

Compatible Spanning Trees in Simple Drawings of K_n ^{*}

Oswin Aichholzer¹, Kristin Knorr², Wolfgang Mulzer²,
Nicolas El Maalouly³, Johannes Obenaus², Rosna Paul¹,
Meghana M. Reddy³, Birgit Vogtenhuber¹, and Alexandra Weinberger¹

¹ Institute of Software Technology, Graz University of Technology, Austria

{oach,ropaul,bvogt,weinberger}@ist.tugraz.at

² Institut für Informatik, Freie Universität Berlin, Germany

{kristin.knorr,wolfgang.mulzer,johannes.obenaus}@fu-berlin.de

³ Department of Computer Science, ETH Zürich, Switzerland

{nicolas.elmaalouly,meghana.mreddy}@inf.ethz.ch

Abstract. For a simple drawing D of the complete graph K_n , two (plane) subdrawings are *compatible* if their union is plane. Let \mathcal{T}_D be the set of all plane spanning trees on D and $\mathcal{F}(\mathcal{T}_D)$ be the *compatibility graph* that has a vertex for each element in \mathcal{T}_D and two vertices are adjacent if and only if the corresponding trees are compatible. We show, on the one hand, that $\mathcal{F}(\mathcal{T}_D)$ is connected if D is a cylindrical, monotone, or strongly c -monotone drawing. On the other hand, we show that the subgraph of $\mathcal{F}(\mathcal{T}_D)$ induced by stars, double stars, and twin stars is also connected. In all cases the diameter of the corresponding compatibility graph is at most linear in n .

Keywords: Compatibility graph · Plane spanning tree · Simple drawing

1 Introduction

A *drawing* D of a graph G is a representation of G in the Euclidean plane such that the vertices of G are distinct points and the edges are Jordan arcs

^{*} This work was initiated at the 6th DACH Workshop on Arrangements and Drawings in Stels, August 2021. We thank all participants, especially Nicolas Grelier and Daniel Perz, for fruitful discussions. O.A., R.P. and A.W. are supported by FWF grant W1230. K.K. is supported by the German Science Foundation (DFG) within the research training group ‘Facets of Complexity’ (GRK 2434). W.M. is partially supported by the German Research Foundation within the collaborative DACH project *Arrangements and Drawings* as DFG Project MU 3501/3-1, and by ERC StG 757609. J.O. is supported by ERC StG 757609. M.M.R. is supported by the Swiss National Science Foundation within the collaborative DACH project *Arrangements and Drawings* as SNSF Project 200021E-171681. (Also note that this author’s full last name consists of two words and is *Mallik Reddy*. However, she consistently refers to herself with the first word of her last name being abbreviated.) B.V. was partially supported by the Austrian Science Fund (FWF) within the collaborative DACH project *Arrangements and Drawings* as FWF project I 3340-N35.

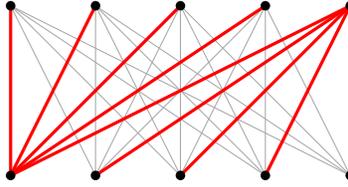


Fig. 1: A simple drawing of the complete bipartite graph with a tree (drawn in red, bold edges) that is an isolated vertex in the corresponding compatibility graph.

connecting their incident vertices such that no edge passes through any other vertex. A drawing is *simple* if any pair of edges intersect at most once - either in a common vertex or a proper *crossing* in the relative interior of the edges. All drawings considered in this paper are simple and the term simple is mostly omitted. A drawing is *plane* if it does not contain any crossing.

For a fixed integer n let D be a simple drawing of the complete graph K_n and let \mathcal{T}_D be the set of all drawings of plane spanning trees which are subdrawings of D . Note that \mathcal{T}_D is non-empty, as it contains at least the n stars in D (where a *star* contains all edges incident to a single vertex). Unless explicitly stated otherwise, the word *tree* always refers to a plane spanning tree in \mathcal{T}_D , where the drawing D is either clear from the context or the statement holds for any simple drawing of K_n . Two (plane) subdrawings H and H' of a simple drawing D are said to be *compatible* if the union of H and H' is still plane.

Let $\mathcal{F}(\mathcal{T}_D)$ be the (abstract) graph that has a vertex for each plane spanning tree in \mathcal{T}_D and two vertices are adjacent if and only if the corresponding trees are compatible. We call $\mathcal{F}(\mathcal{T}_D)$ the *compatibility graph* of \mathcal{T}_D . In this paper, we study properties of $\mathcal{F}(\mathcal{T}_D)$, focusing primarily on connectivity aspects:

Question 1. *Let n be an integer. Is the compatibility graph $\mathcal{F}(\mathcal{T}_D)$ connected for any simple drawing D of the complete graph K_n ?*

Note that the notion of compatibility is closely related to the notion of edge flips: An *edge flip* in a plane spanning tree is the operation of removing an edge and replacing it with a new edge such that the resulting graph is again a plane spanning tree. In our setting, we further require this pair of edges to be non-crossing. In fact, one can simulate transformations via compatible trees in terms of crossing free edge flips: for two compatible trees T_1, T_2 , successively add edges from T_2 to T_1 , while removing an edge that is not in T_2 from the resulting cycle.

We observe that the compatibility graph of simple drawings that are not of the complete graph might not be connected even if the graph is dense. For example, Figure 1 shows a simple drawing of the complete bipartite graph containing a plane tree that crosses all edges of the graph not belonging to the tree. Hence, this tree is an isolated vertex in the corresponding compatibility graph.

Related work. The problem of transforming elements within a class of objects (e.g. plane spanning trees or matchings) into each other via a certain operation

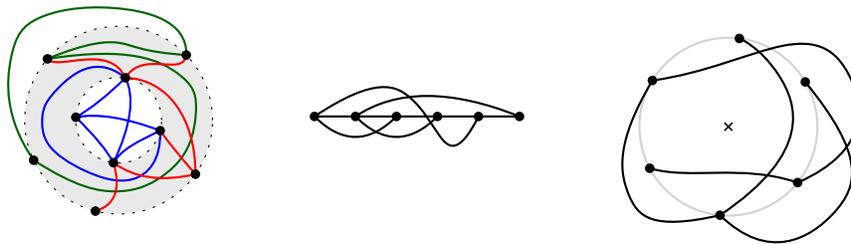


Fig. 2: *Left to right*: cylindrical, monotone, strongly c -monotone drawing.

(e.g. edge flips or compatibility) has been studied extensively in a huge variety of contexts. Considering edge flips, some of the earliest results have been obtained on triangulations: Wagner [15] showed connectivity of the corresponding flip-graph in the combinatorial setting and Lawson [12] in the straight-line setting. For more details we refer the reader to the survey of Bose and Hurtado [8].

Considering the notion of compatibility, most of the work has been done in the straight-line setting, e.g., in the context of perfect matchings with [5,7] or without [1,2] vertex coloring, or for edge-disjoint compatibility [3,11]. Aichholzer et al. [4] showed, in the straight-line setting, that the compatibility graph of plane spanning trees is connected with diameter $O(\log k)$, where k denotes the number of convex layers of the point set. Buchin et al. [9] provided a corresponding worst case lower bound of $\Omega(\log n / \log \log n)$.

It is natural to extend this question to simple drawings, which however are inherently difficult to handle (even the existence of certain plane substructures is still unresolved in simple drawings; see e.g. [13]). On the positive side, García, Pilz and Tejel [10] proved that any maximal plane subgraph is 2-connected, which guarantees for any plane spanning tree the existence of a compatible plane spanning tree. In this paper, we aim to shed some light on this wide open topic of compatibility graphs of trees in simple drawings.

Contribution. We approach Question 1 from two directions, proving a positive answer for special classes of drawings (namely, cylindrical, monotone, and strongly c -monotone drawings) and for special classes of spanning trees (namely stars, double stars, and twin stars). We postpone the precise definitions of these classes of drawings and graphs to the later sections, however, Figure 2 gives an illustration of these notions.

Theorem 1. *Let D be a cylindrical, monotone, or strongly c -monotone drawing of the complete graph K_n . Then, the compatibility graph $\mathcal{F}(\mathcal{T}_D)$ is connected.*

Theorem 2. *Let D be a simple drawing of the complete graph K_n and let \mathcal{T}_D^* be the set of all plane spanning stars, double stars, and twin stars on D . Then, the compatibility graph $\mathcal{F}(\mathcal{T}_D^*)$ is connected.*

Section 2 is devoted to the proof of Theorem 1, while Section 3 is dedicated to the proof of Theorem 2. All results that are marked by a (clickable) (\star) have a full proof in the appendix.

2 Special simple drawings of K_n

In this section we prove connectedness of the compatibility graph for certain classes of drawings. Clearly, for any drawing of K_n that admits a plane spanning tree which is not crossed by any edge of D , the compatibility graph is connected with diameter at most 2. This is, for example, the case for 2-page book drawings, where the vertices are placed along a line and each edge lies entirely in one of the two open halfplanes defined by this line.

2.1 Cylindrical drawings

Following the definition of Schaefer [14], in a *cylindrical drawing* of a graph the vertices are placed along two concentric circles, the *inner* and *outer* circle, and no edge is allowed to cross these circles.

Lemma 1. (\star) *Let D be a cylindrical drawing of K_n . Then $\mathcal{F}(\mathcal{T}_D)$ is connected with diameter at most 4.*

2.2 Monotone drawings

A simple drawing in which no two vertices have the same x -coordinate and every edge is drawn as an x -monotone curve is called *monotone drawing*. Let v_1, v_2, \dots, v_n denote the sequence of vertices in increasing x -order. W.l.o.g. assume that these vertices are on the x -axis. Then, the plane spanning path $\mathcal{S} = v_1, v_2, \dots, v_n$ is called *spine path*. An edge that intersects the spine path is called *twiggy edge*.

We define a relation on the twiggy edges of D as follows: for two twiggy edges e, f we have $e \succ f$ if they are non-intersecting and admit a vertical line intersecting the relative interiors of both edges that intersects e at a larger y -coordinate than f . All other pairs of twiggy edges are incomparable. For a set E of pairwise non-intersecting twiggy edges, an edge $e \in E$ is *maximal* if there is no other edge $f \in E$ s.t. $f \succ e$. Note that this relation is acyclic, i.e., there are no twiggy edges e_1, \dots, e_k such that $e_1 \succ e_2 \succ \dots \succ e_k \succ e_1$. And hence, any non-empty set of twiggy edges admits a maximal element.

Lemma 2. (\star) *For any monotone drawing D of K_n , the compatibility graph $\mathcal{F}(\mathcal{T}_D)$ is connected with diameter $O(n)$.*

Proof (Sketch). We show that any plane spanning tree T in D can be transformed to the spine path \mathcal{S} . If T does not contain any twiggy edge, clearly it is compatible to \mathcal{S} . Otherwise, we proceed as follows. Corresponding to a maximal twiggy edge e of T , we find a path P' connecting the vertices of e (see Figure 3).

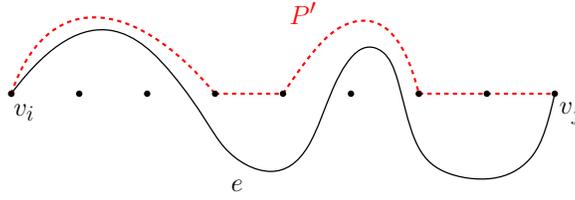


Fig. 3: The (maximal) twiggly edge $e = v_i v_j$ divides the vertices between v_i and v_j into two groups – above and below. The path P' is formed by joining the consecutive vertices lying above e including the vertices of e .

We can show that P' is compatible to T and lies strictly above e . Thus, we can add P' to T , which creates at least one cycle in T . Removing appropriate edges including e , we get a compatible tree with at least one twiggly edge less and repeating this process, we will eventually reach the spine path \mathcal{S} .

2.3 Strongly c -monotone drawings

A curve is called c -monotone (w.r.t. a point x) if every ray emanating from x intersects the curve at most once. A simple drawing is c -monotone, if all vertices are drawn along a circle and every edge is a c -monotone curve w.r.t. the center of the circle. A c -monotone drawing is *strongly* c -monotone if for any pair of edges e, e' there is a ray (rooted at the circle center) that neither intersects e nor e' .

In a (strongly) c -monotone drawing, we label the vertices v_1, v_2, \dots, v_n in cyclic order and denote the center of the circle by c . In the following, we often consider edges and their intersections with rays rooted at c ; unless stated otherwise, any ray is rooted at c and edges are intersected in their relative interiors.

An edge e connecting two consecutive vertices v_i, v_{i+1} is called *cycle edge* and if e is drawn along the “shorter” side of the circle it is called *spine edge* (that is, no ray formed by the center and any vertex intersects e). All spine edges form the *spine* and any path consisting entirely of spine edges is called *spine path*.

Lemma 3. (\star) *Any strongly c -monotone drawing D of K_n either has all cycle edges as spine edges or is isomorphic to a monotone drawing.*

Again, we define *twiggly* edges to be those that intersect a spine edge. A crucial difference to the monotone setting is that an analogue to the relation ‘ \succ ’ (adjusted with respect to the intersection with rays emanating from c) may now be cyclic and hence, we cannot guarantee the existence of a *maximal* twiggly edge anymore. We therefore need a different approach.

For a twiggly edge $e = uw$, let x_1, \dots, x_k be its crossings with the spine (note that these are not vertices of K_n) and assume the labeling to be in such a way that u, x_1, \dots, x_k, w appear in clockwise order. For $i \in \{1, \dots, k\}$ denote the vertex (of K_n) in clockwise order before x_i by x_i^- and the one after by x_i^+ . Furthermore, set $u = x_0^-$ and $w = x_{k+1}^+$. Then, for $i \in \{0, \dots, k\}$, we call the

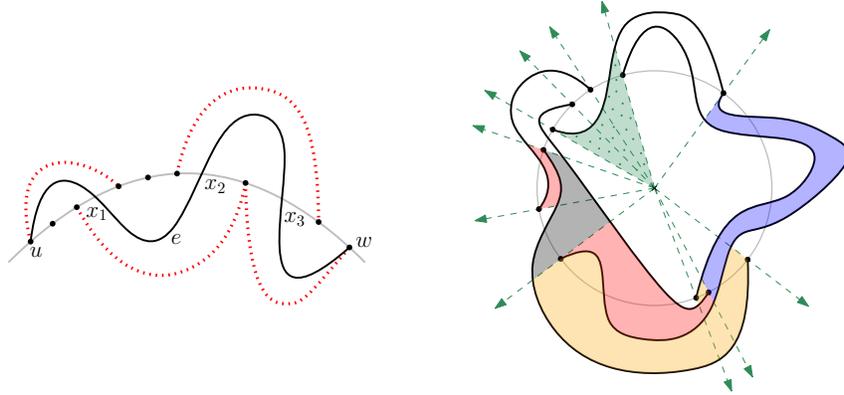


Fig. 4: *Left*: The red dotted edges are bumpy edges of the twiggly edge e . *Right*: A set of twiggly edges and some corridors; the dotted green is an inner corridor.

edges $x_i^- x_{i+1}^+$ *bumpy* edges (see Figure 4 (left)). Note that bumpy edges do not intersect the spine and for any twiggly edge there are at least two bumpy edges.

Clearly, we can identify any ray r with an angle θ , the angle it forms with the vertical ray (upwards). Two edges e, f are called *neighbours on an interval* $[\theta_1, \theta_2]$, if for any ray $r \in [\theta_1, \theta_2]$ the intersections of e and f with r appear consecutively on r . A *corridor* is a maximally connected region bounded by two neighbouring edges (along a maximal interval). Again, we identify corridors by an interval $[\theta_1, \theta_2]$ and usually we speak of corridors defined by the edges of a plane spanning tree. The *twiggly depth* (with respect to a plane spanning tree T) of a ray r is the number of twiggly edges (of T) that r intersects.

We extend our definition of neighbours (along an interval) also to the very inside and very outside by inserting a dummy edge at the circle center and one at infinity. More precisely, an edge e is the neighbor of the circle center c along an interval $[\theta_1, \theta_2]$ if for any ray $r \in [\theta_1, \theta_2]$ the intersection of r and e is closest to c (and furthest in the case of being a neighbor of infinity). We call the corresponding corridors *inner/outer* corridors. Note that the set of all corridors partitions the plane. See Figure 4 (right) for an illustration.

We further remark that for any plane spanning tree T , any corridor $C = [\theta_1, \theta_2]$ (of edges of T) begins and ends at a vertex, i.e., the rays at θ_1 and θ_2 hit a vertex.

Lemma 4. (\star) *For any plane spanning tree T of a strongly c -monotone drawing D and any corridor C of T with start and end vertex s and t , there is a path P in D from s to t staying entirely in C , that does not intersect T . Furthermore, if C is an inner or outer corridor, P does not use any twiggly edge.*

Lemma 5. *For any strongly c -monotone drawing of K_n , the compatibility graph $\mathcal{F}(\mathcal{T}_D)$ is connected with diameter $O(n)$.*

Proof. Let D be a strongly c -monotone drawing of K_n and let T be a plane spanning tree. We show that T can be compatibly transformed to a spine path

(by iteratively decreasing its twiggly depth). By Lemma 3, we may assume that all n spine edges are present in D . Again, if there is no twiggly edge in T , then T is compatible with the spine.

Let E_{twig} be the set of twiggly edges of T and construct the set \mathcal{C} of all corridors. Next, for any corridor $C \in \mathcal{C}$ with start and end vertex s and t , we add the path P_C as guaranteed by Lemma 4 to T .

Clearly, we do not disconnect T when removing E_{twig} now. Indeed, let $e = uw \in E_{twig}$, then the collection of corridor paths below (and also above) e connects u and w . So we remove E_{twig} and potentially some further edges until T forms a spanning tree again (which by Lemma 4 is also plane). Furthermore, any ray r that intersects x previous twiggly edges (i.e., E_{twig}) intersects $x + 1$ corridors, two of which are either an inner or outer corridor. By Lemma 4 and the properties of c -monotone curves, r intersects at most $x - 1$ (new) twiggly edges. Hence, the twiggly depth of any ray decreased by at least one and we recursively continue this process until all rays have twiggly depth 0, in which case T is compatible to a spine path. As we have twiggly depth at most $n - 1$ in the beginning, $\mathcal{F}(\mathcal{T}_D)$ has diameter $O(n)$.

Theorem 1 now follows from Lemma 1, Lemma 2, and Lemma 5.

3 Special Plane Spanning Trees

In this section, we are not restricting our drawing anymore, i.e., D will be a simple drawing of K_n throughout this section. Instead we focus on special classes of spanning trees and show that the subgraph $\mathcal{F}(\mathcal{T}_D^*)$ of $\mathcal{F}(\mathcal{T}_D)$ induced by the set of vertices corresponding to stars, double stars, and twin stars is connected.

A plane spanning tree with a fixed path P of length k such that all other vertices are incident to either the start or end vertex of P is called a k -star. A 0-star (i.e., P consists of a single vertex) is called *star*. A 1-star is called *double star* and a 2-star is called *twin star*.

The following relation, introduced in [6], will be very useful: Given a simple drawing of K_n with vertex set V and two vertices $g \neq r \in V$, for any two vertices $v_i, v_j \in V \setminus \{g, r\}$, we define $v_i \rightarrow_{gr} v_j$ if and only if the edge $v_i r$ crosses $v_j g$. In [6] it is shown that this relation is asymmetric and acyclic.

We start by showing that stars can always be transformed into each other via a sequence of crossing free edge flips.

Lemma 6. *Any two stars in D have distance $O(n)$ in $\mathcal{F}(\mathcal{T}_D^*)$.*

Proof. Given a star T in g (i.e., g is incident to all other vertices of T), we can transform it into a star H in r via a sequence of crossing free edge flips, such that in every step, the graph is a double star with fixed path r, g , in the following way. We label the vertices in $V \setminus \{g, r\}$ such that $v_i \rightarrow_{gr} v_j$ implies $i < j$ (see Figure 5). We iteratively replace an edge $g v_i$ by $r v_i$ starting from $i = n - 2$ and continuing in decreasing order. Clearly, all intermediate trees are double stars (with fixed path r, g) and hence, it remains to argue that the flips are compatible,

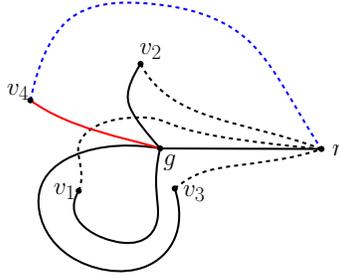


Fig. 5: Proof of Lemma 6: The solid edges represent a star in g , while the dotted edges form a star in r . The vertices are labeled conforming to the relation \rightarrow_{gr} . In order to transform the star in g to the star in r , the first step is adding the dotted blue edge v_4r and deleting the red edge v_4g .

i.e., for $i = n - 2, \dots, 1$ the edge gv_i does not cross any edge of the current T . By construction, in any step i , T contains edges of the form (a) rv_j for $j > i$ and (b) gv_k for $k < i$. The edge gv_i cannot cross edges in (a) by the definition of the relation \rightarrow_{gr} and also not those in (b) due to the properties of simple drawings. As we need at most $n - 2$ steps for the transformation, any two stars have distance $O(n)$ in $\mathcal{F}(\mathcal{T}_D^*)$.

Theorem 2 then follows from Lemma 6 in combination with the following two lemmata.

Lemma 7. (\star) Any double star in D has distance $O(n)$ to any star in $\mathcal{F}(\mathcal{T}_D^*)$.

Lemma 8. (\star) Any twin star in D has distance $O(n)$ to any star in $\mathcal{F}(\mathcal{T}_D^*)$.

References

1. Aichholzer, O., Bereg, S., Dumitrescu, A., García, A., Huemer, C., Hurtado, F., Kano, M., Márquez, A., Rappaport, D., Smorodinsky, S., Souvaine, D., Urrutia, J., Wood, D.: Compatible geometric matchings. *Computational Geometry: Theory and Applications* **42**(6–7), 617–626 (2009). <https://doi.org/10.1016/j.comgeo.2008.12.005>
2. Aichholzer, O., García, A., Hurtado, F., Tejel, J.: Compatible matchings in geometric graphs. In: *Proceedings of the XIV Encuentros de Geometría Computacional (EGC'11)*. pp. 145–148 (2011)
3. Aichholzer, O., Asinowski, A., Miltzow, T.: Disjoint compatibility graph of non-crossing matchings of points in convex position. *Electronic Journal of Combinatorics* **22**(1) (2015). <https://doi.org/10.37236/4403>
4. Aichholzer, O., Aurenhammer, F., Huemer, C., Krasser, H.: Transforming spanning trees and pseudo-triangulations. *Information Processing Letters* **97**(1), 19–22 (2006). <https://doi.org/10.1016/j.ipl.2005.09.003>, <https://www.sciencedirect.com/science/article/pii/S0020019005002486>

5. Aichholzer, O., Barba, L., Hackl, T., Pilz, A., Vogtenhuber, B.: Linear transformation distance for bichromatic matchings. *Computational Geometry: Theory and Applications* **68**, 77–88 (2018). <https://doi.org/10.1016/j.comgeo.2017.05.003>
6. Aichholzer, O., Parada, I., Scheucher, M., Vogtenhuber, B., Weinberger, A.: Shooting Stars in Simple Drawings of $K_{m,n}$. In: Proc. 35th European Workshop on Computational Geometry EuroCG '19. pp. 59:1–59:6. Utrecht, The Netherlands (2019), <http://www.eurocg2019.uu.nl/papers/59.pdf>
7. Aloupis, G., Barba, L., Langerman, S., Souvaine, D.: Bichromatic compatible matchings. *Computational Geometry: Theory and Applications* **48**(8), 622–633 (2015). <https://doi.org/10.1016/j.comgeo.2014.08.009>
8. Bose, P., Hurtado, F.: Flips in planar graphs. *Computational Geometry* **42**(1), 60–80 (2009). <https://doi.org/10.1016/j.comgeo.2008.04.001>, <https://www.sciencedirect.com/science/article/pii/S0925772108000370>
9. Buchin, K., Razen, A., Uno, T., Wagner, U.: Transforming spanning trees: A lower bound. *Computational Geometry* **42**(8), 724–730 (2009). <https://doi.org/10.1016/j.comgeo.2008.03.005>
10. García, A., Pilz, A., Tejel, J.: On plane subgraphs of complete topological drawings. *ARS MATHEMATICA CONTEMPORANEA* **20**, 69–87 (2021), <https://doi.org/10.26493/1855-3974.2226.e93>
11. Ishaque, M., Souvaine, D.L., Tóth, C.D.: Disjoint compatible geometric matchings. *Discrete & Computational Geometry* **49**, 89–131 (2013). <https://doi.org/10.1007/s00454-012-9466-9>
12. Lawson, C.L.: Transforming triangulations. *Discrete mathematics* **3**(4), 365–372 (1972)
13. Rafla, N.H.: The Good Drawings D_n of the Complete Graph K_n . Ph.D. thesis, McGill University, Montreal (1988), https://escholarship.mcgill.ca/concern/file_sets/cv43nx65m?locale=en
14. Schaefer, M.: The graph crossing number and its variants: A survey. *Electronic Journal of Combinatorics* **1000** (2013)
15. Wagner, K.: Bemerkungen zum Vierfarbenproblem. *Jahresbericht der Deutschen Mathematiker-Vereinigung* **46**, 26–32 (1936), <http://eudml.org/doc/146109>

A Missing proofs of Section 2

We need a few more definitions and observations for the proof of Lemma 1. We denote the vertices on the inner/outer circle of a cylindrical drawing by *inner/outer vertices* and the edges connecting any two inner/outer vertices by *inner/outer edges*. The remaining edges are called *side edges*. Furthermore, edges that join consecutive vertices of either of the two circles are called *cycle edges*.

Remark 1. Cycle edges may only be crossed by side edges and either of the two circles contains at most one such cycle edge that is crossed.

In other words, along both circles there is a Hamiltonian path of cycle edges that is not crossed by any other edge.

Further note that any plane spanning tree in a cylindrical drawing contains at least one side edge.

Lemma 1. *(★) Let D be a cylindrical drawing of K_n . Then $\mathcal{F}(\mathcal{T}_D)$ is connected with diameter at most 4.*

Proof. (Disclaimer: we moved a few definitions and remarks, that are relevant for this proof, to the preceding paragraphs.) Let $T_1, T_2 \in \mathcal{T}_D$ and $e_1 = v_s v_r$ be a side edge of T_1 as well as $e_2 = v'_s v'_r$ be a side edge of T_2 . Furthermore, as guaranteed by Remark 1, let $S_1 \in \mathcal{T}_D$ be a plane spanning tree consisting of the two Hamiltonian paths of cycle edges and e_1 (and similarly $S_2 \in \mathcal{T}_D$ consists of the cycle paths and e_2). Clearly, T_1 and S_1 are compatible as well as T_2 and S_2 . If S_1 and S_2 are compatible, we are done. Otherwise, e_1 and e_2 are crossing in D . As the induced 4-tuple $\{v_s, v'_s, v_r, v'_r\}$ has at most one crossing, the two side edges $v_s v'_r$ and $v'_s v_r$ cannot cross e_1 and e_2 . Now consider one of these edges as e' and construct a plane spanning tree S_3 , using only e' and the uncrossed cycle paths. Clearly, S_3 is compatible to S_1 as well as S_2 and hence, T_1, S_1, S_3, S_2, T_2 is a path of length 4 in $\mathcal{F}(\mathcal{T}_D)$.

Lemma 2. *(★) For any monotone drawing D of K_n , the compatibility graph $\mathcal{F}(\mathcal{T}_D)$ is connected with diameter $O(n)$.*

Proof. We prove the connectedness by showing that any plane spanning tree in D can be transformed (via a sequence of compatible trees) to the spine path \mathcal{S} . So, let $T \in \mathcal{T}_D$ be a plane spanning tree. If T does not contain any twiggly edge, then it is compatible to \mathcal{S} . Otherwise, let $e = v_i v_j$ (with $i < j$) be a maximal twiggly edge of T . Note that all twiggly edges in T are non-intersecting, since T is plane. Define

$$V_e^\uparrow = \{v_k : i < k < j, v_k \text{ is above } e\}.$$

Consider the path $P' = v_i, V_e^\uparrow, v_j$ that starts at v_i , ends at v_j and inbetween uses only vertices of V_e^\uparrow (see Figure 3). Note that the vertices in P' are in increasing x -order. Clearly, no edge in P' can intersect the twiggly edge e , since it would have to intersect it twice. Hence, P' lies (strictly) above e . Next we show that P' cannot intersect T . To this end, let $f \in T$ and assume, for the sake

of contradiction, that it intersects an edge $v_x v_y$ of P' . Clearly, f can neither be incident to v_x nor v_y . However, by construction of P' , in order to reach its incident vertices, f has to either (i) intersect $v_x v_y$ twice, or (ii) intersect e , or (iii) intersect a spine edge. (i) yields a contradiction to the properties of simple drawings, (ii) yields a contradiction to the planarity of T , and (iii) yields a contradiction to the maximality of the twiggly edge e .

Hence, we can add P' to T without introducing any crossings. Note that some edges of P' may already be present in T , however, since this insertion creates a cycle, at least one is not. In order to reach a compatible tree, we remove e from this cycle (and potentially more edges until reaching a plane spanning tree again). We created a compatible tree with at least one twiggly edge less and repeating this process, we will eventually reach the spine path \mathcal{S} . As we have at most $n - 1$ twiggly edges, $\mathcal{F}(\mathcal{T}_D)$ has diameter $O(n)$.

Lemma 3. *(\star) Any strongly c -monotone drawing D of K_n either has all cycle edges as spine edges or is isomorphic to a monotone drawing.*

Proof. Assume there is an edge $e = v_i v_{i+1}$ of consecutive vertices which does not form a spine edge, i.e., e intersects any ray which is not in the wedge $v_i c v_{i+1}$. The c -monotonicity of edges implies that any edge intersecting the wedge $v_i c v_{i+1}$ must intersect any ray in this wedge. This, however, would be in contradiction to the property of strongly c -monotone drawings and hence, the wedge $v_i c v_{i+1}$ is not intersected by any edge. Clearly, this implies the drawing to be (strongly) isomorphic to a monotone drawing: any ray r in the wedge $v_i c v_{i+1}$ is uncrossed by all edges and hence, we can cut the spine at the intersection with r and stretch it to a monotone drawing.

Lemma 4. *(\star) For any plane spanning tree T of a strongly c -monotone drawing D and any corridor C of T with start and end vertex s and t , there is a path P in D from s to t staying entirely in C , that does not intersect T . Furthermore, if C is an inner or outer corridor, P does not use any twiggly edge.*

Proof. Start the path P at s and always connect to the next vertex (along the circle) that is in the corridor C . Due to strong c -monotonicity, the corresponding edge must run through C . Note that P may consist of bumpy, spine, and twiggly edges. However, P may only contain a twiggly edge if it is forced to by *both* edges bounding the corridor (e.g., the blue corridor in Figure 4 (right)).

Moreover, since we always connect consecutive vertices along the corridor, any edge e , intersecting an edge f of P , must *surround*⁴ f . We claim that e must also be a twiggly edge, which can be seen as follows. The twiggly edge f together with the spine forms a collection of faces which are only bounded by f and the spine (see Figure 6). Since e surrounds f , e cannot have a vertex within such a face. Hence, if e intersects f , e must also intersect the spine and thus, is a twiggly edge. Therefore, e cannot belong to T (since it is in the corridor). In particular, P does not intersect T .

⁴ We say that an edge $e = uv$ *surrounds* an edge f (all four vertices are distinct), if the wedge ucv (containing e) also contains f ; see Figure 6.

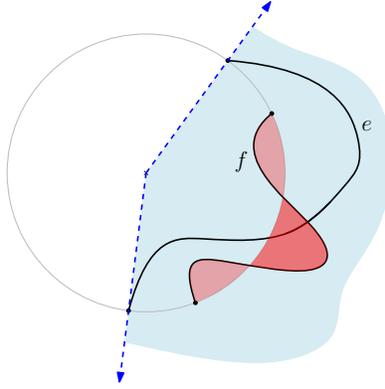


Fig. 6: The faces formed by f and the spine are red. If e intersects f it must leave the corresponding face by intersecting the spine.

Lastly observe that the paths in an inner/outer corridor do not use any twiggly edge, since any inner/outer corridor is bounded by a dummy edge which does not cross the spine.

B Missing proofs of Section 3

Lemma 7. (\star) Any double star in D has distance $O(n)$ to any star in $\mathcal{F}(\mathcal{T}_D^*)$.

Proof. Let T be a double star with fixed path g, r and T' be a star in r' . T can be transformed to a star in r (or g) completely analogous as in the proof of Lemma 6. The only difference is that there are additional edges attached to r , i.e., edges of type (a) in Lemma 6, which do not interfere. Next, using Lemma 6 again, we transform this star in r to T' . This also implies the distance between a double star and a star in $\mathcal{F}(\mathcal{T}_D^*)$ to be $O(n)$.

Lemma 8. (\star) Any twin star in D has distance $O(n)$ to any star in $\mathcal{F}(\mathcal{T}_D^*)$.

Proof. Let T be a twin star with fixed path g, s, r and T' be a star in r' . Note that all edges in T are incident to g or r and hence, we can add the edge gr and remove gs or rs in order to obtain a double star. Using Lemma 7 this double star can be transformed to T' . Thus, a twin star can be transformed to a star in $O(n)$ steps.