cannot be extended to arrangements of pseudocircles. For example, the right-hand diagram in Figure 5 cannot be extended to an arrangement of pseudocircles because the *st*-arc indicated by the dashed curve induces a coherent spiral in any such extension. A similar idea is used in Section 7 to construct a non-extendible drawing of K_{10}.
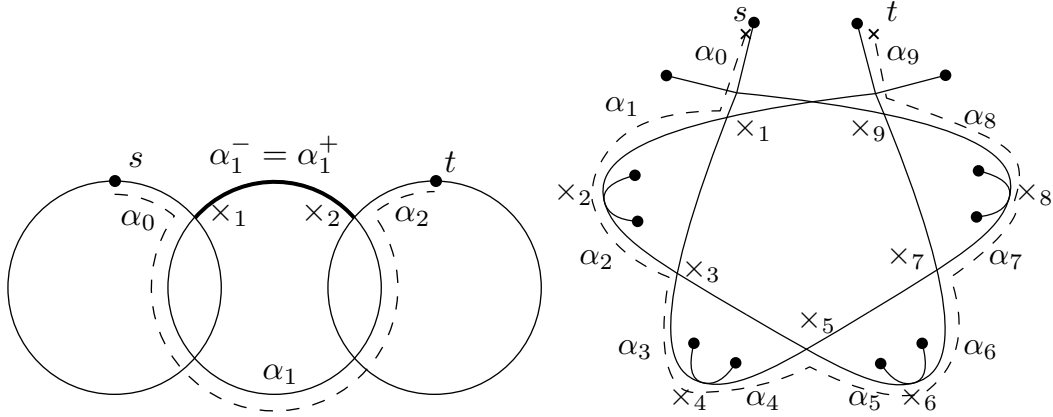
Figure 5: The dotted arcs are spirals. The left-hand one has weight 2 and has an external segment. The right-hand one is coherent and is in a drawing related to the examples in Section 7.

Proof of Theorem 2.3. The implication $(2.3.2) \Rightarrow (2.3.3)$ is trivial. We prove $(2.3.3) \Rightarrow (2.3.1)$ and $(2.3.1) \Rightarrow (2.3.2)$ to complete the proof that $(2.3.1)$, $(2.3.3)$, and $(2.3.2)$ are equivalent. Since $(2.3.4) \Rightarrow (2.3.2)$ is trivial, we finish the proof with $(2.3.2) \Rightarrow (2.3.4)$.

$(2.3.3) \Rightarrow (2.3.1)$. Suppose by way of contradiction that γ_0, γ_1 are distinct curves in Γ such that $|\gamma_0 \cap \gamma_1| > 2$. Then $|\gamma_0 \cap \gamma_1| \geq 4$. Let $s \in \gamma_0 \setminus \gamma_1$ and, traversing γ_0 in one direction starting at s, let p_1, p_2, p_3 be the first three points of $\gamma_0 \cap \gamma_1$ encountered.

Let A be the arc obtained by starting at s, continuing along γ_0 through p_1 to p_2 and then following $\gamma_1 \setminus \{p_1\}$ through p_3 to a point t just beyond p_3. The decomposition of A is $\alpha_0 \alpha_1$, with, for $i = 1, 2$, $\alpha_i \subsetneq A \cap \gamma_i$. Since every arc in $P(\Gamma)$ with weight at most 1 is a spiral, A is a spiral. The fact that the points p_2 and p_3 of $\gamma_0 \cap \gamma_1$ are consecutive in γ_0 imply that α_0^+ is coherent.

On the other hand, the sp_2-subarc of A intersects the $p_1 p_2$-subarc of $\gamma_1 \setminus \{p_3\}$ just in $\{p_1, p_2\}$. This shows that α_1^- is coherent, completing the proof that A is coherent, the required contradiction.

$(2.3.1) \Rightarrow (2.3.2)$. To obtain a contradiction, suppose A is a coherent spiral with least weight. An arc with weight 0 is not coherent, so the decomposition $\alpha_0 \alpha_1 \cdots \alpha_m$ of A has $m \geq 1$.

The first claim imposes constraints on what happens at points "under" a forward jump such as, in Figure 4, $\times_2$ under α_0^+. For an extension α_i^ε of α_i, a_i^ε denotes the other end of α_i^ε.

11

Claim 1. *Suppose that α_j^+ is a coherent extension of α_j. If there is an $\ell \geq j + 2$ such that a_j^+ is in $\alpha_\ell \setminus \{\times_\ell\}$, then:*

1. *for each $i = j + 2, j + 3, \ldots, \ell$, α_i^- is disjoint from $\alpha_j^+ \setminus \{a_j^+\}$; and*

2. *either: for some $i \in \{j + 2, j + 3, \ldots, \ell\}$, $a_j^+ \in \alpha_i^-$; or α_{j+1}^+ intersects α_j^+.*

Likewise, suppose α_j^- is a coherent extension of α_j. If there is an $\ell \leq j - 2$ such that a_j^- is in $\alpha_\ell \setminus \{\times_{\ell+1}\}$, then:

3. *for each $i = \ell, \ell + 1, \ldots, j - 2$, α_i^+ is disjoint from $\alpha_j^- \setminus \{a_j^-\}$; and*

4. *either: for some $i \in \{\ell, \ell - 1, \ldots, j - 2\}$, $a_j^- \in \alpha_i^+$; or α_{j-1}^- intersects α_j^-.*

Proof. We prove the first statement; the "likewise" is the same, but for the traversal of A in the reverse direction.

By way of contradiction, suppose first that, for some $i \in \{j + 2, j + 3, \ldots, \ell\}$, α_i^- intersects $\alpha_j^+ \setminus \{a_j^+\}$; let $\times_{i^-,j^+}$ be the first intersection with α_j^+ as we traverse α_i^- from $\times_i$. The interior of the subarc $\alpha_i^-[\times_i, \times_{i^-,j^+}]$ of α_i^- from $\times_i$ to $\times_{i^-,j^+}$ is on the side of the unique simple closed curve contained in $A \cup \alpha_j^+$ that is opposite to the side that contains $A[s, \times_{j+1}]$.

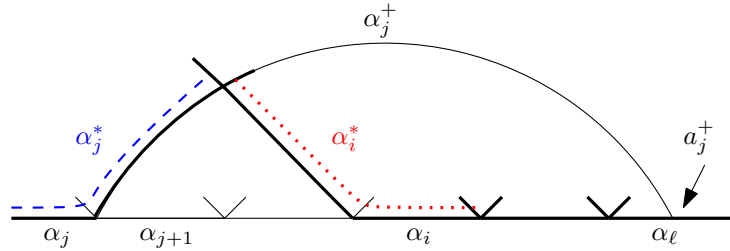See Figure 6 for an illustration of this proof.

Figure 6: Illustration for the proof of Claim 1. (Color online.)

Let α_j^* be the subarc of γ_j consisting of α_j and the portion of α_j^+ from $\times_{j+1}$ to $\times_{i^-,j^+}$. Likewise, let α_i^* be the subarc of γ_i consisting of α_i and $\alpha_i^-[\times_i, \times_{i^-,j^+}]$. The arc A' with decomposition $\alpha_0 \ldots \alpha_{j-1} \alpha_j^* \alpha_i^* \alpha_{i+1} \ldots \alpha_m$ has smaller weight than A. Also, even if $\times_{i^-,j^+} = \times_{j+1}$, A' is a spiral. (In case $\times_{i^-,j^+} = \times_{j+1}$, then γ_j, γ_{j+1}, and γ_i all cross at $\times_{j+1}$. This ensures that

12

the cyclic rotation of these three curves at $\times_{j+1}$ is $\alpha_j, \alpha_{j+1}, \alpha_i^-, \alpha_j^+, \alpha_{j+1}^-, \beta_i$, where β_i is the continuation of γ_i from $a_i^- = \times_{i^-,j^+}$.)

To see that A' is coherent, first let C be the simple closed curve $(\alpha_j^* \setminus \alpha_j) \cup (\alpha_i^* \setminus \alpha_i) \cup A[\times_{j+1}, \times_i]$. For each of the segments α' of A', α' contains some segment α_k of A. Let α_k^ε be a coherent extension of α_k for A. Follow α_k^ε from its end in α_k (or, if $k \in \{j, i\}$, from $\times_{i^-,j^+}$). If we never encounter C, then we arrive at A' on the same side. On the other hand, if we encounter C, it is not at a point in A and so it is in $(\alpha_j^* \setminus \alpha_j) \cup (\alpha_i^* \setminus \alpha_i)$.

Label as the *outside of C* the side of C containing $A[s, \times_{j+1}]$. Since α_k^ε starts on the outside of C, its first intersection with C is from that side. Thus, the portion of α_k^ε up to that first intersection with C is a coherent extension of α', as required.

To complete the proof that A' is coherent, we note that, if, for the segment α of A', both α^- and α^+ are coherent, then they are contained in coherent extensions of the corresponding segment of A. Since these extensions for A are distinct, as extensions for A' they are also distinct. Thus, there is a coherent spiral with weight less than m, a contradiction.

For (2), if, for each $i = \ell, \ell+1, \ldots, j-1$, $a_j^+ \notin \alpha_i^-$ and α_{j+1}^+ is disjoint from α_j^+, then α_{j+1}^+ and (using (1)) each α_i^- is a coherent extension in $A[\times_{j+1}, a_j^+]$. The same argument as in the preceding paragraph shows that $A[\times_{j+1}, a_j^+]$ is a coherent spiral with smaller weight than A, a contradiction. $\qquad\square$

The next claim considers "reverse" coherent extensions. This claim rules out the possibility of an extension such as α_3^- in Figure 4.

Claim 2. *There do not exist $j, \ell \in \{0, 1, 2, \ldots, m\}$ such that $\ell < j$ and α_j^+ is a coherent extension with an end in α_ℓ.*

Likewise, there do not exist $j, \ell \in \{0, 1, 2, \ldots, m\}$ such that $j < \ell$ and α_j^- is a coherent extension with an end in α_ℓ.

Proof. We only prove the first statement. Choose the least j for which such an $\ell < j$ exists. Suppose first that $a_j^+ \in \alpha_{j-1}$. Thus, $\gamma_{j-1} \cap \gamma_j \supseteq \{a_j^+, \times_j\}$. Because $\alpha_j^+ \neq \alpha_j^-$, $a_j^+ \neq \times_j$. Thus, Hypothesis (2.3.1) implies $\gamma_{j-1} \cap \gamma_j = \{a_j^+, \times_j\}$. Let C be the simple closed curve contained in $\alpha_{j-1} \cup \alpha_j \cup \alpha_j^+$. Except for $\times_j$, α_{j-1}^+ is disjoint from C.

Note that a_{j-1}^+ is not in $\alpha_{j-1} \cup \alpha_j$. Therefore, a_{j-1}^+ and the start of α_{j-1}^+ from $\times_j$ are on different sides of C, which is impossible. Thus, $a_j^+ \notin \alpha_{j-1}$.

The choice of j implies that either α_{j-1}^+ is not coherent or it does not intersect $A[a_j^+, \times_j]$. Therefore, α_{j-1}^+ must intersect $\alpha_j \cup \alpha_j^+$ at a point other

13

than $\times_j$. This gives the two crossings of γ_{j-1} with γ_j.

An intersection of α_{j-1}^- with γ_j yields a third intersection of γ_{j-1} with γ_j, and the theorem is proved. Therefore, we assume α_{j-1}^- is disjoint from α_j^-. On the other hand, the choice of j implies that, for each k with $\times_{k+1}$ in the interior of $A[a_j^-, \times_j]$, $a_k^+ \neq a_j^-$. Therefore Claim 1 (4) shows α_{j-1}^- is not disjoint from α_j^-, the final contradiction. $\qquad\square$

We may now suppose that there does not exist an α_j^+ that is a coherent extension with an end in any α_k such that $k < j$. Similarly, we may assume that there does not exist an α_j^- that is a coherent extension with any end in any α_k such that $k > j$.

The final claim combines the first two to completely determine the nature of a coherent extension. Before we get to it, we require one more detail.

Claim 3. *Let $k \in \{1, 2 \ldots, m\}$. Suppose α_{k-1}^- is not a coherent extension and that $\gamma_k \setminus \alpha_k^-$ has an intersection with γ_{k-1}. Then α_k^- is not coherent.*

Likewise, if α_k^+ is not coherent and $\gamma_{k-1} \setminus \alpha_{k-1}^+$ has an intersection with γ_k, then α_{k-1}^+ is not coherent.

Proof. We only prove the first statement. Because A is coherent and α_{k-1}^- is not a coherent extension, α_{k-1}^+ is a coherent extension of α_{k-1}. Let $\times$ be the intersection of $\gamma_k \setminus \alpha_k^-$ with γ_{k-1}; it follows that $\gamma_{k-1} \cap \gamma_k \supseteq \{\times_k, \times\}$. Since $\times \notin \alpha_k^-$, $\times$ is neither $\times_k$ nor a_k^-. Hypothesis (2.3.1) implies $\gamma_{k-1} \cap \gamma_k = \{\times_k, \times\}$.

Since $\alpha_k^- \setminus \{\times_k\}$ is disjoint from γ_{k-1}; in particular, it is disjoint from α_{k-1}. For $k = 1$, the preceding claim implies that α_k^- is not coherent. Thus, we suppose $k \geq 2$.

The union of α_{k-1}^- and $A[\times_{k-1}, a_{k-1}^-]$ is a simple closed curve C. Let p be a point of α_k^- near $\times_k$. From p trace an arc δ alongside α_{k-1}, across α_{k-1}^- and, continuing beside A, on to a point near the end $\times_0$ of A. Thus, δ is along the side of A faced by all the $\times_i$. Because α_{k-1}^- is not coherent, it does not return to A on this side and, therefore, δ crosses α_{k-1}^- only once. Consequently, δ crosses C only once.

Suppose by way of contradiction that α_k^- is a coherent extension of α_k. Because α_k^- does not intersect $\gamma_{k-1} \setminus \{\times_k\}$, it cannot cross C. Therefore, it does not intersect the portion of δ from its crossing with α_{k-1}^- to its end near $\times_0$. In particular, α_k^- has no end in $\alpha_0 \alpha_1 \cdots \alpha_{k-2}$. The first paragraph shows α_k^- is also disjoint from $\alpha_{k-1} \setminus \{\times_k\}$. Thus, a_k^- is in $A[\times_{k+1}, t]$, contradicting the preceding claim. $\qquad\square$

14

We are now ready to get the fine detail of the coherent extensions of A.

Claim 4. *For $0 \leq j \leq m - 1$, if α_j^+ is a coherent extension of α_j, then $a_j^+ \in \alpha_{j+1}$.*

Likewise, for $1 \leq j \leq m$, if α_j^- is a coherent extension of α_j, then $a_j^- \in \alpha_{j-1}$.

Proof. We only prove the first statement. Suppose that j is least such that α_j^+ is a coherent extension of α_j and $a_j^+ \notin \alpha_{j+1}$. We show by induction that, for each $k = 0, 1, \ldots, j$, α_k^- is not a coherent extension of α_k. For $k = 0$, Remark 2.2 shows that α_0^- is not a coherent extension of α_0. Now let $k \geq 1$ and suppose that α_{k-1}^- is not a coherent extension of α_{k-1}.

Since A is coherent and α_{k-1}^- is not a coherent extension of α_{k-1}, we have that α_{k-1}^+ is a coherent extension of α_{k-1}. By the choice of j, $a_{k-1}^+ \in \alpha_k$. The preceding claim implies that α_k^- is not a coherent extension of α_k, completing the proof that, for each $k = 0, 1, \ldots, j$, α_k^- is not a coherent extension of α_k.

Let ℓ be such that $a_j^+ \in \alpha_\ell \setminus \{\times_\ell\}$. The second claim and the choice of j show that $\ell > j + 1$. Part (1) of the first claim implies that, for each $i = j + 2, j + 3, \ldots, \ell$, α_i^- is disjoint from $\alpha_j^+ \setminus \{a_j^+\}$, while the second claim asserts that $a_j^+ \notin \alpha_i^-$. Part (2) of the first claim now implies that α_{j+1}^+ intersects α_j^+ at a point q.

We showed that α_j^- is not a coherent extension of α_j and that α_{j+1}^+ intersects γ_j at $\times_{j+1}$ and q. Consequently, the coherence of A and the preceding claim show that α_{j+1}^+ is a coherent extension. The second claim shows that, for some $r > j + 1$, $a_{j+1}^+ \in \alpha_r \setminus \{\times_r\}$. The first two claims show that, for each $i = j + 3, j + 4, \ldots, r$, α_i^- is disjoint from α_{j+1}^+.

Let α_ℓ^* be the subarc $\alpha_\ell[\times_\ell, a_j^+]$ and let A' be the arc consisting of α_{j+1}, $\alpha_{j+2}, \ldots, \alpha_{\ell-1}, \alpha_\ell^*$. Just above, we showed that α_{j+2}^- is disjoint from α_j^+. Therefore, α_{j+2}^- is a coherent extension of α_{j+2} with respect to A'. The second claim shows that α_{j+2}^- has both ends in α_{j+1}.

Since α_{j+2}^- intersects α_{j+1} in $\times_{j+2}$ and a_{j+2}^-, we see that $a_{j+1}^+ \notin \alpha_{j+2}$; therefore $r > j + 2$. Moreover, it follows that α_{j+2}^- is disjoint from α_{j+1}^+. This, together with the fact that, for each $i = j + 3, j + 4, \ldots, r$, α_i^- is disjoint from α_{j+1}^+, contradicts the first claim. $\qquad\square$

Because A is coherent and Remark 2.2 shows α_0^- is not a coherent extension of α_0, α_0^+ is a coherent extension of α_0. Likewise, α_m^- is a coherent extension of α_m. It follows that there is a $j \geq 1$ such that α_{j-1}^+ is a coherent extension of α_{j-1} and α_j^- is a coherent extension of α_j. The fourth claim

15

implies that both ends of α_{j-1}^+ are in α_j and both ends of α_j^- are in α_{j-1}. This implies that $|\gamma_{j-1} \cap \gamma_j| \geq 3$, completing the proof that $(2.3.1) \Rightarrow (2.3.2)$.

$(2.3.2) \Rightarrow (2.3.4)$. For this argument, we will make use of the following trivial observation.

Observation 2.4. *Let Q, R, S be arcs in the sphere such that R and S both have their ends in Q, but otherwise are disjoint from Q. We assume R and S have finitely many intersections and these are all crossings.*

Assume that short subarcs of S starting at each end of S are on different sides of the unique simple closed curve in $R \cup Q$. Then $(R \cap S) \setminus Q$ has at least one point.

By way of contradiction, let A be a least-weight spiral in $P(\Gamma)$ having no segment external for A. Since the only segment in an arc of weight 0 is external for that arc, A has positive weight m. Let $\alpha_0 \alpha_1 \cdots \alpha_m$ be the decomposition of A.

By assumption, A is incoherent; let α_i be an incoherent segment of A. By definition, both α_i^- and α_i^+ are both incoherent extensions. Thus, short subarcs near their ends a_i^- and a_i^+ are on the side of A that is opposite the side faced by all the crossings $\times_1, \ldots, \times_m$. We remark that one or both of a_i^- and a_i^+ might be in $\{\times_0, \times_{m+1}\}$. If $i = 0$, then only the subarc near a_0^+ is forced by incoherence to be on the side of A not faced by the crossings; the fact that γ_0 is a simple closed curve implies a_0^- is as well. An analogous statement applies if $i = m$.

As we traverse A from $\times_0$ to $\times_{m+1}$, we first encounter a_i^- and then a_i^+ (these could be equal). Therefore, either $a_i^- \in \alpha_0 \alpha_1 \cdots \alpha_{i-1}$ or $a_i^+ \in \alpha_{i+1} \alpha_{i+2} \cdots \alpha_m$. As these are symmetric up to reversal of A, we assume the latter.

Let B be the spiral $\alpha_{i+1} \alpha_{i+2} \cdots \alpha_m$; evidently, its weight is less than that of A. Therefore, B has a segment α_j that is external for B. Notice that the side of B faced by all its crossings is separated from $A[\times_0, \times_{i+1}] \setminus \{\times_{i+1}\}$ by the simple closed curve $A[\times_{i+1}, a_i^+] \cup \alpha_i^+$.

Since no segment of A is external for A, γ_j intersects A at a point outside α_j. This implies that α_j^- and α_j^+ intersect A at points a_j^- and a_j^+, respectively, that are not in α_j. Since α_j is external for B, a_j^- and a_j^+ are in $A \setminus B$.

Claim 5. *Traversing α_j^- and α_j^+ from $\times_j$ and $\times_{j+1}$, respectively, they each have an intersection with α_i^+ and these intersections are distinct points of α_i^+.*

16

Proof. In an extreme case, $j = i + 1$; here $\times_{i+1}$ is the required common point between α_j^- and α_i^+. In the other extreme case, $a_i^+ = \times_{m+1}$ and $j = m$; here $\times_{m+1}$ is the common point between α_j^+ and α_i^+. (These could happen simultaneously.) Otherwise, $\times_j$ and $\times_{j+1}$ are interior points of B.

If we take $Q = A$, $R = \alpha_i^+$ and S to be either α_j^- or α_j^+, then Observation 2.4 implies there are distinct points, one in each of $(\alpha_i^- \cap \alpha_j^+) \setminus A$ and $(\alpha_i^+ \cap \alpha_j^+) \setminus A$. $\qquad\square$

Since $|\gamma_i \cap \gamma_j| \leq 2$, the only points in $\gamma_i \cap \gamma_j$ are the points in each of $\alpha_j^- \cap \alpha_i^+$ and $\alpha_j^+ \cap \alpha_i^+$. In particular, α_j^- and α_j^+ are both disjoint from $\alpha_i^- \setminus \{a_i^-\}$.

Because $a_j^- \notin B$ and α_j^- is disjoint from $\alpha_i^- \setminus \{a_i^-\}$, $a_j^- \in A[\times_0, \times_i]$. On the other hand, the disjointness of α_j^- with $\alpha_i^- \setminus \{a_i^-\}$ implies that a_i^- is not further from $\times_0$ in A than a_j^- is. In turn, this implies that a_i^- is in $A[\times_0, \times_i] \setminus \{\times_i\}$. Therefore, reversing the direction of traversal of A, α_i^- may play the role of α_i^+ in the preceding argument.

Thus, there is a $k \in \{0, 1, \dots, i-1\}$ such that α_k is a segment of $\alpha_0 \alpha_1 \cdots \alpha_{i-1}$, external for $\alpha_0 \alpha_1 \cdots \alpha_{i-1}$, but not external for A. The argument above for α_i^+ and α_j applies to α_i^- and α_k. Thus, α_k^- and α_k^+ have their endpoints a_k^- and a_k^+, respectively, in $A[\times_{i+1}, \times_{m+1}] \setminus \{\times_{i+1}\}$. Observation 2.4 implies that each of α_j^- and α_j^+ has an intersection with each of α_k^- and α_k^+. In particular, $\gamma_j \cap \gamma_k$ has at least four points. This contradicts (2.3.1), which is equivalent to (2.3.2). $\qquad\blacksquare$

3 A pseudocircle transversal

We recall that the proof of Theorem 1.1 has two parts. For our current set Γ of pairwise intersecting pseudocircles, we must (i) find a pair of initial approximations to the pseudocircle for the next edge and (ii) show that the pair of approximations imply the existence of the desired curve crossing all the curves in Γ. In this section, our focus is on the second of these parts.

In particular, in Section 6, we will show how to find two initial approximations that together intersect all the curves in our current Γ. The main result of this section is to use these two approximations to find the single curve that intersects all the curves in Γ.

We are reminded of the theorems that are:

- Helly-type: if a collection of sets is such that every k of the sets admits a transversal, then the whole collection admits a transversal; and
- Gallai-type: if a collection of sets is such that every k of the sets admits a transversal, then the whole admits a small set of partial transversals whose union is a transversal.

Our theorem has the following different character: if a collection of sets admits a small set of partial transversals whose union is a transversal, then it admits a transversal. We do not know of another example of this type of theorem.

Definition 3.1. Let Γ be an arrangement of pseudocircles.

- A set Λ of simple closed curves is a Γ-*transversal* if every curve in Γ intersects at least one of the curves in Λ.

- A simple closed curve γ is a Γ-*pseudocircle* if $\Gamma \cup \{\gamma\}$ is an arrangement of pseudocircles.

Theorem 3.2. *Let Γ be an arrangement of pseudocircles. Let γ_1 and γ_2 Γ-pseudocircles such that $\{\gamma_1, \gamma_2\}$ is a Γ-transversal. Suppose*

(3.2.1) *$\gamma_1 \cap \gamma_2$ is a non-trivial arc and*

(3.2.2) *if $\delta_1, \delta_2 \in \Gamma$ are such that $\delta_1 \cap \gamma_1 = \varnothing$ and $\delta_2 \cap \gamma_2 = \varnothing$, then $\delta_1 \cap \delta_2 \neq \varnothing$.*

Then there exists a Γ-pseudocircle γ containing $\gamma_1 \cap \gamma_2$ and $\gamma \setminus (\gamma_1 \cap \gamma_2)$ is contained in the closure of the face F of $\gamma_1 \cup \gamma_2$ not incident with $\gamma_1 \cap \gamma_2$.

Our proof shows that one can sweep either of γ_1 or γ_2 to the required γ. In their classic paper [25], Snoeyink and Hershberger show how to sweep one curve through the others in an arrangement of pseudoarcs and pseudocircles. In particular, up to sweeping, γ is unique.

In the proof of Theorem 1.1, any two curves in the current Γ intersect, so Hypothesis (3.2.2) holds automatically.

Proof of Theorem 3.2. Let $\sigma = \gamma_1 \cap \gamma_2$ and, for $i = 1, 2$, let

$$\Gamma_i = \{\delta \in \Gamma \ : \ \delta \cap \gamma_i \neq \emptyset\}.$$

Because $\{\gamma_1, \gamma_2\}$ is a Γ-transversal, $\Gamma_1 \cup \Gamma_2 = \Gamma$. Let $n = n(\Gamma, \gamma_1, \gamma_2) = |\Gamma_2 \setminus \Gamma_1|$. Define $k = k(\Gamma, \gamma_1, \gamma_2)$ as the number of crossings in $P(\Gamma)$ included in the face F of $\gamma_1 \cup \gamma_2$ not incident with σ. We proceed by induction on $n + k$.

We can assume that neither $\Gamma_1 \setminus \Gamma_2$ nor $\Gamma_2 \setminus \Gamma_1$ is empty, else we pick γ to be either equal to γ_2 or γ_1.

If there is an arc α of some $\delta \in \Gamma_2 \setminus \Gamma_1$ incident with a face of $P(\Gamma \cup \{\gamma_1, \gamma_2\})$ that is included in F and incident with γ_1, then by shifting some part of γ_1 to cross α via a Reidemeister Type II move, we obtain a curve γ_1' such that the pair (γ_1', γ_2) satisfies the same hypothesis as (γ_1, γ_2). Since $n(\Gamma, \gamma_1', \gamma_2) + k(\Gamma, \gamma_1', \gamma_2) < n(\Gamma, \gamma_1', \gamma_2) + k(\Gamma, \gamma_1', \gamma_2)$, the result follows by induction.

In the alternative, there exists an arc A with ends in $\gamma_2 \setminus \sigma$, but otherwise contained in $F \cap P(\Gamma_1 \cup \{\gamma_1\})$, such that $\gamma_2 \cup A$ separates $\gamma_1 \setminus \sigma$ from $P(\Gamma_2 \setminus \Gamma_1) \cap F$. Let Δ_A be the closure of the component of $F \setminus A$ that is incident with both A and $\gamma_1 \setminus \sigma$. Among the finitely many choices for A, we choose A so that Δ_A is minimal under inclusion. See Figure 7.
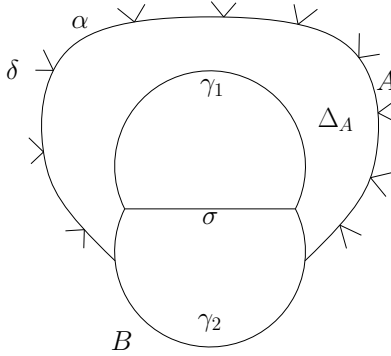

Figure 7: The region Δ_A.

We apply Theorem 2.3, specifically $(2.3.1) \Rightarrow (2.3.2)$, to see that there is a crossing in A facing Δ_A. The proof is by contradiction, assuming that every crossing faces the other side of A. In particular, A is a spiral; the contradiction arises from the following.

Claim 1. *There is a crossing in A facing Δ_A.*

Proof. By way of contradiction, suppose every crossing faces the other side of A; in particular, A is a spiral. We show that A is coherent, contradicting

19

Theorem 2.3 (specifically, (2.3.1) $\Rightarrow$ (2.3.2)). It suffices to show that if α is an arc in the decomposition of A, then α is coherent. Let δ be the curve in Γ_1 containing α.

Let B be the closed arc in $\gamma_2 \setminus \sigma$ such that $A \cup B$ is a simple closed curve. Consider the continuation of δ from one end of α. Since δ crosses γ_1 twice, the continuation must eventually reach γ_1; in particular, it must have a first intersection with $A \cup B$.

We show below that it is impossible for both continuations to have these first intersections in B. Therefore, for one of them, the first intersection is in the interior of A, showing that α is coherent.

So suppose both continuations intersect B for the first time at p_1 and p_2. Since $|\gamma_2 \cap \delta| \leq 2$, $\gamma_2 \cap \delta = \{p_1, p_2\}$. Therefore, δ is contained in the region bounded by $A \cup (\gamma_2 \setminus B)$ and intersects the boundary of this region only at α. This shows that $\delta \notin \Gamma_1$, the required contradiction. $\qquad\square$

Claim 2. *Let $\delta \in \Gamma$. Then every arc in $\delta \cap \Delta_A$ has one end in the interior of $\gamma_1 \setminus \sigma$ and one end not in the interior of $\gamma_1 \setminus \sigma$.*

Proof. Let β be an arc $\delta \cap \Delta_A$. If β has no end in $\gamma_1 \setminus \sigma$, then $\beta \cup A$ contains an arc A' that separates $\gamma_1 \setminus \sigma$ from $P(\Gamma_2) \cap F$ such that $\Delta_{A'}$ is properly contained in Δ_A, contradicting the choice of A.

If β has no end in the complement of the interior of $\gamma_1 \setminus \sigma$ in the boundary of Δ_A, then β is contained in Δ_A and has both ends in the interior of $\gamma_1 \setminus \sigma$. Since $\delta \cap \gamma_1$ has exactly two points, these are the two ends of β. The preceding paragraph shows that β is the only arc in $\delta \cap \Delta_A$, and hence $\delta \setminus \beta$ is included in the side of γ_1 disjoint from $\gamma_2 \setminus \sigma$. Thus, for any $\delta' \in \Gamma_2 \setminus \Gamma_1$, since $\delta' \cap \Delta_A$ and $\delta' \cap \gamma_1$ are empty, this implies that $\delta \cap \delta' = \emptyset$, contradicting (3.2.2). $\qquad\square$

Claim 1 implies there are distinct elements δ, δ' of Γ_1 that have a crossing $\times$ in A through which they proceed into the Δ_A-side of A. Claim 2 shows both extensions from this crossing are arcs ρ and ρ' in δ and δ', respectively, joining $\times$ to their ends a and a', respectively, in the interior of $\gamma_1 \setminus \sigma$. All intersections of δ and δ' with γ_1 are crossings, so this is true in particular for a and a'.

Since ρ and ρ' cross at $\times$, they have at most one other crossing. Let $\times^*$ be that other crossing if it exists; otherwise $\times^*$ is $\times$. The union of the subarcs of ρ from a to $\times^*$, ρ' from a' to $\times^*$, and $\gamma_1 \setminus \sigma$ from a to a' is a simple closed curve λ in the closed disc Δ_A.

We aim to show that the interiors of the arcs $\rho \cap \lambda$ and $\rho' \cap \lambda$ are not crossed by any curve in Γ. If not, then there is a $\mu \in \Gamma$ that crosses, say, $\rho \cap \lambda$. By Claim 2, the component of $\mu \cap \Delta_A$ containing this crossing has a subarc μ' with one end in ρ and one end in the interior of A. We may assume μ' has no other intersection with ρ.

There is an arc $A' \neq A$ in $A \cup \mu' \cup \rho_e$ that separates the interior of $\gamma_1 \setminus \sigma$ from $P(\Gamma_2) \cap F$. However, $\Delta_{A'}$ is a proper subset of Δ_A. This contradiction shows that no curve in Γ intersects either $\rho \cap \lambda$ or $\rho' \cap \lambda$.

It follows that we can perform the equivalent of a Reidemeister III move to shift the portion of γ_1 between a and a' across $\times^*$. (It is possible that $a = a'$. In this case, we still do the Reidemeister III move, starting by shifting γ_1 into the face bounded by λ. In the context of Theorem 1.1, this actually does not occur.) If γ_1' is the resulting curve, then $n'(\Gamma, \gamma_1', \gamma_2) = n$, $k'(\Gamma, \gamma_1', \gamma_2) = k - 1$, and hence the result follows from the induction. ∎

4   h-Convex and pseudospherical drawings

In this section we introduce h-convex drawings and prove the easy implication (1.1.2) $\Rightarrow$ (1.1.3).

The following notions were introduced by Arroyo et al. [12]. (Here, for a drawing D of a graph G, if T is a subgraph of D, then $D[T]$ denotes the subdrawing of D induced by T. If T is just an edge e and its ends or has vertex set S and no edges, then we write simply $D[e]$ or $D[S]$, respectively.)

Definition 4.1. Let D be a drawing in the sphere of the complete graph K_n in which any two edges have at most one point in common and that this point, if it exists, is either a common incident vertex or a crossing point.

(**4.1.1**) Let T be a 3-cycle in K_n. A closed disc Δ bounded by $D[T]$ is *a convex side of* $D[T]$ if, for every two vertices x, y such that $D[\{x, y\}] \subseteq \Delta$, then $D[xy] \subseteq \Delta$.

(**4.1.2**) The drawing D is *convex* if, for every 3-cycle T in K_n, $D[T]$ has a convex side.

(**4.1.3**) If D is convex, then D is *h-convex* (short for "hereditarily convex") if there is a set $\mathcal{C}$ consisting of, for every 3-cycle T of K_n, a convex side Δ_T such that, for two 3-cycles T, T', if $D[T] \subseteq \Delta_{T'}$, then $\Delta_T \subseteq \Delta_{T'}$.

(4.1.4) The drawing D is *f-convex* (short for "face convex") if there is a face F of $D[K_n]$ such that, for every 3-cycle T of K_n, the closed disc Δ_T bounded by $D[T]$ and disjoint from F is convex.

A pseudolinear drawing of K_n in the plane is evidently homeomorphic to an f-convex drawing in the sphere. The converse is proved in [10]: taking a witnessing face of an f-convex drawing of K_n in the sphere to be the outer face yields a pseudolinear drawing of K_n in the plane.

Arrangements of pseudolines naturally correspond to rank 3 oriented matroids [13, Def. 5.3.1]. Theorem 1.1 shows that an h-convex drawing of K_n is also equivalent by Reidemeister III moves to a rank 3 oriented matroid.

Clearly f-convex drawings are h-convex and h-convex drawings are convex. Evidence is given in [12] to support the conjecture that every crossing-minimal drawing of K_n is convex. A polynomial-time algorithm recognizing h-convexity follows from their result that a drawing of K_n is h-convex if and only if it does not contain as a subdrawing any of the three drawings (two of K_5 and one of K_6) shown in Figure 8. We do not need this result here and there is little overlap of this work with [12].

Ábrego et al [6, Sec. 4.3, 4.4] show that every crossing-minimal 2-page drawing D of K_n has, up to symmetry, a certain matrix representation of D. It is assumed that the *spine* of the drawing is the x-axis and each edge is drawn either in the closed upper half plane H^+ having $y \geq 0$ or in the closed lower half plane H^- having $y \leq 0$. For n even, they show there is a unique such drawing, while for n odd, there are some options.

For n even, it is quite straightforward to verify convexity and h-convexity of D by choosing a particular side of each 3-cycle xyz. If all three of xy, xz, yz are in the same one of H^+ and H^-, then the convex side of xyz is the bounded side of xyz. If xy and yz are in, say, H^+ but xz is not, then there are two cases. If y—the vertex incident with both H^+-edges—is in the middle of x, y, z in the linear order on S, then the unbounded side of xyz is the convex side. Otherwise, y is an end of the ordering of x, y, z on S and the bounded side is the convex side.

The verification of convexity is very simple; h-convexity requires some additional thinking. Since it is outside the scope of this work, we omit these discussions. Moreover, for n odd, it is not at all obvious that our arguments hold. In the case n is odd, the locations of some edges are not determined in [6]. We expect to be able to provide a detailed analysis of this case in a separate note, which would also include the analysis above for n even.
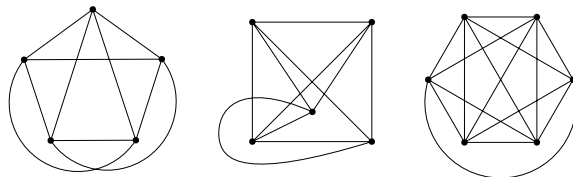
Figure 8: The three obstructions for h-convexity.

Our principal goal is to prove Theorem 1.1. As mentioned in the intro-duction, the implication $(1.1.1) \Rightarrow (1.1.2)$ is trivial. To see the implication $(1.1.2) \Rightarrow (1.1.3)$, let e, f, and g be the three edges of a 3-cycle T, contained in pseudocircles γ_e, γ_f, and γ_g, respectively. For each $x \in \{e, f\}$, (PS3) shows that the open arc $\gamma_x \setminus D[x]$ is contained in one of the two open sides of $D[T]$. Since the γ_x cross at their common vertices, it follows that one open side of $D[T]$ is disjoint from all of γ_e, γ_f, and γ_g; we set Δ_T to be the closure of this side of $D[T]$.

To see that Δ_T is convex, let u, v be vertices in Δ_T. Then u and v are on the same side of each of γ_e, γ_f, and γ_g and so (PS3) shows uv does not cross any of γ_e, γ_f, and γ_g. In particular, uv does not cross $D[T]$, so Δ_T is indeed convex.

For h-convexity, suppose a second 3-cycle uvw is drawn in Δ_T. Then (PS2) shows that, for example, γ_{uv} must cross each of γ_e, γ_f, and γ_g. There-fore, the side of $D[uvw]$ contained in Δ_T is Δ_{uvw}, as required.

The proof of the remaining implication in Theorem 1.1 is given in Section 6. Section 5 provides the necessary discussion of h-convex drawings required to obtain an appropriate initial approximation of the simple closed curve for the "next" edge of K_n to extend a (partial) collection of curves satisfying (PS1)–(PS3).

5   h-convex drawings

Our main goal now is to prove the remaining part of Theorem 1.1: an h-convex drawing of K_n extends to simple closed curves that satisfy (PS1), (PS2), and (PS3). The proof, given in the next section, requires three facts about h-convex drawings of K_n: Lemmas 5.2, 5.3, and 5.4 below. The latter two are straightforward consequences of the first. However, the proof of the

first is elucidated in this section through a fairly long series of lemmas.

The reader may skip the proof of Lemma 5.2 in order to proceed directly to the proof of Theorem 1.1.

Notation 5.1. Let D be an h-convex drawing of K_n and let $\mathcal{C}$ be a particular choice of convex sides of the 3-cycles witnessing the h-convexity of D as in Definition (4.1.3). Let e be any edge of K_n with an arbitrary orientation from one end of e to the other.

(HC1) Set Σ_e^1 to be the set of all vertices v of K_n not incident with e such that the side in $\mathcal{C}$ of the 3-cycle containing v and e is the left side, relative to the given orientation of e. The remaining vertices not incident with e have their convex side that is in $\mathcal{C}$ relative to e on the right and they make up Σ_e^2.

(HC2) For $i = 1, 2$, we set D_e^i to be the subdrawing of D of the complete subgraph induced by Σ_e^i and the ends of e.

Our next lemma is the main point of this section. Its proof follows its two straightforward consequences.

Lemma 5.2. *Let D be an h-convex drawing of K_n with a specified witnessing set of convex sides. For $i = 1, 2$, if Σ_e^i is not empty, then there is a closed disc Δ_e^i containing D_e^i and bounded by a cycle C_e^i of K_n containing e and whose vertices are otherwise contained in Σ_e^i. Furthermore, $\Delta_e^1 \cap \Delta_e^2$ consists of $D[e]$ and its ends.*
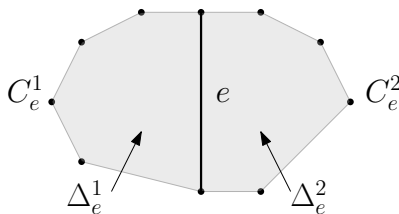

Figure 9: Illustrating the consequence of Lemma 5.2.

Figure 5 illustrates the conclusion of Lemma 5.2. Before the proof, we give two simple consequences, which are also used in the next section.

The first simple consequence is about edges not in either Δ_e^1 or Δ_e^2. Recall that $\mathbb{S}^2$ denotes the sphere. For $e \in E(K_n)$, set $F_e = \mathbb{S}^2 \setminus (\Delta_e^1 \cup \Delta_e^2)$.

Lemma 5.3. *Let D be an h-convex drawing of K_n with a specified witnessing set of convex sides and e and e' be distinct edges of K_n. If $D[e']$ has a point in F_e, then e' has an end in each of Σ_e^1 and Σ_e^2.*

Proof. Let u and v be the ends of e. We prove the contrapositive. Suppose that for some $k \in \{1, 2\}$, e' has both ends in $\{u, v\} \cup \Sigma_e^k$. Then $D[e']$ is contained in Δ_e^k by Lemma 5.2, so $D[e'] \cap F_e = \varnothing$. ∎

We need some notation for the next lemma. For distinct edges e, e' of K_n, label each vertex of C_e^1 (from Lemma 5.2) with 1, 2, 3, respectively, to indicate it is in $\Sigma_{e'}^1$, in $\Sigma_{e'}^2$, or incident with e'.

Lemma 5.4. *Suppose D is an h-convex drawing of K_n with a specified witnessing set of convex sides, e, e' edges of K_n, and the labelling of C_e^1 as in the preceding sentence. Then there is no $1, 2, 1, 2$ pattern in the cyclic order around C_e^1.*

Proof. Otherwise, there are four vertices v_1, v_2, v_3, v_4 of C_e^1 in this cyclic order with v_1, v_3 having label 1 and v_2, v_4 having label 2. As all v_i are in Σ_e^1, the definition of C_e^1 implies $v_1 v_3$ crosses $v_2 v_4$ in D_e^1.

However, the edge $v_1 v_3$ is in $D_{e'}^1$, while $v_2 v_4$ is in $D_{e'}^2$. Lemma 5.3 shows that both $v_1 v_3$ and $v_2 v_4$ are contained in Δ_e^1, so they cross in D. However, Lemma 5.2 implies they do not cross, a contradiction. ∎

The remainder of this section is devoted to the proof of Lemma 5.2. In the arguments below, we will use, without particular reference, the following observation: if e and f are crossing edges, then each of the four 3-cycles in the unique K_4 containing e and f has a side (the one containing the fourth vertex of the K_4) that is definitely not convex. Our drawings are convex, so, for such a 3-cycle, the convex side that is in $\mathcal{C}$ is determined. If C is the 4-cycle bounding a face of this K_4, then the closed disc bounded by $D[C]$ and containing $D[\{e, f\}]$ is the *crossing side* of $D[C]$.

Lemma 5.5. *Let D be an h-convex drawing of K_n with a specified witnessing set of convex sides, and let e be an edge of K_n. Let $i \in \{1, 2\}$ and suppose $|\Sigma_e^i| \geq 1$.*

    (i) In D_e^i, e is not crossed. In particular, there is a face F_e^i of D_e^i incident with e and containing $\Sigma_e^{(3-i)}$.

(ii) If J is any crossing K_4 in D_e^i, then F_e^i is contained in the face of $D[J]$ bounded by a 4-cycle. In particular, no crossing of D_e^i is incident with F_e^i, so F_e^i is bounded by a cycle in D_e^i.

Proof. If e were crossed by an edge f in D_e^i, then the crossing K_4 shows that the two ends of f are not both in the same one of Σ_e^1 and Σ_e^2, a contradiction. Thus, e is not crossed and is incident with exactly two faces of D_e^i; F_e^i is the one containing $\Sigma_e^{(3-i)}$, completing (i).

For (ii), let C be the 4-cycle in J bounding a face of $D[J]$. We rule out one trivial case immediately. If e is in J, then, since e is not crossed in D_e^i, it is in C. The convex sides (these are unique and in $\mathcal{C}$) of each 3-cycle in J are all on the crossing side of $D[C]$. Since the two vertices of J not incident with e are in Σ_e^i, F_e^i is contained in the face of $D[J]$ bounded by C, as required.

Therefore, we may assume there is a vertex u incident with e and not in J. As e is not crossed in D_e^i, $D[e]$ is contained in one of the faces F of $D[J]$. Since $D[u]$ is incident with F_e^i, F is also the face of $D[J]$ containing F_e^i. We are done if F is bounded by a 4-cycle in J, so we assume, by way of contradiction, that F is incident with the crossing of $D[J]$.

Convexity implies that the edges joining u to the vertices of J are all contained on the crossing side of $D[C]$. Thus, C, u, and these four edges constitute a planar embedding of the 4-wheel W. Each of the four 3-cycles in W has its convex side on the crossing side of $D[C]$: three of these 3-cycles are contained in convex sides of 3-cycles of J, so for them the assertion follows from h-convexity; the fourth is in a crossing K_4 on the crossing side of $D[C]$.

If e is one of the edges of W, then the end of e different from u has two neighbours in C that are in different ones of Σ_e^1 and Σ_e^2, a contradiction. If e is not in W, then its other end v is in one of the four faces of $D[W]$ incident with u. This implies that v is on the convex side of the bounding 3-cycle and the two vertices of C in this 3-cycle are in different ones of Σ_e^1 and Σ_e^2, a contradiction. ∎

For $i = 1, 2$, if $\Sigma_e^i \neq \varnothing$, then let C_e^i be the cycle in D_e^i that is the boundary of F_e^i. All its vertices not incident with e are, by definition, in Σ_e^i. Note that, with F_e as defined preceding Lemma 5.3, $F_e = F_e^1 \cap F_e^2$.

We are aiming to show that, for any three vertices in D_e^i, the specified convex side that they bound is contained in the side of $D[C_e^i]$ not containing F_e^i. In particular, this shows that D_e^i is f-convex. This is our next lemma.

We remark that the main result of [10] further implies that D_e^i is pseudo-linear. Thus, any edge e of an h-convex drawing partitions the vertices into

26

two pseudolinear subdrawings D_e^1 and D_e^2. This generalizes the fact that, in a spherical drawing, for each great circle C that contains an edge, the vertices in either closed side of C induce a rectilinear drawing.

Lemma 5.6. *Let D be an h-convex drawing of K_n with witnessing set $\{\Delta_T \mid T$ is a 3-cycle of $K_n\}$ of convex sides. Let e be an edge of K_n and let $i \in \{1, 2\}$. With D_e^i as in Notation 5.1, if $|\Sigma_e^i| \geq 1$, then, for each 3-cycle T in D_e^i, $\Delta_T \cap F_e^i = \varnothing$. In particular, D_e^i is f-convex.*

Proof. Let uvw be a 3-cycle in D_e^i. Suppose first that there is an edge f with both ends in C_e^i that crosses uvw. If f has one end in uvw, then Lemma 5.5 (ii) shows the crossing in the K_4 that includes u, v, w, and f is separated from F_e^i by the face-boundary 4-cycle in the K_4. Therefore, the convex side of each of the four 3-cycles in the K_4 is the side that is disjoint from F_e^i. In particular, this holds for uvw, as required.

In the remaining case, both vertices incident with f are in the side of uvw that contains F_e^i. In this case, Definition 4.1 (4.1.1) of convex side shows it is the other side, the one disjoint from F_e^i, that is convex, as required. Therefore, we can assume no edge having both ends in C_e^i crosses uvw.

Suppose that C_e^i has at least four vertices and let a and b be any two vertices of C_e^i, neither of which is an end of e, and consider the K_4 containing a, b, and e. As a, b, and e are all incident with F_e^i, this K_4 has a face incident with all four of its vertices. It follows that this is a crossing K_4. Lemma 5.5 implies that the crossing in this K_4 is separated from F_e^i by the face-bounding 4-cycle. Thus, all the 3-cycles in this K_4 have their convex side disjoint from F_e^i. Although this was already known for the two 3-cycles containing e, we now know it for the two 3-cycles containing the edge ab.

Finally, if C_e^i has only three vertices, then the result follows from h-convexity. Otherwise, let y be one of the ends of e and consider C_e^i together with all the chords from y. By the preceding paragraph, all of the 3-cycles using two of these edges incident with y and an edge of C_e^i have their convex side disjoint from F_e^i. From the earlier discussion, none of these chords crosses uvw. It follows that uvw is contained in the convex side of one of them; this convex side is disjoint from F_e^i. Thus, the chosen convex side for uvw is, by h-convexity, disjoint from F_e^i. ∎

The remaining detail about h-convex drawings we need is that $D_e^1 \subseteq F_e^2$ and $D_e^2 \subseteq F_e^1$. As mentioned after the proof Lemma 5.5, C_e^i is the boundary

of the face F_e^i of D_e^i. The other closed disc in the sphere bounded by $D[C_e^i]$ is denoted Δ_e^i. Evidently, $D[\Sigma_e^i] \subseteq \Delta_e^i$. We begin by showing that $\Sigma_e^2 \cap \Delta_e^1 = \varnothing$.

Lemma 5.7. *Let D be an h-convex drawing of K_n with a specified witnessing set of convex sides. Let $e = uv$ be an edge of K_n and $i \in \{1,2\}$, and let Σ_e^i and D_e^i be as in Notation 5.1. For $w \in V(G) \setminus \{u,v\}$, if $D[w] \subseteq \Delta_e^i$, then $w \in \Sigma_e^i$.*

Proof. Suppose that $D[w] \subseteq \Delta_e^1$. If C_e^1 is a 3-cycle, then its convex side Δ_e^1 contains w; h-convexity implies that the 3-cycle including w and e is contained in Δ_e^1, and hence $w \in \Sigma_e^1$. Thus, we may assume that C_e^1 has length at least 4.

For each edge $ab \in C_e^1 - e$, let J_{ab} denote the crossing K_4 in D induced by e and ab. The closed disc Δ_e^1 is the union of the crossing sides of the J_{ab}, so $D[w]$ is contained in the crossing side of some J_{ab}. Since J_{ab} is a crossing K_4, the convex sides of all the 3-cycles in J_{ab} are determined.

If $D[w]$ is contained in the convex side of one of the 3-cycles $D[auv]$ or $D[buv]$, then it follows from h-convexity that this side contains the convex side of $D[wuv]$ that is in $\mathcal{C}$, and thus $w \in \Sigma_e^1$. Therefore we may assume that $D[w]$ is contained in the convex sides of both $D[abu]$ and $D[abv]$. Consequently, $D[wu]$ is contained in the convex side of $D[abu]$ and $D[wv]$ is contained in the convex side of $D[abv]$. The K_4 with vertices a, w, u, v has a crossing in D and determines the convex side of the 3-cycle containing w and e, and therefore shows that $w \in \Sigma_e^1$. ∎

Next we move on to edges. The following result is preparatory to showing edges of D_e^1 and D_e^2 do not cross.

Lemma 5.8. *Let D be an h-convex drawing of K_n with a specified witnessing set of convex sides. If, for $i = 1, 2$, $x_i \in \Sigma_e^i$, then the 3-cycles induced by x_1, e and x_2, e do not cross in D.*

Proof. Let J be the K_4 induced by x_1, x_2, e and, for $i = 1, 2$, let T_i be the 3-cycle induced by x_i, e. If $D[J - x_1 x_2]$ has a crossing, then T_1 and T_2 cross but e is not crossed. The contradiction is that x_1 and x_2 are on the same side of e. ∎

We are now ready for the next major step.

Lemma 5.9. *Let D be an h-convex drawing of K_n with a specified witnessing set of convex sides. For $i \in \{1, 2\}$, let D_e^i be as in Notation 5.1. Then no edge of D_e^2 crosses any edge of D_e^1.*

Proof. By way of contradiction, suppose some edge $D[x_2y_2]$ of D_e^2 crosses some edge $D[x_1y_1]$ of D_e^1. Lemma 5.8 implies not both $\{x_1, y_1\}$ and $\{x_2, y_2\}$ can contain an end of e. Without loss of generality, we assume neither x_1 nor y_1 is an end of e. Furthermore, $x_2y_2 \neq e$, so we may choose the labelling such that x_2 is not an end of e. Let J_1 be the K_4 induced by x_1, y_1, e.

Lemma 5.8 implies that x_2y_2 does not cross any edge of J_1 incident with an end of e, so the only crossing of x_2y_2 with J_1 is with x_1y_1. Let F be the face of $D[J_1]$ containing F_e^1. Lemma 5.7 shows that $D[x_2] \in F$, and similarly that $D[y_2] \in F$ if y_2 is not an end of e. As we traverse $D[x_2y_2]$ from $D[x_2]$, we cross $D[x_1y_1]$ once, and cross nothing else in $D[J_1]$. Therefore, $D[x_1y_1]$ is incident with F, as is $D[e]$. Since the face F of $D[J_1]$ is incident with all four vertices of J_1, it follows that J_1 is a crossing K_4 in D.

Now, just after we traverse $D[x_2y_2]$ across $D[x_1y_1]$, we are in a face of $D[J_1]$ incident with the crossing of $D[J_1]$ and with both $D[x_1]$ and $D[y_1]$. This face is not incident with either end of e, nor is it equal to F. But, y_2 is either an end of e or $D[y_2]$ lies in F, so $D[x_2y_2]$ must cross $D[J_1]$ a second time, which is a contradiction. ∎

We conclude our study of h-convex drawings with the proof of Lemma 5.2.

Proof of Lemma 5.2. As in the paragraph immediately following the proof of Lemma 5.5, for $i = 1, 2$, let C_e^i be the cycle in D_e^i that is the boundary of F_e^i. Furthermore, let Δ_e^i be the closed disc bounded by C_e^i that contains $D[\Sigma_e^i]$.

The result is an application of the following simple fact about curves in the sphere.

Observation 5.10. *Let Δ_1 and Δ_2 be closed discs in the sphere bounded by simple closed curves γ_1 and γ_2, respectively. If:*

- *$\gamma_1 \cap \gamma_2$ is an arc (or empty); and*

- *$\gamma_1 \not\subseteq \Delta_2$ and $\gamma_2 \not\subseteq \Delta_1$,*

then $\Delta_1 \cap \Delta_2 = \gamma_1 \cap \gamma_2$.

Lemmas 5.7 and 5.9 imply that, for $\{i, j\} = \{1, 2\}$, the open arc $D[C_e^i] \setminus D[e]$ is disjoint from Δ_e^j. The result is an immediate application of the preceding paragraph. ∎

6 Proof of Theorem 1.1

In this section, we prove (1.1.3)$\Rightarrow$(1.1.1): an h-convex drawing of K_n has simple closed curve extensions of the edges satisfying (PS1), (PS2), and (PS3). This completes the proof of Theorem 1.1.

The proof iteratively constructs the set of simple closed curve extensions of the edges. We assume that, for some $J \subset E(K_n)$ and for all $e \in J$, there exist extensions γ_e satisfying (PS1), (PS2), (PS3), and a fourth property (PS4):

(PS4) For each $e \in J$ and each $e' \in E(K_n) \setminus \{e\}$, γ_e intersects the closed edge $D[e']$ at most once, and, if it exists, the point of intersection is either a crossing or a vertex incident with both e and e'.

Notice that if $J = E(K_n)$, then the extensions of the edges in J automatically satisfy (PS4) provided they satisfy (PS1), (PS2), (PS3). The extra assumption (PS4) is required for inductive purposes.

We pick any $e_0 \in E(K_n) \setminus J$; the extension of γ_{e_0} is obtained as a result of Theorem 3.2. Thus, we need to find the two initial curves $\gamma_{e_0}^1$ and $\gamma_{e_0}^2$ satisfying the hypotheses of Theorem 3.2 with respect to $\{\gamma_e \mid e \in J\}$. The curve $\gamma_{e_0}^i$ contains $D[e_0]$ and is completed by an arc joining the ends of e_0 that is in F_{e_0} and "very near" the path $C_{e_0}^i - e_0$; this is where we use the specified convex sides of an h-convex drawing. Both curves are contained in $D[e_0] \cup F_{e_0}$. How "near" is "very near" will depend on the curves that are already determined. Our next lemma is the crucial point; Corollary 6.2 provides $\gamma_{e_0}^1$ and $\gamma_{e_0}^2$.

Lemma 6.1. *Let D be an h-convex drawing of K_n with witnessing set $\mathcal{C}$ of convex sides. Let $J \subseteq E(K_n)$ and suppose that, for each $e \in J$, there is a simple closed curve γ_e in $D[e] \cup F_e$ containing $D[e]$, and such that the extensions $\{\gamma_e \mid e \in J\}$ satisfy (PS1) and (PS4). If $e_0 \in E(K_n) \setminus J$, then, for any $i \in \{1, 2\}$ and any sufficiently small neighbourhood N of $D[C_{e_0}^i]$ in $F_{e_0} \cup D[C_{e_0}^i]$, there are at most two arcs in $(\gamma_e \cap N) \setminus D[C_{e_0}^i]$. Furthermore, at most one of these segments is contained in $D[e]$.*

Proof. We begin with the central claim. For the sake of definiteness, we assume $i = 1$.

Claim 1. *If u is a vertex incident with e and is in $C^1_{e_0}$, then there is no arc of γ_e contained in the interior of $\Delta^1_{e_0}$ that joins two points in $D[C^1_{e_0}]$ neither of which is $D[u]$.*

Proof. Suppose to the contrary that there is such an arc α. It follows from $\alpha \subseteq \gamma_e$ and (PS4) that there is no closed edge of $D[C^1_{e_0}]$ containing both ends y, z of α. Thus, each component of $D[C^1_{e_0}] \setminus \{y, z\}$ has a vertex of $C^1_{e_0}$. Let $D[w]$ be a vertex in the component of $D[C^1_{e_0}] \setminus \{y, z\}$ not containing $D[u]$.

Both u and w are in $C^1_{e_0}$, so both are drawn in $\Delta^1_{e_0}$. Therefore, $D[uw] \subseteq \Delta^1_{e_0}$, so $D[uw]$ crosses α. Since $\alpha \subseteq \gamma_e$, $D[e]$ cannot cross α, so $uw \neq e$. But the edge $D[uw]$ has the two points $D[u]$ and the crossing with α in γ_e, contradicting (PS4) and completing the proof. □

Suppose that $\gamma_e \cap D[C^1_{e_0}]$ contains a vertex $D[u]$, which must be incident with e since γ_e satisfies (PS1). If both ends of e are in $C^1_{e_0}$, then $D[e] \subseteq D^1_{e_0}$ since $D^1_{e_0}$ is the subdrawing of D induced by $\Sigma^1_{e_0}$. Furthermore, the claim implies that the ends of e are the only intersections of γ_e with C_{e_0}. Therefore $D[e]$ is the only segment of γ_e contained in $\Delta^1_{e_0}$.

We may therefore assume that u is the only end of e in $C^1_{e_0}$. Then there are only two directions from $D[u]$ in γ_e; each of these directions can give an arc in $\gamma_e \cap \Delta^1_{e_0}$ having $D[u]$ as an end. By Claim 1, these are the only possible intersections of γ_e with $\Delta^1_{e_0}$. In this case, $\gamma_e \cap \Delta^1_{e_0}$ is either one arc, with $D[u]$ as either an end or an interior point, or $\gamma_e \cap \Delta^1_{e_0}$ is just $D[u]$.

It follows that if an end of e is in $C^1_{e_0}$, then the two cases above show that $\gamma_e \cap \Delta^1_{e_0}$ is either a non-trivial arc or $D[u]$. In the former case, only the ends of this arc can be the start of a segment of γ_e from a point of $D[C^1_{e_0}]$ into F_{e_0}, as required. In the latter case, $\gamma_e \cap \Delta^1_{e_0} = D[u]$, and there are exactly two arcs of γ_e having an end in $D[u]$ and extending into F_{e_0}.

Thus, we may assume that no point of $\gamma_e \cap D[C^1_{e_0}]$ is a vertex. In this case, (PS4) implies that every intersection of γ_e with $D[C^1_{e_0}]$ is a crossing. To prove that there are at most two segments of γ_e from a point of $D[C^1_{e_0}]$ into F_{e_0}, it suffices to prove that γ_e has at most two crossings with $D[C^1_{e_0}]$.

Suppose by way of contradiction that there are three crossings of γ_e with $D[C^1_{e_0}]$. Traverse $D[C^1_{e_0}]$ in one direction from such a crossing z. The first vertex $D[v]$ reached is incident with an edge f of $C^1_{e_0}$ such that both $D[v]$

31

and z are in the closed edge $D[f]$. Since $z \in \gamma_e$, none of the rest of $D[f]$ (including $D[v]$) is in γ_e. In particular, $D[v]$ is in $\Sigma_e^1 \cup \Sigma_e^2$.

Suppose $D[w]$ is the first vertex reached from z traversing $D[C_{e_0}^1]$ in the other direction, that is, if w is the other vertex incident with f. Because $\gamma_e \subseteq D[e] \cup F_e$, $D[v]$ and $D[w]$ are on different sides of γ_e. Therefore, they are in different ones of Σ_e^1 and Σ_e^2. Thus, every crossing of γ_e with $D[C_{e_0}^1]$ produces a change between 1 and 2 in the "1,2,3"-labelling of Lemma 5.4. Let z_1, z_2, z_3 be three crossings of γ_e with $D[C_{e_0}^1]$, in this cyclic order. Then the vertices of $C_{e_0}^1$ nearest each z_k have labels 1 and 2. Starting at z_1, we find 1 and 2 near it. Up to relabelling, we may assume the 1 occurs between z_1 and z_3 and the 2 between z_1 and z_2. Then choose the 1 near z_2 and the 2 near z_3 to obtain a 1,2,1,2 pattern, contradicting Lemma 5.4. ∎

The following corollary is a straightforward consequence of Lemma 6.1.

Corollary 6.2. *Let D be an h-convex drawing of K_n with witnessing set $\mathcal{C}$ of convex sides. Let $J \subseteq E(K_n)$ and suppose that, for each $e \in J$, there is a simple closed curve γ_e in $D[e] \cup F_e$ containing $D[e]$, and such that the extensions $\{\gamma_e \mid e \in J\}$ satisfy (PS1) and (PS4). For any $i \in \{1, 2\}$ and any sufficiently small neighbourhood N of $D[C_{e_0}^i]$ in $F_{e_0} \cup D[C_{e_0}^i]$, there is a choice of $\gamma_{e_0}^i$ in N such that the curves in $J \cup \{\delta_{e_0}^i\}$ satisfy (PS1), (PS2w), (PS3), and (PS4).* ∎

We are now ready for the proof of Theorem 1.1.

Proof of Theorem 1.1 (1.1.3) implies (1.1.1). Suppose $J \subseteq E(K_n)$ and we have, for each $e \in J$, a simple closed curve γ_e, such that the set $\{\gamma_e \mid e \in J\}$ satisfies (PS1)–(PS4). If $J = E(K_n)$, then we are done; otherwise, let $e_0 \in E(K_n) \setminus J$.

We show there is a curve γ_{e_0} containing $D[e_0]$ and otherwise in the face F_{e_0} of $D_{e_0}^1 \cup D_{e_0}^2$ bounded by $(C_{e_0}^1 - e_0) \cup (C_{e_0}^2 - e_0)$ and such that $\{\gamma_e \mid e \in J \cup \{e_0\}\}$ satisfies (PS1)–(PS4).

Let M consist of those $e \in E(K_n) \setminus (J \cup \{e_0\})$ such that $D[e] \cap F_{e_0} \neq \varnothing$. In any order, repeatedly use Corollary 6.2 to obtain, for all $e \in M$, δ_e^1 so that the curves in the set $\Gamma = \{\gamma_e \mid e \in J\} \cup \{\delta_e^1 \mid e \in M\}$ satisfy (PS1), (PS2w), (PS3), and (PS4).

For each $e \in J \cup M$, $\gamma_e \setminus e$ is in the face F_e of $D_e^1 \cup D_e^2$ bounded by $(C_e^1 - e) \cup (C_e^2 - e)$, Σ_e^1 is on one side of γ_e and Σ_e^2 is on the other side of γ_e. Let $\delta_{e_0}^1$ be as in Corollary 6.2 with respect to e_0 and J; this is our first

approximation to γ_{e_0}. Applying Corollary 6.2 to $C_{e_0}^2$, consider an analogous curve δ_e^2. By Corollary 6.2, both $\Gamma \cup \{\delta_{e_0}^1\}$ and $\Gamma \cup \{\delta_{e_0}^2\}$ satisfy (PS1), (PS3), (PS4). Moreover, $\delta_{e_0}^1$ and $\delta_{e_0}^2$ intersect each curve in Γ at most twice and all intersections are crossings.

For $i = 1, 2$, let $\Gamma_i = \{\delta \in \Gamma : \delta \cap \delta_{e_0}^i \neq \emptyset\}$.

Claim 1. *Either $\Gamma = \Gamma_1 \cup \Gamma_2$ or D is f-convex.*

Proof. Suppose that for some $e \in J \cup M$ the extension $\delta \in \Gamma$ of e is not in $\Gamma_1 \cup \Gamma_2$. From Lemma 5.3 it follows that any edge of M crosses both $\delta_{e_0}^1$ and $\delta_{e_0}^2$. Therefore $e \in J$ and $\delta = \gamma_e$. Since γ_e does not intersect $\delta_{e_0}^1 \cup \delta_{e_0}^2$, we conclude that γ_e is disjoint from F_{e_0}.

Recall that, for $\ell = 1, 2$, $\Delta_{e_0}^\ell$ is the closed disc in $D_{e_0}^\ell$ bounded by $C_{e_0}^\ell$ and disjoint from F_{e_0}. The preceding paragraph implies that, for some $\ell \in \{1, 2\}$, γ_e is contained in $\Delta_{e_0}^\ell$. It follows that: (i) e has both ends in $\Delta_{e_0}^\ell$; and (ii) every vertex of $C_{e_0}^\ell$ is in the same one of Δ_e^1 and Δ_e^2 (because, by assumption, γ_e separates Δ_e^1 from Δ_e^2).

If an edge xy crosses e, then x and y are not both in the same one of Σ_e^1 and Σ_e^2. Therefore, (ii) implies that, if x, y are vertices in $C_{e_0}^\ell$, then xy does not cross e. In particular, letting z be one end of e_0 and letting xy run through the edges of $C_{e_0}^\ell$, the 3-cycles xyz bound convex sides that cover $\Delta_{e_0}^\ell$. It follows that e is contained in one of these; let it be xyz.

Suppose by way of contradiction that e has an end u that is not one of x, y, z. Then h-convexity implies that the convex sides in $\mathcal{C}$ of the 3-cycles uxy, uxz, and uyz are all contained in the convex side of xyz.

Let v be the other end of e. If v is one of x, y, z, then the resulting planar K_4 shows that e has the two vertices in $\{x, y, z\} \setminus \{v\}$ not both in the same one of Σ_e^1 and Σ_e^2, a contradiction. Likewise, if v is not one of x, y, z, then the one of the 3-cycles uxy, uxz, and uyz containing v on its convex side has its two vertices from x, y, z not both in the same one of Σ_e^1 and Σ_e^2, a contradiction. These contradictions show that both ends of e are among x, y, z; that is, both ends of e are in $C_{e_0}^\ell$.

Next, suppose by way of contradiction, that e is not an edge of $C_{e_0}^\ell$. Then it is a chord of $C_{e_0}^\ell$ in $D_{e_0}^\ell$, and so it crosses an edge xy with x and y in $C_{e_0}^\ell$ on different sides of e. But then we have x and y are not both in the same one of Σ_e^1 and Σ_e^2, a contradiction.

Lemma 5.6 shows that $D_{e_0}^\ell$ (using the convex sides in $\mathcal{C}$) is f-convex. Since e is in $C_{e_0}^\ell$, it follows that the vertices of $D_{e_0}^\ell$ not incident with e are, for some $k \in \{1, 2\}$, all in Σ_e^k. Since $\delta_j \subseteq \Delta_{e_0}^\ell$, all vertices of $\Delta_{e_0}^{3-\ell}$ are in the same

33

Σ_e^k as the two vertices incident with e_0. It follows that all vertices of K_n not incident with e are in the same Σ_e^k, showing that D is f-convex, as claimed.
$\square$

In the case D is f-convex, [10, Thm. 1] shows that D is homeomorphic to a pseudolinear drawing in the plane. By definition, the pseudolines intersect once in the plane, and they can be chosen so that they all cross again at the point at infinity that completes the sphere.

Thus we may assume that $\Gamma_1 \cup \Gamma_2 = \Gamma$. In this case we apply Theorem 3.2 to Γ, $\sigma = D[e_0]$, $\gamma_1 = \delta_{e_0}^1$ and $\gamma_2 = \delta_{e_0}^2$ to obtain a curve γ_{e_0} such that $\{\gamma_e \mid e \in J \cup \{e_0\}\}$ satisfies (PS1), (PS2), (PS3), and (PS4), as desired. $\blacksquare$

A given h-convex drawing D may have different choices for the convex sides of the 3-cycles that witness h-convexity. In Section 5, the extensions of D into arrangements of pseudocircles rely on a choice of convex sides witnessing h-convexity. Moreover, the proof of the implication (1.1.2) $\Rightarrow$ (1.1.3) shows how the choice of convex sides can be recovered from such an arrangement of pseudocircles. Define two such arrangements of pseudocircles to be *equivalent* if they determine the same convex sides.

The sweeping theorem of Snoeyink and Hershberger [25] shows that either of two equivalent arrangements of pseudocircles can be shifted to the other by a sequence of Reidemeister II and III moves. Simple examples show that Type II moves may be required.

7 Non-extendible drawings of K_9 and K_{10}

In this section, we present a drawing of each of K_9 and K_{10}. The drawing D_9 of K_9 in Figure 10 has an extension to an arrangement of pseudocircles (that is, (PS2w)) that satisfies (PS1), but no such extension also satisfies (PS2). The drawing D_{10} of K_{10} in Figure 10 does not have an extension to an arrangement of pseudocircles.

The analyses of the drawings D_9 and D_{10} involve spirals in the drawings D_1 and D_2 in Figure 11. For convenience, we extend the notion of spiral to closed spiral. For an arrangement Γ of simple closed curves, a *closed spiral* is a simple closed curve γ in $P(\Gamma)$ with basepoint s such that, for every sufficiently small open interval I in γ containing s, $\gamma \setminus I$ is a spiral.

To see the point of closed spirals, consider the drawing D_1. The unique simple closed curve γ is D_1 is a closed spiral, whose basepoint is the vertex
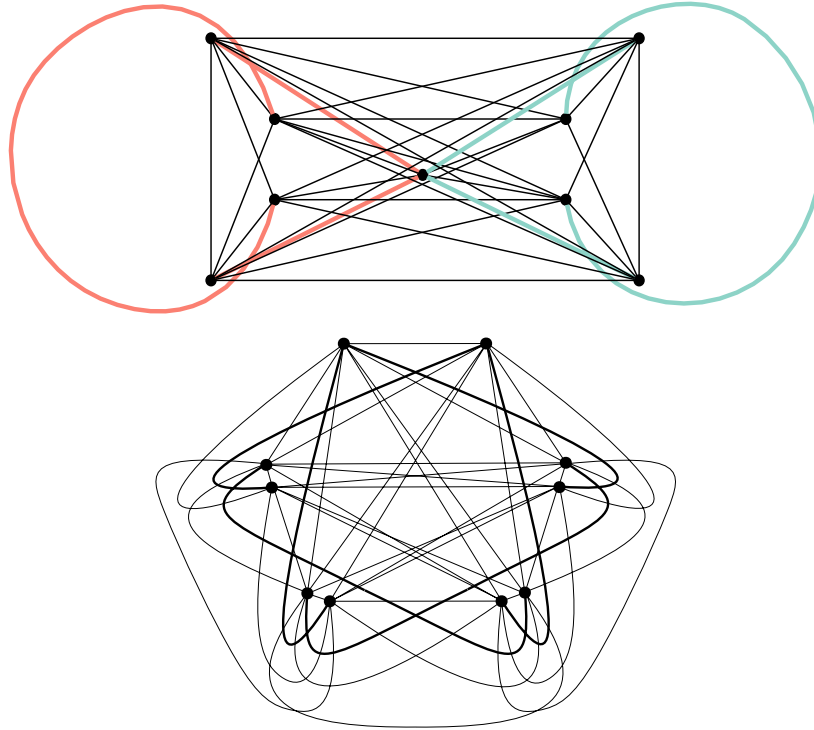
Figure 10: The upper drawing D_9 of K_9 has a pseudocircular extension, but none satisfying (PS2). The dark and light gray (red and green online) subdrawings in D_9 are both isomorphic to D_1 (see Figure 11 below). The lower drawing D_{10} of K_{10} has no pseudocircular extension. In the drawing of K_{10}, the thicker edges are the drawing D_2 (Figure 11).

in γ. Let Γ be an arrangement of pseudocircles extending D_1. Then, for any sufficiently small open interval I in γ containing s, I does not contain any crossing of Γ other than s, and $\gamma \setminus I$ is a spiral in $P(\Gamma)$. By Theorem 2.3, $\gamma \setminus I$ has an external segment, which corresponds to a simple closed curve in $P(\Gamma)$ that is contained in the interior of γ.

It follows that, any extension of D_9 to an arrangement of pseudocircles must have a pseudocircle in the bounded side of each of the coloured copies of D_1, showing it does not satisfy (PS2).

On the other hand, we can extend the straight edges into lines and extend the two curved edges by line segments connecting their vertices. In the

Figure 11: The crucial configurations D_1 on the left and D_2 on the right.

sphere, adding the point at infinity to the straight lines gives an extension of the drawing of K_9 to an arrangement of simple closed curves satisfying (PS1) and (PS2w). Some of the pseudocircles may intersect tangentially at the point at infinity since we did not assume that the lines are in general position, but this can be corrected if we perturb the curves. Thus, without (PS3), (PS1) and (PS2w) do not imply (PS1) and (PS2).

For the drawing D_{10}, we show that there is no arrangement of pseudocircles extending the drawing D_2 on the right in Figure 11. Since D_2 is contained in D_{10}, this implies there is no arrangement of pseudocircles extending D_{10}.

Let γ denote the unique simple closed curve in D_2 containing the ten crossings and none of the vertices. Let s be the upper most crossing in the diagram. Then γ is a closed spiral.

For any sufficently small open interval I in γ containing s, $\gamma \setminus I$ is a spiral in $P(\Gamma)$ that has weight 7 with decomposition $\alpha_0 \alpha_1 \ldots \alpha_7$. The drawing D_2 already shows that the segments $\alpha_1, \ldots, \alpha_7$ are coherent. The segments α_0 and α_7 are symmetric.

The extension α_0^+ is contained in pseudocircle γ_0 containing α_0. We note that γ_0 contains a vertex outside of our original simple closed curve γ. Therefore, the choice of I shows α_0^+ has its end on $\gamma \setminus I$. Consequently, α_0 is also coherent. Likewise, α_7 is coherent, showing $P(\Gamma)$ contains a coherent spiral, contradicting Theorem 2.3. That is, D_2 and, consequently, the drawing of K_{10} do not have extensions to arrangements of pseudocircles.

8   Conclusion

In this section, we mention a few results and questions about drawings of K_n related to this work. We start with the questions.

The drawing D_{10} in Figure 10 has no extension to an arrangement of

pseudocircles. It is natural to wonder if there is a fixed k such that every drawing of K_n has an extension to an arrangement of simple closed curves that pairwise cross at most k times. In fact, for a drawing of any simple graph, there is such an extension with $k \leq 4$. See Figure 12 for an idea of how such an extension may be achieved.
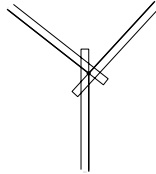

Figure 12: Extensions near a vertex.

Some interesting questions remain unresolved.

Question 1. Does every convex drawing of K_n have an extension to simple closed curves pairwise crossing at most twice? exactly twice? The drawings in Figure 10 are not convex.

Question 2. Arroyo et al [9] characterize drawings of (not necessarily complete) graphs whose edges extend to an arrangement of pseudolines by giving the complete (infinite) list of obstructions. Given the close connection we developed here between pseudospherical and pseudolinear drawings for complete graphs, it is reasonable to wonder if there is an analogous theorem for "arrangements of pseudocircles".

We conjecture that there is a list-of-obstructions characterization of when an arbitrary graph has an extension satisfying (PS1), (PS2) and (PS3). It is not clear to us at this juncture how to proceed with this.

The study of spirals and Theorem 2.3 played an essential role in the proof of Theorem 1.1 and also led to the drawings D_1 and D_2 used in the drawings of K_9 and K_{10}.

Recall that a string is a homeomorph of a compact real interval.

Question 3. Can we characterize those sets of strings that are extendible to an arrangement of pseudocircles?

So far, the authors have not found an example of a set of strings not-extendible to a set of pseudocircles that cannot be explained in terms of spirals and Theorem 2.3. A further study of spirals may be the key to solve this question.

We conclude with two simple results about pseudospherical drawings of K_n. For the first, Lemma 5.2 implies that every edge of a pseudospherical drawing of K_n induces a split of K_n into two pseudolinear drawings of smaller complete graphs (one having k, say, vertices and the other $n+2-k$ vertices). Every pseudolinear drawing of K_n has at least $n^2 + O(n \log(n))$ empty triangles (that is, 3-cycles having a side that does not contain a vertex) [10]. Thus, we can estimate the number of empty triangles on one side or other of the split. Adding in empty triangles involving vertices on different sides of the split yields at least $\frac{3}{4}n^2 + O(n \log(n))$ empty triangles in a pseudospherical drawing of K_n.

Rafla [24] conjectured that every (good) drawing of K_n has a Hamilton cycle with no self-crossing. Ábrego et al. have enumerated all the drawings of K_n with $n \leq 9$ [2] and in this way verified the conjecture for all these drawings of K_n. We extend to pseudospherical drawings the folklore proof that a pseudolinear drawing of K_n has such a Hamilton cycle.

For each edge e of the pseudospherical drawing D of K_n, let $c(e)$ denote the number of pseudocircles that $D[e]$ crosses. Choose a Hamilton cycle H in K_n that minimizes $\sum_{e \in E(H)} c(e)$. If $e_1, f_1 \in E(H)$ cross in D, then let J be the K_4 induced by e_1 and f_1. Then $D[J]$ has exactly one crossing, namely e_1 with f_1. The remaining four edges in J come in two disjoint, non-crossing pairs. For one such pair $\{e_2, f_2\}$, $(H - \{e_1, f_1\}) + \{e_2, f_2\}$ is also a Hamilton cycle H'.

We claim that $\sum_{e \in E(H')} c(e) < \sum_{e \in E(H)} c(e)$, contradicting the choice of H.

1. Clearly, the two contributions of e_1 crossing γ_{f_1} and f_1 crossing γ_{e_1} are counted for the H-sum, but not for the H'-sum.
   Now suppose $e \in E(H')$ crosses the pseudocircle γ_f containing $f \in E(H')$.

2. If e is not $\{e_2, f_2\}$, then the $e\gamma_f$-crossing is counted for both H and H'.

3. If $e = e_2$, say, then γ_f crosses e_2 and must continue through another side of the 4-cycle $D[J - \{e_1, f_1\}]$. In order to get there, it must cross at least one of e_1 and f_1. If γ_f crosses both e_2, f_2, then it also crosses both e_1, f_1. Therefore, γ_f has at least as many crossing with e_1, f_1 that it has with e_2, f_2, 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, T. Hackl, J. Pammer, A. Pilz, P. Ramos, G. Salazar, and B. Vogtenhuber. All good drawings of small complete graphs. In Proc. 31^{st} European Workshop on Computational Geometry EuroCG '15, pages 57-60, Ljubljana, Slovenia, 2015.

[3] B.M. Ábrego, O. Aichholzer, S. Fernández-Merchant, D. McQuillan, B. Mohar, P. Mutzel, P. Ramos, R.B. Richter, B. Vogtenhuber, Bishellable drawings of K_n. SIAM J. Discrete Math. **32** (2018), no. 4, 2482–2492.

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

[5] B.M. Ábrego, O. Aichholzer, S. Fernández-Merchant, P. Ramos, and B. Vogtenhuber, Non-shellable drawings of K_n with few crossings, in: Proc. 26th Annual Canadian Conference on Computational Geometry CCCG 2014.

[6] B.M. Ábrego, O. Aichholzer, S. Fernández-Merchant, P. Ramos, and G. Salazar, The 2-page crossing number of K_n, Disc. Comput. Geom. **49** (2013), no. 4, 747–777.

[7] O. Aichholzer, J. García, D. Orden, and P. Ramos, New lower bounds for the number of $(\le k)$-edges and the rectilinear crossing number of K_n, Discrete Comput. Geom. **38** (2007), no. 1, 1–14.

[8] O. Aichholzer, T. Hackl, A. Pilz, G. Salazar, and B. Vogtenhuber, Deciding monotonicity of good drawings of the complete graph, In: Proc. XVI Spanish Meeting on Computational Geometry (EGC 2015), 33–36, 2015.

[9] A. Arroyo, J. Bensmail, and R.B. Richter, Extending drawings of graphs to arrangements of pseudolines, submitted.

[10] A. Arroyo, D. McQuillan, R.B. Richter, and G. Salazar, Levi's Lemma, pseudolinear drawings of K_n, and empty triangles, J. Graph Theory **87** (2018), no. 4, 443–459.

[11] A. Arroyo, D. McQuillan, R.B. Richter, and G. Salazar, Drawings of K_n with the same rotation scheme are the same up to triangle-flips (Gioan's theorem), Australas. J. Combin. **67** (2017), 131–144.

[12] A. Arroyo, D. McQuillan, R.B. Richter, and G. Salazar, Convex drawings of K_n: topology meets geometry, submitted.

[13] A. Björner, M. Las Vergnas, B. Sturmfelds, N. White, G. Ziegler, Oriented Matroids, Encyclopedia of Mathematics and its Applications Book 46, Cambridge University Press, 2009.

[14] J. Blažek and M. Koman, A minimal problem concerning complete plane graphs, in M. Fiedler (ed), Theory of Graphs and Its Applications, Czech. Acad. of Sci. (1964), 113–117.

[15] S. Felsner and M. Scheucher. Arrangements of pseudocircles: on circularizability, Extended abstract in: Proceedings GD 18 LNCS **11282** (2018) 555–568. To appear Disc. Comp. Geom., DOI: 10.1007/s00454-019-00077-y.

[16] S. Felsner and M. Scheucher, Homepage of Pseudocircles, `https://www3.math.tu-berlin.de/diskremath/pseudocircles/`

[17] E. Gioan, Complete graph drawings up to triangle mutations, submitted. (A preliminary version appears as: Complete graph drawings up to triangle mutations, Proceedings WG 2005 (Metz), LNCS **3787**, pp139–150 (2005) (International Workshop on Graph-Theoretic Concepts in Computer Science).)

[18] F. Harary and A. Hill, On the number of crossings in a complete graph, Proc. Edinburgh Math. Soc. **2** (1962/1963), no. 13, 333–338.

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

[20] J. Kynčl, `https://mathoverflow.net/questions/128878/drawings-of-complete-graphs-with-zn-crossings`

[21] 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.

[22] C. Medina, J. Ramírez-Alfonsín, and G. Salazar. The unavoidable arrangements of pseudocircles, Proc. Amer. Math. Soc. **147** (2019), no. 7, 3165–3175.

[23] J.W. Moon, On the distribution of crossings in random complete graphs, SIAM J. Appl. Math. **13** (1965), 506–510. Erratum, same journal **32** (1977), no. 3, 706.

[24] N. H. Rafla, The good drawings D_n of the complete graph K_n, Ph.D. thesis, McGill University, Montreal, 1988.

[25] J. Snoeyink and J. Hershberger, Sweeping arrangements of curves, Disc. Comp. Geom. (New Brunswick, NJ, 1989/1990), 309–349, DIMACS Ser. Discrete Math. Theoret. Comput. Sci., **6** Amer. Math. Soc. Providence, RI, 1991.