Drawings of K_n with the same rotation scheme are the same up to Reidemeister moves (Gioan's Theorem)

In memory of our friend Dan Archdeacon.

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

LaTeX-ed: December 24, 2015

Abstract

A *good drawing* of K_n is a drawing of the complete graph with n vertices in the sphere such that: no two edges with a common end cross; no two edges cross more than once; and no three edges all cross at the same point. Gioan's Theorem asserts that any two good drawings of K_n that have the same rotations of incident edges at every vertex are equivalent up to Reidemeister moves. At the time of preparation, 10 years had passed between the statement in the WG 2005 conference proceedings and our interest in the proposition. Shortly after we completed our preprint, Gioan independently completed a preprint.

1 Introduction

sec:intro

The main result of this work is the proof of the following result, presented by Gioan at the International Workshop on Graph-Theoretic Concepts in Computer Science 2005 (WG 2005) [7].

th:gioan

Theorem 1.1 (Gioan's Theorem) *Let D_1 and D_2 be good drawings (defined below) of K_n in the sphere that have the same rotation schemes. Then there is a sequence of Reidemeister moves (example below, defined in Section 2) that transforms D_1 into D_2.*

We are only using "Reidemeister III" moves to shift a bit of the interior of an edge across another crossing (without crossing anything else). Figure 1.2 shows a typical example of "before" and "after" the move.

*Supported by CONACYT. (GS by grant no. 222667)
†Supported by NSERC grant no. 41705-2014 057082.
+University of Waterloo, ‡Norwich University, and ×UASLP

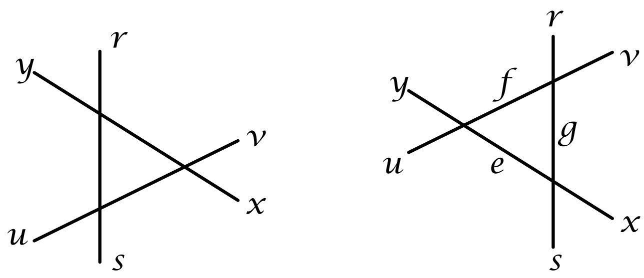

Figure 1.2: A Reidemeister III move that transforms one drawing into another. fg:vxInR

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

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

Throughout this work, all drawings of graphs are *good drawings*:

- no two edges incident with a common vertex cross;
- no three edges cross at a common point; and
- no two edges cross each other more than once.

Some of our interest in this problem derives from Dan Archdeacon's combinatorial generalization of this problem. Since his website may soon be lost and there is no other version that we know of, we reproduce it here.

> *Suppose the vertex set of K_n is $I_n = \{1, ..., n\}$. A local neighborhood of a vertex k in a planar drawing determines a cyclic permutation of the edges incident with k by considering the clockwise ordering in which they occur. Equivalently (looking at the edges' opposite endpoints), it determines a local rotation $\rho(k)$: a cyclic permutation of $I_n - k$. A (global) rotation is a collection of local rotations $\rho(k)$, one for each vertex k in I_n.*
>
> *It is well known that the rotations of K_n are in a bijective correspondence with the embeddings of K_n on oriented surfaces. The rotation arising from a planar drawing also determines which edges cross. Namely, edges ab, cd cross in the drawing if and only if the induced local rotations on the vertices $\{a, b, c, d\}$ give a nonplanar embedding of that induced K_4.* [This is not quite true: the rotation determines the crossing among the six edges in the K_4 induced by a, b, c, d, but it is not necessarily true that it is ab with cd. AMRS]

The stated conjecture on the crossing number of K_n asserts that the minimum number (over all planar drawings) of induced nonplanar K_4's satisfies the given lower bound. We generalize this to all rotations.

Conjecture: *In any rotation of K_n, the number of induced nonplanar K_4's is at least $(1/4)[n/2][(n-1)/2][(n-2)/2][(n-3)/2]$ where $[m]$ is the integer part of m.*

Not every rotation corresponds to a drawing (see the related problem "Drawing rotations in the plane"), so this conjecture is strictly stronger than the one on the crossing number of K_n. However, this conjecture has the advantage of reducing a geometric problem to a purely combinatorial one.

The problem arose from my attempts to prove the lower bound on the crossing number. It is supported by computer calculations. Namely, I wrote a program which started with a rotation of K_n and using a local optimization technique (hill-climbing), randomly swapped edges in a local rotation whenever that swap did not increase the number of induced nonplanar K_4's. The resulting locally minimal rotations tended to resemble the patterns apparent in an optimal drawing of K_n. For small n this minimum was the conjectured upper bound. For larger n it was usually slightly larger.

It is well-known that the rectilinear crossing number (all edges are required to be straight-line segments) of K_n is, for $n \geq 10$, strictly larger than $H(n)$ [4]. In fact, this applies to the more general *pseudolinear* crossing number [2].

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

A drawing of K_n is *pseudolinear* if there is an arrangement Σ of $\binom{n}{2}$ pseudolines such that the edges of K_n are all contained in different pseudolines of Σ. It is clear that a rectilinear drawing (chosen so no two lines are parallel) is pseudolinear.

The arguments (originally due to Lovász et al [11] and, independently, Ábrego and Fernández-Merchant [1]) that show every rectilinear drawing of K_n has at least $H(n)$ crossings apply equally well to pseudolinear drawings.

The proof that every optimal pseudolinear drawing of K_n has its outer face bounded by a triangle [6] uses the "allowable sequence" characterization of pseudoline arrangements of Goodman and Pollack [8]. Our principal result in [5] is that there is another, topological, characterization of pseudolinear drawings of K_n.

A drawing D of K_n is *face-convex* if there is an open face F of D such that, for every 3-cycle T of K_n, if Δ is the closed face of $D[T]$ disjoint from F, then, for any two vertices u, v such that $D[u], D[v]$ are both in Δ, the arc $D[uv]$ is also contained in Δ.

The main result in [5] is that every face-convex drawing of K_n is pseudolinear and conversely. An independent proof has been found by Aichholzer et al [3]; their proof uses Knuth's CC systems [9], which are an axiomatization of sets of pseudolines. Moreover, their statement is in terms of a forbidden configuration. Properly speaking, their result is of the form, "there exists a face relative to which the forbidden configuration does not occur". Their face and our face are the same. However, our proof is completely different, yielding directly a polynomial time algorithm for finding the pseudolines.

Aichholzer et al show that there is a pseudolinear drawing of K_n having the same crossing pairs of edges as the given drawing of K_n. Gioan's Theorem [7] (Theorem 1.1 above) is then invoked to show that the original drawing is also pseudolinear.

The proof in [5] is completely self-contained; in particular, it does not invoke Gioan's Theorem. An earlier version anticipated an application of Gioan's Theorem similar to that in [3]; hence our interest in having a proof.

A principal ingredient in our argument is a consideration of the facial structure of an arrangement of arcs in the plane. An *arrangement of arcs* is a finite set Σ of open arcs in the plane such that, for every $\sigma \in \Sigma$, $\mathbb{R}^2 \setminus \sigma$ is not connected and any two elements of Σ have at most one point in common, which must be a crossing.

Let Σ be an arrangement of arcs. Since Σ is finite, there are only finitely many faces of Σ: these are the components of $\mathbb{R}^2 \setminus (\bigcup_{\sigma \in \Sigma} \sigma)$. As it comes up often, we let $\mathcal{P}(\Sigma)$ be the pointset $\bigcup_{\sigma \in \Sigma} \sigma$.

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

Although we do not need it here, the following simple lemma motivates one that we do use in our proof of Gioan's Theorem. Its simple proof from [5] is included here for completeness.

lm:dualPaths **Lemma 1.3 (Existence of dual paths)** *Let Σ be an arrangement of arcs in the plane and let a, b be points of the plane not in any line in Σ. Then there is an ab-path in Σ^* crossing each arc in Σ at most once.*

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

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

2 Proof of Gioan's Theorem

sec:gioan

In this section, we give a simple, self-contained proof Gioan's Theorem [7]. When we completed the proof in August 2015, we corresponded with Gioan, who was independently preparing his own version. Each version has had some impact on the other. We do not include any of the first order logical considerations that occur in Gioan's version.

For convenience, we restate our main result here. The definition of a Reidemeister move is given just after this statement.

Theorem 1.1 *Let D_1 and D_2 be drawings of K_n in the sphere that have the same rotation schemes. Then there is a sequence of Reidemeister moves that transforms D_1 into D_2.*

In order to define Redemeister move and prove our first intermediate lemmas, we require a small new consideration. Let Σ be an arrangement of arcs in the plane. A *vertex of* Σ is a point that is the intersection of two or more arcs in Σ.

At a vertex v, the rotation of the arcs containing v is of the form $\sigma_1, \sigma_2, \dots, \sigma_k, \sigma_1, \sigma_2, \dots, \sigma_k$; each arc occurs twice here, once for each of the "rays" it contains that start at v. Let $(F_0, F_1, \dots, F_{k-1}, F_k, F_{k+1}, \dots, F_{2k-1})$ the cyclic sequence of faces around v.

Suppose P is a dual path containing the subpath $(F_0, F_1, \dots, F_k)$ such that P crosses each arc in Σ at most once. The path obtained from P by *sliding over the vertex v* is the path P, except $(F_0, F_1, \dots, F_k)$ is replaced by the dual path (of the same length) $(F_0, F_{2k-1}, F_{2k-2}, \dots, F_{k+1}, F_k)$. (None of $F_{2k-1}, F_{2k-2}, \dots, F_{k+1}$ can occur in P, as P crosses each arc of Σ at most once. Thus, the result of the sliding is indeed a new dual path.)

A *Reidemeister move* is a sliding over a vertex v that is in precisely two arcs in Σ. The following may be viewed as a supplement to Lemma 1.3.

lm:reidemeister

Lemma 2.1 *Let Σ be an arrangement of arcs in the plane and let a and b be any two points in the plane not in $\mathcal{P}(\Sigma)$. Let F_a and F_b be the faces of Σ containing a and b, respectively. Then any two F_aF_b-paths in Σ^*, each crossing every arc in Σ at most once, are equivalent up to sliding over vertices.*

Proof. Let P and Q be distinct F_aF_b-paths in Σ^*. Let P_1 and Q_1 be subpaths of P and Q having common end points but being otherwise disjoint. Then (any natural image in the plane of) $P_1 \cup Q_1$ bounds a disc Δ and each curve in Σ that crosses one of P_1 and Q_1 crosses the other. We will show that there is a vertex in Δ over which we can slide P_1.

Since P_1 and Q_1 are distinct dual paths, there is a vertex of Σ in Δ. Let $\sigma \in \Sigma$ have an arc across Δ and contain a vertex of Σ; let v be the first vertex of Σ encountered as we traverse σ across Δ from its P_1-end. Among all the $\sigma \in \Sigma$ that contain v, either all have v as their first encountered vertex or there are two, σ and $\bar{\sigma}$, consecutive in the rotation at v, such that v is the first encountered vertex for σ, but not for $\bar{\sigma}$. In the former case, we can slide P_1 across v.

Suppose $\sigma' \in \Sigma$ has a crossing with $\bar{\sigma}$ between the intersection of $\bar{\sigma}$ with P_1 and v. Let Δ' be the disc bounded by P_1, σ, and $\bar{\sigma}$. Then $\sigma' \cap \Delta'$ intersects the boundary of Δ' at least twice, but not on $\sigma \cap \Delta'$. Thus, σ' crosses P_1 between $\sigma \cap P_1$ and $\bar{\sigma} \cap P_1$.

Let $\bar{v}$ be the first vertex of Σ encountered as we traverse $\bar{\sigma}$ from $\bar{\sigma} \cap P_1$. Then every other arc in Σ that contains $\bar{v}$ intersects P_1 between $\sigma \cap P_1$ and $\bar{\sigma} \cap P_1$.

Letting $b(v)$ denote the number of arcs in Σ that cross P_1 between $\sigma \cap P_1$ and $\bar{\sigma} \cap P_1$, we see that $b(\bar{v}) < b(v)$. Therefore, there is always a vertex w of Σ such that $b(w) = 0$ and we can slide P_1 across w.

After sliding P_1 across w, we get a new P that either has more vertices in common with Q or the disc bounded by the new P_1 and Q_1 has fewer vertices of Σ. In either case, an easy induction completes the proof. ■

Gioan's Theorem considers two drawings D_1 and D_2 of K_n in the sphere that have the same rotation scheme. Let t, u, v, w be four distinct vertices of K_n. Let T be the triangle induced by t, u, v. Then $D_1[T]$ is a simple closed curve in the sphere. The rotations at t, u, and v determine where bits of the edges $D_1[tw]$, $D_1[uw]$, and $D_1[vw]$ go from their ends t, u, and v, respectively, relative to $D_1[T]$. The side of $D_1[T]$ that has the majority (two or three) of these bits of edges is where $D_1[w]$ is. If tw is the minority edge, then $D_1[tw]$ crosses $D_1[uv]$; conversely, a crossing K_4 produces, for each of its triangles, a minority edge. This simple observation immediately yields the following fundamental fact.

it:rotationK4 (F1) Let D_1 and D_2 be two drawings of K_n with the same rotation scheme. If J is any K_4 in K_n, then there is an orientation-preserving homeomorphism of the sphere to itself mapping $D_1[J]$ onto $D_2[J]$ that preserves the vertex-labels of J.

There are some elementary corollaries of (F1):

otationCrossing (F2) the pairs of crossing edges are determined by the rotation scheme;

tionDirCrossing (F3) if the edges of K_n are oriented, then the directed crossings are determined by the rotation scheme; and ngleContainment (F4) if u, v, w, x are distinct vertices of K_n, then the side of the triangle (relative to any of its oriented sides) induced by u, v, w that contains x is determined by the rotation scheme.

By (F3), we mean that, if e and f cross, then, as we follow the orientation of e, the crossing of e by the traversal f is either left-to-right in all drawings or right-to-left in all drawings, depending only on the rotation scheme.

These facts can hardly be new. In fact, variations of some of them appear in Kynčl [10].

gOrderDetermines **Lemma 2.2** *Let D_1 and D_2 be two drawings of K_n with the same rotation scheme. Suppose that, for each edge e, as we traverse e from one end to the other, the edges that cross e occur in the same order in both D_1 and D_2. Then there is an orientation-preserving homeomorphism of the sphere mapping $D_1[K_n]$ onto $D_2[K_n]$ that preserves all vertex- and edge-labels.*

Proof. This is a consequence of the well known fact that a rotation scheme of a connected graph determines a unique (up to surface orientation-preserving homeomorphisms) cellular embedding of a graph in an orientable surface ([12, Thm. 3.2.4]). We construct a planar map from each of D_1 and D_2 by inserting a vertex of degree 4 at each crossing point. The oriented crossings and the orders of the crossings of each edge are the same in both D_1 and D_2, so the rotations at these degree 4 vertices are also the same. Therefore, the planar maps are the same, as claimed. ■

Lemma 2.2 asserts that the orders of crossings determine the drawing. Thus, we need to consider the situation that some edge has two edges crossing it in different orders in the two drawings. The first step, our next lemma, is to identify a special structure that must occur.

Let e, f, and g be three distinct edges in a drawing D of K_n, no two having a common end. Suppose each two of e, f, and g have a crossing, labelled $\times_{e,f}$, $\times_{e,g}$, and $\times_{f,g}$. The union of the segments of each of e, f, and g between their two crossings is a simple closed curve. If one of the two sides of this simple closed curve does not have an end of any of e, f, and g, then this side is the *pre-Reidemeister triangle constituted by e, f, and g*.

Let D_1 and D_2 be drawings of K_n with the same rotation scheme. A *Reidemeister triangle* for D_1 and D_2 is a pre-Reidemeister triangle T for both D_1 and D_2 constituted by the edges e, f, and g but with the clockwise traversal of the three segments between pairs of crossings giving the opposite cyclic ordering of the crossings.

Let J be a crossing K_4 in D_1. Then $D_2[J]$ is also a crossing K_4, with the same pair of edges crossing: there is a label-preserving homeomorphism of the sphere to itself that maps $D_1[J]$ onto $D_2[J]$. Letting $\times$ denote the crossing (in both $D_1[J]$ and $D_2[J]$), $D_1[J]$ has five faces: one *4-face* bounded by a 4-cycle in J; and four *3-faces*, each incident with $\times$. If x and r are two vertices of J incident with a 3-face, then we use $T^1_{x,r}$ to denote this 3-face and $xr\times$ to denote its boundary.

Our next lemma corresponds to Lemma 3.2 of [7]. This result is a central, non-trivial point in the argument.

OrderReidTriang **Lemma 2.3** *Let D_1 and D_2 be two drawings of K_n with the same rotation scheme. Then no Reidemeister triangle in $D_1[K_n]$ for D_1 and D_2 contains a vertex of $D_1[K_n]$.*

Proof. Let R be a Reidemeister triangle in $D_1[K_n]$ for D_1 and D_2. We use the same labelling $e = xy$, $f = uv$, and $g = rs$ as above for the edges determining R; all of r, s, u, v, x, and y are in the same face F of $D_1[R]$. By way of contradiction, suppose there is a vertex a of K_n in the other face F_a of $D_1[R]$. See the left-hand figure in Figure 2.4.

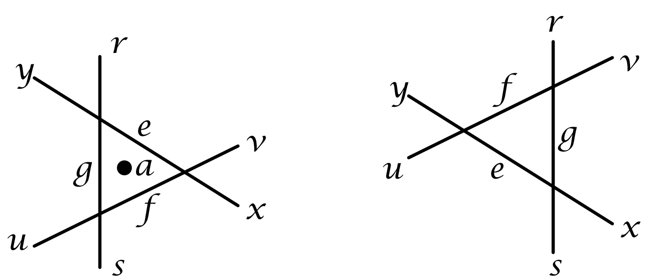

Figure 2.4: The Reidemeister triangle in D_1 and D_2. fg:vxInR

In the K_4 induced by u, v, x, y, a is in the 3-face $T^1_{u,y}$ bounded by $uy\times^1_{e,f}$ and, therefore, in the faces bounded by the 3-cycles uyx and yuv that do not contain $D_1[v]$ and $D_1[x]$, respectively. By (F1), this holds true also for D_2. Analogous statements hold for the other two K_4's involving two of the three edges from e, f, g.

It follows that a is in all of the faces in D_2 bounded by $uy\times^2_{e,f}$, $rv\times^2_{f,g}$, and $xs\times^2_{e,g}$ in D_2 that are disjoint from $D_2[r]$, $D_2[x]$, and $D_2[u]$, respectively. Label these faces as $T^2_{u,y}$, $T^2_{r,v}$, and $T^2_{x,s}$, respectively. Since a is in all three, $T^2_{u,y} \cap T^2_{r,v} \cap T^2_{x,s} \neq \varnothing$.

On the other hand, the edge uy does not cross either uv or xy. Since it joins the two points $D_2[u]$ and $D_2[y]$ and they are on the same side of $D_2[R]$, $D_2[uy]$ crosses $D_2[R]$ an even number of times; we conclude that $D_2[uy]$ does not cross $D_2[R]$. The same applies to $D_2[rv]$ and $D_2[xs]$.

Since the open segment τ of xy in the boundary of $D_2[R]$ is not crossed by the boundary $uy\times^2_{e,f}$ of $T^2_{u,y}$, τ is disjoint from $T^2_{u,y}$. In particular, $\times^2_{f,g}$ is not in $T^2_{u,y}$; it follows that points in the interior of $T^2_{r,v}$ near $\times_{f,g}$ are not in $T^2_{u,y}$. Consequently, $T^2_{r,v}$ is not contained in $T^2_{u,y}$. The symmetry of the situation shows that none of $T^2_{u,y}$, $T^2_{r,v}$, and $T^2_{x,s}$ is contained in the other.

On the other hand, $a \in T^2_{u,y} \cap T^2_{r,v} \cap T^2_{x,s}$, so that each two of $uy\times^2_{e,f}$, $rv\times^2_{f,g}$, and $xs\times^2_{e,g}$ intersect; since they intersect each other an even number of times, they intersect each other at least twice.

Therefore, the 6-cycle $rvuyxs$ has at least nine crossings in D_2, consisting of the three that define R and the at least six mentioned at the end of the preceding paragraph. Since nine is the most crossings a 6-cycle can have in a good drawing, we conclude that it is exactly nine. Thus, any two of $uy\times^2_{e,f}$, $rv\times^2_{f,g}$, and $xs\times^2_{e,g}$ cross exactly twice. Moreover, every pair of non-adjacent edges in the 6-cycle must cross. In particular, rv crosses uy.

By goodness, the uv-segment in $rv\times^2_{f,g}$ does not intersect either uy- or the uv-segment of $uy\times^2_{e,f}$. Since uv crosses xy at $\times^2_{e,f}$, it does not do so a second time. Therefore, the uv-segment of $rv\times^2_{f,g}$ does not cross $uy\times^2_{e,f}$.

We have already mentioned that rv crosses uy. If rs crosses $uy\times^2_{e,f}$, then, as rs already crosses xy and uv, this crossing is not on either of the xy- and uv-segments of $uy\times^2_{e,f}$. Also, rv does not cross the uv-segment of $uy\times^2_{e,f}$. Thus, there are two possibilities: either rv crosses both uy and the xy-segment of $uy\times^2_{e,f}$ or both rv and the rs-segment of $rv\times^2_{f,g}$ cross uy.

The conclusion is that either rv crosses $uy\times^2_{e,f}$ twice or that uy crosses $rv\times^2_{f,g}$ twice. Since these conclusions are symmetric, we may assume the former. The final piece of information that we require is the order in which these two crossings occur. By way of contradiction, suppose that, as we traverse $D_2[rv]$ from $D_2[v]$, we first cross the xy-segment of $uy\times^2_{e,f}$ before crossing uy. See Figure 2.5.

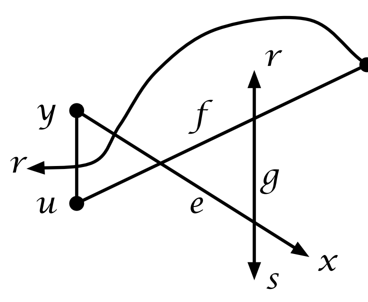

Figure 2.5: $D_2[rv]$ crosses $T_{u,y}$ in the wrong order. fg:wrong

Consider the simple closed curve Ω consisting of the arc in $D_2[uv]$ from $\times^2_{e,f}$ to $D_2[v]$, then along $D_2[rv]$ from $D_2[v]$ to the crossing of $D_2[rv]$ with the xy-segment of $uy\times^2_{e,f}$, and then along $D_2[xy]$ back to $\times^2_{e,f}$.

By goodness, the portion of $D_2[rs]$ from $\times^2_{f,g}$ to $D_2[r]$ cannot cross Ω, so $D_2[r]$ is on the side of Ω that is different from the side containing the crossing of rv with uy. Again, goodness forbids the crossing of Ω with the portion of $D_2[rv]$ from r to the crossing with uy. This contradiction shows that the first crossing of $uy\times^2_{e,f}$ by $D_2[rv]$, as we start at v, is with uy. See Figure 2.6.

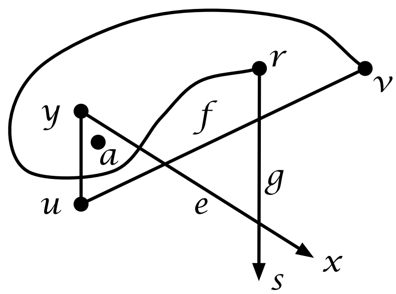

Figure 2.6: This is how $D_2[rv]$ crosses $T_{u,y}$.

fg:right

The vertex a is in $T_{r,v} \cap T_{u,y}$. As $D_1[a]$ and $D_1[y]$ are on different sides of $D_1[R]$, $D_1[ay]$ crosses at least one of $D_1[rs]$, $D_1[uv]$, and $D_1[xy]$. Thus, (F2) implies $D_2[ay] \not\subseteq T_{u,y}$.

Goodness implies that $D_2[ay]$ must cross the uv-segment of $uy\times^2_{e,f}$. In order to do that, it must cross rv first. But now y and the crossing $\times$ of $D_2[ay]$ with $D_2[uv]$ are separated by the simple closed curve Ω' consisting of the portion of uv from $\times^2_{e,f}$ to v, rv from v to its crossing with xy, and the portion of xy between this crossing and $\times^2_{e,f}$.

However, the portion of ay from $\times$ to y cannot cross any of the three parts of Ω', because each part is contained either in an edge incident with y or is crossed by the complementary part of ay. This contradiction completes the proof. ■

We are now ready to prove Gioan's Theorem. The structure of our proof is very much the same as that given by the algorithm in [7].

Proof of Theorem 1.1. Let v be any vertex of K_n. Induction (with $n \leq 4$ as the base) shows that there is a sequence Γ of such Reidemeister moves that converts the drawing $D_1[K_n - v]$ into $D_2[K_n - v]$.

We claim we can realize all the Reidemeister moves of Γ within the drawing $D_1[K_n]$, by interspersing some moves that only move edges incident with v. Suppose $\Gamma = \gamma_1\gamma_2\cdots\gamma_k$ and that, for some $i \geq 1$, we have been able to do all the moves γ_1, γ_2, $\ldots$, γ_{i-1}. At this point, we have a drawing $D_1'[K_n]$ that has the property that, doing the sequence $\gamma_i\gamma_{i+1}\cdots\gamma_k$ on $D_1'[K_n-v]$, we obtain $D_2[K_n-v]$. In particular, the Reidemeister triangle R_i used to perform γ_i is empty relative to $D_1'[K_n-v]$. Moreover, since D_1' and D_2 have the same rotation schemes, Lemma 2.3 implies that v is also not inside $D_1'[R_i]$.

It follows that the only crossings of $D_1'[R_i]$ are by edges incident with v. Such edges cross $D_1'[R_i]$ in exactly two of its three sides and, since no two of the edges incident with v cross, their segments inside $D_1'[R_i]$ are disjoint. It follows that they can be moved out of R_i by Reidemeister moves, creating a new drawing $D_1''[K_n]$ in which R_i is empty. Thus, performing the Reidemeister move γ_i on $D_1''[K_n]$ produces a new drawing of K_n in which the moves γ_1, γ_2, ..., γ_i have all been done.

It follows that we may assume $D_1[K_n-v]$ is the same as $D_2[K_n-v]$. We complete the proof by showing that we can perform Reidemeister moves on the edges incident with v to convert D_1 into D_2. For ease of notation and reference, we will use K_n-v to denote the common drawings $D_1[K_n-v]$ and $D_2[K_n-v]$. We may assume that, for $i=1,2$, $D_i[K_n]$ is obtained from K_n-v by using dual paths for each edge vw incident with v, together with a small segment in the last face to get from the dual vertex in that face to w.

This understanding needs a slight refinement, since, for example, it is possible for two edges incident with v to use the same sequence of faces (in whole or in part). Thus, as dual paths, they would actually use the same segments. We allow this, as long as the two edges do not cross on the common segments. They can be slightly separated at the end to reconstruct the actual drawing.

The triangles of K_n-v and the common rotations determine the face F of K_n-v containing v, so this is the same in both D_1 and D_2. It follows that, for each vertex w of K_n-v, $D_1[vw]\cup D_2[vw]$ is a closed curve C_w with finitely many common segments (which might be just single dual vertices). In particular, the closed curve C_w divides the sphere into finitely many regions.

ertexSeparation **Claim 1** *For each vertex w of K_n-v, all the vertices of $K_n-\{v,w\}$ are in the same region of C_w.*

Proof. Let x and y be vertices of $K_n-\{v,w\}$. If xy does not cross $D_1[vw]$, then it also does not cross $D_2[vw]$, so xy is disjoint from C_w, showing x and y are in the same region of C_w. Letting J be the K_4 induced by v,w,x,y, we may assume that, in each of $D_1[J]$ and $D_2[J]$, vw crosses xy.

For $i=1,2$, the path (x,w,y) in $D_i[J]$ is incident with the face F_i of $D_i[J]$ bounded by the 4-cycle (v,x,w,y,v). In particular, there is an xy-arc γ_i in F_i that goes very near alongside (x,w,y) and is disjoint from $D_i[vw]$. Furthermore, we may choose the arcs γ_1 and γ_2 to be equal. The arc γ_1 shows that x and y are in the same region of C_w. □

For $w \in V(K_n - v)$, a *w-digon* is a simple closed curve in C_w consisting of a subarc of $D_1[vw]$ and a subarc of $D_2[vw]$. If $D_1[vw] \neq D_2[vw]$, then it is easy to see that there is at least one w-digon.

For each w-digon C, Claim 1 shows that one side of C in the sphere has no vertex of $K_n - \{v, w\}$. This closed disc is the *clean side* of C.

Label the vertices of $K_n - v$ as $w_1, w_2, \dots, w_{n-1}$. Suppose $i \in \{1, 2, \dots, n-1\}$ is such that $D_1[vw_1], \dots, D_1[vw_{i-1}]$ are all the same as $D_2[vw_1], \dots, D_2[vw_{i-1}]$, respectively. We show that there is a drawing D_1', obtained from D_1 by Reidemeister moves that move only edges from among $vw_i, vw_{i+1}, \dots, vw_{n-1}$, so that $D_1'[vw_1], \dots, D_1'[vw_i]$ are all the same as $D_2[vw_1], \dots, D_2[vw_i]$, respectively. This will complete the proof.

If $D_1[vw_i] = D_2[vw_i]$, then we set $D_1' = D_1$, and we are done. In the remaining case, there are w_i-digons. We show that we can find a sequence of Reidemeister moves to create a new drawing D_1' such that $D_1'[vw_i]$ has more agreement with $D_2[vw_i]$ than $D_1[vw_i]$ has. Furthermore, the only edges moved are among $vw_i, vw_{i+1}, \dots, vw_{n-1}$. Clearly, this is enough.

Begin by selecting, among all w_i-digons, a w_i-digon C with minimal clean side S; no other w_i-digon has its clean side contained in S. If xy is an edge of $K_n - \{v, w_i\}$ that intersects S, then $xy \cap (S \setminus C)$ consists of a single arc that has one end in $D_1[vw_i]$ and one end in $D_2[vw_i]$. We will not do anything with these arcs, except in various applications of Lemma 2.1 in which only edges from among $D_1[vw_i], D_1[vw_{i+1}], \dots, D_1[vw_{n-1}]$ are adjusted.

`cl:1-(i-1)` **Claim 2** *For $j \in \{1, 2, \dots, i-1\}$, the edge $D_1[vw_j]$ is disjoint from S.*

Proof. If, for some $j \in \{1, 2, \dots, i-1\}$, $D_1[vw_j]$ has a point in S, then, since w_j is not in S, $D_1[vw_j]$ crosses C. But $D_1[vw_j] = D_2[vw_j]$, showing it is disjoint from $(D_1[vw_i] \cup D_2[vw_i]) \setminus \{D_2[v]\}$. □

`cl:(i+1)-(k-1)` **Claim 3** *For each $j \in \{i+1, i+2, \dots, k-1\}$, $D_1[vw_j] \cap (S \setminus C)$ consists of disjoint open arcs, each having both ends in $(C \cap D_2[vw_i]) \setminus D_2[w_i]$.*

Proof. Let p be a point of $D_1[vw_j] \cap (S \setminus C)$. Since w_j is not inside S, as we follow $D_1[vw_j]$ from p towards $D_1[w_j]$, there is a first point q in C. Since $D_1[vw_j]$ is disjoint from $D_1[vw_i]$, q must be in $D_2[vw_i] \setminus D_2[\{v, w_i\}]$. Likewise, in moving from p toward $D_1[v]$, there is a first point reached that is in $D_2[vw_i] \setminus D_2[w_i]$. □

Notice that no two of the arcs described in Claim 3 can intersect. Therefore, among all the arcs in $D_1[vw_j] \cap (S \setminus C)$, there is an arc α which, together with a subarc α' of $D_2[vw_i] \cap C$, makes a minimal digon. We can then use Lemma 2.1 to move α to agree with α'. We repeat this procedure until there are no such arcs left in S, at which time Lemma 2.1 shows we can move $D_1[vw_i]$ onto $D_2[vw_i]$. (Here is the principal place where we allow several edges to share a dual path in D_1. As we shift α, it becomes equal to a subarc of $D_2[vw_i]$. These are both incident with v. Since the two ends of α both cross out of S, these two arcs do not cross.)

The only edges moved are among $D_1[vw_i], \dots, D_1[vw_{n-1}]$, as claimed. ■

References

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

dolinearVsUsual [2] B.M. Ábrego, J. Balogh, S. Fernández-Merchant, J. Leaños, and G. Salazar, An extended lower bound on the number of ($\leq k$)-edges to generalized configurations of points and the pseudolinear crossing number of K_n, J. Combin. Theory Ser. A **115** (2008), no. 7, 1257–1264.

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

tilinearVsUsual [4] O. Aichholzer, J. García, D. Orden, and P. Ramos, New lower bounds for the number of ($\leq k$)-edges and the rectilinear crossing number of K_n, Discrete Comput. Geom. **38** (2007), no. 1, 1–14.

amrs [5] A. Arroyo, D. McQuillan, R.B. Richter, and G. Salazar, Levi's Lemma, pseudolinear drawings of K_n, and empty triangles, submitted.

optTriang [6] J. Balogh, J. Leaños, S. Pan, R.B. Richter, and G. Salazar, The convex hull of every optimal pseudolinear drawing of K_n is a triangle, Australas. J. Combin. **38** (2007), 155–162.

gioan [7] E. Gioan, Complete graph drawings up to triangle mutations, submitted December 2015. (Extended abstract in D. Kratsch (Ed.): WG 2005, LNCS 3787, pp. 139–150, 2005, Springer-Verlag, Berlin-Heidelberg.)

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

knuth [9] D. Knuth, Axioms and Hulls, Vol. 606, Lecture Notes Comput. Sci., Springer, Berlin, 1992.

kyncl [10] J. Kynčl, Simple realizability of complete abstract topological graphs in P, Disc. Comput. Geom. **45** (2011), 383–399.

lovasz [11] 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.

mt [12] B. Mohar and C. Thomassen, Topological Graph Theory, Johns Hopkins University Press, Baltimore, MD, 2001.