

Untangling a Planar Graph

Xavier Goaoc · Jan Kratochvíl · Yoshio Okamoto ·
Chan-Su Shin · Andreas Spillner · Alexander Wolff

Received: 26 November 2007 / Revised: 7 November 2008 / Accepted: 20 November 2008 /
Published online: 9 January 2009
© The Author(s) 2009. This article is published with open access at Springerlink.com

Abstract A straight-line drawing δ of a planar graph G need not be plane but can be made so by *untangling* it, that is, by moving some of the vertices of G . Let $\text{shift}(G, \delta)$ denote the minimum number of vertices that need to be moved to untangle δ . We show that $\text{shift}(G, \delta)$ is NP-hard to compute and to approximate. Our

This paper is based on two preliminary versions: [7] and [24]. Our work was started at the 9th “Korean” Workshop on Computational Geometry and Geometric Networks organized by A. Wolff and X. Goaoc, July 30–August 4, 2006 in Schloss Dagstuhl, Germany. Further contributions were made at the 2nd Workshop on Graph Drawing and Computational Geometry organized by W. Didimo and G. Liotta, March 11–16, 2007 in Bertinoro, Italy.

J. K. was supported by research project 1M0545 of the Czech Ministry of Education.

Y. O. was supported by GCOE Program “Computationism as a Foundation for the Sciences” and Grant-in-Aid for Scientific Research from Ministry of Education, Science and Culture, Japan, and Japan Society for the Promotion of Science.

C.-S. S. was supported by Research Grant 2008 of the Hankuk University of Foreign Studies.

X. Goaoc

LORIA–INRIA Grand Est, Nancy, France
e-mail: goaoc@loria.fr

J. Kratochvíl

Department of Applied Mathematics and Institute of Theoretical Computer Science,
Charles University, Prague, Czech Republic
e-mail: honza@kam.mff.cuni.cz

Y. Okamoto

Graduate School of Information Science and Engineering, Tokyo Institute of Technology, Tokyo,
Japan
e-mail: okamoto@is.titech.ac.jp

C.-S. Shin

School of Electronics and Information Engineering, Hankuk University of Foreign Studies, Yongin,
Korea
e-mail: cssin@hufs.ac.kr

hardness results extend to a version of 1BENDPOINTSEMBEDDABILITY, a well-known graph-drawing problem.

Further we define $\text{fix}(G, \delta) = n - \text{shift}(G, \delta)$ to be the maximum number of vertices of a planar n -vertex graph G that can be fixed when untangling δ . We give an algorithm that fixes at least $\sqrt{((\log n) - 1)/\log \log n}$ vertices when untangling a drawing of an n -vertex graph G . If G is outerplanar, the same algorithm fixes at least $\sqrt{n/2}$ vertices. On the other hand, we construct, for arbitrarily large n , an n -vertex planar graph G and a drawing δ_G of G with $\text{fix}(G, \delta_G) \leq \sqrt{n-2} + 1$ and an n -vertex outerplanar graph H and a drawing δ_H of H with $\text{fix}(H, \delta_H) \leq 2\sqrt{n-1} + 1$. Thus our algorithm is asymptotically worst-case optimal for outerplanar graphs.

Keywords Graph drawing · Straight-line drawing · Planarity · NP-hardness · Hardness of approximation · Moving vertices · Untangling · Point-set embeddability

1 Introduction

A *drawing* of a graph G maps each vertex of G to a distinct point of the plane and each edge uv to an open Jordan curve connecting the images of u and v . A drawing of G is *plane* if no two distinct edges *cross*, that is, intersect. By the famous theorem of Wagner [28], Fáry [6], and Stein [25], any planar graph admits a plane *straight-line* drawing, that is, a drawing that maps edges to straight-line segments. Obviously not every straight-line drawing of a planar graph is plane. In this paper we are exclusively interested in such straight-line drawings. Thus by a drawing we will always mean a straight-line drawing. Since a (straight-line) drawing is completely defined by the position of the vertices, moving a vertex is a natural operation to modify such a drawing. If a drawing is to be made plane—or *untangled*—by successively moving vertices, it is desirable to move as few vertices as possible. The smaller the number of moves, the less likely it is that an observer gets confused, that is, the more likely the observer’s *mental map* [17] is preserved during a sequence of changes. A recreational version of the problem of minimizing the number of moves is given by Tantalo’s popular on-line game *Planarity* [26], where the aim is to untangle a straight-line drawing as quickly as possible, again by vertex moves. Actually, in Tantalo’s game an additional difficulty for the player is the fixed size of the screen; Liske’s [15] version of the game allows rescaling and hence is fully equivalent to untangling.

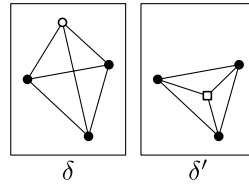
We define the *vertex-shifting distance* d between two drawings δ and δ' of a graph $G = (V, E)$ to be the number of vertices of G whose images under δ and δ' differ:

$$d(\delta, \delta') = |\{v \in V \mid \delta(v) \neq \delta'(v)\}|.$$

A. Spillner
School of Computing Sciences, University of East Anglia, Norwich, UK
e-mail: aspillner@cmp.uea.ac.uk

A. Wolff
Faculteit Wiskunde en Informatica, Technische Universiteit Eindhoven, Eindhoven, The Netherlands
url: <http://www.win.tue.nl/~awolff>

Fig. 1 Two drawings of K_4 :
 δ is not plane, δ' is plane;
 $d(\delta, \delta') = 1$



Given our edit operation, d represents the edit distance for straight-line drawings of graphs (see Fig. 1 for an example). For a drawing δ of a planar graph G , we denote by $\text{shift}(G, \delta)$ the minimum number of vertices that need to be moved in order to untangle δ . In some sense $\text{shift}(G, \delta)$ measures the distance of δ from planarity. This suggests the following computational problem.

MINSHIFTEDVERTICES(G, δ): given a drawing δ of a planar graph G , find a plane drawing δ' of G with $d(\delta, \delta') = \text{shift}(G, \delta)$.

The symmetric point of view is often helpful. Therefore we denote by $\text{fix}(G, \delta)$ the maximum number of vertices that can be fixed when untangling δ ; we refer to such vertices as *fixed* vertices. Clearly, $\text{fix}(G, \delta) = n - \text{shift}(G, \delta)$, where n is the number of vertices of G . We call the corresponding problem, that is, finding a plane drawing of a given planar graph G that maximizes the number of fixed vertices with a given drawing δ , **MAXFIXEDVERTICES**. We denote by $\text{fix}(G)$ the minimum of $\text{fix}(G, \delta)$ over all drawings δ of G . Analogously, we denote by $\text{shift}(G)$ the maximum of $\text{shift}(G, \delta)$ over all drawings δ of G .

Kaufmann and Wiese [11] considered the graph-drawing problem **1BEND-POINTSEMBEDDABILITY** that will turn out to be related to **MINSHIFTEDVERTICES**. They defined a planar graph $G = (V, E)$ to be *k-bend embeddable* if, for any set S of $|V|$ points in the plane, there is a one-to-one correspondence between V and S that can be extended to a plane drawing of G with at most k bends per edge. Kaufmann and Wiese showed that (a) every 4-connected planar graph is 1-bend embeddable, (b) every planar graph is 2-bend embeddable, and (c) given a planar graph $G = (V, E)$ and set S of $|V|$ points on a line, it is NP-complete to decide whether there is a correspondence between V and S that makes it possible to 1-bend embed G on S .

The contributions we present in this paper are three-fold:

- We prove that the decision versions of **MAXFIXEDVERTICES** and **MINSHIFTEDVERTICES** are NP-hard (Theorem 3.1) and lie in \mathcal{PSPACE} (Proposition 3.7). We further prove that **MINSHIFTEDVERTICES** is hard to approximate in the following sense: if there is a real $\varepsilon \in (0, 1]$ and a polynomial-time algorithm that guarantees to untangle any drawing δ of any n -vertex planar graph G with at most $(n^{1-\varepsilon}) \cdot (\text{shift}(G, \delta) + 1)$ moves, then $\mathcal{P} = \mathcal{NP}$ (Theorem 3.3).
- We complement the complexity result of Kaufmann and Wiese [11] on **1BEND-POINTSEMBEDDABILITY** by showing that it is NP-hard to decide whether a *given* one-to-one correspondence between the vertices of a planar graph G and a planar point set S extends into a plane drawing of G with at most one bend per edge (Theorem 3.4). We also show that the problem lies in \mathcal{PSPACE} (Theorem 3.6) and that an optimization version of the problem is hard to approximate (Corollary 3.5).

- We show that $\text{fix}(H) \geq \sqrt{n/2}$ for any n -vertex outerplanar graph H (Corollary 4.9) and that $\text{fix}(G) \geq \sqrt{\frac{(\log n)-1}{\log \log n}}$ for any general planar graph G with $n \geq 4$ vertices (Theorem 4.11), where the base of logarithms is 2. We also give, for arbitrarily large n , examples of an n -vertex outerplanar graph H with $\text{fix}(H) \leq 2\sqrt{n-1} + 1$ (Theorem 6.1) and of an n -vertex planar graph G with $\text{fix}(G) \leq \sqrt{n-2} + 1$ (Theorem 5.3). We prove the two bounds by using drawings where all vertices lie on a line.

2 Previous and Related Work

Arguably, one of the earliest results on untangling, for the n -path in the real line, is the Erdős–Szekeres theorem, which we state here for further reference.

Theorem 2.1 (Erdős and Szekeres [5]) *Any sequence of $n \geq sr + 1$ different real numbers has an increasing subsequence of $s + 1$ terms or a decreasing subsequence of $r + 1$ terms.*

The current best bounds on $\text{fix}(G)$, where G is restricted to certain classes of planar graphs, are summarized in Table 1. Recall that a lower bound of $f(n)$ means that we can untangle any drawing of any n -vertex graph G in the given graph class while fixing at least $f(n)$ vertices, whereas an upper bound of $g(n)$ means that for arbitrarily large n , there exists a drawing δ of an n -vertex graph G in the given graph class such that at most $g(n)$ vertices can stay fixed when untangling δ .

Untangling was first investigated for the n -cycle C_n , following the question by Watanabe [29] of whether $\text{fix}(C_n) \in \Omega(n)$. The answer turned out to be negative: Pach and Tardos [18] showed, by a probabilistic argument, that $\text{fix}(C_n) \in O((n \log n)^{2/3})$. They also showed that $\text{fix}(C_n) \geq \lfloor \sqrt{n} + 1 \rfloor$ by applying the Erdős–Szekeres theorem to the sequence of the indices of the vertices of δ in clockwise order around some specific point. Cibulka [4] recently improved that lower bound to $\Omega(n^{2/3})$ by applying the Erdős–Szekeres theorem not once but $\Theta(n^{1/3})$ times.

Pach and Tardos [18] extended the question to planar graphs and asked whether there is a constant $\gamma > 0$ such that $\text{fix}(G) \in \Omega(n^\gamma)$ for any planar n -vertex graph G . This question was recently answered in the affirmative by Bose et al. [2], who showed that $\text{fix}(G) \geq \sqrt[4]{n/3}$. While their bound improves on our Theorem 4.11, their algorithm uses our algorithm as a subroutine (specifically the result in Corollary 4.9). A recent improvement in our analysis also improves their bound, yielding

Table 1 Best known bounds for $\text{fix}(G)$, where G is a graph of the given graph class with n vertices

Graph class	Lower bound		Upper bound	
Cycles	$\Omega(n^{2/3})$	[4]	$O((n \log n)^{2/3})$	[18]
Trees	$\sqrt{n/2}$	[2]	$3\sqrt{n} - 3$	[2]
Outerplanar graphs	$\sqrt{n/2}$	[20] and Corollary 4.9	$2\sqrt{n-1} + 1$	Theorem 6.1
Planar graphs	$\sqrt[4]{(n+1)/2}$	[2]	$\sqrt{n-2} + 1$	Theorem 5.3

$\text{fix}(G) \geq \sqrt[4]{(n+1)/2}$. Kang et al. [10] showed that for arbitrarily large n there is a planar graph G with n vertices and $\text{fix}(G) \leq 2\sqrt{n} + 1$. For our upper bound of $\sqrt{n-2} + 1$, see Theorem 5.3. Kang et al. [10] also shed some light on how upper bounds on $\text{fix}(G)$ are affected by restricting the possible locations of vertices in the drawings of G . In particular, they showed that initial drawings with all vertices on a line, such as our examples in Theorems 6.1 and 5.3, are the worst case in the sense that any planar graph G has such a drawing δ with $\text{fix}(G) = \text{fix}(G, \delta)$ and that their upper bound holds even in the case where initial drawings are restricted to drawings where vertices correspond to a set of points on the boundary of a convex set. (Note that this generalizes both the vertices-on-a-line case and the vertices-in-convex-position case.)

Verbitsky [27] investigated planar graphs of higher connectivity. He proved linear upper bounds on $\text{fix}(G)$ for three- and four-connected planar graphs. Cibulka [4] gave, for any planar graph G , an upper bound on $\text{fix}(G)$ that is a function of the number of vertices, the maximum degree, and the diameter of G . This latter bound implies, in particular, that $\text{fix}(G) \in O((n \log n)^{2/3})$ for any three-connected planar graph G and that any graph H such that $\text{fix}(H) \geq cn$ for some $c > 0$ must have a vertex of degree $\Omega(nc^2 / \log^2 n)$.

For the class of trees, Bose et al. [2] showed that $\text{fix}(T) \geq \sqrt{n/2}$ for any tree T with n vertices. They further showed that $\text{fix}(T) \leq 3\sqrt{n} - 3$ for a collection of stars with n vertices in total, which, up to adding one vertex to turn these stars into a single tree, implies that the previous bound is asymptotically tight. We have obtained the same lower bound of $\sqrt{n/2}$ for the larger class of outerplanar graphs (Corollary 4.9). This bound was obtained independently by Ravsky and Verbitsky [20] via a finer analysis of sets of collinear vertices in plane drawings.

The hardness of computing $\text{fix}(G, \delta)$ given G and δ was obtained independently by Verbitsky [27] by a reduction from independent set in line-segment intersection graphs. While our proof is more complicated than his, it is stronger as it also yields hardness of approximation and extends to the problem 1BENDPOINTSEMBEDDABILITY with given vertex–point correspondence.

Finally, a somewhat related problem is that of *morphing*, or isotopy, between two plane drawings δ_1 and δ_2 of the same graph G , that is, to define for each vertex v of G a movement from $\delta_1(v)$ to $\delta_2(v)$ such that at any time during the move the drawing defined by the current vertex positions is plane. We refer the interested reader to the survey by Lubiw et al. [16].

3 Complexity

In this section, we investigate the complexity of MINSHIFTEDVERTICES and of 1BENDPOINTSEMBEDDABILITY with given vertex–point correspondence.

Theorem 3.1 *Given a planar graph G , a drawing δ of G , and an integer $K > 0$, it is NP-hard to decide whether $\text{shift}(G, \delta) \leq K$.*

Proof Our proof is by reduction from PLANAR3SAT, which is NP-hard [14]. An instance of PLANAR3SAT is a 3-SAT formula φ whose variable–clause graph is planar. Note that this graph can be laid out (in polynomial time) such that variables

Fig. 2 Embedding of a planar 3-SAT formula

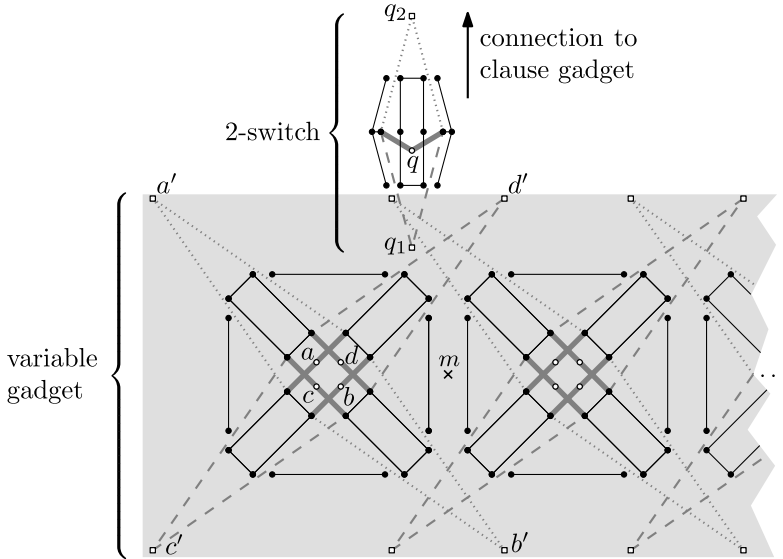
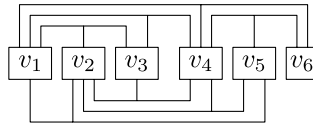


Fig. 3 Edges in the variable gadget: immobile (*solid black*) and mobile (*thick solid gray*). The predestined positions of mobile edges either correspond to true (*dashed gray*) or to false (*dotted gray*)

correspond to rectangles centered on the x -axis and clauses correspond to noncrossing three-legged “combs” completely above or completely below the x -axis [12], see Fig. 2. We refer to this layout of the variable-clause graph as λ_φ . We now construct a graph G_φ with a straight-line drawing δ_φ such that the following holds: δ_φ can be untangled by moving at most K vertices if and only if φ is satisfiable. We fix K later.

Our graph G_φ consists of two types of substructures (or *gadgets*), modeling the variables and clauses of φ . The overall layout of G_φ follows λ_φ (see Fig. 2): the variable gadgets are drawn in the same order along the x -axis as the variable nodes in λ_φ , and the clause gadgets form noncrossing three-legged combs that lie on the same side of the x -axis as the corresponding clause nodes in λ_φ .

In our gadgets, see Figs. 3 and 4, there are two types of vertices and edges; those that *may* move and those that are *meant* not to move. We refer to the two types as *mobile* and *immobile*. Each mobile vertex (but no immobile vertex) is incident to two edges that cross two other edges. The drawing δ_φ that we specify in the following has $2K$ crossings; if φ has a satisfying truth assignment, δ_φ can be untangled by moving K mobile vertices. Otherwise, at least one immobile vertex must move, and thus in total at least $K + 1$ vertices need to move. In the figures, immobile vertices are marked by black disks, mobile vertices by circles, and their predestined positions by

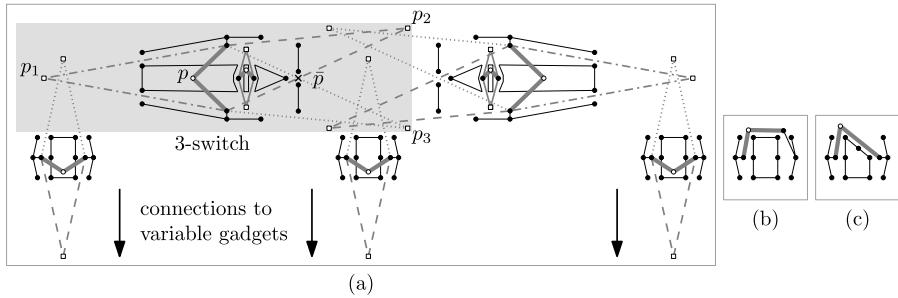


Fig. 4 (a) A clause gadget consists of three big 2-switches (drawn vertically) and two 3-switches (drawn horizontally; one is shaded). Each 3-switch contains another small 2-switch. Note that not all immobile vertices are marked. (b) and (c) Two ways in which originally immobile vertices can move to avoid a crossing if φ is not satisfiable

little squares. Mobile edges—edges incident to a mobile vertex—are drawn as thick solid gray line segments, and their predestined positions as gray line segments that are dashed, dotted, or dashed-dotted (and thus not solid). Immobile edges are drawn as solid black line segments.

Now consider the gadget for some variable x in φ , see the shaded area in Fig. 3. The gadget consists of a horizontal chain of a certain number of roughly square blocks. Each block consists of 28 vertices (four of which are mobile) and 28 edges. In Fig. 3 the four mobile vertices of the leftmost block are labeled in clockwise order a, d, b , and c . Note that the gray edges incident to a and b intersect those incident to c and d . Thus either both a and b or both c and d must be moved to untangle the block. Each mobile vertex $w \in \{a, b, c, d\}$ can move into exactly one position w' (up to small perturbations). The resulting incident edges are drawn by dotted and dashed gray line segments, respectively. Note that neighboring blocks in the chain are placed such that the only way to untangle them simultaneously is to move corresponding pairs of vertices and edges. Thus either all blocks of a variable gadget use the dashed line segments or all use the dotted line segments. These two ways to untangle a variable gadget correspond to the values *true* and *false* of the variable, respectively.

Let C be the numbers of clauses of φ . For each of the $3C$ literals in φ , we connect the gadget of the corresponding variable to the gadget of the clause that contains the literal. Each block of each variable gadget is connected to a specific clause gadget above or below the variable gadget, thus there are $3C$ blocks in total. Each connection is realized by a part of G_φ that we call a 2-switch. A 2-switch consists of 15 vertices and 14 edges. The mobile vertex q of the 2-switch in Fig. 3 is incident to two thick gray edges that intersect two immobile edges of the 2-switch. Thus q must move. There are (up to small perturbations) two possible positions, namely q_1 and q_2 , see Fig. 3.

The 2-switch in Fig. 3 corresponds to a positive literal. For negated literals, the switch must be mirrored either at the vertical or at the horizontal line that runs through the point m . Note that a switch can be stretched vertically in order to reach the right clause gadget. Further note that if a literal is *false*, the mobile vertex of the corresponding 2-switch must move away from the variable gadget and towards the clause

gadget to which the 2-switch belongs. In that case we say that the 2-switch *transmits pressure*.

A clause gadget consists of three vertical 2-switches and two horizontal 3-switches. A 3-switch consists of 23 vertices and 18 edges plus a small “inner” 2-switch, see the shaded area in Fig. 4. Independently from the other, each of the two 3-switches can be stretched horizontally in order to reach vertically above the variable gadget to which it connects via a 2-switch. The mobile vertex p of the left 3-switch in Fig. 4 is incident to two thick gray edges that intersect two immobile edges of the 3-switch. Thus p must move. There are (again up to small perturbations) three possible positions, namely p_1 , p_2 , and p_3 . Note that we need the inner 2-switch, otherwise there would be a fourth undesired position for moving p , namely the one labeled \bar{p} in Fig. 4. By construction, a clause gadget can be made plane by only moving the mobile vertices of all switches if and only if at most two of the three big 2-switches transmit pressure, that is, if at least one of the literals in the clause is *true*.

The graph G_φ that we have now constructed has $O(C)$ vertices, $O(C)$ edges, and $X = 26C$ crossings; $4 \cdot 3C$ in blocks and $2 \cdot 7C$ in switches. Recall that any mobile vertex is incident to two edges that each cross another edge. Thus a mobile vertex corresponds to a pair of crossings. By moving a mobile vertex to any of its predestined positions, the corresponding pair of crossings disappears. If φ is satisfiable, G_φ can be made plane by moving $K = X/2$ mobile vertices since no new crossings are introduced. If φ is not satisfiable, there is at least one pair of crossings that cannot be eliminated by moving the corresponding mobile vertex alone since all its predestined positions are blocked. Thus at least *two* vertices must be moved to eliminate that pair of crossings—and still all the other $K - 1$ pairs of crossings must be eliminated by moving at least one vertex per pair, totaling in at least $K + 1$ moves. Thus φ is satisfiable if and only if G_φ can be made plane by moving exactly K (mobile) vertices.

Recall that G_φ consists of $O(C)$ vertices and edges. We construct δ_φ step by step, starting with the vertices of the variable gadgets and then treating the clauses from innermost to outermost. In order for the 2- and 3-switches to reach far enough, note that each desired position of a mobile vertex is determined by two pairs of immobile vertices. By making the distances of the two vertex pairs (polynomially) small, the desired position can be confined to a region that is small enough to force the mobile vertex of the next switch into one of its remaining positions. Now it is clear that it is possible to place vertices at coordinates whose representation has size polynomial in the length L of a binary encoding of φ . This implies that our reduction is polynomial in L . \square

Remark 3.2 Our proof can be slightly modified to show that the problem is also hard if we are additionally given an axis-parallel rectangle that contains the initial graph drawing, and each move is constrained to that rectangle—in other words Tantaló’s version of the planarity game. In the proof we simply compute from the given planar 3-SAT formula a rectangle that is large enough to accommodate not only the initial drawing, but also the plane drawing that we get in case the formula has a satisfying truth assignment. Note that this rectangle is barely larger than the smallest axis-parallel rectangle that contains all vertices of our initial graph drawing.

We now consider the approximability of MINSHIFTEDVERTICES. Since $\text{shift}(G, \delta) = 0$ for plane drawings, we cannot use the usual definition of an approximation factor unless we slightly modify our objective function. Let $\text{shift}'(G, \delta) = \text{shift}(G, \delta) + 1$ and call the resulting decision problem MINSHIFTEDVERTICES'. Now we can modify the above reduction to get a non-approximability result.

Theorem 3.3 *For any constant real $\varepsilon \in (0, 1]$, there is no polynomial-time $n^{1-\varepsilon}$ -approximation algorithm for MINSHIFTEDVERTICES' unless $\mathcal{P} = \mathcal{NP}$.*

Proof Let n_φ be the number of vertices of the graph G_φ with drawing δ_φ that we constructed above. We add to G_φ for each edge e $n_\varphi^{(3-\varepsilon)/\varepsilon}$ copies, half of them on each side of e , in the close vicinity of e . If one of the endpoints of e is a mobile vertex, then all copies are incident to that vertex. In the following we detail where to place the other (new) endpoints of these edges.

We go through each immobile vertex v of G_φ . Let $\text{deg}_\varphi v$ be the degree of v in G_φ . Note that $1 \leq \text{deg}_\varphi v \leq 3$. If $\text{deg}_\varphi v = 1$, we place the endpoints of the copies of the edge e incident to v on the two rays that are orthogonal to e in v . On each ray we place half of the endpoints and connect them by new edges along the ray, starting with v , see vertex v_1 in Fig. 5.

Otherwise, if $\text{deg}_\varphi v > 1$, let e, e' be two edges that are incident to v and consecutive in the circular ordering around v . Now we add half of the endpoints of e and e' on a ray between e and e' emanating from v , in the same manner as above. The position of the ray depends on whether both e and e' are immobile or one of them is mobile. (Being immobile, vertex v is incident to at most one mobile edge.) In the first case we place the new vertices on the angular bisector of e and e' , see vertex v_2 in Fig. 5.

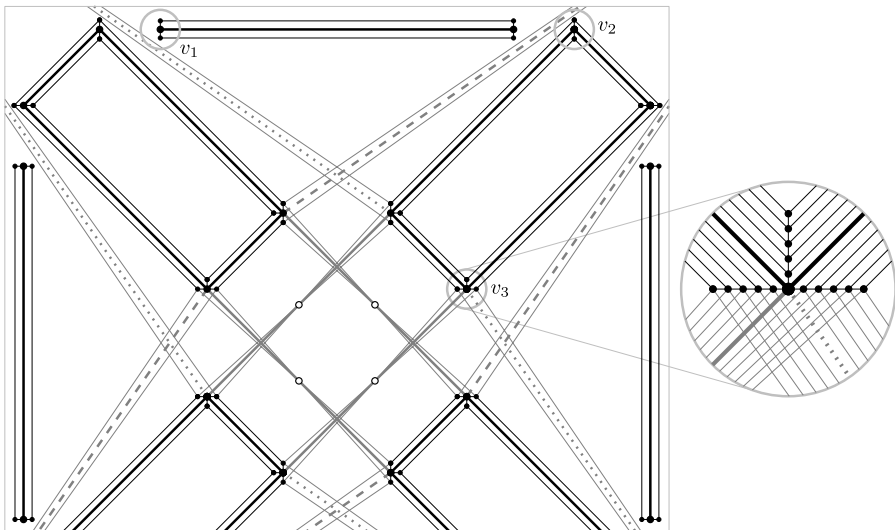


Fig. 5 Clipping of the modified variable gadget for the proof of Theorem 3.3. The old vertices and edges are drawn *thicker* than the new ones. Each old edge has $n_\varphi^{(3-\varepsilon)/\varepsilon}$ new copies

In the second case where one of the edges, say e , is mobile, note that the original and all predestined positions of e lie in an open halfplane bounded by a line ℓ through v . So we place the new vertices on ℓ , half on each side of v , see vertex v_3 in Fig. 5.

Let G be the resulting graph, δ its drawing, and $n \leq (3/2 \cdot n_\varphi^{(3-\varepsilon)/\varepsilon} + 1) \cdot n_\varphi$ the number of vertices of G . Note that φ is satisfiable if and only if $\text{shift}'(G, \delta) = \text{shift}'(G_\varphi, \delta_\varphi) = K + 1$. Otherwise, in the original graph G_φ at least one immobile vertex has to move. This vertex either is incident to a mobile edge or it is not, see Figs. 4(b) and (c), respectively. In the new graph G , which contains G_φ , also at least one (original) immobile vertex v has to move. If v is not incident to a mobile edge, in order to make space, all new vertices in the vicinity of v have to move, too. If v is incident to a mobile edge, a new vertex in the vicinity of v has to move only if it is incident to a new copy of the mobile edge. That is, in both cases, at least $n_\varphi^{(3-\varepsilon)/\varepsilon}$ vertices have to move. In other words, $\text{shift}'(G, \delta) \geq K + 2 + n_\varphi^{(3-\varepsilon)/\varepsilon}$. Note that G can be constructed in polynomial time since we have assumed ε to be a constant.

Suppose there was a polynomial-time $n^{1-\varepsilon}$ -approximation algorithm \mathcal{A} for $\text{MIN-SHIFTEDVERTICES}'$. We can bound its approximation factor by $n^{1-\varepsilon} \leq ((3/2 \times n_\varphi^{(3-\varepsilon)/\varepsilon} + 1) \cdot n_\varphi)^{1-\varepsilon} \leq (2n_\varphi^{(3-\varepsilon)/\varepsilon} \cdot n_\varphi)^{1-\varepsilon} = 2^{1-\varepsilon} n_\varphi^{(3-3\varepsilon)/\varepsilon} \leq 2n_\varphi^{(3-3\varepsilon)/\varepsilon}$. Now let M be the number of moves that \mathcal{A} needs to untangle δ . If φ is satisfiable, then $M \leq \text{shift}'(G, \delta) \cdot n^{1-\varepsilon} = (K + 1) \cdot n^{1-\varepsilon} \leq (n_\varphi + 1) \cdot 2n_\varphi^{(3-3\varepsilon)/\varepsilon} = 2n_\varphi^{(3-2\varepsilon)/\varepsilon} + O(n_\varphi^{(3-3\varepsilon)/\varepsilon})$. On the other hand, if φ is unsatisfiable, then $M \geq \text{shift}'(G, \delta) > n_\varphi^{(3-\varepsilon)/\varepsilon}$. Since we can assume that n_φ is sufficiently large, the result of algorithm \mathcal{A} (that is, the number M) tells us whether φ is satisfiable. So either our assumption concerning the existence of \mathcal{A} is wrong, or we have shown the NP-hard problem PLANAR3SAT to lie in \mathcal{P} , which in turn would mean that $\mathcal{P} = \mathcal{NP}$. \square

We now state a hardness result that establishes a connection between $\text{MIN-SHIFTEDVERTICES}$ and the well-known graph-drawing problem $\text{1BENDPOINT-SETEMBEDDABILITY}$. We define the problem $\text{1BENDPOINTSETEMBEDDABILITY-WITHCORRESPONDENCE}$ as follows. Given a planar graph $G = (V, E)$, a set S of points in the plane with rational coordinates and a one-to-one correspondence ζ between V and S , decide whether ζ can be extended to a plane 1-bend drawing of G , that is, whether G has a plane drawing δ such that $\delta(v) = \zeta(v)$ for all $v \in V$ and such that δ maps each edge of G to a 1-bend polygonal chain.

Theorem 3.4 $\text{1BENDPOINTSETEMBEDDABILITYWITHCORRESPONDENCE}$ is NP-hard.

Proof The proof uses nearly the same gadgets as in the proof of Theorem 3.1: set G'_φ to a copy of G_φ where each length-2 path (u, v, w) containing a mobile vertex v is replaced by the edge uw . We refer to this type of edges as *new* edges. The vertices of G'_φ are mapped to the corresponding vertices in δ_φ . We claim that G'_φ has a 1-bend drawing if and only if the given planar-3SAT formula φ is satisfiable.

In order to see that the claim holds, note the two differences to the proof of Theorem 3.1. First, in 1BENDPOINTSEMBEDDABILITYWITHCORRESPONDENCE all vertices *are* fixed. This makes it even easier to argue correctness. Second, any edge can bend, not only new edges, which are meant to bend. Due to the fact that vertices cannot move, however, all groups of edges that are meant to be obstacles will remain obstacles to the bending of the new edges. The only way to embed the new edges is to route them around the obstacles exactly as in Figs. 3 and 4(a). \square

Now suppose that we already know that G has a plane drawing with at most one bend per edge. Then it is natural to ask for a drawing with as few bends as possible. Let $\beta(G)$ be 1 plus the minimum number of bends over all plane 1-bend drawings of G . The following corollary shows that it is hard to approximate $\beta(G)$ efficiently.

Corollary 3.5 *Given a planar graph $G = (V, E)$, a set $S \subset \mathbb{Q}^2$, a one-to-one correspondence ζ between V and S that can be extended to a plane 1-bend drawing of G , and a constant $\varepsilon \in (0, 1]$, it is NP-hard to approximate $\beta(G)$ within a factor of $n^{1-\varepsilon}$.*

Proof We slightly change the clause gadget in the proof of Theorem 3.4. Apart from the three vertical 2-switches, the clause gadget now consists of two 4-switches and of two stacks of s edges each, see Fig. 6. Let G''_φ be the resulting graph, which depends on the given planar 3SAT formula φ . The 4-switches make sure that G''_φ always has a drawing with at most one bend per edge. Each stack is placed in the vicinity of a 4-switch such that all stack edges have to bend if the central switch edge is forced to bend into the direction of the stack. (In Fig. 6, the central switch edges in the left and right 4-switch are labeled e_C and e'_C , respectively.) If φ is not satisfiable, at least one clause evaluates to *false* and in the corresponding gadget all s edges in the left or all s edges in the right stack need to bend.

The number s of edges per stack can be set to $n'_\varphi^{(3-\varepsilon)/\varepsilon}$, where n'_φ is the number of vertices of the graph G'_φ defined in the proof of Theorem 3.4. Then, the remaining calculations for proving hardness of approximation are similar to those in the proof of Theorem 3.3. \square

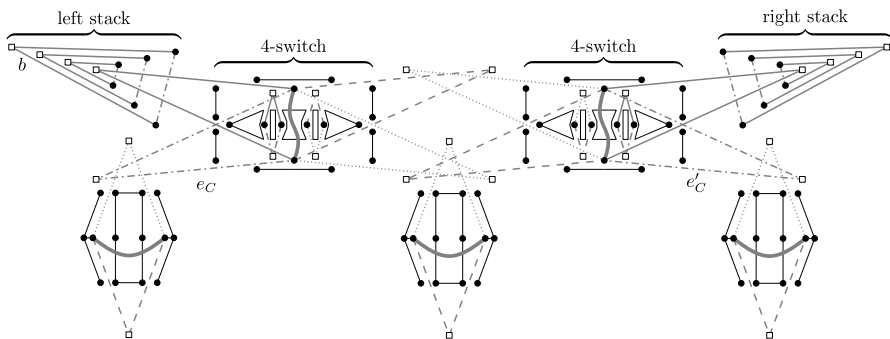


Fig. 6 Gadget of clause C adapted for the proof of Corollary 3.5. Edges e_C and e'_C each belong to a 4-switch, that is, they can be drawn in four combinatorially different ways (drawn in gray; solid vs. dashed-dotted vs. dotted vs. dashed). Note that not all vertices are marked

We do not know whether 1BENDPOINTSEMBEDDABILITYWITHCORRESPONDENCE or MINSHIFTEDVERTICES lie in \mathcal{NP} , but it is not hard to show the following.

Theorem 3.6 1BENDPOINTSEMBEDDABILITYWITHCORRESPONDENCE is in \mathcal{PSPACE} .

Proof Let $G = (V, E)$ be a planar graph, S a set of n points in the plane with rational coordinates, and ζ a one-to-one correspondence between V and S . Any 1-bend drawing of G that extends ζ is uniquely determined by choosing, for each edge e , the position (x_e, y_e) of the bend b_e of e . (If an edge uv is to be drawn without bend, any point in the relative interior of the line segment connecting $\zeta(u)$ and $\zeta(v)$ can be chosen.) Thus, the set of all plane 1-bend drawings of G that extend ζ can be represented by a subset of $\mathbb{R}^{2|E|}$. The bend b_e splits (the drawing of) the edge e into two relative open line segments to which we refer as *half-edges*.

In order to decide the existence of a plane 1-bend drawing, we specify a predicate in polynomial inequalities with integer coefficients and with variables in the set $\mathcal{E} = \{x_e, y_e \mid e \in E\}$. We do this by first expressing the condition that no two half-edges with distinct endpoints may intersect. Given four distinct points A, B, C , and D , the requirement that points C and D lie in different half-planes determined by the line through A and B can be expressed by an inequality $P(A, B, C, D) < 0$, where P is a degree-4 polynomial with integer coefficients and with variables representing the coordinates of the four points [13]. The requirement that the line segments AB and CD are disjoint is described by the disjunction $(P(A, B, C, D) > 0) \vee (P(C, D, A, B) > 0)$.

Second, we add conditions that guarantee that no bend b_e coincides with a point in S , that all bends are distinct, and that no two half-edges overlap if they share an endpoint. All these conditions can also be described as Boolean combinations of polynomial inequalities with integer coefficients and with variables from \mathcal{E} . As a consequence, deciding whether ζ extends to a 1-bend drawing of G recasts into deciding the non-emptiness of a set in $\mathbb{R}^{2|E|}$ defined by a predicate whose atomic formulas are polynomial inequalities with integer coefficients, a problem that is in \mathcal{PSPACE} [3, 21]. \square

For MINSHIFTEDVERTICES and MAXFIXEDVERTICES an additional trick is needed.

Proposition 3.7 MINSHIFTEDVERTICES and MAXFIXEDVERTICES are in \mathcal{PSPACE} .

Proof Obviously, both problems have the same optimal solutions, so it is enough to treat one of them, say MINSHIFTEDVERTICES. We build on the formulation sketched in the proof of Theorem 3.6. Additionally, we introduce a binary variable z_v for each vertex v that encodes whether we move vertex v ($z_v = 1$) or not ($z_v = 0$). In order to restrict z_v to these two values, we introduce the quadratic equation $z_v(z_v - 1) = 0$. The x -coordinate of vertex v in the plane target drawing can then be described by

$(1 - z_v)X_v + z_v x_v$, where X_v is the original x -coordinate of v , and x_v is the x -coordinate of v after a possible movement. The y -coordinate of v is treated analogously. The intersection of edges can be expressed as in the proof of Theorem 3.6. To bound the number of moved vertices by K , we introduce the inequality $\sum_{v \in V} z_v \leq K$. \square

4 Planar Graphs: Lower Bound

Any drawing of a planar graph with $n \geq 3$ vertices, other than K_3 or K_4 , can be untangled while fixing at least three vertices [27]. In this section, we give an algorithm proving that

$$\text{fix}(G) \geq f(n) = \sqrt{\frac{(\log n) - 1}{\log \log n}}$$

for any planar graph G with $n \geq 4$ vertices. Note that f actually grows, albeit very slowly: $f(n) > 3$ only for some $n \approx 6 \cdot 10^{15}$. Partially building on our algorithm, Bose et al. [2] showed that $\text{fix}(G) \geq \sqrt[4]{(n+1)}/2$, a bound greater than 3 for $n > 161$.

We first give some definitions (Sect. 4.1) and sketch the basic idea of our algorithm (Sect. 4.2). Then we describe our algorithm (Sect. 4.3) and prove its correctness (Sect. 4.4). The bound $\text{fix}(G) \geq f(n)$ depends on finding a plane embedding of G that contains a long simple path with an additional property. We show how to find such an embedding in Sect. 4.5.

4.1 Definitions and Notation

Recall that a *plane embedding* of a planar graph is given by the circular order of the edges around each vertex and by the choice of the outer face. A plane embedding of a planar graph can be computed in linear time [9]. If G is triangulated, a plane embedding of G is determined by the choice of the outer face. Further recall that an edge of a graph is called a *chord* with respect to a path (or cycle) Π if the edge does not lie on Π but both its endpoints are vertices of Π .

For a point $p \in \mathbb{R}^2$, let $x(p)$ and $y(p)$ be the x - and y -coordinates of p , respectively. We say that p lies *vertically below* $q \in \mathbb{R}^2$ if $x(p) = x(q)$ and $y(p) \leq y(q)$. For a polygonal path $\Pi = v_1, \dots, v_k$, we denote by $V_\Pi = \{v_1, \dots, v_k\}$ the set of vertices of Π and by $E_\Pi = \{v_1 v_2, \dots, v_{k-1} v_k\}$ the set of edges of Π . We call a polygonal path $\Pi = v_1, \dots, v_k$ *x -monotone* if $x(v_1) < \dots < x(v_k)$. In addition, we say that a point $p \in \mathbb{R}^2$ lies *below* an x -monotone path Π if p lies vertically below a point p' (not necessarily a vertex!) on Π . Analogously, a line segment \overline{pq} lies below Π if every point $r \in \overline{pq}$ lies below Π . We do not always strictly distinguish between a vertex v of G and the point $\delta(v)$ to which this vertex is mapped in a particular drawing δ of G . Similarly, we write vw both for the edge of G and the straight-line segment connecting $\delta(v)$ with $\delta(w)$.

4.2 Basic Idea

Note that in order to establish a lower bound on $\text{fix}(G)$, we can assume that the given graph G is triangulated. Otherwise we can triangulate G arbitrarily (by fixing

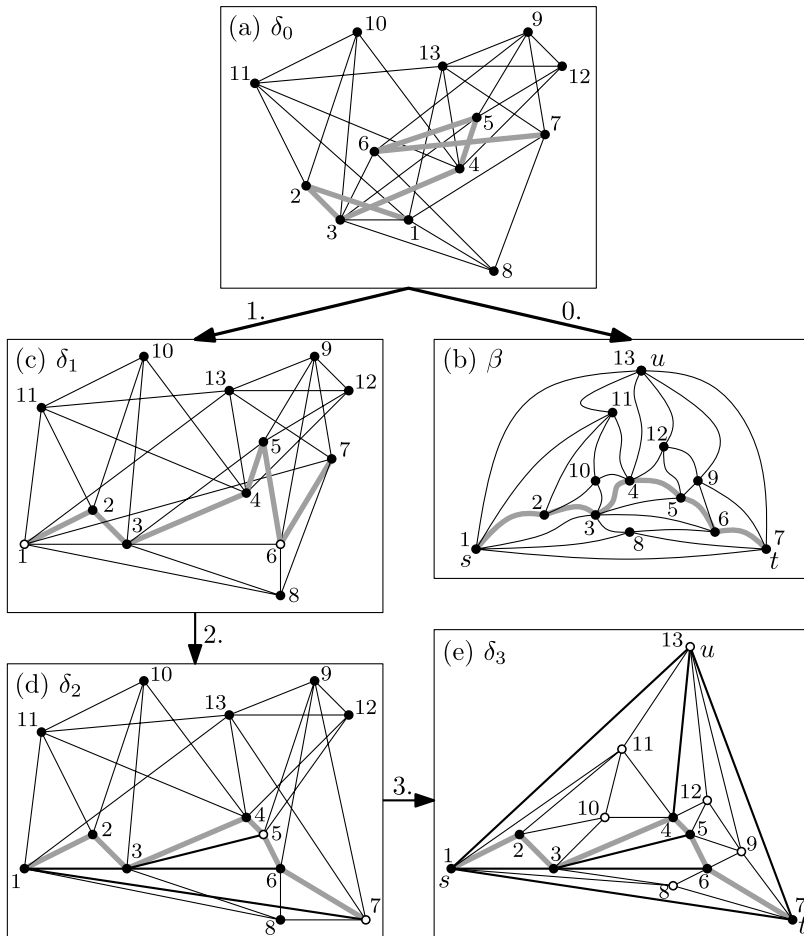


Fig. 7 An example run of our algorithm. **(a)** Input: the given non-plane drawing δ_0 of a triangulated planar graph G . **(b)** Plane embedding β of G with path Π (drawn in gray) that connects two vertices on the outer face. To untangle δ_0 we first make Π x -monotone **(c)**, then we bring all chords (bold segments) to one side of Π **(d)**, move u to a position on the other side of Π where u sees all vertices in V_Π , and finally move the vertices in $V \setminus (V_\Pi \cup \{u\})$ to suitable positions within the faces bounded by the bold gray and black edges **(e)**. Vertices that move from δ_{i-1} to δ_i are marked by circles; those that do not move are marked by black disks

an embedding of G and adding edges until all faces are 3-cycles) and work with the resulting triangulated planar graph. A plane drawing of the latter yields a plane drawing of G . So let G be a triangulated planar graph, and let δ_0 be any drawing of G . It will also be convenient to assume that in the given drawing δ_0 , the vertices of G are mapped to points with pairwise distinct x -coordinates. By slightly rotating the drawing δ_0 we can always achieve this.

The basic idea of our algorithm is to find a plane embedding β of G such that there exists a long simple path Π connecting two vertices s and t of the outer triangle stu with the property that all chords of Π lie on one side of Π (with respect to β) and u

lies on the other. For an example of such an embedding β , see Fig. 7(b). We describe how to find β and Π depending on the maximum degree and the diameter of G in Sect. 4.5. For the time being, we assume they are given. Now our goal is to produce a drawing of G according to the embedding β and at the same time keep many of the vertices of Π at their positions in δ_0 . Having all chords on one side is the crucial property of Π , we use to achieve this. We allow ourselves to move all other vertices of G to any location we like, a process we will occasionally refer to as *drawing* certain subgraphs of G . This gives us a lower bound on $\text{fix}(G, \delta)$ in terms of the number l of vertices of Π . Our method is illustrated in Fig. 7; we give the details in the next subsection.

4.3 Description of the Algorithm

Let C denote the set of chords of Π . We assume that these chords lie to the right of Π when we traverse this path from s to t in the embedding β . (Note that “below” is not defined in an embedding.) Let V_{bot} denote the set of vertices of G that lie to the right of Π in β , and let $V_{\text{top}} = V \setminus (V_{\Pi} \cup V_{\text{bot}})$. Note that u lies in V_{top} . Let I be a subset of the vertices of Π such that no two vertices in I are connected by a chord of Π . We will choose I such that $|I| \geq (l + 1)/2$, and our method tries to fix many of the vertices in I .

In Step 1 of our algorithm we move some of the vertices in V_{Π} from the position they have in δ_0 to new positions such that the resulting ordering of the vertices in V_{Π} according to increasing x -coordinates is the same as the ordering along Π in β . This yields a new (usually non-plane) drawing δ_1 of G that maps Π on an x -monotone polygonal path Π_1 . By Theorem 2.1 we can choose δ_1 such that at least $\sqrt{|I|}$ of the vertices in I remain fixed. Let $F \subseteq I \subseteq V_{\Pi}$ be the set of the fixed vertices. Note that $\delta_1(v) = \delta_0(v)$ for all $v \in V \setminus V_{\Pi}$, see Fig. 7(c).

Once we have constructed Π_1 , we have to find suitable positions for the vertices in $V_{\text{top}} \cup V_{\text{bot}}$. This is simple for the vertices in V_{top} : if we move vertex u , which lies on the outer face, far enough above Π_1 , then the polygon P_1 bounded by Π_1 and by the edges us and ut will be *star-shaped*. Recall that a polygon P is called *star-shaped* if the interior of its kernel is nonempty, and the *kernel* of a clockwise-oriented polygon P is the intersection of the right half-planes induced by the edges of P . Now if P_1 is star-shaped, we have fulfilled one of the assumptions of the following result of Hong and Nagamochi [8] for drawing *triconnected* graphs, that is, graphs that cannot be disconnected by removing two vertices. We will use their result in order to draw into P_1 the subgraph G_{top}^+ of G induced by $V_{\text{top}} \cup V_{\Pi}$ *excluding* the chords in C .

Theorem 4.1 [8] *Given a triconnected plane graph H , every drawing δ^* of the outer facial cycle of H on a star-shaped polygon P can be extended in linear time to a plane drawing of H .*

Observe, however, that G_{top}^+ is *not* necessarily triconnected: vertex u may be adjacent to vertices on Π other than s and t . In order to fix this, we split G_{top}^+ into smaller units along the edges incident to u . Let $(s =)w_1, w_2, \dots, w_l(= t)$ be the sequence of vertices of Π . Let (i, k) be a pair of integers such that $1 \leq i < k \leq l$, vertices w_i

and w_k are adjacent to u and any vertex w_j with $i < j < k$ is not adjacent to u . Consider the subgraph of G_{top}^+ induced by the vertices that lie (with respect to β) inside of or on the cycle $u, w_i, w_{i+1}, \dots, w_j$. In the following we convince ourselves that this subgraph is actually triconnected. Let \mathcal{H}_{top} be the family of all such subgraphs.

Recall that a planar graph H is called a *rooted triangulation* [1] if in every plane drawing of H there exists at most one facial cycle with more than three vertices. According to Avis [1], the following lemma is well known.

Lemma 4.2 [1] *A rooted triangulation is triconnected if and only if no facial cycle has a chord.*

Now it is clear that we can apply Theorem 4.1 to draw each of the subgraphs in \mathcal{H}_{top} . By the placement of u , each drawing region is star-shaped, and by construction, each subgraph is chordless and thus triconnected. However, to draw the graph G_{bot}^+ induced by $V_{\text{bot}} \cup V_{\Pi}$ (including the chords in C), we must work a little harder.

In Step 2 of our algorithm we once more change the drawing of Π . Let $V^* = V_{\Pi} \setminus I$. Note that every chord of Π has at least one of its endpoints in V^* . Now we go through the vertices in V^* in a certain order, moving each vertex vertically down as far as necessary (see vertices 5 and 7 in Fig. 7(d)) to achieve two goals: (a) all chords in C move below the resulting polygonal path Π_2 , and (b) the faces bounded by Π_2 , the edge st , and the chords become star-shaped polygons. This defines a new drawing δ_2 , which leaves all vertices in F and all vertices in $V \setminus V_{\Pi}$ fixed.

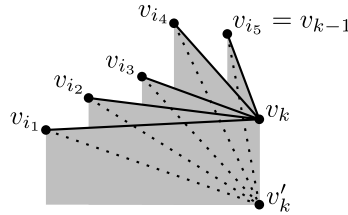
In Step 3 we use the fact that Π_2 is still x -monotone. This allows us to move vertex u to a location above Π_2 where it can see every vertex of Π_2 . Now Π_2 , the edges of type uw_i (with $1 < i < l$) and the chords in C partition the triangle ust into star-shaped polygons with the property that the subgraphs of G that have to be drawn into these polygons are all rooted triangulations, and thus triconnected. This means that we can apply Theorem 4.1 to each of them. The result is our final—and plane—drawing δ_3 of G , see Fig. 7(e).

4.4 Correctness of the Algorithm

We now show that our algorithm indeed produces a plane drawing where many of the vertices on the chosen path Π are fixed. To this end, recall that $F \subseteq I$ is the set of vertices in V_{Π} we fixed in Step 1, that is, in the construction of the x -monotone polygonal path Π_1 . Our goal is to fix the vertices in F when we construct Π_2 , which also is an x -monotone polygonal path but has two additional properties: (a) all chords in C lie below Π_2 and (b) the faces induced by Π , w_1w_l , and the chords in C are star-shaped polygons. The following lemmas form the basis for the proof of the main theorem of this section (Theorem 4.7), which shows that this can be achieved.

Lemma 4.3 *Let $\Pi = v_1, \dots, v_k$ be an x -monotone polygonal path such that (i) the segment v_1v_k lies below Π and (ii) the polygon P bounded by Π and v_1v_k is star-shaped. Let v'_k be any point vertically below v_k . Then the polygon $P' = v_1, \dots, v_{k-1}, v'_k$ is also star-shaped.*

Fig. 8 Illustration of the proof of Lemma 4.4



Proof Only two edges change when we move vertex v_k to its new position v'_k , namely v_1v_k and $v_{k-1}v_k$. Consider the remaining $k - 2$ edges that do not change and let K be the intersection of the corresponding right half-planes. Since the $k - 2$ edges form an x -monotone path, K is not bounded. Let q be a point in the interior of the kernel of P . Then q lies in the interior of K and, moreover, every point that is vertically below q also lies in the interior of K . Let q' be a point vertically below q and sufficiently close to the edge $v_1v'_k$. Then q' lies in the interior of the kernel of P' , and therefore P' is star-shaped by definition. \square

Lemma 4.4 *Let $\Pi = v_1, \dots, v_k$ be an x -monotone polygonal path, and let D be a set of pairwise non-crossing straight-line segments with endpoints in V_Π that all lie below Π . Let v'_k be a point vertically below v_k , let $\Pi' = v_1, \dots, v_{k-1}, v'_k$, and finally let D' be a copy of D with each segment $v_iv_k \in D$ replaced by $v_iv'_k$.*

Then the segments in D' are pairwise non-crossing and all lie below Π' .

Proof Let $v_{i_1}v_k, \dots, v_{i_m}v_k$ be the straight-line segments incident to v_k (both on the monotone path Π and in D), sorted clockwise around v_k such that $v_{i_m} = v_{k-1}$. Note that, since the straight-line segments in D are below Π , the vertices v_{i_1}, \dots, v_{i_m} are also sorted according to increasing x -coordinates, and all of them have smaller x -coordinate than v_k . Hence, the situation is as depicted in Fig. 8.

For $1 \leq j \leq m$, let B_j denote the set of points that lie below the straight-line segment $v_{i_j}v_k$ and define $B = \bigcup_{j=1}^m B_j$, see the shaded region in Fig. 8. Note that the interior of B cannot contain any vertices of Π since this would contradict the fact that Π is x -monotone or the fact that all straight-line segments in D are below Π . But this implies that none of the straight-line segments $v_{i_j}v'_k$ (drawn dotted in Fig. 8) is crossed by a straight-line segment in D' since this would yield a contradiction to the fact that the straight-line segments in D are non-crossing or to the fact that the interior of B does not contain a vertex of Π . No other crossings can occur in D' since the straight-line segments in D are non-crossing. This finishes the proof. \square

Recall that we aim at finding a large set $I \subseteq V_\Pi$ such that no two vertices in I are connected by a chord of Π . The set F of fixed vertices will be a subset of I . Note that I may contain vertices connected by an *edge* of Π . In the following lemma, the set V^* contains all vertices of Π that we have to move in order to draw the chords of Π straight-line; clearly such a set must cover all chords of Π . Thus the set V^* plays the role of the complement of I .

Lemma 4.5 *Let $\Pi = v_1, \dots, v_k$ be an x -monotone polygonal path. Let C_Π be a set of chords of Π that can be drawn as non-crossing curved lines below Π . Let V^**

be a vertex cover of C_Π . Then there is a way to modify Π by decreasing the y -coordinates of the vertices in V^* such that the resulting straight-line drawing δ^* of $G_\Pi = (V_\Pi, E_\Pi \cup C_\Pi)$ is plane, the bounded faces of δ^* are star-shaped, and all edges in C_Π lie below the modified polygonal path, which is also x -monotone. The coordinates of the vertices of the modified path have bit length $O(nL)$, where L is the maximum bit length of the vertex coordinates of Π .

Proof We use induction on the number $m = |C_\Pi|$ of chords. If $m = 0$, we need not modify Π . So, suppose that $m > 0$. We first choose a chord $vw \in C_\Pi$ with $x(v) < x(w)$ such that there is no other edge $v'w' \in C_\Pi$ with the property that $x(v') \leq x(v)$ and $x(w') \geq x(w)$. Clearly, such an edge always exists. Then we apply the induction hypothesis to $C_\Pi \setminus \{vw\}$. This yields a modification Π' of Π such that Π' is x -monotone, all edges in the resulting straight-line drawing δ' of $G_\Pi - vw$ lie below Π' , and all bounded faces in this drawing are star-shaped.

Now consider the chord vw and, without loss of generality, assume that $v \in V^*$. Let Z denote the set of those vertices $z \in V_\Pi$ with the property that no point vertically below z and distinct from z is contained in an edge of the drawing δ' . Note that, since Π' is x -monotone, there must exist a point p vertically below v such that for no vertex $z \in Z$, the straight-line segment pz crosses any edge in the drawing δ' of $G_\Pi - vw$.

Let $i \in \{1, \dots, k\}$ be such that $v = v_i$. We move vertex v to the point p to obtain a new drawing δ'' of $G_\Pi - vw$. Then we apply Lemma 4.4 to the rightmost vertex of the x -monotone subpath of Π' with vertices v_1, \dots, v_i and, similarly, we apply Lemma 4.4 to the leftmost vertex of the x -monotone subpath of Π' with vertices v_i, \dots, v_k . It follows that this does not produce any crossings among the edges in the drawing δ'' . Moreover, by our choice of the chord vw , for each face in the drawing δ' of $G_\Pi - vw$ that has vertex v in its facial cycle, v must be the leftmost or the rightmost vertex in this facial cycle. Hence, we can apply Lemma 4.3 to these faces. This yields that they remain star-shaped in δ'' . By the choice of p we ensure that the bounded face that results from adding the straight-line segment pw to the drawing δ'' is also star-shaped.

Concerning the size of the coordinates we argue as follows. Without loss of generality we can assume that all vertices of Π have negative y -coordinates. Now consider the addition of the i th chord vw . Let y_{i-1} be the minimum y -coordinate of a vertex in the drawing before moving vertex v down. Then it is not hard to check that, in order to add the chord vw without introducing any crossings, it suffices to move v down to a point with y -coordinate $(2R_x)y_{i-1}$, where R_x is the ratio of the maximum over the minimum difference between the x -coordinates of any two distinct vertices in V_Π . Solving the recurrence for y_i yields $|y_i| \leq |(2R_x)^i y_0|$. Therefore, since there are only $O(n)$ chords, the y -coordinates in the resulting x -monotone path can be encoded using $O(nL)$ bits. □

Remark 4.6 Unfortunately, there are indeed instances where our algorithm actually needs $\Theta(n^2)$ bits for representing all y -coordinates of the modified path. Let $k > 0$ be an odd integer, and let Π be a path with $n = 2k + 1$ vertices v_1, \dots, v_n , where $v_i = (i, 0)$ for $1 \leq i \neq k + 1 \leq n$ and $v_{k+1} = (k + 1, -1)$, see the thick light-gray path

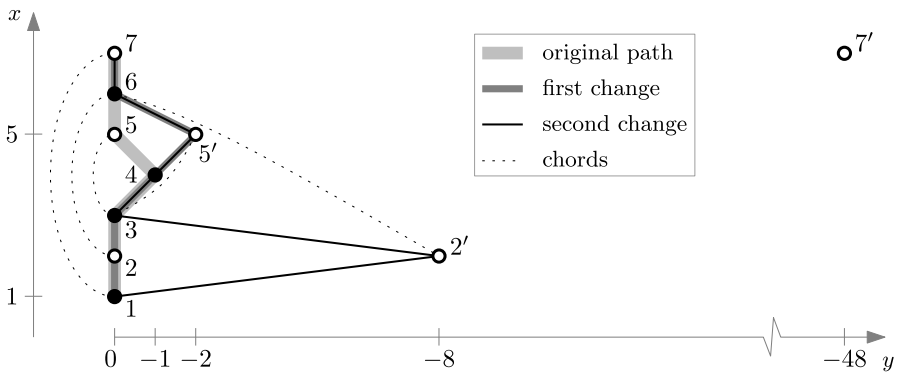


Fig. 9 When applying the algorithm that proves Lemma 4.5 to the *thick light-gray* n -vertex path with the *dotted chords* indicated on the left and with the vertex cover V^* indicated by *circles*, some of the y -coordinates of the resulting path need more than n bits. Note that the x -axis is vertical

in Fig. 9. We set $C_\Pi = \{v_1 v_n, v_2 v_{n-1}, \dots, v_k v_{k+2}\}$ (drawn with dotted arcs in Fig. 9) and $V^* = \{v_2, v_4, \dots, v_{k-1}, v_{k+2}, \dots, v_{n-2}, v_n\}$ (marked with circles in Fig. 9).

Our algorithm straightens the chords in the order from innermost to outermost, that is, vertices are moved in the order $v_{k+2}, v_{k-1}, v_{k+4}, \dots, v_2, v_n$. To simplify presentation, let w_1, w_2, \dots, w_k denote the vertices of V^* in this order, and let $w_0 = v_{k+1}$. For $i = 0, \dots, k$, denote the final position of w_i by $(x_i, -y_i)$. Then clearly $|x_i - x_{i-1}| = 2i - 1$ for $i = 1, \dots, k$. The edges incident to w_{i-1} have slope $\pm y_{i-1}$ (with the exception of the irrelevant edge $w_0 w_1$), thus $y_i > y_{i-1} + y_{i-1} \cdot |x_i - x_{i-1}| = y_{i-1} \cdot 2i$. The recursion solves to $y_i > 2^i i!$.

Now suppose that we have modified the x -monotone path Π_1 according to Lemma 4.5. Then the resulting x -monotone path Π_2 admits a straight-line drawing of the chords in C below Π_2 such that the bounded faces are star-shaped polygons, see the example in Fig. 7(d). Recall that $u \in V_{\text{top}}$ is the vertex of the outer triangle in β that does not lie on Π . We now move vertex u to a position above Π_2 such that all edges $uw \in E$ with $w \in V_\Pi$ can be drawn without crossing Π_2 and such that the resulting faces are star-shaped polygons. Since Π_2 is x -monotone, this can be done. As an intermediate result, we obtain a plane straight-line drawing of a subgraph of G where all bounded faces are star-shaped. It remains to find suitable positions for the vertices in $(V_{\text{top}} \setminus \{u\}) \cup V_{\text{bot}}$. For every star-shaped face f , there is a unique subgraph G_f of G that must be drawn inside this face. Note that by our construction every edge of G_f that has both endpoints on the boundary of f must actually be an edge of the boundary. Therefore, G_f is a rooted triangulation where no facial cycle has a chord. Now Lemma 4.2 yields that G_f is triconnected. Finally, we can use the result of Hong and Nagamochi [8] (see Theorem 4.1) to draw each subgraph of type G_f and thus finish our construction of a plane straight-line drawing of G , see the example in Fig. 7(e). We summarize.

Theorem 4.7 *Let G be a triangulated planar graph that contains a simple path $\Pi = w_1, \dots, w_l$ and a face $uw_1 w_l$. If G has an embedding β such that $uw_1 w_l$ is the*

outer face, u lies on one side of Π , and all chords of Π lie on the other side, then $\text{fix}(G) \geq \sqrt{(l+1)/2}$.

Proof We continue to use the notation introduced earlier in this section. Recall that we aim at finding a large set $I \subseteq V_\Pi$ such that no two vertices in I are connected by a chord of Π . The complement $V_\Pi \setminus I$ of this set I is the vertex cover V^* of C that we need for applying Lemma 4.5.

Further, $F \subseteq I$ is the set of vertices that we fixed in the first step, that is, in the construction of the x -monotone path Π_1 . It follows from Proposition 1 in the paper by Pach and Tardos [18] that we can make sure that $\text{fix}(G) = |F| \geq \sqrt{|I|}$. Consider the graph G_C with vertex set V_Π and edge set C . An independent set in G_C has exactly the property that we want for I . Thus it suffices to show that the l -vertex graph G_C has an independent set I of size at least $(l+1)/2$. We do this by giving a simple algorithm.

Our algorithm is greedy: we always take a vertex v of smallest degree, put it in the independent set I under construction, remove v and the neighbors of v from V_Π , and remove the edges incident to these vertices from C . We repeat this until G_C is empty.

Note that G_C initially has at least one isolated (that is, degree-0) vertex and that the bound is obvious if G_C is a forest—the algorithm first picks all isolated vertices and then repeatedly picks leaves. Even if G_C contains cycles, the algorithm always picks vertices of degree at most 2. This is due to the fact that all chords lie on one side of Π , and thus G_C is and remains outerplanar, and any outerplanar graph has a vertex of degree at most 2.

Let n_i be the number of vertices that have degree i when they are put in I . As observed above, $|I| = n_0 + n_1 + n_2$. Whenever we put a vertex of degree i into I , we remove $i+1$ vertices from V_Π , thus $l = n_0 + 2n_1 + 3n_2$. Let f be the number of bounded faces of G_C . Whenever the algorithm removes a degree-2 vertex, the number of bounded faces of G_C decreases by one, thus $f = n_2$. We claim—and will prove below—that $f+1 \leq n_0$. Now adding $1 \leq n_0 - n_2$ to the above expression for l yields $l+1 \leq 2n_0 + 2n_1 + 2n_2 = 2|I|$, or $|I| \geq (l+1)/2$, which proves the theorem.

It remains to prove our claim, that is, $n_0 - 1 \geq f$. In other words, we need to show that G_C contains at least one isolated vertex more than bounded faces. Recall that G_C does not include the edges of Π . For a chord $c = w_i w_j$ in C , we define $\{w_i, w_{i+1}, \dots, w_j\} \subseteq V_\Pi$ to be the *span* of c . Now consider a face F of G_C with vertices $w_{i_1}, w_{i_2}, \dots, w_{i_k}$ and $i_1 < i_2 < \dots < i_k$. The edges of F are $w_{i_1} w_{i_2}, w_{i_2} w_{i_3}, \dots, w_{i_k} w_{i_1}$. Note that the span of $w_{i_k} w_{i_1}$ contains the span of every other edge of F . We define the span of F to be the span of the edge $w_{i_k} w_{i_1}$.

We prove our claim by induction on f . As noted above, G_C contains at least one isolated vertex. This establishes the base of the induction. Now suppose that $f > 0$. Consider the set M of all faces of G_C whose span is maximal with respect to set inclusion. If $|M| > 1$, we apply the induction hypothesis to the subgraphs of G_C induced by the spans of the faces in M . Otherwise, let F^* be the only face in M , and let e_1, \dots, e_{k-1} be the edges of F^* whose span is properly contained in the span of F^* . We apply the induction hypothesis to the subgraphs of G_C induced by the spans of e_1, \dots, e_{k-1} . Since $k \geq 3$, there are at least two such subgraphs. Each of them contains at least one isolated vertex more than bounded faces. Taking F^* into

account, we conclude that G_C also contains at least one isolated vertex more than bounded faces. This completes the proof of our claim. \square

4.5 Finding a Suitable Path

We finally present two strategies for finding a suitable path Π . Neither depends on the geometry of the given drawing δ_0 of G . Instead, they exploit the graph structure of G . The first strategy works well if G has a vertex of large degree and, even though it is very simple, yields asymptotically tight bounds for outerplanar graphs.

Lemma 4.8 *Let G be a triangulated planar graph with maximum degree Δ . Then $\text{fix}(G) \geq \sqrt{(\Delta + 1)/2}$.*

Proof Let u be a vertex of degree Δ and consider a plane embedding β of G where vertex u lies on the outer face. Since G is planar, such an embedding exists. Let $\{w_1, \dots, w_\Delta\}$ be the set of neighbors of u sorted clockwise around u in β . This gives us the desired polygonal path $\Pi = w_1, \dots, w_\Delta$ that has no chords on the side that contains u . Thus Theorem 4.7 yields $\text{fix}(G) \geq \sqrt{(\Delta + 1)/2}$. \square

Lemma 4.8 yields a lower bound for outerplanar graphs that is asymptotically tight as we will see in Sect. 6.

Corollary 4.9 *Let G be an outerplanar graph with n vertices. Then $\text{fix}(G) \geq \sqrt{n/2}$.*

Proof We select an arbitrary vertex u of G . Since G is outerplanar, we can triangulate G in such a way that in the resulting triangulated planar graph G' vertex u is adjacent to every other vertex in G' . Thus the maximum degree of a vertex in G' is $n - 1$, and the result follows by Lemma 4.8. \square

Our second strategy works well if the diameter d of G is large.

Lemma 4.10 *Let G be a triangulated planar graph of diameter d . Then $\text{fix}(G) \geq \sqrt{d}$.*

Proof We choose two vertices s and v such that a shortest s – v path has length d . We compute any plane embedding of G that has s on its outer face. Let t and u be the neighbors of s on the outer face. Recall that a *Schnyder wood* (or *realizer*) [23] of a triangulated plane graph is a (special) partition of the edge set into three spanning trees each rooted at a different vertex of the outer face. Edges can be viewed as being directed to the corresponding roots. The partition is special in that the cyclic pattern in which the spanning trees enter and leave a vertex is the same for all inner vertices. Schnyder [23] showed that this cyclic pattern ensures that the three unique paths from a vertex to the three roots are vertex-disjoint and chordless. Let π_s , π_t , and π_u be the “Schnyder paths” from v to s , t , and u , respectively. Note that the length of π_s is at least d , and the lengths of π_t and π_u are both at least $d - 1$. Let Π be the path that goes from s along π_s to v and from v along π_t to t . The length of Π is at least

$2d - 1$. Note that, due to the existence of π_u , the path Π has no chords on the side that contains u . Thus, Theorem 4.7 yields $\text{fix}(G, \delta) \geq \sqrt{d}$. \square

Next we determine the trade-off between the two strategies above.

Theorem 4.11 *Let G be a planar graph with $n \geq 4$ vertices. Then $\text{fix}(G) \geq \sqrt{\frac{(\log n) - 1}{\log \log n}}$, where the base of logarithms is 2.*

Proof Let G' be an arbitrary triangulation of G . Note that the maximum degree Δ of G' is at least 3 since $n \geq 4$ and G' is triangulated. To relate Δ to the diameter d of G' , we use a very crude counting argument—Moore’s bound: starting from an arbitrary vertex of G , we bound the number of vertices we can reach by a path of a certain length. Let j be the smallest integer such that $1 + (\Delta - 1) + (\Delta - 1)^2 + \dots + (\Delta - 1)^j \geq n$. Then $d \geq j$. By the definition of j we have $n \leq (\Delta - 1)^{j+1} / (\Delta - 2)$, which we can simplify to $n \leq 2(\Delta - 1)^j$ since $\Delta \geq 3$. Hence we have $d \geq j \geq \frac{(\log n) - 1}{\log(\Delta - 1)}$.

Now, if $\Delta \geq (\log n) + 2$, Lemma 4.8 yields $\text{fix}(G') \geq \sqrt{((\log n) + 3)/2}$. Otherwise $d \geq \frac{(\log n) - 1}{\log \log n}$, and we can apply Lemma 4.10. Observing that $\text{fix}(G) \geq \text{fix}(G')$ yields the desired bound. \square

Remark 4.12 The proof of Theorem 4.11 (together with the auxiliary results stated earlier) yields an $O(n^2)$ -time algorithm for untangling a given straight-line drawing of a planar graph G with n vertices by moving some of its vertices to new positions. The first step, that is, computing the x -monotone path Π_1 , takes $O(n \log n)$ time [22]. Moving the vertices of Π_1 such that the faces induced by the path and its chords become star-shaped takes $O(\gamma(n)n)$ time (Lemma 4.5), where $\gamma(n) = O(n)$ is an upper bound on the time needed to perform an elementary operation involving numbers of bit length $O(n)$. The remaining steps of our method can be implemented to run in $O(n)$ time. This includes calling the algorithm of Hong and Nagamochi [8] and computing the Schnyder wood [23], which we need in the proof of Lemma 4.10.

5 Planar Graphs: Upper Bound

We now give an upper bound for general planar graphs that is better than the upper bound $O((n \log n)^{2/3})$ of Pach and Tardos [18] for cycles. Our construction uses the following sequence, which we call σ_q and which we reuse in Section 6:

$$((q - 1)q, (q - 2)q, \dots, 2q, q, \underline{0}, 1 + (q - 1)q, \dots, 1 + q, \underline{1}, \dots, q^2 - 1, \dots, (q - 1) + q, \underline{q - 1}).$$

Note that σ_q can be written as $(\sigma_q^0, \sigma_q^1, \dots, \sigma_q^{q-1})$, where $\sigma_q^i = ((q - 1)q + i, (q - 2)q + i, \dots, 2q + i, q + i, i)$ is the subsequence of σ_q that consists of all q numbers in σ_q that are congruent to i modulo q . To stress this, the last element in each of

these subsequences is underlined in σ_q . Thus σ_q consists of q^2 distinct numbers. It is not hard to see the following.

Observation 5.1 The longest increasing or decreasing subsequence of σ_q has length q .

We call two subsequences $\Sigma = s_1, s_2, \dots, s_l$ and $\Sigma' = s'_1, s'_2, \dots, s'_l$ of σ_q *separated* if:

- (i) s_l comes before s'_1 or s'_l comes before s_1 in σ_q , and
- (ii) $\max(\Sigma) < \min(\Sigma')$ or $\max(\Sigma') < \min(\Sigma)$.

Observation 5.2 Let Σ and Σ' be two separated decreasing or two separated increasing subsequences of σ_q . Then $|\Sigma \cup \Sigma'| \leq q + 1$.

Proof First consider the case that Σ and Σ' are both decreasing. Since they are separated we can assume without loss of generality that $\max(\Sigma) < \min(\Sigma')$. We define $V_i = \{iq + j : 0 \leq j \leq q - 1\}$ for $i = 0, \dots, q - 1$. Then, since Σ and Σ' are both decreasing, they can each have at most one element in common with every V_i . Now suppose that they have both one element in common with some V_{i_0} . Then, since $\max(\Sigma) < \min(\Sigma')$, Σ cannot have an element in common with any V_i , $i > i_0$, and Σ' cannot have an element in common with any V_i , $i < i_0$. Therefore, $|\Sigma \cup \Sigma'| \leq q + 1$.

If Σ and Σ' are both increasing, then, similarly as above, every subsequence σ_q^i with $0 \leq i \leq q - 1$ can have at most one element in common with each Σ and Σ' . Moreover, at most one subsequence σ_q^i can have an element in common with both Σ and Σ' . This implies that $|\Sigma \cup \Sigma'| \leq q + 1$, as required. □

Theorem 5.3 *For any integer $n_0 > 0$, there exists a planar graph G with $n \geq n_0$ vertices and $\text{fix}(G) \leq \sqrt{n - 2} + 1$.*

Proof Let $q = \lceil \sqrt{n_0} \rceil$. We define the graph G as a path of q^2 vertices $1, 2, \dots, q^2$ all connected to the two endpoints of an edge $\{a, b\}$ with $a, b \notin \{1, 2, \dots, q^2\}$, see Fig. 10(a). Hence G has $n = q^2 + 2$ vertices. Let δ_G be the drawing of G where vertices $1, 2, \dots, q^2$ are placed on a vertical line ℓ in the order given by σ_q . We place vertices a and b below the others on ℓ , see Fig. 10(b). Let δ'_G be an arbitrary plane drawing of G obtained by untangling δ_G . Since all faces of G are 3-cycles, the outer face in δ'_G is a triangle. All faces of G contain a or b . This has two consequences. First, a and b must move to new positions in δ'_G , otherwise all other vertices would have to move. Second, at least one of them, say a , appears on the outer face.

Case 1. Vertex b also lies on the outer face. Then there are just two possibilities for the embedding of G : as in Fig. 10(a) or with the indices of all vertices reversed, that is, vertex i becomes $q^2 - i - 1$. Now let $0 \leq i < j < k \leq q^2 - 1$ be three fixed vertices. By symmetry we can assume that j lies in $\Delta(a, b, i)$. Then k also lies in $\Delta(a, b, i)$ since the path connecting j to k does not intersect the sides of this triangle. Note that k cannot lie between i and j on ℓ as otherwise one of the edges $\{a, k\}$ and

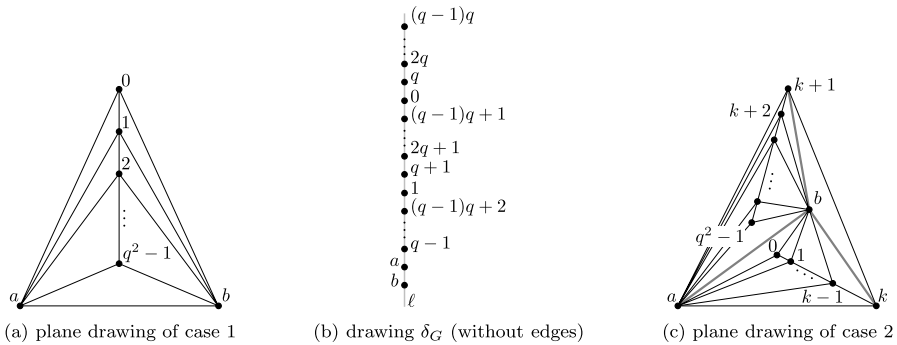


Fig. 10 Drawings of the graph G that we use in the proof of Theorem 5.3

$\{b, k\}$ would intersect the polygonal path connecting i to j . Thus, each triplet of fixed vertices forms a monotone sequence along ℓ . This in turn yields that *all* fixed vertices in $\{0, \dots, q^2 - 1\}$ form a monotone sequence along ℓ . Due to the construction of σ_q , such a sequence has length at most $q = \sqrt{n - 2}$.

Case 2. Vertex b does not lie on the outer face. Then the outer face is of the form $\Delta(a, k, k + 1)$ with $0 \leq k \leq q^2 - 2$. The three edges $\{b, a\}$, $\{b, k\}$, and $\{b, k + 1\}$ incident to b split $\Delta(a, k, k + 1)$ into the three triangles $\Delta(a, k, b)$, $\Delta(a, b, k + 1)$, and $\Delta(b, k, k + 1)$, see Fig. 10(c). Every vertex of δ'_G lies in one of them. Since δ'_G is plane, vertex $k - 1$ must belong to $\Delta(a, k, b)$, and, by induction, so do all vertices $i \leq k$; similarly, all vertices $i \geq k + 1$ lie in $\Delta(a, b, k + 1)$. We can thus apply the argument of Case 1 to each of the two subgraphs contained in $\Delta(a, b, k)$ and $\Delta(a, b, k + 1)$. This yields two separated monotone sequences of length at most q each. Note, however, that both are increasing or both are decreasing since one type forces a to the left and b to the right of ℓ and the other does the opposite. Due to Observation 5.2, the length of two separated monotone subsequences of σ_q sums up to at most $q + 1 = \sqrt{n - 2} + 1$.

To summarize, Case 2 yields a larger number of potentially fixed vertices, and thus $\text{fix}(G, \delta_G) \leq q + 1 = \sqrt{n - 2} + 1$.

Note that actually $\text{fix}(G, \delta_G) = q + 1$ as we can fix, for example, the vertices $0, q, 2q, \dots, (q - 1)q$, and $(q - 1)q + 2$. □

6 An Upper Bound for Outerplanar Graphs

In this section we show that the lower bound $\text{fix}(H) \geq \sqrt{n/2}$ that holds for any outerplanar graph H with n vertices (see Corollary 4.9) is asymptotically tight in the worst case.

Theorem 6.1 *For any integer n_0 , there exists an outerplanar graph H with $n \geq n_0$ vertices and $\text{fix}(H) \leq 2\sqrt{n - 1} + 1$.*

Proof Let $q = \lceil \sqrt{n_0} \rceil$. We define the outerplanar graph H as a path of q^2 vertices $0, 1, \dots, q^2 - 1$ and an extra vertex $c = q^2$ that is connected to all other vertices, see

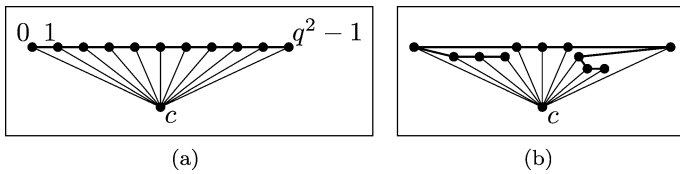


Fig. 11 The outerplanar graph H that we use in the proof of Theorem 6.1

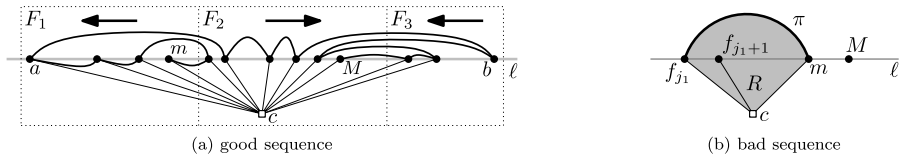


Fig. 12 Analyzing the sequence of fixed vertices along the line ℓ

Fig. 11(a). Hence H has $n = q^2 + 1$ vertices. Let δ_H be the drawing of H where all vertices are placed on a horizontal line ℓ as follows. Vertices $0, \dots, q^2 - 1$ are arranged in the order σ_q introduced in Sect. 5, and vertex c can go to an arbitrary (free) spot on ℓ .

In the following we show that $\text{fix}(H, \delta_H) \leq 2q + 1 = 2\sqrt{n-1} + 1$. To this end, let δ'_H be an arbitrary plane drawing of H obtained by untangling δ_H , and let F be the set of fixed vertices. Note that H has many plane embeddings—for example, Fig. 11(b)—but only two outerplane embeddings: Fig. 11(a) and its mirror image. Our proof exploits the fact that the simple structure of H forces the left-to-right sequence of the fixed vertices to also have a very simple structure.

Consider the drawing δ'_H . If vertex c lies on ℓ in δ'_H , then, since c is connected by an edge to every other vertex of H and all these vertices lie on ℓ in the drawing δ_H , at most two of these other vertices can be fixed. Hence, the interesting case is that c does not lie on ℓ in δ'_H and, therefore, $c \notin F$. Hence, $F \subseteq \{0, 1, \dots, q^2 - 1\}$. We only consider the interesting case that $|F| \geq 2$. Let m and M be the minimum and maximum in F , respectively. Without loss of generality we assume that c lies below ℓ and that m lies to the left of M (otherwise we reflect δ'_H on the x/y -axis). Let a and b be the left- and rightmost vertices in F , see Fig. 12(a).

Let $F_0 = f_1, f_2, \dots, f_{|F|}$ be the vertices in F ordered as we meet them along ℓ from left to right. Let $F_1 = f_1, f_2, \dots, f_{j_1}$ be the longest subsequence of F_0 starting at f_1 such that $f_{i-1} > f_i$ for $2 \leq i \leq j_1$. Note that by definition $f_1 = a$. We claim that $f_{j_1} = m$. Assume to the contrary that $f_{j_1} \neq m$. Then $f_{j_1} > m$ and, clearly, F_1 does not contain m . Thus m lies to the right of f_{j_1+1} .

Consider the path $\pi = f_{j_1}, f_{j_1} - 1, \dots, m$ in H . Since $f_{j_1+1} > f_{j_1} > m$, f_{j_1+1} is not a vertex of π . Let R be the polygon bounded by π and by the edges cf_{j_1} and cm . Since δ'_H is plane, R is simple. Note that f_{j_1+1} lies in the interior of R , and M lies in the exterior of R , as indicated in Fig. 12(b), where the interior of R is shaded. To see this, first note that M , which lies to the right of m , cannot lie in the interior of R since otherwise the path in H with vertices $M, M - 1, \dots, f_{j_1}$ would intersect π or one of the edges that connect a vertex on π with c . But then, since π can intersect neither

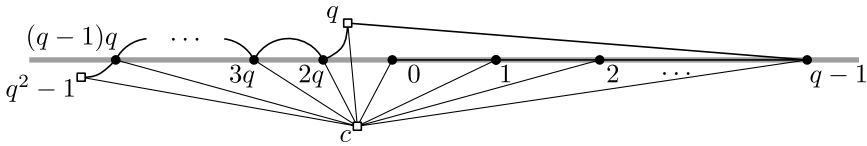


Fig. 13 This plane drawing of the graph H defined in the proof of Theorem 6.1 shows that $\text{fix}(H, \delta_H) \geq 2q - 2$ since it fixes the vertices $0, 1, 2, \dots, q - 1, 2q, 3q, \dots, (q - 1)q$ of δ_H . The curved arcs indicate chains of vertices that have been moved

edge cM nor edge cf_{j_1+1} , vertex f_{j_1+1} must lie in the interior of R , as required. This yields a contradiction since the path $M, M - 1, \dots, f_{j_1+1}$ in H must now intersect the boundary of R . Thus our assumption $f_{j_1} \neq m$ is wrong, and we have indeed $f_{j_1} = m$.

Now let $F_2 = f_{j_1+1}, f_{j_1+2}, \dots, f_{j_2}$ be the longest subsequence of F_0 starting at f_{j_1+1} such that $f_{i-1} < f_i$ for each i with $j_1 + 1 \leq i \leq j_2$. With similar arguments as above we can show that $f_{j_2} = M$. Moreover, let $F_3 = f_{j_2+1}, f_{j_2+2}, \dots, f_{j_3}$ be the subsequence of F_0 starting at f_{j_2+1} such that $f_{i-1} > f_i$ for each i with $j_2 + 1 \leq i \leq j_3$. Again, with similar arguments as above we can show that either F_3 is empty or $f_{j_3} = f_{|F|}$. In addition, we can show in an analogous way that $f_1 < f_{|F|}$ holds.

Thus, the set F is partitioned into F_1, F_2 , and F_3 . The sequence F_2 is increasing, and both F_1 and F_3 are decreasing (or empty). Thus, by Observation 5.1, $|F_2| \leq q$ and, by Observation 5.2, $|F_1| + |F_3| \leq q + 1$, since $f_1 < f_{|F|}$ implies that F_1 and F_3 are separated. Hence, $|F| \leq 2q + 1$, as required.

Note that this upper bound is almost tight: $\text{fix}(H, \delta_H) \geq 2q - 2$ as indicated in Fig. 13. □

7 Conclusions

In this paper, we have presented several new results on the problem of untangling a given drawing of a graph, a problem originally introduced by Watanabe [29] for the special case of cycles.

On the computational side, we have proved that `MINSHIFTEDVERTICES` is NP-hard and also hard to approximate; we also showed that our proof technique extends to another graph drawing problem, namely `1BENDPOINTSEMBEDDABILITY` with given vertex–point correspondence. Related questions that remain open are the inapproximability of `MAXFIXEDVERTICES` and the hardness of `MAXFIXEDVERTICES` and `MINSHIFTEDVERTICES` for special classes of graphs such as cycles. We have shown that all these problems lie in \mathcal{PSPACE} , but do they also lie in \mathcal{NP} ? Also, we are not aware of any result in the direction of parameterized complexity.

On the combinatorial side, Table 1 summarizes the best currently known worst-case bounds for untangling several important classes of planar graphs. It reveals that the gap for general planar graphs is probably the most interesting remaining open problem in the field.

Acknowledgements We thank Vida Dujmović for her comments on an earlier version of this paper. She came up with an idea that directly improved our lower bounds in Theorem 4.11 and Corollary 4.9 by a factor of $\sqrt{3}$. Thanks to Tomáš Gavenčiak for interesting discussions about the results of Kang et al. [10] and to Ignaz Rutter for the idea behind Theorem 3.6. We thank all anonymous and the non-anonymous referee for their detailed comments that have helped us to improve the presentation of the paper.

Open Access This article is distributed under the terms of the Creative Commons Attribution Noncommercial License which permits any noncommercial use, distribution, and reproduction in any medium, provided the original author(s) and source are credited.

References

1. Avis, D.: Generating rooted triangulations without repetitions. *Algorithmica* **16**, 618–632 (1996)
2. Bose, P., Dujmović, V., Hurtado, F., Langerman, S., Morin, P., Wood, D.R.: A polynomial bound for untangling geometric planar graphs. In: P. Ossona de Mendez, D. Poulalhon, M. Pocchiola, J.L. Ramírez Alfonsín, G. Schaeffer (eds.) *Proc. Internat. Conf. Topological Geom. Graph Theory (TGGT'08)*. *Electr. Notes Discrete Math.*, pp. 205–210 (2008). Long version available at <http://arxiv.org/abs/0710.1641>
3. Canny, J.: Some algebraic and geometric computations in PSPACE. In: *Proc. of the 20th ACM Sympos. Theory Comput. (STOC'88)*, pp. 460–469 (1988)
4. Cibulka, J.: Untangling polygons and graphs. In: P. Ossona de Mendez, D. Poulalhon, M. Pocchiola, J.L. Ramírez Alfonsín, G. Schaeffer (eds.) *Proc. of Internat. Conf. Topological Geom. Graph Theory (TGGT'08)*, *Electr. Notes Discrete Math.* pp. 200–204 (2008). Long version available at <http://arxiv.org/abs/0802.1312>
5. Erdős, P., Szekeres, G.: A combinatorial problem in geometry. *Compos. Math.* **2**, 463–470 (1935)
6. Fáry, I.: On straight-line representation of planar graphs. *Acta Sci. Math. (Szeged)* **11**, 229–233 (1948)
7. Goac, X., Kratochvíl, J., Okamoto, Y., Shin, Ch.-S., Wolff, A.: Moving vertices to make drawings plane. In: *Proc. of the 15th International Symposium on Graph Drawing (GD'07)*. *Lecture Notes in Computer Science*, vol. 4875, pp. 101–112. Springer, Berlin (2008)
8. Hong, S.-H., Nagamochi, H.: Convex drawing of graphs with non-convex boundary. *Discrete Appl. Math.* **156**, 2368–2380 (2008)
9. Hopcroft, J., Tarjan, R.E.: Efficient planarity testing. *J. ACM* **21**, 549–568 (1974)
10. Kang, M., Pikhurko, O., Ravsky, A., Schacht, M., Verbitsky, O.: Obfuscated drawings of planar graphs. ArXiv report (March 2008). Available at <http://arxiv.org/abs/0803.0858>
11. Kaufmann, M., Wiese, R.: Embedding vertices at points: Few bends suffice for planar graphs. *J. Graph Algorithms Appl.* **6**(1), 115–129 (2002)
12. Knuth, D.E., Raghunathan, A.: The problem of compatible representatives. *SIAM J. Discrete Math.* **5**(3), 422–427 (1992)
13. Kratochvíl, J., Matoušek, J.: Intersection graphs of segments. *J. Combin. Theory Ser. B* **62**, 289–315 (1994)
14. Lichtenstein, D.: Planar formulae and their uses. *SIAM J. Comput.* **11**(2), 329–343 (1982)
15. Liske, N.: Planarity game. <http://www.creativecouple.de/en/Projects/Planarity.html> (accessed 8 September 2008)
16. Lubiw, A., Petrick, M., Spriggs, M.: Morphing orthogonal planar graph drawings. In: *Proc. of the 17th Annu. ACM-SIAM Symposium on Discrete Algorithms (SODA'06)*, pp. 222–230 (2006)
17. Misue, K., Eades, P., Lai, W., Sugiyama, K.: Layout adjustment and the mental map. *J. Visual Lang. Comput.* **6**(2), 183–210 (1995)
18. Pach, J., Tardos, G.: Untangling a polygon. *Discrete Comput. Geom.* **28**(4), 585–592 (2002)
19. Pach, J., Wenger, R.: Embedding planar graphs at fixed vertex locations. *Graphs Comb.* **17**(4), 717–728 (2001)
20. Ravsky, A., Verbitsky, O.: On collinear sets in straight line drawings. ArXiv report (June–July 2008). Available at <http://arxiv.org/abs/0806.0253>
21. Renegar, J.: On the computational complexity and geometry of the first-order theory of the reals, part I: Introduction. Preliminaries. The geometry of semi-algebraic sets. The decision problem for the existential theory of the reals. *J. Symb. Comput.* **13**(3), 255–300 (1992)

22. Schensted, C.: Longest increasing and decreasing subsequences. *Can. J. Math.* **13**, 179–191 (1961)
23. Schnyder, W.: Planar graphs and poset dimension. *Order* **5**(4), 323–343 (1989)
24. Spillner, A., Wolff, A.: Untangling a planar graph. In: Proc. of the 34th International Conference on Current Trends in Theory and Practice of Computer Science (SOFSEM'08). Lecture Notes in Computer Science, vol. 4910, pp. 473–484. Springer, Berlin (2008)
25. Stein, S.K.: Convex maps. *Proc. Am. Math. Soc.* **2**, 464–466 (1951)
26. Tantalo, J.: Planarity. Web site at <http://planarity.net/> (accessed 21 May 2007)
27. Verbitsky, O.: On the obfuscation complexity of planar graphs. *Theor. Comput. Sci.* **396**(1–3), 294–300 (2008) (first published at <http://arxiv.org/abs/0705.3748>)
28. Wagner, K.: Bemerkungen zum Vierfarbenproblem. *Jahresbericht Deutsch. Math. Verein.* **46**, 26–32 (1936)
29. Watanabe, M.: Open problem. In: 5th Czech–Slovak Symposium on Combinatorics (1998)