Intersection Patterns of Edges in Topological Graphs

THÈSE N^{O} 5334 (2012)

PRÉSENTÉE LE 11 MAI 2012

À LA FACULTÉ DES SCIENCES DE BASE
CHAIRE DE GÉOMÉTRIE COMBINATOIRE
PROGRAMME DOCTORAL EN MATHÉMATIQUES

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

POUR L'OBTENTION DU GRADE DE DOCTEUR ÈS SCIENCES

PAR

Radoslav FULEK

acceptée sur proposition du jury:

Prof. K. Hess Bellwald, présidente du jury
Prof. J. Pach, directeur de thèse
Prof. S. Felsner, rapporteur
Prof. G. Tóth, rapporteur
Prof. E. Welzl, rapporteur

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Suisse
2012

Abstract

This thesis is devoted to crossing patterns of edges in topological graphs. We consider the following four problems:

1. A thrackle is a graph drawn in the plane such that every pair of edges meet exactly once: either at a common endpoint or in a proper crossing. Conway's Thrackle Conjecture says that a thrackle cannot have more edges than vertices. By a computational approach we improve the previously known upper bound of $1.5n$ on the maximal number of edges in a thrackle with n vertices to $1.428n$. Moreover, our method yields an algorithm with a finite running time that for any $\epsilon > 0$ either verifies the upper bound of $(1+\epsilon)n$ on the maximum number of edges in a thrackle or disproves the conjecture.
2. It is not hard to see that any simple graph admits a poly-line drawing in the plane such that each edge is represented by a polygonal curve with at most three bends, and each edge crossings realizes a prescribed angle α. We show that if we restrict the number of bends per edge to two and allow edges to cross in k different angles, a graph on n vertices admitting such a drawing can have at most $O(nk^2)$ edges. This generalizes a previous result of Arikushi et al., in which the authors treated a special case of our problem, where $k = 1$ and the prescribed angle has 90 degrees.
3. The classical result known as Hanani-Tutte Theorem states that a graph is planar if and only if it admits a drawing in the plane in which each pair of non-adjacent edges crosses an even number of times. We prove the following monotone variant of this result, conjectured by J.Pach and G.Tóth. If G has an x-monotone drawing in which every pair of independent edges crosses evenly, then G has an x-monotone embedding (i.e. a drawing without crossings) with the same vertex locations. We show several interesting algorithmic consequences of this result.
4. In a drawing of a graph, two edges form an odd pair if they cross each other an odd number of times. A pair of edges is independent (or non-adjacent) if they share no endpoint. For a graph G we let $\mathrm{ocr}(G)$ be the smallest number of odd pairs in a drawing of G and let $\mathrm{iocr}(G)$ be the smallest number of independent odd pairs in a drawing of G. We construct a graph G with $\mathrm{iocr}(G) < \mathrm{ocr}(G)$, answering a question by Székely, and—for the first time—giving evidence that crossings of adjacent edges may not always be trivial to eliminate.

Keywords. Topological graph, Graph drawing, Polyline drawing, Planar graph, Thrackle, Crossing number

Zusammenfassung

In dieser Dissertation werden Kreuzungsstrukturen von Kanten in topologischen Graphen untersucht. Insbesondere beschäftigen wir uns mit folgenden vier Problemen:

1. Unter einem Thrackle versteht man einen Graphen, der so in der Ebene gezeichnet werden kann, dass sich jedes Kantenpaar genau einmal trifft: Entweder in einem gemeinsamen Endpunkt, oder sie haben eine echte Kreuzung. Conway's Thrackle Hypothese besagt, dass ein Thrackle nicht mehr Kanten als Knoten haben kann. Mit einem rechnergestützten Ansatz verbessern wir die bisher beste bekannte obere Schranke an das Verhältnis von Kanten zu Knoten von 1.5 auf 1.428. Außerdem liefert unser Ansatz einen Algorithmus mit endlicher Laufzeit, welcher für beliebiges $\epsilon > 0$ entweder die obere Schranke von $1 + \epsilon$ an das Verhältnis verifiziert, oder Conway's Hypothese durch ein Gegenbeispiel widerlegt.
2. Es ist nicht schwer zu sehen dass jeder einfache Graph so in der Ebene gezeichnet werden kann, dass jede Kante als Polygonzug mit höchstens drei Krümmungen repräsentiert wird, bei der jede Kantenkreuzung in einem festen vorgegebenem Winkel α stattfindet. Wir zeigen dass, falls wir die Anzahl Krümmungen pro Kante auf zwei reduzieren, gleichzeitig aber k verschiedene Winkel für die Kantenkreuzungen zulassen, ein Graph der auf solche Art gezeichnet werden kann höchstens $O(nk^2)$ viele Kanten haben kann, wobei n die Anzahl an Knoten des Graphen ist. Dies verallgemeinert ein Resultat von Arikushi et al., in welchem die Autoren den Spezialfall behandeln wo $k = 1$ und der vorgegebene Winkel 90 Grad ist.
3. Ein klassisches Resultat bekannt als das Hanni-Tutte Theorem besagt, dass ein Graph planar ist genau dann wenn er so in der Ebene gezeichnet werden kann, dass für jedes Paar nicht-adjazenter Kanten die Anzahl an Kreuzungen gerade ist. Wir zeigen die folgende monotone Variante des Theorems, eine Hypothese von J.Pach and G.Tóth. Wenn ein Graph G eine x-monotone Zeichnung erlaubt in welcher die Anzahl an Kreuzungen jedes unabhängigen Kantenpaares gerade ist, dann besitzt G eine x-monotone Einbettung (also eine Zeichnung ohne Kreuzungen) mit denselben Knotenpositionen. Außerdem diskutieren wir mehrere interessante algorithmische Konsequenzen dieses Resultates.
4. In einer Zeichnung eines Graphen bilden zwei Kanten ein ungerades Paar, falls die Anzahl Kreuzungen dieses Paares (mit sich selbst) ungerade ist. Ein Kantenpaar ist unabhängig, falls es keinen gemeinsamen Knoten besitzt. Für einen Graphen G bezeichnet $\text{ocr}(G)$ die kleinste

Zahl an ungeraden Paaren, die eine Zeichnung von G haben kann. Analog bezeichnet $\mathrm{iocr}(G)$ die kleinste Zahl von ungeraden unabhängigen Paaren, die eine Zeichnung von G haben kann. Wir konstruieren einen Graphen G mit der Eigenschaft dass $\mathrm{iocr}(G) < \mathrm{ocr}(G)$. Dies beantwortet eine Frage von Székely und liefert —zum ersten Mal— einen Beleg dafür dass Kreuzungen von adjazenten Kanten nicht immer trivial zu eliminieren sind.

Schlagwörter Topologischen Graphen, Zeichnen von Graphen, Polylinie Zeichnung, Planaren Graphen, Thrackle, Kreuzungszahl

I would like to thank all those who directly or indirectly contributed to the creation of the thesis.

First and foremost, I would like to thank my advisor János Pach for giving me the opportunity to pursue research in his group at EPFL, nice discussions and for inviting many interesting and clever people to Lausanne, many of which became my friends and co-authors.

Equally, I would like to thank all of my co-authors that have contributed to my thesis. In particular, I would like to thank Csaba Tóth for exposing several weak points in my arguments, to Eyal Ackerman for joining our two-member team consisting of me and Csaba in writing our joint paper. I would like to thank Michael Pelsmajer for a nice fruitful collaboration that lead to our joint work with Marcus Schaefer and Daniel Štefankovič.

Special thanks go to Dömötör for sending me the source files of his thesis what allowed me to start writing mine, to Martin for helping me with the German version of the abstract, and to Filip for a valuable assistance at the last stage of the thesis preparation.

I am also grateful to the members and visitors of the dcg and the disopt group at EPFL, for creating such a wonderful working and social environment. In particular, I would like to thank Adrian, Andreas, Andres, Andrew, Balázs, Carsten, Dave, Dömötör, Fabrizio, Filip, Fritz, Gabriel, Genna, János, Jarek, Jocelyne, Judit, Laura, Martin, Marco, Nabil, Nicolai, Nicola, Padmini, Rom, Saurabh and Thomas.

Finally, I would like to thank to my mom, dad and brother for their emotional support, to Hanka for coming to Lausanne for some time, and to Juraj for many interesting mathematical discussions.

Contents

Part I
Introduction

1 Introduction and Organization

In this thesis we study drawings of graphs in the plane, where a graph is a system of two elements sets called edges over a finite set of vertices. In a drawing of a graph vertices are represented by points in the plane and edges are continuous arcs joining corresponding vertices.

The problem of Seven Bridges of Königsberg is widely regarded as the first graph theoretical question. The problem was stated and solved in the 18th century by Euler. However, it was only at the turn of the 19th and 20th century when graphs started to be systematically studied mainly due to the most prominent graph theoretical question at the time: The Four-color Problem. This problem is not just about graphs, but about graphs drawn in the plane without any edge crossings. A graph drawn in the plane (with possibly crossing edges) is usually called a topological graph* and it is the primary object of study in this thesis.

One of the first milestones in the investigation of drawings of graphs in the plane is the famous result of Kuratowski from 1934 that characterizes those graphs which admit a crossing free drawing in the plane, i.e. planar graphs. Nevertheless, this discovery had remained for many years also one of the last milestones in the area, as the mathematical questions about drawings of graphs attracted in the middle of the 20th century only a scant handful of researchers. In the 1960s, embeddings of graphs, i.e. drawings without edge crossings on surfaces, started becoming a more popular subject in mathematics, but the results from this era are mostly of somewhat different flavor to what we want to focus on in this thesis. The research on topological graphs became even more active in the 1980s, and since then its frontiers have been advancing in various research directions, many of them as parts of the theoretical computing science. In this thesis we also decided to approach graph drawing problems from a more algorithmic and computational side.

From the broader perspective, a drawing of a graph in the plane can be seen as a continuous mapping of a 1-dimensional topological space into the surface with genus 0. Problems regarding how a topological space "fits" into other topological space are mostly studied in the area of algebraic topology.

*Even though topological graph theory studies drawings of graphs in surfaces and spatial embeddings of graphs, in this thesis we are concerned only with drawings of graphs in the plane.

The area of mathematics motivated at the initial stage by intuitive questions that took a turn towards the development of the machinery of ever-increasing complexity. Even though algebraic topology offers a lot of the powerful tools for tackling problems involving complicated topological spaces of any dimension, still seemingly simple and intuitively explainable topological problems about drawing a graph in the plane remain unsolved.

One of these problems Conway's Thrackle Conjecture, which is treated in this thesis, was asked more than 40 years ago by John Conway. A thrackle is a graph drawn in the plane such that every pair of edges meet exactly once: either at a common endpoint or in a proper crossing. The conjecture says that a thrackle cannot have more edges than vertices. Since then some progress on the conjecture was done by topologists as well as by non-topologists. A variant of the conjecture that suits better the available methods was first resolved by means of algebraic topology, but later a more intuitive proof was found.

By using an intuitive approach and computational power of a computer we show that the number of edges in a thrackle on n vertices can be at most $1.42n$ thereby improving the best previously known upper bound of $1.5n$ by Cairns and Nikolayevsky [10]. Moreover, our method can in principle lead to an upper bound of $(1+\epsilon)n$ for any $\epsilon > 0$ or a refutation of the conjecture. Our main result in this area can be stated as follows.

Theorem 1.1. *For any $\epsilon > 0$ there exists a finite collection of graphs $\mathcal{G}$ such that if none of the graphs in $\mathcal{G}$ can be drawn as a thrackle, the maximum number of edges in a thrackle on n vertices is at most $(1+\epsilon)n$.*

Beside its esthetic qualities, the conjecture is believed to represent the tip of an "iceberg," obstructing our understanding of crossing patterns of edges in topological graphs.

The motivation for the other problem we study comes from the area of graph visualizations. Here, we still want to represent a graph in the plane, but its edges are restricted to be poly-line segments (i.e. polygonal curves) that may cross only at certain angles. Since every graph admits a drawing in which every edge is a poly-line segment with at most three bends and all crossing pairs of edges cross at the same angle, we restrict the maximum number of bends per edge to two. We show that a graph admitting a poly-line drawing where all crossing pairs of edges cross at one of the finitely many angles cannot have more than linearly many edges and is thus quite sparse. More precisely, we proved the following.

Theorem 1.2. *A graph on n vertices that can be drawn in the plane so that*

(i) each edge is represented by a poly-line segment with at most two bends;

(ii) no two edges cross at a bend;

(iii) and each edge crossing realizes one of k prescribed angles,

can have at most $O(k^2 n)$ edges.

The previous theorem generalizes the collection of previous results about drawings of graphs with right angle crossings [6, 17], and along the way we study more general, not necessarily poly-line, drawings of graphs about which we obtain results that might be of independent interest.

The two above problems can be also regarded as Turán-type problems for topological graphs. In a Turán-type problem we ask for the maximal density of a combinatorial object, in our case a topological graph, not containing certain forbidden configuration. Thus, in case of Thrackle Conjecture our combinatorial object is a topological graph, whose any pair of edges meets at most once and the forbidden configuration is a pair of disjoint non-crossing edges. For the second problem the combinatorial object is a poly-line drawing of a graph with at most two bends per edge and the forbidden configuration is a pair of edges crossing in an unprescribed angle.

A great deal of research on graph drawing is about drawings which do not contain any edge crossings, or in other words about planar drawings. The point of departure for the second part of this thesis is an algebraic condition for planarity proved for the first time by Hanani in 1930s [26] and reproved by Tutte in 1970 [73]. The condition says that a graph is planar, if we can draw it in the plane so that every pair of edges, not sharing a vertex, crosses an even number of times. We prove a variant of this results, in which edges are drawn as x-monotone curves and the x-coordinates of the vertices are fixed.

Theorem 1.3. *If G has an x-monotone drawing in which every pair of independent edges crosses evenly, then G has an x-monotone embedding with the same vertex locations.*

This theorem answers in affirmative a question asked by Pach and Tóth [50], who proved the weaker result, in which we require that every pair of edges crosses an even number of times.

One can have doubts about significance of such an improvement. However, we show interesting algorithmic consequences of our result, which do not follow from the weaker version. In particular, the theorem gives an algorithm for level-planarity testing and layout. An algorithm which appears to be the first practical one for the problem with the complete proof of its correctness.

The most common concept for measuring non-planarity of a graph is perhaps the crossing number, which is defined as the minimum number of edge crossings in a drawing of the graph. Several other natural variants of the crossing number were also introduced and systematically studied, for example the pair crossing number is defined to be the minimum number of pairs of crossing edges in a drawing of the graph. Many of these variants were proved to be related to the actual crossing number. However, not much is known about how closely they are related and if some of them happen to coincide. Whether the pair crossing number and crossing number is the same for all graphs is perhaps the most tantalizing question about crossing numbers that has not been so far, to the best of our knowledge, resolved.

It is very easy to see that in any drawing of a graph minimizing the crossing number the adjacent edges cannot cross. It is possible that it was this observation, which lead Tutte in his paper from 1970 [73] about algebraic crossing numbers to claim "We are taking the view that crossings of adjacent edges are trivial and easily got rid of." Thirty-four years later Székely [70] commented on this sentence as follows: "We interpret this sentence as a philosophical view and not a mathematical claim."

We conclude the thesis with a theorem that refutes Tutte's intuition about crossings of edges that share an endpoint. Our result is about the odd crossing number, which is defined as the minimum number of pairs of edges crossing an odd number of times in a drawing of the graph. It used to be an open problem whether the odd and pair crossing number give actually for every graph the same value. Finally, in 2005 this was resolved in negative by Pelsmajer et al. [53].

For each variant of a crossing number we define its independent version, in which we ignore crossings between adjacent edges. We show that there exist graphs for which the odd (resp. algebraic) crossing number and the independent odd (resp. algebraic) crossing number are not the same, which is the first result indicating that counting or not-counting crossings between edges sharing a vertex might make a difference.

Theorem 1.4. *For every $n > 0$, there exists a graph G having the independent odd (resp. algebraic) crossing number less than the odd (resp. algebraic) crossing number by more than n.*

The thesis is organized as follows. Section 2 is reserved for the definitions and auxiliary claims used throughout the thesis. The central part of the thesis contains four section each of which is devoted to one of the problems described above. Section 3 contains material from the author's joint work with János Pach [22], and is devoted to the upper bound on the maximum number of edges in a thrackle. In Section 4 we reproduce the result from [2] about poly-line drawings of graphs with edges crossing in prescribed angles. The leading motive of Chapter 5 is the strong version of Hanani-Tutte theorem for monotone drawings of graphs proved in [24]. In Chapter 6 we continue along similar lines as in Chapter 5 and separate odd and independently odd crossing number, which was done in [23]. The third part of the thesis contains open problems and conjectures, bibliography and CV.

2 Definitions and Preliminaries

Definition 2.1. *By a* graph $G = (V, E)$ *we understand a pair consisting of a finite set of* vertices V *and a system of two-elements sets* E *over* V *called* edges.

The *degree* of a vertex $v \in V$ is defined as the number of edges of G that v is contained in. Throughout the thesis we assume that E is not a multi-set and that no edge in E contains one vertex twice unless stated otherwise. Thus, G will be usually a *simple* graph.

We say that a graph $G = (V, E)$ is *bipartite* if its vertex set can be partitioned into two parts A and B; $A \cup B = V$ such that $E \subseteq A \times B$. For the other graph notions used in the thesis we point the reader to the book of R. Diestel [18].

Definition 2.2. *A topological graph (resp. a drawing of a graph) is a representation of the graph in the plane such that*

(i) its vertices are represented by distinct points;

(ii) its edges are represented by Jordan arcs connecting its corresponding endpoints;

(iii) the arcs representing edges can meet either at a common vertex are in a proper crossing

If it leads to no confusion, we make no notational distinction between a drawing and the underlying abstract graph G. In the same vein, $V(G)$ and $E(G)$ will stand for the vertex set and edge set of G as well as for the sets of points and curves representing them.

Then a graph is *planar* if it can be drawn in the plane without any edge crossing. For a non-planar graph G it makes sense to ask what is the minimum number of crossings in a drawing of G. This brings us to the notion of a *crossing number* $cr(G)$ of a graph G, which is defined as the minimum number of crossings in a topological graph whose underlying graph is G.

A fundamental result about the crossing number is so called Crossing Lemma first proved by Ajtai et al. [4] and Leighton [41] which gives a lower bound on this graph parameter. The current best version is by Pach et al. [47].

Lemma 2.1 (Crossing Lemma).

$$cr(G) \geq \frac{0.032|E|^3}{|V|^2} - 1.06|V|$$

If $|E| \geq \frac{103}{16}|V|$, *then*

$$cr(G) \geq \frac{0.032|E|^3}{|V|^2}$$

We will use this result in Section 4 to prove an upper bound on the number of edges of a graph admitting a specific poly-line drawing.

In Section 6 we introduce and study other versions of crossing numbers in which instead of the total number of edge crossings we count pairs of edges that cross.

In this thesis we prove an analogue of the following classical result by Hanani [26] and Tutte [73].

Theorem 2.2 (**Hanani-Tutte** (1934-1970))**.** *If a graph can be drawn in the plane such that every pair of independent edges crosses an even number of times, then it is planar.*

We will often use the following claim, intuitively quite obvious, which is called Jordan Curve Theorem in the literature (see e.g. [46]).

Theorem 2.3 (**Jordan Curve Theorem**)**.** *A complement of a continuous closed non-self-intersecting curve in the plane consists of exactly two (topologically) connected parts.*

The proof of the variant of Hanani-Tutte Theorem is based on the following well-known fact.

Lemma 2.4. *Connected components of the complement of a closed continuous curve in the plane can be two-colored so that no two components sharing a non-trivial part of the boundary receive the same color.*

Part II
Results

3 Thrackle

In the present section we will investigate the maximum number of edges in a graph admitting a special type of drawing in which all the edges intersect. The material from this section was published in [22].

A drawing of G is called a *thrackle* if every pair of edges meet precisely once, either at a common vertex or at a proper crossing. (A crossing p of two curves is *proper* if at p one curve passes from one side of the other curve to its other side.) More than *forty* years ago Conway [63, 8, 66] conjectured that every thrackle has at most as many edges as vertices, and offered a bottle of beer for a solution. Since then the prize went up to a thousand dollars. In spite of considerable efforts, Conway's thrackle conjecture is still open. It is believed to represent the tip of an "iceberg," obstructing our understanding of crossing patterns of edges in topological graphs. If true, Conway's conjecture would be tight as any cycle of length at least *five* can be drawn as a thrackle, see [76]. Two thrackle drawings of C_5 and C_6 are shown in Figure 1.

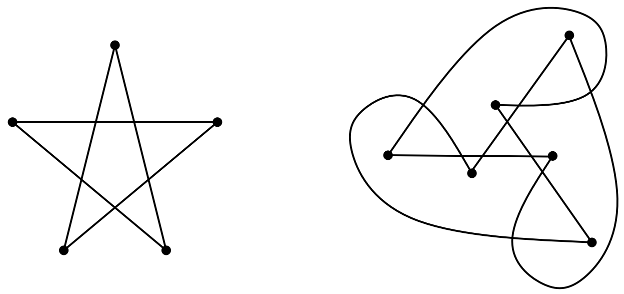

Figure 1: C_5 and C_6 drawn as thrackles

Obviously, the property that G can be drawn as a thrackle is *hereditary*: if G has this property, then any subgraph of G does. It is very easy to verify (cf. [76]) that C_4, a cycle of length *four*, cannot be drawn in a thrackle. Therefore, every "thrackleable" graph is C_4-free, and it follows from extremal graph theory that every thrackle of n vertices has at most $O(n^{3/2})$ edges [18]. The first linear upper bound on the maximum number of edges of a thrackle of n vertices was given by Lovász et al. [43]. This was improved to a $\frac{3}{2}(n-1)$ by Cairns and Nikolayevsky [10].

We provide a finite approximation scheme for estimating the maximum number of edges that a thrackle of n vertices can have. We apply our technique to improve the best known upper bound for this maximum.

To state our results, we need a definition. Given three integers $c', c'' > 2$, $l \geq 0$, the *dumbbell* $\mathrm{DB}(c', c'', l)$ is a simple graph consisting of two disjoint cycles of length c' and c'', connected by a path of length l. For $l = 0$, the two cycles share a vertex. It is natural to extend this definition to negative values of l, as follows. For any $l > -\min(c', c'')$, let $\mathrm{DB}(c', c'', l)$ denote the graph consisting of two cycles of lengths c' and c'' that share a path of length $-l$. That is, for any $l > -\min(c', c'')$, we have

$$|V(\mathrm{DB}(c', c'', l))| = c' + c'' + l - 1.$$

The three types of dumbbells (for $l < 0$, $l = 0$, and $l > 0$) are illustrated in Figure 2.

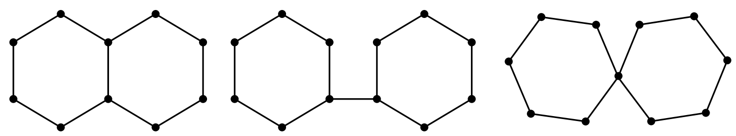

Figure 2: Dumbbells $\mathrm{DB}(6, 6, -1), \mathrm{DB}(6, 6, 1)$, and $\mathrm{DB}(6, 6, 0)$

Our first theorem shows that for any $\varepsilon > 0$, it is possible to prove Conway's conjecture up to a multiplicative factor of $1 + \varepsilon$, by verifying that no dumbbell smaller than a certain size depending on ε is thrackleable.

Theorem 3.1. *Let $c \geq 6$ and $l \geq -1$ be two integers, such that c is even, with the property that no dumbbell $\mathrm{DB}(c', c'', l')$ with $-c'/2 \leq l' \leq l$ and with even $6 \leq c', c'' \leq c$ can be drawn in the plane as a thrackle. Let $r = \lfloor l/2 \rfloor$. Then the maximum number of edges $t(n)$ that a thrackle on n vertices can have satisfies $t(n) \leq \tau(c, l)n$, where*

$$\tau(c, l) = \begin{cases} \frac{47c^2+116c+80}{35c^2+68c+32} & \text{if } l = -1, \\ 1 + \frac{2c^2r+4cr^2+22cr+7c^2+22c+8r^2+24r+16}{2c^2r^2+14c^2r+4cr^2+16cr+24c^2+12c} & \text{if } l \geq 0, \end{cases}$$

as n tends to infinity.

As both c and l get larger, the constant $\tau(c, l)$ given by the second part of Theorem 3.1 approaches 1. On the other hand, assuming that Conway's conjecture is true for all bipartite graphs with up to 10 vertices, which will be verified in Section 3.3, the first part of the theorem applied with $c = 6, l = -1$

yields that $t(n) \le \frac{617}{425}n < 1.452n$. This bound is already better than the bound $\frac{3}{2}n$ established in [10].

By a more careful application of Theorem 3.1, i.e. taking $c = 6$ and $l = 0$, we obtain an even stronger result.

Theorem 3.2. *The maximum number of edges $t(n)$ that a thrackle on n vertices can have satisfies the inequality $t(n) \le \frac{167}{117}n < 1.428n$.*

Our method is algorithmic. We design an $e^{O((1/\varepsilon^2)\ln(1/\varepsilon))}$ time algorithm to prove, for any $\varepsilon > 0$, that $t(n) \le (1+\varepsilon)n$ for all n, or to exhibit a counterexample to Conway's conjecture. The proof of Theorem 3.2 is computer assisted: it requires testing the planarity of certain relatively small graphs. As the upper bound on the maximum number of edges in a thrackle was used to obtain various density-type results, our improvement yields a better bound in several other claims, see e.g. [1, 21].

For thrackles drawn by straight-line edges, Conway's conjecture had been settled in a slightly different form by Hopf and Pannwitz [34] and by Sutherland [68] before Conway was even born, and later, in the above form, by Erdős and Perles. Assuming that Conway's conjecture is true, Woodall [76] gave a complete characterization of all graphs that can be drawn as a thrackle. He also observed that it would be sufficient to verify the conjecture for dumbbells. This observation is one of the basic ideas behind our arguments.

A *generalized thrackle* is a drawing of a graph in the plane with the property that any pair of edges share an odd number of points at which they properly cross or which are their common endpoints. Obviously, every thrackle is a generalized thrackle but not vice versa: although C_4 is not thrackleable, it can be drawn as a generalized thrackle, which is not so hard to see. The corresponding question about the maximum number of edges in a generalized thrackle was completely resolved in [10]. Several interesting special cases and variants of the conjecture are discussed in [9, 11, 25, 43, 60, 61].

In Section 3.1, we describe a crucial construction of Conway and summarize some earlier results needed for our arguments. The proofs of Theorems 3.1 and 3.2 are given in Sections 3.2 and 3.3. The analysis of the algorithm for establishing the upper bound of $(1+\varepsilon)n$ is also given in Section 3.3 (Theorem 3.7). In the last section of the present chapter, we discuss some related Turán-type extremal problems for planar graphs.

3.1 Conway's doubling and preliminaries

In this section, we review some earlier results that play a key role in our arguments.

We need the following simple observation.

Lemma 3.3. [43] *A (generalized) thrackle cannot contain two vertex disjoint odd cycles.*

Proof. Let G denote a thrackle. Two vertex disjoint odd cycles in G have to cross each other an even number of times by Lemma 2.4. On the other hand, since G is a thrackle, two vertex disjoint odd cycles in G have to cross each other an even number of times. $\square$

Lovász, Pach, and Szegedy [43] gave a somewhat counterintuitive characterization of generalized thrackles containing no odd cycle: a bipartite graph is a generalized thrackle if and only if it is *planar*. Moreover, it follows immediately from Lemma 3 and the proof of Theorem 3 in Cairns and Nikolayevsky [10] that this statement can be strengthened as follows.

Lemma 3.4. [10] *Let G be a bipartite graph with vertex set $V(G) = A \cup B$ and edge set $E(G) \subseteq A \times B$. If G is a generalized thrackle then it can be redrawn in the plane without crossing so that the cyclic order of the edges around any vertex $v \in V(G)$ is preserved if $v \in A$ and reversed if $v \in B$.*

We recall a construction of Conway for transforming a thrackle into another one. It can be used to eliminate odd cycles.

Let G be a thrackle or a generalized thrackle that contains an *odd* cycle C. In the literature, the following procedure is referred to as *Conway's doubling*: First, delete from G all edges incident to at least one vertex belonging to C, including all edges of C. Replace every vertex v of C by two nearby vertices, v_1 and v_2. For any edge vv' of C, connect v_1 to v_2' and v_2 to v_1' by two edges running very close to the original edge vv', as depicted in Figure 3. For any vertex v belonging to C, the set of edges incident to v but not belonging to C can be divided into two classes, $E_1(v)$ and $E_2(v)$: the sets of all edges whose initial arcs around v lie on one side or the other side of C. In the resulting topological graph G', connect all edges in $E_1(v)$ to v_1 and all edges in $E_2(v)$ to v_2 so that every edge connected to v_1 crosses all edges connected to v_2 exactly once in their small neighborhood. See Figure 3. All other edges of G remain unchanged. Denote the vertices of the original odd cycle C by $v^1, v^2, \ldots, v^k$, in this order. In the resulting drawing G', we obtain an *even* cycle $C' = v_1^1 v_2^2 v_1^3 v_2^4 \ldots v_2^1 v_1^2 v_2^3 v_1^4 \ldots$ instead of C. It is easy to verify that G' is drawn as a thrackle, which is stated as part (ii) of the following lemma (see also Lemma 2 in [10]).

Lemma 3.5. (Conway, [76, 10]) *Let G be a (generalized) thrackle with at least one odd cycle C. Then the topological graph G' obtained from G by Conway's doubling of C is*

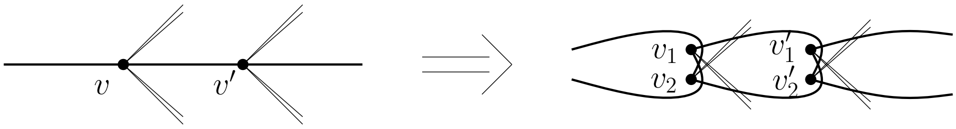

Figure 3: Conway's doubling of a cycle

(i) bipartite, and
(ii) a (generalized) thrackle.

Proof. It remains to verify part (i). Let k denote the length of the (odd) cycle $C \subseteq G$, and let C' stand for the doubled cycle in G'. The length of C' is $2k$. Let π denote the inverse of the doubling transformation. That is, π identifies the opposite pairs of vertices in C', and takes C' into C.

Suppose for a contradiction that G' is not bipartite. In view of Lemma 3.3, no odd cycle of G' is disjoint from C'. Let D' be an odd cycle in G' with the smallest number of edges that do not belong to C'. We can assume that D' is the union of two paths, P_1 and P_2, connecting the same pair of vertices u, v in C', where P_1 belongs to C' and P_2 has no interior points on C'.

If $\pi(u) \neq \pi(v)$, that is, the length of P_1 is not 0 or k, then $\pi(D') = \pi(P_1) \cup \pi(P_2)$ is a simple cycle in G. Notice that the lengths of P_1 and P_2 have different parities. If the length of P_1 is even, say, then, according to the rules of doubling, the initial and final pieces of P_2 in small neighborhoods of u and v are on the same side of the (arbitrarily oriented) cycle C'. Consequently, the initial and final pieces of $\pi(P_2)$ in small neighborhoods of $\pi(u)$ and $\pi(v)$ are on the *same* side of C. On the other hand, using the fact that G is a generalized thrackle, the total number of intersection points between the odd path $\pi(P_2)$ and the odd cycle C is odd (see the proof of Lemma 2.2 from [43]). Thus, if we two color the regions of the plane bounded by pieces of C, so that any pair of neighboring regions receive different colors, the initial and final pieces of $\pi(P_2)$ in small neighborhoods of $\pi(u)$ and $\pi(v)$ must lie in the regions colored with different colors. Since C is odd and drawn as a generalized thrackle, it follows that the initial and final pieces of $\pi(P_2)$ in small neighborhoods of $\pi(u)$ and $\pi(v)$ must lie on *different* sides of C, a contradiction.

The cases when P is odd and when $\pi(u) = \pi(v)$ can be treated analogously. □

Finally, we recall an observation of Woodall [76] mentioned in the introduction, which motivated our investigations.

As thrackleability is a hereditary property, a minimal counterexample to the thrackle conjecture must be a connected graph G with exactly $|V(G)|+1$

edges and with no vertex of degree *one*. Such a graph G is necessarily a dumbbell $\mathrm{DB}(c', c'', l)$. If $l \neq 0$, then G consists of two cycles that share a path or are connected by a path uv. In both cases, we can "double" the path uv, as indicated in Figure 4, to obtain another thrackle G'. It is easy to see that G' is a dumbbell consisting of two cycles that share precisely one vertex (the vertex v in the figure). Moreover, if any of these two cycles is not even, then we can double it and repeat the above procedure, if necessary, to obtain a dumbbell $\mathrm{DB}(b', b'', 0)$ drawn as a thrackle, where b' and b'' are even numbers.

Thus, in order to prove the thrackle conjecture, it is enough to show that no dumbbell $\mathrm{DB}(c', c'', 0)$ consisting of two even cycles that share a vertex is thrackleable.

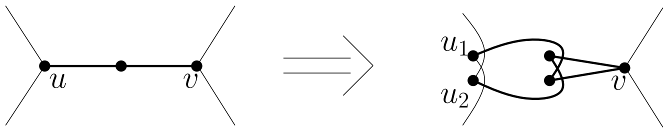

Figure 4: Doubling the path uv

3.2 Proof of Theorem 3.1

Let $c \geq 6$ and $l \geq -1$ be two integers, and suppose that no dumbbell $\mathrm{DB}(c', c'', l')$ with $-c'/2 \leq l' \leq l$ and with even $6 \leq c', c'' \leq c$ can be drawn in the plane as a thrackle. For simpler notation, let $r = \lfloor l/2 \rfloor$.

Let $G = (V, E)$ be a thrackleable graph with n vertices and m edges. We assume without loss of generality that G is connected and that it has no vertex of degree *one*. Otherwise, we can successively delete all vertices of degree *one*, and argue for each connected component of the resulting graph separately.

As usual, we call a graph *two-connected* if it is connected and it has no *cut vertex*, i.e., it cannot be separated into two or more parts by the removal of a vertex [18].

We distinguish three cases:

(A) G is bipartite;

(B) G is not bipartite, and the graph G' obtained by performing Conway's doubling of a shortest odd cycle $C \subset G$ is 2-connected;

(C) G is not bipartite, and the graph G' obtained by performing Conway's doubling of a shortest odd cycle $C \subset G$ is not 2-connected.

In each case, we will prove that $m \leq \tau(c,l)n$.

(A) By Lemma 3.4, in this case G is planar. We fix an embedding of G in the plane. According to the assumption of our theorem, G contains no subgraph that is a dumbbell $\mathrm{DB}(c', c'', l')$, for any even $6 \leq c' \leq c'' \leq c$, and $-c'/2 \leq l' \leq l$. We also know that G has no C_4. We are going to use these conditions to bound the number of edges $m = |E(G)|$.

Notice that we also exclude dumbbells $\mathrm{DB}(c', c'', l')$ with $-c' \leq l' < -c'/2$. Indeed, in this case $\mathrm{DB}(c', c'', l')$ is isomorphic to $\mathrm{DB}(c', d, k)$, where $d = (c' + c'' + 2l')$, $k = (-c' - l')$, and $d < c'' \leq c$, $\max(-c'/2, -d/2) \leq k < 0$.

Suppose first that G is *two-connected*. Let f denote the number of faces, and let f_c stand for the number of faces with at most c sides. By double counting the edges, we obtain

$$2m \geq 6f_c + (c+2)(f - f_c). \tag{1}$$

If $l = -1$, then applying the condition on forbidden dumbbells, we obtain that no two faces of size at most c share an edge, so that $6f_c \leq m$. If $l \geq 0$, Menger's theorem (see e.g. [18], Sec 2.) implies that any two faces of size at most c are connected by two vertex disjoint paths. Since any such path must be longer than l, to each face we can assign its vertices as well as the $r = \lfloor l/2 \rfloor$ closest vertices along two vertex disjoint paths leaving the face, and these sets are disjoint for distinct faces. Thus, we have $f_c(2r+6) \leq n$. In either case, we have

$$f_c \leq \begin{cases} \frac{m}{6} & \text{if } l = -1, \\ \frac{n}{2r+6} & \text{if } l \geq 0. \end{cases} \tag{2}$$

Combining the last two inequalities, we obtain

$$f \leq \frac{(c-4)f_c + 2m}{c+2} \leq \begin{cases} \frac{(c-4)\frac{m}{6} + 2m}{c+2} & \text{if } l = -1, \\ \frac{(c-4)\frac{n}{2r+6} + 2m}{c+2} & \text{if } l \geq 0. \end{cases}$$

In view of Euler's polyhedral formula $m + 2 = n + f$, which yields

$$m \leq \begin{cases} \frac{6c+12}{5c+4}n - \frac{12c+24}{5c+4} & \text{if } l = -1, \\ \frac{2cr+4r+7c+8}{2cr+6c}n - \frac{2c+4}{c} & \text{if } l \geq 0. \end{cases} \tag{3}$$

It can be shown by routine calculations that the last estimates, even if we ignore their negative terms independent of n, are stronger than the ones claimed in the theorem. (In fact, they are also stronger than the corresponding bounds (5) and (4) in Case (B); see below.) This concludes the proof of the case (A) when G is 2-connected.

If G is not 2-connected, then consider a block decomposition of G, and proceed by the induction on the number of blocks. The base case, i.e when G is 2-connected, is treated above. Otherwise G can be obtained as a union of two bipartite graphs $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ sharing exactly one vertex. By induction hypothesis we can use (3) to bound the number of edges in G_i, for $i = 1, 2$, by substituting $|E_i|$ and $|V_i|$ for m and n, respectively. We obtain the claimed bound on the maximum number of edges in G by adding up the bounds on $|E_1|$ and $|E_2|$ as follows.

$$|E(G)| = |E(G_1)|+|E(G_2)| \leq k_1|V(G_1)|+k_1|V(G_2)|-2k_2 = k_1|V(G)|+k_1-2k_2$$

where $k_1 = k_1(c, l)$ and $k_2 = k_2(c, l)$ represent the constants in (3). Induction goes through, because $k_1 < k_2$ for all considered values of c and l.

(B) In this case, we establish two upper bounds on the maximum number of edges in G: one that decreases with the length of the shortest odd cycle $C \subseteq G$ and one that increases. Finally, we balance between these two bounds.

By doubling a shortest odd cycle $C \subseteq G$, as before, we obtain a bipartite thrackle G' (see Lemma 3.5). Let C' denote the doubled cycle in G'. By Lemma 3.4, G' is a two-colorable planar graph. Moreover, it can be embedded in the plane without crossing so that the cyclic order of the edges around each vertex in one color class is preserved, and for each vertex in the other color class reversed. A closer inspection of the way how we double C shows that as we traverse C' in G', the edges incident to C' start on alternating sides of C'. This implies that, after redrawing G' as a plane graph, all edges incident to C' lie on one side, that is, C' is a *face*.

Slightly abusing the notation, from now on let G' denote a crossing-free drawing with the above property, which has a $2|C|$-sided face C'. Denoting the number of vertices and edges of G' by n' and m', the number of faces and the number of faces of size at most c by f' and f'_c, respectively, we have $n' = n + |C| = |V(G')|$, $m' = m + |C| = |E(G')|$, and, as in Case (A), inequality (2),

$$f'_c \leq \begin{cases} \frac{1}{6}m' & \text{if } l = -1, \\ \\ \frac{n'}{2r+6} & \text{if } l \geq 0. \end{cases}$$

Double counting the edges of G', we obtain

$$2m' \geq 6f'_c + (c+2)(f' - 1 - f'_c) + 2|C|.$$

In case $l \geq 0$, combining the last two inequalities, we have

$$f' \leq \frac{(c-4)f'_c + 2(m' - |C|) + c + 2}{c+2} \leq \frac{(c-4)\frac{n'}{2r+6} + 2(m' - |C|) + c + 2}{c+2}.$$

By Euler's polyhedral formula, $f' = m' - n' + 2$. Thus, after ignoring the negative term, which depends only on c and l, the last inequality yields

$$|E(G)| \leq \frac{2cr + 4r + 7c + 8}{2cr + 6c} n + |C| \frac{c-4}{2cr + 6c}. \tag{4}$$

The case $l = -1$ can be treated analogously, and the corresponding bound on $E(G)$ becomes

$$|E(G)| \leq \frac{6c + 12}{5c + 4} n + |C| \frac{c-4}{5c+4}. \tag{5}$$

We now establish another upper bound on the number of edges in G: one that decreases with the length of the shortest odd cycle C in G. As in [43], we remove from G the vertices of C together with all edges incident to them. Let G'' denote the resulting thrackle. By Lemma 3.3, G'' is bipartite. By Lemma 3.4, it is a planar graph. From now on, let G'' denote a fixed (crossing-free) embedding of this graph. According to our assumptions, G'' has no subgraph isomorphic to $\mathrm{DB}(c', c'', l')$, for any even numbers c' and c'' with $6 \leq c' \leq c'' \leq c$, and for any integer l' with $-c'/2 \leq l' \leq l$.

We can bound $|E(G'')|$, as follows. By the minimality of C, each vertex $v \in V(G)$ that does not belong to C is joined by an edge of G to at most one vertex on C. Indeed, otherwise, v would create either a C_4 or an odd cycle shorter than C. Hence, if $l \geq 0$, inequality (3) implies that

$$|E(G)| \leq |E(G'')| + |C| + (n - |C|) \leq \frac{2cr + 4r + 7c + 8}{2cr + 6c}(n - |C|) + n. \tag{6}$$

In the case $l = -1$, we obtain

$$|E(G)| \leq |E(G'')| + |C| + (n - |C|) \leq \frac{6c + 12}{5c + 4}(n - |C|) + n. \tag{7}$$

It remains to compare the above upper bounds on $|E(G)|$ and to optimize over the value of $|C|$. If $l > -1$, then the value of $|C|$ for which the right-hand sides of (4) and (6) coincide is

$$|C| = \frac{cr + 3c}{cr + 2r + 4c + 2} n.$$

The claimed bound follows by plugging this value into (4) or (6).

In the case $l = -1$, the critical value of $|C|$, obtained by comparing the bounds (5) and (7), is

$$|C| = \frac{5c+4}{7c+8}n.$$

Plugging this value into (5) or (7), the claimed bound follows.

(C) As before, let C be a shortest odd cycle in G, and let G' be the graph obtained from G after doubling C. The doubled cycle is denoted by $C' \subset G'$. Let $G_0 \supseteq C$ denote a *maximal* subgraph of G, which is turned into a *two*-connected subgraph of G' after performing Conway's doubling on C. Let G_1 stand for the graph obtained from G by the removal of all *edges* in G_0.

It is easy to see that G_1 is bipartite, and each of its connected components shares exactly one vertex with G_0. Indeed, if a connected component $G_2 \subseteq G_1$ were not bipartite, then, by Lemma 3.3, G_2 would share at least one vertex with C, which belongs to an odd cycle of G_2. By the maximal choice of G_0, after doubling C, the component G_2 must turn into a subgraph $G_2' \subset G'$, which shares precisely *one* vertex with the doubled cycle C'. Thus, G_2 must also share precisely *one* vertex with C, which implies that $G_2' \subseteq G'$ has an odd cycle. This contradicts Lemma 3.5(i), according to which G' is a bipartite graph.

Therefore, G_1 is the union of all blocks of G, which are not entirely contained in G_0. Since each connected component G_2 of G_1 is bipartite, the number of edges of G_2 can be bounded from above by (3), just like in Case (A).

In order to bound the number of edges of G, we proceed by adding the connected components of G_1 to G_0, one by one. As at the end of the discussion of Case (3.2), using the fact that the last terms in (3), which do not depend on n, are smaller than -2, we can complete the proof by induction on the number of connected components of G_1.

3.3 A better upper bound

As was pointed out before, if we manage to prove that for any l', $-3 \le l' \le -1$, the dumbbell $\mathrm{DB}(6,6,l')$ is not thrackleable, then Theorem 3.1 yields that the maximum number of edges that a thrackle on n vertices can have is at most $\frac{617}{425}n < 1.452n$. This estimate is already better than the currently best known upper bound $\frac{3}{2}n$ due to Cairns and Nikolayevsky [10].

In order to secure this improvement, we have to exclude the subgraphs $\mathrm{DB}(6,6,-1)$, $\mathrm{DB}(6,6,-2)$, and $\mathrm{DB}(6,6,-3)$. The fact that $\mathrm{DB}(6,6,-3)$

cannot be drawn as a thrackle was proved in [43] (Theorem 5.1). Here we present an algorithm that can be used for checking whether a "reasonably" small graph G can be drawn as a thrackle. We applied our algorithm to verify that DB$(6, 6, -1)$ and DB$(6, 6, -2)$ are indeed not thrackleable. In addition, we show that DB$(6, 6, 0)$ cannot be drawn as thrackle, which leads to the improved bound in Theorem 3.2.

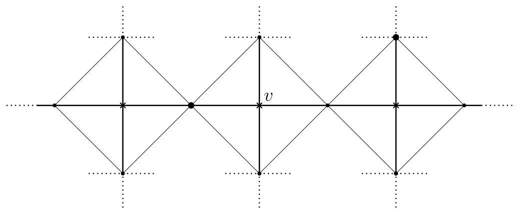

Figure 5: 4-cycle around a vertex v of G', which was a crossing point in G

Let $G = (V, E)$ be a thrackle. Direct the edges of G arbitrarily. For any $e \in E$, let $E_e \subseteq E$ denote the set of all edges of G that do not share a vertex with e, and let $m(e) = |E_e|$. Let $\pi_e = (\pi_e(1), \pi_e(2), \ldots, \pi_e(m(e))$ stand for the $m(e)$-tuple (permutation) of all edges belonging to E_e, listed in the order of their crossings along e.

Construct a planar graph G' from G, by introducing a new vertex at each crossing between a pair of edges of G, and replacing each edge by its pieces. In order to avoid G' having an embedding in which two paths corresponding to a crossing pair of edges of G do not *properly* cross, but touch each other, we introduce a new vertex in the interior of every edge of G', whose both endpoints are former crossings. For each former crossing point v, we add a cycle of length *four* to G', connecting its neighbors in their cyclic order around v, as illustrated in Figure 5. In the figure, the thicker lines and points represent edges and vertices or crossings of G, while the thinner lines and points depict the *four*-cycles added at the second stage.

Obviously, G' is completely determined by the directed abstract underlying graph of G and by the set of permutations $\Pi(G) := \{\pi_e \in E_e^{m(e)} | e \in E\}$. Thus, a graph $G = (V, E)$ can be drawn as a thrackle if and only if there exists a set Π of $|E|$ permutations of $E_e, e \in E$, such that the abstract graph G' corresponding to the pair (G, Π) is *planar*. In other words, to decide whether a given abstract graph $G = (V, E)$ can be drawn as a thrackle, it is enough to consider all possible sets of permutations Π of $E_e, e \in E$, and to check if the corresponding graph $G' = G'(G, \Pi)$ is planar for at least one of them. The first deterministic linear time algorithm for testing planarity was found by Hopcroft and Tarjan [33]. However, in our implementation we used an improved algorithm for planarity testing by Fraysseix et al. [14], in particular, its implementation in the library P.I.G.A.L.E. [13]. We leave the pseudocode of our routine for the abstract. The source code can be found here : http://dcg.epfl.ch/webdav/site/dcg/users/183292/public/Thrackle.zip.

It was shown in [43] (Lemma 5.2) that in every drawing of a directed cycle C_6 as a thrackle, either every oriented path $e_1e_2e_3e_4$ is drawn in such a way that $\pi_{e_1} = (e_4, e_3)$ and $\pi_{e_4} = (e_1, e_2)$, or every oriented path $e_1e_2e_3e_4$ is drawn in such a way that $\pi_{e_1} = (e_3, e_4)$ and $\pi_{e_4} = (e_2, e_1)$. Using this observation (which is not crucial, but saves computational time), we ran a backtracking algorithm to rule out the existence of a set of permutations Π, for which $G'(\mathrm{DB}(6,6,0),\Pi)$, $G'(\mathrm{DB}(6,6,-1),\Pi)$, or $G'(\mathrm{DB}(6,6,-2),\Pi)$ is planar. Our algorithm attempts to construct larger and larger parts of a potentially good set Π, and at each step it verifies if the corresponding graph still has a chance to be extended to a planar graph. In the case of $\mathrm{DB}(6,6,0)$, to speed up the computation, we exploit Lemma 2.2 from [43].

Summarizing, we have the following

Lemma 3.6. *None of the dumbbells* $\mathrm{DB}(6,6,l')$, $-3 \le l' \le 0$ *can be drawn as a thrackle.*

According to Lemma 3.6, Theorem 3.1 can be applied with $c = 6, l = 0$, and Theorem 3.2 follows.

For any $\varepsilon > 0$, our Theorem 3.1 and the above observations provide a deterministic algorithm with bounded running time to prove that all thrackles with n vertices have at most $(1+\varepsilon)n$ edges or to exhibit a counterexample to Conway's conjecture.

In what follows, we estimate the dependence of the running time of our algorithm on ε. The analysis uses the standard random access machine model. In particular, we assume that all basic arithmetic operations can be carried out in constant time.

Theorem 3.7. *For any* $\varepsilon > 0$*, there is a deterministic algorithm with running time* $e^{O((1/\varepsilon^2)\ln(1/\varepsilon))}$ *to prove that all thrackles with* n *vertices have at most* $(1+\varepsilon)n$ *edges or to exhibit a counterexample to Conway's conjecture.*

Proof. First we estimate how long it takes for a given c and l, satisfying the assumptions in Theorem 3.1, to check whether there exists a dumbbell $\mathrm{DB}(c', c'', l')$ with c' and c'' even, $6 \le c' \le c'' \le c$, and with $-c'/2 \le l' \le l$, that can be drawn as a thrackle. Clearly, there are

$$\sum_{\substack{c'=6 \\ c' \text{ is even}}}^{c} \frac{(\frac{c'}{2}+l+1)(c-c'+2)}{2} = \frac{1}{8}lc^2+\frac{1}{48}c^3-\frac{3}{4}lc+l+\frac{1}{4}c^2-\frac{25}{12}c+3 \le \kappa(lc^2+c^3)$$

dumbbells to check, for some $\kappa > 0$. In order to decide, whether a fixed dumbbell with m edges can be drawn as a thrackle, we construct at most $(m-2)!^m$ graphs, each with at most $O(m^2)$ edges, and we test each of them for planarity. Thus, the total running time of our algorithm is $O((lc^2 + c^3)(2c + l - 2)!^{2c+l}(2c+l)^2)$. Approximating the factorials by Stirling's formula, we can conclude that the running time is $O((2c+l)^{(2c+l)^2+\frac{1}{2}(2c+l)+5}e^{-(2c+l)})$.

Now, for any $1 > \epsilon > 0$ we show how big values of l and c we have to take so that Theorem 3.1 gives the upper bound $(1+\epsilon)n$ on the maximum number of edges in a thrackle. We remind the reader that $r = \lfloor l/2 \rfloor$. It can be shown by routine calculation that there are three constants κ, κ_r and κ_c so that the following holds. Given $\epsilon > 0$, for $r = \lceil \frac{\kappa_r}{\varepsilon} \rceil$, and c such that

$$c \geq \frac{\kappa_c}{\varepsilon} \geq \frac{\kappa r^2}{\epsilon(2r^2 + 14r + 24) - 2r - 7}$$

the value of $\tau(c,l)$ introduced in Theorem 3.1 is at most $1+\varepsilon$. For the sake of completeness we give the sufficient condition for c only in terms of r and ϵ:

$$\begin{aligned} c \quad \geq \quad & \frac{r^2(2-2\epsilon) + r(11-8\epsilon) + 11 - 6\epsilon}{\epsilon(2r^2+14r+24) - 2r - 7} + \\ & + \frac{(r+3)\sqrt{(r^2(4+8\epsilon+4\epsilon^2) + r(4+36\epsilon+8\epsilon^2) + 1 + 28\epsilon + 4\epsilon^2)}}{\epsilon(2r^2+14r+24) - 2r - 7} \end{aligned}$$

Thus, for these values of c and r Theorem 3.1 gives the required bound, i.e. at most $(1+\epsilon)n$. Plugging $\frac{\kappa_c}{\varepsilon}$ and $2\frac{\kappa_r}{\varepsilon}$ as c and l, respectively, in $O((2c+l)^{(2c+l)^2+\frac{1}{2}(2c+l)+5}e^{-(2c+l)})$, the theorem follows. $\square$

4 Poly-line drawings and angle crossings

Even though the title of this section suggests that we will restrict our attention to the poly-line drawings of graphs, i.e. drawings of graphs in which edges are represented by poly-line segments, the main result is proved by relaxing geometric conditions. Thus, we will study topological graphs, whose edges are not necessarily poly-line segments and can be two-colored so that certain properties are satisfied (Theorems 4.3 and 4.4). The material from this section was published in [2].

Graphs that admit polyline drawings with few bends per edge and such that every crossing occurs at a large angle have received some attention lately, since cognitive experiments [35, 36] indicate that such drawings are almost as readable as planar drawings. That is, one can easily track the edges in such drawings, even though some edges may cross.

A *polyline drawing* of a graph G is a topological graph where each edge is drawn as a simple polygonal arc between the incident vertices but not passing through any bend point of other arcs. In a polyline drawing, every crossing occurs in the relative interior of two segments of the two polygonal arcs, and so they have a well-defined crossing angle in $(0, \frac{\pi}{2}]$.

Didimo *et al.* [17] introduced *right angle crossing* (*RAC*) drawings, which are polyline drawings where all crossings occur at right angle. They proved that a graph with $n \geq 3$ vertices that admits a straight line RAC drawing has at most $4n - 10$ edges, and this bound is the best possible. A different proof of the same upper bound was later found by Dujmović *et al.* [19]. It is not hard to show that every graph admits a RAC drawing with three bends per edge (see Figure 6 for an example). Arikushi *et al.* [6] have recently proved, improving previous results by Didimo *et al.* [17], that if a graph with n vertices admits a RAC drawing with at most two bends per edge, then it has $O(n)$ edges.

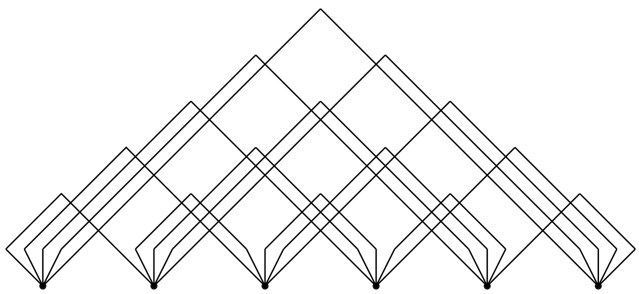

Figure 6: A RAC drawing of K_6 with 3 bends per edge.

Dujmović *et al.* [19] generalized RAC drawings, allowing crossings at a range of angles rather than at right angle. They considered αAC drawings, which are polyline drawings where every crossing occurs at some angle *at least* α. They showed that a graph with n vertices and a straight line αAC drawing has at most $\frac{\pi}{\alpha}(3n-6)$ edges, by partitioning the graph into $\frac{\pi}{\alpha}$ planar graphs. They also proved that their bounds are essentially optimal for $\alpha = \frac{\pi}{k} - \varepsilon$, with $k = 2, 3, 4, 6$ and sufficiently small $\varepsilon > 0$.

We first consider polyline drawings where every crossing occurs at the same angle $\alpha \in (0, \frac{\pi}{2}]$. An $\alpha AC_b^=$ drawing of a graph is a polyline drawing where every edge is a polygonal arc with at most b bends and every crossing occurs at angle *exactly* α. If $b = \infty$, the number of bends on an edge can be arbitrarily big, but finite. It is easy to see that every graph with $n > 2$ vertices that admits an $\alpha AC_0^=$ drawing has at most $3(3n-6)$ edges (see Lemma 4.5 below). Every graph admits an $\alpha AC_3^=$ drawing for every $\alpha \in (0, \frac{\pi}{2}]$: Didimo *et al.* [17] constructed a RAC drawing of the complete graph with three bends per edge (see also Figure 6), where every crossing occurs between a pair of orthogonal segments of the same orientation, so an affine transformation deforms all crossing angles uniformly. It remains to consider graphs that admit $\alpha AC_1^=$ or $\alpha AC_2^=$ drawings. We prove the following.

Theorem 4.1. *For every $\alpha \in (0, \frac{\pi}{2}]$, a graph on n vertices that admits an $\alpha AC_2^=$ drawing has $O(n)$ edges. Specifically, a graph on n vertices has*

(a) *at most $27n$ edges if it admits an $\alpha AC_1^=$ drawing; and*
(b) *at most $383.7n$ edges if it admits an $\alpha AC_2^=$ drawing.*

It is not difficult to draw hexagonal tiling of the plane with all diagonals inside each hexagon as $\frac{\pi}{3}AC_1^=$ drawing. Hence, the part (a) of the previous theorem cannot be improved below $6n$. On the other hand, in [6] it was shown that for every sufficiently large n there exists a graph on n vertices with $7.83n$ edges admitting a RAC drawing with two bends per edge. Thus, the part (b) of the previous theorem cannot be improved below $7.83n$.

For $\alpha = \frac{\pi}{2}$, slightly better bounds have been derived by Arikushi *et al.* [6]: they proved that if a graph on n vertices admits a RAC drawing with at most one (resp., two) bends per edge, then it has at most $6.5n$ (resp., $74.2n$) edges. Their proof techniques, however, do not generalize to all $\alpha \in (0, \frac{\pi}{2}]$.

A straightforward generalization of $\alpha AC_1^=$ and $\alpha AC_2^=$ drawings are polyline drawings where each crossing occurs at an angle from a list of k distinct angles.

Theorem 4.2. *Let $A \subset (0, \frac{\pi}{2}]$ be a set of k angles, $k \in \mathbb{N}$, and let G be a graph on n vertices that admits a polyline drawing with at most b bends per edge such that every crossing occurs at some angle from A. Then,*

(a) *G has $O(kn)$ edges if $b = 1$;*

(b) *G has $O(k^2 n)$ edges if $b = 2$.*

We remark that the bound in the part (a) of the previous theorem is tight. The corresponding lower bound is discussed in Section 4.9. We believe that the part (b) can be still improved to $O(kn)$.

Suppose that every edge in a topological graph is partitioned into edge segments, such that all crossings occur in the relative interior of the segments. The bends in polyline drawings, for example, naturally define such edge partitions. An *end segment* is an edge segment incident to a vertex of the edge, while a *middle* segment is an edge segment not incident to any vertex. The key idea in proving Theorems 4.1 and 4.2 is to consider the crossings that involve either two end segments, or an end segment and a middle segment. This idea extends to topological graphs whose edge segments satisfy a few properties, which automatically hold for polyline drawings with same angle crossings (perhaps after removing a constant fraction of the edges). We obtain the following results, which might be of independent interest.

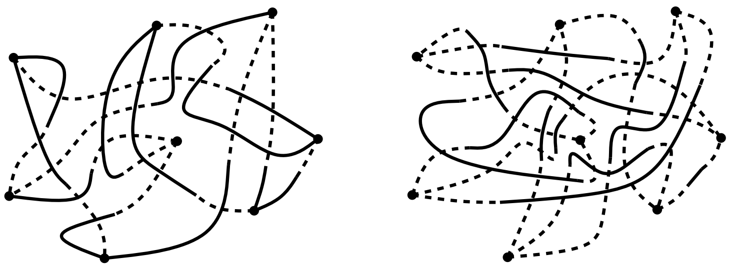

(a) (b)

Figure 7: (a) A 3-regular topological graph satisfying the conditions of Theorem 4.3 for $k = 2$; blue segments are dashed. (b) A 3-regular topological graph satisfying the conditions of Theorem 4.4 for $k = 2$.

Theorem 4.3. *Let $G = (V, E)$ be a topological graph on n vertices, in which every edge can be partitioned into two end segments, one colored red and the other colored blue, such that (see Figure 7(a))*

(1) *no two end segments of the same color cross;*

(2) *every pair of end segments intersects at most once; and*

(3) *no blue end segment is crossed by more than k red end segments that share a vertex.*

Then G has $O(kn)$ edges.

We show that the above theorem implies the following stronger result.

Theorem 4.4. *Let $G = (V, E)$ be a topological graph on n vertices. Suppose that every edge of G can be partitioned into two end segments and one middle segment such that (Figure 7(b))*

(1) *each crossing involves one end segment and one middle segment;*

(2) *each middle segment and end segment intersect at most once; and*

(3) *each middle segment crosses at most k end segments that share a vertex.*

Then G has $O(kn)$ edges.

We remark that the bound in Theorem 4.3 and 4.4 is tight, which follows from the tightness of the part (a) of the Theorem 4.2.

Note that Theorem 4.4 implies Theorem 4.3. Indeed, given a graph that satisfies the constraints in Theorem 4.3, one can partition every edge e into three parts as follows: its two end segments are the red segment and a crossing-free portion of the blue segment incident to a vertex, while the rest of the blue segment is the middle segment of e. Such a partition clearly satisfies the constraints in Theorem 4.4 with the same parameter k.

We begin with a few preliminary observations in Section 4.1. In Section 4.2, we consider polyline drawings with one possible crossing angle and prove Theorem 4.1. Then we extend the proof of Theorem 4.1(a) allowing up to k possible crossing angles and prove Theorem 4.2(a). We also show that Theorem 4.1(b) can be generalized to a weaker version of Theorem 4.2(b) with an upper bound of $O(k^4 n)$ (rather than $O(k^2 n)$). In Section 4.5, we generalize the crossing conditions from angle constraints to colored segments in topological graphs, and prove Theorems 4.3 and 4.4. Theorem 4.2(b) is derived from these general results at the end of Section 4.5. We conclude with some lower bound constructions and open problems in Section 4.9.

4.1 Preliminaries

In a polyline drawing of a graph, the edges are simple polygonal paths, consisting of line segments. We start with a few initial observations about line segments and polygonal paths. We say that two line segments *cross* if their relative interiors intersect in a single point. (In our terminology, intersecting segments that share an endpoint or are collinear do not cross.)

The following lemma is about the crossing pattern of line segments: if any two crossing segments cross at the same angle $\alpha \in (0, \frac{\pi}{2}]$, then a constant fraction of the segments are pairwise non-crossing. This lemma will be instrumental when applied to specific edge segments of an $\alpha AC_\infty^=$ drawing D: if we mark on each edge in an $\alpha AC_\infty^=$ drawing D one segment the lemma allows us to partition the set of edges in D into at most three parts so that no two marked segments of two edges in one part cross.

Lemma 4.5. *Let $\alpha \in (0, \frac{\pi}{2}]$ and let S be a finite set of line segments in the plane such that any two segments may cross only at angle α. Then S can be partitioned into at most three subsets of pairwise noncrossing segments. Moreover, if $\frac{\pi}{\alpha}$ is irrational or if $\frac{\pi}{\alpha} = \frac{q}{p}$, where $\frac{q}{p}$ is irreducible and q is even, then S can be partitioned into at most two subsets of pairwise noncrossing segments. It is easy to see that if q is odd three colors might be necessary.*

Proof. Partition S into maximal subsets of pairwise parallel line segments. Let $\mathcal{S}$ denote the subsets of S. We define a graph $G_{\mathcal{S}} = (\mathcal{S}, E_{\mathcal{S}})$, in which two subsets $S_1, S_2 \in \mathcal{S}$ are joined by an edge if and only if their respective directions differ by angle α. Clearly, the maximum degree of a vertex in $G_{\mathcal{S}}$ is at most two, and so $G_{\mathcal{S}}$ is 3-colorable. In any proper 3-coloring of $G_{\mathcal{S}}$, the union of each color class is a set of pairwise noncrossing segments in S, since they do not meet at angle α.

If $\frac{\pi}{\alpha}$ is irrational, then $G_{\mathcal{S}}$ is cycle-free. If $\frac{\pi}{\alpha} = \frac{p}{q}$, where $\frac{p}{q}$ is irreducible and q is even, then $G_{\mathcal{S}}$ can only have even cycles. In both cases, $G_{\mathcal{S}}$ is 2-colorable, and S has a partition into two subsets of pairwise noncrossing segments. □

The first claim in Lemma 4.5 can easily be generalized to finite sets of crossing angles [7].

Lemma 4.6 ([7])**.** *Let $A \subset (0, \frac{\pi}{2}]$ be a set of k angles, $k \in \mathbb{N}$, and let S be a finite set of line segments in the plane such that any two segments may cross only at an angle in A. Then S can be partitioned into at most $2k+1$ subsets of pairwise noncrossing segments.*

Proof. Partition S into maximal subsets of pairwise parallel line segments. Let $\mathcal{S}$ denote the subsets of S. We define a graph $G_{\mathcal{S}} = (\mathcal{S}, E_{\mathcal{S}})$, in which two subsets $S_1, S_2 \in \mathcal{S}$ are joined by an edge if and only if their respective directions differ by an angle in A. Clearly, the maximum degree of a vertex in $G_{\mathcal{S}}$ is at most $2k$, and so $G_{\mathcal{S}}$ is $(2k+1)$-colorable. In any proper $(2k+1)$-coloring of $G_{\mathcal{S}}$, the union of each color class is a set of pairwise noncrossing segments in S, since they do not meet at an angle in A. □

In the proof of Theorem 4.1, rather than counting the edges in an graph with an $\alpha AC_1^=$ (resp., $\alpha AC_2^=$) drawing, we estimate the number of edges in an auxiliary multigraph, called a *red* graph. The edges of the red graph closely follow the edges of the $\alpha AC_1^=$ (resp., $\alpha AC_2^=$) drawing, and each bend lies at a crossing point. This ensures that the red graph has a polyline drawing where the angle between any two consecutive segments of an edge is exactly α. Since we direct the red edges, it will be necessary to distinguish between counterclockwise angles α and clockwise angles $-\alpha$.

Consider a simple open polygonal path $\gamma = (v_0, v_1, \ldots, v_n)$ in the plane. Refer to Figure 8(a). At every interior vertex v_i, $1 \leq i \leq n-1$, the *turning angle* $\angle(\gamma, v_i)$ is the directed angle in $(-\pi, \pi)$ (the counterclockwise direction is positive) from ray $\overrightarrow{v_{i-1}v_i}$ to $\overrightarrow{v_i v_{i+1}}$. The *turning angle* of the polygonal path γ is the sum of turning angles over all interior vertices $\sum_{i=1}^{n-1} \angle(\gamma, v_i)$. We say that two line segments have a *common tail* if they share an endpoint and one of them is contained in the other (*e.g.*, segments pu_1 and pv_1 have a common tail in Figure 8(c)).

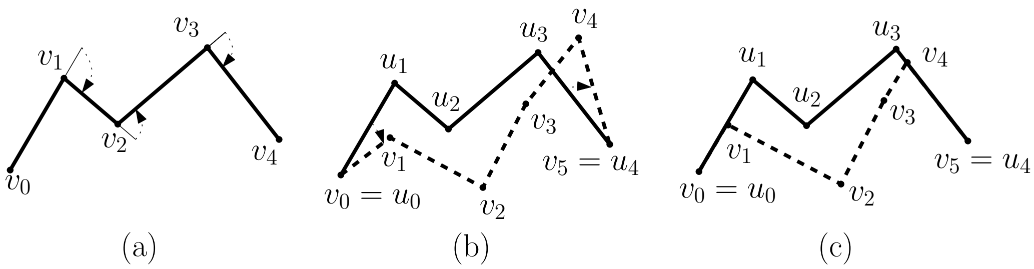

Figure 8: (a) The turning angles of a polygonal path. (b) Two crossing polygonal paths with the same turning angle between p and q. (c) Two noncrossing polygonal paths with the same turning angle between p and q.

In the proof of Theorem 4.1 instead of bounding directly the number of edges in an $\alpha AC_1^=$ (resp. $\alpha AC_2^=$) drawing we bound the number of edges in a polyline crossing free drawing of a multigraph, in which the set of turning angles of the edges has a constant size. We will use the next lemma to bound the maximal multiplicity of an edge in such multigraph.

Lemma 4.7. *Let p and q be two points in the plane. Let γ_1 and γ_2 be two directed simple polygonal paths from p to q. If γ_1 and γ_2 have the same turning angle and they do not cross, then the first segment of γ_1 shares a common tail with the first segment of γ_2 and the last segment of γ_1 shares a common tail with the last segment of γ_2.*

Proof. Let $\gamma_1 = (u_0, u_1, \ldots, u_m)$ and $\gamma_2 = (v_0, v_1, \ldots, v_n)$, with $p = u_0 = v_0$ and $q = u_m = v_n$. Let β be their common turning angle. Since γ_1 and γ_2 do not cross, they enclose a weakly simple polygon P (i.e. we allow line segments to overlap) with $m+n$ vertices (Figure 8(c)). Suppose w.l.o.g. that the vertices of P in clockwise order are $v_0 = u_0, u_1, \ldots, u_m = v_n, v_{n-1} \ldots, v_1$. Every interior angle of P is in $[0, 2\pi]$, and the sum of interior angles is $(m+n-2)\pi$, since P is an $(m+n)$-gon. The sum of interior angles at the vertices

$u_1, \ldots, u_{m-1}$ is $(m-1) \cdot \pi + \beta$; and the sum of interior angles at $v_1, \ldots, v_{n-1}$ is $(n-1) \cdot \pi - \beta$. Hence the interior angles at p and q are both 0. $\square$

Let G be a topological multigraph. We say that two edges *overlap* if their intersection contains a connected set of more than one point. A maximal connected component of the intersection of two edges is called an *overlap* of the two edges. A *common tail* is an overlap of two edges that contains a common endpoint of the two edges. In Sections 4.2 and 4.5, we construct topological multigraphs whose edges may overlap, but only in common tails.

Lemma 4.8. *Let G be a topological multigraph in which some edges may overlap, but only in common tails. Then the edges of G can be slightly perturbed such that all overlaps are removed and no new crossings are introduced.*

Proof. We successively perturb G and decrease the number of edge pairs that have a common tail. Let $e = (u, v)$ be an edge in G, and let $e_1, e_2, \ldots, e_k$ be edges in G that have a common tail with e such that their overlaps with e contain the vertex u. Direct all these edges away from u. Then every edge e_i, $i = 1, 2, \ldots, k$, follows an initial portion of e, and then turns either right or left at some *turning point* p_i. Assume without loss of generality that there is at least one right turning point, and let p_j be the last such point. (Observe that the common tails of e incident to u and v, respectively, are disjoint, otherwise two edges that share common tails with e would overlap in an arc that is not a common tail. It follows that the part of e between u and p_j is disjoint from any common tail incident to v.)

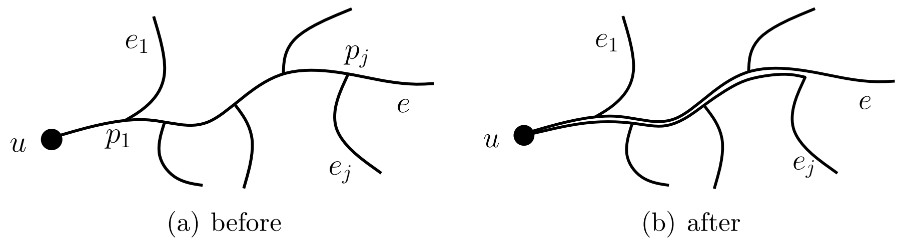

Figure 9: Removing overlaps.

Redraw all the edges e_i with a right turning point such that they closely follow e on the right. See Figure 9. We have removed the overlap between e and e_j, and decreased the number of edge pairs that have a common tail. $\square$

In the sequel we will use the following upper bound (Theorem 4.9) for the maximum number of edges in a simple quasi-planar graph by Ackerman and Tardos [3]. A topological graph is *simple* if any two of its edges meet at most once, either at a common endpoint or at a crossing. A topological graph is *quasi-planar* if it has no three pairwise crossing edges.

Theorem 4.9 ([3]). *A simple quasi-planar graph on $n \geq 4$ vertices has at most $6.5n - 20$ edges.*

4.2 Polyline drawing with one crossing angle

In this section, we prove Theorem 4.1. Our proof technique can be summarized as follows. Consider an $\alpha AC_\infty^=$ drawing D of a graph $G = (V, E)$, where each edge has an arbitrary number of edge segments, and any two edges cross at angle α. For a constant fraction of the edges $(u, v) \in E$, we draw a new directed "red" edge that connects u to another vertex in V (which is not necessarily v). The red edges follow some edges in D, and they only turn at edge crossings of D. Some of the red edges may be parallel (even though G is a simple graph), but none of them is a loop, and some of them may have a common tail. The vertex set V and the red edges form a topological multigraph, which we call the "red graph." Every edge in the red graph is a polyline where the turning angles at each bend is $\pm\alpha$ or $\pm(\pi - \alpha)$. The multiplicity of the red edges can be bounded using Lemma 4.7. By Lemma 4.5, a constant fraction of the red edges form a crossing-free multigraph, and overlaps can be removed using Lemma 4.8. We continue with the details.

An $\alpha AC_\infty^=$ drawing of a graph G is a polyline drawing with an arbitrary number of bends where every crossing occurs at angle α. Every edge is a polygonal arc that consists of line segments. The first and last segments of each edge are called *end segments*, all other segments are called *middle segments*. Note that each end segment is incident to a vertex of G. Let $G = (V, E)$ be a graph with an $\alpha AC_\infty^=$ drawing. It is clear that G has at most $3n - 6$ crossing-free edges, since they form a plane graph. All other edges have some crossings. We distinguish several cases below depending on whether the edges have crossings along their end segments.

4.3 Crossings between end segments

Lemma 4.10. *Let $\alpha \in (0, \frac{\pi}{2}]$ and $G = (V, E)$ be a graph on $n \geq 4$ vertices that admits an $\alpha AC_\infty^=$ drawing such that an end segment of every edge $e \in E$ crosses an end segment of some other edge in E. Then $|E| \leq 36n$. Moreover, the number of edges in E whose both end segments cross some end segments is at most $18n$.*

Proof. Let D be an $\alpha AC_{\infty}^{=}$ drawing of G as above. Let S be the set of end segments that cross some other end segments in D. We have $|E| \leq |S| \leq 2|E|$. Direct each segment $s \in S$ from an incident vertex in V to the other endpoint (which is either a bend point or another vertex in V). For a straight line edge, choose the direction arbitrarily.

We construct a directed multigraph $G' = (V, \Gamma)$. We call the edges in Γ *red*, to distinguish them from the edges of E. For every end segment $s \in S$, we construct a red edge $\gamma(s)$, which is a polygonal path with one bend between two vertices in V. For a segment $s \in S$, the path $\gamma(s)$ is constructed as follows (refer to Figure 10).

Let $u_s \in V$ denote the starting point of s (along its direction). Let c_s be the first crossing of s with an end segment, which we denote by t_s. Let $v_s \in V$ be a vertex incident to the end segment t_s. Now let $\gamma(s) = (u_s, c_s, v_s)$.

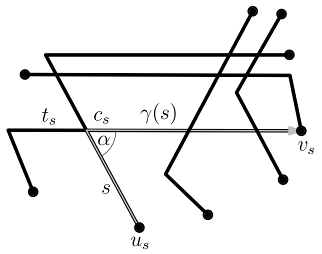

Figure 10: Construction of a red edge $\gamma(s) = (u_s, c_s, v_s)$.

Note that for every $s \in S$, the first segment of $\gamma(s)$ is part of the segment s and does not cross any segment in S. Hence the first segments of the red edges $\gamma(s)$ are distinct and do not cross other red edges. However, the second segment of $\gamma(s)$ may cross other red edges. Since the edges of G cross at angle α and c_s is a crossing, the turning angle of $\gamma(s)$ is $\pm\alpha$ or $\pm(\pi - \alpha)$. Note also that red edges may have common tails (which can be removed using Lemma 4.8).

We show that for any two vertices $u, v \in V$, there are at most 4 directed red edges from u to v. The red edges from u to v cannot cross, since their first segments are crossing-free, and their second segments are all incident to the same point v. By Lemma 4.7, any two noncrossing paths of the same turning angle between u and v must overlap in the first and last segments, however, the first segments of the red edges are pairwise non-overlapping. Since the red edges may have up to 4 distinct turning angles, there are at most 4 red edges from u to v.

We distinguish two types of red edges. Let $\Gamma_1 \subseteq \Gamma$ be the set of red edges whose second segment crosses some other red edge, and let $\Gamma_2 = \Gamma \setminus \Gamma_1$ be the set of red edges where both segments are crossing-free.

Note that two edges in Γ_1 cannot follow the same path γ in opposite directions because the first segments of every red edge is crossing-free. Hence, there are at most 4 red edges in Γ_1 between any two vertices in V. Let S_1 be the set of second segments of the red edges in Γ_1. By Lemma 4.5, there is a subset $S_1' \subseteq S_1$ of pairwise noncrossing segments of size at least $\frac{1}{3}|\Gamma_1|$. Let Γ_1' be the set of red edges containing the segments S_1', with $|\Gamma_1'| \geq \frac{1}{3}|\Gamma_1|$.

If Γ_2 contains two edges that follow the same path γ in opposite directions, then pick one arbitrarily, and let $\Gamma_2' \subseteq \Gamma_2$ be the selected red edges, with $|\Gamma_2'| \geq \frac{1}{2}|\Gamma_2|$. Now $(V, \Gamma_1' \cup \Gamma_2')$ is a crossing-free multigraph with maximum multiplicity 4, with at most $4(3n-6)$ edges. Note that any overlap between red end segments can be removed using Lemma 4.8, and so $(V, \Gamma_1' \cup \Gamma_2')$ becomes a planar multigraph with maximum multiplicity 4. It follows that $|\Gamma_1' \cup \Gamma_2'| \leq 4(3n-6)$, hence $|\Gamma| \leq 3 \cdot 4(3n-6) = 36n - 72$ for $n \geq 3$.

For the last part of the statement observe that in the above argument, an edge in E is counted twice if both of its end segments are in S. □

We are now ready to prove part (a) of Theorem 4.1.

Lemma 4.11. *For any angle $\alpha \in (0, \frac{\pi}{2}]$, a graph on n vertices that admits an $\alpha AC_1^=$ drawing has at most $27n$ edges.*

Proof. Let $G = (V, E)$ be a graph with $n \geq 4$ vertices drawn in the plane with an $\alpha AC_1^=$ drawing. Let $E_1 \subseteq E$ denote the set of edges in E that have at least one crossing-free end segment. Let $G_1 = (V, E_1)$ and $G_2 = (V, E \setminus E_1)$.

It is easy to see that if $\alpha \neq \frac{\pi}{3}$, then G_1 is a simple quasi-planar graph and so it has at most $6.5n - 20$ edges by Theorem 4.9. If $\alpha = \frac{\pi}{3}$, let S_1 be the set of crossed end segments of edges in E_1. By Lemma 4.5, there is a subset $S_1' \subseteq S_1$ of pairwise noncrossing segments of size $\frac{1}{3}|E_1|$. The graph G_1' corresponding to these edges is planar, with at most $3n-6$ edges. Hence E_1 contains at most $3 \cdot (3n-6) = 9n - 18$ edges.

By Lemma 4.10 applied to the subgraph of G containing only crossing edges, G_2 has at most $18n$ edges. Hence, G has at most $24.5n$ edges if $\alpha \neq \frac{\pi}{3}$ and at most $27n$ edges otherwise. □

Remark. It is easy to generalize the proof of Lemma 4.10 to the case that every two polyline edges cross at one of k possible angles. The only difference is that the red edges may have up to $2k$ different turning angles.

Lemma 4.12. *Let $G = (V, E)$ be a graph on $n \geq 4$ vertices that admits a polyline drawing such that an end segment of every edge $e \in E$ crosses an end segment of some other edge in E at one of k possible angles. Then $|E| \leq 36kn$. Moreover, the number of edges in E whose both end segments cross some end segments is at most $18kn$.* □

Corollary 4.13 (Theorem 4.2(a))**.** *Let $A \subset (0, \frac{\pi}{2}]$ be a set of k angles. If a graph G on n vertices admits a drawing with at most one bend per edge such that every crossing occurs at some angle from A, then G has at most $(18k + 3(2k+1))n = (24k+3)n$ edges.* □

4.4 Poly-line drawings with at most two bends per edge

In the proof of Lemma 4.10, we constructed red edges in an $\alpha AC_1^=$ drawing of a graph G such that each red edge had a crossing-free first segment and one bend at a crossing. A similar strategy works for $\alpha AC_2^=$ drawings, but the red edges may now have up to two bends.

Lemma 4.14. *Let $\alpha \in (0, \frac{\pi}{2}]$ and $G = (V, E)$ be a graph on $n \geq 4$ vertices that admits an $\alpha AC_2^=$ drawing such that every crossing occurs between an end segment and a middle segment. Then $|E| \leq 116.14n$.*

Proof. Let D be an $\alpha AC_2^=$ drawing of G where every crossing occurs between an end segment and a middle segment. Every middle segment crosses at most two end segments incident to the same vertex in V since every crossing occurs at the same angle α. Let M be the set of middle segments that cross at least 3 end segments, and let S be the set of end segments that cross some middle segment in M. We distinguish two types of edges in G: let $E_1 \subseteq E$ be the set of edges with at least one end segment in S, and let $E_2 = E \setminus E_1$ be the set of edges with no end segment in S. Then $G_2 = (V, E_2)$ has at most $2|E_2|$ crossings in this drawing. The crossing number of a graph with n vertices and m edges is by the second part of Crossing Lemma (Lemma 2.1) at least

$$0.032\frac{m^3}{n^2}$$

Applying this to (V, E_2), we have

$$2|E_2| \geq 0.032\frac{|E_2|^3}{n^2}$$

which gives $|E_2| \leq 7.9n$. Note that if the condition of the second part of Crossing Lemma is not satisfied we have even $|E_2| < 7n < 7.9n$. In the remainder of the proof, we derive an upper bound for $|E_1|$.

We have $|E_1| \leq |S| \leq 2|E_1|$. Direct each segment $s \in S$ from an incident vertex in V to the other endpoint (which is either a bend or another vertex in V). We construct a directed multigraph (V, Γ), which we call the *red* graph. For every end segment $s \in S$, we construct a red edge $\gamma(s) \in \Gamma$, which is a polygonal path with two bends between two vertices in V. It is constructed as follows. Refer to Figure 11.

Let $u_s \in V$ denote the starting point of s (along its direction). Let c_s be the first crossing of s with a middle segment in M, which we denote by m_s. Recall that m_s crosses at least three end segments, at most two of which are incident to u_s. Let $d_s \in m_s$ be the closest crossing to c_s with an end segment that is not incident to u_s. Let $v_s \in V$ be a vertex incident to the end segment containing d_s. If c_s and d_s are consecutive crossings along m_s, then let $\gamma(s) = (u_s, c_s, d_s, v_s)$, see Figure 11(a). Otherwise, there is exactly one crossing x_s between c_s and d_s such that $u_s x_s$ is part of some end segment, and $\angle(u_s c_s, u_s x_s) = \pm(\pi - 2\alpha)$. In this case, let $\gamma(s) = (u_s, x_s, d_s, v_s)$, see Figure 11(b).

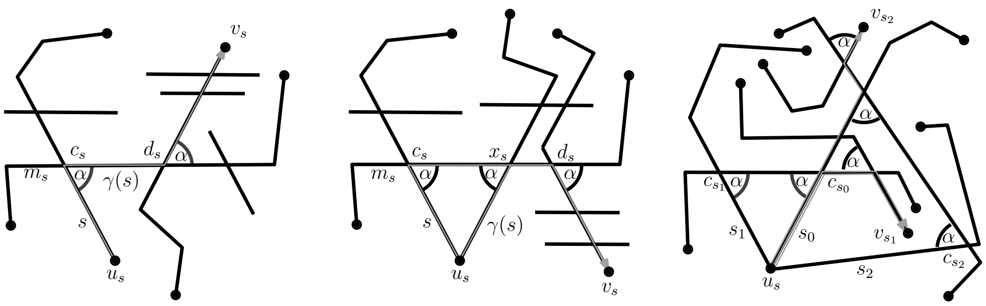

Figure 11: Construction of a red edge $\gamma(s)$. (a) c_s and d_s are consecutive crossings along m_s. (b) there is a crossing x_s between c_s and d_s. (c) The first segments of three red edges corresponding to s_0, s_1 and s_2, respectively, may overlap.

Every edge $\gamma(s) \in \Gamma$ has three segments: the first and third segments of $\gamma(s)$ lie along some end segments of edges in E, and the second segment of $\gamma(s)$ lies along a middle segment in M. By construction, the middle segment of $\gamma(s)$ is between two consecutive crossings along a middle segment in M, and so it does not cross any red edges. The two end segments of $\gamma(s)$ can cross only middle segments of red edges, however, the red middle segments are crossing-free. We conclude that no two red edges cross.

Since the bends c_s, d_s, and x_s are at crossings in an $\alpha AC_2^=$ drawing of G, the turning angle of $\gamma(s)$ must be among the 9 angles in $\{0, \pm\pi, \pm 2\alpha, \pm 2(\pi - \alpha), \pm(\pi - 2\alpha)\}$. Indeed, all the turning angles along $\gamma(s)$ belong to the set $\{\pm\alpha, \pm(\pi - \alpha)\}$. Note also that the red edges may have common tails (which can be removed using Lemma 4.8). Furthermore, the first segments of at most three red edges may overlap because the angle between s and the first segment of $\gamma(s)$ is 0 or $\pm(\pi - 2\alpha)$. However, if the first segments of three red edges overlap, then at most two of these edges are parallel (that is, join the same two vertices in V), see Figure 11(c).

We show that there are at most 36 directed red edges between any two vertices $u, v \in V$. By Lemma 4.7, any two noncrossing paths of the same turning angle between u and v must overlap in the first and last segments. As noted above, the first segments of at most two parallel red edges overlap. Since the red edges may have up to 9 distinct turning angles, there are at most 18 red edges from u to v by Lemma 4.7. Hence there are at most 36 red edges between u and v (in either direction).

Since (V, Γ) is a planar multigraph with edge multiplicity at most 36, it has at most $|\Gamma| \leq 36(3n-6) < 108n$ edges. Altogether, we have $|E| = |E_1| + |E_2| \leq |E_1| + |\Gamma| \leq 7.9n + 108n = 115.9n$. $\square$

We can now prove part (b) of Theorem 4.1.

Lemma 4.15. *For any angle $\alpha \in (0, \frac{\pi}{2}]$, a graph $G = (V, E)$ on n vertices that admits an $\alpha AC_2^=$ drawing has less than $385n$ edges.*

Proof. Consider an $\alpha AC_2^=$ drawing of G. Let E_0 be the set of edges which have an end segment crossing the end segment of another edge. By Lemma 4.10, we have $|E_1| \leq 36n$.

Consider the edges $E_1 = E \setminus E_0$. By Lemma 4.5, there is a partitioned $E_1 = E_{11} \cup E_{12} \cup E_{13}$ such that the middle segments of the edges in each subset are pairwise noncrossing. Suppose without loss of generality that $|E_{11}| = \max(|E_{11}|, |E_{12}|, |E_{13}|)$. By Lemma 4.14, we have $|E_{11}| \leq 115.9$. It follows that $|E| = |E_0| + |E_1| \leq |E_0| + 3|E_{11}| \leq (36 + 3 \cdot 115.9)n = 383.7n$. $\square$

Remark. It is not difficult to generalize the proof of Lemma 4.14 to the case that every two polyline edges cross at one of k possible angles. The only difference is that the red edges may have up to $(4k)^2$ different turning angles, and that the first segment of a red edge may overlap at most $(2k-1)$ first segments of other red edges.

Lemma 4.16. *Let $\alpha \in (0, \frac{\pi}{2}]$ and $G = (V, E)$ be a graph on $n \geq 4$ vertices that admits a polyline drawing such that every crossing occurs between an end segment and a middle segment at one of k possible angles. Then $|E| = O(k^3 n)$.* $\square$

Corollary 4.17. *Let $A \subset (0, \frac{\pi}{2}]$ be a set of k angles. If a graph G on n vertices admits a drawing with at most two bends per edge such that every crossing occurs at some angle from A, then G has $O(k^4 n)$ edges.* $\square$

The dependence on k can be improved. In the next section, we reduce the upper bounds in Lemma 4.16 and Corollary 4.17 to $O(nk)$ and $O(nk^2)$, respectively.

4.5 Crossing between end segments in topological graphs

In this section, we prove Theorems 4.3 and 4.4, and then deduce part (b) of Theorem 4.2 from these general results. Our proof techniques are similar to the method in the previous section: we construct a topological multigraph (V, Γ) whose edges are drawn along some edges in a given topological graph (V, E). The key difference is that we do not assume anything about the crossing angles, and so we cannot use Lemma 4.7 for bounding the edge multiplicity in (V, Γ). The greatest challenge in this section is to bound the edge multiplicity in the auxiliary graph (V, Γ) using solely combinatorial and topological conditions.

4.6 Proof of Theorem 4.3

We start with the proof of Theorem 4.3, which is the topological analogue of our result for $\alpha AC_1^=$ drawings.

Proof of Theorem 4.3: Let $G = (V, E)$ be a topological graph on n vertices, and assume that every edge in E is partitioned into a red end segment and a blue end segment, such that:

(1) no two end segments of the same color cross;

(2) every pair of end segments intersects at most once; and

(3) no blue end segment is crossed by more than k red end segments that share a vertex.

Assume further that G is drawn so that the number of edge crossings is minimized subject to the conditions (1)–(3). We show that G has $O(kn)$ edges.

By Theorem 4.9 the graph G has at most $6.5n - 20$ edges with a crossing-free end segment, since these edges form a simple quasi-planar graph. Denote by $E_1 \subseteq E$ the set of the remaining edges of G.

For every edge $e \in E_1$ we draw a new edge $\gamma(e)$ as follows. Let s denote the red end segment of e, and let $u_s \in V$ be the vertex incident to s. Direct s from u_s to its other endpoint, and let c_s be the first crossing point along s. By condition (2), c_s is a crossing of s with a blue end segment s' of some edge e', where s' is incident to a unique vertex v_s. It is clear that $u_s \neq v_s$, since otherwise we can redraw the portion of e' between v_s and c_s so that it closely follows e and thereby reduce the total number of crossings in G without violating conditions (1)–(3). Let $\gamma(e)$ be the Jordan arc between u_s and v_s that follows the red segment s from u_s to c_s, and the blue segment s' from

c_s to v_s. The new edges form a topological multigraph graph $G' = (V, \Gamma)$, where $\Gamma = \{\gamma(e) : e \in E_1\}$. We call the edges in Γ *red-blue* to distinguish them from the edges in E.

Note that G' is a plane multigraph. Indeed, crossings may occur only between a red end segment and a blue end segment, however, the red end segment of every edge in G' is crossing-free. G' might contain edges with a common tail, however, these overlaps may be removed using Lemma 4.8.

We define a *bundle of edges* in G' as a maximal set of parallel edges such that the interior of the region enclosed by the edges does not contain any vertex of V. Recall that a plane multigraph on n vertices has at most $3n$ edges if it has no face of size 2. Therefore, G' has at most $3n$ bundles.

Proposition 4.1. *Every bundle of edges of G' contains at most $4k+6$ edges.*

Proof. Let B be a bundle of edges between vertices $u, v \in V$. Let $B_1 \subseteq B$ be the set of red-blue edges in B whose red segment is incident to u, and assume without loss of generality that $|B_1| \geq |B|/2$.

Label the red-blue edges in B_1 by $\gamma_1, \gamma_2, \ldots, \gamma_\ell$ in the order they appear in the rotation system at u such that the closed region R enclosed by γ_1 and γ_ℓ contains all other edges of B_1. For $i = 1, 2, \ldots, \ell$, let $e_i \in E$ be the edge of the original graph G incident to u whose red segment contains the red segment of γ_i.

Let $Q = \{e_2, e_3, \ldots, e_{\ell-1}\} \subseteq E$ be a set of $\ell - 2$ edges in E_1 containing the red segments of $\gamma_2, \gamma_3, \ldots, \gamma_{\ell-1}$. By property (3), the red segments of at most k edges in Q cross the blue segment of γ_1. Similarly, the red segments of at most k edges in Q cross the blue segment of γ_ℓ. So the red segments of at least $\ell - (2k+2)$ edges in Q cross neither the red nor the blue segment of γ_1 and γ_ℓ. The relative interiors of these red segments lie in the interior of region R. The blue segment of such an edge in Q cannot cross the boundary of R (since the red segments of γ_1 and γ_ℓ are crossing-free, and blue segments do not cross each other), so this edge must connect u and v. The graph G has at most one edge between u and v, and so $\ell - (2k+2) \leq 1$. It follows that $|B_1| = \ell \leq 2k+3$, as required. ☐

Therefore, G' has at most $3(4k+6)n$ edges. We conclude that $|E| \leq |E_1| + 6.5n = |E'| + 6.5n \leq (4k+6)3n + 6.5n = (12k+24.5)n$. ☐

4.7 Proof of Theorem 4.4

We now prove Theorem 4.4.

Proof of Theorem 4.4: Let $G = (V, E)$ be a topological graph on n vertices. Assume that every edge of G is partitioned into two end segments and one middle segment such that

(1) each crossing involves one end segment and one middle segment;

(2) each middle segment and end segment intersect at most once; and

(3) each middle segment crosses at most k end segments that share a vertex.

Assume further that G is drawn in the plane so that the number of edge crossings is minimized subject to the constraints (1)–(3). We show that G has $O(kn)$ edges.

Observe that G has at most $3n - 6$ edges whose both end segments are crossing-free, since two such edges cannot cross each other. Let $E_1 \subset E$ denote the set of edges with at least one crossed end segment. Let S be the set of end segments with at least one crossing each. It is clear that $|E_1| \leq |S| \leq 2|E_1|$. We construct a red edge for every end segment in S.

For every end segment $s \in S$, let $u_s \in V$ be the vertex incident to s. Direct s from u_s to its other endpoint, and let c_s be the first crossing along s. Direct every middle segment arbitrarily. For every end segment $s \in S$, we construct a directed red edge $\gamma(s)$, which is a Jordan arc from u_s to another vertex in V. These edges form a directed topological multigraph (V, Γ) with $\Gamma = \{\gamma(s) : s \in S\}$. The edges in Γ are called *red* to distinguish them from the edges of E.

For $s \in S$, the red edge $\gamma(s)$ is constructed as follows. See Figure 12(a) for an example. Point c_s is the crossing of s with some middle segment m_s. Let d_s be the first intersection point along m_s after c_s (following the direction of m_s) with an end segment s' which is not adjacent to u_s. That is, d_s is either a crossing of m_s with an end segment or it is the endpoint of m_s (if c_s is the last crossing along m_s or all segments that m_s crosses after c_s are incident to u_s). At any rate, d_s lies on a unique end segment, which is incident to a unique vertex $v_s \in V$. Now let the directed edge $\gamma(s)$ follow segment s from u_s to c_s, the middle segment m_s from c_s to d_s, and the end segment s' from d_s to v_s.

Since G is drawn with the minimal number of crossings, we have $u_s \neq v_s$. Indeed, suppose that $u_s = v_s$. If d_s is the endpoint of the middle segment m_s, then we could redraw the edge $e_m \in E$ containing m_s so that the middle segment of the edge e_m ends right before reaching point c_s and then e_m continues to u_s closely following along s without crossings. By redrawing e_m this way, we reduce the total number of crossings without violating conditions (1)–(3).

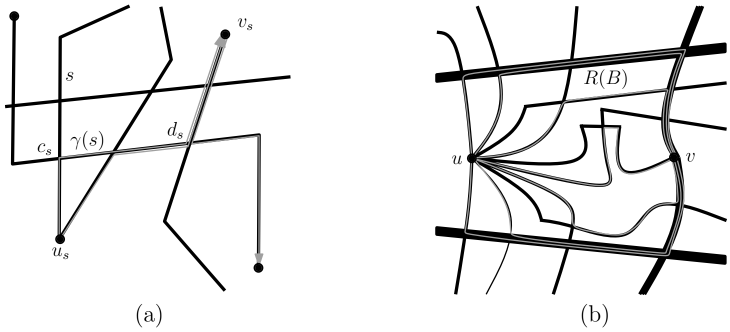

Figure 12: (a) Construction of a red edge $\gamma(s)$. (b) A bungle of 7 red edges from u to v.

Every red edge $\gamma(s)$ is naturally partitioned into three segments: a first, a middle, and a third segment. We briefly summarize the properties we have established for the three segments of the red edges.

(i) The first segments of the red edges are distinct, they lie along the end segments of G, and they are crossing-free.

(ii) The middle segment of each red edge lies along a middle segment in G, following its prescribed direction. A middle segment of a red edge $\gamma(s)$ may be crossed by the last segment of some other red edge $\gamma(s')$ if $\gamma(s')$ is incident to vertex $u_s \in V$.

(iii) The last segment of each $\gamma(s)$ lies along an end segment of G, and it possibly has a common tail with other red edges. The last segment of $\gamma(s)$ may cross middle segments of other red edges.

Observe that a crossing in the red graph can occur only between two red edges sharing a vertex (see Figure 13). Note also that two red edges in Γ cannot follow the same Jordan arc in opposite directions (*e.g.*, (u, v) and (v, u)), since every red edge follows a prescribed direction along its middle segment. We show that (V, Γ) contains a plane subgraph having at least $|\Gamma|/4$ edges.

Label each vertex in V by either 0 or 1 as described below, and let $\Gamma_1 \subseteq \Gamma$ denote the set of red edges directed from a vertex labeled 0 to one labeled 1. If the labels are distributed uniformly at random, then every edge in Γ is in Γ_1 with probability $1/4$. Thus, the expected number of edges in Γ_1 can be expressed as follows.

$$\mathbb{E}(|\Gamma_1|) = \sum_{\gamma \in \Gamma} \frac{1}{4} = \frac{1}{4}|\Gamma|$$

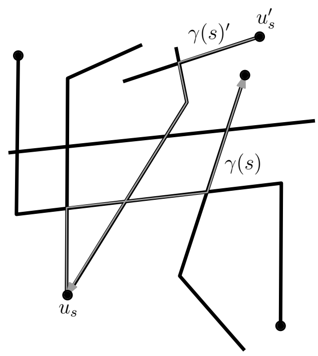

Figure 13: Two red crossing edges $\gamma(s)$ and $\gamma(s)'$.

Hence there is a labeling such that $|\Gamma_1| \geq |\Gamma|/4$. Fix such a labeling for the remainder of the proof. By the properties of red edges noted above, no two edges in Γ_1 cross. If two red edges in Γ_1 overlap, then they have a common tail. By Lemma 4.8, overlaps along common tails can be removed, and so (V, Γ_1) is a directed plane multigraph.

In (V, Γ_1), we define a *bundle* as a maximal set of directed parallel edges such that the interior of the region enclosed by the edges does not contain any vertex of V. See Figure 12(b). Let $\mathcal{B}$ denote the set of bundles of (V, Γ_1). For a bundle $B \in \mathcal{B}$, let $R(B)$ denote the region enclosed by the edges in B. Since Γ_1 is planar and each edge goes from a vertex labeled 0 to one labeled 1, the interior of the regions $R(B)$, $B \in \mathcal{B}$, are pairwise disjoint. Recall that a plane multigraph on n vertices has at most $3n$ edges if it has no face of size 2. Therefore, there are at most $3n$ bundles in $\mathcal{B}$.

Proposition 4.2. *Let $B = \{\gamma_1, \gamma_2, \ldots, \gamma_\ell\} \subseteq \Gamma_1$ be a bundle in $\mathcal{B}$ from u to v appearing in the rotation system at u in this order. For $i = 1, 2, \ldots, \ell$, let $e_i \in E$ be the edge of the original graph G incident to u whose end segment contains the first segment of γ_i. Then there are at least $\ell - (2k+3)$ edges e_i, $1 \leq i \leq \ell$, whose first segment lies entirely in the region $R(B)$.*

Proof. Let $E_{uv} = \{e_2, e_3, \ldots, e_{\ell-1}\} \subseteq E$ be a set of $\ell - 2$ edges in E containing the first segments of $\gamma_2, \gamma_3, \ldots, \gamma_{\ell-1} \in \Gamma_1$. By property (3), the first segments of at most k edges in E_{uv} cross the middle segment of γ_1. Similarly, the first segments of at most k edges in E_{uv} cross the middle segment of γ_ℓ. So at least $\ell - (2k+2)$ edges in E_{uv} cross neither the first nor the middle segment of γ_1 and γ_ℓ. At most one of these edges joins u and v, since G is a simple graph. So at least $\ell - (2k+3)$ edges in E_{uv} has to cross the last segment of γ_1 or the last segment of γ_ℓ. However by property (1), only the middle segments of the edges in E_{uv} can cross the last segment of γ_1 or γ_ℓ. Hence, for at least $\ell - (2k + 3)$ edges in E_{uv}, the first segment lies entirely in region $R(B)$. $\square$

Partition Γ_1 into two two subsets $\Gamma_1 = \Gamma_2 \cup \Gamma_3$ as follows. Let Γ_2 contain all edges of all bundles of size at most $2k + 3$, as well as those edge $\gamma_i \in B$ of any larger bundle $B \in \mathcal{B}$ such that the first segment of the corresponding edge $e_i \in E$ is *not* contained in the region $R(B)$. Let $\Gamma_3 = \Gamma_1 \setminus \Gamma_2$, that is, Γ_3 contains all edges $\gamma_i \in \Gamma_1$ such that γ_i is part of some bundle B of size at least $2k + 4$, and the first segment of the corresponding edge $e_i \in E$ lies entirely in the region $R(B)$. By Proposition 4.2, each bundle in $\mathcal{B}$ contains at most $2k + 3$ edges of Γ_2. Since there are at most $3n$ bundles, we have $|\Gamma_2| \le (2k + 3)3n$. It remains to bound the number of edges in Γ_3.

Label each region $R(B)$ enclosed by a bundle $B \in \mathcal{B}$ by either 0 or 1 as described below. Let $\Gamma_4 \subseteq \Gamma_3$ be the set of edges $\gamma \in \Gamma_3$ such that the end segment of the edge $e \in E$ containing the first segment of γ lies in a region labeled 0, and the other end segment of e either lies in a region labeled 1 (including its boundary) or it does not lie in any such region. If the labels are distributed uniformly at random, then every edge in Γ_3 will be in Γ_4 with probability at least 1/4. Then the expected number of edges in Γ_4 is

$$\mathbb{E}(\Gamma_4) = \frac{1}{4}|\Gamma_3|$$

Hence there is a labeling such that $|\Gamma_4| \ge |\Gamma_3|/4$. Fix such a labeling for the remainder of the proof. Let $E_4 \subseteq E$ denote the edges containing the first segments of the red edges in Γ_4. In what follows instead of bounding the size of Γ_4 we in fact bound the size of E_4. We can proceed in this fashion as each end segment of an edge in E_4 gave rise to at most one edge in Γ_4.

Consider a bundle $B \in \mathcal{B}$ of edges from u to v, and suppose that the region $R(B)$ is enclosed by the edges $\gamma_1, \gamma_\ell \in B$. We say that an edge $e \in E$ *traverses* $R(B)$ if a connected component of $e \cap R(B)$ intersects the interior of $R(B)$, but this component is incident to neither u nor v. See Figure 14(a). Suppose that $e \in E$ traverses $R(B)$. Then e crosses $\gamma_1 \cup \gamma_\ell$ twice. The first segments of γ_1 and γ_ℓ are crossing-free by property (i) of the red edges. The middle segments of γ_1 and γ_ℓ can cross only end segments incident to u by property (ii) of the red edges. By condition (3), at most $2k$ edges traverse $R(B)$ such that they intersect the middle segment of γ_1 or γ_ℓ. Denote by $T_1 \subset E$ the set of edges in E that traverse some bundle $B \in \mathcal{B}$ such that they cross a middle segment on the boundary of $R(B)$. Summing over all bundles, we have $|T_1| \le (3n)(2k) = 6kn$.

Assume now that $e \in E$ crosses the last segment of γ_1 and γ_ℓ. By conditions (1) and (2), the middle segment of e crosses the last segment of γ_1

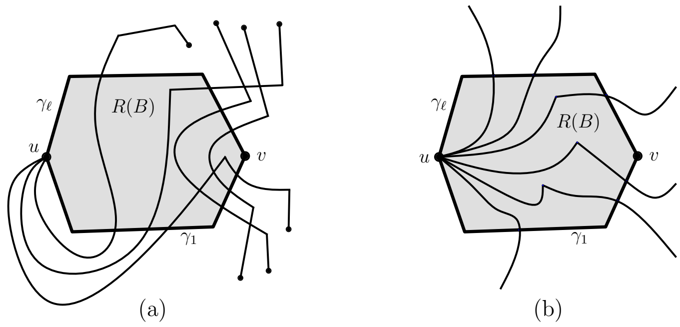

Figure 14: (a) Four edges traverse $R(B)$, two of them cross some middle segments on the boundary of $R(B)$, and two of them cross end segments only. (b) Construction of two edges in $\widehat{E}$ by redrawing their portions within a region $R(B)$.

and γ_ℓ. That is, a connected component of $e \cap R(B)$ is part of the middle segment of e. By condition (1) and (2), the middle segments that traverse $R(B)$ cannot cross each other inside $R(B)$. Among all middle segments that traverse $R(B)$, let $m(B)$ be the one whose intersection with γ_1 (and γ_ℓ) is closest to u. Recall that for a red edge $\gamma_i \in B \cap \Gamma_4$, we denote by $e_i \in E_4$ the edge containing the first segment of γ_i. By the choice of $\Gamma_4 \subset \Gamma_3$, an end segment of e_i lies in the region $R(B)$. If the first segment of e_i crosses any middle segment that traverses $R(B)$, then e_i must cross $m(B)$. By condition (3), however, at most k edges $e_i \in E_4$ cross $m(B)$. Denote by $T_2 \subset E$ the set of edges in E such that one end segment lies entirely in the region $R(B)$ of some bundle $B \in \mathcal{B}$, and this end segment crosses a middle segment that traverses $R(B)$. Summing over all bundles, we have $|T_2| \leq (3n)k = 3kn$.

Let $\Gamma_5 \subseteq \Gamma_4$ be the set of edges $\gamma_i \in \Gamma_4$ such that the corresponding edge e_i is neither in T_1 nor in T_2. Let $E_5 \subseteq E$ denote the edges containing the first segments of the red edges in Γ_5. By the above argument, we have $|\Gamma_4| \leq |\Gamma_5| + |T_1| + |T_2| \leq |\Gamma_5| + 9kn$. Indeed, even though an edge in $T_1 \cup T_2$ can give rise to two edges in Γ_4, if that is the case we counted it already twice in $T_1 \cup T_2$. It remains to derive an upper bound for $|\Gamma_5|$.

In order to bound $|\Gamma_5|$, we construct the new topological graph $(V, \widehat{E})$. For each $\gamma_i \in \Gamma_5$, we construct an edge $\hat{e}_i \in \widehat{E}$ as follows. Suppose that γ_i is in a bundle $B \in \mathcal{B}$ from u to v. Let $e_i \in E_5$ denote the edge that contains the first segment of γ_i. Suppose that $e_i = (u, w)$. By construction, the first segment of e_i lies in the region $R(B)$, and it does not cross any edge in E_5 that traverses $R(B)$. Indeed, the first segment of e_i cannot cross any edge in E_5 that traverses $R(B)$ since $E_5 \cap (T_1 \cup T_2) = \emptyset$, and the middle segment of e_i cannot cross any edge in E_5 that traverses $R(B)$ since $E_5 \cap T_1 = \emptyset$. Draw the

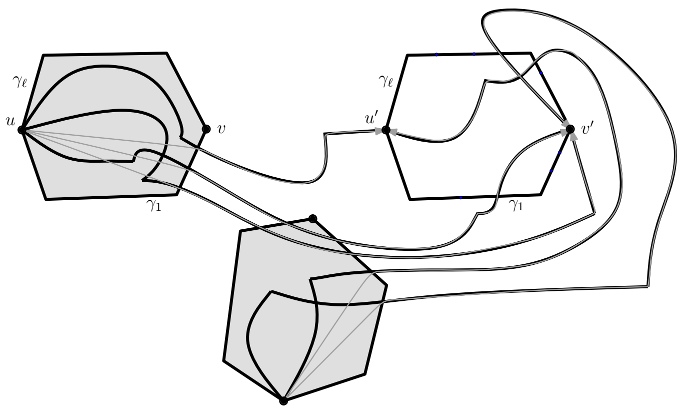

Figure 15: Construction of $\widehat{E}$; grey regions are labeled with 0.

edge $\hat{e}_i = (u, w)$ as follows (refer to Figure 14(b)): $\hat{e}_i$ starts from the vertex u, it goes to the first intersection point $e_i \cap \partial R(B)$ inside the region $R(B)$ as described bellow, then it follows e_i to the endpoint w outside of $R(B)$. Due to the 0-1 labeling of the regions $R(B)$, all edges of $\widehat{E}$ that intersect the interior of a region $R(B)$ labeled with 0 are incident to only one vertex of the bundle B (see Figure 15). Therefore, the portion of the edges in $\widehat{E}$ in each region $R(B)$; $B \in \mathcal{B}$, can be drawn without crossings. We can now partition each edge $\hat{e}_i \in \widehat{E}_5$ into *two* segments: its *blue* segment consists of its part inside a region $R(B)$ and its part along the middle segment of e_i; its *red* segment is the last segment of e_i. As noted above, $\hat{e}_i$ does not cross any other edge of $\widehat{E}$ inside the region $R(B)$. By property (3), every blue segment crosses at most k red segments incident to the same vertex. We can apply Theorem 4.3 for the graph $(V, \widehat{E})$. It follows that $|\widehat{E}| = O(kn)$, hence $|\Gamma_5| = O(kn)$.

In summary, we have $|E| < 3n + |E_1| \leq 3n + |\Gamma| \leq 3n + 4|\Gamma_1| = 3n + 4(|\Gamma_2| + |\Gamma_3|) \leq 3n + 4(2k+3)3n + 4|\Gamma_3| \leq (8k+39)n + 16|\Gamma_4| \leq (8k+39)n + 16(9kn + |\Gamma_5|) = (152k+39)n + 16|\Gamma_5| \leq (152k+39)n + 16|\widehat{E}| \leq (152k+39)n + 16(12k+24.5)n = (344k+431)n$, as required. $\square$

4.8 Completing the proof of Theorem 4.2

We are now ready to prove Theorem 4.2(b).

Proof of Theorem 4.2(b): Let $A \subset (0, \frac{\pi}{2}]$ be a set of k angles. Let G be a graph on n vertices that admits a drawing with at most two bends per edge such that every crossing occurs at some angle from A. Partition the edges into two subsets $E = E_1 \cup E_2$, where E_1 is the set of edges which have an end segment crossing some other end segment, and E_2 contains all other edges in E. By Lemma 4.12, we have $|E_1| \leq 36kn$.

Let S_2 be the set of *middle* segments of all edges in E_2. By Lemma 4.6, there is a subset $S_2' \subseteq S_2$ of at least $\frac{1}{2k+1}|S_2| = \frac{1}{2k+1}|E_2|$ pairwise noncrossing segments. Let $E_2' \subseteq E_2$ be the set of edges whose middle segments are in S_2'. Note that in the graph (V, E_2'), every crossing is between an end segment and a middle segment. Moreover, no middle segment crosses more than $2k$ end segments that share a vertex, since there are k possible crossing angles, and for each angle $\alpha \in A$, two end segments incident to a vertex that meet a middle segment at the angle α form an isosceles triangle. Therefore, it follows from Theorem 4.4 that $|E_2'| = O(kn)$. Altogether, we have $|E| = |E_1| + |E_2| \leq 36kn + (2k+1)|E_2'| = O(k^2n)$. $\square$

4.9 Discussion and Concluding Remarks

We have shown that for every list A of k angles, a graph on n vertices that admits a polyline drawing with at most one (resp., two) bends per edge in which all crossings occur at an angle from A has at most $O(kn)$ (resp., $O(k^2n)$) edges. It is easy to construct a *straight line* graph with n vertices and $\Omega(kn)$ edges such that the edges cross in at most k different angles: Let the vertices $v_1, \ldots, v_n$ be equally spaced points along a circle in this order, add a straight line edge v_iv_j if and only if $|i-j| \leq k+1$.

With one bend per edge, one can construct slightly larger graphs on n vertices, but the number of edges remains $O(kn)$ by Theorem 4.2(a). However, we do not know whether the upper bound of $O(k^2n)$ in Theorem 4.2 is the best possible for polyline drawings with two bends per edge and k possible crossing angles.

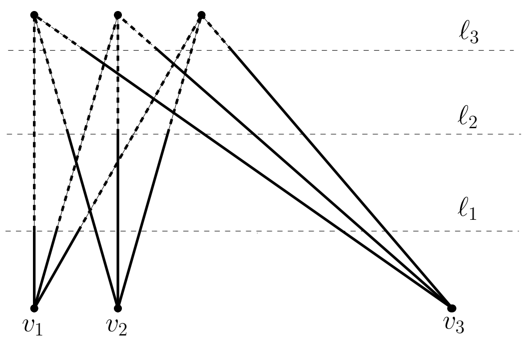

Figure 16: A straight line drawing of $K_{3,3}$ satisfying the conditions of Theorem 4.3 for $k = 2$.

In Theorems 4.3 and 4.4, the upper bound $O(kn)$ cannot be improved.

For every $k \geq 0$ there is a straight line drawing of $K_{k+1,k+1}$ satisfying the conditions of Theorem 4.3, by the following construction which is due to Rom Pinchasi [62]. Place $k+1$ vertices on a horizontal line. Then add $k+1$ vertices $v_1, v_2, \ldots, v_{k+1}$ one by one on another horizontal line below the first line, and connect each of them to all the first $k+1$ vertices, as follows. Place v_1 arbitrarily on the bottom line, and partition each of its adjacent edges into red and blue segments such that the red segments are adjacent to v_1. For $i = 2, \ldots, k+1$, add v_i far enough to the right of v_{i-1} such that the edges between v_i and the first set of vertices cross only blue segments of previous edges. Let (x_i, y_i) be the highest crossing point on edges adjacent to v_i. Fix a horizontal line ℓ_i slightly above the line $y = y_i$ and partition every edge e adjacent to v_i into red and blue segments, such that endpoints of the red segment are v_i and $e \cap \ell_i$. See Figure 16 for an example. It is not hard to verify that this drawing satisfies the condition of Theorem 4.3. By taking $\left\lfloor \frac{n}{2(k+1)} \right\rfloor$ disjoint copies of $K_{k+1,k+1}$ drawn as above, we obtain a graph on n vertices and $\Omega(kn)$ edges satisfying the conditions of Theorem 4.3.

5 Hanani-Tutte and Monotone drawings

In what follows we introduce a monotone variant of Hanani-Tutte theorem and then show some interesting algorithmic consequences of this result. A considerable part of the material from the present section was published in [24].

5.1 Introduction

The classic Hanani-Tutte theorem (Theorem 2.2) states that if a graph can be drawn in the plane so that no pair of independent edges crosses an odd number of times, then it is planar [26, 73]. (Two edges are *independent* if they do not have a shared endpoint.) There are many ways to look at this result; for example, in algebraic topology it is seen as a special case of the van Kampen-Flores theorem [44, Chapter 5] which classifies obstructions to embeddability in topological spaces. This point of view leads to challenging open questions (see, for example, [45]), but even in 2-dimensional surfaces the problem is not understood well (see [67] for a survey of what we do know).

Here, we study a variant of the problem which was introduced by Pach and Tóth [50]. A curve is *x-monotone* if intersects every vertical line at most once. A drawing of a graph is *x-monotone* if every edge is x-monotone and every vertical line contains at most one vertex. The natural analogue of the Hanani-Tutte theorem in this context is that for any x-monotone drawing in which no pair of independent edges crosses an odd number of times, there is an x-monotone *embedding* (that is, a crossing-free drawing) where the x-coordinate of every vertex is unchanged. The truth of this result was left as an open problem by Pach and Tóth, and conjectured in [67]. We prove it as Theorem 5.7 in Section 5.4.

The weak version of the classic Hanani-Tutte theorem states that if a graph can be drawn so that *no* pair of edges crosses oddly, then it is planar. The analogue for x-monotone drawings states that if there is an x-monotone drawing in which no pair of edges crosses an odd number of times, then there is an x-monotone embedding where the x-coordinate of every vertex is unchanged. This variant of the weak Hanani-Tutte theorem was first proved by Pach and Tóth.* We give a new proof of this result as Theorem 5.1 in Section 5.3, which extends an elementary topological approach found in some earlier papers on the Hanani-Tutte theorem, such as [54].

A traditional approach to Hanani-Tutte style results is via obstructions; this sometimes leads to very slick proofs, like Kleitman's proof of the Hanani-

*There is a gap in the original argument; an updated version is now available [50, 51].

Tutte theorem for the plane [40], but there are two drawbacks: complete obstruction sets are not always known, e.g. for the torus or, in spite of several attempts, for x-monotone embeddings (as discussed in [20]); and this approach is of little help algorithmically. Pach and Tóth took another approach, building on a proof of the weak Hanani-Tutte theorem for surfaces by Cairns and Nikolayevsky [10].

Somewhat surprisingly, Theorem 5.7 leads to a practical solution to a well-known graph drawing problem. In Section 5.6 we will see how Theorem 5.7 can be used to recognize level-planar graphs and find an embedding in quadratic time. While the best-known algorithms for this problem run in linear time [37], they are quite complicated, while our algorithms are quite simple. There are previous claims for simple quadratic time algorithms for level-planarity testing, which we discussed above.

The condition that edges are x-monotone in Theorem 5.1 and Theorem 5.7 can be replaced, somewhat surprisingly, by a weaker notion. Let us say that an edge uv in a drawing is *x-bounded* if every interior point p of uv satisfies $x(u) < x(p) < x(v)$. That is, an edge is x-bounded if it lies strictly between its endpoints; it need not be x-monotone within those bounds. (See Corollary 5.6 and Remark 5.8.) It also works for level planarity, which we use to show that the definition of level planarity as given in the literature can be relaxed (Corollary 5.15).

The usual (not monotone) Hanani-Tutte theorem and weak Hanani-Tutte theorems have extensions that relate crossing number to odd crossing number and independent odd crossing number [49, 57, 58]. In Section 5.7, we investigate analogous possibilities for monotone notions of the crossing number. Our results show that the monotone Hanani-Tutte theorems cannot be strengthened in ways that are possible in the non-monotone case.

Before continuing, we discuss some basics including terminology and notation.

5.2 Basics, Terminology and Notation

Given an x-monotone drawing, by stretching and compressing the plane horizontally we can change the x-coordinates of the vertices while maintaining x-monotone edges, arbitrarily except for the restriction that the relative order of the x-coordinates of the vertices is unchanged. We can also alter the y-coordinate of any vertex by stretching the plane vertically near that vertex, so that edges remain x-monotone and other vertices are fixed. Thus, we can modify an x-monotone drawing to relocate vertices arbitrarily, so long as the relative x-order of vertices is unchanged. As a result, if we are given an x-monotone drawing and an x-monotone redrawing where the relative order of the x-coordinates is unchanged, we may assume without loss of generality that the location of every vertex is unchanged. Alternatively, we may assume without loss of generality that the vertices in an x-monotone drawing are located at the points $(1, 0)$, …, $(|V|, 0)$. Moreover, all of this still works if the edges are x-bounded rather than x-monotone.

Usually we consider graphs, but we will also have cause to study *multigraphs*, which allow the possibility of having more than one edge between each pair of vertices. Since we only consider x-monotone edges and x-bounded edges, there will never be any loops.

For any graph G and $S \subseteq V(G)$, let $G[S]$ denote the *subgraph induced by S*; that is, the graph on vertex set S with edge set $\{uv \in E(G) : u \in S, v \in S\}$. The *rotation* at a vertex is the clockwise ordering of edges at that vertex, in a drawing of a graph. The *rotation system* of a graph is the collection of rotations at its vertices. In an x-monotone drawing, the *right (left) rotation* at a vertex is the clockwise order of the edges leaving the vertex towards the right (left). Note that the right rotation is ordered from top to bottom, and the left rotation is ordered from bottom to top. By a *wedge* at a vertex v of G in a planar representation of G we understand a pair of edges (e, e') incident to v that are consecutive in its rotation system. A face f of a plane graph G contains a wedge (e, e'), $e, e' \in E$, if f contains v, e and e' on the boundary. We will not always carefully distinguish between an abstract graph and a topological (drawn or embedded) graph; "vertex" and "edge" are used in both contexts. We use $x(v)$ to denote the x-coordinate of a vertex v located in the plane.

Rather than give the level-planarity definitions here, we postpone the discussion until Section 5.6.

5.3 Weak Hanani-Tutte for Monotone Drawings

An edge is *even* if it crosses every other edge an even number of times (possibly 0 times). A drawing is *even* if all its edges are even.

Theorem 5.1 (Pach, Tóth [50, 51])**.** *If G has an x-monotone even drawing, then G has an x-monotone embedding with the same vertex locations and rotation system.*

The weak Hanani-Tutte theorem states that every graph with an even drawing is planar (and it is also known that there is an embedding with the same rotation system). For background and variants of the weak Hanani-Tutte theorem, see [67].

Theorem 5.1 may prompt the reader familiar with Hanani-Tutte style results (in particular [49, Theorem 1] and [54, Theorem 2.1]) to ask whether something stronger is true: a “removing even crossings” lemma which would say that all even edges can be made crossing-free in the drawing of a graph which contains odd edges (while maintaining x-monotonicity and vertex locations). We will see in Section 5.7 that there cannot be any such lemma for x-monotone drawings.

Nearly the same result is claimed by Pach and Tóth in [50, Theorem 1.1], but instead of maintaining the rotation, they state that we can find an *equivalent* x-monotone embedding, where two drawings are equivalent if no edge changes whether it passes above or below a vertex. However, there are simple examples that show that one cannot hope to maintain equivalence in this sense, such as the graph on the left in Figure 17.

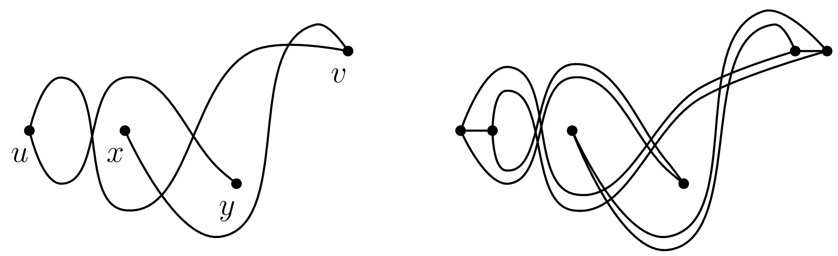

Figure 17: (*left*) An x-monotone even drawing. Since x is above uv and y is below uv, any equivalent x-monotone embedding with the same relative x-ordering of the vertices will have uv below x and above y. But then xv is above uy, so it is not equivalent. (*right*) Essentially the same argument applies to this 2-connected example.

The proof in [50] contains a gap: it is not immediately clear how multiple faces that share a boundary can be embedded simultaneously.* Eliminating the gap requires dropping equivalence. Pach and Tóth have prepared an updated version of the paper that includes a more detailed argument [51].†

As the graph on the right in Figure 17 shows, the counterexample can be made 2-connected, so equivalence cannot be obtained by assuming 2-connectedness. On the other hand, see Corollary 5.3 for a positive result.

We will repeatedly make use of a simple topological observation: suppose we are given two curves (not necessarily monotone) that start at the line $x = x_1$ and end at the line $x = x_2$, and that lie entirely between $x = x_1$ and

*In the text after Lemma 2.1 on page 42 of [50], D_κ cannot necessarily be glued together without changing equivalence.

†In this newer version, equivalence is redefined to mean having the same rotation system.

$x = x_2$. The two curves cross an even number of times if and only if they have the same vertical order at $x = x_1$ and $x = x_2$. (If they start or end in the same point p, the *vertical order at* p is determined by the vertical order in which they enter p).

We will also find the following redrawing tool useful.

Lemma 5.2. *Suppose that a multigraph G has an x-monotone embedding and f is an inner face of the embedding, with m_f and M_f being the leftmost and the rightmost vertex of f. Now suppose that we add an edge $m_f M_f$ to the embedding so that it lies in f. (Note that $m_f M_f$ is not required to be x-monotone and that there may be multiple ways of inserting $m_f M_f$ into the rotations at m_f and M_f.) Then the resulting graph $G \cup \{m_f M_f\}$ has an x-monotone embedding with the same vertex locations and rotation system.*

Note that the redrawing in Lemma 5.2 does not preserve equivalence in the sense of Pach and Tóth [50]. See Figure 18.

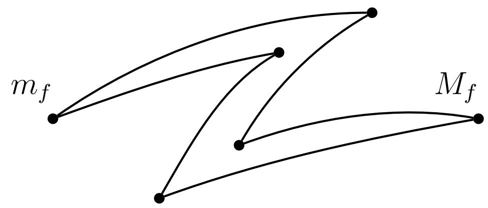

Figure 18: Although we can draw the edge $m_f M_f$ within the Z-shaped face, any subsequent x-monotone redrawing that maintains relative vertex x-order and rotation system will not be equivalent.

Proof of Lemma 5.2. If G consists of multiple components, it is sufficient to prove the result for the component containing f and shift its embedding vertically so that it does not intersect any other component. This allows us to assume that G is connected. Then every face is bounded by a closed walk.* The boundary of f can be broken into two m_f, M_f-walks B_1, B_2, with B_1 starting above $m_f M_f$ in the rotation at m_f, and B_2 starting below.

Let D_f be the drawing of G intersected with $U_f := \{(x, y) \in \mathbb{R}^2 : x(m_f) < x < x(M_f)\}$. ($D_f$ is a subset of the plane, not a graph.) We will locally redraw G in U_f so that $m_f M_f$ can be inserted as a straight-line segment. For each (topologically) connected component Z of D_f, either (i) for every

* *Walks* are like paths except that vertices and edges can be repeated. In a *closed walk* the last vertex is the same as the first vertex.

x between $x(m_f)$ and $x(M_f)$, there is a y-value of B_1 at x that is below all y-values of Z at x, or (ii) for any x between $x(m_f)$ and $x(M_f)$, there is a y-value of B_2 at x that is above all y-values of Z at x.

Let Z_1 be the union of all components of the first type, and Z_2 be the union of all components of the second type. Let L be the line through m_f and M_f. We will show how to move Z_1 to the half-plane above L, without changing the x-value of any point in Z_1 while fixing the points on the boundary of U_f. Let P be an x-monotone curve with endpoints m_f and M_f that lies strictly below Z_1 in U_f (note that m_f and M_f do not belong to U_f). Now move every point v of Z_1 up by the vertical distance between P and L at $x = x(v)$. We proceed similarly to move Z_2 strictly below L, after which we can embed $m_f M_f$ as L. The overall embedding is as desired. □

Proof of Theorem 5.1. We prove the following statement by induction on the number of vertices and then the number of edges:

> If G is a multigraph that has an x-monotone drawing in which all edges are even, then G has an x-monotone embedding with the same vertex locations and rotation system.

In the base case G consists of a single vertex, so the result is immediate. If G consists of multiple components, we can apply induction to each component and combine the drawings by stacking them vertically, that is, translate each component vertically so no two components intersect. Thus, we may assume that G is connected.

We first consider the case that there is more than one edge between the two leftmost vertices of G, $x_1 < x_2$. If there are several edges between x_1 and at x_2, say $e_1, \ldots, e_k$, these have to be consecutive in the rotations at both x_1 and x_2: This is trivial for the rotation at x_2, since all edges incident to it on the left have to go to x_1. Now suppose there is an edge $f = x_1 x_\ell$, $\ell > 2$ so that f falls between two edges e_i and e_j in the rotation at x_1, $1 \leq i, j \leq k$. It is easy to see that f must cross either e_i or e_j oddly which contradicts f being even, so such an edge does not exist. Hence all edges between x_1 and x_2 are consecutive and, moreover, have mirror rotations at x_1 and x_2 (again a consequence of them being even). We can then replace them with a single edge e between x_1 and x_2. By induction, that reduced graph has the required embedding, and we can replace e with the multiple edges $e_1, \ldots, e_k$ obtaining the desired embedding of G.

Now, consider the case that there is only a single edge $x_1 x_2$ between x_1 and x_2. We can contract $x_1 x_2$ by moving x_2 along the edge towards x_1 and inserting the right rotation of x_2 into the rotation at x_1 (see Figure 19). Note that all edges remain even (since $x_1 x_2$ is even), so by induction the new graph has an x-monotone embedding in which $x_1 = x_2 < x_3 < \cdots < x_n$. We can now split the merged vertex into two vertices again and insert a crossing-free edge x_1x_2, obtaining an embedding of the original graph (since we kept the rotation) with the original rotation.

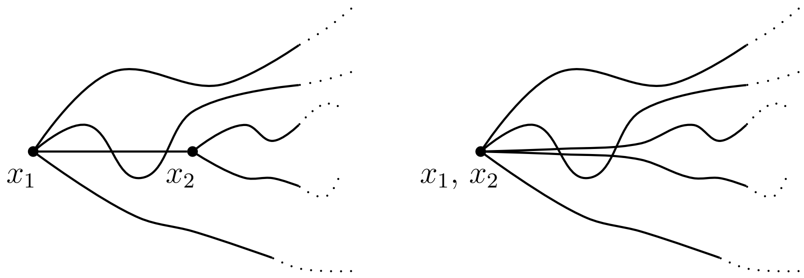

Figure 19: How to contract edge x_2x_1 towards x_1 and merging rotations.

This leaves us with the case that there is no edge between x_1 and x_2. If $G - x_1$ consists of a single component, consider all edges $e_1, \ldots, e_k$ leaving x_1. Each of these edges passes x_2 either above or below x_2. We claim that it is not possible that there are edges e and f so that e leaves x_1 above f but passes under x_2, while f passes above x_2: assume for a contradiction that this is the case; pick a cycle C that contains both e and f (this cycle exists, since we assumed $G - x_1$ has a single component). Let M_C be the rightmost vertex of C. Consider the following two curves within C: C_e, which starts just below x_2 on e and leads to M_C, and C_f, which starts just above x_2 on f and leads to M_C. Note that since e leaves x_1 above f and C is a cycle consisting of even edges lying entirely between x_1 and M_C, curve C_e enters M_C above C_f. Pick a shortest path P_{x_2} from x_2 to C (such a path exists, since G is connected). We distinguish two cases (illustrated in Figure 20).

(i) *P_{x_2} lies strictly to the left of M_C.* Without loss of generality, suppose that P_{x_2} ends on C_f. Let P'_{x_2} be the x_2, M_C-subpath of $P_{x_2} \cup C_f$. Since C_e and P'_{x_2} share no edges, and C_e passes below x_2, C_e must enter M_C below P'_{x_2} (all edges are even). However, the last part of P'_{x_2} belongs to C_f, so C_e enters M_C below C_f which we know to be false.

(ii) *P_{x_2} contains a vertex at or to the right of $x = M_C$.* Let P'_{x_2} be the shortest subpath of P_{x_2} starting at x_2 and ending at or to the right of $x = M_C$. Since P_{x_2} has no edges in common with either C_e or C_f, P'_{x_2} enters M_C above C_e and below C_f if P'_{x_2} ends in M_C. Otherwise, P'_{x_2} passes M_C above C_e and below C_f. Since we know that C_e enters M_C above C_f, Case (ii) also leads to a contradiction.

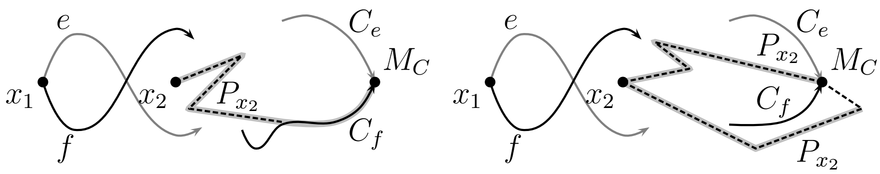

Figure 20: (*left*) Case (i): P_{x_2} is dashed, P'_{x_2} is the thick gray path from x_2 to M_C; (*right*): Case (ii): Both subcases are displayed: the top P_{x_2} stays to the left of M_C, while the bottom P_{x_2} passes to the right of M_C. P'_{x_2} as thick gray path in both cases.

This establishes the claim that if e leaves x_1 above f, then it is not possible that e passes below x_2 while f passes above x_2. In other words, if some edge e starting at x_1 passes below x_2, then all edges starting below e at x_1 also pass below x_2. Hence, all edges passing below x_2 are consecutive at x_1 and so, perforce, are the edges passing above x_2. We can now add a new edge e from x_2 to x_1 that attaches in the rotation between the group of edges passing above x_2 and the edges passing below x_2. This new edge will then be even, so we are in an earlier case that we know how to solve (by contracting the new edge, which reduces the number of vertices, then applying induction).

It remains to deal with the case that $G - x_1$ separates into multiple components. Let H'_i, $i = 1, 2$ be two of those components and let H_i, $i = 1, 2$ be H'_i together with its edges of attachment to x_1, that is $H_i = G[\{x_1\} \cup V(H'_i)]$. Note that the edges of H_1 and H_2 attaching to x_1 cannot interleave, that is, at x_1 we cannot have edges e_1, e_2, f_1, f_2 in that order so that $e_i, f_i \in E(H_i)$ for $i = 1, 2$; the reason is that e_i and f_i can be extended to a cycle $C_i \subset H_i$ and C_1 would cross C_2 an odd number of times for e_1, f_1 to interleave with e_2, f_2 at x_1. This implies that we can define a partial ordering $\prec$ on these components, where $H_1 \prec H_2$ if the edges attaching H_1 to x_1 are surrounded (in the right rotation at x_1) by the edges attaching H_2 to x_1. Now let H be a minimal element of $\prec$; then the edges of H attaching to x_1 are consecutive at x_1. If H contains the rightmost vertex of G, then H is also a maximal element in $\prec$, so H cannot be the only minimal element of $\prec$; in this case, reassign H to another minimal element of $\prec$ that does not contain the rightmost vertex of G. Let $H' = H - \{x_1\}$.

Consider $G - V(H')$. By induction, there is an embedding of $G - V(H')$ which maintains the vertex locations and the rotation system. Let f be the face incident to x_1 into which H has to be reinserted (so that we recover the original rotation system). We can assume that f is not the outer face: if it is, we can make it an inner face by adding an edge from x_1 to the rightmost vertex of G. By Lemma 5.2, we can assume that the embedding has an x-monotone edge from x_1, starting where H' was attached in its rotation, to the rightmost vertex incident to f, which we call M_f. We can find an x-monotone embedding of H by induction. Note that all vertices of H must lie to the left of M_f, since otherwise an edge of H must have crossed an edge on the boundary of f oddly before $G - V(H')$ we redrawn using Lemma 5.2. But then we can insert the new embedding of H into the embedding of $(G - V(H')) \cup \{x_1 M_f\}$ near the edge $x_1 M_f$, such that there are no crossings, which gives us the desired embedding of G. □

Note that in the proof of Theorem 5.1, all redrawing steps maintain equivalence except for applications of Lemma 5.2. This part of the proof, however, only arises in the case that $G - \{x_1\}$ is not connected. Hence, if we can make an assumption on G so that this case never occurs, we can conclude that the resulting embedding is equivalent to the original drawing in the sense of Pach and Tóth [50]. We already saw that 2-connectedness is not sufficient, however, another notion is: a graph in which the vertices are ordered (from left to right, say) is a *hierarchy* if every vertex except the rightmost one has an edge leaving it towards the right [16].

Corollary 5.3. *If G has an x-monotone even drawing and G is a hierarchy, then G has an equivalent x-monotone embedding with the same vertex locations and rotation system.*

Proof. This follows from the proof of Theorem 5.1. The only operation that changes equivalence of edges and vertices in that proof is the application of Lemma 5.2. If G is a hierarchy, $G - x_1$ consists of a single component, since any two vertices in $G - x_1$ are connected by a path (in a hierarchy any two vertices must have a common ancestor). Since contracting the leftmost edge of a hierarchy results in another hierarchy the result follows by induction. □

With an eye towards the algorithmic discussion in Sections 5.5 and 5.6, we analyze the situation more closely and ask how hard it is to find an actual embedding. We are given an x-monotone even drawing, but what is the encoding?

To decide whether each independent pair of edges crosses evenly or oddly, it is enough if for every edge/vertex pair e, v with $x(e_1) \leq x(v) \leq x(e_2)$, we know whether whether e passes above or below v. For two edges that share an endpoint, we also need to know the rotation at that endpoint, but no more. Thus, we will assume that an x-monotone even drawing is encoded by the above/below information for all relevant edge/vertex pairs, the rotation system, and the x-coordinate order of the vertices. In what follows we assume that the above/below information for edge/vertex pairs and rotation systems are stored in an $|E| \times |V|$ matrix M, i.e. an entry corresponding to a vertex v and an edge e encodes one the following three states: e passes above v, e passes below v, or e is not passing by v. We store the rotation systems of the vertices in linked lists, i.e. each vertex stores its left and right rotation system in a separate linked list. For a graph or a multigraph $G = (V, E)$, this representation has size $O(|V||E|)$.

If we have no odd pair in a monotone drawing of our ordered multigraph G, there exists by Theorem 5.1 a monotone embedding of G with the same rotation systems. Thus, the only thing of our representation of G that could change in an embedding of G is the above-below relationship of the edges and vertices. Note that once we have computed the above-below relationship of the edges and vertices in G yielding an actual embedding, it is easy to construct for each vertex v the order in which the edges pass above and below v in our embedding in $O(|V|^2)$ time. Indeed, we can sweep the embedding from left to right while keeping the top-bottom order of the edges which is updated each time we visit a vertex. Clearly, the top-bottom order of the edges and vertices gives us also the face structure of the embedding of G. Hence, by the next lemma we can say that Theorem 5.1 can be made algorithmic and its running time is $O(|V|^2)$.

Lemma 5.4. *Given an x-monotone even drawing of a multigraph G, we can find an x-monotone embedding with the same vertex locations and rotation system in time $O(|V|^2)$.*

Proof. By the remark preceding the lemma, it is enough to update the matrix M so that the above-below relationship of the vertices, which M defines, yields an embedding. Note that the only place in the proof of Theorem 5.1 where the above-below relation of the edges and vertices is changed is the application of Lemma 5.2. It turns out that a straightforward implementation of the proof would give a super quadratic running time, because each application of Lemma 5.2 could possibly change $O(|V|^2)$ entries of M. Hence, the main obstacle to an efficient algorithm lies in the repeated applications of Lemma 5.2, which re-embeds parts of the graph. However, in a careful implementation of Theorem 5.1 for connected graphs we can avoid using Lemma 5.2 altogether.

Let w denote a closed walk which bounds a face of an embedding of G incident to its leftmost vertex x_1, such that x_1 is a cut-vertex of G. Let M_f denote the vertex of w with the maximal x coordinate. We avoid using the lemma by inserting an x-monotone edge x_1M_f at x_1 and M_f into the wedges contained in f in the case there exists a subgraph H (as defined in the proof of Theorem 5.1) of G containing x_1 that will be inserted into f. The fact that each such edge can be added in the way that it does not create an odd crossing pair is guaranteed by Lemma 5.5. Since we store rotation systems at vertices in linked lists, we can traverse the boundary of a face in a linear time in order to find M_f. As we have not enough time to update M for all such additional edges at the moment of their insertion, we will update the entries in M corresponding to these edges later.

Next, we analyze the proof of Theorem 5.1 and discuss the running time of its individual steps. Thus, we assume that we are given a monotone drawing of an ordered graph G encoded in the way as discussed above. We want to produce a matrix M giving an embedding of G with the same rotation systems if the condition of Theorem 5.1 is satisfied. Note that to check this condition can be done easily in a quadratic time as the number of edges of G cannot be more than $O(|V|)$. Each pair of edges not sharing a vertex can be checked for the parity of crossings in a constant time, and for each vertex v we can check every pair of edges sharing v for the parity of crossings in $O(d_v^2)$ time.

First, we discuss the case when we can contract the edge between two leftmost vertices of G. If such an edge already exists we simply insert the right rotation system at v_2 to the right rotation system of x_1. We identify new edges with the corresponding columns of M and put on the stack the information needed to recover G, as it was before the contraction. This can be clearly carried out in $O(|V|)$ time. Analogously, we perform a de-contraction of an edge by splitting the leftmost vertex into two leftmost vertices.

If there exists no edge between two leftmost vertices we need to check whether such an edge can be added to G. In other words, we need to know if there exists a pair of edges e and f, resp., incident to x_1 passing above and below, resp., x_2 such that e is below f in the rotation at x_1. This can be done in $O(|V|)$ by using matrix M.

Second, we assume that there exists no edge between x_1 and x_2 and that such an edge cannot be even added without violating the condition of Theorem 5.1, which means that x_1 is a cut vertex of G at this step. Thus, we need to identify all connected components $H_1', \ldots H_k'$ of $G \setminus \{x_1\}$. We recall that $H_i' = H_i \setminus \{x_1\}$. We construct a directed forest F having H_i-s as the vertices, in which we join H_i and H_j by an edge directed from H_i to H_j if $H_i \prec H_j$. For a pair $H_i \prec H_j$ let w denote the walk bounding the face f of H_j, into which we will insert H_i. Let M_f denote the vertex of w with the largest x-coordinate. We add into the rotation system at x_1 and M_f an edge x_1M_f so that its corresponding wedges in the rotation systems belong to f. We call the edge x_1M_f *special*. Let H_i'' denote the union H_i with the special edges.

We recursively run our algorithm for each H_i, $1 \leq i \leq k$. Note that

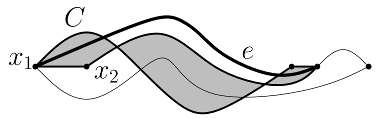

Figure 21: Deciding the above-below relationship of x_2 and the special edge e.

each time we contract the edge x_1x_2, we have to decide for each special edge whether it passes above or below x_2 in order to merge correctly the rotation systems at x_1 and x_2. This is easy, if an edge between x_1 and x_2 already exists. Refer to Figure 21. Otherwise, we add an edge between x_1 and x_2 while ignoring the special edges. There exists a path P from x_2 to x_1 in $G \setminus \{x_1\}$ not using the added edge x_1x_2 and giving rise to a cycle $C = Px_1x_2$. By Lemma 2.4, we can two color the complement C' of C in $\mathbb{R}^2$ such that no two adjacent connected parts of C' receives the same color. By a color of a vertex of $G \setminus C$ we understand the color of the connected part of C' it belongs to. We can compute in $O(|V|)$ time the color of each vertex of $G \setminus C$, since the color of a vertex v is the same as the color of the connected part of $\mathbb{R}^2 \setminus C$ containing the initial piece of a path from C to v. Then by the color of the right endpoint of a special edge e (resp. the rightmost piece of e in case the right endpoint is on C) we know to determine whether e passes above or below x_2.

It remains to combine the even drawings of H_i'' by using the special edges. Note that this can be done in a way so that each entry of M is changed at most once. Indeed, for a pair $H_i \prec H_j$ all the entries corresponding to pairs consisting of a vertex v of H_j and an edge e of H_i are updated so that the corresponding entry is the same as the entry for the pair consisting of v and the special edge corresponding to H_i. Similarly, the entries corresponding to pairs consisting of a vertex v of H_i and an edge e of H_j are opposite to the entry of the pair consisting of the special edge corresponding to H_i and an endpoint of e. The corresponding entries between the components not joined by a directed edge are updated according to the rotation system at x_1.

To prove that the running time of our algorithm is $O(|V|^2)$ is straightforward as we proceed in $2|V|$ contraction and de-contraction steps each of which has the time complexity of $O(|V|)$ and updating the matrix M takes at most $O(|V|^2)$. $\square$

The assumption in Theorem 5.1 can be weakened, somewhat surprisingly, replacing x-monotonicity of edges by a weaker notion. Recall that an edge uv in a drawing is *x-bounded* if every interior point p of uv satisfies $x(u) < x(p) < x(v)$. That is, an edge is x-bounded if it lies strictly between its endpoints; it need not be x-monotone within those bounds.

Lemma 5.5. *Suppose we are given a drawing of a graph G with an x-bounded edge e. Then e can be redrawn, without changing the remainder of the drawing or the position of e in the rotations of its endpoints, so that e is x-monotone and the parity of crossing between e and any other edge of G has not changed.*

Proof. Suppose that $e = ab$ and let v be an arbitrary vertex between a and b: $x(a) < x(v) < x(b)$. Since e connects a to b it has to cross the line $x = x(v)$ an odd number of times. Consequently, e crosses one of the two parts into which v splits $x = x(v)$ evenly, and e crosses the other part oddly. In a small neighborhood of $x = x(v)$, redraw G by pushing all crossings of e with $x = x(v)$ from the even side across v to the odd side (see top and middle part of Figure 22). Note that the odd side of $x = x(v)$ remains odd and there are no crossing with e left on the even side. Moreover, the parity of crossing between e and any other edge does not change since e is moved an even number of times across v. Repeat this for all v between a and b; now e only passes above or below each such v, never both. We can now deform e into an x-monotone edge connecting a and b, without having the edge pass over any vertices. Since the deformation does not pass over any vertex, it does not affect the parity of crossing between e and any other edge. This means we have found the redrawing required by the lemma (see middle and bottom part of Figure 22).

In hindsight we see that the redrawing in Lemma 5.5 can be done quite efficiently: for each vertex v between a and b we only need to know whether e passes oddly above or below it, and we can build a polygonal arc from a to b that passes each vertex on the odd side. □

Redrawing one edge at a time using Lemma 5.5 gives us the following strengthening of Theorem 5.1. Later, we will use that result to strengthen Theorem 5.7, and to show that x-monotone edges can be replaced by x-bounded edges in the definition of level planarity (see Corollary 5.15 in Section 5.6).

Corollary 5.6. *If G has an even drawing in which every edge is x-bounded, then G has an x-monotone embedding with the same vertex locations and rotation system.*

5.4 Strong Hanani-Tutte for Monotone Drawings

Pach and Tóth [50] wrote "It is an interesting open problem to decide whether [the conclusion of Theorem 5.1] remains true under the weaker as-

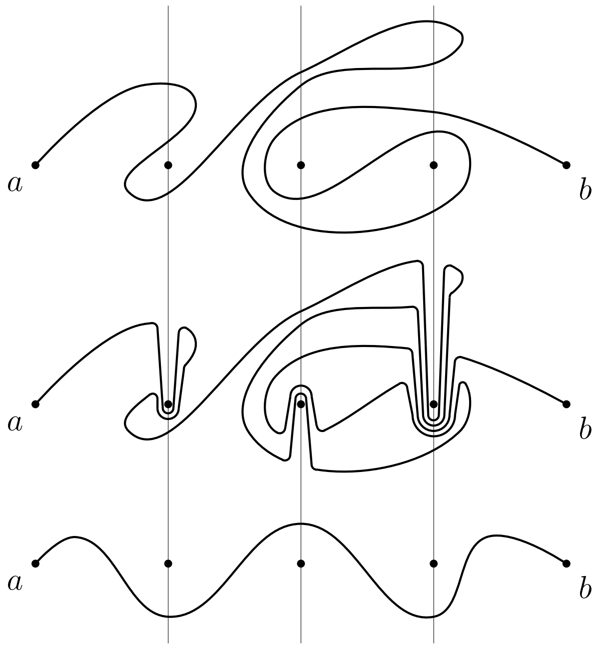

Figure 22: How to redraw an x-bounded edge. (*top*) Before the redrawing. (*middle*) After pushing e off the odd parts. (*bottom*) After deforming e into an x-monotone drawing.

sumption that any two *non-adjacent* edges cross an an even number of times." The goal of this section is to establish this result.

Theorem 5.7. *If G has an x-monotone drawing in which every pair of independent edges crosses evenly, then G has an x-monotone embedding with the same vertex locations.*

Remark 5.8. *As in the case of Theorem 5.1, the statement of Theorem 5.7 remains true if we only require edges to be x-bounded rather than x-monotone: simply redraw edges one at a time using Lemma 5.5, before applying Theorem 5.7.*

In a proof of the standard Hanani-Tutte theorem, it is obvious that a minimal counterexample has to be 2-connected, since embedded subgraphs can be merged at a cut-vertex. Unfortunately, the merge requires a redrawing that does not maintain monotonicity, so here we must use structural properties that are more tailored to x-monotone redrawings. For a subgraph H of G let $N(H)$ denote the set of neighbors of vertices of H in $G - V(H)$, that is, $N(H) := \{u : uv \in E(G), v \in V(H), u \in V(G) - V(H)\}$.

Lemma 5.9. *Suppose that G is a smallest (fewest vertices) counterexample to Theorem 5.7. Then:*

(i) *G is connected.*

(ii) *G has no connected subgraph H and vertices $a, b \in V(G) - V(H)$ such that $x(a) < x(v) < x(b)$ for all $v \in V(H)$, $N(H) = \{a, b\}$, and $V(G) - (V(H) \cup \{a, b\}) \neq \emptyset$.*

(iii) *If G has a cut-vertex a and $G - \{a\}$ has a component H such that $x(a) < x(v)$ for all $v \in V(H)$, then H has only one vertex b, and G has no edge ac with $x(b) < x(c)$. Also, in this case G has no connected subgraph $H' \neq \emptyset$ so that $x(a) < x(v) < x(b)$ for all $v \in V(H')$, $a \in N(H') \neq \{a\}$, and $x(v) > x(b)$ for all $v \in N(H') - \{a\}$.*

Proof. If a smallest counterexample G is not connected, none of its components are counterexamples to Theorem 5.7. But then we could embed each component separately and stack the drawings vertically so they do not intersect each other, yielding an embedding of G. This contradiction establishes (i).

Consider case (ii). Since G is a smallest counterexample, both $G - V(H)$ and $G[V(H) \cup \{a, b\}]$ have embeddings (both graphs are smaller than G by assumption). We can deform the crossing-free drawing of $G[V(H) \cup \{a, b\}]$ so that it becomes very flat. If $ab \in E(G)$ we can then insert this drawing into the drawing of $G - V(H)$ near the edge ab, without adding crossings. This gives us a crossing-free drawing of G, which is a contradiction. If $ab \notin E(G)$ then we add ab to the drawing of $G - V(H)$ so that it has no independent odd crossings (we will presently see how this can be done); the resulting $G - V(H) \cup \{ab\}$ has fewer vertices than G so it also has an embedding, and we can proceed as in the case that $ab \in E(G)$, removing the edge ab in the end.

When $ab \notin E(G)$, here is how we draw the edge ab with no independent odd crossings: Let P be any a, b-path with interior vertices in H. By suppressing the interior vertices of P, we can consider it an x-bounded edge (in the sense defined earlier) between a and b, so Lemma 5.5 tells us that we can draw an x-monotone edge that has the same parity of crossing with all edges of $G - V(H)$ not incident to a and b as does P. Hence, we do not create any independent odd crossing pair of edges.

Finally, we consider (iii), where H is a component of $G - \{a\}$ so that $x(a) < x(v)$ for all $v \in V(H)$. Let b be the vertex with the largest x-value in H. If $|V(H)| > 1$, then we have case (ii) using $H := H - b$. Therefore $|V(H)| = 1$ and $V(H) = \{b\}$. If G has an edge ac with $x(b) < x(c)$, we can first embed $G - \{b\}$ (since it is smaller than G), and then add ab and b to the embedding alongside of ac without crossings.

It remains to consider a connected subgraph $H' \neq \emptyset$ so that $x(a) < x(v) < x(b)$ for all $v \in V(H')$, $a \in N(H') \neq \{a\}$, and $x(v) > x(b)$ for all $v \in N(H') - \{a\}$. If there is an edge e not in H' with endpoints in H', we can replace H' be $H' \cup \{e\}$ and it still satisfies all the conditions; thus we may assume that H' contains all such edges, i.e., that H' is an induced subgraph of G. By minimality, $G - \{b\}$ has an embedding. Of all the edges from a to H', let au be the one that is lowest in the rotation at a. Let f be the face in the drawing of G that lies immediately below au. Follow the boundary of f from a to u until it exits H' to a vertex c not in H'. If $c = a$ then H' could not have any neighbors v with $x(v) > x(b)$, a contradiction. The only other possibility is that $x(c) > x(b)$. Then by Lemma 5.2, we can add the edge ac to $G - \{b\}$ and obtain an embedding without introducing crossings. Since $x(a) < x(b) < x(c)$, we can instead add ab to the drawing without crossings, so G has an embedding which is a contradiction. $\square$

The proof of Theorem 5.7 now proceeds by induction on the number of *odd pairs* (pairs of edges that cross an odd number of times). Roughly speaking: If we encounter an odd pair (by necessity its edges are adjacent), we can either make it cross evenly or we are in a situation which has been excluded by Lemma 5.9. To realize this goal, we need additional intermediate results. These results are not about smallest counterexamples, but are true in general.

For the lemmas we introduce some new terminology generalizing our usual notion of lying above or below a curve to curves with self-intersections: Let C be a curve in the plane with endpoints p and r so that for every point $c \in C - \{p, r\}$, $x(p) < x(c) < x(r)$. (This is similar to the definition of an x-bounded edge except that we allow self-intersections.) Suppose that q is a point for which $x(p) \leq x(q) \leq x(r)$. Extend C via a horizontal ray from p to $x = -\infty$ and a horizontal ray from r to $x = \infty$, and consider the plane $\mathbb{R}^2$ minus that extended curve. By Lemma 2.4, we can 2-color its faces so that adjacent faces (faces whose boundaries intersect in a nontrivial curve) have opposite colors. We say that q is *above* (*below*) C if q lies in a face with the same color as the upper (lower) unbounded region. We will proceed with the proof of Theorem 5.7.

Proof of Theorem 5.7. Let G be a smallest (fewest vertices) counterexample to the theorem. By Lemma 5.9(i), G is connected. Fix an x-monotone drawing of G with the same vertex locations, which minimizes the number of odd pairs (that is, the number of pairs of edges crossing oddly). If there are no odd pairs, then Theorem 5.1 completes the proof.

Suppose that there are edges e_1 and e_2 that cross oddly. Then e_1 and e_2 have a shared endpoint v_0, and we may assume that v_0 is the left endpoint of e_1 and e_2. Choose e_1 and e_2 so that their ends at v_0 have minimum distance in the right rotation at v_0, with e_1 above (that is, preceding) e_2. Then e_1 and e_2 are not consecutive in the rotation at v_0; if they were, they could be redrawn so that they cross once more near v_0, by switching their order in the rotation at v_0; this contradicts the choice of drawing of G. So there is at least one edge incident to v_0 that lies between e_1 and e_2 in the rotation at v_0, and by minimality, all such edges cross each other evenly and cross both e_1 and e_2 evenly. Pick one such edge, e_3. Let v_1, v_2, v_3 be the right endpoints of e_1, e_2, e_3, respectively, and let G_0 be the subgraph of G induced by all vertices v fulfilling $x(v) \geq x(v_0)$.

We inductively construct three subgraphs G_i of G_0 that contains e_i for $i = 1, 2, 3$. We consider these indices to be in $\mathbb{Z}_3$, in particular $3 + 1 \equiv 1$ and $1 - 1 \equiv 3$.

At each step, let v_i denote the vertex of G_i with the maximum x-coordinate, for $i = 1, 2, 3$, and let $\{l, c, b\} = \{1, 2, 3\}$ such that $x(v_b) < x(v_c) < x(v_l)$. We call G_l the *leading* graph, G_c the *chasing* graph, and G_b the *left behind* graph. Thus, b, c and l are not fixed throughout our process.

We construct G_i, $i = 1, 2, 3$, in the two consecutive stages. The second stage might be omitted, though. We maintain properties (1), (2), (3), (4), and (5) throughout Stage 1. During Stage 2, we maintain properties (1), (4), and (5).

(1) Each vertex of $G_b - v_0$ is below every v_0, v_{b+1}-path (in the sense as defined above) in G_{b+1} and above every v_0, v_{b-1}-path (in the sense as defined above) in G_{b-1}.

(2) $G_1 \cap G_2 \cap G_3 = \{v_0\}$.

(3) v_c is below the edge in G_l incident to v_l if $l = c + 1$ and v_c is above the edge in G_l incident to v_l if $l = c - 1$, and there is no vertex in $G_c \cup G_l$ between v_c and v_l.

(4) For every vertex v in $G_1 \cup G_2 \cup G_3$, $x(v) \geq x(v_0)$.

(5) $G_i - v_0$ is connected, for all i.

Initialize each G_i to have e_i and both its endpoints; (2), (4) and (5) are obvious, and we can check that (1) and (3) hold by case analysis.

Stage 1: (see Figure 23 for an illustration)

Refer to Figure 24. At each step, we pick an edge uw such that $u \neq v_0$, $u \in V(G_b)$, and $uw \notin E(G_b)$. We add uw and w to G_b. (We do not need to

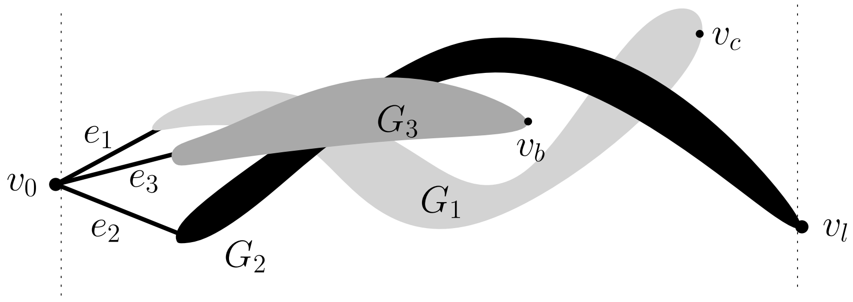

Figure 23: Stage 1

add w to G_b if G_b already contained w.) If $w = v_l$, we proceed to Stage 2. Otherwise, in the next step we continue with Stage 1.

If such an edge does not exist, we add recursively to G_c and G_l all the neighbors of their vertices, except for v_0, with the x coordinate less than $x(v_b)$, i.e. we add to G_i; $i \in \{c, l\}$, every edge uw, such that $u \neq v_0$, $u \in V(G_i)$, $x(v_0) < x(u) < x(v_b)$, and $uw \notin E(G_i)$, and do not proceed with Stage 2.

After each step of Stage 1, property (5) clearly remains true. We have to show that after each step of Stage 1, properties (1) and (4) remain true, and if $w \neq v_l$ we also need to show that properties (2) and (3) remain true:

First, consider the case $x(w) \leq x(v_c)$. Then property (3) clearly remains true.

To prove (2), suppose for a contradiction that $w \in G_j - v_0$ with $j \neq b$. Let k be the unique index in $\{1, 2, 3\} \setminus \{b, j\}$, and let P_k be a v_0, v_k-path in G_k. By induction on property (1), u is below P_k if $k = b + 1$ and u is above P_k if $k = b - 1$. Since uw is not incident to P_k, w is also below P_k if $k = b + 1$ and w is above P_k if $k = b - 1$. There exists a path P_j in $G_j \setminus \{v_0\}$ from w to v_j. Hence, if $j = c$, as $G_j \cap G_k = \{v_0\}$ (induction on property (2)), v_j is below P_k, if $k = b + 1$ and v_j is above P_k, if $k = b - 1$. However, this contradicts (3). On the other hand, if $j = l$, $v_c = v_k$ is above the last edge on P_j, if $k = b + 1$ and $v_c = v_k$ is below the last edge on P_j, if $k = b - 1$, which again contradicts (3).

To prove (1), let P_j be a v_0, v_j-path in G_j for $j \in \{b - 1, b + 1\}$. By induction on property (1), u is below P_{j+1} and above P_{j-1}. If $x(w) > x(v_0)$, then since uw must cross each of those paths evenly, w must also be below P_{j+1} and above P_{j-1}. If $x(w) < x(v_0)$, then let w' be a point (not vertex) on uw near the line $x = x(v_0)$, such that the uw'-portion of the curve uw lies entirely to the right of $x = x(v_0)$. Then, by the same argument as before, w' must be below P_{j+1} and above P_{j-1}. But then uw must pass below v_0 and above v_0, a contradiction. This proves (1), and (4) as well.

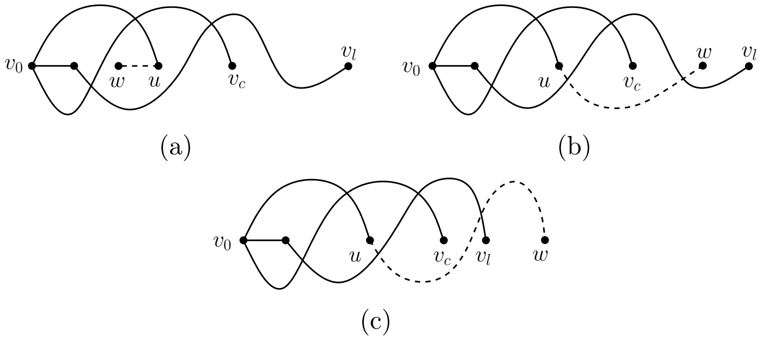

Figure 24: A step during Stage 1: (a) The indices b, c and l will not be changed; (b) The vertex w will become v_c; (c) The vertex w will become v_l.

Now, consider the case $x(w) > x(v_c)$. Property (4) is obviously holds.

In the next round G_c will become G_b, so we must prove (1) for it: Let P_c be a v_0, v_c-path in G_c. By induction on property (1), u lies below P_c if $c = b+1$ and u lies above P_c if $c = b-1$. Since w is not on P_c, uw crosses P_c evenly, and (by a similar argument to the above argument using w') uw passes below v_c if $c = b+1$ and uw passes above v_c if $c = b-1$. Therefore, once we relabel G_c as G_b by induction on (3), its vertex v_c (relabeled as v_b) will satisfy property (1).

Since G_c will become G_b, all of its vertices are to the left of the maximal vertices of G_{c+1} and G_{c-1} after uw and w are added to G_b (which equals G_{c+1} or G_{c-1}). By induction on (4), all vertices of $G_c - v_0$ lie to the right of $x(v_0)$. Let P'_{c+1} and P'_{c-1} be paths from v_0 to the maximal vertices of G_{c+1} and G_{c-1} after uw and w are added; they are disjoint from $G_c - v_0$ by induction on (2). Therefore any path in $G_c - v_0$ is even with respect to P'_{c+1} and P'_{c-1}, and all its vertices are either above P'_{c+1} or below P'_{c+1}. $G_c - v_0$ is connected by induction on (5), so every vertex of $G_c - v_0$ is joined to v_c by a path in $G_c - v_0$. Then, since v_c will satisfy (1) after relabeling, every vertex of $G_c - v_0$ will satisfy (1) after relabeling.

Still within the case $x(w) > x(v_c)$, suppose that $w \neq v_l$. We need to show that (2) and (3) hold. By induction on (3), $G_1 \cup G_2 \cup G_3$ has no vertices between $x = x(v_c)$ and $x = x(v_l)$, so w does not lie in $G_1 \cup G_2 \cup G_3$; so (2) will remain true.

If $x(w) < x(v_l)$, then after relabeling w will become the new v_c. If $x(w) > x(v_l)$, then after relabeling w will become the new v_l, and v_l will become the new v_c. Let P_l be a v_0, v_l-path. By induction on (1), u is below P_l if $l = b+1$ and u is above P_l if $l = b-1$. If $x(w) < x(v_l)$, then uw crosses P_l evenly, so w is also below P_l if $l = b+1$ and above P_l if $l = b-1$, which implies (3) (for the relabeled graphs). If $x(w) > x(v_l)$, then let w' be a point on uw slightly to the left of $x(v_l)$, so that the uw'-portion of the curve uw lies entirely to the left of $x = x(v_l)$. The same sort of argument shows that w' lies below P_l if $l = b+1$ and above P_l if $l = b-1$. Therefore uw must pass below v_l if $l = b+1$ and above v_l if $l = b-1$; which implies (3) (for the relabeled graphs).

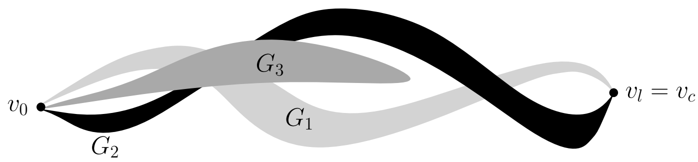

(a)

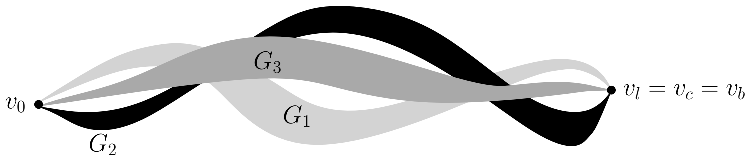

(b)

Figure 25: (a) The situation at the beginning of Stage 2; (b) The end of Stage 2 if $G_1 \cap G_2 \cap G_3 = \{v_0, v_l\}$.

Stage 2: (see Figure 25 for an illustration)

At the end of Stage 1, we had $w = v_l$. After uw and w are added to G_b, the value of v_b will update and it will equal w. At that point we have $x(v_c) < x(v_l) = x(v_b)$, at which point b is updated to take the value of c, (1) is shown for the updated value of b, and then Stage 1 finally ends.

For convenience of intuition, let's also update the value of c so that it takes the old value of b, so we still have $\{b, c, l\} = \{1, 2, 3\}$. Then we have $x(v_b) < x(v_c) = x(v_l)$ and $v_c = v_l$. (So c and l do not have the same meaning that they had during Stage 1.)

Since (2) was true until the end of Stage 1, we now have $G_c \cap G_l = \{v_0, v_l\}$ and $G_b \cap G_c = \{v_0\} = G_b \cap G_l$. Note that stronger version of (4) is true:

(4') For every vertex v in $G_1 \cup G_2 \cup G_3$, $x(v_0) \leq x(v) \leq x(v_l)$.

Also by the fact that (3) holds in the end of Stage 1 (using the same argument as above) we get the following modification of (3') easily maintained during the second stage.

(3') Let $v_l \in G_i, G_j$, $i \neq j$. Any vertex of $G_i \setminus \{v_0, v_l\}$, is below (resp. above) a v_0, v_l-path in G_j if $i = j-1$ (resp. $i = j+1$).

(1), (3') and (4') imply the following modification of (2) maintained during Stage 2.

(2') $G_i \cap G_j \subseteq \{v_0, v_l\}$ for each $i, j \in \{1, 2, 3\}, i \neq j$.

At each step of Stage 2, we pick an edge uw such that $u \in V(G_b) \setminus \{v_0, v_l\}$ and $uw \notin E(G_b)$, add uw and (if w is not already in G_b) w to G_b. If there exists no vertex that could be added to G_b and $v_l \in G_b$, we continue Stage 2 with G_c and G_l, i.e. at each step we pick an edge uw such that $u \in V(G_i) \setminus \{v_0, v_l\}$ and $uw \notin E(G_i)$, add uw (if w is not already in G_i) and w to G_i, $i \in \{l, c\}$. Note that (5) clearly remains true, and that (1) immediately implies that (4') remains true. Thus, we must check that (1) remains true.

The previous proof of (1) in the case $x(w) < x(v_c)$ is easily modified to prove (1) in the case that $x(v_0) < x(w) < x(v_l)$, and to show a contradiction when $x(w) < x(v_0)$ or $x(w) > x(v_l)$. For the contradiction, let w' be a point close to $x = x(v_0)$ or $x = x(v_l)$ such that the u, w'-portion of the curve uw lies entirely between $x = x(v_0)$ and $x = x(v_l)$. However, similarly as above we again get that uw must pass below and above v_l (resp. v_0), which is a contradiction. This completes our discussion of the steps in Stage 2.

After the process is over: Let G_c, G_l and G_b denote the chasing, leading and left-behind graph, respectively, at the end of the process.

If the process ends with Stage 2 then at least one of the graphs G_l or G_c, let's say G_l, is not just a single edge. Hence, by property (4') the graph G_l violates Lemma 5.9(ii).

Otherwise the process ends with Stage 1 and we are done by Lemma 5.9(iii) with G_b playing the role of H and the subgraph of G_c (or G_l) induced by the vertices with the x coordinate less than $x(v_b)$ playing the role of H'. □

Similarly as in case of Theorem 5.1 we are also interested in the algorithmic version of Theorem 5.7. In other words, given a monotone drawing of an ordered graph G in which every pair of independent edges cross an even number of times, we want to find a monotone embedding of G. We encode the drawing of G in the same way as before, i.e. our input consists of $G = (V, E)$ where each vertex stores separate linked lists of its left and right neighbors, an $|E| \times |V|$ matrix M storing the above-below relation of the edges and vertices, and the order of the vertices $x_1, \ldots, x_n$ along the x-axis. Note that we are free to choose rotation systems at the vertices of G in the input graph. Then by the discussion preceding the proof of Lemma 5.4 it is sufficient that our output consists only of the updated matrix M and the rotation systems at vertices of G yielding an embedding.

Lemma 5.10. *Given an x-monotone independently even drawing of G, we can find an x-monotone embedding with the same vertex locations and rotation system in time $O(|V|^2)$.*

Proof. Without loss of generality we assume that G is connected. Since checking the condition of Theorem 5.7 can be done in $O(|V|^2)$ by the same token as in case of Theorem 5.1 we can proceed to the the construction of the embedding. We will closely follow the proof of Theorem 5.7.

First, we need to choose the rotation systems at the vertices of G so that no two consecutive edges in the left or right rotation system at any vertex form an odd pair. To this end we construct for the right (resp. left) rotation system at each vertex v a complete directed graph $G_v = (V_v, E_v)$, so called tournament, the vertices of G_v are the edges leaving v towards right (resp. left), and we join $e_1, e_2 \in V_v$ by an edge oriented from e_1 to e_2 if and only if e_1 must be in the left (resp. right) rotation system at v above e_2 in order to form with e_2 an even pair. The desired rotation system at v corresponds to a directed Hamiltolian path in G_v whose existence is guaranteed by a result of Rédei [65]. Moreover, such a path can be found in $O(|V_v|^2)$ time (by the proof in [65])*. Hence, we can choose the desired rotation systems for G in $O(|V|^2)$ time.

Second, we will produce a set $\mathcal{G}$ of even drawings of subgraphs of G with some additional special edges. Thus, for each element $G' = (V', E') \in \mathcal{G}$ we have $V' \subseteq V$, and $E' = E_0 \cup E_s$, where $E_0 \subseteq E$ and for each $e = x_i x_j \in E_s$, $i < j$, there exists a *corresponding graph* $G'' = (V'', E'') \in \mathcal{G}$ such that the leftmost (resp. rightmost) vertex of G'' is x_i and the rightmost (resp. the leftmost) vertex of G'' is to the left of x_j or equal to x_j. Moreover,

(i) if $V(G') \cap V(G'') \neq \emptyset$, then $|V(G') \cap V(G'')| \leq 2; V(G') \cap V(G'') \neq V(G'), V(G'')$, and all vertices in the intersection are the endpoints of a special edge of G' (or G'') that corresponds to G'' (or G');

(ii) $\bigcup_{G' \in \mathcal{G}} V(G') = V = V(G)$ and $\bigcup_{G' \in \mathcal{G}} E(G') \supseteq E = E(G)$.

By the convexity of a quadratic function once we obtain even drawings of all graphs in $\mathcal{G}$ by Lemma 5.4, we can get their embeddings in $O(|V|^2)$ time. Finally, we can obtain an embedding of G from the embedded graphs in $\mathcal{G}$ in $O(|V|^2)$ time by the technique used in the proof of Lemma 5.4. To this end we replace special edges by corresponding graphs, which amounts to updating the matrix M and rotation systems at vertices. The required time complexity follows, since we perform $O(|V|)$ many replacements of special edges, each of which takes $O(|V|)$ time, if we do not count the update of the matrix M. Moreover, each entry of M is accessed at most once. It remains to show how to obtain $\mathcal{G}$ and even drawings of graphs in $\mathcal{G}$ in $O(|V|^2)$ time.

*There exists $O(n \log n)$ algorithm for this task, but this algorithm would not help us to reduce the overall running time.

If G does not contain an odd edge pair we are done as $\mathcal{G} = \{G\}$. Otherwise, by the proof of Theorem 5.7 using a depth-first search we obtain graphs G_1, G_2 and G_3; $\{v_0, v_l\} \supset G_1 \cap G_2 \cap G_3$, such that either to one of them, let's say, G_1 we can apply Lemma 5.9(iii) or to $G_1 \cup G_2 \cup G_3$ (resp. w.l.o.g. $G_1 \cup G_2$) we can apply Lemma 5.9(ii).

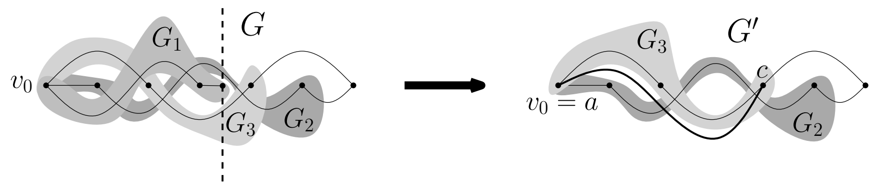

Figure 26: Constructing G' from G.

Ref. to Figure 26. In the former, we recurse on G_1 and $G \setminus (G_1 \setminus v_0)$, where v_0 is the leftmost (resp. rightmost) vertex of G_1. At this point we do not add the special edge $ac = v_0c$ as defined in Lemma 5.9(iii), partially because we do not know yet which vertex can play the role of c and partially because we want it to form an even pair with every other edge. We also need to shorten the right (resp. left) rotation system of $G \setminus (G_1 \setminus v_0)$. Here, we do not need to copy the matrix M. We store for G_1 only the rotation systems at its vertices and the list of edges.

Note that $G \setminus (G_1 \setminus v_0)$ might contain consecutive edges in the rotation system at v_0 crossing an odd number odd times. Hence, before the recursive step we, first, recursively switch the order of such edge pairs. Since each time we perform such a switch, we decrease the number of odd pairs of edges by one, overall this adjustment takes at most $O(|V|^2)$ steps.

Let $G' \in \mathcal{G}'$ denote the unique graph not corresponding to any special edge in the set of graphs $\mathcal{G}'$ returned for $G \setminus (G_1 \setminus x_i)$. Now, we can add to G' the special edge ac as defined in Lemma 5.9(iii) that G_1 will correspond to (if G' does not already contain this edge), since we have an even drawing of G' and by Theorem 5.1 from now on rotation systems in the drawing of G' remain fixed. Indeed, we can simply traverse the edges on the boundary of the face immediately below G_2 (or G_3) until we encounter c. This can be done in linear time, if G' still contains G_2 or G_3 (constructed above). Otherwise, G' contains a subgraph corresponding to G_2 and G_3. Thus, we need to find a wedge at v_0 contained in the face immediately below the subgraph corresponding to either G_2 or G_3 instead. To this end we need to keep for every edge of $G \setminus G'$ a pointer to an edge, it corresponds to in G'. Thus, for each edge of G that belongs to G_1 we store a pointer to $v_0c = ac$.

Now, we need to find the last edge in the right (resp. left) rotation system at v_0 in G', which either belongs to G_2 or to which an edge of G_2 is pointing. Clearly, this can be carried out in $O(|V|)$ time.

In what follows we show how to obtain the above-below relation of the vertices of G' with respect to ac, so that the resulting drawing of G' will be even, and update the matrix M in $O(|V|)$ time. For this we need to find out, which vertices of G' between a and c are reachable by a topological path P of G' that remains between a and c, such that P starts either at a point with the x-coordinate equal to $x(a)$ or at a point with x-coordinate equal to $x(c)$ and its initial piece is above ac. Our approach is justified by the fact that the drawing of G' is even. Thus, P starts above ac if and only if it ends above ac. Hence, the task amounts to performing a depth-first search in G. Finally, we return the union of $\mathcal{G}'$ and $\mathcal{G}''$, such that $\mathcal{G}''$ corresponds to G_1, with the corresponding even drawings of the graphs in this union.

In case of an application of Lemma 5.9(ii) we proceed analogously. At least one, let's say G_1, of the graphs G_1, G_2 and G_3 is not a single edge and its rightmost (resp. leftmost) vertex corresponds to v_l at the end of the Stage 2. Thus, in this case we recurse on G_1 and $G \setminus (G_1 \setminus v_0)$ and proceed as in the previous case so that v_0v_l will be a special edge added to G' corresponding to G_1.

Note that the time complexity of the above routine can be upper bounded by a function f satisfying: $f(n) \leq f(k) + f(n-k) + O(n)$, $k \geq 1$. Hence, we get the total time complexity $O(|V|^2)$ and that concludes the proof. □

5.5 The x-Monotonicity Testing and Layout

The strong Hanani-Tutte theorem can be viewed as an algebraic characterization of planarity: testing whether a graph is planar can be recast as solving a system of linear equations. Unfortunately, the system has $|E| \cdot |V| = O(|V|^2)$ variables which leads to an impractical $O(|V|^6)$ running time. There are linear-time algorithms for planarity testing based on a Hanani-Tutte-like characterization, but they do not take the algebraic route [15, 12].

Similarly, Theorem 5.7 can be viewed as an algebraic criterion for testing whether a graph has an x-monotone embedding, for a given x-coordinate order of the vertices. However, unlike the system of linear equations for planarity, the equations for x-monotonicity are so simple that solvability can be checked directly in quadratic time. In what follows we present the details of this algorithm. In Section 5.6 we will see how to extend the algorithm to recognizing level-planar graphs, so we obtain a very simple, quadratic-time algorithm for level-planarity testing. Linear time algorithms for this task are known, but are quite complex. We will also discuss the rather confusing situation of algorithms for level-planarity testing in more detail. Finally, we also show embedding algorithms for both x-monotone and level-planar graphs.

How can we use Theorem 5.7 to test whether a given graph G with x-coordinates assigned to the vertices has an x-monotone embedding? Let D be an x-monotone embedding of G and let D' be an x-monotone drawing of G on the same vertex set. Pick any edge e in D' and continuously transform it into its drawing in D; we can assume that the edge remains x-monotone during the transformation. As the edge changes, its parity of intersection with any independent edge only changes when it passes over a vertex v (at which point its parity of intersection with every edge incident to v changes). The same effect can be achieved by making an (e, v)-*move*: Take e, and close to $x = x(v)$ deform it into a spike that passes around v. In other words: if G has an x-monotone embedding then there is a set of (e, v)-moves that turns D' into a drawing in which every pair of independent edges crosses evenly. Since the reverse is also true, by Theorem 5.7, we now have an efficient test.

Theorem 5.11. *Given a graph G and a placement of the vertices of G in the plane, we can test in time $O(|V|^2)$ whether G has an x-monotone drawing on that vertex set.*

Proof. If G has an x-monotone embedding on the given vertex set, then no two vertices lie on a vertical line. As discussed in Section 5.2, we can deform the plane so that the vertices are located at $(1, 0)$, …, $(|V|, 0)$, and the drawing will remain x-monotone—but it will remain an embedding as well. Thus, we can assume that the vertices are located at $(1, 0)$, …, $(|V|, 0)$.

Now draw each edge as a monotone arc above $y = 0$. Note that two edges cross oddly in this drawing if and only if their endpoints alternate in the order along the x-axis. By the discussion preceding the theorem, it is sufficient to decide whether there is a set of (e, v)-moves that turns this drawing into a drawing in which every pair of independent edges crosses evenly. We can model this using a system of equations: We introduce variables $x_{e,v}$ for each $e \in E$ and $v \in V$; $x_{e,v} = 1$ means an (e, v)-move is made, $x_{e,v} = 0$ means it is not. For two edges $e = (e_1, e_2)$ and $f = (f_1, f_2)$ to intersect, their intervals on the x-axis have to overlap. And there are two cases: the endpoints alternate (and the edges cross oddly in the initial drawing) or they do not (and the edges cross evenly). Let us first consider the case $e_1 < f_1 < e_2 < f_2$. In the initial drawing, e and f cross oddly, so we must have $x_{e,f_1} = 1 - x_{f,e_2}$ for e and f to cross evenly. If $e_1 < f_1 < f_2 < e_2$, then e and f cross evenly, and we must have $x_{e,f_1} = x_{e,f_2}$ for e and f to cross evenly. Note that these equalities are the only conditions that affect whether e and f cross evenly. Hence, it is sufficient to set up this system of equations for all such pairs of edges e and f and solve it. This can be done using a simple depth-first search: build a graph F on vertex set $E \times V$. Consider every pair of independent edges $e = (e_1, e_2)$ and $f = (f_1, f_2)$ in G, If $e_1 < f_1 < e_2 < f_2$, then add a red edge $((e, f_1), (f, e_2))$ to F. If $e_1 < f_1 < f_2 < e_2$, add a green edge $((e, f_1), (e, f_2))$ to F. Now perform a depth-first traversal of (the not necessarily connected) graph F. When starting the traversal at a new root arbitrarily assign a value of 0 to the root variable. When following a green edge, assign the parent value to the child vertex, when following a red edge, swap 0 to 1 and vice versa. Whenever encountering a back-edge verify that the value assignment to the endpoints of the edge is consistent with its color (green for equal, red for different). If this test fails, the graph cannot be embedded. Otherwise, the depth-first search succeeds and the graph has an x-monotone embedding.

Since we can assume that G is planar, we know that $|E| \leq 3|V|$, and our algorithm runs in time $O(|V|^2)$ with a small constant factor. □

Since the proof of Theorem 5.11 contains an algorithm to test whether a graph has an x-monotone drawing in which every pair of independent edges crosses evenly, and, moreover, produces such a drawing, Lemma 5.10 completes the proof of the following result.

Corollary 5.12. *We can test whether a graph G has an x-monotone embedding and find such an embedding (if it exists) in time $O(|V|^2)$.*

5.6 Level-Planarity Testing and Layout

The definition of an x-monotone drawing does not allow two vertices to have the same x-coordinate. If we remove this restriction we enter the realm of leveled graphs: a *leveled graph* is a graph $G = (V, E)$ together with a level function $\ell : V \to \mathbb{Z}$. A *leveled drawing* of (G, ℓ) is a drawing in which edges are x-monotone and $x(v) = \ell(v)$ for every $v \in V$. (G, ℓ) is *level-planar* if it has a leveled embedding. Some papers have considered *proper* levelings, in which each edge's endpoints are on consecutive levels; we typically do not require our leveling to be proper.

Although testing x-monotonicity may be of little practical interest, of course, however, the work we did can easily be extended to cover level-planarity testing, an important case of layered graph drawing [16, 31, 32, 42, 37].

Level-planar graphs can be recognized in linear time using PQ-trees [42, 37]; this work is based on earlier work for the special case of hierarchies [16]. There had been an earlier attempt at extending this to general graphs [31, 32], but there were gaps in the algorithm as pointed out in [38]. So there may be some value in finding simpler algorithms for the problem. Alternative routes have included identifying Kuratowski-style obstruction sets for level-planarity [30], characterizations via vertex-exchange graphs [29, 27] and reductions to 2-satisfiability [64]. It appears that all of these approaches have subtle problems: currently known obstruction sets for the general case are not complete and are known to be infinite (for standard notions of obstruction containment); only special cases, like trees, are understood [20]. The testing [29] and layout [27] algorithms based on vertex-exchange graphs rest on a characterization of level-planarity that is not fully established at this point, the case when the vertex exchange graph is disconnected remains open [28]; this is unfortunate, since both algorithms are relatively fast, $O(|V|^2)$ for both testing and layout, and very simple (and somewhat similar to ours, even if the characterization they are based on is different). Finally, there also seems a gap in the suggested reduction to 2-satisfiability (which, if correct, would also result in a quadratic time algorithm).

Thus, although the algorithm we are about to describe may not be the first simple, quadratic-time algorithm for level-planarity testing, it appears to be the first with a complete correctness proof.

If G has an x-monotone embedding, then (G, ℓ) is level-planar with $\ell(v) = |\{u : x(u) \leq x(v)\}|$.

Our interest in this section is the reverse direction; how can we reduce testing level-planarity to testing x-monotonicity? The answer is a simple construction: Take a leveled drawing of (G, ℓ). Perturb all vertices slightly, so 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 extends slightly beyond all the perturbed vertices from the same level. If there is a vertex with both left and right rotation empty, remove it.

Suppose that the resulting graph G' has an x-monotone embedding with the same vertex locations. By the construction of G', every vertex v that used to have level $\ell(v) = x^*$ is now incident to an edge that passes over the line $x = x^*$. Since all these curves may not intersect each other, we can perturb the drawing slightly (while keeping it x-monotone) to move every vertex of G back to its original level. Also, if (G, ℓ) is level-planar, then G' is obviously x-monotone, so we can use the algorithm from Theorem 5.11 on G' to test level-planarity of (G, ℓ). Since we only added at most $|V(G)|$ vertices and edges to G, the resulting algorithm still runs in quadratic time—with a small constant factor.

Corollary 5.13. *Given a leveled graph (G, ℓ) we can test in time $O(|V|^2)$ whether G is level-planar.*

Note that this result does not require the leveling of G to be proper and thus improves on the algorithm by Healy and Kuusik [29] (assuming it is correct) which requires the leveling to be proper. Turning an improper leveling into a proper leveling (by subdividing edges) can increase the number of vertices by a quadratic factor.

The reduction from level-planar graphs to x-monotone graphs also allows us to find an embedding of the level-planar graph.

Corollary 5.14. *Given a level-planar graph G we can find a level-planar embedding in time $O(|V|^2)$.*

This improves on the algorithm of Harrigan and Healy [27] (if it is correct), since that algorithm runs in quadratic time for proper levelings.

There is one final conclusion we want to draw from the reduction of level planarity to x-monotonicity: when defining a level planar drawing we required edges to be x-monotone (in the literature one also finds the equivalent requirement that edges are straight-line segments between levels). Similarly to Corollary 5.6 it is now easy to see that the x-monotonicity requirement is stronger than necessary, it suffices to require edges in a level planar drawing to be x-bounded.

Corollary 5.15. *If (G, ℓ) can be embedded so that $x(v) = \ell(v)$ for every $v \in V$ and every edge is x-bounded, then G is level planar.*

Proof. Fix an embedding of (G, ℓ) so that $x(v) = \ell(v)$ for every $v \in V$ and every edge is x-bounded. Consider the ordered graph G' constructed before Corollary 5.13. Then G' has an ordered embedding in which every edge is x-bounded. By Corollary 5.6, G' has an x-monotone embedding in which each vertex keeps its x-coordinate (and the rotation system remains unchanged). As above, from this embedding we can obtain a level-planar embedding of G which completes the proof. □

5.7 Monotone Crossing Numbers

Our Hanani-Tutte results can be recast as results about monotone crossing numbers of leveled graphs. For a leveled graph (G, ℓ) let mon-cr(G, ℓ) be the smallest number of crossings in any leveled drawing of (G, ℓ). Similarly, we can define mon-ocr(G, ℓ) as the smallest number of pairs of edges that cross oddly in any leveled drawing of (G, ℓ). Finally, mon-iocr(G, ℓ) is the smallest number of pairs of non-adjacent edges that cross oddly in any leveled drawing of (G, ℓ). We suppress ℓ and simply write mon-cr(G, ℓ), mon-ocr(G, ℓ), and mon-iocr(G, ℓ). With this notation we can restate the original result by Pach and Tóth, our Theorem 5.1 as saying that mon-ocr$(G, \ell) = 0$ implies mon-cr$(G, \ell) = 0$. Similarly, our Theorem 5.7 can be restated as mon-iocr$(G, \ell) = 0$ implies mon-cr$(G, \ell) = 0$.

From this point of view we can now ask questions that parallel analogous problems for the regular (non-monotone) crossing number variants: cr, ocr, and iocr. For example, we know that $\mathrm{ocr}(G) = \mathrm{cr}(G)$ for $\mathrm{ocr}(G) \leq 3$ [56] and $\mathrm{iocr}(G) = \mathrm{cr}(G)$ for $\mathrm{iocr}(G) \leq 2$ [58]. Pach and Tóth showed that $\mathrm{cr}(G) \leq \binom{2\,\mathrm{ocr}(G)}{2}$ [49, 54]. The core step in this result is a "removing even crossings" lemma, in this particular case: if G is drawn in the plane and E_0 is the set of its even edges, then G can be redrawn so that all edges in E_0 are free of crossings. It immediately implies $\mathrm{cr}(G) \leq \binom{2\,\mathrm{ocr}(G)}{2}$, since only non-even edges can be involved in crossings (and every pair of non-even edges needs to cross at most once). A similar result for monotone drawings fails dramatically:

Theorem 5.16. *For every n there is a graph G so that* mon-cr$(G, \ell) \geq n$ *and* mon-ocr$(G, \ell) = 1$.

In other words: even if there are only two edges crossing oddly and all other edges are even, then any x-monotone drawing of G with the given leveling may require an arbitrary number of crossings. Thus we cannot hope to establish a "removing even crossings" lemma in the context of x-monotone drawings since it would imply a bound on mon-cr(G, ℓ) in terms of mon-ocr(G, ℓ).

In the analysis below we use repeatedly the following simple observation.

Observation 5.1. *A pair of x-monotone edges uv and wz such that $x(u) < x(w) < x(v) < x(z)$ (resp. $x(u) < x(w) < x(z) < x(v)$) is odd if and only if either uv is passing below w and wz is passing below v, or uv is passing above w and wz is passing above v (resp. uv is passing below w and above z or vice versa).*

The example we use only uses 8 vertices, allowing multiple edges which we bundle into a single weighted edge. Consider the graph on 8 vertices with edges 36 and 57 of weight 1 and edges 12, 13, 25, 26, 37, 46, 47, 68, and 78 of weight $n > 1$. Weighted edges can be replaced by paths of length 2 turning the example into a simple graph.

We next argue that mon-cr$(G, \ell) \geq n$. Suppose there is a drawing D with mon-cr$(D) < n$. Then the only pair of edges that may intersect is 36 and 57. Without loss of generality, we can assume that 12, 13 and 78 are drawn exactly as they are in Figure 27. We distinguish two cases depending on whether 46 passes below 5 (as in Figure 27) or above 5. Let us first consider

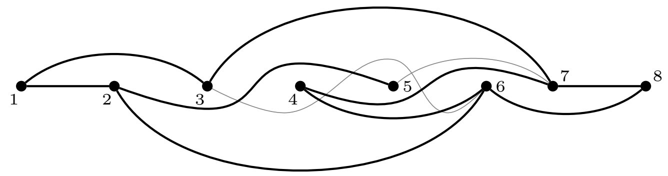

Figure 27: The drawing showing mon-ocr$(G, \ell) \leq 1$. The thick edges have weight n, the thin, gray edges have weight 1.

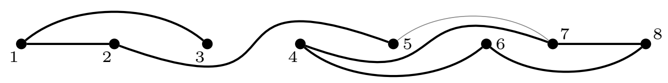

Figure 28: The unique way of drawing edges 25, 57, 68, and 27, assuming 46 passes below 5 and mon-cr$(D) < n$.

the case that 46 passes below 5. Adding edges 25, 57, we see that they are forced to be drawn as in Figure 27. At this point, edge 68 has to pass below 7 and then 47 is forced. That is, if we assume that 46 passes below 5, then the edges we added have to be drawn as shown in Figure 28. By inspection it is clear that adding edge 36 to this drawing will cause at least n crossings, either with edge 25 or edge 47.

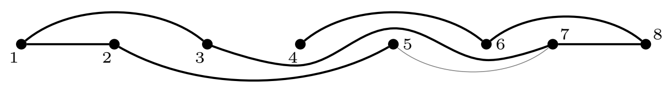

Figure 29: The unique way of drawing edges 25, 57, 68 and 37, assuming 46 passes above 5 and mon-cr$(D) < n$.

On the other hand, if 46 passes above 5, then edge 25 is forced to pass below 3 and 4 and edge 57 is forced below 6. This forces 68 above 7 which in turn forces 37 below 4 and 6 and above 5. However, now it is impossible to add edge 26 without having it cross either 13 or 37, see Figure 29.

6 Separating crossing numbers

In this section we show that certain definitions of crossing numbers do not happen to coincide. A considerable part of the material from the present section was published in [23].

6.1 Introduction

When drawing a graph, some assumptions are natural: there are only finitely many crossings, no more than two edges cross in a point, edges do not pass through vertices, and edges do not touch.* Sometimes these assumptions are relaxed (*degenerate drawings* allow more than two edges to cross in a point), and sometimes more restrictions are added, for example adjacent edges may not be allowed to cross.

The *crossing number* $\mathrm{cr}(G)$ of a graph G is the smallest number of crossings in a drawing of G. It is easy to see that in an optimal drawing, adjacent edges of G do not cross (such crossings can always be removed). This may have led researchers on crossing numbers to think that adjacent crossings are irrelevant or even to prohibit them in drawings.† Another source for ignoring adjacent crossings may be the fact that graph drawings are often straight-line drawings in which adjacent edges naturally cannot cross.

Pach and Tóth point out in "Which Crossing Number is It Anyway?" that there have been many different ideas on how to define a notion of crossing number, including the following (see [49, 70]):

pair crossing number: $\mathrm{pcr}(G)$, the smallest number of pairs of edges crossing in a drawing of G,

odd crossing number: $\mathrm{ocr}(G)$, the smallest number of pairs of edges crossing oddly (*odd pairs*) in a drawing of G.

Tutte introduced another type of crossing number by orienting edges arbitrarily, then letting $\lambda(e, f)$ be the difference in the number of crossings where e is pointed to the left of f and the number of crossings where e is pointed to the right of f. Changing the orientation of e or f will only change the sign of $\lambda(e, f)$, so one can define:

algebraic crossing number: $\mathrm{acr}(G)$, the minimum of $\sum |\lambda(e, f)|$ in a drawing of G, where the sum is taken over pairs of edges e, f.

*For a detailed discussion see [70].
†Székely discusses this issue in [70].

By definition we have $\mathrm{ocr}(G) \leq \mathrm{pcr}(G) \leq \mathrm{cr}(G)$ and $\mathrm{ocr}(G) \leq \mathrm{acr}(G) \leq \mathrm{cr}(G)$.

For each of these notions, one can ask whether adjacent crossings matter. In [48], Pach and Tóth suggest a systematic study of this issue (see also [8, Section 9.4]) by introducing two rules: "Rule +" restricts the drawings to drawings in which adjacent edges are not allowed to cross. "Rule −" allows crossings of adjacent edges, but does not count them towards the crossing number. Each parameter ocr, pcr, acr, and cr can be modified by either rule, but since $\mathrm{cr}_+ = \mathrm{cr}$ (implied by the discussion at the beginning of the section), this yields up to eleven possible distinct variants.

The tables below are based on a figure from [8]. The notion of ocr_- was introduced as the *independent odd crossing number*, iocr, by Székeley [70].[*]

Rule +	ocr_+	pcr_+	cr
	ocr	pcr	
Rule −	$\mathrm{iocr} = \mathrm{ocr}_-$	pcr_-	cr_-

	ocr_+	acr_+	cr
	ocr	acr	
	$\mathrm{iocr} = \mathrm{ocr}_-$	acr_-	cr_-

It immediately follows from the definitions that the values in each table increase monotonically as one moves from the left to the right and from the bottom to the top. Not much more is known about the relationships between these crossing number variants. In [48], Pach and Tóth write, "We cannot prove anything else about $\mathrm{iocr}(G)$, $\mathrm{pcr}_-(G)$, and $\mathrm{cr}_-(G)$. We conjecture that these values are very close to $\mathrm{cr}(G)$, if not the same. That is, we believe that by letting pairs of incident edges cross an arbitrary number of times, we cannot effectively reduce the total number of crossings between independent pairs of edges."[†] Tutte [73] seems to have had a similar opinion, when he explained his choice to study acr_-, writing, "We are taking the view that crossings of adjacent edges are trivial, and easily got rid of." Székely [70] later commented "We interpret this sentence as a philosophical view and not a mathematical claim." West [75] and Székely [69] mention the specific question of whether there are graphs with $\mathrm{iocr}(G) < \mathrm{ocr}(G)$.

There are situations when the entire system of crossing numbers collapses. The classic Hanani-Tutte theorem states that if a graph can be drawn in the plane so that no pair of independent edges crosses an odd number of times, then it is planar [26, 73]. In other words, $\mathrm{iocr}(G) = 0$ implies that

[*]Székeley credits Tutte [73] with the (implicit) definition of iocr, but Tutte is really concerned with the algebraic crossing numbers only, acr and acr_-; he does not consider parity.

[†]Some authors write *incident edges* to mean two edges that share an endpoint, but we will only use *adjacent edges*. Non-adjacent edges are also called *independent edges*.

$\mathrm{cr}(G) = 0$ and, thus, that all of the eleven variants are equal (to zero). This was extended to show that all eleven variants are equal as long as $\mathrm{iocr}(G) \leq 2$ [58]. Székely gave an explicit criterion for when all variants are equal [71]. It is also known that all eleven variants are within a square of each other, since $\mathrm{cr}(G) \leq \binom{2\,\mathrm{iocr}(G)}{2}$ [58]. For drawings of G on the projective plane N_1, we know that $\mathrm{iocr}_{N_1}(G) = 0$ implies that $\mathrm{cr}_{N_1}(G) = 0$, so again all variants are equal (to zero) in this case [52].

Setting aside the Rule $-$ variants, there are some strong results for the remaining seven variants, ocr, ocr_+, acr, acr_+, pcr, pcr_+ and cr. If $\mathrm{ocr}(G) \leq 3$ then all these seven variants are equal [55]. For drawings on any surface S, if $\mathrm{ocr}_S(G) = 0$ then all seven variants are equal (to zero) [57]. Valtr [74] showed that $\mathrm{cr}(G) = O(\mathrm{pcr}^2(G)/\log \mathrm{pcr}(G))$, which Tóth [72] improved to $\mathrm{cr}(G) = O(\mathrm{pcr}^2(G)/\log^2 \mathrm{pcr}(G))$.

On the other hand, we know that ocr and pcr differ: there is an infinite family of graphs with $\mathrm{ocr}(G) < 0.867 \cdot \mathrm{pcr}(G)$ [56]. Tóth improved this by giving a family of graphs with $\mathrm{acr}(G) < 0.855 \cdot \mathrm{pcr}(G)$ [72] (so $\mathrm{ocr}(G) < 0.855 \cdot \mathrm{pcr}(G)$ as well). For such G it immediately follows that $\mathrm{ocr}(G) < \mathrm{cr}(G)$ and $\mathrm{acr}_-(G) < \mathrm{cr}(G)$, answering questions of Pach and Tóth [49] and Tutte [73]; additional consequences can be deduced from the tables above. However, none of these results address the intuitions expressed by Tutte and by Pach and Tóth about how Rule $-$ may or may not affect cr, pcr, ocr, or acr.

We can finally give a result of this nature.

Theorem 6.1. *For every n, there exists a (vertex) two-connected a graph G with* $\mathrm{iocr}(G) < \mathrm{ocr}(G) - n$.*

In short, adjacent crossings matter.†

To prove Theorem 6.1, we will first prove a separation for monotone drawings of ordered graphs. An *ordered graph* is a graph with a total ordering of its vertices. For our purposes, we will assume that the vertex set of an ordered graph is a subset of the integers, and we will only consider drawings where each vertex n has x-coordinate equal to n. As in the previous section, a drawing of a graph is *x-monotone* if every edge intersects every vertical line at most once and every vertical line contains at most one vertex. We can generalize each crossing number variant to x-monotone drawings of ordered graphs G, which we denote $\text{mon-cr}(G)$, $\text{mon-ocr}(G)$, $\text{mon-iocr}(G)$, etc.

Pach and Tóth proved that $\text{mon-ocr}(G) = 0$ implies $\text{mon-cr}(G) = 0$ [50]. We strengthened this by showing that $\text{mon-iocr}(G) = 0$ implies $\text{mon-cr}(G) =$

*If we did not require G to be two-connected it would be enough to prove that $\mathrm{iocr}(G) < \mathrm{ocr}(G) - 1$, since we can take n disjoint copies of G.

†Among other things, Theorem 6.1 justifies the rather baroque **NP**-completeness proof for iocr in [59]. **NP**-completeness of ocr is simpler in comparison [50].

0 [24], which had been left as an open problem in [50]. On the other hand, in the same paper we showed that for every n there is a graph G such that $\text{mon-cr}(G) \geq n$ and $\text{mon-ocr}(G) = 1$. In this paper, we will show that there can also be an arbitrary gap between mon-ocr and mon-iocr.

Theorem 6.2. *For every $n > 2$ there is an ordered graph G with* $\text{mon-iocr}(G) = 2 < n = \text{mon-ocr}(G)$.

We will use Theorem 6.2 to prove Theorem 6.1. We note that by using our technique we can separate other crossing number variants excluding acr_+ and ocr_+ given that we can separate the corresponding monotone variants.

6.2 Separating Monotone Crossing Numbers

We generalize the crossing number definitions for graphs with weighted edges. Suppose that G is a graph and each edge e has weight $w(e)$. A crossing between edges e and f is assigned *crossing weight* equal to the product $w(e)w(f)$. Let D be an arbitrary drawing of G, and define

$\text{cr}(D) =$ the sum of crossing weights, taken over all crossings in D,

$\text{cr}_-(D) =$ the sum of crossing weights, taken over all crossings between non-adjacent edges in D,

$\text{pcr}(D) =$ the sum of $w(e)w(f)$, taken over all crossing pairs e, f in D,

$\text{pcr}_-(D) =$ the sum of $w(e)w(f)$, taken over all independent crossing pairs e, f in D,

$\text{pcr}_+(D) =$ the sum of $w(e)w(f)$, taken over all crossing pairs e, f in D free of crossings created by adjacent edges,

$\text{ocr}(D) =$ the sum of $w(e)w(f)$, taken over all odd crossing pairs e, f in D,

$\text{iocr}(D) =$ the sum of $w(e)w(f)$, taken over all independent odd crossing pairs e, f in D.

Analogously we can also define $\text{acr}(D)$ and $\text{acr}_-(D)$. Let $\text{cr}(G) = \min_D \text{cr}(D)$, $\text{ocr}(G) = \min_D \text{ocr}(D)$, and $\text{iocr}(G) = \min_D \text{iocr}(D) \ldots$ etc., with each minimum taken over all drawings D of G. If we assign every edge weight equal to 1, then these definitions revert back to their original, unweighted versions.

Consider an ordered graph $G = (\{1, \ldots, 16\}, E)$, $E = \{(1,3), (2,4), (2,15), (3,4), (3,5), (4,5), (4,6), (4,7), (5,8), (6,9), (7,10), (8,11), (9,12), (10,13),$

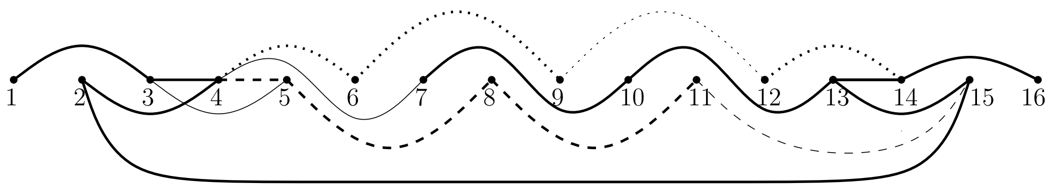

(a)

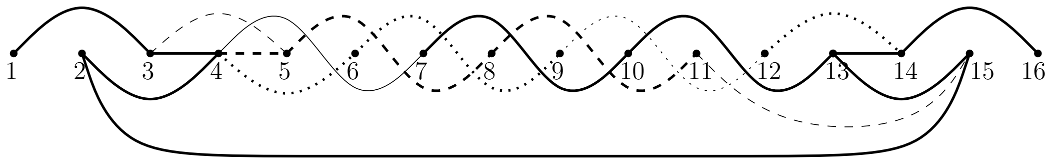

(b)

Figure 30: (a) Drawing of the ordered graph G witnessing that its monotone odd crossing number is at most x; (b) Drawing of the ordered graph G witnessing that its independent monotone odd crossing number is at most 2.

$(12, 14), (13, 14), (13, 15), (14, 16)\}$ (see Figure 30(a)), with its bold edges of weight x and the thin edges of weight 1.

Theorem 6.3. *For the weighted ordered graph G with $x \geq 3$, we have*

$$\text{mon-iocr}(G) = 2 < x = \text{mon-ocr}(G). \tag{8}$$

Proof. The drawing of G in Figure 30(b) gives almost the first part of the claim: mon-iocr$(G) \leq 2$. The following analysis is based on Observation 5.1.

Let G_1 denote the ordered subgraph of G depicted in Figure 31(a). We show that there exists only one way (up to flipping the drawing upside down) how the above-below relation of the vertices and edges in a drawing of G_1 with the monotone odd crossing number less than x looks like: Ref. to Figure 31(a). The edge $(2, 15)$ passes either above all the vertices $4, 7, 10$ and 13 or below them. Hence, from now one, without loss of generality, we assume that the edge $(2, 15)$ passes below all the other vertices of the cycle $\{2, 4, 7, 10, 13, 15\}$. Now, the edge $(1, 3)$ passes above 2, as otherwise either the edge $(1, 3)$ or $(3, 4)$ forms an odd pair with $(2, 15)$. Similarly, the edge $(14, 16)$ passes above the vertex 15. At this point we already determined the position of the vertex 3 with respect to the edge $(2, 4)$ and $(2, 15)$ and 14 with respect to $(13, 15)$ and $(2, 15)$, i.e. 3 and 14, resp., is above $(2, 4)$ and $(2, 15)$, and $(13, 15)$ and $(2, 15)$, resp.

If we extend G_1 to the ordered subgraph G_2 of G depicted on Figure 31(b), the above-below relationship of the edges and vertices is again forced in a drawing of G_2 with the monotone odd crossing number less than x. To this end note that the edge $(12, 14)$ must pass above the vertex 13. This forces the vertices 6,9 and 12, resp., above the edges $(4, 7), (7, 10)$ and $(10, 13)$, respectively. Since, the vertex 14 is above the edge $(2, 15)$ the vertices of $G_2 \setminus G_1$ are above it as well.

Finally, by extending G_2 to the subgraph G_3 (see Figure 31(c)) of G we see that in any monotone drawing of G we get an edge of the weight x crossed by another edge an odd number of times: Observe that the edge $(11, 15)$ has to pass below all the intermediate vertices. Hence, the vertex 11 is below the edge $(10, 13)$. On the other hand, the edge $(3, 5)$ is passing above the vertex 4 forcing the vertex 5 above the edge $(4, 6)$. Since the vertex 11 is below the edge $(10, 13)$, the vertex 5 has to be below the edge $(4, 7)$. For if not, either the edge $(5, 8)$ or $(8, 11)$ participates in an odd crossing pair. Similarly, the vertex 11 must be above the edge $(9, 12)$, as the vertex 5 is above the edge $(4, 6)$. However, this means that the edge $(4, 5)$ creates an odd crossing pair.

The analysis also shows that mon-iocr$(G) \geq 2$. Thus, the claim follows.

$\square$

6.3 From Weighted Edges to Unweighted Edges

Suppose that G is a graph or ordered graph with edges of positive integer weight. Let G' be the graph obtained by replacing each edge of weight w with w edges of weight 1, equivalently, with w unweighted edges. Choose any of the eleven crossing variants mentioned in Section 6.1, and consider a drawing of G' (which is x-monotone if G is an ordered graph) that optimizes that crossing variant. Suppose that e_1 and e_2 are copies of the same edge e of G. Without loss of generality, we may assume that e_1 contributes less than or equal to what e_2 contributes to the chosen crossing parameter. We can redraw e_2 along the side of e_1 so that they do not cross; then e_2 will contribute the same to the crossing parameter as e_1, so the new drawing is still optimal. Hence, we may assume that in an optimal drawing of G', multiple edges are drawn in a bundle, all with essentially the same behavior.* It follows that all crossing parameters are the same for G and G'.

*This argument was probably first made in Kainen [39] for the standard crossing number.

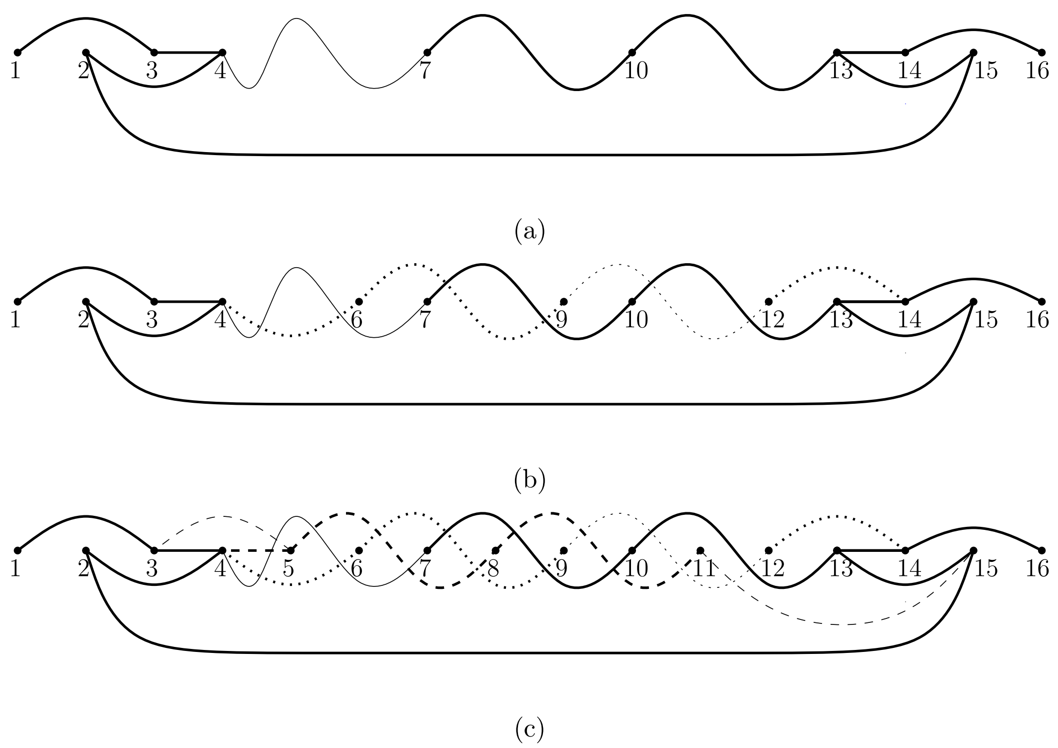

Figure 31: (a) Above-below relationship of the vertices and edges in a drawing of the ordered graph G_i with the odd crossing number less than x: (a) $G_i = G_1$; (b) $G_i = G_2$ and (c) $G_i = G_3$.

Lemma 6.4. *Subdividing an edge of a graph does not change* ocr *or* iocr*. Subdividing an edge of an ordered graph near one of its endpoints does not change* mon-ocr *or* mon-iocr*. These results hold for graphs with multiple edges as well.*

Proof. Let G be a graph or ordered graph, possibly with multiple edges. If G is an ordered graph, we will restrict all drawings to be x-monotone drawings.

Fix an ocr-optimal (iocr-optimal) drawing of G, and choose any edge uv. Subdivide uv with a vertex z, which is added to the drawing of uv near the endpoint u. Then for each edge $e \neq uv$, e will cross zv oddly if and only if e crossed uv oddly, and e does not cross uz at all. Hence ocr is unchanged; iocr is also unchanged unless e shares an endpoint with uv but not with zv, which means that e is incident to u but not v. In this case, we can deform a small section of e until it passes over z (while maintaining its monotonicity, if G is ordered); do this for all such e. This yields a drawing with iocr no bigger than in the initial drawing.

Now consider any drawing of the new graph. We can erase z from that drawing to obtain a drawing of the original, unsubdivided graph. If G is ordered, then we erase z from an x-monotone drawing where z lies strictly between u and v, so we obtain an x-monotone drawing of G. Erasing z moves all odd pairs of edges with uz or zv to become odd pairs with uv (and if an edge crosses both uz and zv oddly, then these cancel and it crosses uv evenly). Hence the number of odd pairs and independent odd pairs does not increase. □

Consider any integer $x \geq 3$. Replace the weighted edges of the graph in Figure 30(a) by multiple edges, and then apply Lemma 6.4 to every edge. We obtain an unweighted ordered graph H with mon-iocr$(H) = 2 < x =$ mon-ocr(H). Thus, Theorem 6.2 is proved.

Before moving on, note that for any drawing of a graph G, we can remove self-intersections of edges without adding any crossing or odd pair, by redrawing locally near the crossing as shown in Figure 32 (originally from [56]).

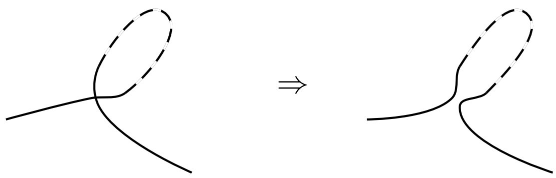

Figure 32: Removing a self-intersection.

6.4 Adjacent Crossings are Not Trivial

Given an ordered graph $G = (V, E)$ with $V = \{v_1 < v_2 \cdots < v_n\}$ let G' be obtained from G by adding the following framework: start with a cycle C_{2n+2} formed from two paths $s, u_1, \ldots, u_n, t$ and $s, w_1, \ldots, w_n, t$; call this the *outer framework*. Add paths $Q_i = u_i v_i w_i$ for $1 \leq i \leq n$; call this the *inner framework*. Assign a weight of $w_I = n^4 + 1$ to the edges in the inner framework and a weight of $w_O = n^4 + n^3 w_I + 1$ to the edges in the outer framework. Edges originally in G remain at weight 1 (unweighted). From the weighted graph G' we will obtain the unweighted graph G'' by replacing each edge of weight $w > 1$ in G' by w copies of P_3.

Lemma 6.5. *With G' as defined above we have $\psi(G'') = \text{mon-}\psi(G) + c$ for any connected graph G, where ψ is one of the crossing numbers* $\{\text{iocr}, \text{ocr}, \text{pcr}, \text{pcr}_-, \text{pcr}_+, \text{cr}_-, \text{cr}\}$ *and $c = w_I \sum_{v_i v_j \in E(G), i<j} (j - i - 1)$.*

Now, we are in the position to prove Theorem 6.1.

Proof of Theorem 6.1: Lemma 6.5 and Theorem 6.2 immediately yield Theorem 6.1. □

Let us show an immediate corollary of the previous lemma about the algebraic crossing number.

Corollary 6.6. *With G' as defined above we have $\psi(G'') = \text{mon-}\psi(G) + c$ for any connected graph G, where ψ is one of the crossing numbers $\{\text{acr}_-, \text{acr}\}$ and $c = w_I \sum_{v_iv_j \in E(G), i<j} (j - i - 1)$.*

Proof. Note that in the monotone setting the algebraic crossing number and the odd crossing number coincide. Thus, we have $\text{acr}(G'') \geq \text{ocr}(G'') = \text{mon-ocr}(G) + c = \text{mon-acr}(G) + c$. The other direction follows by taking a monotone drawing realizing $\text{mon-acr}(G)$ and overlaying it with a planar drawing of the framework (see the first paragraph of the proof of Lemma 6.5 for the details). The same argument works also for acr_-. □

Thus, Lemma 6.5 works for all the introduced crossing number variants except ocr_+ and acr_+. Moreover, we show in the end of this section that it indeed cannot work for these two crossing number variants.

In [24] we showed that for every n there is an ordered graph G such that $\text{mon-cr}(G) \geq n$ and $\text{mon-ocr}(G) = 1$. Together with Lemma 6.5, this yields a new graph G' with $\text{ocr}(G') < \text{cr}(G')$, joining the earlier examples from [56] and [72].Note that our lemma can lead to further separations given that we separate the corresponding monotone variants.

For the proof of Lemma 6.5, we need the following lemma. An *even* edge is an edge that crosses every other edge an even number of times (possibly zero times).

Lemma 6.7 (Pelsmajer, Schaefer, and Štefankovič [55])**.** *If D is a drawing of G in the plane and E is the set of even edges in D, then G has a redrawing in which all edges in E are crossing-free, there are no new pairs of edges that cross an odd number of times, and the cyclic order of edges at each vertex does not change.*

Proof of Lemma 6.5. First note that $\psi(G'') \leq \text{mon-}\psi(G) + c$ is immediate: take a monotone drawing realizing $\text{mon-}\psi(G)$ and overlay it with a planar drawing of the framework, call the resulting drawing D' (see Figure 33 for an example). Then $\psi(D') = \text{mon-}\psi(G) + c$ since the only crossings are single crossings between pairs of non-adjacent edges that count the same whatever ψ is. From D' we can obtain a drawing D'' of G'' by replacing the weighted edges in the drawing by parallel P_3s; then $\psi(D'') = \psi(D')$ (since the framework edges are not involved in any adjacent crossings), so $\psi(G'') \leq \psi(D'') = \text{mon-}\psi(G) + c$.

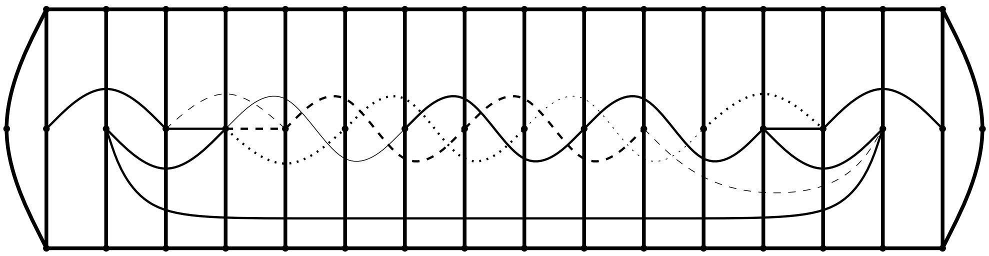

Figure 33: Overlay of G from Figure 30(a) with framework.

It remains to prove $\psi(G'') \geq \text{mon-}\psi(G)+c$ for $\psi \in \{\text{cr}, \text{ocr}, \text{pcr}, \text{pcr}_-, \text{pcr}_+, \text{cr}_-, \text{iocr}\}$. It is easy to see that $\psi(G'') \geq \psi(G')$: fix an ψ-optimal drawing of G''. Consider w parallel paths P_3 that were used to replace an edge of weight w in G'. Pick one of these paths P that contributes the smallest amount to $\psi(G'')$. Now redraw the remaining $w-1$ paths to run very close to P and without crossing each other. This redrawing cannot increase the value of ψ of the drawing. But now we can bundle the parallel paths into a single weighted edge to obtain a drawing D' of G' with $\psi(D') \leq \psi(G'')$. So $\psi(G') \leq \psi(G'')$.

Hence, to establish the lemma it is sufficient to show that $\psi(G') \geq \text{mon-}\psi(G) + c$. We proceed in three steps; we first show that there is a ψ-minimal drawing of G' in which the edges of the outer framework are crossing-free. In the second step we show that we can assume that the edges of the inner framework do not cross each other. In the third step we show that from such a drawing of G', we can construct a monotone drawing of G with at most $\psi(G') - c$ crossings. It follows that $\text{mon-}\psi(G) \leq \psi(G') - c$.

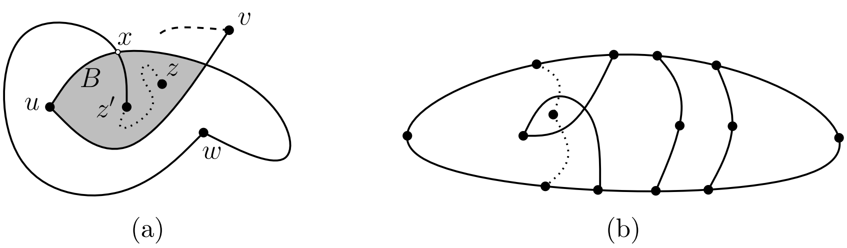

Figure 34: (a) A path avoiding u and w from v to z cannot enter the bigon B; (b) Non-empty bigon created by two adjacent inner framework edges.

For the first step, fix an ψ-minimal drawing of G'. For $\psi \in \{\text{pcr}, \text{pcr}_+, \text{cr}\}$ the claim is immediate: any edge crossing an edge of the outer framework contributes at least w_O to $\psi(G')$. However, we already proved that $\psi(G') \leq \text{mon-}\psi(G) + c \leq n^4 + n^3 w_I < w_O$, so all edges of the outer framework must be crossing-free. If $\psi = \text{ocr}$ then edges of the outer framework cannot be involved in any odd pairs, since any such odd pair would contribute w_O to ocr and, as above, $\psi(G') \leq \text{mon-}\psi(G) + c \leq n^4 + n^3 w_I < w_O$. So all the edges in the outer framework are even. We can then apply Lemma 6.7 to make all edges in the outer framework crossing-free without introducing any new pair of edges crossing oddly (in particular, ψ does not increase). For the case $\psi = \text{iocr}$ the argument is similar to ocr. In any iocr-minimal drawing, edges of the outer framework cannot be involved in any independent odd pairs, so all odd pairs involving these edges must have adjacent edges. However, all vertices in the outer framework have degree 2 or 3, so we can modify the drawing near each of these vertices to ensure that all the edges in the outer framework are actually even. We then proceed as in the case of ocr. This leaves $\psi \in \{\text{pcr}_-, \text{cr}_-\}$. In both cases we have to rule out the possibility that an outer framework edge is crossed by an adjacent edge in every optimal drawing, i.e. a drawing witnessing $\text{pcr}_-(G')$ or $\text{cr}_-(G')$. Suppose that an outer framework edge uv is crossed by an inner or outer framework edge uw. Refer to Figure 34(a). Let z denote a vertex inside a bigon B not containing v and w created by uv and uw. If z does not exist, we can reduce the number of crossings between uv and uw. The first edge of a path from w to z using only framework edges and avoiding u and v has to have the second endpoint, let's call it z', inside B as no other subsequent edge on this path is incident to w,u or v. Note that the edge wz' cannot cross uv. Let x denote the first crossing point on this path intersecting wu. Then the first edge of a path from v to z using only framework edges and avoiding u and w cannot cross the parts of wu and wz' between w and x. However, no other subsequent edge on this path is incident to v,u or w, which means that the path from v to z cannot enter B (unless we have an independent crossing pair consisting of an outer framework edge) containing z (contradiction).

This completes the first step: we know that we can assume that the outer framework is entirely free of crossings. Since we assumed that G is a connected graph, all vertices of G must lie in the same face of C_{2n+2}, without loss of generality, the inner face. Since every edge not in the outer framework is incident to a vertex of G this also implies that all edges lie in the inner face and the outer face is therefore empty.

In the second step we show that we can assume that edges of the inner framework do not cross each other. Recall that $Q_i = u_i v_i w_i$ is the inner framework path passing through v_i with endpoints u_i and w_i on C_{2n+2}, for

$1 \leq i \leq n$.

For $\psi \in \{\mathrm{cr}, \mathrm{pcr}, \mathrm{pcr}_+\}$ the claim is immediate again, since any such crossing would contribute $w_I^2 = w_I(n^4+1) = n^4 w_I + w_I > n^3 w_I + n^4 + 1 = w_O$ to $\psi(G')$, but we already know that $\psi(G') \leq w_O$. If $\psi \in \{\mathrm{cr}_-, \mathrm{pcr}_-\}$, it is enough to observe that any bigon created by a pair of adjacent crossing inner framework edges has to be empty. For if not, a vertex inside such a bigon is connected by an inner framework edge with the outerface yielding two crossing non-adjacent inner framework edges (see Fig. 34(b)) contributing more than w_O to $\psi(G')$.

For $\psi = \mathrm{ocr}$, we can similarly conclude that any two edges of the inner framework cross evenly, and for $\psi = \mathrm{iocr}$, we know that any independent pair of edges in the inner framework crosses evenly. Suppose that $\psi = \mathrm{iocr}$ and two adjacent edges of the inner framework, u_iv_i and v_iw_i, cross oddly. In that case, we perform a (u_iv_i, v_i)-*move* (that is, we deform a small section of u_iv_i, bring it close to v_i and then make it pass over v_i); this does not affect iocr and ensures that u_iv_i and v_iw_i cross evenly. We conclude that for $\psi \in \{\mathrm{ocr}, \mathrm{iocr}\}$ any two edges of the inner framework cross an even number of times. We next show how to remove crossings between edges of the inner framework.

To this end, let us consider $Q_1 = u_1v_1w_1$. Let e be an edge of the inner framework that crosses u_1v_1 (we allow the case $e = v_1w_1$). Deform e near each such crossing so that it follows along u_1v_1 toward v_1 and then over v_1. Since e must have crossed u_1v_1 an even number of times, this procedure will not change the value of ψ for the drawing. Performing this for all such edges e of the inner framework leaves u_1v_1 free of crossings with edges of the inner framework. This redrawing process may have introduced self-crossings of v_1w_1 which can be removed without affecting ψ, as described at the end of Section 6.2. So u_1v_1 crosses no edge of the inner framework and v_1w_1 crosses every other edge of the inner framework evenly. Without loss of generality, we can assume that t is in the exterior of $su_1v_1w_1s$. Then the interior of $su_1v_1w_1s$ does not contain any vertices: every vertex (other than t) has a path consisting of edges of weight at least w_I to t, contributing at least w_I^2 to ψ, which we know to be impossible. Now cut each edge e of the inner framework where it crosses v_1w_1. We can partition the crossings of e and v_1w_1 into pairs since they cross evenly, and then for each pair we add curves that run along each side of v_1w_1 that connect the severed ends of e. Thus, e is replaced by a curve that may have more than one component, all but one of which are closed curves with no vertex, and none of the components intersect v_1w_1. Because of the way the connecting curves are added in pairs, the value of ψ is unchanged. The components lying within $su_1v_1w_1s$ are all closed curves without vertices. Moreover, since there is no vertex within that region, they can be deleted without affecting ψ. Any two of the curves on the other side of Q_1 can be merged by erasing a tiny bit of each curve and adding two parallel curves within the region that join the erased bits of opposite curves, giving a wide berth to all vertices, which ensures that ψ is unchanged. Repeating this process merges all curve components in that region into a single curve, and after removing self-intersections we obtain a valid drawing of e within that region. We can now repeat this argument with Q_2 and $su_1u_2v_2w_2w_1s$, and so on, to establish that none of the Q_i, $1 \leq i \leq n$ have crossings with any edges of the inner framework. This completes the second step.

Hence, for the third step, we can assume that every crossing is between two edges of G or between an edge of G and an edge of the inner framework.

At this point, let us deform the whole drawing so that $C_{2n+2} \cup \{Q_1, Q_n\} - \{s, t\}$ is a rectangle and all the Q_i are parallel straight-line segments orthogonal to the outer framework.

For $\psi = \{\text{cr}, \text{cr}_-, \text{pcr}_-\}$ we are nearly done: a G-edge e connecting v_i to v_j must cross all Q_k with $i < k < j$, forcing at least c crossings. This leaves $\psi(G') - c \leq \text{mon-}\psi(G) \leq n^4 < w_I$ crossings counting towards $\psi(G')$. Since a crossing with an edge of the inner framework contributes at least w_I to $\psi(G')$ this accounts for all crossings with edges of the inner framework, except possibly with edges of Q_i and Q_j in case of $\{\text{cr}_-, \text{pcr}_-\}$. However, similarly as before a non-empty bigon created by Q_i (resp. Q_j) and v_iv_j forces an additional non-adjacent crossing pair involving inner framework edge. So, an edge $e = v_iv_j$ crosses all Q_k with $i < k < j$ and no other Q_ks. The actual behavior of e between two neighboring Q_ks is irrelevant and within each such region we can replace e by a straight-line segment connecting its crossings between neighboring Q_ks. This does not affect ψ and results in a monotone drawing of G with $\psi(G') - c$ crossings, proving that $\text{mon-}\psi(G) \leq \psi(G') - c$ which is what we had to prove.

For $\psi \in \{\text{ocr}, \text{iocr}\}$ we need to do a bit more work. A G-edge e connecting v_i to v_j must cross all Q_k with $i < k < j$ oddly. So the crossings of G-edges with the inner framework contribute at least c to the value of ψ. This leaves at most $\psi(G') - c \leq \text{mon-}\psi(G) < w_I$ in $\psi(G')$ unaccounted for. So there are no non-adjacent odd pairs with edges of the inner framework except those absolutely necessary to connect the endpoints of every edge in G. The only case in which odd pairs with inner framework edges can still occur is in the iocr case (where such crossings do not count) if an edge v_iv_j, $i < j$ crosses an adjacent inner framework edge (u_iv_i, v_iw_i, u_jv_j, or v_jw_j) oddly. In this case we redraw v_iv_j near each endpoint (if necessary) so that the ends of v_iv_j at v_i and v_j lie between Q_i and Q_j; this does not affect iocr and results in v_iv_j crossing both Q_i and Q_j an even number of times. It is possible at this

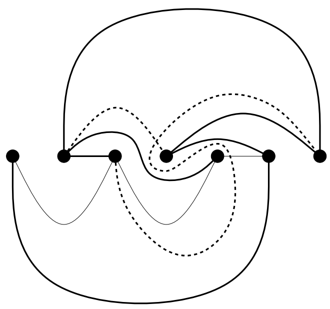

Figure 35: Ordered multigraph and its optimal drawing with respect to ocr_+ and acr_+ when overlayed with our framework; the thin edges have the weight 1; the thick edges x^2; and the dashed edges x.

point that v_iv_j crosses both u_kv_k and v_kw_k oddly, where $k \in \{i, j\}$. In that case we perform a (v_iv_j, v_k)-move; this does not affect iocr and ensures that v_iv_j crosses both u_kv_k and v_kw_k evenly.

Thus for $\psi \in \{\text{ocr}, \text{iocr}\}$ we can now assume that if an edge $e = v_iv_j$ crosses u_kv_k or v_kw_k with $k \leq i$ or $k \geq j$ it must do so evenly. As we did above for the inner framework edges, we push all crossings of e with u_kv_k along u_kv_k and over v_k to v_kw_k so that u_kv_k does not cross e at all; pushing e off u_kv_k does not affect ψ, since e crossed u_kv_k evenly. For all $k \leq i$ and $k \geq j$ cut e at v_kw_k; pair up crossings of e with v_kw_k and reconnect severed ends of e on both side of v_kw_k for all $k \leq i$, $k \geq j$. Closed components of e between Q_i and Q_j can be reconnected to the arc-component of e without affecting ψ. Every other closed component of e is entirely contained in a region which does not contain a vertex, so all such components are even and can be dropped without affecting ψ. In the end, all of e lies in the region formed by C_{2n+2} and Q_i and Q_j.

Now for any $i < k < j$ we have either $\text{ocr}(e, u_kv_k) = 0$ and $\text{ocr}(e, v_kw_k) = w_I$ or $\text{ocr}(e, u_kv_k) = w_I$ and $\text{ocr}(e, v_kw_k) = 0$ (since we have already accounted for all crossings with edges of weight at least w_I). For every k push all crossings of e with Q_k from the edge with $\text{ocr} = 0$ to the other edge (not affecting the value of ψ); that is, e avoids one of the edges of Q_k for every $i < k < j$. Let e' be any other curve in the region in C_{2n+2} bounded by Q_i, Q_j that shares ends with e (here, an end is an endpoint together with a small, crossing-free part of the edge incident to the endpoint); furthermore, suppose that e' avoids the same edge in each Q_k as does e. Then $\text{ocr}(e, g) = \text{ocr}(e', g)$ for every edge g (other than e), since e can be continuously deformed to e' without passing over any vertex. In particular, we can replace e with a monotone polygonal arc without changing the value of ψ. Repeating this for all edges of G gives us a monotone drawing of G with mon-ψ crossings. This completes the argument for $\psi \in \{\mathrm{ocr}, \mathrm{iocr}\}$. $\square$

We illustrate the impossibility of an extension of Lemma 6.5 to ocr_+ and acr_+ without additional ideas by Figure 35. It is likely, in case the separation is possible, that a framework construction similar to one used in Lemma 6.5 works also for ocr_+ and acr_+, but we were unable to find it.

Part III

Questions, Bibliography and CV

7 Summary of Interesting Open Questions

We conclude the thesis with a couple of open problems identifying the points where we stopped in our investigation.

In Chapter 3, we provided for any $\epsilon > 0$ an algorithm with a finite running time that either improves the upper bound on the maximum number of edges in a thrackle to $(1 + \epsilon)n$ or disproves the Conway's conjecture. The next natural step for an improvement of our upper bound would be to test dumbbells involving cycles of length eight (or higher) for thrackleability. Since there are many ways how to draw the cycle of length eight as a thrackle, not to mention bigger cycles, the running time of our testing algorithm was too long (using the computational power at our disposal) for testing such dumbbells. Hence, it would be useful to prove some additional properties of hypothetical thrackle drawings of these small dumbbells, which could reduce the running time of our algorithm.

Question 7.1. *Are the dumbbells* $\mathrm{DB}(8,6,l)$*,* $\mathrm{DB}(8,6,l)$ *and* $\mathrm{DB}(8,8,l)$ *for* $l = -4 \ldots 4$ *thrackleable?*

For the poly-line drawings of graphs with at most two bends per edge we could not resolve the following question.

Question 7.2. *Let* $A \subset (0, \frac{\pi}{2}]$ *be a set of* k *angles,* $k \in \mathbb{N}$*, and let* G *be a graph on* n *vertices that admits a polyline drawing with at most* 2 *bends per edge such that every crossing occurs at some angle from* A*. Is the upper bound of* $O(k^2 n)$ *on the maximum number of edges in* G *tight?*

How hard is it to determine whether a graph admits a polyline drawing with few bends per edge and few crossing angles? Recently, it was shown that recognizing straight line RAC graphs is NP-hard [5], so it is likely that recognizing graphs that admit αAC_2 drawings for a given α or just some α are hard as well. It might also be interesting to find or approximate the minimum value t for a given graph G such that G admits a polyline drawing with at most two bends per edge and t possible crossing angles.

Regarding x-monotone drawings let us define y-monotonicity like x-monotonicity after a 90-degree rotation; not very exciting by itself, but what happens if we want embeddings that are *bi-monotone*, that is, both x and y-monotone?

Question 7.3. *If there exists a drawing of a graph which is bi-monotone, is there a straight-line drawing with the same ordering of the vertices along the x-axis and y-axis?*

Question 7.4. *Can we test bi-monotone planarity efficiently? In other words, given a planar graph G with its vertices drawn in the plane in a general position can we decide in a polynomial time if we can extend the given drawing of the vertices of G to a bi-monotone embedding of G?*

We note that there is no straightforward generalization even of the weak version of Hanani-Tutte Theorem to bi-monotone drawings (see Figure 36). Thus, there exists no algorithm for the previous question analogous to our algorithm for finding a monotone embedding of an ordered graph.

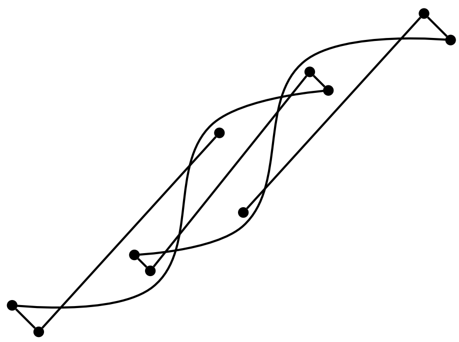

Figure 36: An ordered path that does not admit a bi-monotone embedding drawn so that every pair of edges crosses an even number of times.

We conclude the thesis with open problems regarding relations of various crossing numbers. Lemma 6.5 allowed us to resolve the question whether counting or not counting crossings of adjacent edges can sometimes make a difference, and we saw that it indeed can make a difference. However, this is only a tiny step towards a satisfactory understanding of all the mutual relations of the variants of crossing numbers we defined in this thesis. Nevertheless, Lemma 6.5 brings us a hope that to separate other crossing number variants can be within the reach, as long as for the corresponding monotone versions a separation is also possible.

A negative answer to one the most intriguing problems in the area "Does the crossing number and the pair crossing number always coincide ?", would follow from the negative answer to the following question.

Question 7.5. *Does* $\text{mon-pcr}(G) = \text{mon-cr}(G)$ *hold for all ordered graphs?*

Despite a considerable effort we were not able to answer the other similar questions about monotone crossing numbers as well. We are particularly interested in the following problem.

Question 7.6. *Does* $\text{mon-cr}_{-}(G) = \text{mon-cr}(G)$ *hold for all ordered graphs?*

So far, it seems that the monotone versions of crossing numbers give larger separations. Thus, there is a reason to believe that if the answer to one of the last two questions turns out to be positive, the answer for the corresponding crossing number problem is positive as well. On the other hand, to the best of our knowledge, there exists no non-trivial proof showing that two crossing number variants (out of those we defined including also the monotone variants) always coincide. Thus, we do not have any evidence suggesting that to obtain a positive answer in the monotone variant is easier than in the general case.

8 Bibliography

References

[1] E. Ackerman, J. Fox, J. Pach, and A. Suk. On grids in topological graphs. In *Symposium on Computational Geometry*, pages 403–412, 2009.

[2] E. Ackerman, R. Fulek, and C. D. Tóth. On the size of graphs that admit polyline drawings with few bends and crossing angles. In *Graph Drawing*, pages 1–12, 2010.

[3] E. Ackerman and G. Tardos. On the maximum number of edges in quasi-planar graphs. *Journal of Combinatorial Theory, Series A*, 114(3):563–571, 2007.

[4] M. Ajtai, V. Chvátal, M. Newborn, and E. Szemerédi. Crossing-free subgraphs. *Theory and Practice of Combinatorics*, 60:9–12, 1982.

[5] E. N. Argyriou, M. A. Bekos, and A. Symvonis. The straight-line rac drawing problem is NP-hard. In *Proceedings of the 37th international conference on Current trends in theory and practice of computer science*, SOFSEM'11, pages 74–85, Berlin, Heidelberg, 2011. Springer-Verlag.

[6] K. Arikushi, R. Fulek, B. Keszegh, F. Moric, and C. D. Tóth. Graphs that admit right angle crossing drawings. In *WG*, pages 135–146, 2010.

[7] K. Arikushi and C. D. Tóth, 2010. manuscript.

[8] P. Brass, W. Moser, and J. Pach. *Research Problems in Discrete Geometry.* Springer, New York, 2005.

[9] G. Cairns, M. Mcintyre, and Y. Nikolayevsky. The thrackle conjecture for k_5 and $k_{3,3}$. *In: Towards a theory of Geometric Graphs, Contemp. Math.*, 342(2):35–54, 2004.

[10] G. Cairns and Y. Nikolayevsky. Bounds for generalized thrackles. *Discrete Comput. Geom.*, 23(2):191–206, 2000.

[11] G. Cairns and Y. Nikolayevsky. Generalized thrackle drawings of non-bipartite graphs. *Discrete Comput. Geom.*, 41(1):119–134, 2009.

[12] H. de Fraysseix. Trémaux trees and planarity. In *The International Conference on Topological and Geometric Graph Theory*, volume 31 of *Electron. Notes Discrete Math.*, pages 169–180. Elsevier Sci. B. V., Amsterdam, 2008.

[13] H. de Fraysseix and P. O. de Mendez. Public implementation of a graph algorithm library and editor. http://pigale.sourceforge.net/.

[14] H. de Fraysseix, P. O. de Mendez, and P. Rosenstiehl. Trémaux trees and planarity. *Int. J. Found. Comput. Sci.*, 17(5):1017–1030, 2006.

[15] H. de Fraysseix and P. Rosenstiehl. A characterization of planar graphs by Trémaux orders. *Combinatorica*, 5(2):127–135, 1985.

[16] G. Di Battista and E. Nardelli. Hierarchies and planarity theory. *IEEE Trans. Systems Man Cybernet.*, 18(6):1035–1046 (1989), 1988.

[17] W. Didimo, P. Eades, and G. Liotta. Theoretical computer science. *Drawing graphs with right angle crossings*, 412(39):5156–5166, 2011.

[18] R. Diestel. *Graph Theory.* Springer, New York, 1997.

[19] V. Dujmović, J. Gudmundsson, P. Morin, and T. Wolle. Notes on large angle crossing graphs. *Chicago Journal of Theoretical Computer Science*, 2011.

[20] 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, 2009.

[21] J. Fox and B. Sudakov. Density theorems for bipartite graphs and related ramsey-type results. *Combinatorica*, 29:153–196, March 2009.

[22] R. Fulek and J. Pach. A computational approach to conway's thrackle conjecture. *Computational Geometry*, 44(6-7):345–355, 2011.

[23] R. Fulek, M. J. Pelsmajer, M. Schaefer, and D. Štefankovič. Adjacent crossings matter. Accepted for GD 2011.

[24] R. Fulek, M. J. Pelsmajer, M. Schaefer, and D. Štefankovič. Hanani-tutte, monotone drawings, and level-planarity. Accepted for WG 2011.

[25] J. E. Green and R. D. Ringeisen. Combinatorial drawings and thrackle surfaces. *Graph Theory, Combinatorics, and Algorithms*, 2:999–1009, 1995.

[26] C. C. H. Hanani). Über wesentlich unplättbare Kurven im dreidimensionalen Raume. *Fundamenta Mathematicae*, 23:135–142, 1934.

[27] M. Harrigan and P. Healy. Practical level planarity testing and layout with embedding constraints. In *Graph drawing*, volume 4875 of *Lecture Notes in Comput. Sci.*, pages 62–68. Springer, Berlin, 2008.

[28] P. Healy. Personal communication (January 2011).

[29] P. Healy and A. Kuusik. Algorithms for multi-level graph planarity testing and layout. *Theoret. Comput. Sci.*, 320(2-3):331–344, 2004.

[30] P. Healy, A. Kuusik, and S. Leipert. A characterization of level planar graphs. *Discrete Math.*, 280(1-3):51–63, 2004.

[31] L. S. Heath and S. V. Pemmaraju. Recognizing leveled-planar dags in linear time. In F. J. Brandenburg, editor, *Graph Drawing, GD'95*, volume 1027 of *Lecture Notes in Comput. Sci.*, pages 300–311. Springer-Verlag, Berlin, 1996.

[32] L. S. Heath and S. V. Pemmaraju. Stack and queue layouts of directed acyclic graphs. II. *SIAM J. Comput.*, 28(5):1588–1626 (electronic), 1999.

[33] J. Hopcroft and R.E. Tarjan Efficient planarity testing. *Journal of the Association for Computing Machinery*, 21(4):549–568.

[34] H. Hopf and E. Pannwitz. Aufgabe nr. 167. *Jahresbericht Deutsch. Math.-Verein.*, 43:114, 1934.

[35] W. Huang. Using eye tracking to investigate graph layout effects. In *6th Asia-Pacific Symp. Visualization(APVIS),IEEE*, pages 97–100, 2007.

[36] W. Huang, S.-H. Hong, and P. Eades. Effects of crossing angles. In *IEEE Pacific Visualization Symp.*, pages 41–46, 2008.

[37] M. Jünger, S. Leipert, and P. Mutzel. Level planarity testing in linear time. In *Graph drawing (Montréal, QC, 1998)*, volume 1547 of *Lecture Notes in Comput. Sci.*, pages 224–237. Springer, Berlin, 1998.

[38] M. Jünger, S. Leipert, and P. Mutzel. Pitfalls of using *PQ*-trees in automatic graph drawing. In G. DiBattista, editor, *Proceedings of the 5th International Symposium on Graph Drawing, GD'97 (Rome, Italy, September 18-20, 1997)*, volume 1353 of *LNCS*, pages 193–204. Springer-Verlag, Berlin, 1998.

[39] P. C. Kainen. A lower bound for crossing numbers of graphs with applications to K_n, $K_{p,q}$, and $Q(d)$. *J. Combinatorial Theory Ser. B*, 12:287–298, 1972.

[40] D. J. Kleitman. A note on the parity of the number of crossings of a graph. *J. Combinatorial Theory Ser. B*, 21(1):88–89, 1976.

[41] T. Leighton. Complexity issues in VLSI. *Foundations of Computing Series*, 1983.

[42] S. Leipert. *Level Planarity Testing and Embedding in Linear Time*. PhD thesis, Universität zu Köln, Köln, 1998.

[43] L. Lovász, J. Pach, and M. Szegedy. On conway's thrackle conjecture. *Discrete & Computational Geometry*, 18:369–376, 1997.

[44] J. Matoušek. *Using the Borsuk-Ulam theorem*. Universitext. Springer-Verlag, Berlin, 2003. Lectures on topological methods in combinatorics and geometry, Written in cooperation with Anders Björner and Günter M. Ziegler.

[45] J. Matousek, M. Tancer, and U. Wagner. Hardness of embedding simplicial complexes in $\mathbb{R}^d$. In C. Mathieu, editor, *Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2009, New York, NY, USA, January 4-6, 2009*, pages 855–864. SIAM, 2009.

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

[47] J. Pach, R. Radoicic, G. Tardos, and G. Toth. Improving the crossing lemma by finding more crossings in sparse graphs. *Discrete & Computational Geometry*, 36:527–552, 2006.

[48] J. Pach and G. Tóth. Thirteen problems on crossing numbers. *Geombinatorics*, 9(4):194–207, 2000.

[49] J. Pach and G. Tóth. Which crossing number is it anyway? *J. Combin. Theory Ser. B*, 80(2):225–246, 2000.

[50] J. Pach and G. Tóth. Monotone drawings of planar graphs. *J. Graph Theory*, 46(1):39–47, 2004.

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

[52] 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.

[53] M. J. Pelsmajer, M. Schaefer, and D. Štefankovič. Odd crossing number is not crossing number. Technical Report TR05-005, DePaul University, April 2005.

[54] M. J. Pelsmajer, M. Schaefer, and D. Štefankovič. Removing even crossings. *J. Combin. Theory Ser. B*, 97(4):489–500, 2007.

[55] M. J. Pelsmajer, M. Schaefer, and D. Štefankovič. Removing even crossings. *J. Combin. Theory Ser. B*, 97(4):489–500, 2007.

[56] M. J. Pelsmajer, M. Schaefer, and D. Štefankovič. Odd crossing number and crossing number are not the same. *Discrete Comput. Geom.*, 39(1):442–454, 2008.

[57] M. J. Pelsmajer, M. Schaefer, and D. Štefankovič. Removing even crossings on surfaces. *European Journal of Combinatorics*, 30(7):1704–1717, 2009.

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

[59] M. J. Pelsmajer, M. Schaefer, and D. Štefankovič. Crossing numbers of graphs with rotation systems. *Algorithmica*, 60:679–702, 2011.

[60] A. Perlstein and R. Pinchasi. Generalized thrackles and geometric graphs in $\mathbb{R}^3$ with no pair of strongly avoiding edges. *Graphs and Combinatorics*, 24:373–389, 2008.

[61] B. Piazza, R. Ringeisen, and S. Stueckle. Subthrackleable graphs and four cycles. *Discrete Mathematics*, 127(1-3):265 – 276, 1994.

[62] R. Pinchasi, 2010. private communication.

[63] U. problems. Chairman: P. erdos. *In:* Combinatorics (Proc. Conf. Combinatorial Math., Math. Inst., Oxford), pages 351–363, 1972.

[64] B. Randerath, E. Speckenmeyer, E. Boros, P. L. Hammer, A. Kogan, K. Makino, B. Simeone, and O. Cepek. A satisfiability formulation of problems on level graphs. *Electronic Notes in Discrete Mathematics*, 9:269–277, 2001.

[65] L. Rédei. Ein kombinatorischer satz. *Acta Litteraria Szeged*, 7:39–43.

[66] R. D. Ringeisen. Two old extremal graph drawing conjectures: progress and perspectives. *Congressus Numerantium*, 115:91–103, 1996.

[67] M. Schaefer. Hanani-Tutte and related results. To appear in Bolyai Memorial Volume.

[68] J. Sutherland. Lösung der aufgabe 167. *Jahresbericht Deutsch. Math.-Verein.*, 45:33–35, 1935.

[69] L. Székely. Progress on crossing number problems. In P. Vojtáš, M. Bieliková, B. Charron-Bost, and O. Sýkora, editors, *SOFSEM 2005: Theory and Practice of Computer Science*, volume 3381 of *Lecture Notes in Computer Science*, pages 53–61. Springer Berlin / Heidelberg, 2005. 10.1007/978-3-540-30577-4_8.

[70] L. A. Székely. A successful concept for measuring non-planarity of graphs: the crossing number. *Discrete Math.*, 276(1-3):331–352, 2004.

[71] L. A. Székely. An optimality criterion for the crossing number. *Ars Math. Contemp.*, 1(1):32–37, 2008.

[72] G. Tóth. Note on the pair-crossing number and the odd-crossing number. *Discrete Comput. Geom.*, 39(4):791–799, 2008.

[73] W. T. Tutte. Toward a theory of crossing numbers. *J. Combinatorial Theory*, 8:45–53, 1970.

[74] P. Valtr. On the pair-crossing number. In *Combinatorial and Computational Geometry*, volume 52 of *Math. Sci. Res. Inst. Publ.*, pages 569–575. Cambridge University Press, Cambridge, 2005.

[75] D. West. Open problems - graph theory and combinatorics. `http://www.math.uiuc.edu/~west/openp/` (accessed April 7th, 2005).

[76] D. R. Woodall. Thrackles and deadlock. *Combinatorial Mathematics and its Applications (Proc. Conf., Oxford, 1969)*, 333–347, 1971.

9 CURRICULUM VITÆ

PERSONAL DATA

First name: **Radoslav**
Surname: **Fulek**
Date of birth: 30^{th} March, 1982.
Nationality: Slovak
Other languages: English (fluent), French (good), German (basic)
e-mail: *radoslav.fulek@epfl.ch*

STUDIES

2008– Ecole Polytechnique Fédérale de Lausanne (EPFL)
Assistant-doctorant of János Pach

2006–2008 Simon Fraser University (SFU), Burnaby, BC, Canada
Graduate Studies in Theoretical Computer Science
-Supervisor: Gábor Tardos.
-Master thesis in Computational Geometry 2008.

2000–2005 Comenius University (UK), Bratislava, Slovakia
Undergraduate and Graduate Studies at Faculty of Mathematics, Physics and Informatics, Department of Informatics
-Supervisor: Ondrej Sýkora.
-Master thesis in Graph Drawing 2005.

1996–2000 Gymnázium Veľká Okružná, Žilina, Slovakia
Class specialized in mathematics

TEACHING EXPERIENCE

2008–
-Teaching assistant for Linear Algebra, Graph Theory, Introduction to Combinatorics, Geometry and Mathematics for Architects.

2006–2008
-Teaching assistant for Introduction to Computing Science/Programming,

Database Systems at SFU

PUBLICATIONS

To appear:

-*On the size of graphs that admit polyline drawings with few bends and crossing angles* (with Eyal Ackerman, and Csaba D. Tóth), Siam J. on Discr. Math.
-*Graphs that Admit Right Angle Crossing Drawings* (with Karin Arikushi, Radoslav Fulek, Balázs Keszegh, Filip Morić, and Csaba D. Tóth), Computational Geometry
-*On disjoint crossing–families in geometric graphs* (with Andrew Suk), Thirty Essays in Geometric Graph Theory, J. Pach ed.
-*Hanani-Tutte, Monotone Drawings and Level-Planarity* (with Michael Pelsmajer, Marcus Schaefer and Daniel Štefankovič), Thirty Essays in Geometric Graph Theory, J. Pach ed.
-*Convex Obstacle Numbers of Outerplanar Graphs and Bipartite Permutation Graphs* (with Noushin Saeedi, and Deniz Sariöz), Thirty Essays in Geometric Graph Theory, J. Pach ed.

2011

-*Hanani-Tutte and Monotone Drawings* (with Michael Pelsmajer, Marcus Schaefer and Daniel Štefankovič), in Proc. 37th Workshop on Graph Theoretic Concepts in Comp. Sci. (Teplá, 2011), LNCS, Springer, 283–294
-*Orthogeodesic Point–Set Embeddability* (with E. Di Giacomo, F. Frati, L. Grilli, M. Krug), in Proc. 19th International Symposium on Graph Drawing (Eindhoven, 2011), LNCS, Springer, 52–63
-*On the Page Number of Upward Planar Directed Acyclic Graphs* (with F. Frati and A. Ruiz Vargas), in Proc. 19th International Symposium on Graph Drawing (Eindhoven, 2011), LNCS, Springer, 391–402
-*Adjacent crossings matter* (with M. Pelsmajer, M. Schaefer and D. Štefankovič), in Proc. 19th International Symposium on Graph Drawing (Eindhoven, 2011), LNCS, Springer, 343–354
-*Bounding diameter of planar graphs* (with Filip Morić and David Pritchard), Discrete Mathematics, vol. 311, issue 5, 2011, 327–335
-*A computational approach to Conway's thrackle conjecture* (with János Pach), Computational Geometry 44, Issues 6–7, 2011, 345–355

2010

-*Graphs that Admit Right Angle Crossing Drawings* (with Karin Arikushi,

Balázs Keszegh, Filip Morić, and Csaba D. Tóth), in Proc. 36th Workshop on Graph Theoretic Concepts in Comp. Sci. (Zarós, 2010), LNCS, Springer, 135–146
-*Polygonizations avoiding a set of points* (with Balázs Keszegh, Filip Morić, and Igor Uljarević), CCCG 2010, 273–276
-*Coloring geometric hypergraph defined by an arrangement of half–planes*, CCCG 2010, 71–74
-*A computational approach to Conway's thrackle conjecture* (with János Pach), in Proc. 18th International Symposium on Graph Drawing (Konstanz, 2010), LNCS, Springer, 226–237
-*On the size of graphs that admit polyline drawings with few bends and crossing angles* (with Eyal Ackerman, and Csaba D. Tóth), in Proc. 18th International Symposium on Graph Drawing (Konstanz, 2010), LNCS, Springer, 1–12
-*A tight lower bound for convexly independent subsets of the Minkowski sums of planar point sets* (with with Ondřej Bílka, Kevin Buchin, Masashi Kiyomi, Yoshio Okamoto, Shin–ichi Tanigawa, and Csaba D. Tóth), The Electronic Journal of Combinatorics 17 (1), 2010

2009
-*A tight lower bound for convexly independent subsets of the Minkowski sums of planar point sets* (Kevin Buchin, Masashi Kiyomi, Yoshio Okamoto, Shin–ichi Tanigawa, and Csaba D. Tóth), in Proc. 7th Japan Conference on Computational Geometry and Graphs (Kanazawa, 2009), JAIST

2008
-*Intersecting convex sets by rays* (with Andreas Holmsen and János Pach), Symposium on Computational Geometry 2008, 385–391

2005
-*Outerplanar Crossing Numbers of 3-Row Meshes, Halin Graphs and Complete p-Partite Graphs* (with Hongmei He, Ondrej Sýkora, Imrich Vrťo), in Proc. 31st Conference on Current Trends in Theory and Practice of Computer Science SOFSEM 2005, LNCS, Springer, 376–379

ATTENDED CONFERENCES AND WORKSHOPS

2011
-*Séminaire du 3ème cycle romand de Recherche Opérationnelle*, Zinal, Switzerland, Jan. 16-20, 2011

-*2nd Emléktábla Workshop*, Gyöngyöstarján, Hungary, Jan 24-27, 2011
-*9th Gremo's Workshop on Open Problems*, Wergenstein, Switzerland, Feb. 7-10, 2011
-*Bertinoro Workshop on Graph Drawing*, Bertinoro, Italy, March 6-11, 2011
-*European Workshop on Computational Geometry*, Morschach, Switzerland, March 28-30, 2011
-*STTI'11 - Současné trendy teoretické informatiky*, Prague, Czech Rep., June 10-11, 2011
-*37th Workshop on Graph Theoretic Concepts in Comp. Sci.*, Teplá, Czech Rep., June 21-24, 2011
-*European Conference on Combinatorics, Graph Theory and Applications*, Budapest, Hungary, Aug. 29 - Sept. 2, 2011
-*19th International Symposium on Graph Drawing*, Eindhoven, Netherlands, Sept. 21-23, 2011

2010

-*Séminaire du 3ème cycle romand de Recherche Opérationnelle*, Zinal, Switzerland, Jan. 17-21, 2010
-*36th Workshop on Graph Theoretic Concepts in Comp. Sci.*, Záros, Greece, June 28-30, 2010
-*1st Emléktábla Workshop*, Gyöngyöstarján, Hungary, July 26-29, 2010
-*18th International Symposium on Graph Drawing*, Konstanz, Germany, Sept. 21-24, 2010
-*Special Semester on Discrete and Computational Geometry - Culminating Workshop*, Lausanne, Switzerland, Nov. 29 - Dec. 3, 2010

2009

-*Séminaire du 3ème cycle romand de Recherche Opérationnelle*, Zinal, Switzerland, Jan. 18-22, 2009
-*Algorithmic and Combinatorial Geometry Conference*, Budapest, Hungary, June 15 - 19, 2009
-*7th Gremo's Workshop on Open Problems*, Stels, Switzerland, July 6-10, 2009

2008

-*Symposium on Computational Geometry*, College Park, MD, USA, June 9-11, 2008
-*SIAM Conference on Discrete Mathematics, The University of Vermont, Burlington*, VT, USA, June 16-19, 2008
-*Sixth Joint Operations Research Days*, Lausanne, Switzerland, Sept. 11, 2008

2007

-*1st Canadian Discrete and Algorithmic Mathematics (CanaDam)*, Banff, AB, Canada, May 28-31, 2007

-*Workshop on Algorithms, Combinatorics, and Geometry*, University of North Texas, Denton, TX, USA, Nov. 29 - Dec. 1, 2007

2005

-*31st Annual Conference on Current Trends in Theory and Practice of Informatics (SOFSEM)*, Liptovský Ján, Slovak Republic, Jan. 22-28, 2005

Appendix

A Backtracking algorithm

For sake of completeness in this section we describe a backtracking algorithm checking, whether a given dumbbell $G = (V, E)$ can be drawn as a thrackle. We orient the edges of G, so that we can traverse them by a single walk, so called Euler's walk, during which we visit each edge just once. We use the notation from Section 3.3.

Let us start with a description of the routines used by our algorithm.

The routine UPDATE(π_e, e', pos) returns the updated permutation $\pi_e \in E'^{m'(e)}_e$, which corresponds to adding one more crossing vertex to an already constructed part of (G', Π') corresponding to a subgraph $G' = (V', E')$ of G, where $e \in E$, $e' \in E'$, $m'(e)$ returns the number of crossings of e already modeled by (G', Π'), and $\Pi'(G') := \{\pi_e \in E'^{m'(e)}_e | e \in E'\}$. UPDATE($\pi(e)$, e', pos) returns the permutation π'_e whose length is by one longer than π_e, such that

$$\pi'_e(i) = \begin{cases} \pi_e(i) & \text{if } i < pos \\ e' & \text{if } i = pos \\ \pi_e(i-1) & \text{if } i > pos \end{cases}$$

REVERSE_UPDATE(π_e, e', *pos*) corresponds to the reverse operation of the operation UPDATE(π_e, e', pos). PICK_NEXT_EDGE(G) returns a next edge in our Euler's walk. In order to check, whether G can be drawn as a thrackle the algorithm just calls the procedure BACKTRACKING(e) for an edge $e \in E$. The algorithm returns *true* if G can be drawn as a thrackle, and it returns *false* if G cannot be drawn as a thrackle. In our description of the algorithm we restrain from all optimization details, which were mentioned in Section 3.3. The pseudocode of the backtracking routine follows.

Algorithm 1: Thrackleabilty testing

```
BACKTRACKING (e ∈ E(G)) ;
begin
    if (G', Π') cannot be extended then
        return true

    if e = -1 then
        e = PICK_NEXT_EDGE(G)

    if e has crossed all edges in E'_e then
        BACKTRACKING(-1)

    else
        forall the e' ∈ E'_e which e has not already crossed do
            for pos = 1 to length(π_e') do
                π_e' = UPDATE(π_e', e, pos) ;
                π_e = UPDATE(π_e, e', LENGTH(π_e)+1) ;
                if IS_PLANAR((G', Π')) then
                    if BACKTRACKING(e) then
                        return true
                    else
                        REVERSE_UPDATE(π_e', e, pos) ;
                        REVERSE_UPDATE(π_e, e', LENGTH(π_e))

                else
                    REVERSE_UPDATE(π_e', e, pos) ;
                    REVERSE_UPDATE(π_e, e', LENGTH(π_e))

            end
        end
    return false
end
```