Monotone drawings of planar graphs

János Pach[*]
City College, CUNY and
Hungarian Academy of Sciences

Géza Tóth[†]
Rényi Institute of the
Hungarian Academy of Sciences

Abstract

Let G be a graph drawn in the plane so that its edges are represented by x-monotone curves, any pair of which cross an even number of times. We show that G can be redrawn in such a way that the x-coordinates of the vertices remain unchanged and the edges become non-crossing straight-line segments.

1 Introduction

A *drawing* $\mathcal{D}(G)$ of a graph G is a representation of the vertices and the edges of G by points and by possibly crossing simple Jordan arcs between them, resp. When it does not lead to confusion, we make no notational or terminological distinction between the vertices (resp. edges) of the underlying abstract graph and the points (resp. arcs) representing them. Throughout this paper, we assume that in a drawing

1. no edge passes through any vertex other than its endpoints;
2. no three edges cross at the same point;
3. if two edges of a drawing share an interior point p then they properly cross at p, i.e., one arc passes from one side of the other arc to the other side.

A drawing is called *x-monotone* if every vertical line intersects every edge in at most one point. For simplicity, we assume that no two vertices in an x-monotone drawing have the same x-coordinate. We call a drawing *even* if any two edges cross an even number of times.

Hanani (Chojnacki) [Ch34] (see also [T70]) proved the remarkable theorem that if a graph G permits an even drawing, then it is *planar*, i.e., it can be redrawn without any crossing. On the other hand, by Fáry's theorem [F48], [W36], every planar graph has a straight-line drawing. We can combine these two facts by saying that every even drawing can be *"stretched"*.

[*]Supported by NSF grant CR-00-98246, PSC-CUNY Research Award 63382-0032 and OTKA-T-032452.
[†]Supported by NSF grant OTKA-T-038397, OTKA-T-032452.

The aim of this note is to show that if we restrict our attention to *x-monotone* drawings, then every even drawing can be stretched without changing the x-coordinates of the vertices.

Consider an x-monotone drawing $\mathcal{D}(G)$ of a graph G. If the vertical ray starting at $v \in V(G)$ and pointing upward (resp. downward) crosses an edge $e \in E(G)$, then v is said to be *below* (resp. *above*) e. Two drawings of the same graph are called *equivalent*, if the above-below relationships between the vertices and the edges coincide.

In the next two sections we establish the following two results.

Theorem 1. *For any x-monotone even drawing of a connected graph, there is an equivalent x-monotone drawing in which no two edges cross each other and the x-coordinates of the corresponding vertices are the same.*

Theorem 2. *For any non-crossing x-monotone drawing of a graph G, there is an equivalent non-crossing straight-line drawing, in which the x-coordinates of the corresponding vertices are the same.*

2 Proof of Theorem 1

We follow the approach of Cairns and Nikolayevsky [CN00]. Consider an x-monotone drawing $\mathcal{D}$ of a graph on the xy-plane, in which any two edges cross an even number of times. Let u and v denote the leftmost and rightmost vertex, respectively. We can assume without loss of generality that $u = (-1, 0)$ and $v = (1, 0)$. Introduce two additional vertices, $w = (0, 1)$ and $z = (0, -1)$, each connected to u and v by arcs of length $\pi/2$ along the unit circle C centered at the origin, and suppose that every other edge of the drawing lies in the interior of C. Denote by G the underlying abstract graph, including the new vertices w and z.

For each crossing point p, attach a *handle* (or bridge) to the plane in a very small neighborhood $N(p)$ of p. Assume that (1) these neighborhoods are pairwise disjoint, (2) $N(p)$ is disjoint from every other edge that does not pass through p, and that (3) every vertical line intersects every handle only at most once. For every p, take the portion belonging to $N(p)$ of one of the edges that participate in the crossing at p, and lift it to the handle without changing the x- and y-coordinates of its points. The resulting drawing $\mathcal{D}_0$ is a crossing-free embedding of G on a surface S_0 of possibly higher genus.

Let S_1 be a very small closed neighborhood of the drawing $\mathcal{D}_0$ on the surface S_0. Note that S_1 is a compact, connected surface, whose boundary consists of a finite number of closed curves. Attaching a disc to each of these closed curves, we obtain a surface S_2 with no boundary. According to Cairns and Nikolayevsky [CN00], S_2 must be a 2-dimensional *sphere*. To verify this claim, consider two closed curves, α_2 and β_2, on S_2. They can be deformed into closed walks, α_1 and β_1, respectively, along the edges of $\mathcal{D}_0$. The projection of these two walks into the (x, y)-plane are closed curves, α and β, in $\mathcal{D}$, which must cross an even number of times. Every crossing between α and β occurs either at a vertex of $\mathcal{D}$ or between two of its edges. By the assumptions, any two edges in $\mathcal{D}$ cross an even number of times. (The same assertion is trivially true in $\mathcal{D}_0 \subset S_2$, because there no two edges cross.) Using the fact that in $\mathcal{D}_0 \subset S_2$ the cyclic order of the edges incident to a vertex is the same as the cyclic order of the corresponding edges in $\mathcal{D}$, we can conclude that α_1 and β_1 cross an even number of times, and the same is true for α_2 and β_2. Thus, S_2 is a surface with no boundary, in which any two closed curves cross an even number of times. This implies that S_2 is a sphere. Consequently, $\mathcal{D}_0$, a crossing-free drawing of G on S_2, corresponds to a plane drawing.

Next, we argue that $\mathcal{D}_0$ can also be regarded as an *x-monotone* plane drawing of G, in which the x-coordinates of the vertices are the same as the x-coordinates of the corresponding vertices in $\mathcal{D}$.

For any point q (either in the plane or in 3-space), let $x(q)$ denote the x-coordinate of q. As before, every boundary curve of S_1 corresponds to a cycle of G. Since in the original drawing the cycle $vwuz$ encloses all other edges and vertices of G, one of the boundary curves of S_1, say γ, corresponds to the cycle $vwuz$. Consider another boundary curve, $\kappa \neq \gamma$, which corresponds to a cycle $v_1v_2\ldots v_i$ in G. Let D_κ be a closed disc in the plane bounded by an equivalent non-crossing x-monotone drawing of the cycle $v_1v_2\ldots v_i$. Glue the boundary of D_κ to κ so that a boundary point b of D_κ will be glued to a point $k \in \kappa$ if and only if b and k correspond to the same point of the cycle $v_1v_2\ldots v_i$ in $\mathcal{D}$. Consequently, $x(b) = x(k)$. Repeat the same procedure for each $\kappa \neq \gamma$.

Finally, we obtain a new surface S containing $\mathcal{D}_0$, which is topologically isomorphic to the unit disk bounded by C, and a natural extension of the x-coordinate function from S_1 to S, which is a continuous real function with no local minimum or maximum. Therefore, $\mathcal{D}_0$ can be regarded as a crossing-free x-monotone drawing of G, equivalent to $\mathcal{D}$. This completes the proof of Theorem 1.

Remark. Theorem 1 cannot be extended to disconnected graphs. To see this, consider a pair of edges, e_1 and e_2, intersecting twice, and place a vertex below e_1 and above e_2, and another one above e_1 and below e_2. Clearly, there exists no equivalent crossing-free x-monotone drawing. On the other hand, if we drop the condition that the new drawing must be equivalent to the original one, then the connected components can be treated separately and their drawings can be shifted in the vertical direction so as to avoid any crossing between them.

3 Proof of Theorem 2

Let $\mathcal{D} = \mathcal{D}(G)$ be a non-crossing x-monotone drawing of a graph G. First, we show that it is sufficient to prove Theorem 2 for *triangulated* graphs. Deleting all vertices (points) and edges (arcs) of $\mathcal{D}$ from the plane, the plane falls into connected components, called *faces*. The x-coordinate of any vertex v will be denoted by $x(v)$.

Lemma 3.1. *By the addition of further edges and an extra vertex, if necessary, every non-crossing x-monotone drawing $\mathcal{D}$ can be extended to an x-monotone triangulation.*

Proof. Consider a face F, and assume that it has more than 3 vertices. It is sufficient to show that one can always add an x-monotone edge between two non-adjacent vertices of F, which does not cross any previously drawn edges.

For the sake of simplicity, we outline the argument only for the case when F is a bounded face. The proof in the other case is very similar, the only difference is that we may also have to add an extra vertex.

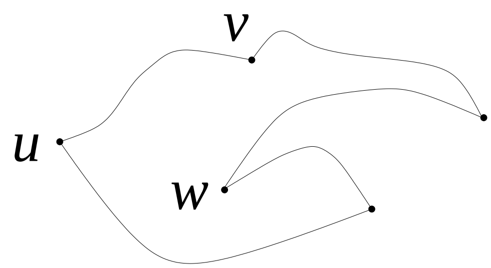

Figure 1. *The vertex w is extreme, u and v are not.*

A vertex w of F is called *extreme* if it is not the left endpoint of any edge or not the right endpoint of any edge in $\mathcal{D}$, and a small neighborhood of w on the vertical line through w belongs to F. In particular, if the boundary of F is not connected, the leftmost (and the rightmost) vertex of each component of the boundary other than the exterior component, is extreme. See Fig. 1.

Suppose first that F has an extreme vertex w. We may assume, by symmetry, that w is not the right endpoint of any edge in $\mathcal{D}$. Starting at w, draw a horizontal ray in the direction of the negative x-axis. Let p be the first intersection point of this ray with the boundary of F. If p is a vertex, then the segment wp can be added to $\mathcal{D}$. Otherwise, one can add an x-monotone edge joining w to the left endpoint of the edge that p belongs to.

Suppose next that none of the vertices of F are extreme. In this case, the boundary of F is connected and any two vertices of F can be joined by an x-monotone curve inside F. However, an edge can be added to $\mathcal{D}$ only if the corresponding two vertices do not induce an edge in the exterior of F. Clearly, letting v_1, v_2, v_3, and v_4 denote four consecutive vertices of F, at least one of the pairs (v_1, v_3) and (v_2, v_4) has this property. $\square$

Now we turn to the proof of Theorem 2. The proof is by induction on the number of vertices. If G has at most 4 vertices, the assertion is trivial. Suppose that G has $n > 4$ vertices and that we have already established the theorem for graphs having fewer than n vertices. By Lemma 3.1, we can assume without loss of generality that the original x-monotone drawing $\mathcal{D}$ of G is triangulated.

CASE 1. There is a triangle $T = v_1v_2v_3$ in $\mathcal{D}$, which is not a face.

Then there is at least one vertex of $\mathcal{D}$ in the interior and at least one vertex in the exterior of T. Consequently, the drawings $\mathcal{D}_{\text{in}}$ and $\mathcal{D}_{\text{out}}$ defined as the part of $\mathcal{D}$ induced by v_1, v_2, v_3, and all vertices *inside* T and *outside* T, resp., have fewer than n vertices. By the induction hypothesis, there exist straight-line drawings $\mathcal{D}'_{\text{in}}$ and $\mathcal{D}'_{\text{out}}$, equivalent to $\mathcal{D}_{\text{in}}$ and $\mathcal{D}_{\text{out}}$, resp., in which all vertices have the same x-coordinates as in the original drawing. Notice that there is an affine transformation A of the plane, of the form

$$A(x, y) = (x, ax + by + c),$$

which takes the triangle induced by v_1, v_2, v_3 in $\mathcal{D}_{\text{in}}$ into the triangle induced by v_1, v_2, v_3 in $\mathcal{D}_{\text{out}}$. Since the image of a drawing under any affine transformation is equivalent to the original drawing, we conclude that $A\left(\mathcal{D}'_{\text{in}}\right) \cup \mathcal{D}'_{\text{out}}$ meets the requirements.

In the sequel, we can assume that $\mathcal{D}$ has no triangle that is not a face. Fix a vertex v of $\mathcal{D}$ with minimum degree. Since every triangulation on $n > 4$ vertices has $3n - 6$ edges, the degree of v is $3, 4$, or 5. If the degree of v is 3, the neighbors of v induce a triangle in $\mathcal{D}$, which is not a face, contradicting our assumption.

There are two more cases to consider.

Case 2. The degree of v is 4.

Let v_1, v_2, v_3, v_4 denote the neighbors of v, in clockwise order. There are three substantially different subcases, up to symmetry. See Fig. 2.

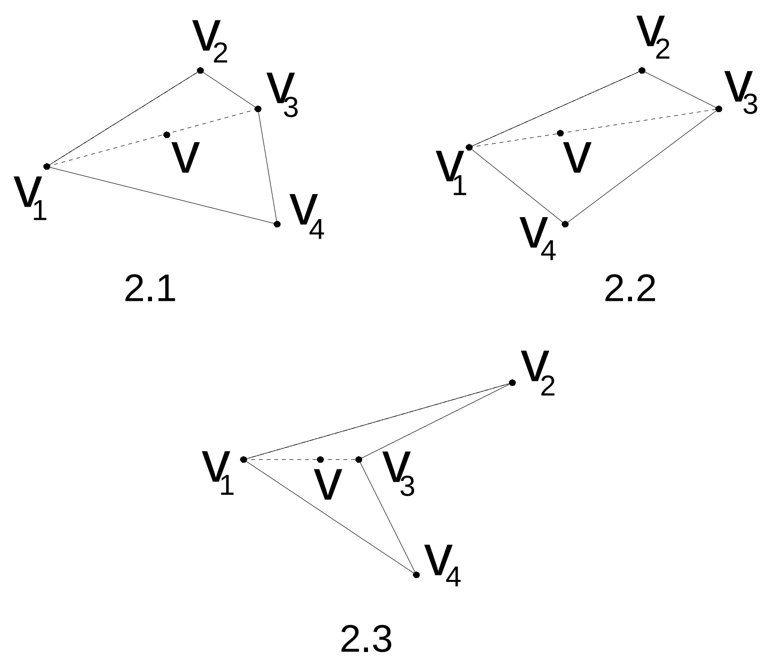

Figure 2. Case 2.

Subcase 2.1: $x(v_1) < x(v_2) < x(v_3) < x(v_4)$

Clearly, at least one of the inequalities $x(v) > x(v_2)$ and $x(v) < x(v_3)$ is true. Suppose without loss of generality that $x(v) < x(v_3)$. If v_1 and v_3 were connected by an edge, then vv_1v_3 would be a triangle with v_2 and v_4 in its interior and in its exterior, resp., contradicting our assumption. Remove v from $\mathcal{D}$, and add an x-monotone edge between v_1 and v_3, running in the interior of the face that contains v. Applying the induction hypothesis to the resulting drawing, we obtain that it can be redrawn by straight-line edges, keeping the x-coordinates fixed. Subdivide the segment v_1v_3 by its (uniquely determined) point whose

x-coordinate is $x(v)$. In this drawing, v can also be connected by straight-line segments to v_2 and to v_4. Thus, we obtain an equivalent drawing which meets the requirements.

SUBCASE 2.2: $x(v_1) < x(v_2) < x(v_3) > x(v_4) > x(v_1)$

SUBCASE 2.3: $x(v_1) < x(v_2) > x(v_3) < x(v_4) > x(v_1)$

In these two subcases, the above argument can be repeated *verbatim*. In Subcase 2.3, to see that $x(v_1) < x(v) < x(v_3)$, we have to use the fact that in $\mathcal{D}$ both vv_2 and vv_4 are represented by x-monotone curves.

CASE 3. The degree of v is 5.

Let v_1, v_2, v_3, v_4, v_5 be the neighbors of v, in clockwise order. There are four substantially different cases, up to symmetry. See Fig. 3.

SUBCASE 3.1: $x(v_1) < x(v_2) < x(v_3) < x(v_4) < x(v_5)$

SUBCASE 3.2: $x(v_1) < x(v_2) < x(v_3) < x(v_4) > x(v_5) > x(v_1)$

SUBCASE 3.3: $x(v_1) < x(v_2) < x(v_3) > x(v_4) < x(v_5) > x(v_1)$

SUBCASE 3.4: $x(v_1) < x(v_2) > x(v_3) > x(v_4) < x(v_5) > x(v_1)$

In all of the above subcases, we can assume, by symmetry or by x-monotonicity, that $x(v) < x(v_4)$. Since $\mathcal{D}$ has no triangle which is not a face, we obtain that v_1v_3, v_1v_4, and v_2v_4 cannot be edges. Delete from $\mathcal{D}$ the vertex v together with the five edges incident to v, and let $\mathcal{D}_0$ denote the resulting drawing. Furthermore, let $\mathcal{D}_1$ (and $\mathcal{D}_2$) denote the drawing obtained from $\mathcal{D}_0$ by adding two non-crossing x-monotone diagonals, v_1v_3 and v_1v_4 (resp. v_2v_4 and v_1v_4), which run in the interior of the face containing v. By the induction hypothesis, there exist straight-line drawings $\mathcal{D}'_1$ and $\mathcal{D}'_2$ equivalent to $\mathcal{D}_1$ and $\mathcal{D}_2$, resp., in which the x-coordinates of the corresponding vertices are the same.

Apart from the edges v_1v_3, v_1v_4, and v_2v_4, $\mathcal{D}'_1$ and $\mathcal{D}'_2$ are non-crossing straight-line drawings equivalent to $\mathcal{D}_0$ such that the x-coordinates of the corresponding vertices are the same. Obviously, the convex combination of two such drawings is another non-crossing straight-line drawing equivalent to $\mathcal{D}_0$. More precisely, for any $0 \le \alpha \le 1$, let $\mathcal{D}'_\alpha$ be defined as

$$\mathcal{D}'_\alpha = \alpha\mathcal{D}'_1 + (1-\alpha)\mathcal{D}'_2.$$

That is, in $\mathcal{D}'_\alpha$, the x-coordinate of any vertex $u \in V(G) - v$ is equal to $x(u)$, and its y-coordinate is the combination of the corresponding y-coordinates in $\mathcal{D}'_1$ and $\mathcal{D}'_2$ with coefficients α and $1-\alpha$, resp.

Observe that the only possible concave angle of the quadrilateral $Q = v_1v_2v_3v_4$ in $\mathcal{D}'_1$ and $\mathcal{D}'_2$ is at v_3 and at v_2, resp. In $\mathcal{D}'_\alpha$, Q has at most one concave vertex. Since the shape of Q changes continuously with α, we obtain that there is a value of α for which Q is a *convex* quadrilateral in $\mathcal{D}_\alpha$. Let $\mathcal{D}'$ be the straight-line drawing obtained from $\mathcal{D}'_\alpha$ by adding v at the unique point of the segment v_1v_4, whose x-coordinate is $x(v)$, and connect it to $v_1, \ldots, v_5$. Clearly, $\mathcal{D}'$ meets the requirements of Theorem 2.

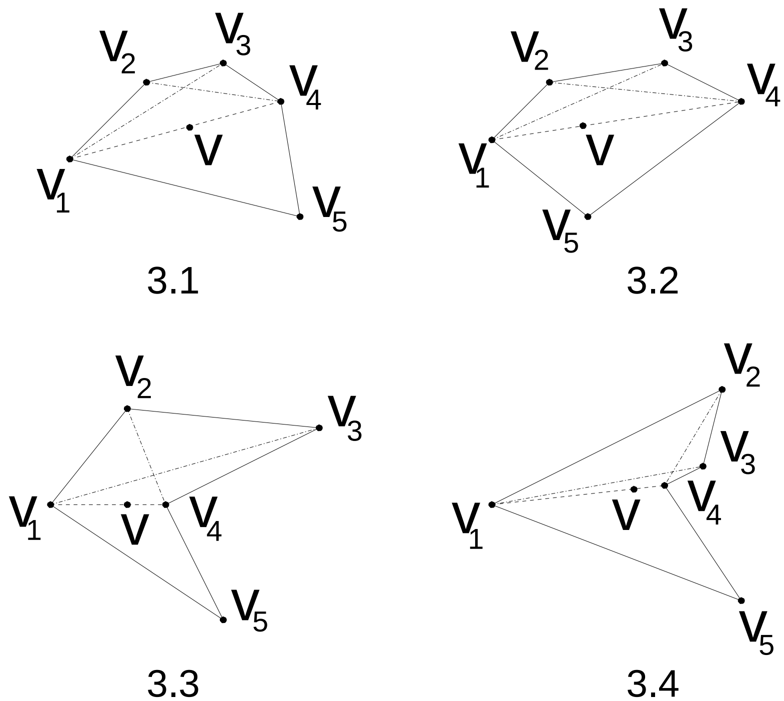

Figure 3. CASE 3.

Added in proof: We are grateful to Professor P. Eades for calling our attention to his paper [EFL97], sketching a somewhat more complicated proof for a result essentially equivalent to our Theorem 2.

References

[Ch34] Ch. Chojnacki (A. Hanani), Über wesentlich unplättbare Kurven im dreidimensionalen Raume, *Fund. Math.* **23** (1934), 135–142.

[CN00] G. Cairns and Y. Nikolayevsky, Bounds for generalized thrackles, *Discrete Comput. Geom.* **23** (2000), 191–206.

[DETT99] G. Di Battista, P. Eades, R. Tamassia, and I. G. Tollis, *Graph Drawing,* Prentice Hall, Upper Saddle River, NJ, 1999.

[EFL97] P. Eades, Q.-W. Feng, and X. Lin, Straight-Line drawing algorithms for hierarchical graphs and clustered graphs, in: *Graph Drawing (s. North, ed.) Lecture Notes in Computer Science* **1190,** 113-128.

[EET76] G. Ehrlich, S. Even, and R. E. Tarjan, Intersection graphs of curves in the plane, *Journal of Combinatorial Theory, Series B* **21** (1976), 8–20.

[F48] I. Fáry, On straight line representation of planar graphs, *Acta Univ. Szeged. Sect. Sci. Math.* **11** (1948), 229–233.

[T70] W. T. Tutte, Toward a theory of crossing numbers, *J. Combinatorial Theory* **8** (1970), 45–53.

[W36] K. Wagner, Bemerkungen zum Vierfarbenproblem, *Jber. Deutsch. math. Vereinigung* **46** (1936), 26–32.