Journal of Graph Algorithms and Applications
http://jgaa.info/ vol. 17, no. 4, pp. 367–440 (2013)
DOI: 10.7155/jgaa.00298

Toward a Theory of Planarity: Hanani-Tutte and Planarity Variants

Marcus Schaefer[1]

[1]School of Computing
DePaul University
Chicago, IL 60604, US

Abstract

We study Hanani-Tutte style theorems for various notions of planarity, including partially embedded planarity and simultaneous planarity. This approach brings together the combinatorial, computational and algebraic aspects of planarity notions and may serve as a uniform foundation for planarity, as suggested earlier in the writings of Tutte and Wu.

Submitted:	Reviewed:	Revised:	Accepted:	Final:
November 2012	March 2013	April 2013	May 2013	May 2013
		Published:		
		July 2013		
	Article type:		Communicated by:	
	Regular paper		W. Didimo and M. Patrignani	

An abstract of this paper was published in the proceedings of Graph Drawing, 2012.
E-mail address: mschaefer@cdm.depaul.edu (Marcus Schaefer)

1 Introduction

Planarity of graphs is a well-studied and well-understood topic, but as soon as we modify planarity in any one of many different ways such as allowing crossings, clustering vertices, requiring monotone drawings, simultaneously drawing multiple graphs, or partially embedding the graph, we very quickly lose the ground under our feet; some problems become **NP**-complete (upward planarity, book embeddability), for others feasible algorithmic solutions are unknown (c-planarity, constrained planarity, simultaneous planarity).

In 1972 Tutte published his paper "Toward a Theory of Crossing Numbers" in which he suggested an algebraic treatment of crossing numbers.[1] This approach has led to some research on crossing number variants, but it had little impact on the crossing number itself. Our plan is to investigate how the algebraic approach fares for crossing number zero. The most famous result in Tutte's paper (found earlier by Hanani) is the Hanani-Tutte theorem which states that a graph is planar if and only if it can be drawn in the plane so that every pair of independent edges crosses an even number of times (including not at all). We rephrase this as a crossing number result: given a drawing D of G, let $\text{iocr}(D)$ be the number of pairs of independent edges of G that cross oddly in D. The *independent odd crossing number* of G, $\text{iocr}(G)$, is defined as the minimum of $\text{iocr}(D)$ over all drawings of G. In crossing number terminology, the Hanani-Tutte theorem states that a graph G is planar if and only if $\text{iocr}(G) = 0$.

The Hanani-Tutte theorem opens up an algebraic approach to planarity; the condition $\text{iocr}(G) = 0$ can be written as a system of linear equations over GF(2), leading to a simple polynomial-time algorithm for planarity testing; we discuss this well-known algebraic criterion in Section 3. This algorithm is not very efficient—it takes $O(n^6)$ time, where $n = |V(G)| + |E(G)|$, but it can be implemented so it is feasible for small graphs.

Remark 1.1 (Running Times)**.** For many of the planarity problems we solve using Hanani-Tutte there are existing algorithms running in linear or quadratic time. Algorithms based on Hanani-Tutte will not be able to compete with these running times unless we find a clever way to implement Hanani-Tutte style systems (see Remark 3.5 for a positive example). My current implementation is written in Python 3.2 and runs on an Intel I7 processor. The program can solve a simultaneous planarity problem on 50 vertices and 175 edges in about 40 seconds. The linear system for this example has dimension 9376×8400 (it's a sparse system though), and takes 8 seconds to create and 32 seconds to solve. We hope to be able to report more detailed results in a later article.

In this paper we begin a systematic study of whether and how the Hanani-Tutte theorem extends to variants of planarity. The theorem turns out to be very versatile and adaptable, giving rise to a uniform approach to many of the variants of planarity considered in the literature. Table 1 summarizes known and new results on Hanani-Tutte theorems. All planarity notions will be defined and discussed in the next section. We describe the complexity of problems using

[1] There were precursors to his approach, notably the paper by Hanani [16], but also work by Flores, van Kampen, and Wu. Some of the history can be found in [67].

standard classes from computational complexity, including **P** for polynomial time and **NP** for non-deterministic polynomial time. **NP**-hard means as hard as any problem in **NP**, and **NP**-complete, or **NPC**, means **NP**-hard and in **NP** [36].

All of our Hanani-Tutte characterizations are based on redrawing results, collected in Section 4. The following conjecture gives a flavor of what these redrawing results look like.

Conjecture 1.2. *Suppose a graph G with subgraph H can be drawn so that every edge of H crosses every independent edge of G evenly. Then there is a drawing of G in which edges of H do not cross each other, and there are no new pairs of independent edges crossing oddly.*

The truth of this conjecture would imply a single polynomial-time algorithm for nearly all known planarity variants, including the infamous case of c-planarity. The reason for this is that, (i), Conjecture 1.2 implies that simultaneous planarity of two graphs can be tested using a simple algebraic criterion (see Corollary 6.23 and Lemma 6.24), and, (ii), nearly all planarity variants are special cases of simultaneous planarity of two graphs, including c-planarity (Theorem 6.17), and variants of constrained, book and level planarity.[2] The results on relationships between different planarity variants will be discussed in the next section; Figure 2 summarizes the results.

Algorithmically, the algebraic approach cannot (currently) compete with PQ-trees and SQPR-trees which give linear-time algorithms in many cases. It does, however, lead to a deeper understanding of planarity, by offering a unified view of planarity, and, in some cases, yields algorithmic solutions where no other algorithms are currently available.

The paper proceeds as follows: In Section 2 we introduce all relevant notions of planarity, and summarize what is known about them with respect to our focus of interest: complexity of the recognition problem, Hanani-Tutte characterizations, obstruction sets, and relationships between them. In Section 3 we explain our approach in some detail for the case of standard planarity, reviewing some well-known results. Section 4 collects redrawing tools used in the remainder of the paper. The core of the paper then consists of Section 5 on partially embedded planarity, which is a powerful notion to have control over, and Section 6 on simultaneous planarity, which we do not yet fully control.

2 Notions of Planarity

We introduce various notions of planarity, and summarize known and new results. For an overview of variants of planarity, see the survey by Maurizio

[2] To clarify the status of the polynomial-time algorithm for simultaneous planarity: the algorithm exists and runs in polynomial time, we just do not know whether it is correct. However, it can be modified (while still running in polynomial time) so it either gives a correct answer, or, presents a counterexample to Conjecture 1.2. We have to leave the details for a later paper.

planarity notion	recognition	obstr.	Hanani-Tutte result	Hanani-Tutte algorithm
standard	linear [44, 61]	[52]	[16, 73]	Cor 3.4 (folklore)
outer	linear [44]	[13]		
partially embedded	linear [3]	[46]	Thm 5.6	Cor 5.12
partial rotation	linear, Cor 5.10	open	Thm 5.9	
partial rotation (with flips)	in **P**, Cor 6.43	open	open	open
partially constrained PQ-planarity	special cases in linear [11]	open	open	open
ec-planarity	linear [40]	open	open	open
ec-planarity (with free edges)	open	open	open	open
x-monotone	linear [48]	open	[56, 57, 34]	quadratic [34]
level	linear [48]	open	Thm 6.8, [34]	quadratic [34]
$\mathcal{T}$-coherent level	special cases in **P**	open	open	special case, Cor 6.11
radial level	linear [7]	open	open	open
upward	**NPC** [37]		Ex 2.2	
projective	linear	[39, 6]	[62]	open
book	**NPC** [17], special cases in **P**		special case, Sec 6.1.2	special cases
partitioned book	**NPC** [45]		open	open
c (clustered)	in **NP**, special cases in **P**	open	special cases	open
cl (clustered level)	open (in **NP**)	open	open	open
simultaneous	in **NP** [38], special cases in **P**	open	special cases, Sec 6	special cases

Table 1: Summary of known and new results on planarity variants.

Patrignani in the "Handbook of Graph Drawing" [61]. A *drawing* of a graph maps the vertices of the graph to distinct points in the plane and edges to simple arcs connecting the endpoints of the edge. We reserve the word *embedding* for crossing-free drawings, that is, drawings in which no two edges have any points in common (other than a common endpoint).[3]

2.1 Planarities and Hanani-Tutte

Planarity A graph is *planar* if it can be *embedded* in the plane, that is, drawn so that no two edges cross each other. There are several linear-time algorithms for recognizing planar graphs and embedding them, the first being due to Hopcroft and Tarjan [44]. For surveys, see [55, 61]. Kuratowski [52] showed that planar graphs can be characterized by a set of excluded (topological) minors, namely $\{K_{3,3}, K_5\}$. The classic Hanani-Tutte theorem was phrased for this standard notion of planarity [16, 73]. We will review this material in Section 3.

Outerplanarity. Outerplanarity easily reduces to planarity by adding an apex vertex[4] to the graph, so testing and embedding are special cases of planarity. There is a finite obstruction set, $\{K_4, K_{2,3}\}$ [13]. There is no separate Hanani-Tutte theorem, but van der Holst has found a homological characterization of outerplanar graphs [74].

Constrained Planarity Constrained planarity is a somewhat vague term describing various planarity restrictions [61], but it mostly seems to be associated with restricting the rotation system of an embedding. A *rotation* at a vertex is a cyclic ordering of the edges incident to the vertex. A drawing *realizes* the rotation if the cyclic clockwise ordering of the edges leaving the vertex in the drawing corresponds to the rotation. A *rotation system* is a collection of rotations for each vertex of the graph. Embeddings are typically described by rotation systems (though that does not work so well for disconnected graphs, see Remark 5.1; that is the reason that partially embedded planarity warrants a separate entry below). In the reverse direction, we can ask whether a graph has an embedding given a rotation system. This becomes interesting if instead of specifying a particular rotation at each vertex, there is a more general scheme determining a set of admissible rotations. For example, instead of specifying the cyclic order of all edges incident to a vertex, we could specify a *partial rotation* in which the ordering is fixed only for a subset of the edges incident to v. We are aware of two other schemes[5]: ec-planarity[6], due to Gutwenger, Klein, and

[3] For simultaneous embeddings, which we will see later, we do allow some crossings in the drawing, but only between edges belonging to different graphs. These crossings do not count.

[4] An apex vertex is adjacent to every every other vertex of the graph.

[5] We will not be able to do full justice to ec-planarity or partially constrained PQ-planarity, there is not time and space enough to include detailed definitions; we have to refer the reader to the original papers [40, 11].

[6] The acronym "ec" abbreviates "embedding constraint".

Mutzel [40], and partially constrained PQ-planarity, introduced by Bläsius and Rutter [11]. These notions combine a couple of common elements in different ways: a tree is used to limit admissible rotations at a vertex. The nodes of the tree can: allow an arbitrary ordering of their children (P-nodes in [11], gc-nodes in [40]), fix the clockwise order of their children (oc-nodes in [40]), or allow the rotation of the children to be either as specified or reversed (mirrored, thus mc-node in [40], Q-nodes in [11]). It may appear that ec-planarity is more powerful, than partially constrained PQ-planarity, since it allows oc-nodes, but that is not the case, since ec-planarity does not allow for partial rotations, which partially constrained PQ-planarity does. Gutwenger, Klein, and Mutzel mention this option, calling it ec-planarity *with free edges*. To the extent of my knowledge, this is the most general constrained planarity notion available at this point.

We know that ec-planarity (without free edges) can be tested in linear time [40], and partially constrained PQ-planarity can be tested in linear time for 2-connected graphs [11] (and it is conceivable that this result can be extended to allow oc-nodes as well). We consider the case of partial rotation systems (where the rotation at each vertex is constrained by an oc-node with free edges) and partial rotation systems with flips (where the rotation at a vertex is constrained by a single mc- or oc-node with free edges). For partial rotation systems, there is a Hanani-Tutte characterization, Theorem 5.9, and a linear-time algorithm, Corollary 5.10. Partial rotation systems with flips can be decided in polynomial time, Corollary 6.43, as a special case of a simultaneous planarity problem. We do not know about obstruction sets or appropriate notions of minor-orderings for the various constrained planarity notions.

Partially embedded planarity Given a graph G and an embedding $\mathcal{H}$ of a subgraph H of G, we can ask whether $\mathcal{H}$ can be extended to an embedding of G, or in other words, whether there is a planar drawing of G that contains $\mathcal{H}$. There is an SPQR-tree algorithm for testing partially embedded planarity in linear time due to Angelini, Di Battista, Frati, Jelínek, Kratochvíl, Patrignani and Rutter [3]. There also is a finite obstruction set found by Jelínek, Kratochvíl and Rutter [46]. We will establish a Hanani-Tutte characterization based on this obstruction set. This yields a new (if not particularly efficient) polynomial-time algorithm for testing partially embedded planarity.

x-monotonicity. An *ordered* graph is a graph $G = (V, E)$ together with a linear ordering of V. An *x-monotone drawing* of G is a drawing in which the left-to-right ordering of vertices agrees with the linear ordering and each edge is *x-monotone*, that is, it is drawn like the graph of a function (every vertical line crosses each edge at most once). The x-monotone graphs are a special case of leveled graphs (which we discuss shortly), and as such they are known to be recognizable and embeddable in linear time [48, 47]. In spite of several attempts, a complete obstruction set for

x-monotone graphs is not currently known. Pach and Tóth [56, 57] proved a weak Hanani-Tutte theorem for x-monotone graphs.[7] A strong Hanani-Tutte theorem for x-monotone graphs is established in [34]; it yields a simple, quadratic-time algorithm for testing x-monotonicity based on the Hanani-Tutte characterization. If one requires a drawing to be monotone with respect to both x and y, one gets the notion of *bi-monotonicity*, introduced in [34]; an example in that paper shows that even the weak Hanani-Tutte theorem fails for bi-monotonicity.

Level planarity. A *leveled* graph is a graph $G = (V, E)$ with a *leveling* ℓ of the vertices, where $\ell : V \to \mathbb{N}$. A *leveled embedding* of a leveled graph is an embedding in which $u \in V$ is to the left (right) of $v \in V$ if $\ell(u) < \ell(v)$ ($\ell(u) > \ell(v)$), and each edge is x-monotone. The difference to x-monotone graphs is that we may have $\ell(u) = \ell(v)$ in which case u and v have to lie on the same vertical line. Level planarity testing and embedding are solvable in linear time due to work by Jünger, Leipert, and Mutzel [48, 47]. No complete obstruction sets are known [25]. The Hanani-Tutte characterizations of x-monotonicity can be extended to level planarity (this follows from work in [34], see Theorem 6.8). The quadratic-time Hanani-Tutte algorithm for testing x-monotonicity can be adapted to test level planarity [34].

$\mathcal{T}$-coherent level planarity. A *generalized k-ary tanglegram* is a triple $(G, \ell, \mathcal{T})$ consisting of a leveled graph (G, ℓ) and a family $\mathcal{T}$ of trees $T_1, \ldots, T_n$ so that the leaves of T_i are exactly the vertices at level i in (G, ℓ).[8] We say $(G, \ell, \mathcal{T})$ can be *embedded* if (G, ℓ) has a *$\mathcal{T}$-coherent* level planar embedding, that is, if (G, ℓ) has a level planar embedding in which the ordering of the vertices at level i is consistent with an ordering of the leaves of T_i in a facial walk along some embedding of T_i (this corresponds to a PQ-tree without Q-nodes). The complexity of the general problem is open, though there are polynomial-time results for special cases in work by Wotzlaw, Speckenmeyer and Porschen [75]. No obstruction sets or Hanani-Tutte style characterizations are known.

Radial level planarity Radial level planarity is a variant of level planarity in which vertices are placed on concentric circles rather than parallel lines. Edges have to keep moving away from the common center c, they cannot double back; more formally, an edge is *radial monotone* if it crosses every cycle with center c at most once. All edges in a radial level planar drawing must be radial monotone. Radial level planarity can be recognized in linear time [7]. We are not aware of any work on obstruction sets or Hanani-Tutte style characterizations.

[7] Remark 3.2 explains the difference between weak and strong variants of Hanani-Tutte.

[8] Generalized k-ary tanglegrams were introduced by Wotzlaw, Speckenmeyer and Porschen [75]. In the original definition, (G, ℓ) is required to be *proper*, that is all edges have to be between adjacent levels. We drop that restriction here.

Upward planarity. A directed graph is *upward planar* if it can be drawn so that the tail of each edge is drawn to the left of its head, and each edge is x-monotone. Testing a directed graph for upward planarity is **NP**-complete [37].[9] That would seem to make a Hanani-Tutte style characterization of upward planarity unlikely, since such characterizations typically lead to polynomial-time algorithms, but this is not necessarily the case. We'll argue in Example 2.2 that there is a Hanani-Tutte theorem for upward planarity. We do not expect a nice obstruction set, though.

Projective planarity. A graph is *projective planar* if it can be embedded in the projective plane. The testing and embedding problem for projective planarity are in linear time [53]. A finite obstruction set is known [39, 6], and, based on that, a Hanani-Tutte characterization [62]. It is open whether the Hanani-Tutte theorem generalizes to any other surface, e.g. the Klein Bottle or the torus.

Book embeddings. A *k-page book* consists of k half-planes identified along the boundary (the *spine*). In an embedding of a graph in a k-page book all vertices have to lie on the spine, and edges may not cross each other or the spine, so each edge is embedded in a particular page. Graphs embeddable in a single page are just the outerplanar graphs. Graphs embeddable in a 2-page book are exactly the planar subgraphs of Hamiltonian graphs, so the recognition problem is **NP**-complete [17, Corollary 4.4]. Variants include assigning edges to pages, called *partitioned book embeddings* in [4], and restricting the ordering of the vertices along the spine to be consistent with the ordering of the leaves of a given tree T, called *T-coherent* in [4]. Partitioned book embeddability is known to be **NP**-complete [45], but is is open whether this remains true for a fixed number of pages; partitioned 2-page book embeddability can be tested in **P** [43, 1].

c-planarity. Roughly speaking, a *clustered graph* is a graph in which certain subsets of the vertices are identified as belonging together (clustered). The clustered graph is then called c-planar if it can be drawn so that vertices belonging to the same cluster can be drawn in the same region. We will give a formal definition in Section 6.1.6. There have been partial results on testing and embedding c-planar graphs in polynomial time, but the general problem remains open [61, Section 1.8.2]. We are not aware of any obstruction sets or a Hanani-Tutte style characterization of c-planarity, but there is recent work by Fulek, Kynčl, and Pálvölgyi for special cases (e.g. the case of two clusters) [33].

Clustered level planarity. Given a graph which is equipped with a leveling and a clustering, one can ask whether there is a level planar drawing of the graph which is also c-planar so that every cluster is bounded by two x-monotone curves. This notion was introduced in [30, 31], the complexity

[9] Chimani and Zeranski [15] recently suggested an exact algorithm for solving upward planarity that seems to perform well in practice.

of the recognition problem seems to be open (we remark that c-planarity is not obviously a special case, because clusters are not allowed to double-back between levels).

Simultaneous planarity. Two graphs $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ that have some vertices and edges in common are *simultaneously planar* if there are planar drawings of each, in which the common graph $G_1 \cap G_2 := (V_1 \cap V_2, E_1 \cap E_2)$ is drawn identically. This problem is also known as *Simultaneous Embeddability with Fixed Edges* for two graphs *(SEFE$_2$)*. A more flexible, equivalent, definition, requires a drawing of $G_1 \cup G_2 = (V(G_1) \cup V(G_2), E(G_1) \cup E(G_2))$ in which no two edges belonging to the same graph, G_1 or G_2, cross each other. This second definition has the advantage that it easily generalizes to an arbitrary number of graphs. We write SEFE$_k$ if we want to emphasize that we are considering the variant for k graphs, and SEFE if we allow arbitrary many graphs. The complexity of testing and embedding two graphs simultaneously is open, but several special cases have been settled [41, 4], also see the survey by Bläsius, Kobourov, and Rutter [9]. No obstruction sets are known, and there is no Hanani-Tutte style characterization. Conjecture 6.20 proposes a very natural Hanani-Tutte style characterization of which we prove several special cases, which in turn give polynomial-time algorithms for some recognition problems. Example 2.3 shows that we cannot expect a Hanani-Tutte theorem for simultaneous planarity of three or more graphs.

Weak realizability. A *topological graph* is a graph $G = (V, E)$ equipped with a symmetric relation $R \subseteq E \times E$. We say (G, R) is *weakly realizable* if there is a drawing of G in which only pairs of edges in R are allowed to cross. In a sense, weak realizability is *the* universal planarity problem, since it encodes the topological inference problem in the form of the region connection calculus RCC8 [69].[10] Testing weak realizability is **NP**-complete [50, 68]. Therefore, we do not expect there to be obstruction sets or a Hanani-Tutte style characterization. It was observed in [38] that simultaneous planarity (for arbitrary many graphs) and weak realizability are equivalent.[11] Example 2.3 then implies that there will not be a traditional Hanani-Tutte style characterization for weak realizability.

Remark 2.1 (Geometric Planarity)**.** Our list of planarity notions does not contain any geometric (straight-line, rectilinear) drawing variants. Hanani-Tutte does not seem to be the right tool to approach problems with a geometric flavor: it cannot capture rectifiability (stretchability) of a drawing. Or if it does, that seems accidental,

[10] The region connection calculus allows predicates for how two different regions can relate to each other including equality, disjointness, overlap, and containment. Satisfiability of Boolean formulas over the domain of simple closed regions turns out to be equivalent to the weak realizability problem.

[11] The simultaneous planarity of $G_1, \dots, G_k$ is equivalent to (G, R) being weakly realizable, where $G = \bigcup_{i=1}^{k} G_i$ and $R = \{(e, f) :$ there is no i so that $e, f \in G_i\}$. On the other hand, weak realizability of (G, R) can be modeled by creating a new graph G_i for every pair of edges $(e, f) \notin R$.

as in the case of straight-line planarity which is equivalent to planarity by Fary's theorem. So while Fary's theorem implies that there is a Hanani-Tutte theorem for straight-line planarity, there is no direct connection between the Hanani-Tutte theorem and the existence of a straight-line drawing. Similarly, there is a Fary theorem for level-planarity [24, 56, 23], so the Hanani-Tutte theorem captures straight-line level-planarity, but again, not immediately.[12]

If we modify partially embedded planarity to require $\mathcal{H}$ and the final drawing of G to be straight-line embeddings, the problem becomes **NP**-complete [60].[13] Similarly, simultaneous geometric planarity for two graphs is **NP**-hard [26], while it is conjectured that simultaneous planarity lies in polynomial time (as would be implied by our Conjecture 6.20 for example).

So we will not treat the geometric case here, but it is an interesting open question for which notions of planarity (geometric or not) and under which assumptions analogues of Fary's theorem can be established (Thomassen [72] found a characterization of rectifiable 1-planar graphs). Apparently, even the case of geodesic embeddings on surfaces is open.

The results and conjectures suggest a generic form of the Hanani-Tutte theorem. Let X-planar be any notion of planarity. Say a drawing satisfies X' if it is X-planar with one modification: any requirement that two independent edges do not cross is replaced by the requirement that they cross evenly and adjacent edges are allowed to cross arbitrarily. Since there are different ways of defining notions of planarity, X'-planarity is not necessarily well-defined.

G is X-planar if and only if there is a drawing of G satisfying X'.

This generic Hanani-Tutte scheme can be made to match any of the planarity notions in Table 1. For example, by the classical Hanani-Tutte theorem, G is planar, if and only if there is a drawing of G in which every two independent edges cross evenly.

Example 2.2 (Upward planarity)**.** Let us consider the case of upward planarity. The generic Hanani-Tutte result would then be: G is upward planar if and only if there is a drawing of G in which all edges are x-monotone and directed the same way, and every two independent edges cross evenly. This is true by Theorem 6.8, the Hanani-Tutte theorem for level planarity.

Example 2.3 (Simultaneous Planarity)**.** The generic Hanani-Tutte theorem is not true for all possible planarity notions. For simultaneous planarity it would state that a family of graphs $(G_i)_{i=1}^k$ is simultaneously planar if and only if there is a drawing of $\bigcup_{i=1}^k G_i$ in which any two independent edges that belong to the same graph G_i, for some i, cross evenly. Figure 1 shows a simultaneous drawing of three graphs in which all pairs of edges belonging to the same graph cross evenly, but the graphs are not simultaneously planar. This means that

[12] We'd be amiss not to mention a Fary theorem for c-planarity: c-planarity is equivalent to straight-line rectangular c-planarity [5], so c-planarity would be another case in point if a Hanani-Tutte theorem can be established for it.

[13] Partially embedded planarity first entered the graph drawing literature in its geometric variant [12, Problem 9].

Hanani-Tutte fails for simultaneous planarity of three graphs and thus for weak realizability.

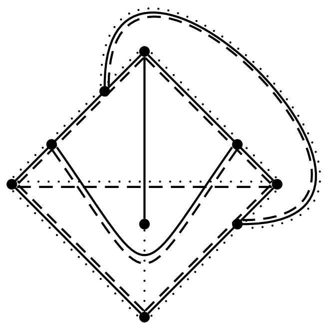

Figure 1: A simultaneous drawing of three graphs—with solid, dashed and dotted edges—that are not simultaneously planar.

The generic Hanani-Tutte theorem also fails for c-planarity [33] and, as mentioned above, for bi-monotonicity [34]. There may still be Hanani-Tutte characterizations for these cases, but they will not fit the generic model described above.

2.2 Relationships among Planarities

There are strong relationships between many of the planarity notions we saw. We borrow a notion from computational complexity to express this: a problem A *reduces* to problem B if there is a polynomial-time computable function f so that $x \in A$ if and only if $f(x) \in B$. If two problems reduce to each other, we call them *equivalent*. This notion of reduction relates the complexity of two problems rather than their inherent structure: for example, c-planarity reduces to book embeddability, since c-planarity lies in **NP** and book embeddability is **NP**-complete. However, we cannot (directly) read off a c-planar embedding of G from a book embedding of $f(G)$. Our reductions will be "natural" in the context of graph drawings, in that an embedding of $f(G)$ will encode an embedding of G—admittedly "natural" is a somewhat vague notion. Figure 2 summarizes known and new reductions.

Question 2.4. We know that SEFE_3 is **NP**-complete [38], so any SEFE problem can be translated into an SEFE_3 problem. Is there a natural construction that achieves this?

Figure 2 suggests that SEFE_2 deserves special attention, since it is a universal problem for many planarity variants, and may turn out to be tractable,

[15] Because of layout restrictions, some planarity variants are missing, including projective planarity, partially constrained PQ-planarity (a special case of ec-planarity with free edges), and the different variants of SEFE_2.

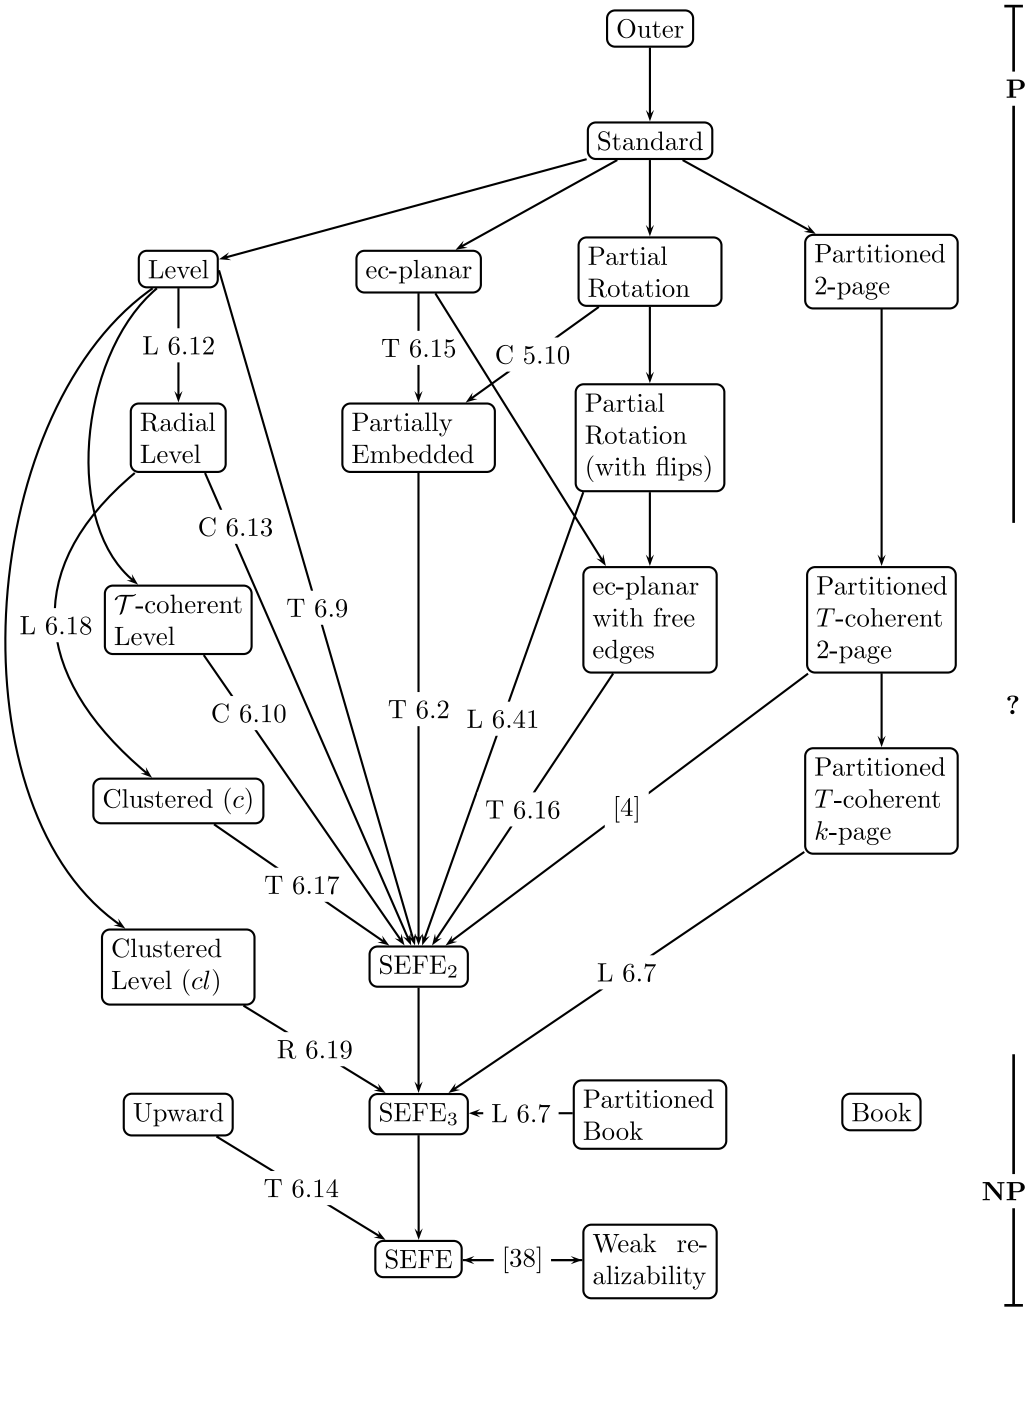

Figure 2: A directed edge denotes a natural reduction between two problems. Reductions for which no reference is given are folklore or straightforward. Abbreviations are (T)heorems, (C)orollaries, (L)emmas and (R)emarks.[15]

as opposed to $SEFE_3$ and weak realizability. As we mentioned earlier, Conjecture 1.2 would be sufficient to show that $SEFE_2$ can be recognized in polynomial time.

Remark 2.5 (Embedding Problems)**.** We concentrate on the recognition problem, mostly ignoring the embedding issue. For most notions of planarity, the embedding problem polynomial-time Turing reduces to the recognition problem: an embedding of a graph G can be constructed by asking a sequence of recognition type questions, typically for subgraphs of G, that allow one to construct the embedding step by step. This leads to an additional factor in the running time, of course.

3 Planarity

This section introduces the Hanani-Tutte theorem in its classical setting: planarity. Most of the material is well-known. Our goal is to illustrate the Hanani-Tutte approach in the simplest, and best-understood, case. Remember that $\mathrm{iocr}(G)$ is the minimum of $\mathrm{iocr}(D)$ over all drawings D of G, where $\mathrm{iocr}(D)$ is the number of pairs of independent edges of G that cross oddly in D. We say a drawing D of G is iocr-0 if $\mathrm{iocr}(D) = 0$. We call an edge in a drawing *(independently) even* if it crosses every other (independent) edge an even number of times.

Theorem 3.1 (Hanani-Tutte)**.** *A graph G is planar if and only if* $\mathrm{iocr}(G) = 0$.

A typical proof of the theorem proceeds as follows: if G is planar, then $\mathrm{iocr}(G) = 0$. For the reverse direction it is sufficient to show that, (i), $\mathrm{iocr}(G) > 0$ for $G = K_{3,3}$ and $G = K_5$, and, (ii), $\mathrm{iocr}(G) > 0$ if $\mathrm{iocr}(H) > 0$ for some minor H of G (we will see proofs of these easy facts in Section 5.2). By Kuratowski's theorem every non-planar graph G contains $K_{3,3}$ or K_5 as a minor, so $\mathrm{iocr}(G) > 0$ by (i) and (ii). Even after completing the details, the use of the obstruction set for planarity leads to a very short, slick proof of the Hanani-Tutte theorem, but this approach has two disadvantages as we look ahead to adapting it to other notions of planarity or embeddability: explicit obstruction sets are rarely known, and they do not typically guide the way to an embedding algorithm.

Remark 3.2 (Weak Hanani-Tutte)**.** There is a weak version of the Hanani-Tutte theorem: if G has a drawing in which all edges are even, then G is planar. Letting $\mathrm{ocr}(D)$ be the number of pairs of edges of G that cross oddly in D, and $\mathrm{ocr}(G)$ be the minimum $\mathrm{ocr}(D)$ for all drawings D of G, we can state the weak version as: G is planar if and only if $\mathrm{ocr}(G) = 0$. Weak Hanani-Tutte characterizations are often easier to prove and sometimes yield stronger conclusions. For planarity, for example, $\mathrm{ocr}(G) = 0$ implies that G has a planar embedding with the same rotation system as the drawing realizing $\mathrm{ocr}(G) = 0$. See [67] for a survey on weak and strong Hanani-Tutte theorems for planarity. We will discuss weak Hanani-Tutte characterizations only rarely in this paper, since they typically cannot be turned into polynomial-time testing algorithms. The strong variants are superior in that respect, as we are about to see in the case of planarity.

The Hanani-Tutte theorem is an algebraic criterion. Start with an arbitrary drawing D of G, and let $i_D(e, f)$ be the number of times that e and f cross in

D. By the Hanani-Tutte theorem planarity of G is equivalent to the existence of a drawing D' in which $i_{D'}(e,f) = 0$ for all pairs of independent edges e and f. If G is planar then such a D' can be obtained (as we will see) from D by a set of (e,v)-moves. An (e,v)-*move* consists of taking a small section of e and deforming it in a narrow tunnel to make it pass over v (while avoiding passing over vertices other than v). See Figure 3 for an illustration. The effect of an (e,v)-move in a drawing D is that the *crossing parity* (the parity of the number of crossings) between e and every edge incident to v changes, while the crossing parity of no other pair of edges is affected.

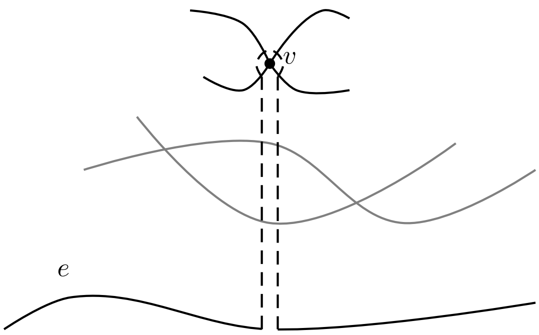

Figure 3: An (e,v)-move. Parity of crossing changes between e and every edge incident to v.

This leads us to the following system P1(D) of equations over GF(2). Create a variable $x_{e,v}$ for every $e \in E(G)$ and $v \in V(G)$. For every pair (e,f) of independent edges in G we require that

$$i_D(e,f) + x_{e,h(f)} + x_{e,t(f)} + x_{f,h(e)} + x_{f,t(e)} = 0 \bmod 2,$$

where $t(g)$ and $h(g)$ denote the two endpoints of edge g (in an arbitrary orientation of G). If we let D' be the drawing obtained from D by making all the (e,v)-moves for which $x_{e,v} = 1$, then $i_{D'}(e,f) = i_D(e,f) + x_{e,h(f)} + x_{e,t(f)} + x_{f,h(e)} + x_{f,t(e)} \bmod 2$ for all pairs of independent edges (e,f).

Lemma 3.3 (Tutte [73], Wu [76]). *Let D be a drawing of G. G is planar if and only if the system P1(D) has a solution over* GF(2).

Proof. Let D be a drawing of G. If G is planar, then there is a planar drawing D' of G. Without loss of generality, we can assume that every vertex $v \in V(G)$ has the same location in D and D'. Hence, only the drawings of edges differ between D and D'. Let D_t, $t \in [0,1]$ be a sequence of drawings changing from D to D' continuously and smoothly. For two independent edges e and f, the value $i_{D_t}(e,f)$ can only change parity if e passes over an endpoint of f or f passes over an endpoint of e. Set $x_{e,v} = 1$ if e passes over v an odd number of times and 0 otherwise, this yields a solution to P1(D).

On the other hand, a solution to P1(D) gives us a way to turn D into an iocr-0 drawing: for every (e, v) for which $x_{e,v} = 1$ in the solution, perform an (e, v)-move. In the resulting drawing D', we will have $i'_D(e, f) = i_D(e, f) + x_{e,h(f)} + x_{e,t(f)} + x_{f,h(e)} + x_{f,t(e)} = 0 \bmod 2$ for every pair of independent edges e and f. Then G is planar by Theorem 3.1. □

Corollary 3.4 (Folklore)**.** *Planarity of a graph can be tested in polynomial time.*

Proof. By Lemma 3.3 planarity of a graph G can be phrased as a linear system of equations over GF(2). Choose an initial drawing D for which $i_D(e, f)$ is easily computable (e.g. place the vertices in convex position). Then solvability of the linear system can be decided in polynomial time. □

Remark 3.5. The algorithm from Corollary 3.4 runs in $O(n^6)$ time for $n = |V(G)| + |E(G)|$, since there are $O(n^2)$ variables and solving the linear system takes cubic time. There are better results on solving linear systems over GF(2), but not significantly better than cubic time, so that trying to improve that part is unlikely to lead to even a quadratic-time algorithm for planarity. However, one may be able to rewrite the linear system to a point where solvability or unsolvability is easily decided. For level planarity this is done in [34], leading to a very simple quadratic algorithm.

What makes the algebraic approach exciting is that many types of constraints that are often placed on planar drawings can be expressed in it. Take, for example, the rotation at a vertex. Suppose we ask whether G is planar with the rotation of one of its vertices fixed. Let w be that vertex, and let D be a drawing in which w has the required rotation. We build a system P2(D) as follows: Create a variable $x_{e,v}$ for every $e \in E(G)$ and $v \in V(G)$. We require that

$$i_D(e, f) + x_{e,h(f)} + x_{e,t(f)} + x_{f,h(e)} + x_{f,t(e)} = 0 \bmod 2$$

for every pair of independent edges (e, f) and every pair of edges (e, f) that share endpoint w. Then G has a planar embedding in which v has the required rotation, if and only if the system P2(D) has a solution over GF(2). We will not argue this special case, but develop a more general result for partial rotations in Sections 5.3 and 6.6.

Remark 3.6 (Obstruction Sets from Hanani-Tutte)**.** The planarity criterion of Hanani-Tutte brings together computational, algebraic and combinatorial aspects of the planarity problem, but that relationship is far from being well-understood. One issue that, as far as we know, has not been investigated, is whether there is any relationship between unsolvability of the system P(D) and obstructions to planarity. This relationship would have to explain obstructions in algebraic terms (that is where the Hanani-Tutte theorem started in the work of Hanani). In other words, the chosen notion of obstruction (minor, subdivision, subgraph) has to be found in the algebraic system. With partially embedded planarity we will see a system where the obstruction set is infinite even for rather strong notions of minor, so if there is a result relating algebraic unsolvability to obstruction sets it will have to account for that.

So far this has only been achieved for planarity and outerplanarity by van der Holst [74]. Since there are proofs of Hanani-Tutte that do not require Kuratowski's theorem [63] this gives an alternative route to Kuratowski's theorem.

4 Redrawing Tools

We establish various redrawing tools that are used in the remainder of the paper. The tools center around redrawing substructures of a graph G, such as edges, cycles, and subgraphs, typically under the assumption that the substructure already has a reasonably good drawing, but we want a better one and so that the overall drawing of G is not affected badly. For example, reasonably good can mean even, in which case better would be crossing-free. Or it could mean independently even, in which case better may mean even. The Hanani-Tutte theorem assumes that *all* edges are independently even, but we need redrawing results that work under weaker, local assumptions, and that allow us to redraw those pieces that are reasonably good. The redrawing results in this section are mostly of this nature: we are given a drawing of G with subgraph H. Edges in H may be independently even, even, free of crossings with H-edges, or entirely free of crossings (each category is better than the previous one). We try to improve the drawing of H in this hierarchy without introducing odd crossings with edges in $G - H$. Redrawing results for simultaneous planarity are covered later, in Section 6.2.3.

The results in this section are not yet systematic and are only a small first step towards establishing a set of tools to tackle Conjecture 6.20 which would imply nearly all the results in this paper. It may be better to skip the proofs in this section in a first reading.

We begin with a basic tool for removing crossings between edges crossing evenly; this idea has also been used in [35, 64] though it may be older than that.

Lemma 4.1. *If an edge e is crossed an even number of times by another edge f, then we can redraw f so as to remove all crossings of e with f, and without affecting the crossing parity of any pair of edges. Crossing-free edges remain crossing-free after the redrawing, but f may consist of more than one component (one of them an arc connecting the endpoints of f; any additional components are closed curves).*

Proof. Imagine e is drawn as a horizontal straight line segment (if e has a self-crossing, it can be redrawn close to the crossing so as to remove the self-crossing), see Figure 4(a). Since e and f cross evenly, we can match up the crossings of f with e in consecutive pairs along e. Cut f at all those crossings, creating two ends for each crossing. Connect paired ends on each side of e, see Figure 4(b). This process does not change the crossing parity of any pair of edges, removes all crossings of f with e, and, since the only crossings introduced in the redrawing are with edges already crossing e, crossing-free edges remain crossing-free after the redrawing. Since f has only two ends, those two ends remain connected by an arc-component, but there may now be additional components of f which have to be closed curves as there are no ends of f remaining. □

Lemma 4.1 can be used to prove the following result, which in turn can be used to prove the strong Hanani-Tutte theorem in the plane without using Kuratowski's theorem..

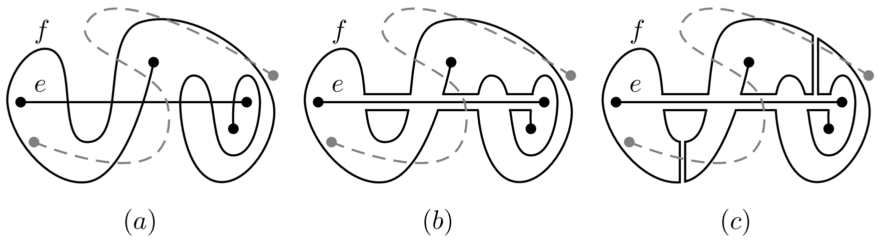

Figure 4: Removing crossings of an edge f crossing e evenly. (a) Initial drawing. (b) After cutting f and reconnecting its ends in consecutive pairs. (c) After reconnecting closed components using narrow tunnels.

Lemma 4.2 (Pelsmajer, Schaefer, Štefankovič [63])**.** *Let E_0 be the set of even edges in a drawing of a graph G. Then there is a drawing of G with the same rotation system in which all edges of E_0 are free of crossings and there are no new pair of edges crossing oddly.*

A proof of this lemma along the lines of Lemma 4.1 can be found in [35], it is shorter and simpler than the original proof from [63]. The proof does require a bit of care: simply dropping all closed components after applying Lemma 4.1 will not work, since this can change the crossing parity between edges; dropping the large closed component of f in Figure 4(b) changes the crossing parity between f and the dashed edge.

Lemma 4.3 is a useful tool for removing crossings with (well-behaved) forests in the graph.

Lemma 4.3. *Suppose H is a subgraph of G all of whose edges are free of crossings, F is a subgraph of G, and E_0 is a set of edges in $E(G) - E(H)$ so that*

(i) $E(F) \subseteq E_0$,

(ii) *every edge in E_0 crosses every edge in F evenly,*

(iii) *any cycle in $H \cup F$ is a cycle in H.*

Then there is a drawing of G in which the edges of F do not cross any of the edges in E_0. In this redrawing, H remains crossing free, the rotation system of G remains the same, and there are no new pairs of edges crossing oddly. In particular, if the drawing of G was iocr-0, it remains so.

Note that (iii) implies that F is a forest, explaining the remark introducing the lemma.

Proof. The goal is to clear edges in F of crossings with edges in E_0 while maintaining several properties of the initial drawing. Let F' be the (possibly empty) subgraph of F containing all edges of F that have no crossings with any edge in E_0. Note that edges in F' cannot separate the face of H they lie in, because of (*iii*) (and the fact that H is free of crossings).

If $F' = F$, we are done. So there must be an edge e in F not yet in F'. Using Lemma 4.1 we can remove all crossings of e with edges of E_0 without changing rotation or crossing parity between any two edges and without creating crossings with crossing-free edges.

Some (or all) edges of E_0 may consist of multiple components now. Reconnect any closed components of an edge to its arc-part by a tunnel, as in Figure 4(c), that avoids edges of $F' \cup \{e\}$ and H. This can be done, since the drawing of $F' \cup \{e\} \subseteq F$ cannot separate the face of H it lies in, by condition (*iii*) and because no two edges in $F' \cup \{e\} \cup H$ cross each other (edges in H remained crossing-free, edges in F' do not cross each other, and we removed crossings of e with edges in E_0 which includes all of F, by (i) and thus F'). We can now add e to F', eventually reaching $F' = F$, which proves the result. □

Lemma 4.4 shows that we can free a subgraph H of G of crossings between edges of H in the right circumstances.

Lemma 4.4. *Let H be a subgraph of G so that every two edges in H cross each other evenly, and every edge of H crosses every independent edge in G evenly. Then there is a drawing of G in which edges of H do not cross each other, the rotation system of H remains the same, and there are no new pairs of independent edges crossing oddly. In particular, if the drawing of G was iocr-0, it remains so.*

Proof. Let e be some edge of H. We want to use Lemma 4.1 to remove crossings with e and then reconnect closed components, but this can lead to problems if H contains cycles as we may be separating closed components from their arc-parts without being able to reconnect them.

We therefore start with a maximal spanning forest F of H. Let $E_0 = E(H)$ and apply Lemma 4.3 to find a drawing of G in which no edge in $E' := E(F) \subseteq E_0$ crosses any edge in E_0 (when applying Lemma 4.3 we choose H to be the empty graph). We will add edges to E' until $E' = E(H)$ while maintaining the property that no edge in E' crosses any edge in H. Once we have reached $E' = E(H)$ the lemma is proved as long as we make sure not to create a new pair of independent edges crossing oddly.

So assume there is an edge $e \in E(H) - E'$ that still has crossings with some edge f in H. By assumption every such edge $f \in E(H)$ crosses e evenly. Let H' be H restricted to $E' \cup \{e\}$. Then H', by itself, is embedded, since edges in E' do not cross e. Pick a cycle C in H', so that e lies on C, and the edges of C bound a face of the embedding of H' (e is not a cut-edge of $E' \cup \{e\}$, since F is a maximal spanning forest of H).[16]

The cycle C (as a curve) is free of self-crossings, since it is part of H', so we can speak of the inside and outside region of C. If there is an edge $g \in E(G)$

[16] C is not necessarily a facial cycle in the strict sense: it may not be equal to the facial walk bounding the same face, since H' need not be 2-connected. However, the edges of C are contained in such a facial walk of H'.

that crosses any edge of C oddly, then g must have an endpoint on C (incident to the edge or edges of C it crosses oddly). We can then move the end of g in the rotation at that endpoint so that g is even with respect to all edges of C. Repeating this for all such edges g, we can make all edges in $E(C)$ even. We then apply Lemma 4.1 for every edge in $E(C)$ and every edge crossing it, removing all crossings with C. If any edge consists of multiple components now, reconnect the closed components to the arc-part of that edge as long as that is possible without introducing crossings with $E' \cup \{e\}$. Let $E' := E' \cup \{e\}$. Note that no edge in E' crosses any edge in H, and the crossing parity of any pair of edges has not changed (in particular, any two H-edges still cross evenly). However, some edges now consist of multiple components: edges in $E(G)-E(H)$ that crossed C and edges of H that crossed e. We want to argue that dropping all closed components simultaneously does not lead to two edges crossing oddly that currently cross evenly. Suppose there are two edges $f, g \in E(G)$ that cross evenly, but if we drop all their closed components, they cross oddly. But then the arc-parts of f and g must cross oddly and, in particular, lie in the same face σ of H'. Any closed components of f and g that could not be reconnected to their arc-parts cannot lie in σ (otherwise they would have been reconnected), and, in particular, cannot intersect f or g. Since any two closed components cross each other evenly, dropping all of them cannot change the crossing parity of f and g. Hence, the new drawing does not contain any new pairs of independent edges crossing oddly. Since e now belongs to E' we have made the required progress. □

Lemma 4.4 is a useful tool, but we would really like to be able to prove the stronger result we stated as Conjecture 1.2 earlier:

> Suppose a graph G with subgraph H can be drawn so that every edge of H crosses every independent edge of G evenly. Then there is a drawing of G in which edges of H do not cross each other, and there are no new pairs of independent edges crossing oddly.

Because of Lemma 4.4 it would be sufficient to find a drawing in which every two edges of H cross each other evenly (and there are no new pairs of independent edges). There is a partial result along the lines of Conjecture 1.2 in [65, Lemma 2.3], but it avoids dealing with cut-vertices and cut-edges of H, which is where the problem gets difficult. In the absence of cut-vertices and cut-edges, we can settle the conjecture.

Lemma 4.5. *Let H be a subgraph of G whose connected components are all 2-connected. If G can be drawn so that every edge of H crosses every independent edge in G evenly, then there is a drawing of G in which edges of H are free of crossings, and there are no new pairs of independent edges crossing oddly.*

The induction in the following proof uses ideas from [63].

Proof. Fix a drawing of G in which every edge of H crosses every independent edge in G evenly. We need to show that there is a drawing of G in which all edges of H are free of crossings. To that end, we induct over $w(G,H) := \sum_{v \in V(H)} d(v)^3$, where $d(v)$ is the degree of v in G. For two graphs with the same weight, we induct over the number of *unprocessed* edges in H. We only label an H-edge *processed* if it is free of crossings, and we guarantee that it remains so. Initially, we call all edges in $E(H)$ *unprocessed*. We write $(G', H') < (G, H)$ if $w(G', H') < w(G, H)$ or $w(G', H') = w(G, H)$ and the number of unprocessed edges in H' is smaller than in H. Our goal is to prove the following claim, which is sufficient to establish the lemma inductively:

> If the statement of the lemma is true for every $(G', H') < (G, H)$, then it is true for (G, H) as well.

It is hard to remove crossings from edges in H as long as H has vertices of H-degree 4 or more. This problem is hard in general (we do not know how to do it), but in this particular case we can deal with it, since H is 2-connected: we clear cycles of H of crossings, which allows us to turn H into a subcubic[17] graph.

Suppose H contains a cycle C and a vertex v belonging to C which has degree at least 4 in H. Modifying the rotations at vertices of C, we can make edges of C even (edges of C can only be crossed oddly by adjacent edges, by assumption on H, hence this can be done by adjusting the rotations of vertices belonging to C). Applying Lemma 4.2 gives us a drawing of G in which C is free of crossings; note that the redrawing procedure does not introduce pairs of independent edges crossing oddly. Since v was incident to at least four edges of H, v is either incident to at least one H-edge outside and at least one H-edge inside of C, or v is incident to at least two H-edges on the same side of C (both inside or both outside). In the first case, we split v into v and v' lengthening C so that v remains incident to the inside edges, and v' is incident to the outside edges. In the second case, we split v into v and v' by pushing v' off of C on the side that the incident H-edges are, and add an edge vv' to H. In both cases, the component containing C remains 2-connected. Consider (G', H'). All components of H' are 2-connected, and no odd crossings between pairs of independent edges were introduced. The overall weight decreases: let $a \leq b$ be the number of edges v be incident to on the two sides of C before the split. We know that $a + b \geq 2$, and v contributes $(a + b + 2)^3$ to the overall weight. Let us consider v and v' after the split. In the first case, they contribute $(a+1)^3 + (b+1)^3 < (a+b+2)^3$ for $a, b \geq 1$. In the second case, they contribute $3^3 + (b + 1)^3 < (b + 2)^3 = (a + b + 2)^3$ for $b \geq 2$ ($a = 0$ in this case). In both cases $w(G', H') < w(G, H)$. Hence, by induction, we can assume that there is a a drawing of G' in which H' is free of crossings, and there are no new pairs of independent edges crossing oddly. Since H' includes the new edge vv', we can contract vv' obtaining a drawing G in which H is free of crossings, and in which there are no new pairs of independent edges crossing oddly.

We can therefore assume that every vertex of H has degree at most 3 in H. Suppose there is an edge in H that is crossed oddly by some other edge; this

[17] A graph is *subcubic* if all its vertices have degree at most 3.

edge is part of some cycle C in H, since H is 2-connected. Let f be an arbitrary edge on C that is crossed oddly by some edge (f has to be unprocessed). If f is not incident to a processed edge, we can make f even by modifying the rotation at its endpoints without introducing crossings with processed edges. Otherwise, f is incident to a processed edge, say $f = uv$ and v is incident to a processed edge. Then v lies on a previously processed cycle C'. So v has H-degree exactly 3 (since f is unprocessed), and f does not cross either of the other two H-edges incident to v (since they both belong to the crossing-free cycle C'). Now f can be made even with respect to edges incident to v by moving their ends in the rotation at v. Since C' is crossing free, any edge crossing f must be on the same side of C' as f, so we do not need to introduce crossings along already processed edges. Repeating this process, we obtain a drawing in which all edges of C are even, and all processed edges have remained crossing free. Applying Lemma 4.2 gives us a drawing of G in which C is free of crossings, all previously processed edges remain crossing-free (since they are even), and there are no new pairs of independent edges crossing oddly. We now call all edges on C processed. Since we know that C contained at least one unprocessed edge, this increases the number of processed edges by at least 1. □

The remaining results in this section start with the assumption that a subgraph H of G is already embedded free of self-crossings, and show how to remove additional crossings. Call a crossing between two edges *independent* if the two crossing edges are independent (i.e. not adjacent).

Lemma 4.6. *Suppose we are given an iocr-*0 *drawing of a graph G containing a planar embedding $\mathcal{H}$ of a subgraph $H \subseteq G$. Let C be a cycle in H. Then we can find an iocr-*0 *drawing of G containing $\mathcal{H}$ in which all edges of C are free of crossings. If C is contained in a facial walk of $\mathcal{H}$, then we can assume that the iocr-*0 *redrawing does not contain any new independent crossings with edges in H.*

Proof. We first make all edges of C even. If some edge f crosses an edge of C oddly, then f must be an edge in $E(G)-E(H)$ incident to a vertex in C. We can then move the end of f at that vertex, so f crosses both edges of C it is incident to evenly. We apply Lemma 4.1 for all edges $f \in E(G)$ and $e \in E(C)$ for which f crosses e (necessarily, $f \notin E(H)$, so this is possible). At this point, edges of C are free of crossings. Closed components of an edge that lie on the same side of C as their arc-components can be reconnected to their arc-components without changing the crossing parity of any two edges or introducing crossings with C. All remaining closed components are dropped. This does not change the crossing parity of any two edges: suppose e and f cross oddly now, whereas they used to cross evenly before we dropped the closed components. Since C is free of crossings, both e and f must lie on the same side of C. The closed components of e and f we dropped must have been on the other side of C, but any two closed components cross evenly, so dropping them cannot have changed the crossing parity of e and f.

If C is contained in a facial walk of $\mathcal{H}$, then we can redraw more carefully so that there are no new independent crossings with edges in H: First, note that the applications of Lemma 4.1 only lead to local redrawings along edges of C, which consists of H-edges and thus cannot be crossed by other H-edges, so no new crossings with edges in H are created. We do introduce new crossings when making edges of C even, but those crossings are between adjacent edges, since we are moving ends of edges. We are thus left with reconnecting the closed components. But this can be done by routing the connecting tunnels of the components close to the facial walk that contains C; that introduces no crossings with H. □

The following corollary shows that one can get rid of independent crossings with edges of H entirely.

Corollary 4.7. *Suppose we are given an iocr-0 drawing of a graph G containing a planar embedding $\mathcal{H}$ of a subgraph $H \subseteq G$. Then we can find an iocr-0 drawing of G containing $\mathcal{H}$ in which edges of H are not involved in independent crossings.*

Proof. Apply Lemma 4.6 for every (simple) cycle contained in a facial walk. If an H-edge e crosses an independent edge f in the resulting drawing, then e must be a cut-edge of H, so both sides of e lie on a facial walk of $\mathcal{H}$. We can then apply Lemma 4.1 to remove crossings of f with e and then, just as in the proof of Lemma 4.6 reconnect components of f by routing the tunnels close to the facial walk. □

The next lemma allows us to clear a single even edge of crossings; we only need and state the result for edges in $E(G) - E(H)$, but using ideas from the proof of Lemma 4.6 it is easily seen to be true for edges in $E(H)$ as well.

Lemma 4.8. *Suppose we are given an iocr-0 drawing of a graph G containing a planar embedding $\mathcal{H}$ of a subgraph $H \subseteq G$. Let e be an even edge in $E(G) - E(H)$. Then we can find an iocr-0 drawing of G containing $\mathcal{H}$ in which e is free of crossings.*

Proof. Apply Corollary 4.7 so that the only crossings with H are between adjacent edges. We need to argue that this can be done so that e remains even. The only problem occurs when in an application of Lemma 4.6 the end of an edge f at a vertex u is moved, where u is an endpoint of e. This situation occurs when we have to make f even with respect to a cycle C containing u; Figure 5 illustrates all possible cases: if f and e are on opposite sides of C, then f crosses either one, (a), or both, (b), edges of C it is adjacent to oddly; if f and e are on opposite sides, then f either crosses both C-edges oddly, (e), or only one of them, in which case we distinguish whether that edge is on the same side of e as f, as in (c), or on the far side, (d). If the end of f is moved past e, this may result in e and f crossing oddly. If we can, we move the end of f in the other direction—cases (a), (c), and (e) in Figure 5. There are two cases in which we cannot: if the end of f is on the other side of C from e and f crosses both C-edges incident to u oddly, case (b), and if f is on the same side of C as e, but crosses the far C-edge oddly, while crossing the close C-edge evenly, case (d). In both cases we allow the end of f to move over e; the point is that in both cases, the end of f ends up on the opposite side of C from e, so after making C free of crossings, e and f no longer cross at all, so it does not matter that we temporarily introduced an odd crossing between them.

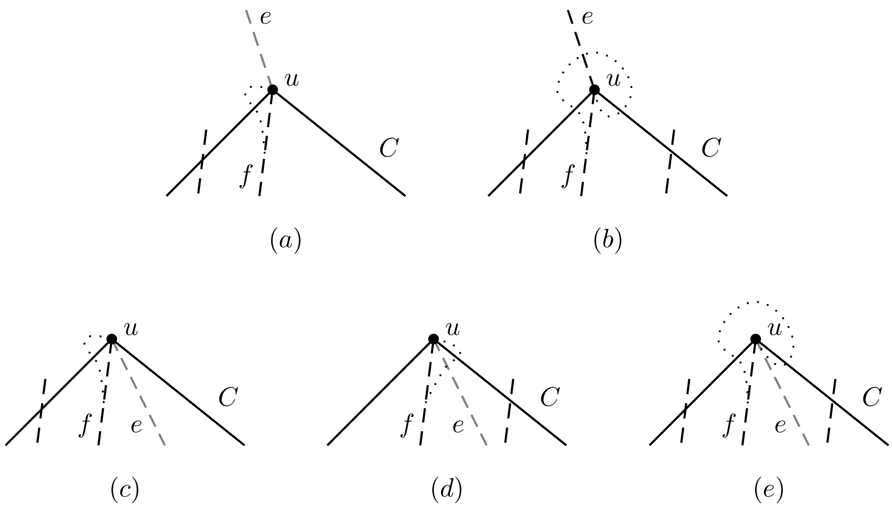

Figure 5: Redrawing at u. Edge f is black and dashed, e gray and dashed; changes to f are drawn as dotted lines.

So e is still even, and it no longer crosses any edges of H except for edges it is incident to. We just argued that Lemma 4.6 can be applied so as to keep e even. Apply the lemma for all cycles in H that contain an endpoint of e. Since e is (and remained) even, its ends never have to be moved in the redrawing, so we do not create new crossings of e with edges in H, so e only crosses cut-edges of H it is incident to. Those crossings we can remove as in the proof of Corollary 4.7: we cut e and reconnect it following a facial walk of H containing both sides of the cut-edge.

Any remaining edges crossing e belong to $E(G) - E(H)$, and we can apply Lemma 4.1 to remove their crossings with e (reconnecting their closed components arbitrarily, just avoiding e itself). At this point e is free of crossings. □

Removing crossings is much easier if we do not have to worry about edges in $E(H)$ as the following lemma shows. Lemma 4.9 and Lemma 4.8 will be the main tools in establishing the Hanani-Tutte theorem for partially embedded planarity.

Lemma 4.9. *Suppose we are given an iocr-*0 *drawing of a graph G containing a planar embedding $\mathcal{H}$ of a subgraph $H \subseteq G$. Let E_0 be the set of even edges of*

G and $E(H) \subseteq E_0$. Then we can find an iocr-0 drawing of G containing $\mathcal{H}$ in which all edges of E_0 are free of crossings.

Remark 4.10. Lemma 4.9 extends Lemma 4.2 from planarity to partially embedded planarity, but it needs the stronger assumption that the initial drawing is iocr-0. The lemma will not remove crossings with even edges in the presence of pairs of independent edges crossing oddly. However, this version suffices for our purposes; we suspect that a stronger version of Lemma 4.9 analogous to Lemma 4.2 can be proved. We would also like to drop the condition $E(H) \subseteq E_0$, but Lemma 4.8 shows that even clearing a single even edge is non-trivial if we do not assume that edges of H are even.[18]

Proof of Lemma 4.9. We want to use Lemma 4.1 to remove crossings of an edge f with an edge e. There is a problem, however, if $f \in E(H)$ since we may not redraw edges of H. Instead we proceed as follows: let $e \in E(H)$ and $f \in E(G)$ so that f crosses e and crosses it evenly (in particular $f \notin E(H)$ since two edges in H cannot cross).

Apply Lemma 4.1 to remove crossings of f with e (note that connecting the severed ends of f does not introduce crossings with edges in H, since the connections are routed along $e \in E(H)$, and e does not cross edges in H). Repeat this for all such pairs $e \in E(H)$ and $f \in E(G)$ for which f crosses e evenly (and at least once). Note that we can apply Lemma 4.1 even if f already consists of multiple components due to earlier cuts. Reattach closed components to their corresponding arcs via thin tunnels as long as this can be done without crossing any edges in H. As always, the crossing parity between no two edges changes, since tunnels introduce two crossings every time they pass through an edge. This may still leave closed components that are separated from their arcs by a cycle in H. In other words, some edge f has an arc component that lies (without loss of generality) within a region bounded by a cycle C for which $E(C) \subseteq E(H)$ and there is a closed curve belonging to f that lies outside the cycle C. If the arc-component of f crosses some arc-component of some edge g oddly, then the arc-component of g lies within the region bounded by C since C consists of H-edges and thus is free of crossings. But then a closed curve belonging to f that lies outside C can only cross g by crossing a closed component of g. Since any two closed components cross evenly, we can drop all remaining closed components without affecting the crossing parity between any two edges. Edges in $E(H)$ are free of crossings.

We now repeat a similar process for edges $e \in E_0$ and $f \in E(G)$ to remove crossings with edges of E_0 one edge at a time. Let $E' \subseteq E_0$ be those edges in E_0 that are already free of crossings. We already established that $E(H) \subseteq E'$. If $E' = E_0$ we are done; otherwise there is some $e \in E_0 - E'$. Note that $e \notin E(H)$. Now e crosses some edges in $E(G)$. However, e cannot cross any edges in $E(H)$ since edges in H are free of crossings. So let f be an arbitrary edge crossing e, then $f \in E(G) - E(H)$. Perform the cutting move from Lemma 4.1 to remove crossings of f with e (note that connecting the severed ends of f cannot lead to

[18] We should clarify: we know that these stronger versions of Lemma 4.9 are true, since they are directly implied by Theorem 5.6, but we are looking for direct redrawing proofs of these results, not requiring the obstruction set machinery.

crossings with an edge in E' since we reroute close to e and edges in E' cannot cross e). We repeat this for all edges $f \in E(G)$ which cross e. We reattach closed components to their corresponding arcs as long as we can do so without crossing edges in E'.

We want to argue that if at this point we drop all closed curve components, the remaining drawing is iocr-0. If this were not the case, there would have to be two independent edges f and g whose arc components cross oddly after dropping all closed components. However, since the drawing was iocr-0 to start with, f and g (as independent edges) crossed evenly before dropping the closed components. Since closed components cross each other evenly, this means that some closed component of f crossed the arc-part of g or vice versa. Let us assume (without loss of generality) that f had a closed component that crossed the arc-part of g. But we know that the arc-part of f also crosses the arc-part of g. Since g crosses none of the edges in E' we could have connected the arc-part of f to its closed component by routing a tunnel close to g. Hence, we can drop all the closed curve components to get an iocr-0 drawing of G without closed components. Moreover, in this drawing e is now free of crossings. This proves the claim inductively. □

Lemma 4.9 is sufficient to prove a weak Hanani-Tutte theorem for partially embedded planarity. The goal of Section 5.2 is the strong version of this lemma, Theorem 5.6. Call a drawing ocr-0 if every two edges of the drawing cross evenly.

Corollary 4.11. *If G has an ocr-0 drawing containing a planar embedding $\mathcal{H}$ of a subgraph $H \subseteq G$, then there is an embedding of G extending $\mathcal{H}$.*

We conclude this section with a nearly trivial result that will nevertheless turn out to be quite useful. It is easy to redraw edges incident to a vertex of degree 3 so that they cross each other evenly: suppose two such edges cross oddly. Then their ends at the vertex must be consecutive, so we can move one of the ends beyond the other in the rotation, changing their crossing parity (and leaving the third edge unaffected). The next lemma shows that this remains true if part of the graph is (and has to remain) embedded.

Lemma 4.12. *Suppose we are given a graph G and a planar embedding $\mathcal{H}$ of a subgraph $H \subseteq G$. Let v be a vertex of degree 3 in G. Then we can redraw edges of $E(G) - E(H)$ locally at v so that all edges incident to v cross each other evenly and so that the crossing parity of all other pairs of edges remains the same.*

Proof. If all edges incident to v belong to H, then there is nothing to show; otherwise, we can moves the ends of $E(G) - E(H)$-edges at v so that all edges incident to v cross evenly. □

5 Partially Embedded Planarity

A *partially embedded graph (PEG)* is a triple $(G, H, \mathcal{H})$ consisting of a graph G, a subgraph H of G, and an embedding $\mathcal{H}$ of H in the plane. We consider two embeddings *topologically equivalent* in the plane if there is an orientation-preserving homeomorphism of the plane that takes one to the other.

Remark 5.1 (Rotation Systems and Embeddings)**.** The definition of a partially embedded graph requires an embedding $\mathcal{H}$ rather than just a rotation system of H. The reason is that we do not require H to be connected, so two embeddings of H may have the same rotation system without being topologically equivalent: a rotation system does not restrict which face of a connected component another component lies in. If two embeddings of a connected graph H have the same rotation system, they are equivalent in the sense that there is an orientation preserving homeomorphism of the sphere that takes one to the other; on the plane, the two embeddings may still look different, since they can differ in which face is the outer face. If H consists of a single connected component that is acceptable, since one can correct for it with a single homeomorphism of the sphere, but if H consists of multiple connected components, each component may require a different homeomorphism, e.g. if the outer face of one component remains the same, but changes for another. In short, we have to be careful if we work with rotation systems, in particular if the graph is not connected. Section 5.3 discusses a weakening of partially embedded planarity that captures only the rotation system. Hoffman and Richter [42] worked out the details of a combinatorial description of an embedding of a disconnected graph on a surface.

A PEG $(G, H, \mathcal{H})$ is *planar* if there is a planar embedding of G that contains $\mathcal{H}$. If $\mathcal{H}$ and $\mathcal{H}'$ are two topologically equivalent embeddings of H, then $(G, H, \mathcal{H})$ is planar if and only if $(G, H, \mathcal{H}')$ is planar, so we can redraw $\mathcal{H}$ as long as we maintain topological equivalence.

Partially embedded planarity can be tested in linear time using SPQR-trees, a result due to Angelini, Di Battista, Frati, Jelínek, Kratochvíl, Patrignani, and Rutter [3]. If $\mathcal{H}$ is a straight-line embedding, and we require G to have a straight-line embedding extending $\mathcal{H}$, then the problem is **NP**-complete [60].

5.1 Minors and Obstructions

Jelínek, Kratochvíl, and Rutter [46] give a forbidden obstruction characterization of partially embedded graphs. Since we are dealing with a partially embedded graph $(G, H, \mathcal{H})$, the usual minor operations need some modification; for example, we cannot arbitrarily contract edges in $E(G) - E(H)$ since the effect on $\mathcal{H}$ can be ambiguous, and there are new operations available to us: if we delete an edge in H, do we delete it in G as well? The following operations are considered in [46].

(ia) Edge Removal. Removing an edge $e \in E(G)$.

(ib) Vertex Removal. Removing a vertex $v \in V(G)$ and all edges incident to it.

(iia) Edge relaxation. *Relaxing an edge* $e \in E(H)$, that is, removing e from $E(H)$ but keeping it in $E(G)$.

(iib) Vertex Relaxation. *Relaxing a vertex* $v \in V(H)$, that is, removing v and edges incident to it from $E(H)$, but keeping them in $E(G)$.

(iii) H**-Edge Contraction** Contracting an edge in H. This requires modifying $\mathcal{H}$, but the result is unique (up to topological equivalence), since the embedding of H is given.

(iva) Simple G**-Edge Contraction.** Contracting an edge in $E(G) - E(H)$ which has at least one endpoint not in $V(H)$.

(ivb) Complicated G**-Edge Contraction.** Contracting an edge in $E(G) - E(H)$ for which (i) the endpoints belong to different components of H, but (ii) there is a face of $\mathcal{H}$ so that both endpoints are incident to that face, and (iii) both endpoints are incident to exactly one H-edge. This requires modifying $\mathcal{H}$, but this can be done, by (ii), and uniquely so since the contraction does not create a cycle in H, by (i), and there is a unique way of joining the rotations of H-edges at the endpoints of the contracted edge, by (iii).

(ivc) Stronger G**-Edge Contraction.** Contracting an edge $E(G) - E(H)$ for which (i) there is a unique face of $\mathcal{H}$ that is incident to both endpoints, and (ii) both endpoints occur exactly once on a facial walk of that face. This requires modifying $\mathcal{H}$, with a unique result, since the edge is uniquely located in the rotation system of H at its endpoints.

Operations in group (i) are traditional, but operations in groups (ii)-(iv) are new. Jelínek, Kratochvíl, Rutter observe that (ivc) implies (ivb).

Definition 5.2 (Jelínek, Kratochvíl, Rutter [46])**.** We say that $(G, H, \mathcal{H})$ is a *PEG-minor* of $(G', H', \mathcal{H})$, and write $(G, H, \mathcal{H}) \preceq (G', H', \mathcal{H})$, if $(G, H, \mathcal{H})$ can be obtained from $(G', H', \mathcal{H})$ by a sequence of operations in groups (i)-(iv).

Remark 5.5 explains why this definition differs slightly from the original definition given by Jelínek, Kratochvíl, and Rutter in [46]. The obstructions for partially embedded planarity with respect to this notion of minor are pictured in Figure 6. We need to explain the infinite family of obstructions A_k called *k-fold alternating chains*.

Definition 5.3 (Jelínek, Kratochvíl, Rutter [46])**.** A PEG $(G, H, \mathcal{H})$ is a *k-fold alternating chain*, an A_k, $k \geq 3$, if H consists of a cycle C on $k+1$ vertices and two vertices u and v that lie on opposite sides of the cycle if k is odd and on the same side otherwise. There are also k paths $P_1, \ldots, P_k$ with endpoints on C so that P_1 has length 2 and passes through u and P_k has length 2 and passes through v, and all other paths are single edges (length one). The endpoints of two consecutive paths alternate on C. Finally, all vertices on C have degree 4 except two which are endpoints of P_2 and P_{k-1} which have degree 3.

Note that A_k does not denote a single graph, there are multiple k-fold alternating chains for $k \geq 5$.

Theorem 5.4 (Jelínek, Kratochvíl, Rutter [46])**.** *A PEG-graph* $(G, H, \mathcal{H})$ *is planar if and only if it does not contain any of the obstructions* K_5*,* $K_{3,3}$*, or* D_1*,* D_2*,* D_3*,* D_4*,* D_{11}*,* D_{14}*,* D_{16}*,* D_{17} *or* A_k*,* $k \geq 3$ *as a PEG-minor.*

Remark 5.5. Jelínek, Kratochvíl, and Rutter [46] do not allow operation (*ivc*) in their definition of PEG-minor, but discuss it as an option, and show that allowing (*ivc*) reduces the set of obstructions to the ones listed in Theorem 5.4. Our goal, different from theirs, is to make the set of obstructions as small as possible, so we use the more powerful notion of PEG-minor that allows (*ivc*). We keep the original numbering of the obstructions, which explains the gaps in the sequence of D_is. Jelínek, Kratochvíl, and Rutter introduce an additional operation that reduces the infinite set of alternating chain obstructions to a finite set. We avoid that operation, since it is somewhat awkward, and it is rather easy to deal with the infinite family of A_ks directly.

5.2 Hanani-Tutte for Partially Embedded Planarity

Recall that a drawing of a graph is iocr-0 if every two independent edges cross each other an even number of times.

Theorem 5.6. *Suppose we are given a graph G and a planar embedding $\mathcal{H}$ of a subgraph $H \subseteq G$. Then G has a planar embedding that extends $\mathcal{H}$ if and only if there is an iocr-0 drawing of G containing $\mathcal{H}$.*

We hope to give a direct proof of Theorem 5.6 at some point, however, in the current paper we base the proof of Theorem 5.6 on the set of obstructions for partially embedded graphs. To do so, we need two results.

Lemma 5.7. *Suppose $(G, H, \mathcal{H})$ is a PEG-minor of $(G', H', \mathcal{H}')$ and there is an iocr-0 drawing of G' extending $\mathcal{H}'$. Then there is an iocr-0 drawing of G extending $\mathcal{H}$.*

Proof. Fix an iocr-0 drawing of $(G', H', \mathcal{H}')$. We need to verify that all the minor operations in groups (*i*)–(*v*) can be performed without making a pair of independent edges cross oddly. This is clear for operations in groups (*i*) and (*ii*).

For (*iii*) and (*iv*) let uv be the edge being contracted. We deal with (*iii*) first, so $uv \in E(H)$. Since H-edges do not cross each other at all, only edges in $E(G) - E(H)$ can cross uv. Any such edge crossing uv will do so evenly, unless it is incident to u or v, let us say edge $f \in E(G) - E(H)$ incident to u. Performing an (f, u)-move ensures that f crosses e evenly. Repeating this for all such edges f turns uv into an even edge. We contract uv by moving v along uv towards u and identifying u and v. Since uv is crossed evenly by all edges that cross it, this does not change the crossing parity between any pair of edges. Merging the rotations of u and v when identifying them leads to a unique rotation of H-edges at $u = v$, since $uv \in E(H)$. Finally, we did not introduce crossings between edges in H, so the induced drawing of H is still an embedding. This explains (*iii*).

For operations in group (*iv*) we proceed similarly. We first consider (*iva*). Let u be the endpoint not belonging to $V(H)$. Then u is not incident to any edges belonging to H. By changing the rotation at u we can make uv even with respect to all edges incident to u. Contract uv by moving u along uv to v and

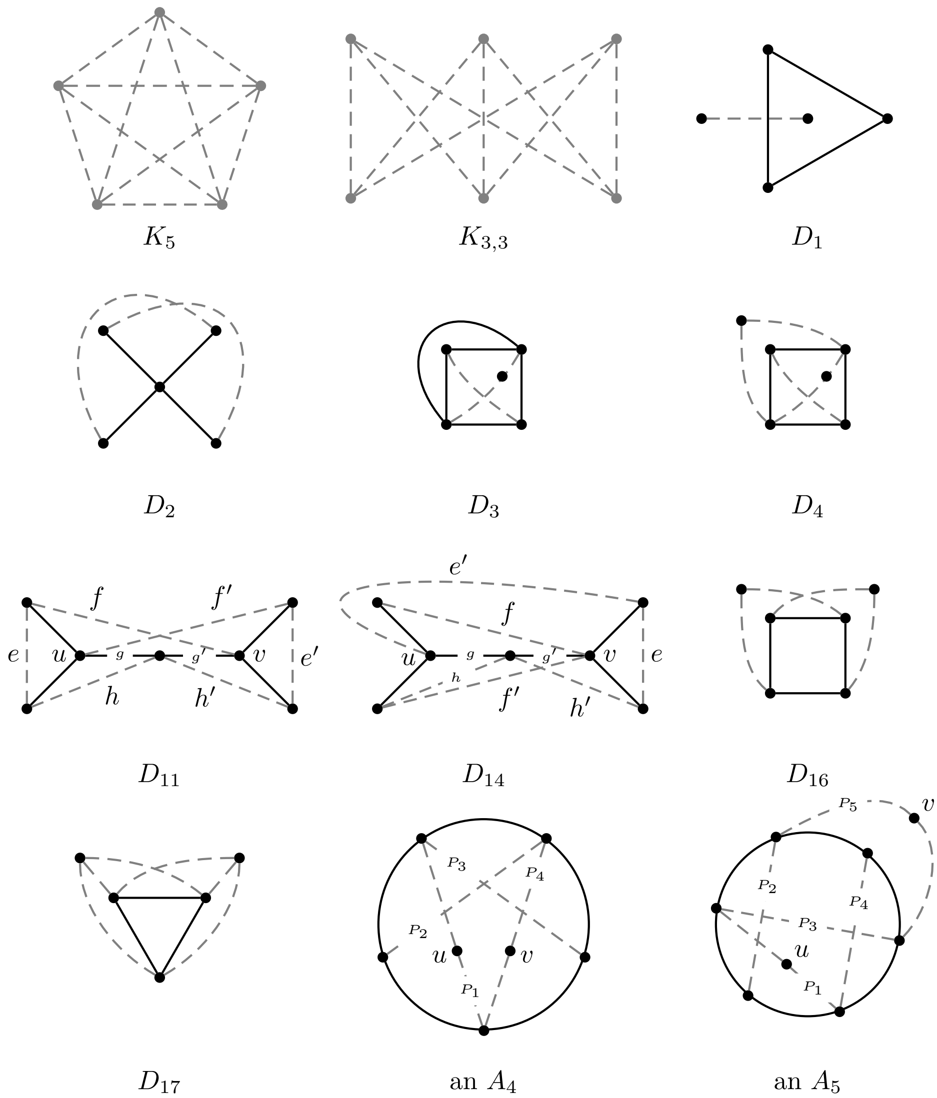

Figure 6: Illustrations of obstructions to partially embedded planarity: K_5, $K_{3,3}$, or D_1, D_2, D_3, D_4, D_{11}, D_{14}, D_{16}, D_{17}. Vertices and edges of H are black, vertices and edges in $G - H$ are gray. For space reasons, some labels are placed *on* the edges (or paths, in the case of the A_k) they label. Based on [46, Figures 1 and 2]

identifying u and v. The only way the contraction can create a new pair of edges that cross oddly, is if one of the edges, say f, crossed uv oddly, while the other edge, say g, is incident to u, so it picks up an odd number of crossings with f. However, f must have u or v as an endpoint if it crossed uv oddly, so it is now adjacent to g, so the contraction did not create a pair of *independent* edges crossing oddly, which is all that matters. Finally the contraction does not affect $\mathcal{H}$ at all, so (*iva*) is fine.

We can ignore (*ivb*) since it is subsumed by (*ivc*), so let us consider that operation. Suppose an endpoint of uv has H-degree larger than 1, say $x \in \{u, v\}$. Then there is a cycle C in H which contains x so that the two C-edges incident to x belong to the unique face of $\mathcal{H}$ that contains both u and v (which exists by assumption). Make C free of crossings using Lemma 4.6. Then uv does not cross any H-edges incident to x: not the two H-edges belonging to C since those are free of crossings, and not any other H-edges incident to x, since those have to lie on the other side of C (remember that x occurs only once in the boundary walk). We can then make uv even with respect to all edges in $E(G) - E(H)$ incident to x by moving their ends in the rotation at x. If, on the other hand, uv has H-degree 1 at $x \in \{u, v\}$ we can perform an (uv, x)-move, if necessary, so that uv is even with respect to the H-edge incident to x. We then modify the rotation of the remaining edges incident to x so they cross uv evenly (this is easy, since they belong to $E(G) - E(H)$). In all combinations of cases, uv is now an even edge. Then Lemma 4.8 allows us to clear uv of all crossings, and we can contract it, identifying u and v, without changing the number of crossings between any pair of edges. Note that all the redrawing could be done keeping $\mathcal{H}$ embedded. □

Lemma 5.8. *There is no iocr-*0 *drawing of any of the obstructions* K_5*,* $K_{3,3}$*, or* D_1*,* D_2*,* D_3*,* D_4*,* D_{11}*,* D_{14}*,* D_{16}*,* D_{17} *and* A_k*,* $k \geq 3$*.*

Proof. K_5 and $K_{3,3}$ do not have iocr-0 drawings. This is well-known (as part of the traditional proofs of the Hanani-Tutte theorem), but we include a short argument: Suppose there were an iocr-0 drawing of K_5 or $K_{3,3}$. Take a Hamiltonian cycle C of either K_5 or $K_{3,3}$, and make all its edges even (by moving ends of edges at the vertices of the cycle). By Lemma 4.6 there is an iocr-0 drawing in which C is free of crossings (H is the empty graph). In the case of $K_{3,3}$ this means that the drawing is ocr-0, since any pair of adjacent edges (the only pairs that can cross oddly in an iocr-0 drawing) includes at least one crossing-free edge of C, so by Corollary 4.11, $K_{3,3}$ is planar, which is a contradiction. For K_5 any pair of adjacent edges that crosses oddly must lie on the same side of C and the ends of the two edges are thus consecutive at their common end. Hence, we can turn the iocr-0 drawing of K_5 into an ocr-0 drawing, which, by Corollary 4.11 yields a planar embedding of K_5, again a contradiction.

A similar argument works for the A_k, $k \geq 3$. Consider an iocr-0 drawing of an A_k graph $(G, H, \mathcal{H})$. Let C be the cycle in H. By Lemma 4.6 we can assume that C is free of crossings. Since the drawing is iocr-0 only adjacent edges can cross each other oddly. As C is free of crossings, such a crossing has to occur between two edges sharing an endpoint on C or between two edges sharing an endpoint at the two isolated vertices of H. The isolated vertices of H have degree 2, so we can assume that the edges incident to them cross evenly. Now two edges incident to the same vertex v of C must lie on the same side of C to cross. Those two edges are the only edges incident to v, inside or outside of C, so we can swap their ends at v to make them cross evenly. We obtain an ocr-0 drawing of A_k, which gives us a planar drawing of A_k (via Corollary 4.11) contradicting Theorem 5.4.

Suppose D_1, D_2, or D_{16} has an iocr-0 drawing. Since all these graphs have max-degree 3, we can use Lemma 4.12 to get an ocr-0 drawing of these graphs. But then, by Corollary 4.11 these graphs are planar, which contradicts Theorem 5.4.

Let $(G, H, \mathcal{H})$ be D_3, D_4, and let C be the 4-cycle contained in H, or let $(G, H, \mathcal{H})$ be D_{17} and C the 3-cycle contained in H. Suppose one of these $(G, H, \mathcal{H})$ has an iocr-0 drawing. Use Lemma 4.6 to remove all crossings with C. Since C contains all edges of H in all cases except for D_3 (which contains an H-edge not in C), we conclude that all edges of H are free of crossings unless $(G, H, \mathcal{H}) = D_3$. But the conclusion is also true for D_3, since the only edge in $E(G) - E(H)$ that could cross the remaining H-edge would have to do so oddly, which is not possible in an iocr-0 drawing. Hence, we can assume in all cases that all edges of H are free of crossings. For D_3 and D_4 the only remaining odd crossings can be between edges that share a common endpoint of degree 2, so they can easily be made even with respect to each other; but then we would have an ocr-0 drawing of D_3 or D_4 which, by Corollary 4.11, implies that these graphs are planar, contradicting Theorem 5.4. For D_{17} we note that if C is free of crossings, it must bound an empty face; we can add a vertex into this face and connect it to every vertex of C, yielding a graph that contains an iocr-0 drawing of $K_{3,3}$ which we know is not possible.

To simplify the discussion of the remaining two graphs, D_{11} and D_{14}, we use the following claim; the reader can visualize the situation using D_{11}.

Claim. Suppose we have an iocr-0 drawing of a PEG $(G, H, \mathcal{H})$ for which H is a tree, and there are three vertices x, y, and u so that: $e = xy$ is an even edge, $xu, yu \in E(H)$, and u is a cut-vertex of H so that $G - \{u, x, y\}$ is connected. Then there is an iocr-0 drawing of $(G, H, \mathcal{H})$ in which

- (*i*) xy, xu, and yu do not cross each other,
- (*ii*) all ends of $E(G) - E(H)$-edges incident to u lie outside xuy,
- (*iii*) all $E(G) - E(H)$-edges incident to u cross both xu and yu evenly.

The redrawing can be performed without changing the crossing parity between any pair of edges or changing the rotation at any vertex.

Let us first prove the claim. Use Lemma 4.8 to remove all crossings with e, establishing (*i*), since H-edges xu and yu do not cross each other, in particular xuy bounds a triangular region. Note that since H is a tree, the redrawing performed in Lemma 4.8 does not change the crossing parity for any pair of edges, and the rotation of all vertices remains the same. Since $G - \{u, x, y\}$ is connected, there is a path connecting any two vertices in $V(G) - \{u, x, y\}$. Any edge of such a path must cross the boundary edges of xuy evenly (since these pairs of edges do not share endpoints and the drawing is iocr-0), so any two vertices in $V(G) - \{u, x, y\}$ lie on the same side of the region xuy, let us say the outside. Let $f \in E(G) - E(H)$ be incident to u. Move the end of f to lie outside of xuy. Then f either crosses both xu and yu evenly, or both oddly. In the second case, perform an (f, u)-move, so that in either case f crosses both xu and yu evenly, establishing (iii), and the end of f' lies outside xuy, which was required by (ii). This completes the proof of the claim.

For D_{11} we proceed as follows: make $e = xy$ even using $(e, \cdot)$, $(f, \cdot)$ and $(h, \cdot)$-moves. By the claim, we can assume that f' lies outside xuy and crosses both xu and yu evenly. Move the end of f', if necessary, so it also crosses g evenly (keeping it outside xuy). At this point, all four edges incident to u cross each other evenly. Repeat the same argument with $e' = xy$ and v to ensure that all four edges incident to v cross each other evenly (this does not affect the crossing parity between f' and edges incident to u). We can now move the ends of h and h' at their shared endpoint so both g and g' are even. The remaining vertices have degree 3, so we can use Lemma 4.12 to ensure that any two edges sharing one of the remaining vertices as an endpoint cross evenly. At this point all pairs of edges cross each other evenly with the possible exception of h and h'. Apply Lemma 4.9 to obtain a drawing of D_{11} in which all edges except possibly h and h' are free of crossings. If h and h' do cross, we can route one along the other starting at one of the crossings to the common endpoint and make sure they cross each other evenly. At this point, another application of Lemma 4.9 yields a planar embedding of D_{11} which contradicts Theorem 5.4.

This leaves us with D_{14}. We apply the claim to $e = xy$ and v to ensure that both f and f' cross vx and vy evenly. Since the ends of f and f' at v now lie outside xvy we can move them, if necessary, so both f and f' cross g' evenly (it is possible that f and f' cross oddly at this point). Now apply the claim to $h = xy$ and u to make e' cross both xu and yu evenly. Move the end of e' so it crosses the third H-edge incident to u evenly as well, so now any two edges incident to u cross evenly. Move the ends of h and h' at their common endpoint so they cross both g and g' evenly. Use Lemma 4.12 to ensure that any two edges sharing one of the remaining vertices as an endpoint cross evenly. At this point all edges in H are even; if there are any odd pairs left, they must be among (h, h') and (f, f'). Apply Lemma 4.9 to obtain a drawing of D_{14} in which all edges except edges in the set $\{h, h', f, f'\}$ are free of crossings. Together with the edges of H, e' separates f from f' (the rotation at u is determined by h), so f and f' cannot cross each other. Hence, after another application of Lemma 4.9, the only remaining pair of edges that cross are h and h'. We deal with them just as we dealt with them in the case of D_{11}, showing that there is a planar drawing of D_{14} which contradicts Theorem 5.4. □

We can now complete the proof of Theorem 5.6.

Proof of Theorem 5.6. A planar embedding of a PEG $(G, H, \mathcal{H})$ is iocr-0, so it remains to show that if there is an iocr-0 drawing of G that extends $\mathcal{H}$, then G has a planar embedding extending $\mathcal{H}$. Suppose not, then $(G, H, \mathcal{H})$ contains one of the obstructions listed in Theorem 5.4 as a PEG-minor. By Lemma 5.7, there is an iocr-0 drawing of that minor, which contradicts Lemma 5.8. □

5.3 Partial Rotation Systems

When we ask whether a PEG $(G, H, \mathcal{H})$ is planar we start with a fixed planar embedding $\mathcal{H}$ of H. What if we only have a rotation system for H, or even weaker than that, a partial rotation system? Say ρ is a *partial rotation system* for a graph G if ρ specifies a cyclic order of a subset E_v^ρ of the edges E_v incident to v for every vertex v. Specifying a partial rotation system is more general than specifying the rotation system of a subgraph H, since the partial rotation system does not necessarily fix both ends of an edge in their respective rotations. It does not capture partial embeddability, though, unless H in $(G, H, \mathcal{H})$ is connected. Corollary 5.10 shows that embeddability of graphs with partial rotation systems can be reduced to partial embeddability.

Let $\mathrm{cr}(G, \rho)$ be the minimum $\mathrm{cr}(D)$ over all drawings D of G that respect the partial rotation system ρ, where we say that D *respects* ρ if the cyclic rotation of edges E_v^ρ at v is as prescribed by ρ. For $\mathrm{iocr}(G, \rho)$ we use a modified definition which counts odd crossings between adjacent edges if they are part of the same E_v^ρ. Given a drawing D respecting ρ, we define $\mathrm{iocr}(D, \rho) := \mathrm{iocr}(D) + \sum_{v \in V(H)} \sum_{e,f \in E_v^\rho} (i_D(e, f) \bmod 2)$. Let $\mathrm{iocr}(G, \rho)$ be the minimum of $\mathrm{iocr}(D, \rho)$ where D ranges over all drawings of G respecting ρ.

Theorem 5.9. *If* $\mathrm{iocr}(G, \rho) = 0$ *then* $\mathrm{cr}(G, \rho) = 0$.

Proof. Fix a drawing D of G with $\mathrm{iocr}(D, \rho) = 0$. Let G' be the graph obtained from G by subdividing each edge of G twice (adding two vertices to each edge), and let ρ' be the rotation induced by ρ on G'. We claim that $\mathrm{iocr}(G, \rho) = \mathrm{iocr}(G', \rho')$. To see that this is true, we subdivide the edges of G one edge at a time, which is sufficient to establish the claim by induction. So let uv be the edge about to be subdivided. Edges crossing uv fall into three disjoint categories: edges incident to u, edges incident to v, and edges incident to neither u nor v. Push all crossings with edges adjacent to u along uv close to u, and crossings with edges adjacent to v close to v. Note that this does not change the crossing parity between any two edges, since we introduce two new crossings as we push one edge past another along uv. We can now split uv into three parts: the part which is crossed by edges incident to u (which starts at u), the part which is crossed by edges incident to v (which starts at v), and the middle part which is crossed by all other edges. Introduce new vertices to separate these parts; the value of iocr of the drawing does not increase since crossings between adjacent edges remain crossings between adjacent edges, so $\mathrm{iocr}(G', \rho) \leq \mathrm{iocr}(G, \rho)$. On the other hand, $\mathrm{iocr}(G, \rho) \leq \mathrm{iocr}(G', \rho')$, since we can suppress the new vertices of G', and the truth of the claim follows.

Hence, we now have a drawing D' of G' satisfying $\text{iocr}(D', \rho') = 0$. Then, by definition, $i_{D'}(e, f) = 0 \bmod 2$ for all $e, f \in E_v^{\rho'}$ for all $v \in V(G)$. We now apply Lemma 4.4 with the graph H on edges $E(H) = \bigcup_{v \in V(G)} E_v^{\rho'}$ (we could also use Lemma 4.3), to obtain a drawing D'' of G' in which H by itself is embedded, say it has embedding $\mathcal{H}$, and the rotation system of H has not changed. Since the overall drawing is iocr-0, Theorem 5.6 allows us to conclude that there is a planar drawing of G' containing $\mathcal{H}$. In that drawing, suppress all vertices in $V(G') - V(G)$ to obtain a planar drawing of G which respects ρ. □

From this theorem we can obtain an algebraic criterion for testing whether a graph G has an embedding respecting a partial rotation system ρ, but it is easier to just use the reduction from embedding (G, ρ) to partially embedded planarity which is implicit in the proof of Theorem 5.9. Note that up to topological equivalence, $\mathcal{H}$ is the unique embedding of H respecting ρ, since after the subdivision, H is a forest.

Corollary 5.10. *Embeddability of a graph with partial rotation system reduces to partially embedded planarity.*

Since partially embedded planarity can be tested in linear time using the algorithm of Angelini, Di Battista, Frati, Jelnek, Kratochvl, Patrignani, and Rutter [3], embeddability of graphs with partial rotation systems can be tested in linear time as well.

One can imagine various variants of this problem. For example, what happens if rotations are allowed to flip? We take that problem up again in Section 6.6. Or, we could allow ρ to specify a partial (cyclic) order of the edges E_v incident to v. The resulting problem is **NP**-complete problem, since Angelini, Di Battista and Frati [2] showed that it is **NP**-complete to test whether 14 embedded graphs (on the same vertex set) have a simultaneous embedding in which the embedding of each graph is respected (so the partial order at each vertex is the disjoint union of 14 total orders). They also show that for three graphs the problem can be solved in polynomial time. Is there a Hanani-Tutte characterization for that case?

5.4 Testing Partially Embedded Planarity

In this section we show how the algebraic criterion for partially embedded planarity given in Theorem 5.6 can be turned into a polynomial-time algorithm. There is, of course, a linear-time algorithm for this task by Angelini, Di Battista, Frati, Jelínek, Kratochvíl, Patrignani, and Rutter [3].

For a given drawing D of a PEG $(G, H, \mathcal{H})$ let $i_D(e, f)$ be the number of times that e and f cross in D. Consider the following system $\text{PEP}(D)$ of equations over GF(2).[19] We have variables $x_{e,v}$ for every $e \in E(G)$ and $v \in V(G)$. For every pair (e, f) of independent edges we require that

$$i_D(e, f) + x_{e,h(f)} + x_{e,t(f)} + x_{f,h(e)} + x_{f,t(e)} = 0 \bmod 2,$$

[19] The particular system $\text{PEP}(D)$ is based on the suggestion of one of the referees and significantly simplifies the original system which used more complicated moves.

and $x_{e,v} = 0$ for $e \in E(H)$ and $v \in V(H)$. Intuitively, this last restriction ensures that the embedding of H does not change.

Theorem 5.11. *Let D be a drawing of PEG $(G, H, \mathcal{H})$. Then $(G, H, \mathcal{H})$ is planar if and only if the system PEP(D) has a solution over* GF(2).

Theorem 5.11 implies that recognizing whether a PEG is planar can be tested in polynomial time. Start with a drawing D for which i_D is easy to calculate (e.g. draw edges in $E(G) - E(H)$ as straight-line segments). Then PEP(D) is a linear system with $|E(G)| \cdot |V(G)|$ variables and $|E(G)| \cdot |E(G)|$ equations over GF(2) which can be solved in $O(|E(G)|^6)$ time. For another approach via simultaneous planarity, see the discussion after Corollary 6.39.

Corollary 5.12. *Partially embedded planarity can be recognized in polynomial time.*

For the interesting direction in Theorem 5.11, one would like to argue as follows: if PEP(D) has a solution, then there is an iocr-0 drawing D of G such that the rotation system of H in D agrees with the rotation system of $\mathcal{H}$, and every two edges in H cross each other evenly (for adjacent edges this is true because we do not allow H-edges to move over H-vertices). Lemma 4.4 allows us to remove crossings between edges of H, without changing the rotation system, and Theorem 5.6 then gives us an embedding of G. This proof works as long as H is connected, since in that case the rotation system of H determines the embedding $\mathcal{H}$. If H consists of multiple components, however, Lemma 4.4 does not guarantee that the embedding of H it produces is topologically equivalent to $\mathcal{H}$. It cannot guarantee this, in general, but in the special case of PEP(D) in which we do not allow edges of H to move over vertices of H, the embedding does not change. Rather than creating a modified version of Lemma 4.4, we will prove this directly.

Proof of Theorem 5.11. Let D be as in the statement. If $(G, H, \mathcal{H})$ is planar, then there is a drawing D' of $(G, H, \mathcal{H})$ which is planar and which contains $\mathcal{H}$. We can assume that $\mathcal{H}$ is in the same location in both D and D'.

Since the locations of all vertices in $V(H)$ are fixed, the two drawings differ only in the locations of vertices in $V(G) - V(H)$ and edges in $E(G) - E(H)$. Let D_t, $t \in [0, 1]$ be a sequence of drawings changing continuously from D to D' (without changing $\mathcal{H}$).

For two independent edges e and f, the value $i_{D_t}(e, f)$ can only change parity if one of the edges passes over an endpoint of the other edge, or if an endpoint of an edge passes over the other edge. For every vertex $v \in V(G)$ let $x_{e,v} = 1$ if edge e passes over v an odd number of times (as t goes from 0 to 1) and 0 otherwise. Note that $x_{e,v} = 0$ if $e \in E(H)$ and $v \in V(H)$ since neither e nor v move, and so the $x_{e,v}$ satisfy PEP(D).

We next argue that a solution to PEP(D) gives us a way to turn D into an iocr-0 drawing of $(G, H, \mathcal{H})$, since then, by Theorem 5.6, $(G, H, \mathcal{H})$ is planar, and we are done. Starting with D, create a new drawing D' as follows: for every $x_{e,v} = 1$, where $e \in E(G) - E(H)$ and $v \in V(G)$, perform an (e, v)-move. For every $x_{e,v} = 1$, where $e \in E(H)$, add a closed component of e around v (not containing any other vertices). In this second case, we know that $v \in V(G) - V(H)$, since $x_{e,v} = 0$ if $e \in E(H)$ and $v \in V(H)$. We now have an iocr-0 drawing of G which contains $\mathcal{H}$, except that some edges of H have additional closed components. We will remove those closed components using a strategy that is based on the proofs of Lemma 4.3 and Lemma 4.4. During the redrawing it can happen that a vertex is surrounded by multiple closed components belonging to the same edge. Since removing an even number of these closed components does not affect the crossing parity, we can always assume that each vertex is surrounded by at most one closed component from each edge.

We make use of a new move which we call a (v, e)-move. A (v, e)-move requires a simple curve c that connects v to some interior point of e (without crossing e otherwise) and has at most finitely many crossings with other parts of the drawing. Then a (v, e)-*move along* c moves v along c to e and then just beyond e. Edges incident to v are moved with v in a narrow tunnel surrounding c. All closed components that surrounded v at its original position are moved to v's new location. Let $L \subseteq E$ be the list of edges that v crosses oddly in this redrawing. For every $f \in L - E(H)$ perform an (f, v)-move. For every $f \in L \cap E(H)$ add a closed component belonging to f around v (and then reduce as described above, if necessary). A (v, e)-move does not affect the crossing parity between any pair of edges, since we compensated for any crossings we introduced while moving v. Moreover, if v was surrounded by a closed component of e, then after the (v, e)-move, that component will be gone. However, the overall number of closed components may have increased.

Let F be a maximal spanning forest of H. Suppose some edge $f \in F$ has an associated closed component around vertex v. Perform a (v, f)-move along a curve c from v to f that avoids F (this is possible, since F is a forest). After this move, the closed component of f around v will have been removed, and we have not introduced any new closed components of edges in F (since c avoided F). Repeating this procedure, we can ensure that no edge of F has any associated closed components.

Suppose there is an edge $e \in E(H) - F$ which still has associated closed components. As above, perform (v, e)-moves for every vertex v which is surrounded by a closed component belonging to e. Choose curves c_v avoiding F and e (since F together with the arc-component of e has at most two faces, with e on the boundary, this is always possible). At the end of all these moves, e will no longer have any associated closed components, however, we may have introduced closed components of edges in $E(H) - F$ that are crossed by a c_v. If this happened, however, note that the vertex v now lies in a different face of $F \cup \{e\}$—recall that e consists of a single arc-component at this point—than the arc-components of the edges of $E(H) - F$ it passed through. We can then proceed as follows: Let C be the unique cycle in $F \cup \{e\}$. Edges of C do not cross each other (or other edges of H), so we can make all edges of C even by moving the ends of edges in $E(G) - E(H)$ at the rotation of their endpoints in

C. We then use Lemma 4.1 to remove all crossings with edges in C, resulting in edges of $E(G) - E(H)$ with closed components. Reconnect closed components of $E(G) - E(H)$ to their arc-components if this can be done without crossing C. Any remaining closed components of edges in $E(G) - E(H)$ lie on the other side of C from their arc-components. The same, as we argued earlier, is true for those edges in $E(H)$ for which closed components were added during the initial (v, e)-moves. Hence, we can drop all closed components belonging to edges in $E(G) - E(H)$ and all closed components added during the (e, v)-moves, without changing the crossing parity between any pair of edges. As a result, we have reduced the overall number of closed components of the drawing before performing the (v, e)-moves by at least one.

Repeating this procedure, we can therefor remove all closed components of edges in $E(H) - F$, without introducing new closed components. Thus, the final drawing we obtain is an iocr-0 drawing of $(G, H, \mathcal{H})$ without any closed components, which is what we needed. □

5.5 Combinatorial Complexity

If we know that $(G, H, \mathcal{H})$ *is* planar, can we say anything about how complex the drawing of G may have to be—given a natural measure of complexity for G and $\mathcal{H}$ such as the number of bends in a poly-line drawing? In terms of computational complexity, this problem is hard: Patrignani [60] showed that if $\mathcal{H}$ is a straight-line embedding of H, then it is **NP**-complete to tell whether there is a straight-line embedding of G containing $\mathcal{H}$. Since we know that we can check in polynomial time whether $(G, H, \mathcal{H})$ is planar, this means that even if $(G, H, \mathcal{H})$ is planar, it is **NP**-complete to check whether G has a straight-line drawing extending $\mathcal{H}$. But what happens if we allow a poly-line drawing of edges in $E(G) - E(H)$?

The special case in which $E(H) = \emptyset$ and vertex locations are given was solved by Pach and Wenger [59] when they showed that we can find a poly-line drawing of a planar graph with at most a linear number of bends along each edge.

Question 5.13. Given a planar PEG $(G, H, \mathcal{H})$ with $\mathcal{H}$ a straight-line embedding, is there a poly-line drawing of G containing $\mathcal{H}$ that contains at most a polynomial (linear) number of bends along each edge?

We conjecture that the answer is yes; there is related work by Di Giacomo, Didimo, Liotta, Meijer, and Wismath [19] if H is a tree and vertices in $V(G) - V(H)$ have to be chosen from a given set of available locations. The general problem seems to be open.

It is tempting to conjecture that a greedy algorithm can find a poly-line embedding of G without too many bends: add edges to the poly-line drawing one at a time so that the drawing remains planar and the new edge minimizes some reasonable parameter such as the number of bends. This will not work for every ordering of edges, an example found by Kratochvíl and Matoušek [51] for string graphs can be adapted to show that adding the edges in a bad order can force an exponential number of bends (even if $E(H) = \emptyset$). See for example Dudeney's Chinese Railway puzzle in Figure 7.

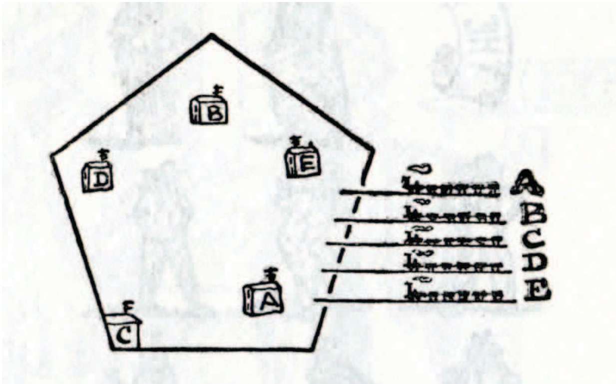

Figure 7: Dudeney's Chinese Railway Puzzle: Can you route each train to its depot without any of the tracks crossing? Problem 80 in Dudeney's "The Canterbury Puzzles" [22] of 1907.

If we replace the depots by points and lay tracks in the order D, E, B, A, C, we can solve the problem with 9 bends. On the other hand, if we first connect C, then there is no solution. If we start with D and A, then B will require 4 bends going back and forth.

Question 5.14. Given a planar $(G, H, \mathcal{H})$ so that $\mathcal{H}$ is a straight-line embedding, is there an ordering of the edges in $E(G) - E(H)$ that leads to a poly-line drawing of G with at most polynomially many bends if we add edges in the given order to $\mathcal{H}$ while minimizing the number of bends along each edge in each step?

Even the case where $E(H)$ is empty should be of interest. Finding a good ordering of $E(G) - E(H)$ may be computationally hard.

6 Simultaneous Graph Drawing

A *simultaneous drawing (with fixed edges)* of two graphs $G_i = (V(G_i), E(G_i))$, $i \in \{1, 2\}$, is a drawing of $G_1 \cup G_2 = (V(G_1) \cup V(G_2), E(G_1) \cup E(G_2))$. It is a *simultaneous embedding (with fixed edges)* of the two graphs if the drawings of G_1 and G_2 considered by themselves are planar embeddings. In other words, if any two edges e and f cross in the drawing, we must have $e \in E(G_1) - E(G_2)$ and $f \in E(G_2) - E(G_1)$ or vice versa.[20] We also say (G_1, G_2) is, or G_1 and

[20]The qualification "with fixed edges" refers to the fact that every edge belonging to both graphs, a *common edge*, is represented by a single curve. For a simultaneous embedding without fixed edges there is no requirement to draw common edges as the same curve. It is

G_2 are, *simultaneously planar*. We earlier saw the *common graph* $G_1 \cap G_2 = (V(G_1) \cap V(G_2), E(G_1) \cap E(G_2))$ of G_1 and G_2 which will play an in important role.

Remark 6.1. Simultaneous planarity is often defined for two graphs on the same vertex set, $V(G_1) = V(G_2)$, leading to a slightly restricted version of the simultaneous embeddability problem. For example, if we assume that $V(G_1) = V(G_2)$ it is known that simultaneous planarity can be tested for two graphs whose common graph is a star [4, 43]. If we do not assume that $V(G_1) = V(G_2)$, then we need to phrase this result differently, and require that the common graph is a star *and* contains all vertices of $G_1 \cup G_2$, as we do in item (vi) below.

Simultaneous planarity generalizes to arbitrarily many graphs; we will use SEFE_k to refer to simultaneous planarity of k graphs, and simply SEFE for the general problem. SEFE_k is known to be **NP**-complete for $k \geq 3$ [38]. The complexity of simultaneous planarity for two graphs remains open, but it is tempting to conjecture that it is in **P**. In fact, we will state a (combinatorial) conjecture later, that implies that simultaneous planarity is in **P**. Several special cases have been solved recently. One can test whether (G_1, G_2) is simultaneously planar

- (i) in linear time if $G_1 \cap G_2$ is 2-connected (Haeupler, Jampani, Lubiw [41]),
- (ii) in linear time if $G_1 \cap G_2$ consists of disjoint cycles (Bläsius, Rutter [10], for an arbitrary number of graphs),
- (iii) in quadratic time if a rotation system of $G_1 \cap G_2$ is given (Bläsius, Rutter [10]),
- (iv) in quadratic time if G_1 and G_2 are 2-connected, and $G_1 \cap G_2$ is connected (Bläsius, Rutter [11]),
- (v) in linear time if G_1 is a pseudoforest and G_2 is planar (Fowler, Gutwenger, Jünger, Mutzel, and Schulz [32]),
- (vi) in linear time if $G_1 \cap G_2$ is a star on all vertices of $G_1 \cup G_2$ (Angelini, Di Battista, Frati, Patrignani, Rutter [4], Hong, Nagamochi [43]).

Results on more than two graphs are rarer, though Angelini, Di Battista, and Frati showed that we can test in polynomial time whether three embedded graphs have a simultaneous embedding that respects the original embeddings; for 14 graphs, the problem is **NP**-complete [2]. For a survey on simultaneous planarity, see [9].

We proceed as follows. In Section 6.1 we will see that several well-known graph drawing problems are special cases of simultaneous planarity, including some whose complexity status was only resolved recently (partially embedded planarity) or is still open (the infamous c-planarity problem). We show that

generally agreed that the nomenclature is unfortunate, but it has become standard. Since we will not discuss the model without fixed edges in this paper, we typically drop "with fixed edges".

there are Hanani-Tutte style theorems for simultaneous planarity in the following cases:

- $G_1 \cap G_2$ consists of disjoint 2-connected components and subcubic components (Theorem 6.32),
- at least one of G_1 or G_2 is the disjoint union of subdivisions of 3-connected graphs (Theorem 6.38).

In particular, there are polynomial-time algorithms for testing simultaneous planarity in all of these cases (Corollary 6.34, Corollary 6.39). These generalize the results listed under (*i*) and (*ii*) above (without achieving the same running time).

We also give algebraic characterizations of simultaneous planarity if each connected component of $G_1 \cap G_2$ has a fixed embedding or a rotation system of $G_1 \cap G_2$ is given (Corollary 6.31), extending (*iii*) above. We do not know whether (*iv*)–(*vi*) have a Hanani-Tutte style characterization, though we suspect they do.

We can apply the results on simultaneous planarity to show a Hanani-Tutte style result for graphs with a partial rotation system where flipping of rotations is allowed (Lemma 6.41, Corollary 6.43).

6.1 The Power of Simultaneous Planarity

In this section we investigate the connections between simultaneous planarity and other notions of planarity, including partially embedded planarity, book embeddings, (radial) level planarity, upward planarity, constrained, and (leveled) c-planarity.

6.1.1 Partially Embedded Planarity

Theorem 6.2. *Partially embedded planarity reduces to the SEFE$_2$ problem.*

Proof. Suppose we are given $(G, H, \mathcal{H})$, where $H \subseteq G$ and $\mathcal{H}$ is an embedding of H. Let H' be an extension of H to a triangulation of the plane (the sphere, to be more precise—we require the outer face to be a triangle as well). We create the triangulation so that $E(H') \cap E(G) = E(H)$, that is, H' and G do not have any edges in common, except the edges of H. This may require adding vertices. We can also ensure that $|V(H')| \geq 4$. Then H', as a maximal planar graph on at least 4 vertices, is 3-connected (e.g. [54, Lemma 2.3.3]) and thus has a unique embedding up to (topological) equivalence by Whitney's theorem [21, Theorem 4.3.2]. Now $\mathcal{H}$ can be extended to an embedding of G if and only if H' and G have a simultaneous embedding with fixed edges (at this point we use that H' shares only edges in H with G). □

In Theorem 6.38 we will see that this allows us to test for partially embedded planarity via an SEFE$_2$ algorithm.

Remark 6.3. Jelínek, Kratochvíl and Rutter [46] point out that their obstructions to partially embedded planarity are obstructions for a simultaneous embedding of (G_1, G_2) because if $\mathcal{H}$ is an embedding of $H := G_1 \cap G_2$, then both of $(G_i, H, \mathcal{H})$ need to be planar.

6.1.2 Book Embeddings

We do not know whether the general 2-page book embedding can be encoded as a simultaneous planarity problem, even allowing an arbitrary number of graphs, not to mention k-page book embeddings. However, some special cases can be treated. In the *partitioned* book embedding problem every edge is assigned to a specific page in which it must be drawn. Given a tree T whose leaves are the vertices of graph G, in a T*-coherent* book embedding the order of the vertices along the spine must be a possible ordering of the leaves of T in a facial walk of a planar embedding of T. The standard book embedding problem is the special case in which T is a star.

We call a graph H in $G_1 \cup G_2$ *spanning* if $V(H) = V(G_1) \cup V(G_2)$. Typically, H will be a connected subgraph of $G_1 \cap G_2$ which forces $V(G_1) = V(G_2)$.

Lemma 6.4 (Angelini, Di Battista, Frati, Patrignani, Rutter [4])**.** *Simultaneous planarity of two graphs on the same vertex set whose intersection is a connected, spanning graph is equivalent to the partitioned* T*-coherent* 2*-page book embedding problem.*

This is pretty easy to see if the intersection is a spanning tree: the partitioned T-coherent 2-page book embedding problem is easily seen to be a special case of this simultaneous planarity problem (with T as the spanning tree, and the two partitions corresponding to the G_1 and G_2-only edges). Given the simultaneous planarity problem, one first ensures that all G_1 and G_2-only edges occur between leaves of the tree (by inserting edges into T). At this point the problem can be viewed as a 2-page book embedding problem, with T being the tree and edges partitioned as $(E(G_1) - E(G_2), E(G_2) - E(G_1))$.

Hong and Nagamochi [43] had shown earlier, that the partitioned T-coherent 2-page book embedding problem is solvable in linear time if T is a (spanning) star, so using Lemma 6.4 yields the following result.

Theorem 6.5 (Angelini, Di Battista, Frati, Patrignani, Rutter [4]; Hong, Nagamochi [43])**.** *Simultaneous planarity of two graphs on the same vertex set whose intersection is a spanning star can be solved in linear time.*

As far as we know, it is open whether the result remains true if we drop the condition that the star be spanning. Hong and Nagamochi [43] also show that the partitioned T-coherent 2-page book embedding problem is (linear time) equivalent to c-planarity for a flat clustering with two internal cluster (that is, one root cluster containing all vertices with two children partitioning the vertices).[21]

[21] A solution to the c-planarity problem for two clusters is also implicit in the work of Biedl, Kaufmann, and Mutzel [8].

More recently, Bläsius and Rutter [11] proved that simultaneous planarity of two 2-connected graphs (not necessarily on the same vertex set) whose intersection is connected can be tested in quadratic time. Lemma 6.4 now immediately implies the following corollary, which can also be proved directly using PQ-trees [45, Theorem 4.2]

Corollary 6.6. *The partitioned T-coherent 2-page book embedding problem can be solved in quadratic time if the graphs assigned to each page are connected.*

A Hanani-Tutte style characterization of the partitioned T-coherent 2-page book embedding problem would be implied by Conjecture 6.33. For the special case that T is a binary tree, and thus subcubic, Theorem 6.32 can be used to get a Hanani-Tutte style characterization (and a polynomial-time algorithm), even if the partitions are not connected. Polynomial-time solvability in general remains open if at least one of the partitions is not connected.

Can we say anything about T-coherent k-page book embeddings, where $k > 2$? Returning to the proof sketch of Lemma 6.4 we see that it easily extends to the case of arbitrarily many graphs: partitioned T-coherent k-page book embeddability is equivalent to the sunflower case of SEFE_k, where the common graph is a spanning tree. It has been conjectured that the general sunflower case is in polynomial time—see Remark 6.25—but it turns out that a much more restricted problem is sufficient to capture T-coherent book embeddings.

Lemma 6.7. *The partitioned T-coherent (k-page) book embedding problem reduces to the sunflower case of the SEFE_3 problem where the common graph is a (maximal) spanning forest.*

To simplify the construction in the following proof we will use the color model of simultaneous planarity to describe multiple graphs: We construct a single graph G in which every edge can have multiple colors (corresponding to the graphs G_i it belongs to). In this model a simultaneous drawing of G is a drawing of G in which no two edges that cross have a color in common.

Proof of Lemma 6.7. Take two copies T_1 and T_2 of T, and connect each vertex in T_1 to its corresponding vertex in T_2 by an edge. In any embedding of the resulting graph, the embeddings of T_1 and T_2 are mirrors of each other: if T_1 has rotation ρ at vertex v, then the rotation at T_2's copy of v is the same as ρ, just flipped. Using this observation, we can construct an SEFE problem on three graphs; we will think of the three graphs as the red, green, and blue graph. Edges that share a color may not cross in a simultaneous embedding. Take four copies T_1, T_2, T_3, T_4 of T, all edges of these graphs are simultaneously red, green, and blue. Connect T_1 to T_2 as described above using red edges, and T_1 to T_3 using green edges, and T_1 to T_4 using blue edges. Let the resulting graph be G. In any simultaneous embedding of G, T_2, T_3 and T_4 are embedded the same way. Now take k copies $G_1, \ldots, G_k$ of G and identify T_4 in G_i with T_2 in G_{i+1}, $1 \leq i < k$. We can find an embedding of the resulting graph F, since edges leaving T_4 are blue, while edges leaving T_2 are red. All T_3 are embedded the same way, and all edges leaving a copy of T_3 are green. Since the book embedding is partitioned, we can assign one copy of T_3 to each page, coloring the edges in each page red (blue would also work). This final graph has an SEFE_3 embedding if and only if the original book embedding problem is solvable. □

Can the construction be improved to yield an SEFE_2 problem? If there were such a reduction which depends polynomially on the number of pages k (like the reduction in Lemma 6.7), then SEFE_2 would be **NP**-complete, since the partitioned book embedding problem is **NP**-complete [45].

6.1.3 Level and Radial Level Planarity

We claimed earlier that there is a Hanani-Tutte theorem for level planarity implicit in [34]. For the sake of completeness, we include a proof.

Theorem 6.8 (Fulek, Pelsmajer, Schaefer, Štefankovič [34])**.** *Let G be a leveled graph with leveling ℓ. Then (G, ℓ) is level planar if and only if there is an iocr-0, leveled drawing of (G, ℓ).*

Proof. Let D be an iocr-0, leveled drawing of (G, ℓ). We use the reduction from [34]: "Perturb all vertices slightly, so that no two vertices are at the same level. If there is a vertex whose left or right rotation is empty, insert a new edge and vertex on its empty side so that the edges extend slightly beyond all the perturbed vertices from the same level." Call the resulting graph G' and its drawing D'. Then D' is an iocr-0 drawing of G', and it is x-monotone, since no two vertices lie on the same vertical line. By the Hanani-Tutte theorem for x-monotone drawings [Theorem 3.1][34], G' has an x-monotone embedding with the same vertex locations. Because of the edges we added to G, we can undo the perturbations of the vertices at the same level to obtain a leveled embedding of (G, ℓ). □

The proof of the Hanani-Tutte theorem for x-monotone drawings is direct, that is, it does not use obstruction sets; the problem of determining the obstruction set for level planarity is still open. The theorem makes it tempting to conjecture that there is a Hanani-Tutte theorem for radial level planarity as well. We have to leave that question open.

Theorem 6.9. *Level planarity reduces to the $SEFE_2$ problem.*

For the proof we will make use of a gadget that allows us to make sure edges pass a particular level only once and cannot double-back. We call it the *gate gadget*, Γ_k; it is shown in Figure 8(a). Assume that all the u_i and v_i, $1 \leq i \leq k$, lie within the outer dashed circle (we have to enforce that in upcoming constructions). The vertices $u_1, \dots u_k$ may attach to u in any order, but that order is the same (from left to right) in which $v_1, \dots, v_k$ attach to v: the reason is that we have k parallel paths uu_iv_iv (enclosed by the dashed circle). Figure 8(b) shows one way of using the gadget: the three edges entering the gadget from the left are in the same order as the corresponding edges leaving it to the right. If we restrict the drawing to the solid edges, and move u and v to the center of the middle solid edge (Figure 8(c)), delete u/v and separate the edges from the left and right into three separate edges, this gives us a drawing in which those edges cross the middle edge exactly once (Figure 8(d)).

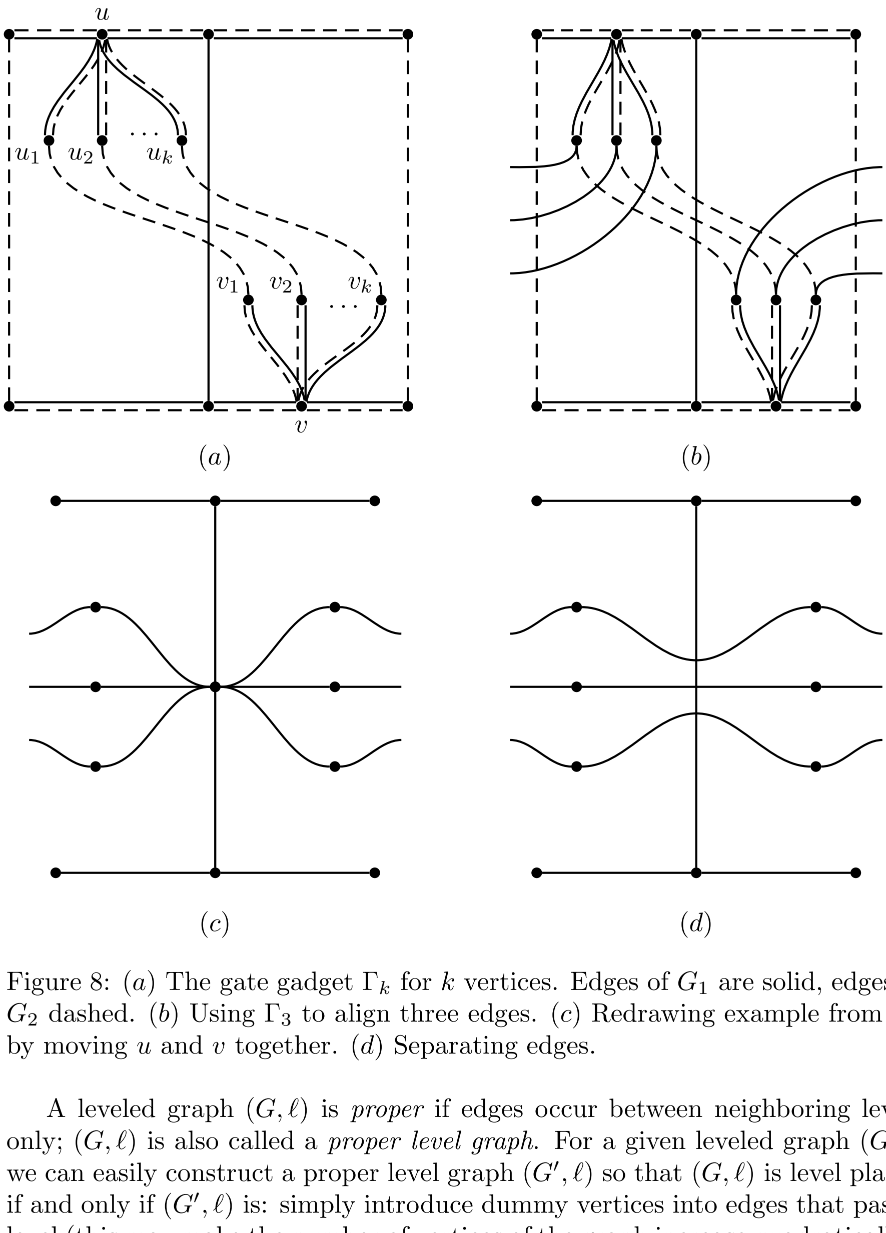

Figure 8: (a) The gate gadget Γ_k for k vertices. Edges of G_1 are solid, edges of G_2 dashed. (b) Using Γ_3 to align three edges. (c) Redrawing example from (b) by moving u and v together. (d) Separating edges.

A leveled graph (G, ℓ) is *proper* if edges occur between neighboring levels only; (G, ℓ) is also called a *proper level graph.* For a given leveled graph (G, ℓ) we can easily construct a proper level graph (G', ℓ) so that (G, ℓ) is level planar if and only if (G', ℓ) is: simply introduce dummy vertices into edges that pass a level (this may make the number of vertices of the graph increase quadratically). Note that for proper level graph the requirement that edges are x-monotone is unnecessary, since edges cannot change order between levels, so they can be straightened out (e.g. by connecting the endpoints of the edge directly). This is not generally true for leveled graphs, where edges could double-back if they were not required to be monotone.

Proof of Theorem 6.9. As we explained, we can assume that (G, ℓ) is proper. Let V_i be the vertices at level i. Create a gate gadget $\Gamma_{|V_i|}$ for the vertices at level i, call it Γ^i. We chain the gadgets as follows: Identify the right G_2-edge (and its endvertices) on the outer face of Γ^i with the left G_2-edge on the outer face of Γ_{i+1}. Make the left edge of the left-most gadget, and the right edge of the right-most gadget belong to $G_1 \cap G_2$; this creates a cycle of $G_1 \cap G_2$-edges surrounding the gadget; to ensure that all edges lie on the same side of the cycle, add a new vertex z and connect it by $G_1 \cap G_2$-edges to every vertex on the cycle. For every edge xy with $\ell(x) = i$ and $\ell(y) = i+1$ add an edge between v_x^i, associated with x in Γ^i and u_y^{i+1}, associated with y in Γ^{i+1}.

It is easy to see that if (G, ℓ) is level planar, then (G_1, G_2) is simultaneously planar. For the reverse direction, observe that G "nearly" occurs as a subdivision of the simultaneous drawing of (G_1, G_2). There is only one problem: as we go from the left copy of a vertex to the right copy of a vertex within the gate gadget (along the solid edges, not the dashed edges), all these paths pass through the same common path (connecting the u and v vertices within the gadget). However, since the gadget was built so that the vertices on the left occur in the same order as the vertices on the right, we can replace the common path by parallel (non-crossing paths) as we saw in Figure 8(b)–(d). □

For an example of the reduction described in the proof, see the leveled graph (G, ℓ) in Figure 9. Figure 10 shows the result of applying the reduction.

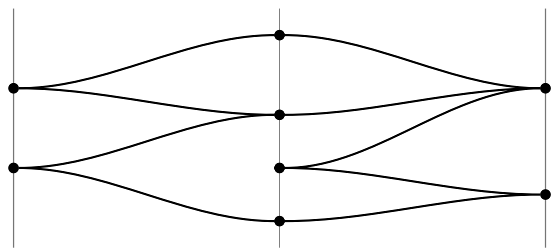

Figure 9: A proper level graph (with three levels).

We saw that level planarity of a leveled graph (G, ℓ) can be generalized by restricting the ordering of the vertices at each level by a family $\mathcal{T}$ of trees T_i so that the leaves of T_i are exactly the vertices at level i in (G, ℓ). In a $\mathcal{T}$-coherent embedding, the ordering of the vertices at each level must correspond to an ordering of the leaves of T_i. We called $(G, \ell, \mathcal{T})$ a generalized k-ary tanglegram and said it is embeddable if (G, ℓ) has a $\mathcal{T}$-coherent embedding.

Level planarity then becomes the special case where each T_i is a star whose leaves are the vertices v with $\ell(v) = i$. What happens in the reduction of

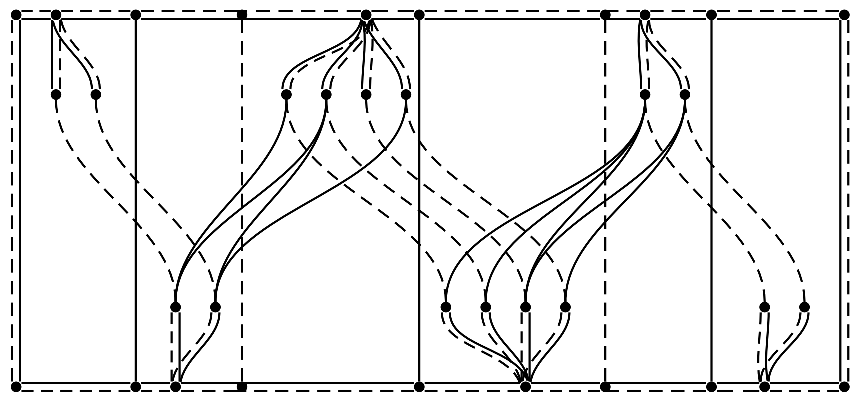

Figure 10: The graph from Figure 9 after applying the reduction in Theorem 6.9 (without the vertex z).

Theorem 6.9 if we replace the star with center u and leaves $u_1, \ldots, u_k$ in a gate gadget Γ_k with a tree T_i by choosing u to be an arbitrary vertex of T_i and identifying the $u_1, \ldots, u_k$ with the leaves of T_i? Then the ordering of the u_i (and thus the v_i) will be forced to be consistent with a facial walk of T_i. (Note that it does not matter which vertex we choose as u, since every vertex in a tree is on the outer face.) This means that we can modify the reduction in Theorem 6.9 to establish the following result:

Corollary 6.10. *Let $(G, \ell, \mathcal{T})$ a generalized k-ary tanglegram. Deciding $\mathcal{T}$-coherent level planarity reduces to the SEFE$_2$ problem.*

The complexity of deciding the embeddability of a generalized k-ary tanglegram is open, but it is known to be solvable in quadratic time if the number of vertices at each level is bounded by a fixed k [75]. We can deal with another special case: if $\mathcal{T}$ consists of rooted binary trees only, the reduction in Corollary 6.10 yields two graphs (G_1, G_2) for which $G_1 \cap G_2$ is nearly subcubic: only the vertices w_O of the region gadget have degree higher than 3. We can easily remedy this by creating two new vertices connected to w_O, one connected to the two edges leading to the top of the gadget, and the other to the two edges leading to the bottom corners of the gadget, making $G_1 \cap G_2$ subcubic. By Corollary 6.34, this SEFE$_2$ problem can be solved in polynomial time.

Corollary 6.11. *Let $(G, \ell, \mathcal{T})$ a generalized k-ary tanglegram so that all trees in $\mathcal{T}$ are rooted binary trees. Then we can test $\mathcal{T}$-coherent level planarity of (G, ℓ) in polynomial time.*

Level planarity is a special case of radial level planarity.

Lemma 6.12. *Level planarity reduces to radial level planarity.*

Proof. Let (G, ℓ) be a leveled graphs with levels in the range $1, \ldots, n$. Extend to a new leveled graph (G', ℓ') by adding an edge between two new vertices, one at level 0 and the other at level $n+1$. Then (G, ℓ) is level planar if and only if (G', ℓ') is radial level planar. □

The construction from Theorem 6.9 can be modified to capture radial level planarity. The difference is that in radial level planarity edges between levels lie in an annulus, not in a disk. Figure 11 shows how to connect (parts of) two gate gadgets so the edges between levels can lie in an annulus. This modification is sufficient to show Corollary 6.13. In Section 6.1.6 we will see that radial level planarity (and thus level planarity, by Lemma 6.12) can be reduced to c-planarity. That is a stronger result, since c-planarity reduces to the SEFE$_2$ problem (Theorem 6.17).

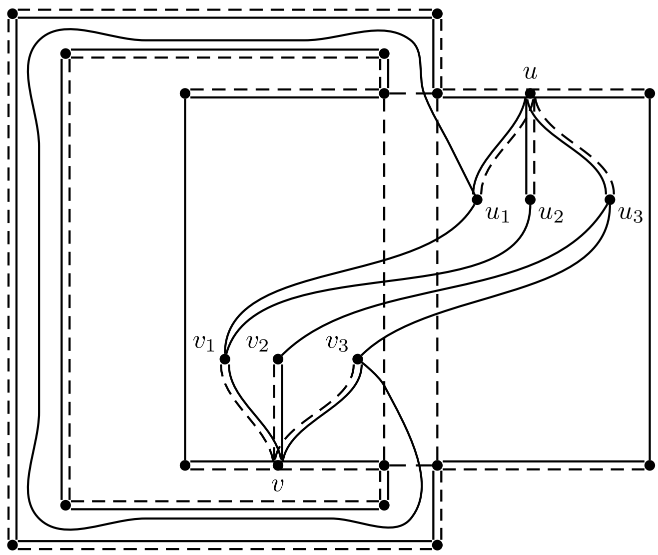

Figure 11: Adding an annulus between two gate gadgets.

Corollary 6.13. *Radial level planarity reduces to the SEFE$_2$ problem.*

Proof. Given a proper level graph (G, ℓ) construct (G_1, G_2) just as in Theorem 6.9 with two differences: do not create the vertex z and instead of identifying the right G_2-edge of Γ^i with the left G_2-edge of Γ^{i+1} insert a tunnel with sides consisting of $G_1 \cap G_2$-edges and ends made of G_2-edges, as shown in Figure 11. The tunnels realizing the annuli can all be realized in parallel, so it is easy to see that if (G, ℓ) is radial level planar, then (G_1, G_2) is simultaneously planar.

If (G_1, G_2) is simultaneously planar, we proceed similarly as in Theorem 6.9 to build a drawing of (G, ℓ) in which vertices lie on concentric circles around a common center. The edges between levels need not be radially monotone, but since (G, ℓ) is proper we can redraw them so that they are (applying Dehn twists to the annulus if necessary[22], we can assume that at least one edge does not wind around the inner circle it attaches to; we can then make it radially monotone; at this point the remaining edges lie in a plane, and we can redraw them so they become radially monotone). This gives us a radial planar drawing of (G, ℓ). □

6.1.4 Upward Planarity

In upward planarity, we are given a directed acyclic graph (dag) G and are asked whether there is an x-monotone drawing of G (every edge crosses every vertical line at most once) that respects the implicit ordering of G, that is, for $uv \in E(G)$ we require that $x(u) < x(v)$ in an upward planar drawing of G, where $x(u)$ is the x-coordinate of the vertex u in the drawing.

To the extent that we believe that SEFE_2 is polynomial-time solvable (as would be implied by Conjecture 6.20, for example), it is unlikely that upward planarity reduces to the SEFE_2 problem since upward planarity testing is **NP**-complete. However, upward planarity does reduce to the SEFE problem. To simplify the language in the following proof, we switch to the color model for simultaneous graphs.

Theorem 6.14. *Upward planarity reduces to the SEFE problem.*

Proof. Let G be a dag. Create a wheel W_4 with center x_c and four vertices x_w, x_n, x_e, x_s on all colors needed for the construction. For each $v \in V(G)$ add a black path $x_n v x_s$. Give each edge $uv \in E(G)$ its own color c_{uv}, and add that color c_{uv} to both paths $x_n u x_s$ and $x_n v x_s$. Moreover, add a path $x_w uv x_e$ in color c_{uv}. See Figure 12 for an example. Suppose the resulting graph has a simultaneous embedding. We can assume that x_c lies on the outside of the wheel cycle. Since the paths $x_n v x_s$ are all black, they force a linear ordering of the vertices $v \in V(G)$ inside the wheel. Moreover, the paths $x_w uv x_e$ force u to occur before v in that linear ordering, since the edges of this path cannot cross $x_n u x_s$ or $x_n v x_s$. The $x_n v x_s$ paths have the same order at x_n and y_n, so we can separate their endpoints turning them into $|V(G)|$ parallel paths $x_n^v v x_s^v$, each drawn as a line segment. An edge uv occurs between $x_n^u u x_s^u$ and $x_n^v v x_s^v$ (since it cannot cross either), and u is to the left of v. This gives us a drawing of G in which every edge is *x-bounded*: $x(u) < x(p) < x(v)$ for every point p on uv. By Corollary 2.7 in [34], the graph has an x-monotone embedding, which is an upward planar embedding of G. □

The proof uses $|E(G)|+1$ colors to encode a dag G. Are three colors sufficient to encode upward planarity?

[22] To perform a Dehn twist, take an essential curve in the annulus (a circle around the hole), cut the annulus at the circle, and give one of the parts a full twist close to the circle, so that after the twist, curves reconnect. This changes, for each curve, the number of times it winds around the inner hole of the annulus. See [70] for a more rigorous definition.

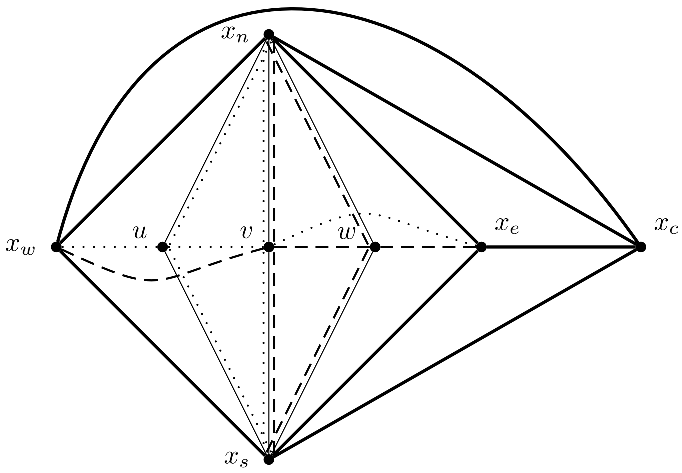

Figure 12: Construction for Theorem 6.14 if G is a directed path uvw. The color c_{uv} is represented by dotted lines, c_{vw} by dashed lines, and black by thin black lines. Thick black lines are edges having all colors.

6.1.5 Constrained Planarity

We already discussed various models of constrained planarity earlier, here we will follow the embedding constraint model suggested by Gutwenger, Klein, and Mutzel [40], since it seems to be the most general one. A *partial embedding constraint* at a vertex v is a rooted and ordered tree T_v whose inner nodes are of three types: gc-nodes, whose children may have arbitrary order, oc-nodes, whose children have a fixed (clockwise) order, and mc-nodes whose children either have the given order, or the reverse (mirror) of that order. The leaves of the tree are a subset of the edges incident to v. A partial embedding constraint describes admissible orderings at v by the possible orderings of leaves of the tree T around its root. If the subset of edges ordered by a partial embedding constraint includes all edges incident to v, we speak of a *(total) embedding constraint*. A graph with a collection of partial embedding constraints is *ec-planar with free edges* or *pec-planar* if it has an embedding in which the rotation at each vertex satisfies the partial embedding constraint. If all embeddings constraints are total, we simply say the graph is *ec-planar*.[23] The following result follows from work in Gutwenger, Klein, and Mutzel [40], but predates the formal notion of partially embedded planarity.

Theorem 6.15. *ec-planarity reduces to partially embedded planarity.*

Gutwenger, Klein, and Mutzel [40] showed that ec-planarity can be solved in

[23] Gutwenger, Klein, and Mutzel [40] only mention the possibility of partial embedding constraints in passing, speaking of allowing “free edges”.

linear time, Theorem 6.15 combined with the linear-time algorithm for partially embedded planarity [3] gives a new proof of this result.

The theorem follows from a construction described in [40]. We will only review that construction very cursorily, assuming that the reader is familiar with the ec-expansion as described in that paper.

Theorem 6.15. Suppose we are given a graph G together with a collection of embedding constraints T_v, $v \in V(G)$ (we can model the absence of an embedding constraint using a gc-node). For each tree T_v construct a gadget S_v in [40] roughly as follows: gc-nodes remain nodes, oc-nodes and mc-nodes with k children turn into a wheel W_{2k} with a cycle on $2k$ vertices, with alternating vertices becoming endpoints of the edges incident to the original node. The last level of the gadget consists of leaves representing the edges incident to v. From G and the embedding constraint, we construct a new graph G', the *ec-expansion*, by taking all gadgets S_v and identifying leaves representing the same edge (there will be at most two such leaves of course). Then G is ec-planar if and only if G' has a planar embedding in which the wheels corresponding to oc-nodes are embedded with the right orientation. The reason is that we can find an ec-planar embedding of G as a subdivision of G': v is the root of S_v, and edges correspond to paths through S_v. The condition on oc-nodes ensures that their children occur on the right order, and for mc-nodes the children are either in the correct or in the reverse order, based on whether the wheel flips (and whether the inner node of the wheel is inside or outside its cycle).[24] We can easily enforce the condition on the wheels corresponding to oc-nodes by using partially embedded planarity: let H be the subgraph of G' consisting of all the wheels corresponding to og-nodes, and fix an embedding $\mathcal{H}$ of H in which all wheels have the right orientation. Then G is ec-planar if and only if $(G', H, \mathcal{H})$ is planar. □

This construction does not easily seem to accommodate the presence of free edges, however, this can be done if we relax partially embedded planarity to SEFE.

Theorem 6.16. *ec-planarity with free edges reduces to the SEFE$_2$ problem.*

Proof. We are given a graph G and a collection of partial embedding constraints, T_w, $w \in V(G)$. We need to construct a pair of graphs (G_1, G_2) such that G is ec-planar (with free edges) if and only if (G_1, G_2) is simultaneously planar. We start with G_1 and G_2 being the empty graph. Construct the gadget S_w from T_w as in the proof of Theorem 6.15 and make it part of G_1. As in the proof of Theorem 6.15, each leaf of S_w corresponds to an edge incident to w. If the embedding constraint at w is total, we are done with S_w. If there are free edges at w, we proceed as follows: take a copy of the gate gadget Γ_k, as shown in Figure 8, where k is the number of edges at w whose rotation is constrained

[24] Gutwenger, Klein, and Mutzel [40, Section 4.2] impose additional restrictions on the embedding of the ec-expansion, but those are not actually necessary to extract an ec-planar drawing of G.

by T_w. Modify the gadget by making the top two $G_1 \cap G_2$-edges part of a C_3, and do the same for the bottom two $G_1 \cap G_2$ edges. Add G_1-edges between the leaves of S_w and the u_i vertices of Γ_k. Create a new vertex w_e for each edge e incident to w. If e is free, then connect w_e to v via a G_1-edge. If e is constrained, then connect it to the v_i corresponding to the u_i that was connected with the vertex belonging to e in S_w, using a G_1-edge. This new gadget, call it S_w again, now has vertices w_e associated with edge e incident to w. After we have created gadgets S_w for all $w \in V(G)$, we identify the two vertices in S_w and $S_{w'}$ corresponding to edge $ww' \in E(G)$. Finally, make all wheels belonging to oc-nodes part of $G_1 \cap G_2$. To ensure that all wheels have the same orientation, we connect them by 3-connected pieces, two at a time, until they are all connected. This completes the construction of (G_1, G_2). We claim that G is ec-planar with free edges if and only if (G_1, G_2) is simultaneously planar.

An ec-planar embedding of G easily leads to a simultaneously planar embedding of (G_1, G_2), the 3-connected G_2 pieces connecting the oc-nodes do not hinder the G_1 edges. In the other direction, observe that for each gate gadget, the uu_i and vv_i edges lie in the same G_2 face formed by the C_8 of G_2 surrounding the gadget (that is why we added the two C_3s). This forces the rotation of the uu_i at u to be a mirror of the vv_i at v which means that the gate gadget orders the constrained edges correctly, and allows the free edges to occur at any point in the rotation. □

We do not know the complexity of testing ec-planarity with free edges; Bläsius and Rutter [11] showed that the special case in which G is 2-connected and there are no oc-nodes, can be solved in linear time (they call this problem partially constrained PQ-planarity, their approach can probably be extended to also allow oc-nodes). Corollary 5.10 showed that the problem reduces to partially embedded planarity if each embedding constraint consists of a single oc-node, but allows free edges, and Corollary 6.43 shows that the problem can be solved in polynomial time if we allow partial embedding constraints consisting of single oc-nodes or single mc-nodes.

6.1.6 Clustered Planarity

Planarity of clustered graphs entered the graph drawing literature in papers by Feng, Cohen and Eades [29, 28] under the name c-planarity.[25] We base the following definition on the one given by Cortese and Di Battista in their very readable survey paper [18].

A *clustered graph* is a graph $G = (V, E)$ together with a rooted tree T whose leaves are the vertices of G. Every internal vertex ν of T corresponds to the *cluster* $V(\nu)$, the set of vertices of G occurring as leaves in the subtree of T rooted at ν. Let $G(\nu)$ be the subgraph of G induced on $V(\nu)$ (not necessarily connected). A *drawing* of a clustered graph (G, T) is a drawing of G together with a simple closed region $R(\nu)$ for each internal vertex ν or T so that: $G(\nu)$

[25] Feng's thesis [27] makes it clear that c-planarity was meant to abbreviate compound planarity, but it has generally been taken to mean clustered planarity.

is drawn inside $R(\nu)$ and $R(\nu) \subseteq R(\mu)$ if and only if ν is a descendant of μ in T. Regions which are not contained within each other are required to be disjoint. We say an edge e is *incident* to ν if e has one endpoint in $V(\nu)$ and the other in $V - V(\nu)$. An *edge-region crossing* occurs if an edge e that is not incident to ν crosses the boundary of $R(\nu)$ or if an edge e that is incident to ν crosses the boundary of $R(\nu)$ more than once. A drawing of a clustered graph is *c-planar* if it contains no edge crossings and no edge-region crossings.

Theorem 6.17. *c-planarity reduces to the SEFE$_2$ problem.*

Proof. Suppose we are given a clustered graph (G, T); without loss of generality, we can assume that for every $v \in V(G)$ there is a vertex ν of T so that $V(\nu) = \{v\}$ (each vertex is contained in its own cluster). We have to create two graphs G_1 and G_2 so that (G, T) is c-planar if and only if G_1 and G_2 have a simultaneous embedding with fixed edges. We will use G_1 to ensure a planar drawing of G and both G_1 and G_2 to enforce the clustering constraints. We start by letting $G_1 = G$ (this will change during the construction).

For every cluster $V(\nu)$, where ν is an interior vertex of T, we split all edges of G_1 incident to ν into two halves. We then add the region gadget C_ν as shown in Figure 13 and connect the severed ends to corresponding vertices in the gadget (if one end is connected to u_i, the other end is connected to v_i). Note that C_ν contributes edges to both G_1 and G_2. If μ is a direct descendant (child) of ν in T we add a G_2-edge from the inner hook in C_ν to the outer hook in C_μ. If μ is a leaf of T, so $V(\nu)$ consists of a single vertex, we add a G_2-edge from u_1 to w_I.

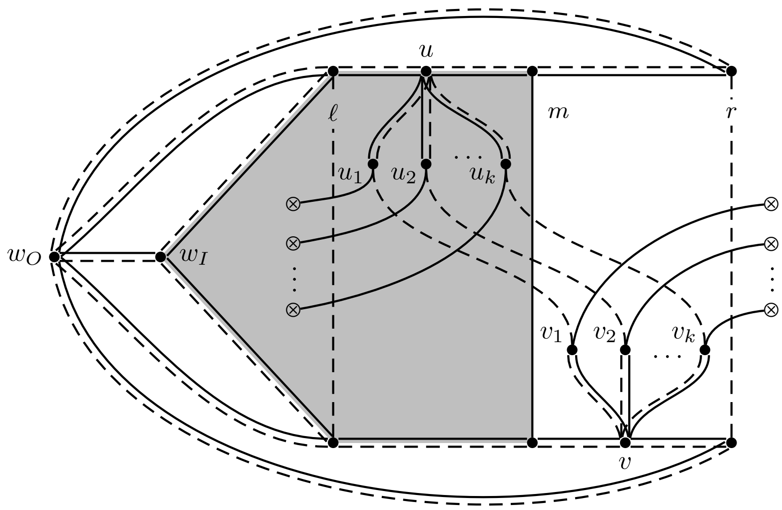

Figure 13: The region gadget C_ν drawn in the case that ν is incident to k edges. G_1-edges are solid, G_2-edges are dashed. The connector half-edges terminate in a $\otimes$. The inner hook is w_I, the outer hook w_O. The drawings of any gadgets C_μ belonging to a descendant μ of ν are forced to lie in the triangle in the gray region. We will define R_ν to be the gray area of C_ν.

If (G, T) is c-planar, then G_1 and G_2 have a simultaneous embedding with fixed edges, indeed, we can construct such an embedding in parallel to how we constructed G_1 and G_2 from (G, T). For the reverse direction, we let R_ν be the gray area of C_ν; let S_ν be the triangular (gray) subarea of R_ν. First note that every vertex of $V(\nu)$ lies in $S_\nu \subset R_\nu$: since every vertex was enclosed in its own region, there is a path from each vertex $u \in V(\nu)$ to w_I consisting of G_2 and $G_1 \cap G_2$-edges only (since we hooked up the region gadgets via G_2-edges). But a path of G_2-edges cannot cross the boundaries of S_ν, so v has to lie inside that triangle. From G_1 we can recover an embedding of G as we did with the gate gadget (Figure 8(b)–(d)): we move u and v together into the center of the black G_1-edge between them and split the vertices into parallel edges, reconnecting the original ends. This is possible, since the gate gadget as part of the region gadget enforces that the ordering of the u_i is the same as of the v_i. □

We return briefly to radial level planarity; we saw earlier that it is a special case of the SEFE$_2$ problem, but we already mentioned that something stronger is true.

Lemma 6.18. *Radial level planarity reduces to c-planarity.*

Proof. Let (G, ℓ) be a leveled graph. We can assume that at every level there is either a single vertex, or all vertices at the level are adjacent to vertices at both lower and higher levels: Suppose there is a vertex u at some level n which is not adjacent to any vertex at a lower (or higher) level. Insert a new level between $n-1$ (or $n+1$) and n, and add a new vertex v to that level together with the edge uv. Clearly, the new graph is radial level planar if and only if the original graph is. Repeating this construction shows that we can make the claimed assumption on (G, ℓ). We can now go further and remove levels with multiple vertices; suppose some level n contains more than one vertex. By assumption all these vertices are adjacent to vertices at both lower and higher levels. Perturb all vertices at that level slightly, creating a new level for each, close to n. The resulting graph is radial level planar if and only if the original graph is; to see that we can go from a radial level embedding of the new graph to a radial level embedding of the original graph observe that we can undo the perturbations: we can move a vertex to the same level as another vertex from the same original level by moving it along one of the edges incident to it (this is why we made these vertices adjacent to vertices at higher and lower levels).

We now have a leveled graph (G, ℓ) with exactly one vertex at each of the levels $1, \ldots, n = |V(G)|$. We create a clustering of the vertices of G so that the regions satisfy $R_{i-1} \subseteq R_i$ and $R_i - R_{i-1}$ contains the unique vertex at level i, where $i = 1, \ldots, n-1$. Then (G, ℓ) is radial level planar if and only if G with this clustering is c-planar. The forward direction is immediate, for the reverse direction we use the fact that there is at most one vertex per level; this ensures that we can make the drawings of edges radial monotone. □

Remark 6.19. Given a graph G equipped with both a leveling $\ell : V(G) \to \mathbb{N}$ and a clustering T, Forster and Bachmaier [30, 31] asked whether the graph is *clustered* *level planar (cl-planar)*: is there a level planar drawing of the graph which is also *c*-planar, and in which every region corresponding to a cluster intersects every horizontal line corresponding to a level in an interval (regions cannot have gaps at any level). The *cl*-planarity problem reduces to the SEFE_2 problem as long as (G, T) is *level-connected*, that is, every cluster contains an edge between any two consecutive levels at which the cluster is present. We do not include the details here, but the construction is a reasonably straightforward modification of gadgets we have already seen. Level-connectedness can be simulated using an additional color, so *cl*-planarity reduces to the SEFE_3 problem via a natural reduction (again the details are straightforward). This leaves open the question whether *cl*-planarity reduces to the SEFE_2 problem without the assumption of level-connectedness. Or may *cl*-planarity be **NP**-complete?

6.2 Simultaneous Planarity and Hanani-Tutte

For a simultaneous drawing D of (G_1, G_2), we let $D[G_i]$ be the drawing D restricted to the vertices and edges in G_i. The *simultaneous crossing number*, $\text{scr}(G_1, G_2)$, as the minimum of $\text{cr}(D[G_1]) + \text{cr}(D[G_2])$ over all simultaneous drawings D of G_1 and G_2 as introduced by Chimani, Jünger and Schulz [14]. So $\text{scr}(G_1, G_2) = 0$ if and only if (G_1, G_2) is simultaneously planar.

Similarly, we define $\text{socr}(G_1, G_2)$, the *simultaneous odd crossing number* as the minimum of $\text{ocr}(D[G_1]) + \text{ocr}(D[G_2])$ over all simultaneous drawings D of G_1 and G_2 and $\text{siocr}(G_1, G_2)$, the *simultaneous independent odd crossing number* as the minimum of $\text{iocr}(D[G_1]) + \text{iocr}(D[G_2])$ over all simultaneous drawings D of G_1 and G_2. We say a simultaneous drawing D of G_1 and G_2 is siocr-0 if $\text{iocr}(D[G_1]) + \text{iocr}(D[G_2]) = 0$. In that case we will speak of an siocr-0 drawing (dropping "simultaneous").

We start with Section 6.2.1 conjecturing a strong Hanani-Tutte theorem for simultaneous planarity, following by the short Section 6.2.2 establishing the weak Hanani-Tutte theorem for simultaneous planarity. Both sections require some specialized redrawing tools for simultaneous drawings, which are collected in Section 6.2.3.

6.2.1 Strong Hanani-Tutte for Simultaneous Planarity?

We conjecture that a strong Hanani-Tutte result for simultaneous embeddability of two graphs holds.

Conjecture 6.20. *If* $\text{siocr}(G_1, G_2) = 0$*, then* $\text{scr}(G_1, G_2) = 0$.

Remark 6.21. We have stated the conjecture in its most striking form, but some of the redrawing results we will see in Section 6.2.3 show that we can weaken the conjecture. By Lemma 6.30 below we can assume that G_1 and G_2 are connected, and by Lemma 6.29 it is sufficient to find an siocr-0 drawing in which edges of the the common graph $G_1 \cap G_2$ cross each other evenly.

If the conjecture were true, then simultaneous planarity could be tested in polynomial time, since $\text{siocr}(G_1, G_2)$ can be encoded as a linear system as follows: Given a simultaneous drawing D of (G_1, G_2) consider the system $\text{SEFE}(D)$

of equations over GF(2). As usual, let $i_D(e,f)$ denote the number of times e and f cross in D. We have variables $x_{e,v}$ for every $e \in E(G_1 \cup G_2)$ and $v \in V(G)$. For every pair (e,f) of independent edges so that $e, f \in E(G_1)$ or $e, f \in E(G_2)$ we require that $i_D(e,f) + x_{e,h(f)} + x_{e,t(f)} + x_{f,h(e)} + x_{f,t(e)} = 0 \bmod 2$.

Lemma 6.22. *Let D be a simultaneous drawing of (G_1, G_2). Then SEFE(D) is solvable if and only if* $\mathrm{siocr}(G_1, G_2) = 0$.

Proof. Fix D. Assume SEFE(D) is solvable. For each $x_{e,v} = 1$ perform an (e,v)-move. Let the resulting drawing be D'. For any pair of edges (e,f) we have $i_{D'}(e,f) = i_D(e,f) + x_{e,h(f)} + x_{e,t(f)} + x_{f,h(e)} + x_{f,t(e)}$. If (e,f) is a pair of independent edges with $e, f \in E(G_1)$ or $e, f \in E(G_2)$, then the last term is 0, so we have $i_{D'}(e,f) = 0$ for these edges, proving that $\mathrm{siocr}(G_1, G_2) = 0$. The other direction is a straightforward adaptation of the proof of Lemma 3.3. □

Corollary 6.23. *If Conjecture 6.20 is true, then simultaneous planarity of two graphs can be decided in polynomial time.*

Conjecture 6.20 is specifically about simultaneous drawings, but one can bring the problem back to drawings of single graphs. Recall that Conjecture 1.2 requires that if G has a subgraph H which is not involved in independent odd crossings with edges of G, then G can be redrawn so that edges of H do not cross each other, and no new pairs of independent edges crossing oddly are introduced.

Lemma 6.24. *If Conjecture 1.2 is true, then so is Conjecture 6.20.*

The proof uses Lemma 6.28 from Section 6.2.3 below.

Proof. Suppose $\mathrm{siocr}(G_1, G_2) = 0$ and let $H = G_1 \cap G_2$ and $G = G_1 \cup G_2$. Then G and H fulfill the assumptions of Conjecture 1.2, so if that conjecture is true, we can find a drawing of G in which edges of H no longer cross each other, and there are no new pairs of edges crossing oddly. In other words, an siocr-0 drawing of G_1 and G_2 containing an embedding of H (if considered by itself). By Lemma 6.28 there is a simultaneous embedding of G_1 and G_2, establishing the (conditional) truth of Conjecture 6.20. □

Remark 6.25 (Multiple Graphs)**.** We already know that a naïve generalization of Conjecture 6.20 to more than 2 graphs fails, recall the example in Figure 1. Can we identify special cases where the multiple graph version is true? One candidate is the *sunflower* scenario in which k graphs $G_1, \ldots, G_k$ all share the same common graph H, that is $H = G_i \cap G_j$ for all $1 \leq i < j \leq k$ (we do not require that $V(G_i) = V(H)$). Haeupler, Jampani, and Lubiw [41] conjecture that this case is polynomial-time solvable. This would be implied by Conjecture 1.2, using an argument similar to the one made in Lemma 6.24; it is not clear that Conjecture 6.20 would be strong enough to settle the sunflower case. By Lemma 6.7 polynomial-time solvability of the sunflower case for three graphs implies polynomial-time solvability of the partitioned (T-coherent) k-page book embedding problem for any k. Theorem 6.32 and Corollary 6.34 below can easily be extended to the sunflower case.

The truth of Conjecture 6.20 would establish SEFE_2 as a benevolent incarnation of the weak realizability problem, and an important one, since, just like weak realizability it is powerful enough to encode many other graph drawing problems. The conjecture is open, but we settle some special cases in the following sections.

6.2.2 Weak Hanani-Tutte for Simultaneous Planarity

Establishing a weak Hanani-Tutte theorem for simultaneous drawings is quite easy; unfortunately, it does not give us any algorithmic handle on the SEFE_2 problem. We first deal with the trivial case that G_1 and G_2 have no edges in common.

Lemma 6.26. *If* $\mathrm{socr}(G_1, G_2) = 0$ *and* $E(G_1) \cap E(G_2) = \emptyset$*, then* $\mathrm{scr}(G_1, G_2) = 0$ *and this can be realized without changing the rotation system of the drawing.*

Proof. Fix a drawing D realizing $\mathrm{socr}(G_1, G_2) = 0$. This drawing also witnesses $\mathrm{ocr}(G_1) = 0$ and $\mathrm{ocr}(G_2) = 0$, so by Lemma 4.2 both G_1 and G_2 have embeddings with their original rotation systems. If we choose embeddings for which the vertex locations and the ends of edges in G_1 and G_2 are as in the original drawing D, then we obtain a drawing witnessing $\mathrm{scr}(G_1, G_2) = 0$. □

The following result is the weak Hanani-Tutte theorem for simultaneous planarity. We do not know whether it is true for arbitrary surfaces, since it relies on Lemma 4.2 which fails for surfaces of higher genus [64, Example 5.3].

Theorem 6.27. *If* $\mathrm{socr}(G_1, G_2) = 0$*, then* $\mathrm{scr}(G_1, G_2) = 0$.

Proof. Fix a simultaneous drawing of G_1 and G_2 realizing $\mathrm{socr}(G_1, G_2) = 0$. Then all edges in $E(G_1) \cap E(G_2)$ are even, so we can apply Lemma 4.2 to redraw $G_1 \cup G_2$ so that no edge in $E(G_1) \cap E(G_2)$ is involved in any crossings, and any two edges that cross oddly in the new drawing already crossed oddly in the original drawing (so there must be one edge from each of G_1 and G_2). Take a maximal spanning forest F of $E(G_1) \cap E(G_2)$ and contract the edges in F (keeping track of the rotation system by merging the rotations of the vertices that are identified). If the resulting graph still contains edges of $E(G_1) \cap E(G_2)$, they must be loops; pick such a loop and split the drawing into two parts: the drawing within the loop and the drawing outside the loop; recursively redraw each part without changing the rotation system and so that each part is simultaneously planar, then recombine the two embeddings and add back the loop. In the base case, there are no edges of $E(G_1) \cap E(G_2)$ left, and we can apply Lemma 6.26 to find a simultaneous embedding of the graphs without changing the rotation system. We can then uncontract the edges of F without introducing crossings (since we kept the rotation system), obtaining a simultaneous embedding of G_1 and G_2. □

The gap between Theorem 6.27 about socr and Conjecture 6.20 seems quite formidable; we start exploring some of that ground in the next few sections. Before we do so, we create some tools for that exploration.

6.2.3 Redrawing Tools for Simultaneous Drawings

We first establish two simple redrawing results for the cases that the common graph $G_1 \cap G_2$ of two graphs G_1, G_2 is drawn so that no two of its edges cross each other, Lemma 6.28, or every two of its edges cross evenly, Lemma 6.29. The final redrawing tool, Lemma 6.30 explains why we can assume that G_1 and G_2 are connected when asking whether $\mathrm{siocr}(G_1, G_2) = 0$.

Lemma 6.28. *If there is an siocr-0 drawing of G_1 and G_2 in which no two edges belonging to $G_1 \cap G_2$ cross each other, then $\mathrm{scr}(G_1, G_2) = 0$ and there is a simultaneous embedding of (G_1, G_2) in which the drawing of $G_1 \cap G_2$ is the same as in the original drawing.*

Proof. Fix D as in the statement of the theorem and let $H = G_1 \cap G_2$. From G_1, G_2, and H construct G_1', G_2' and H' by adding all vertices of $G = G_1 \cup G_2$. Let $\mathcal{H}'$ be the drawing of H' as part of D. The current siocr-0 drawing D of (G_1, G_2) can be restricted to an iocr-0 drawing of PEG $(G_1', H', \mathcal{H}')$ and an iocr-0 drawing of PEG $(G_2', H', \mathcal{H}')$. By Theorem 5.6 there are planar drawings of G_1' and G_2' that extend $\mathcal{H}'$. Since $\mathcal{H}'$ contains all the vertices of G_1 and G_2, the two drawings together are a simultaneous planar drawing of G_1 and G_2 showing that $\mathrm{scr}(G_1, G_2) = 0$. Moreover, the drawing of H is the same as in the original drawing D. □

Lemma 6.29. *If there is an siocr-0 drawing of G_1 and G_2 in which every two edges of $H = G_1 \cap G_2$ cross each other evenly, then $\mathrm{scr}(G_1, G_2) = 0$ and there is a simultaneous embedding of (G_1, G_2) in which the drawing of $G_1 \cap G_2$ has the same rotation system as in the original drawing.*

Proof. We can apply Lemma 4.4 with $G = G_1 \cup G_2$ and $H = G_1 \cap G_2$ since every two edges in H cross evenly, and the given drawing is siocr-0. Since we do not introduce new pairs of independent edges crossing oddly, we obtain a new siocr-0 drawing D of (G_1, G_2) in which no two edges of H cross each other. At this point, we apply Lemma 6.28 to obtain the conclusion. Note that the rotation system stays the same. □

Note that we do not generally assume that G_1 and G_2 are connected. Bläsius and Rutter [10] showed that given (G_1, G_2) one can find a pair of connected graphs (G_1', G_2') in linear time so that (G_1, G_2) is simultaneously planar if and only if (G_1', G_2') is. This means that one can assume that G_1 and G_2 are connected for simultaneous planarity problems (at the price of some linear time preprocessing).

We prove a similar result for siocr. Say (G_1', G_2') *extends* (G_1, G_2) if $V(G_i) = V(G_i')$ and $E(G_i) \subseteq E(G_i')$, for $i \in \{1, 2\}$.

Lemma 6.30. *If $\mathrm{siocr}(G_1, G_2) = 0$, then we can build in linear time a pair of connected graphs (G_1', G_2') extending (G_1, G_2) with $\mathrm{siocr}(G_1', G_2') = 0$.*

Proof. Fix an siocr-0-drawing of (G_1, G_2) and let $V = V(G_1) \cup V(G_2)$. We can assume that $G_1 \cup G_2$ is connected: let u and v be vertices belonging to different components of $G_1 \cup G_2$. Imagine the siocr-0-drawing of the component containing v on a sphere and project it on the plane so that v lies on the outer face. Now re-insert that drawing into the original drawing so that u and v lie in the same face. We can then connect them with a $G_1 \cap G_2$-edge. The resulting drawing is still an siocr-0-drawing.

Let $uv \in E(G_1) - E(G_2)$ so that u and v belong to different connected components of G_2. Let $U \subseteq V$ be the set of vertices of the connected component of G_2 containing u. Any edges between U and $V - U$ must belong to $E(G_1) - E(G_2)$. Our goal is to make uv even with respect to all edges while keeping the drawing siocr-0. By wrapping edges adjacent to uv around their common endpoint with uv, we can make all edges adjacent to uv even with respect to it (see beginning of Section 6.5 for a definition of edge wrapping). Since the overall drawing is siocr-0, this means that uv crosses all G_1-edges evenly. So if there is an edge f that crosses uv oddly, we must have $f \in E(G_2) - E(G_1)$. Perform (f, u')-moves for all vertices $u' \in U$. Now f crosses uv evenly. Moreover, the drawing remains siocr-0: by construction, f can only change parity with edges that have one endpoint in U and the other in $V - U$. But all those edges belong to $E(G_1) - E(G_2)$. Since $f \in E(G_2) - E(G_1)$, changing these parities does not affect siocr-0 of the drawing.

At this point uv is an even edge, and we can apply Lemma 4.1 to remove all crossings with uv; we reconnect closed components avoiding uv, the drawing remains siocr-0. We now have an siocr-0-drawing of (G_1, G_2) in which uv is free of crossings. For $(G'_1, G'_2) = (G_1, G_2 \cup \{uv\})$ we have shown that $\text{siocr}(G'_1, G'_2) = 0$.

This process works similarly for $uv \in E(G_2) - E(G_1)$, so we can repeatedly apply it until G'_1 and G'_2 are both connected. By construction (G'_1, G'_2) extends (G_1, G_2), and $\text{siocr}(G'_1, G'_2) = 0$. Finally, the construction (though not the redrawings) can be performed in linear time. □

6.3 Simultaneous Partially Embedded Planarity

Suppose we are given an embedding $\mathcal{H}$ of the common graph $H = G_1 \cap G_2$, and we ask whether there is a simultaneous embedding of (G_1, G_2) in which the common graph is embedded as $\mathcal{H}$. This is easily seen to be equivalent to testing whether PEGs $(G_i, H, \mathcal{H})$ are planar for $i = 1, 2$ (Jünger and Schulz prove a similar result [49, Theorem 4]). Using the linear-time algorithm for partially embedded planarity due to Angelini, Di Battista, Frati, Jelínek, Kratochvíl, Patrignani and Rutter [3], the simultaneous partially embedded planarity problem can be solved in linear time.

Our main goal in this section is to relax the embedding condition to allow different connected components of $G_1 \cap G_2$ to move with respect to each other, that is, we do not restrict their relative location. Suppose, for example, $G_1 \cap G_2$ consists of a $K_{2,n}$ and a $K_{2,m}$, $n, m \geq 3$. If we assume the embeddings of $K_{2,n}$ and $K_{2,m}$ are fixed up to orientation-preserving topological equivalence, then there are still $(n + m - 1)$ different ways the two graphs can (in the absence of other restrictions) be located relative to each other. Specifying an embedding of all of H restricts this to a single choice, so instead we now want to investigate what happens if instead we specify the embedding of each connected component of the common graph, or as Bläsius and Rutter [10] put it, the SEFE problem in which *the embedding of each connected component of the common graph is fixed*. There is one more subtlety here. We may not want to restrict the embedding of each connected component in the plane, but rather in the sphere. In other words, we may only want to restrict the rotation system of a connected component (we discussed this difference earlier in Remark 5.1). Then $K_{2,n}$ by itself, for example, has n different embeddings in the plane, depending on which face we choose to be the outer face. If we have two graphs, say $K_{2,n}$ and $K_{2,m}$ then we have nm choices for a particular relative embedding of $K_{2,n}$ and $K_{2,m}$ so $(n+m-1)nm$ different embeddings overall. This is the model investigated by Bläsius and Rutter [10, 66].

We study the hybrid problem, where for each connected component of the common graph we either specify an embedding in the sphere (a rotation system) or an embedding in the plane. We say "embedding in the sphere" rather than rotation system, since they are equivalent for connected graphs and it allows us to still call this problem the "SEFE problem in which the embedding of each connected component of the common graph is fixed".

The characterization work has already been done in Lemma 6.29, so we are left with rephrasing the characterization as an algebraic system. Let (G_1, G_2) be given together with an embedding (plane or sphere) for each connected component of $G_1 \cap G_2$. Let D be a simultaneous drawing of (G_1, G_2) in which all connected components of $G_1 \cap G_2$ are drawn so they satisfy the embedding and constraints. Consider the following system SPEP(D) of equations over GF(2). We have variables $x_{e,v}$ for every $e \in E(G_1 \cup G_2)$ and $v \in V(G)$. For every pair (e, f) of independent edges so that $e, f \in E(G_1)$ or $e, f \in E(G_2)$ we require that $i_D(e, f) + x_{e,h(f)} + x_{e,t(f)} + x_{f,h(e)} + x_{f,t(e)} = 0 \bmod 2$. Moreover, if H' is a connected component of $H = G_1 \cap G_2$ for which an embedding on the sphere has been fixed, we require that $i_D(e, f) + x_{e,h(f)} + x_{e,t(f)} + x_{f,h(e)} + x_{f,t(e)} = 0 \bmod 2$ for *every* pair of edges $e, f \in E(H')$. If H' is a connected component of H for which an embedding in the plane has been specified, we require that $x_{e,v} = 0 \bmod 2$ for every $e \in E(H')$ and $v \in V(H')$.

Corollary 6.31. *Suppose we are given (G_1, G_2) and an embedding (plane or sphere) for each connected component H_i of $H = G_1 \cap G_2$. Let D be a simultaneous drawing of (G_1, G_2) in which the embedding of H_i satisfies the embedding constraint specified for H_i. Then (G_1, G_2) has a simultaneous embedding satisfying the embedding constraint of each connected component of $G_1 \cap G_2$ if and only if SPEP(D) is solvable.*

It may be easiest to think of the two extreme cases: specifying a rotation system of the common graph (an embedding on the sphere for each component), or specifying a planar embedding of each connected component of the common graph. The corollary implies that these and the more general hybrid SEFE problem can be tested in polynomial time. Bläsius and Rutter [10] showed that the problem can be solved in quadratic time for the rotation system variant, and their algorithm can probably be adapted to deal both with the planar embedding variant and the hybrid variant [66].

Proof of Corollary 6.31. Fix D and suppose SPEP(D) is solvable. For each (e, v) for which $x_{e,v} = 1$ perform an (e, v)-move. Let D' be the resulting drawing. The conditions for SPEP(D) imply that D' is an siocr-0 drawing of (G_1, G_2). If H' is a connected component of H for which an embedding on the sphere was given, we required that $i_{D'}(e, f) = i_D(e, f) + x_{e,h(f)} + x_{e,t(f)} + x_{f,h(e)} + x_{f,t(e)} = 0 \bmod 2$ for every pair $e, f \in E(H')$ (independent or not), so again every two edges in $E(H')$ cross evenly. We conclude that any two edges in H cross each other evenly: if they belong to different components of $G_1 \cap G_2$ because they are independent and the drawing is siocr-0, and if they belong to the same component, by the argument we just made. We can now apply Lemma 6.29 to obtain a simultaneous embedding of (G_1, G_2). The rotation system of no connected component of H changes in the application of Lemma 6.29. So we know that each connected component of H has the same rotation system in D and D'. This means that the simultaneous embedding already satisfies all the embedding constrains for connected components for which an embedding on the sphere was given. Consider a component H' for which an embedding in the plane was specified. At least in principle, an application of Lemma 6.29 could have changed the outer face of the embedding of H', but we want to argue that this did not happen in the redrawings we performed. Let $\mathcal{H}'$ be the given planar embedding of H'. Imagine a curve γ from a point at infinity to a vertex v on the outer face of $\mathcal{H}'$ in D so that γ does not cross edges of $\mathcal{H}'$. During the redrawing in Lemma 6.29 we cut edges, create closed components, drop closed components or reattach closed components. None of those operations changes the crossing parity of any edge of H' with γ. The only operations that could would be (e, v)-moves, where $e \in E(H')$ and v is the vertex in H' to which γ attaches, but we required that $x_{e,v} = 0 \bmod 2$, since e and v both belong to H'. Hence, after the redrawing, γ still attaches to the same location in the rotation at v, and the point at infinity is still in the same outer face of $\mathcal{H}'$ which means that the new drawing of H' is (orientation-preserving) topologically equivalent to $\mathcal{H}'$ in the plane.

For the reverse direction, suppose that there is a simultaneous embedding D' of (G_1, G_2) in which every connected component of H satisfies its embedding constraint. Let D_t, $t \in [0, 1]$ be a continuous and smooth transformation from $D = D_0$ to $D' = D_1$. We want to argue that this transformation is equivalent to a set of (e, v)-moves satisfying $x_{e,v} = 0 \bmod 2$ for every $e \in E(H')$, $v \in V(H')$ if H' is a connected component of H for which an embedding in the plane has been specified. So let us first consider such a component H'. Since the embedding of H' in D and D' is the same (though not necessarily in the same location), we can assume that the transformation simply moves the embedding of H' from its original to its new location. This shows that $x_{e,v} = 0 \bmod 2$ for every $e \in E(H')$, $v \in V(H')$, since two edges belonging to H' need never cross during the transformation. We are left with the argument that the transformation from D to D' is equivalent to a set of (e, v)-moves. This is immediate for pairs of independent edges (e, f), since their crossing parity is only affected by one of them passing over an endpoint of the other. However, we also require $i_{D'}(e, f) = 0 \bmod 2$ for pairs of adjacent edges $e, f \in H'$ if H' is a connected component for which an embedding in the sphere is given. This means we need to also track the movement of the ends of H'-edges at vertices of H'. Since the rotation of H' in D and D' is the same, we can assume that the ends of H'-edges at a vertex $v \in V(H')$ are short line segments at the same angles in both D and D'. This means that as we go from D to D', an H'-edge e at v winds around v a certain number of times. The effect of a single (clockwise or counterclockwise) turn of e around v on the other H'-edges is the same as that of an (e, v)-move. Hence, the effect of the transformation from D to D' on $i_D(e, f)$ for $e, f \in G_1 \cup G_2$ is captured by (e, v)-moves. □

We can try to take this characterization of partially embedded simultaneous planarity further: can we restrict which face of a component H' another component H'' ends up in? The answer is, yes, this can be done by not allowing arbitrary combinations of (e, v)-moves between $e \in E(H')$ and $v \in V(H'')$, but instead introducing complex moves: (E', v)-moves for $E' \subseteq E$, which are equivalent to performing (e, v)-moves for every $e \in E'$, and E' is the set of edges that would have to be crossed to move H'' into the particular face of H'. We have to leave details for a later time.

6.4 Simultaneous Planarity and 2-connectivity

As a small step towards Conjecture 6.20 we deal with the case that the common graph consists of 2-connected and subcubic components.

Theorem 6.32. *If* $\operatorname{siocr}(G_1, G_2) = 0$ *and* $G_1 \cap G_2$ *consists of disjoint* 2-*connected components and subcubic components, then* $\operatorname{scr}(G_1, G_2) = 0$.

We would like to prove Theorem 6.32 under the condition that both G_1 and G_2 are 2-connected and no conditions on $G_1 \cap G_2$. Recall that Bläsius and Rutter [11] have a quadratic-time testing algorithm assuming one also knows that $G_1 \cap G_2$ is connected. We think an even stronger result is true: Call a pair (G_1, G_2) *well-connected* if for every two adjacent edges $e, f \in E(G_1) \cap E(G_2)$, there are cycles $C_i \in E(G_i)$, $i \in \{1, 2\}$ so that e and f lie on both C_1 and C_2. Note that (G_1, G_2) is well-connected if both G_1 and G_2 are 2-connected, or if every connected component of $G_1 \cap G_2$ is 2-connected (or an isolated vertex or an isolated edge).

Conjecture 6.33. *If* $\operatorname{siocr}(G_1, G_2) = 0$ *for a well-connected pair* (G_1, G_2)*, then* $\operatorname{scr}(G_1, G_2) = 0$.

This conjecture may be a reasonable next stepping stone on the way from Theorem 6.32 to Conjecture 6.20.

Proof of Theorem 6.32. Let H_1 consist of all 2-connected components in $H = G_1 \cap G_2$, and let H_2 contain all remaining components in H. So $H = H_1 \dot\cup H_2$,

H_1 is the disjoint union of 2-connected components, and H_2 is the disjoint union of subcubic components. Now apply Lemma 4.5 with H_1 and $G = G_1 \cup G_2$. Note that the assumptions of the lemma on H_1 are satisfied, since we start with an siocr-0 drawing of G. We obtain a drawing of (G_1, G_2) which is siocr-0, since no new pair of independent edges crosses oddly, and in which H_1 is free of crossings. Since H_2 is subcubic, we can make all its edges even with respect to each other: pick a vertex v in H_2, since it is incident to at most three H_2-edges, we can move the ends of those edges in the rotation at v so they cross each other evenly; repeating this for all $v \in V(H_2)$ ensures that all H_2 edges cross each other evenly. This local redrawing does not affect the drawing of H_1 at all. Apply Lemma 4.4 to $G = G_1 \cup G_2$ and $H = H_1 \cup H_2$. After the redrawing we have an siocr-0 drawing of G_1 and G_2 in which H is embedded (though it may still cross edges in $E(G) - E(H)$). Lemma 6.28 yields $\mathrm{scr}(G_1, G_2) = 0$. □

As usual, the characterization in Theorem 6.32 gives us a naïve polynomial-time algorithm for testing whether two graphs whose intersection is the disjoint union of 2-connected graphs and subcubic graphs, have a simultaneous embedding. A linear-time algorithm for the case that the intersection graph is 2-connected was recently given by Haeupler, Jampani, and Lubiw [41] using PQ-trees. Angelini, Di Battista, Frati, Patrignani, and Rutter [4] also solved this problem under the stronger assumption that the intersection graph is spanning and 2-connected. Haeuper, Jampani, and Lubiw [41] observe that their result extends to arbitrarily many graphs that all have the same graph in common (the sunflower case). Our characterization also extends to the sunflower case where the common graph is the disjoint union of 2-connected and subcubic components.

Corollary 6.34. *Simultaneous planarity can be tested in polynomial time for pairs of graphs (G_1, G_2) for which $G_1 \cap G_2$ consists of disjoint 2-connected components and subcubic components only.*

Proof. By Theorem 6.32 simultaneous planarity for (G_1, G_2) can be verified by testing whether $\mathrm{siocr}(G_1, G_2) = 0$, which can be done in polynomial time by Lemma 6.22. □

See Corollary 6.11 for an application of this result.

Remark 6.35. Theorem 6.32 seems very close to covering the case of level planarity or even c-planarity, recall the gate-gadget in Figure 8(a) or the region gadget in Figure 13. The components of $G_1 \cap G_2$ are nearly 2-connected, except for some stars that occur in pairs (the uu_i and vv_i edges). In spite of some effort, we have not yet been able to prove Hanani-Tutte style results that cover the gadgets used in the reductions, even for level planarity.

6.5 Simultaneous Planarity, 3-connectivity, and Partially Embedded Planarity

An *edge wrap* is an (e, v)-move in which v is an endpoint of e. We can think of deforming e close to v and wrapping it once around v before continuing. A *v-flip* means flipping the rotation at v (reversing the cyclic order of edges incident to v) so that every pair of edges incident to v changes crossing parity. Both moves are illustrated in Figure 14.

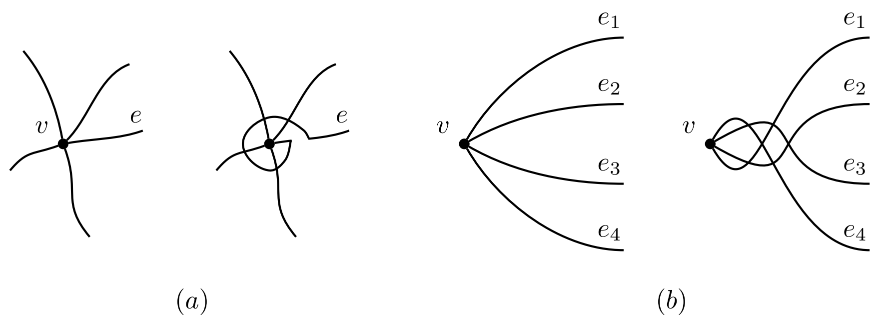

Figure 14: (a) An (e, v)-wrap. (b) A v-flip.

By Whitney's theorem, a subdivision of a planar 3-connected graph has a unique embedding in the plane (up to topological equivalence) [21, Theorem 4.3.2]. In other words, there are at most two planar rotation systems for such a graph (in some cases, such as cycles, the two rotation systems cannot be distinguished). Recall that a drawing is ocr-0 if every two edges cross evenly.

Lemma 6.36. *Suppose G is the subdivision of a 3-connected planar graph and we are given an iocr-0 drawing of G in which the rotation system is the same as in a planar drawing of G. Then the drawing can be made ocr-0 by a sequence of edge wraps and vertex flips.*

Proof. Start with an iocr-0 drawing of G with the rotation system as described in the statement of the lemma. Let v be a vertex at which two edges incident to v cross oddly, and suppose there is no sequence of edge wraps and vertex flips so that all pairs of edges incident to v cross each other evenly.

We claim that we can then perform a sequence of edge wraps and vertex flips so that there are four edges e, f, g, and h incident to v such that e, f, and g cross pairwise evenly, and h crosses exactly one of e, f, and g oddly.

To see that the claim is true, start with two arbitrary edges e and f at v. Performing an (f, v)-wrap if necessary, we can assume that e and f cross evenly, so there must be a third edge g incident to v. If g crosses both e and f evenly, we are fine, if it crosses both oddly, perform a (g, v)-wrap, so again g crosses both e and f evenly. Otherwise g crosses exactly one of e and f oddly, say e, and the other, f, evenly. Perform a v-flip. Then among the three pairs only g and e cross evenly. But then a (f, v)-wrap ensures that e, f, and g cross each other pairwise evenly. This shows that v must be incident to a fourth edge h. If every additional edge crosses either all of $\{e, f, g\}$ evenly, or all of them oddly, an (h, v)-wrap, if necessary, makes each additional edge even with respect to all previous edges. Since we assumed that this is not possible, there must be some edge h that is even with exactly one or exactly two of $\{e, f, g\}$. Performing an (h, v)-move if necessary, we can assume that h is even with respect to exactly one of $\{e, f, g\}$, say h crosses e oddly, and f and g evenly.

Now let C be a cycle containing e and f. We can make all edges of C even (with respect to all edges) by modifying rotations at vertices of C. For v we need to change the rotation to do so: the end of h has to be moved past e, so h and e cross evenly, but not past f, since h and f already cross evenly. So the new rotation at v differs (as a cyclic order) from the original rotation at v since the cyclic order of e, h, and f has changed. Moreover, since g does not have to move in the rotation to make the cycle even, the rotation at v differs from the flip of the rotation at v as well.

Once C is even, we can apply Lemma 4.2 to find an iocr-0 drawing of G in which C is free of crossings. Now let $H = (V(G), E(C))$, the subgraph of H which has all vertices of G and the edges of C. By Theorem 5.6 we conclude that there is a planar drawing of G extending H. This is not possible: the embedding of C and the vertex locations are the same as in H, so the rotation at v does not agree with either of the two possible rotation systems of G when embedded in the plane. □

Remark 6.37. Instead of invoking Theorem 5.6 in the proof of Lemma 6.36, we could have also reproduced a proof of the Hanani-Tutte theorem.

We are now in a position to resolve Conjecture 6.20 in case one of the graphs consists of subdivisions of 3-connected graphs.

Theorem 6.38. *If* $\mathrm{siocr}(G_1, G_2) = 0$ *and at least one of* G_1 *or* G_2 *is the disjoint union of subdivisions of* 3*-connected graphs, then* $\mathrm{scr}(G_1, G_2) = 0$.

In the main case of interest one of G_1 or G_2 is 3-connected, but the proof has enough flexibility to establish the stronger result.

Proof. Begin with an siocr-0 drawing of G_1 and G_2, assume G_2 is the disjoint union of subdivisions of 3-connected graphs. Change the rotation of edges of each connected component of G_2 so that its rotation corresponds to the rotation system of a planar embedding of that component of G_2 (the drawing remains siocr-0). By Lemma 6.36 we can perform a sequence of G_2-edge wraps and vertex flips for each connected component of G_2 that turn the drawing of G_2 into an ocr-0 drawing (while keeping the overall drawing siocr-0).

Let $H = (V(G_1) \cup V(G_2), E(G_1) \cap E(G_2))$, so H is the intersection graph of G_1 and G_2 with all vertices of G_1 and G_2 added to it. Lemma 4.4 allows us to redraw G_1 and G_2 so that no two edges of H cross each other, and the drawing remains siocr-0. If we restrict the drawing to G_2, we have an iocr-0 drawing of $(G_2, H, \mathcal{H})$, where $\mathcal{H}$ is the planar embedding of H. By Theorem 5.6 there is a planar drawing of G_2 extending $\mathcal{H}$. Combining this planar drawing with the current drawing of G_1 (which we can do, since $\mathcal{H}$ fixed all vertex positions), we obtain an siocr-0 drawing of (G_1, G_2) in which G_2 is planar, although edges of G_2 may be crossed by edges in $E(G_1) - E(G_2)$.

Consider the PEG $(G_1, H, \mathcal{H})$. The current drawing shows that this PEG has an iocr-0 drawing, and thus, by Theorem 5.6 a planar drawing. However, that drawing, together with the drawing of G_2 shows that $\text{scr}(G_1, G_2) = 0$. □

Corollary 6.39. *Simultaneous planarity can be tested in polynomial time for pairs of graphs (G_1, G_2) at least one of which is the disjoint union of subdivisions of 3-connected graphs.*

Proof. This immediately follows from Theorem 6.38 together with Lemma 6.22. □

Theorem 6.2 showed that partially embedded planarity is a special case of the SEFE_2 problem. Since G_2 in that reduction is 3-connected, Corollary 6.39 implies that we can test for partially embedded planarity in polynomial time by asking whether $\text{siocr}(G_1, G_2) = 0$. In the reverse direction, Corollary 6.39 implies that if we are given (G_1, G_2) and an embedding $\mathcal{G}_2$ of G_2, we can test in polynomial time whether (G_1, G_2) has an SEFE embedding in which G_2 is embedded as $\mathcal{G}_2$: add edges and vertices to $\mathcal{G}_2$ turning G_2 into a 3-connected graph (and without adding edges that belong to G_1, that may require adding new vertices). There is a faster solution due to Angelini, Di Battista, Frati, Jelínek, Kratochvíl, Patrignani and Rutter [3, Theorem 5.1] who showed that this variant of SEFE reduces to partially embedded planarity.

Theorem 6.40 (Angelini, Di Battista, Frati, Jelínek, Kratochvíl, Patrignani and Rutter [3])**.** *Suppose we are given (G_1, G_2) an an embedding $\mathcal{G}_2$ of G_2. Let $\mathcal{G}_{1,2}$ be the embedding of $G_1 \cap G_2$ in $\mathcal{G}_2$. Then (G_1, G_2) has an SEFE embedding where G_2 is embedded as $\mathcal{G}_2$ if and only if the PEG $(G_1, G_1 \cap G_2, \mathcal{G}_{1,2})$ is planar.*

6.6 Simultaneous Planarity and Rotation Systems

In Section 5.3 we considered graphs G with a partial rotation system ρ. In this section we want to add one small twist: we allow rotations at certain vertices to flip, that is, we specify a set $U \subseteq V(G)$ so that for each vertex $v \in U$ the cyclic order of edges E_u^ρ at u is either as specified by ρ or reversed. This corresponds to the special case of ec-planarity with free edges in which all embedding constraints are single oc-nodes or mc-nodes. Partially embedded planarity is not flexible enough to capture this variant, but simultaneous planarity is.

Lemma 6.41. *Given a graph G with partial rotation system ρ and a set of vertices $U \subseteq V(G)$, the problem of whether G can be embedded so that the rotation at all vertices is as specified by ρ, with flips allowed for vertices in U, can be rephrased as a simultaneous planarity problem, where one of the graphs is the disjoint union of 3-connected components.*

Proof. Let G_1 be G after each edge has been subdivided twice; in particular, $V(G) \subseteq V(G')$. With each edge vu in E_v^ρ we can uniquely associate an edge in G': pick the first edge on the path of length three from v to u in G'. Turn the star of edges with center v whose rotation is determined by ρ into a wheel, so that the additional cycle respects the rotation at v. Add all these wheels to a new graph G_2. A simultaneous embedding of (G_1, G_2) contains a subdivision of G which realizes ρ except that the rotation at every vertex may be flipped. Take all the wheels in G_2 associated with vertices belonging to $V(G) - U$, and add edges to that subgraph until it is triangulated; subdivide each of the new edges once (so we do not create common edges with G_1); add all the new edges and vertices to G_2. Now in a simultaneous embedding of (G_1, G_2) either all rotations of vertices in $V(G) - U$ are flipped compared to ρ or none of them are. Vertices associated with vertices in U, however, can still flip independently. □

Remark 6.42. One can extend the construction in Lemma 6.41 to force groups of vertices to flip simultaneously or not at all.

Corollary 6.43. *Suppose we are given a graph G with partial rotation system ρ and a set of vertices $U \subseteq V(G)$. We can test in polynomial time whether G can be embedded so that the rotation at all vertices is as specified by ρ, except for vertices of U where rotations are also allowed to flip.*

Proof. The result follows immediately from Lemma 6.41 and Corollary 6.39. □

7 Questions on Crossing Numbers

While planarity and its variants are important parts of graph drawing, in practice many visualization tasks will have to allow crossings of some type, both to make visualizations possible and to improve other aesthetic criteria. This is already built into some of the drawing models: simultaneous planarity pointedly ignores certain types of crossings in the drawing, and weak realizability gives full flexibility (at the expense of **NP**-completeness).

In a next step, we should consider crossing number variants. We already saw that the Hanani-Tutte theorem can be stated as saying that $\text{iocr}(G) = 0$ implies $\text{cr}(G) = 0$. Indeed, $\text{iocr}(G) = \text{cr}(G)$ as long as $\text{iocr}(G) \leq 2$ and $\text{cr}(G) \leq \binom{2\,\text{iocr}(G)}{2}$ [65], and there are graphs for which $\text{iocr}(G) < \text{cr}(G)$.

Question 7.1. Is there a function f for which $\text{scr}(G_1, G_2) \leq f(\text{siocr}(G_1, G_2))$?

For leveled graphs we saw that $\text{iocr}(G, \ell) = 0$ implies that $\text{cr}(G, \ell) = 0$ [34] (see Theorem 6.8). Examples from [35] show that there is no function f with $\text{cr}(G, \ell) \leq f(\text{iocr}(G, \ell))$. What happens for upward planarity? Let us use $\text{cr}_{\preceq}$ and $\text{iocr}_{\preceq}$ for the variants of crossing number and independent crossing number in which drawings of a directed graph have to be x-monotone (each edge crosses every vertical line at most once) and all edges point in the same direction. By definition, $\text{cr}_{\preceq}(G) = 0$ if and only if G is upward planar. As we mentioned in Example 2.2, $\text{iocr}_{\preceq}(G) = 0$ implies $\text{cr}_{\preceq}(G) = 0$.

Question 7.2. Is there a function f so that $\text{cr}_{\preceq}(G) \leq f(\text{iocr}_{\preceq}(G))$ for every directed graph G?

If we remove the leveling (ordering) from the graph and only require that the drawing of the graph be x-monotone, there is such a result due to Pach and Tóth [58], $\text{mon-cr}(G) \leq \binom{2\,\text{mon-iocr}(G)}{2}$, where mon-cr is the monotone crossing number, and mon-iocr is the monotone independent crossing number.

Acknowledgements

Several people were kind enough to read through the long manuscript draft for this paper and send me feedback. Rado Fulek spotted several errors, now fixed, I hope, Karsten Klein made many useful suggestions (in particular on graph drawing issues), as did Michael Pelsmajer, who read an earlier version of the paper. I also want to thank the two anonymous referees who diligently worked through this paper at break-neck speed and were still able to give me very detailed and high-quality comments that led to many important clarifications and simplifications (most noteworthy, Section 5.4 which has been replaced in its entirety).

References

[1] P. Angelini, M. Di Bartolomeo, and G. Di Battista. Implementing a partitioned 2-page book embedding testing algorithm. In Didimo and Patrignani [20], pages 79–89. doi:10.1007/978-3-642-36763-2_8.

[2] P. Angelini, G. Di Battista, and F. Frati. Simultaneous embedding of embedded planar graphs. In T. Asano, S.-i. Nakano, Y. Okamoto, and O. Watanabe, editors, *Algorithms and Computation*, volume 7074 of *Lecture Notes in Computer Science*, pages 271–280. Springer Berlin / Heidelberg, 2011. doi:10.1007/978-3-642-25591-5_29.

[3] P. Angelini, G. Di Battista, F. Frati, V. Jelínek, J. Kratochvíl, M. Patrignani, and I. Rutter. Testing planarity of partially embedded graphs. In M. Charikar, editor, *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2010, Austin, Texas, USA, January 17-19, 2010*, pages 202–221. SIAM, 2010. URL: http://dl.acm.org/citation.cfm?id=1873601.

[4] P. Angelini, G. Di Battista, F. Frati, M. Patrignani, and I. Rutter. Testing the simultaneous embeddability of two graphs whose intersection is a biconnected or a connected graph. *J. Discrete Algorithms*, 14:150–172, 2012. doi:10.1016/j.jda.2011.12.015.

[5] P. Angelini, F. Frati, and M. Kaufmann. Straight-line rectangular drawings of clustered graphs. *Discrete Comput. Geom.*, 45(1):88–140, 2011. doi:10.1007/s00454-010-9302-z.

[6] D. Archdeacon. A Kuratowski theorem for the projective plane. *J. Graph Theory*, 5(3):243–246, 1981. doi:10.1002/jgt.3190050305.

[7] C. Bachmaier, F. J. Brandenburg, and M. Forster. Radial level planarity testing and embedding in linear time. *J. Graph Algorithms Appl.*, 9(1):53–97, 2005. doi:10.7155/jgaa.00100.

[8] T. Biedl, M. Kaufmann, and P. Mutzel. Drawing planar partitions. II. HH-drawings. pages 124–136. doi:10.1007/10692760_11.

[9] T. Bläsius, S. G. Kobourov, and I. Rutter. Simultaneous embeedings of planar graphs. In Tamassia [71], chapter 11, pages 349–382. To appear.

[10] T. Bläsius and I. Rutter. Disconnectivity and relative positions in simultaneous embeddings. In Didimo and Patrignani [20], pages 31–42. doi:10.1007/978-3-642-36763-2_4.

[11] T. Bläsius and I. Rutter. Simultaneous PQ-ordering with applications to constrained embedding problems. In S. Khanna, editor, *SODA 2013*, pages 1030–1043. SIAM, 2013.

[12] F. Brandenburg, D. Eppstein, M. T. Goodrich, S. Kobourov, G. Liotta, and P. Mutzel. Selected open problems in graph drawing. In G. Liotta, editor, *Graph drawing*, volume 2912 of *Lecture Notes in Computer Science*, pages 515–539. Springer-Verlag, Berlin, 2004. doi:10.1007/978-3-540-24595-7_55.

[13] G. Chartrand and F. Harary. Planar permutation graphs. *Ann. Inst. H. Poincaré Sect. B (N.S.)*, 3:433–438, 1967.

[14] M. Chimani, M. Jünger, and M. Schulz. Crossing minimization meets simultaneous drawing. In *PacificVis*, pages 33–40. IEEE, 2008. doi:10.1109/PACIFICVIS.2008.4475456.

[15] M. Chimani and R. Zeranski. Upward planarity testing via SAT. In W. Didimo and M. Patrignani, editors, *Graph Drawing*, volume 7704 of *Lecture Notes in Computer Science*, pages 248–259. Springer Berlin Heidelberg, 2013. doi:10.1007/978-3-642-36763-2_22.

[16] C. Chojnacki (Haim Hanani). Über wesentlich unplättbare Kurven im dreidimensionalen Raume. *Fundamenta Mathematicae*, 23:135–142, 1934.

[17] F. R. K. Chung, F. T. Leighton, and A. L. Rosenberg. Embedding graphs in books: a layout problem with applications to VLSI design. *SIAM J. Algebraic Discrete Methods*, 8(1):33–58, 1987.

[18] P. F. Cortese and G. Di Battista. Clustered planarity. In *Computational geometry (SCG'05)*, pages 32–34. ACM, New York, 2005. doi:10.1145/1064092.1064093.

[19] E. Di Giacomo, W. Didimo, G. Liotta, H. Meijer, and S. K. Wismath. Point-set embeddings of trees with given partial drawings. *Comput. Geom.*, 42(6-7):664–676, 2009. doi:10.1016/j.comgeo.2009.01.001.

[20] W. Didimo and M. Patrignani, editors. *Graph Drawing - 20th International Symposium, GD 2012, Redmond, WA, USA, September 19-21, 2012, Revised Selected Papers*, volume 7704 of *Lecture Notes in Computer Science*. Springer, 2013.

[21] R. Diestel. *Graph theory*, volume 173 of *Graduate Texts in Mathematics*. Springer, Heidelberg, 4th edition, 2010.

[22] H. E. Dudeney. *The Canterbury puzzles and other curious problems*. Thomas Nelson and Sons, London, Edinburgh, New York, 1907.

[23] P. Eades, Q. Feng, X. Lin, and H. Nagamochi. Straight-line drawing algorithms for hierarchical graphs and clustered graphs. *Algorithmica*, 44(1):1–32, 2006. doi:10.1007/s00453-004-1144-8.

[24] P. Eades, Q.-W. Feng, and X. Lin. Straight-line drawing algorithms for hierarchical graphs and clustered graphs. In *Graph drawing*, volume 1190 of *Lecture Notes in Comput. Sci.*, pages 113–128, London, UK, 1997. Springer. doi:10.1007/3-540-62495-3_42.

[25] A. Estrella-Balderrama, J. J. Fowler, and S. G. Kobourov. On the characterization of level planar trees by minimal patterns. In D. Eppstein and E. R. Gansner, editors, *Graph drawing*, volume 5849 of *Lecture Notes in Computer Science*, pages 69–80. Springer, Berlin, 2010. doi:10.1007/978-3-642-11805-0_9.

[26] A. Estrella-Balderrama, E. Gassner, M. Jünger, M. Percan, M. Schaefer, and M. Schulz. Simultaneous geometric graph embeddings. In S.-H. Hong, T. Nishizeki, and W. Quan, editors, *Graph drawing*, volume 4875 of *Lecture Notes in Computer Science*, pages 280–290. Springer, Berlin, 2008. doi:10.1007/978-3-540-77537-9_28.

[27] Q.-W. Feng. *Algorithms for Drawing Clustered Graphs.* PhD thesis, Department of Computer Science and Software engineering, University of Newclastle, Apr. 1997.

[28] Q.-W. Feng, R. F. Cohen, and P. Eades. How to draw a planar clustered graph. In D.-Z. Du and M. Li, editors, *Proceedings of the 1st Annual International Conference (COCOON '95) held in Xi'an, August 24–26, 1995*, volume 959 of *Lecture Notes in Computer Science*, pages 21–30. Springer-Verlag, Berlin, 1995. doi:10.1007/BFb0030816.

[29] Q.-W. Feng, R. F. Cohen, and P. Eades. Planarity for clustered graphs. In P. G. Spirakis, editor, *Algorithms—ESA '95, Third Annual European Symposium*, volume 979 of *Lecture Notes in Computer Science*, pages 213–226, Corfu, Greece, 25–27 Sept. 1995. Springer. doi:10.1007/3-540-60313-1_145.

[30] M. Forster. *Crossings in Clustered Level Graphs.* PhD thesis, Universität Passau, 2004.
http://www.opus-bayern.de/uni-passau/volltexte/2005/48/.

[31] M. Forster and C. Bachmaier. Clustered level planarity. In P. van Emde Boas, J. Pokorný, M. Bieliková, and J. Stuller, editors, *SOFSEM*, volume 2932 of *Lecture Notes in Computer Science*, pages 218–228. Springer, 2004. doi:10.1007/978-3-540-24618-3_18.

[32] J. J. Fowler, C. Gutwenger, M. Jünger, P. Mutzel, and M. Schulz. An SPQR-tree approach to decide special cases of simultaneous embedding with fixed edges. In I. G. Tollis and M. Patrignani, editors, *Graph Drawing*, volume 5417 of *Lecture Notes in Computer Science*, pages 157–168. Springer Berlin Heidelberg, 2009. doi:10.1007/978-3-642-00219-9_16.

[33] R. Fulek, J. Kynčl, and D. Pálvölgyi. Efficient c-planarity testing algebraically, 2012. Unpublished manuscript.

[34] R. Fulek, M. Pelsmajer, M. Schaefer, and D. Štefankovič. Hanani-Tutte, monotone drawings, and level-planarity. In J. Pach, editor, *Thirty Essays on Geometric Graph Theory*, pages 263–287. Springer, 2012. doi:10.1007/978-1-4614-0110-0_14.

[35] R. Fulek, M. J. Pelsmajer, M. Schaefer, and D. Štefankovič. Adjacent crossings do matter. *J. Graph Algorithms Appl.*, 16(3):759–782, 2012. doi:10.7155/jgaa.00266.

[36] M. R. Garey and D. S. Johnson. *Computers and intractability.* W. H. Freeman and Co., San Francisco, Calif., 1979.

[37] A. Garg and R. Tamassia. On the computational complexity of upward and rectilinear planarity testing. *SIAM J. Comput.*, 31(2):601–625, 2001. doi:10.1137/S0097539794277123.

[38] E. Gassner, M. Jünger, M. Percan, M. Schaefer, and M. Schulz. Simultaneous graph embeddings with fixed edges. In F. V. Fomin, editor, *Graph-theoretic concepts in computer science*, volume 4271 of *Lecture Notes in Computer Science*, pages 325–335. Springer-Verlag, Berlin, 2006. doi:10.1007/11917496_29.

[39] H. H. Glover, J. P. Huneke, and C. S. Wang. 103 graphs that are irreducible for the projective plane. *J. Combin. Theory Ser. B*, 27(3):332–370, 1979. doi:10.1016/0095-8956(79)90022-4.

[40] C. Gutwenger, K. Klein, and P. Mutzel. Planarity testing and optimal edge insertion with embedding constraints. *J. Graph Algorithms Appl.*, 12(1):73–95, 2008. doi:10.7155/jgaa.00160.

[41] B. Haeupler, K. Jampani, and A. Lubiw. Testing simultaneous planarity when the common graph is 2-connected. In O. Cheong, K.-Y. Chwa, and K. Park, editors, *Proceedings of the 21st Annual International Symposium (ISAAC 2010) held in Jeju, December 15–17, 2010*, volume 6507 of *Lecture Notes in Computer Science*, pages 410–421, Berlin, 2010. Springer. doi:10.1007/978-3-642-17514-5_35.

[42] P. Hoffman and B. Richter. Embedding graphs in surfaces. *J. Combin. Theory Ser. B*, 36(1):65–84, 1984. doi:10.1016/0095-8956(84)90014-5.

[43] S.-H. Hong and H. Nagamochi. Two-page book embedding and clustered graph planarity. Technical Report Technical Report 2009-004, Kyoto University, 2009.

[44] J. Hopcroft and R. Tarjan. Efficient planarity testing. *J. Assoc. Comput. Mach.*, 21:549–568, 1974. doi:10.1145/321850.321852.

[45] D. Hoske. Book embedding with fixed page assignments. Bachelor thesis, Karlsruhe Institute of Technology, Karlsruhe, Germany, 2012.

[46] V. Jelínek, J. Kratochvíl, and I. Rutter. A Kuratowski-type theorem for planarity of partially embedded graphs. *Comput. Geom.*, 46(4):466–492, 2013. doi:10.1016/j.comgeo.2012.07.005.

[47] M. Jünger and S. Leipert. Level planar embedding in linear time. *J. Graph Algorithms Appl.*, 6(1):67–113, 2002. doi:10.7155/jgaa.00045.

[48] M. Jünger, S. Leipert, and P. Mutzel. Level planarity testing in linear time. In S. H. Whitesides, editor, *Proceedings of the 6th International Symposium (GD '98) held in Montréal, QC, August 13–15, 1998*, volume 1547 of *Lecture Notes in Computer Science*, pages 224–237. Springer-Verlag, Berlin, 1998. doi:10.1007/3-540-37623-2_17.

[49] M. Jünger and M. Schulz. Intersection graphs in simultaneous embedding with fixed edges. *J. Graph Algorithms Appl.*, 13(2):205–218, 2009. doi:10.7155/jgaa.00184.

[50] J. Kratochvíl. String graphs. II. Recognizing string graphs is NP-hard. *J. Combin. Theory Ser. B*, 52(1):67–78, 1991. doi:10.1016/0095-8956(91)90091-W.

[51] J. Kratochvíl and J. Matoušek. String graphs requiring exponential representations. *Journal of Combinatorial Theory, Series B*, 53:1–4, 1991.

[52] C. Kuratowski. Sur les problèmes des courbes gauches en Topologie. *Fund. Math.*, 15:271–283, 1930.

[53] B. Mohar. Projective planarity in linear time. *J. Algorithms*, 15(3):482–502, 1993. doi:10.1006/jagm.1993.1050.

[54] B. Mohar and C. Thomassen. *Graphs on surfaces.* Johns Hopkins Studies in the Mathematical Sciences. Johns Hopkins University Press, Baltimore, MD, 2001.

[55] T. Nishizeki and N. Chiba. *Planar graphs: theory and algorithms*, volume 140 of *North-Holland Mathematics Studies.* North-Holland Publishing Co., Amsterdam, 1988.

[56] J. Pach and G. Tóth. Monotone drawings of planar graphs. *J. Graph Theory*, 46(1):39–47, 2004. doi:10.1002/jgt.10168.

[57] J. Pach and G. Tóth. Monotone drawings of planar graphs. *ArXiv e-prints*, Jan. 2011. arXiv:1101.0967.

[58] J. Pach and G. Tóth. Monotone crossing number. In M. van Kreveld and B. Speckmann, editors, *Graph drawing*, volume 7034 of *Lecture Notes in Computer Science*, pages 278–289, Heidelberg, 2012. Springer. doi:10.1007/978-3-642-25878-7_27.

[59] J. Pach and R. Wenger. Embedding planar graphs at fixed vertex locations. *Graphs Combin.*, 17(4):717–728, 2001. doi:10.1007/PL00007258.

[60] M. Patrignani. On extending a partial straight-line drawing. *Internat. J. Found. Comput. Sci.*, 17(5):1061–1069, 2006. doi:10.1142/S0129054106004261.

[61] M. Patrignani. Planarity testing and embedding. In Tamassia [71], chapter 1, pages 1–42. To appear.

[62] M. J. Pelsmajer, M. Schaefer, and D. Stasi. Strong Hanani–Tutte on the projective plane. *SIAM Journal on Discrete Mathematics*, 23(3):1317–1323, 2009. doi:10.1137/08072485X.

[63] M. J. Pelsmajer, M. Schaefer, and D. Štefankovič. Removing even crossings. *J. Combin. Theory Ser. B*, 97(4):489–500, 2007. doi:10.1016/j.jctb.2006.08.001.

[64] M. J. Pelsmajer, M. Schaefer, and D. Štefankovič. Removing even crossings on surfaces. *European J. Combin.*, 30(7):1704–1717, 2009. doi:10.1016/j.ejc.2009.03.002.

[65] M. J. Pelsmajer, M. Schaefer, and D. Štefankovič. Removing independently even crossings. *SIAM Journal on Discrete Mathematics*, 24(2):379–393, 2010. doi:10.1137/090765729.

[66] I. Rutter. Personal communication, April 2013.

[67] M. Schaefer. Hanani-Tutte and related results. In I. Bárány, K. J. Böröczky, G. F. Tóth, and J. Pach, editors, *Geometry—Intuitive, Discrete, and Convex—A Tribute to László Fejes Tóth*, Bolyai Society Mathematical Studies, Berlin. Springer. To appear.

[68] M. Schaefer, E. Sedgwick, and D. Štefankovič. Recognizing string graphs in NP. *J. Comput. System Sci.*, 67(2):365–380, 2003. doi:10.1016/S0022-0000(03)00045-X.

[69] M. Schaefer and D. Štefankovič. Decidability of string graphs. *J. Comput. System Sci.*, 68(2):319–334, 2004. doi:10.1016/j.jcss.2003.07.002.

[70] J. Stillwell. *Classical topology and combinatorial group theory*, volume 72 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, second edition, 1993.

[71] R. Tamassia, editor. Discrete Mathematics and Its Applications. Chapman and Hall/CRC, 2013. To appear.

[72] C. Thomassen. Rectilinear drawings of graphs. *J. Graph Theory*, 12(3):335–341, 1988. doi:10.1002/jgt.3190120306.

[73] W. T. Tutte. Toward a theory of crossing numbers. *J. Combinatorial Theory*, 8:45–53, 1970. doi:10.1016/S0021-9800(70)80007-2.

[74] H. van der Holst. Algebraic characterizations of outerplanar and planar graphs. *European J. Combin.*, 28(8):2156–2166, 2007. doi:10.1016/j.ejc.2007.04.005.

[75] A. Wotzlaw, E. Speckenmeyer, and S. Porschen. Generalized k-ary tanglegrams on level graphs: A satisfiability-based approach and its evaluation. *Discrete Appl. Math.*, 160(16-17):2349–2363, 2012. doi:10.1016/j.dam.2012.05.028.

[76] W. J. Wu. On the planar imbedding of linear graphs. I. *J. Systems Sci. Math. Sci.*, 5(4):290–302, 1985.