Hanani-Tutte, Monotone Drawings, and Level-Planarity*

Radoslav Fulek†

Ecole Polytechnique Fédérale de Lausanne
Lausanne, Switzerland
radoslav.fulek@epfl.ch

Michael J. Pelsmajer‡

Department of Applied Mathematics
Illinois Institute of Technology
Chicago, Illinois 60616, USA
pelsmajer@iit.edu

Marcus Schaefer

Department of Computer Science
DePaul University
Chicago, Illinois 60604, USA
mschaefer@cs.depaul.edu

Daniel Štefankovič

Computer Science Department
University of Rochester
Rochester, NY 14627-0226
stefanko@cs.rochester.edu

October 18, 2011

Abstract

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. Pach and Tóth showed that if a graph has an x-monotone drawing in which every pair of edges crosses an even number of times, then the graph has an x-monotone embedding in which the x-coordinates of all vertices are unchanged. We give a new proof of this result and strengthen it by showing that the conclusion remains true even if adjacent edges are allowed to cross each other oddly. This answers a question posed by Pach and Tóth. We show that a further strengthening to a "removing even crossings" lemma is impossible by separating monotone versions of the crossing and the odd crossing number.

Our results extend to level-planarity, which is a well-studied generalization of x-monotonicity. We obtain a new and simple algorithm to test level-planarity in quadratic time, and we show that x-monotonicity of edges in the definition of level-planarity can be relaxed.

---

*An earlier version of this paper appeared in the proceedings of WG 2011.

†The first author gratefully acknowledges support from the Swiss National Science Foundation Grant No. 200021-125287/1.

‡The second author gratefully acknowledges the support from NSA Grant H98230-08-1-0043 and the Swiss National Science Foundation Grant No. 200021-125287/1.

1 Introduction

The classic Hanani-Tutte theorem states that if a graph can be drawn in the plane so every pair of independent edges crosses an even number of times, then it is planar [9, 32]. (Two edges are *independent* if they do not share an endpoint.) There are many ways to look at this result: in algebraic topology it is seen as a special case of the van Kampen-Flores theorem [22, Chapter 5] which identifies obstructions to embeddability in topological spaces. This point of view leads to challenging open questions (see, for example, [23]); even in two dimensions—for surfaces—the problem is not understood well. (See [31] for a survey of what we do know.)

Here, we study a variant of the problem which was introduced by Pach and Tóth [25]. A curve is *x-monotone* if it 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. In this context, the natural analogue of the Hanani-Tutte theorem is that for any x-monotone drawing in which every pair of independent edges crosses an even number of times, there is an x-monotone *embedding* (that is, a crossing-free drawing) with the same vertex locations. The truth of this result was left as an open problem by Pach and Tóth. We prove this monotone Hanani-Tutte theorem as Theorem 3.1 in Section 3.

The weak version of the classic Hanani-Tutte theorem states that if a graph can be drawn so that *every* pair of edges crosses evenly, then it is planar. (For background and variants of the weak Hanani-Tutte theorem, see [31].) For x-monotone drawings this translates to the claim that if there is an x-monotone drawing in which every pair of edges crosses an even number of times, then there is an x-monotone embedding with the same vertex locations. This weak monotone Hanani-Tutte theorem was first proved by Pach and Tóth.[1] We give a new proof of this result as Theorem 2.1 in Section 2.

Traditionally, Hanani-Tutte style results are obtained via characterizations by obstructions. This can lead to very slick proofs, like Kleitman's proof of the Hanani-Tutte theorem for the plane (using obstructions K_5 and $K_{3,3}$) [19]. There are two drawbacks to this approach: complete obstruction sets are often unknown, e.g., for the torus or, in spite of several attempts (as discussed in [7]), for x-monotone embeddings. Secondly, this approach is of little help algorithmically. Pach and Tóth proved the weak monotone Hanani-Tutte theorem using an approach of Cairns and Nikolayevsky [2] with which they proved the weak Hanani-Tutte theorem for orientable surfaces. Our approach is based on a different line of attack which began in [28].

In Section 4 we establish a connection between x-monotonicity and another well-known graph drawing notion, level-planarity. Through this connection, the monotone Hanani-Tutte theorem (Theorem 3.1) leads to a simple, quadratic-time algorithm for recognizing level-planar graphs. While the best-known algorithm for this problem runs in linear time [17], it is quite compli-

[1] There is a gap in the original argument; an updated version is now available [25, 26].

cated. There have been previous claims for simple quadratic time algorithms for level-planarity testing, which we discuss in Remark 4.3.

The condition that edges are x-monotone in the monotone Hanani-Tutte theorems can be replaced by a weaker notion we call x-bounded. Let $x(v)$ denote the x-coordinate of a vertex v located in the plane. 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 2.7 and Remark 3.2). As a consequence, we obtain a relaxed, yet equivalent, definition of level-planarity (Corollary 4.5). We also describe an even weaker condition (nearly x-bounded) in Section 2.1 and show that we can get similar results for it as well (see Lemma 2.8 and Remark 3.2).

The classic Hanani-Tutte theorems have extensions that bound crossing number in terms of odd crossing number and independent odd crossing number, with equality for very small values [24, 28, 29]. We will see in Section 5 that such extensions fail for monotone odd and monotone independent odd crossing numbers. Also, Theorem 2.1 may prompt the reader familiar with Hanani-Tutte style results (in particular [24, Theorem 1] and [28, Theorem 2.1]) to ask whether a stronger result is true: a "removing even crossings" lemma which would say that all even edges can be made crossing-free even in the presence of odd edges (while maintaining x-monotonicity and vertex locations). We will see in Section 5 that there cannot be any such lemma for x-monotone drawings.

We end this section by stating a few definitions. 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 order of the edges leaving the vertex towards the right (left). 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. Our multigraphs will never have loops. For any graph G and $S \subseteq V(G)$, let $G[S]$ denote the *subgraph induced by S*, which is the graph on vertex set S with edge set $\{uv \in E(G) : u \in S, v \in S\}$.

2 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 2.1 (Weak Monotone Hanani-Tutte; Pach, Tóth [25, 26])**.** *If G has an x-monotone even drawing, then G has an x-monotone embedding with the same vertex locations and rotation system.*

Our goal in this section is to give a new proof of Theorem 2.1.

Remark 2.2. By stretching and compressing and x-monotone drawing in the plane horizontally we can change the x-coordinates of vertices arbitrarily as long as their relative x-order remains the same, and all the edges will remain x-monotone. We can also alter the y-coordinate of any vertex by stretching the plane vertically near that vertex, so that all edges remain x-monotone and all other vertices are fixed. Thus, we can modify an x-monotone drawing to relocate vertices arbitrarily, as long as the relative x-order of vertices is unchanged.

As a result, in any redrawing of an x-monotone drawing in which the relative x-order of the vertices does not change, we may as well assume that the location of every vertex has remained unchanged. Alternatively, we may instead assume that the vertices in an x-monotone drawing are located at the points $(1, 0), \ldots, (|V|, 0)$. The same argument applies if the edges are x-bounded or nearly x-bounded (defined in Section 2.1) rather than x-monotone. (We briefly consider drawings with straight-line edges in Remark 2.4, and in that context the argument no longer works.)

The result claimed by Pach and Tóth in [25, Theorem 1.1] is almost the same as Theorem 2.1, but instead of maintaining rotations, they state that one can find an *equivalent* x-monotone embedding. Here, two drawings are equivalent if no edge changes whether it passes above or below a vertex. However, the example on the left in Figure 1 shows that one cannot hope to maintain equivalence in this sense.

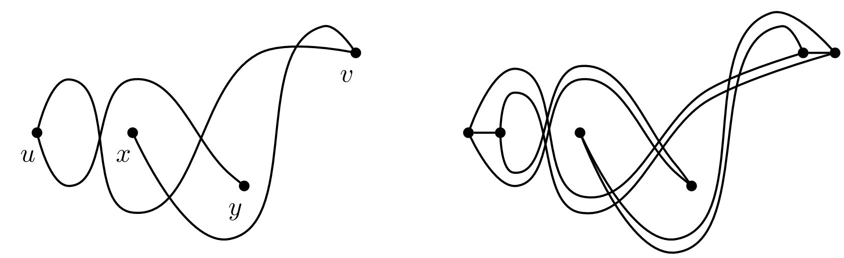

Figure 1: *(Left)* An x-monotone even drawing. Since x is above uv and y is below uv, any equivalent x-monotone embedding with the same relative x-ordering of the vertices will have uv below x and above y. But then xv is above uy, so it is not equivalent. *(Right)* Essentially the same argument applies to this 2-connected example.

The proof in [25] contains a gap: it is not immediately clear how multiple faces that share a boundary can be embedded simultaneously.[2] Eliminating the gap requires dropping equivalence. Pach and Tóth have prepared an updated version of the paper that includes a more detailed argument [26].[3] As the graph on the right in Figure 1 shows, the counterexample can be made 2-connected,

[2]On page 42 of [25], in the text after Lemma 2.1, D_κ cannot necessarily be glued together without changing equivalence.

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

so equivalence cannot be obtained by assuming 2-connectedness. On the other hand, see Corollary 2.5 for a positive result about equivalent redrawing.

In our proof of Theorem 2.1, 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 $x = x_2$. 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 use the following redrawing tool.

Lemma 2.3. *Let f be an inner face of an x-monotone embedding of a multigraph G, with m_f and M_f being the leftmost and the rightmost vertex of f. Add an edge m_fM_f to the embedding so that it lies in f. (Note that m_fM_f is not required to be x-monotone and and that there may be multiple ways of inserting m_fM_f into the rotations at m_f and M_f.) Then the resulting graph $G \cup \{m_fM_f\}$ has an x-monotone embedding with the same vertex locations and rotation system.*

Note that the redrawing in Lemma 2.3 destroys equivalence in the sense of Pach and Tóth [25]. This is necessarily the case; see Figure 2 for an example.

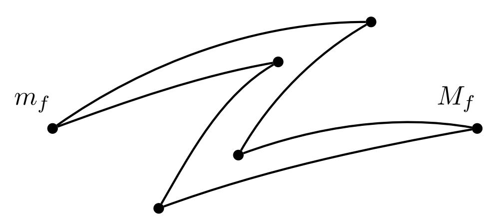

Figure 2: Although we can draw the edge m_fM_f within the Z-shaped face, any subsequent x-monotone redrawing that maintains relative vertex x-order and rotation system will not be equivalent.

Proof of Lemma 2.3. 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.[4] The boundary of f can be broken into two m_f, M_f-walks B_1, B_2, with B_1 starting above m_fM_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

[4] *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.

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 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 2.1. 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 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 3). 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 $x_1 x_2$, obtaining an embedding of the original graph (since we kept the rotation) with the original rotation.

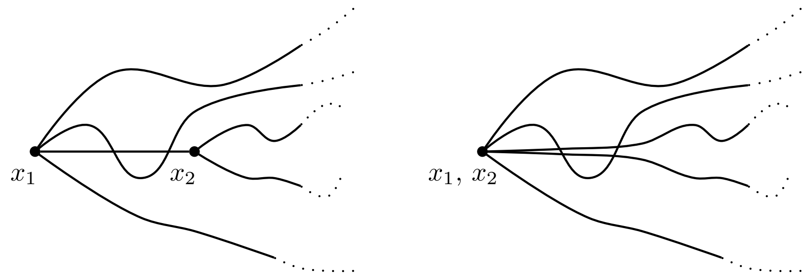

Figure 3: How to contract edge x_2x_1 towards x_1 and merging rotations.

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 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 $G - x_1$ is 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 4).

- (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.

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,

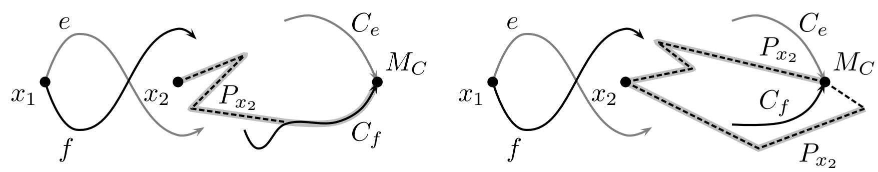

Figure 4: *(Left)* Case (i): P_{x_2} is dashed, P'_{x_2} is the thick gray path from x_2 to M_C; *(Right)* Case (ii): Both subcases are displayed: the top P_{x_2} stays to the left of M_C, while the bottom P_{x_2} passes to the right of M_C. P'_{x_2} as thick gray path in both cases.

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, meaning that 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 if e_1, f_1 interleaves 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 (or edge) 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 2.3, 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 $G - V(H')$ was redrawn using Lemma 2.3. 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. □

Remark 2.4. There is a straight-line variant of Theorem 2.1 if we allow the y-coordinates of vertices to change. This has nothing to do with the Hanani-Tutte part of the result; it is entirely due to the fact that any x-monotone embedding can be turned into a straight-line embedding in which every vertex keeps its x-coordinate [6, 25]. This redrawing can lead to an exponential blow-up in the area required for the drawing [21].[5]

All redrawing steps in the proof of Theorem 2.1 maintain equivalence except for applications of Lemma 2.3. This part of the proof only arises if $G - \{x_1\}$ 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 [25]. 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 [5].

Corollary 2.5. *If a hierarchy G has an x-monotone even drawing, then G has an equivalent x-monotone embedding with the same vertex locations and rotation system.*

Proof. This follows from the proof of Theorem 2.1. The only operation that changes equivalence of edges and vertices in that proof is the application of Lemma 2.3. 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. □

2.1 From x-monotone to x-bounded

The x-monotonicity assumption in Theorem 2.1 can be replaced by a weaker condition. 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 2.6. *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

[5] The examples in the paper allow multiple vertices in each layer, but these can be replaced by the requirement that vertices are not too close to edges they are not incident to.

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 5). 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 5).

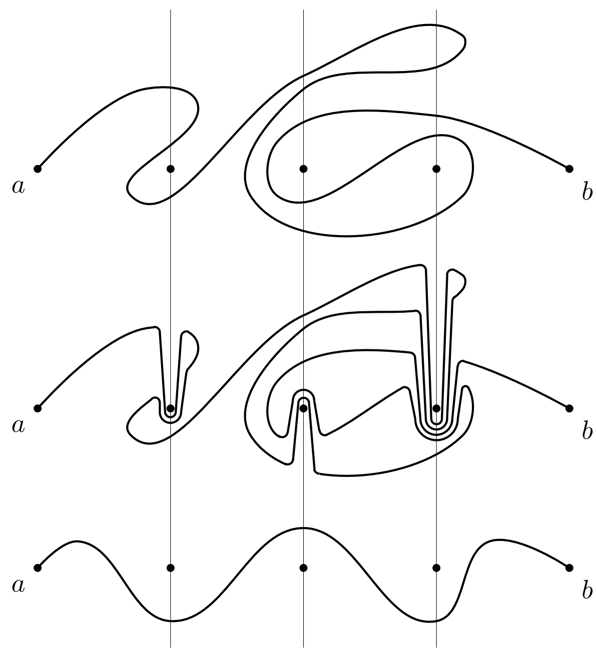

Figure 5: How to redraw an x-bounded edge. *(Top)* Before the redrawing. *(Middle)* After pushing e off the odd parts. *(Bottom)* After deforming e into an x-monotone drawing.

In hindsight we see that the redrawing in Lemma 2.6 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 2.6 gives us the following strengthening of Theorem 2.1. Later, we will use that result to strengthen Theorem 3.1, and to show that x-monotone edges can be replaced by x-bounded edges in the definition of level planarity (see Corollary 4.5 in Section 4.2).

Corollary 2.7. *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.*

Next, we give a condition weaker than x-bounded, for which we can prove some of the same results. Consider an edge e in a given drawing of a graph G with endpoints u, v such that $x(u) < x(v)$. Let C_e be the concatenation of e with the line segment from v to u. We say that e is *nearly x-bounded* if for every vertex z with $x(z) < x(u)$ or $x(z) > x(v)$, the winding number of C_e with respect to z is even.

Lemma 2.8. *Suppose we are given a drawing of a graph G with a nearly x-bounded edge e. Then e can be redrawn, without changing the remainder of the drawing, so that e is x-bounded and the parity of crossing does not change between e and any edge of G that is independent of e.*

Proof. We can gradually deform e to the line segment $\overline{uv}$, which causes e to become x-bounded. Suppose that e passed over the vertex z an odd number of times during the deformation. Since e is nearly x-bounded, it must be that $x(u) \leq x(z) \leq x(v)$. If $x(u) < x(z) < x(v)$, we can stretch e near the line $x = x(z)$ so that it passes over z once more, and e remains x-bounded. In the end, since e has passed over every vertex other than u and v an even number of times, the parity of crossing with e and any edge e' of G remains unchanged unless e' shares an endpoint with e. □

Although this does not allow us to directly generalize Corollary 2.7 to drawings with nearly x-bounded edges, we will apply Lemma 2.8 presently, in Remark 3.2.

3 Strong Hanani-Tutte for Monotone Drawings

Pach and Tóth [25] wrote "It is an interesting open problem to decide whether [the conclusion of Theorem 2.1] remains true under the weaker assumption that any two *non-adjacent* edges cross an an even number of times." The goal of this section is to establish this result.

Theorem 3.1 (Monotone Hanani-Tutte)**.** *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 3.2. Similar to Theorem 2.1 and Corollary 2.7, the statement of Theorem 3.1 remains true if we only require edges to be x-bounded or nearly x-bounded rather than x-monotone: simply redraw edges one at a time using Lemma 2.6 and/or Lemma 2.8, before applying Theorem 3.1.

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 3.3. *Suppose that G is a smallest (fewest vertices) counterexample to Theorem 3.1. 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 3.1. 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 (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 2.6 tells us that we can draw an x-monotone edge that has the same parity of crossing with all edges of $G - V(H)$ as does P.

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' by $H' \cup \{e\}$ 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 2.3, 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. □

The proof of Theorem 3.1 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 3.3. 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. 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.

In the following two lemmas, let G satisfy the assumption of Theorem 3.1, that is, we assume an even x-monotone drawing in which every pair of independent edges in G crosses evenly. Both lemmas deal with the following scenario: G contains three edges $e_i = v_0v_i$, $i \in \{1, 2, 3\}$ so that e_3 lies between e_1 and e_2 in the right rotation of v_0, with e_1 above e_2 at v_0, e_1 and e_2 cross oddly, and e_3 crosses each of the other two edges evenly.

Lemma 3.4. *For arbitrary $x_R > x(v_0)$, define G' as the graph induced by G on vertices v with $x(v_0) < x(v) \leq x_R$. Let G'_i be the component of G' that contains v_i. (If $x(v_i) > x_R$, then $G'_i = \emptyset$.)*

Suppose that G'_1, G'_2, G'_3 are pairwise disjoint and that for every i there is a path P_i (in G) from v_0 through e_i to some vertex v'_i satisfying $x(v'_i) \geq x_R$ so *that all vertices v of P_i satisfy $x(v) \geq x(v_0)$. (If $G_i = \emptyset$, then let $E(P_i) = \{e_i\}$.) Then each G'_i has no neighbors (in G) to the left of $x(v_0)$, for $i \in \{1, 2, 3\}$.*

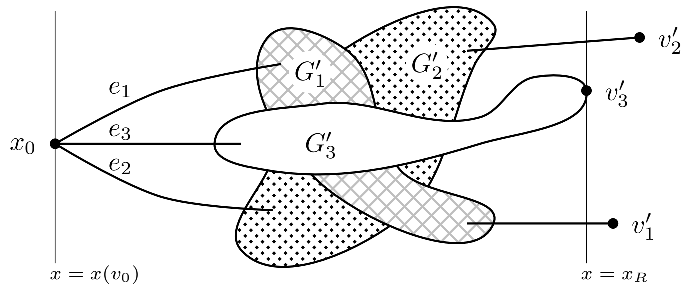

Figure 6: Lemma 3.4.

Proof. By choosing each P_i to be minimal, we can assume that for every vertex v of P_i other than its final endpoint v'_i, we have $x(v_0) \leq x(v) < x_R$, and also $x(v'_i) \geq x_R$. Note that for any point in the plane q with $x(v_0) \leq x(q) \leq x_R$ that does not lie on the curve P_i, q is either *above* or *below* P_i in the sense defined just before Lemma 3.4.

Suppose, for a contradiction, that G'_i has a neighbor v' to the left of $x(v_0)$. Then we may let P'_i be a path from v_i to v' such that every vertex of $P'_i - v'$ is in G'_i. Fix j, k so that $\{i, j, k\} = \{1, 2, 3\}$.

The paths P'_i and P_j are disjoint, so every edge of P'_i crosses P_j evenly (as every pair $(e, f) \in (P'_i, P_j)$ crosses evenly). Every vertex of $P'_i - v'$ is between $x(v_0)$ and x_R, so if v_i is above P_j, then every vertex of $P'_i - v'$ is above P_j, and when the last edge of P'_i reaches the line $x = x(v_0)$, it must be above v_0. Likewise, if v_i is below P_j, then the last edge of P'_i must pass below v_0. Similarly, v_i is above (below) P_k if and only if the last edge of P'_i passes above (below) v_0.

We split the proof into cases. Suppose that $(i, j, k) = (1, 2, 3)$. Then e_i begins in the rotation at v_0 above e_j and e_k, and e_i crosses e_j oddly and e_k evenly. Since e_i crosses other edges of P_j and P_k evenly, v_i must be below P_j and above P_k. Then by the above/below argument in the previous paragraph, the last edge of P'_i must pass both below and above v_0, a contradiction. The case $(i, j, k) = (1, 3, 2)$ is the same, and the cases with $i = 2$ are symmetric.

Suppose that $i = 3$; without loss of generality, $(j, k) = (1, 2)$. Then e_i begins in the rotation at v_0 below e_j and above e_k, and e_i crosses e_j and e_k evenly. Since e_i crosses other edges of P_j and P_k evenly, v_i must be below P_j and above P_k. Then by the earlier above/below argument, the last edge of P'_i must pass both below and above v_0, a contradiction. ∎

Lemma 3.5. *Suppose that for some distinct $j, k \in \{1, 2, 3\}$, there is a cycle C that contains e_j and e_k such that every vertex v of C satisfies $x(v) \geq x(v_0)$. Let*

v_R be the vertex on C with largest x-value. Let i be the unique index such that $\{i, j, k\} = \{1, 2, 3\}$. Suppose that v_i is not in C.

Let G'_i be the component of $G - V(C)$ that contains v_i. Then every vertex v of G'_i satisfies $x(v_0) < x(v) < x(v_R)$.

Proof. Let P_j and P_k be the v_0, v_R-paths in C that contain e_j, e_k, respectively.

Suppose that $x(v_i) > x(v_R)$. First, consider the case $(i, j, k) = (1, 2, 3)$: Since e_j and e_k begin in the rotation at v_0 below e_i, and e_i crosses e_j oddly and e_k evenly, it must be that v_j is above e_i and v_k is below e_i. (See Figure 7.) Every other edge of P_j crosses e_i evenly, so all its other vertices are also above e_i; likewise, every other vertex of P_k is below e_i. But then v_R lies above and below e_i; contradiction. The case $(i, j, k) = (1, 3, 2)$ is the same, and the cases with $i = 2$ are symmetric. Suppose that $i = 3$; without loss of generality, $(j, k) = (1, 2)$. Then in the rotation at v_0, e_j is above e_i and e_k is below e_i, and e_i crosses e_j and e_k evenly. Then v_j is above e_i and v_k is below e_i. Then (as seen earlier), every vertex of $P_j - v_0$ is above e_i and every vertex of $P_k - v_0$ is below e_i, a contradiction since v_R is in both.

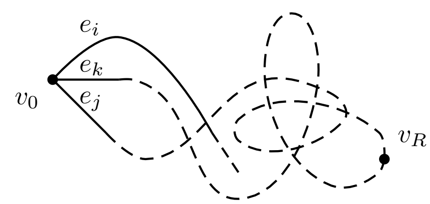

Figure 7: e_i crosses e_j oddly and e_k evenly.

Thus, we may assume that $x(v_0) < x(v_i) < x(v_R)$. As argued in the proof of Lemma 3.4, v_i is below P_j and above P_k, or v_i is above P_j and below P_k, depending on the order of values assigned to i, j, k.

Suppose that there is a path P'_i in $G - V(C)$ from v_i to a vertex v' with $x(v') < x(v_0)$ or $x(v') > x(v_R)$. Let P'_i be a minimal such path, so that it exits the region between lines $x = x(v_0)$, $x = x(v_R)$ with its last edge e'. P'_i is disjoint from P_j, so e' passes above (below) v_0 or v_R if and only if every vertex of P'_i is above (below) P_j. Likewise if we replace P_j by P_k. But then the vertices of P'_i are either all above P_j and P_k or they are all below P_j and P_k, including v_i, which contradicts what we have already shown for v_i. □

We are finally in a position to prove Theorem 3.1. We need one more piece of terminology: consider two edges e and f that share the same right (or left) endpoint. The *distance* between e and f is the number of edge ends between the ends of e,f in the left (or right) rotation at their shared vertex. (We do not measure distance within the entire rotation; only within the right or left rotation.)

Proof of Theorem 3.1. Let G be a smallest (fewest vertices) counterexample to the theorem. By Lemma 3.3(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 2.1 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 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 $x(v) \geq x(v_0)$.

Case 1. Vertices v_1, v_2, v_3 are in different components of $G_0 - v_0$.

In Case 1, for each $i \in \{1, 2, 3\}$, consider the component of $G_0 - v_0$ that contains v_i and let v_i' be its vertex with largest x-value. Assign i, j, k so that $\{i, j, k\} = \{1, 2, 3\}$, and $x(v_i')$ is smaller than $x(v_j')$ and $x(v_k')$. Let $x_R = x(v_i')$ and apply Lemma 3.4, which defines G_i', G_j', G_k' and shows that G_i' has no neighbors to the left of $x(v_0)$. Then by Lemma 3.3(iii) (with $a = v_0$ and $H = G_i'$), G_i' has just one vertex $v_i = v_i'$ $(= b)$ and $x(v_i) > x(v_j)$ and $x(v_i) > x(v_k)$. Then G_j' and G_k' are non-empty, so they also have no neighbors to the left of $x(v_0)$. This contradicts the second part of Lemma 3.3(iii) with H' equal to G_j' (or G_k') restricted to vertices v with $x(v) \leq x_R$.

If we are not in Case 1, let x_L be smallest such that the subgraph induced by $\{v \in V(G) : x(v_0) < x(v) \leq x_L\}$ has a component that contains at least two vertices of v_1, v_2, v_3. Then there is a cycle C that contains e_j and e_k for some distinct $k, j \in \{1, 2, 3\}$ and a vertex v_L such that $x(v_L) = x_L$ and $x(v_0) \leq x(v) \leq x(v_L)$ for all $v \in V(C)$. If $vv_L \in \{e_1, e_2, e_3\}$, then we may assume that C contains vv_L.

Let i be the unique index for which $\{i, j, k\} = \{1, 2, 3\}$. By the previous assumption, $v_i \neq v_L$. By Lemma 3.5, $x(v_i) < x(v_L)$ or $v_i \in V(C) - v_L$.

Suppose that there is a path Q from v_i to C so that $x(v_0) < x(v) < x(v_L)$ for all $v \in V(Q)$. Then $Q \cup e_i \cup C - v_L$ contains a cycle C' with e_i and either e_j or e_k. But every vertex v of C' satisfies $x(v_0) \leq x(v) < x(v_L)$ for all v in C', contradicting the choice of v_L.

It immediately follows that v_i is not in $V(C) - v_L$; also $v_i \neq v_L$, so we may let G_i' be the component of $G - V(C)$ that contains v_i. By Lemma 3.5, G_i' lies between $x = x(v_0)$ and $x = x(v_L)$ (since $v_i \neq v_L$). The previous paragraph also implies that G'_i has no neighbors in $V(C) - \{v_0, v_L\}$. Let v'_i be the vertex of G'_i with largest x-value, let $x_R = x(v'_i)$, and define G'_i, G'_j, G'_k according to Lemma 3.4 (and note that this doesn't alter G'_i).

Case 2. G'_i is not adjacent to v_L.

(Same as Case 1:) By Lemma 3.3(iii), G'_i has only the one vertex $v_i = v'_i$, and G'_j and G'_k are non-empty because $x(v_i)$ is greater than $x(v_j)$ and $x(v_k)$. Then we can apply Lemma 3.3(iii) with H' equal to G'_j (or G'_k) restricted to the vertices with the x-coordinate smaller than $x(v'_i)$, and we are done.

Case 3. There is an edge from G'_i to v_L.

Apply Lemma 3.3(ii) with $H = G'_i$. This completes the proof of the theorem. □

4 Level-Planarity Testing

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.[6] Unfortunately, the system has $|E| \cdot |V| = O(|V|^2)$ variables which leads to an impractical $O(|V|^6)$ running time.[7]

Similarly, Theorem 3.1 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. We present the details of this algorithm in Section 4.1. In Section 4.2 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 discuss the rather confusing situation of algorithms for level-planarity testing in more detail in Remark 4.3.

4.1 Testing x-Monotonicity

How can we use Theorem 3.1 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

[6] Tutte presented his theorem as an algebraic characterization of planarity, but he did not investigate algorithmic implications [32]. Algebraic planarity testing based on the Hanani-Tutte characterization was probably first described by Wu [34, 35] in a sequence of papers first published in Chinese in the 70s.

[7] There are linear-time algorithms for planarity testing based on a Hanani-Tutte-like characterization, but they do not take the algebraic route [4, 3].

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 3.1, we now have an efficient test.

Theorem 4.1. *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 Remark 2.2, 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. □

Remark 4.2. The $O(|V|^2)$ bound in Theorem 4.1 can be far from optimal since only (e, v)-moves for which v lies between the endpoints of e are possible. If we define the *layout complexity* of a graph with assigned x-coordinates as $|\{(e_1e_2, v) : x(e_1) < x(v) < x(e_2), v \in V(G), e_1e_2 \in E(G)\}|$, then the algorithm in Theorem 4.1 runs in linear time in the size of the layout. This measure seems fair if we actually want to draw the graph (since we have to know in what order edges pass a vertex).

4.2 Testing Level-Planarity

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 *leveling* $\ell : V \to \mathbb{Z}$. 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.

Our results can easily be extended to handle level-planarity testing, an important case of layered graph drawing [5, 14, 15, 20, 17].

Remark 4.3. Level-planar graphs can be recognized and embedded in linear time using PQ-trees [20, 17, 16]; this work is based on earlier work for the special case of hierarchies [5]. There had been an earlier attempt at extending this to general graphs [14, 15], but there were gaps in the algorithm as pointed out in [18]. Alternative routes have included identifying Kuratowski-style obstruction sets for level-planarity [13], characterizations via vertex-exchange graphs [12, 10] and reductions to 2-satisfiability [30]. 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 [7]. The testing [12] and layout [10] 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 [11]; this is unfortunate, since both algorithms are relatively fast, $O(|V|^2)$ for both testing and layout, and very simple (the testing algorithm is somewhat similar to ours, even if the characterization it is 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 testing 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.

Level-planar graphs do not directly generalize x-monotone graphs since an x-monotone graph can have vertices at non-integer levels. However, 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 4.1 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 4.4. *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 [12] (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.

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). As with Corollary 2.7, it is now easy to see that the x-monotonicity requirement is stronger than necessary.

Corollary 4.5. *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 leveled graph G' constructed before Corollary 4.4. Then G' has a leveled embedding in which every edge is x-bounded. By Corollary 2.7, 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. ☐

5 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), mon-ocr(G), and mon-iocr(G). With this notation we can restate the original result by Pach and Tóth, our Theorem 2.1, as saying that mon-ocr$(G) = 0$ implies mon-cr$(G) = 0$. Similarly, our Theorem 3.1 can be restated as mon-iocr$(G) = 0$ implies mon-cr$(G) = 0$.

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 $\mathrm{ocr}(G) \leq 3$ [28] and $\mathrm{iocr}(G) = \mathrm{cr}(G)$ for $\mathrm{iocr}(G) \leq 2$ [29]. Pach and Tóth showed that $\mathrm{cr}(G) \leq \binom{2\,\mathrm{ocr}(G)}{2}$ [24, 28]. 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 5.1. *For every n there is a graph G so that* mon-cr$(G) \geq n$ *and* mon-ocr$(G) = 1$.

In other words: even if there are only two edges crossing each other 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) in terms of mon-ocr(G).

Proof. Our example 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.

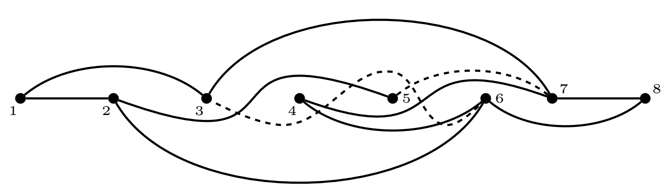

Figure 8: The drawing showing mon-ocr$(G) \leq 1$. The solid edges have weight n, the dashed edges have weight 1.

We next argue that mon-cr$(G) \geq n$. 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 8. We distinguish two cases depending on whether 46 passes below 5 (as in Figure 8) or above 5. Let us first consider the case that 46 passes below 5. Adding edges 25, 57, we see that they are forced to be drawn as in Figure 8. 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 9. 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.

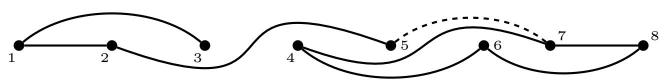

Figure 9: The unique way of drawing edges 25, 57, 68, and 27, assuming 46 passes below 5 and mon-cr$(D) < n$.

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 10.

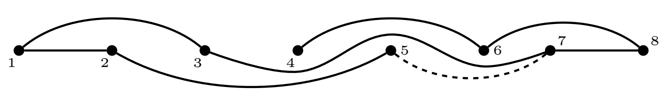

Figure 10: The unique way of drawing edges 25, 57, 68 and 37, assuming 46 passes above 5 and mon-cr$(D) < n$.

□

6 Future Directions

The following questions were first included in the conference version of this paper. Since that time there has already been some progress, which we include as annotations.

Planarity Testing Can the running time of our level-planarity testing algorithms be improved? There are several obstacles to this, most fundamentally the problem that to describe the effect of (e, v)-moves we need a system with a quadratic number of variables. It is not obvious (to us at least), how to reduce the size of this system. Other problems in the algorithm like the linear overhead in the "conquering" steps of the algorithms may be overcome with better data structures.

Monotone Crossing Numbers The *monotone crossing number* of a leveled graph G is the smallest number of crossings in any x-monotone drawing of the leveled graph. This problem is known to be **NP**-hard (even for two levels [8]) and the monotone crossing number can be arbitrarily large, even for a planar graph.[8]

We get a more interesting question if we define the *monotone crossing number for unleveled graphs* as the smallest crossing number of any x-monotone drawing for any leveling of the graph. Is this monotone crossing number bounded in the crossing number? For comparison, $\mathrm{rcr}_2(G)$ is at most $\binom{\mathrm{cr}(G)}{2}$, where $\mathrm{rcr}_2(G)$ allows straight-line edges with one bend [1].

Pach and Tóth [27] recently announced that the (unleveled) monotone crossing number of a graph G can be bounded by $2\,\mathrm{cr}(G)^2$ and that there are graphs for which the monotone crossing number is at least $7/6\,\mathrm{cr}(G) - 6$. We should also mention that Valtr [33] showed that the monotone crossing number is bounded by $4k^{4/3}$ where k is the monotone pair crossing number (again for unleveled graphs).

Bi-monotonicity 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?

- If a graph has both an x-monotone embedding and a y-monotone embedding, does it always have a bi-monotone embedding?
- If there is a drawing of a graph which is bi-monotone, is there a straight-line drawing with the same x and y ordering?
- What about bi-level-planarity?

As far as we know, bi-monotonicity and bi-level-planarity are new concepts, however, they are quite natural: If we specify the relative locations of objects on a map, we specify them in terms of "west/east of" and "north/south of" which is exactly what bi-monotonicity models. Imagine specifying the stations for a subway map: actual distances do not matter, what matters is relative location in terms of x and y.

As it turns out it is possible that a leveled graph has both an x-monotone and a y-monotone embedding without having a bi-monotone embedding, see the leveled graph in Figure 11. By Theorem 2.1 the graph is x-monotone and (applying the Theorem at an angle of 90 degrees) y-monotone. However, it can be shown that the graph is not bi-monotone.

This means that an algebraic bi-monotonicity criterion has to be more sophisticated than just requiring a bi-monotone even drawing. It also opens up the question of what is the complexity of recognizing bi-monotone or bi-level planar graphs?

Acknowledgment. We thank the anonymous referee for careful proofreading.

[8]The leveled graph 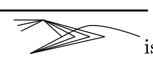 is such an example.

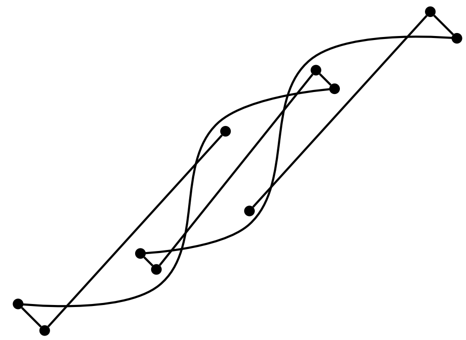

Figure 11: An leveled path that is not bi-monotone.

References

[1] Daniel Bienstock and Nathaniel Dean. Bounds for rectilinear crossing numbers. *J. Graph Theory*, 17(3):333–348, 1993.

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

[3] Hubert 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.

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

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

[6] Peter Eades, Qingwen Feng, Xuemin Lin, and Hiroshi Nagamochi. Straight-line drawing algorithms for hierarchical graphs and clustered graphs. *Algorithmica*, 44(1):1–32, January 2006.

[7] Alejandro Estrella-Balderrama, J. Joseph Fowler, and Stephen G. Kobourov. On the characterization of level planar trees by minimal patterns. In David Eppstein and Emden R. Gansner, editors, *Graph Drawing*, volume 5849 of *Lecture Notes in Computer Science*, pages 69–80. Springer, 2009.

[8] Michael R. Garey and David S. Johnson. Crossing number is NP-complete. *SIAM Journal on Algebraic and Discrete Methods*, 4(3):312–316, 1983.

[9] Chaim Chojnacki (Haim Hanani). Über wesentlich unplättbare Kurven im drei-dimensionalen Raume. *Fundamenta Mathematicae*, 23:135–142, 1934.

[10] Martin Harrigan and Patrick 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.

[11] Patrick Healy. Personal communication (January 2011).

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

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

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

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

[16] Michael Jünger and Sebastian Leipert. Level planar embedding in linear time. *J. Graph Algorithms Appl.*, 6:no. 1, 67–113 (electronic), 2002. Graph drawing and representations (Prague, 1999).

[17] Michael Jünger, Sebastian Leipert, and Petra 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.

[18] Michael Jünger, Sebastian Leipert, and Petra Mutzel. Pitfalls of using *PQ*-trees in automatic graph drawing. In Giuseppe 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.

[19] Daniel 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.

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

[21] Xuemin Lin and Peter Eades. Towards area requirements for drawing hierarchically planar graphs. *Theor. Comput. Sci.*, 292(3):679–695, 2003.

[22] Jiří Matoušek. *Using the Borsuk-Ulam theorem.* Universitext. Springer-Verlag, Berlin, 2003.

[23] Jiří Matousek, Martin Tancer, and Uli Wagner. Hardness of embedding simplicial complexes in $\mathbb{R}^d$. In Claire 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.

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

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

[26] János Pach and Géza Tóth. Monotone drawings of planar graphs. *ArXiv e-prints*, January 2011.

[27] János Pach and Géza Tóth. Monotone crossing number. In *Graph drawing 2011*, to appear.

[28] Michael J. Pelsmajer, Marcus Schaefer, and Daniel Štefankovič. Removing even crossings. *J. Combin. Theory Ser. B*, 97(4):489–500, 2007.

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

[30] Bert Randerath, Ewald Speckenmeyer, Endre Boros, Peter L. Hammer, Alexander Kogan, Kazuhisa Makino, Bruno Simeone, and Ondrej Cepek. A satisfiability formulation of problems on level graphs. *Electronic Notes in Discrete Mathematics*, 9:269–277, 2001.

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

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

[33] Pavel 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.

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

[35] Wen Jun Wu. On the planar imbedding of linear graphs (continued). *J. Systems Sci. Math. Sci.*, 6(1):23–35, 1986.