

Convexity-increasing morphs of planar graphs [☆]

Linda Kleist ^a, Boris Klemz ^{b,*}, Anna Lubiw ^c, Lena Schlipf ^d, Frank Staals ^e, Darren Strash ^f

^a Technische Universität Braunschweig, Germany

^b Freie Universität Berlin, Germany

^c University of Waterloo, Canada

^d FernUniversität in Hagen, Germany

^e Utrecht University, the Netherlands

^f Hamilton College, USA



ARTICLE INFO

Article history:

Received 20 June 2018

Received in revised form 20 December 2018

Accepted 20 May 2019

Available online 26 July 2019

Keywords:

Morphing

Convex graph drawing

Convexifying

Tutte

ABSTRACT

We study the problem of *convexifying* drawings of planar graphs. Given any planar straight-line drawing of an internally 3-connected graph, we show how to morph the drawing to one with strictly convex faces while maintaining planarity at all times. Our morph is *convexity-increasing*, meaning that once an angle is convex, it remains convex. We give an efficient algorithm that constructs such a morph as a composition of a linear number of steps where each step either moves vertices along horizontal lines or moves vertices along vertical lines. Moreover, we show that a linear number of steps is worst-case optimal. To obtain our result, we use a well-known technique by Hong and Nagamochi for finding redrawings with convex faces while preserving y -coordinates. Using a variant of Tutte's graph drawing algorithm, we obtain a new proof of Hong and Nagamochi's result which comes with a better running time. This is of independent interest, as Hong and Nagamochi's technique serves as a building block in existing morphing algorithms.

© 2019 Elsevier B.V. All rights reserved.

1. Introduction

A *morph* between two planar straight-line drawings Γ_0 and Γ_1 of a plane graph G is a continuous movement of the vertices from their positions in Γ_0 to their positions in Γ_1 , with the edges following along as straight-line segments between their endpoints. A morph is *planar* if it preserves planarity of the drawing at all times.

Motivated by applications in animation and in reconstruction of 3D shapes from 2D slices, the study of morphing has focused on finding a morph between two given planar drawings. The existence of planar morphs was established long ago [7, 38], followed by algorithms that produce good visual results [21,23], and algorithms that find “piece-wise linear” morphs with a linear number of steps [2]. Our focus is somewhat different, and more aligned with graph drawing goals—our input is a planar graph drawing and our aim is to morph it to a better drawing, in particular to a convex drawing. A morph *convexifies* a given straight-line graph drawing if the result is a (strictly) *convex graph drawing*, i.e. a planar straight-line graph drawing in which every face is a (strictly) convex polygon. For a survey on convex graph drawing, see [35].

[☆] A preliminary version of this paper appeared in the proceedings of WG 2018 [26].

* Corresponding author.

E-mail addresses: kleist@ibr.cs.tu-bs.de (L. Kleist), klemz@inf.fu-berlin.de (B. Klemz), alubiw@uwaterloo.ca (A. Lubiw), lena.schlipf@fernuni-hagen.de (L. Schlipf), f.staals@uu.nl (F. Staals), dstrash@hamilton.edu (D. Strash).




Fig. 1. A sequence of convexity-increasing morphs (horizontal, vertical, horizontal) that morph a straight-line drawing of a graph G (1) into a strictly convex drawing of G (4). The dashed segments are auxiliary edges added by our algorithm before each morphing step.

We first observe that it is easy, using known results, to find a planar morph that convexifies a given drawing—we can just create a convex drawing with the same faces (assuming such a drawing exists), and morph to that specific drawing using the known planar morphing algorithms. (For a discussion of the techniques used, see the section on related work.) In this paper, we are interested in a stronger condition: we want to find a convexifying morph which is also *convexity-increasing*, meaning that once an angle of an inner face is convex, it remains convex. We illustrate a convexity-increasing morph in Fig. 1.

Besides the theoretical goal of studying continuous motion that is monotonic in some measure (e.g. edge lengths [25]), another motivation comes from visualization—a morph of a graph drawing should maintain the user’s “mental map” [33] which means changing as little as possible, and making observable progress towards a goal. Most previous morphing algorithms fail to provide convexity-increasing morphs even if the target is a convex drawing because they start by triangulating the drawing. Therefore, an original convex angle may be subdivided by new triangulation edges, so there is no constraint that keeps it convex.

Related work. To the best of our knowledge, previous work on convexity-increasing morphs only considers the case when the input graph is a cycle (or a path). Connelly et al. [13] and Canterella et al. [8] gave algorithms to convexify a simple polygon while preserving edge lengths. Since their motions are “expansive”, they are convexity-increasing. Aichholzer et al. [1] gave an algorithm to find a “visibility-increasing” morph of a simple polygon to a convex polygon; this condition implies the condition of being convexity-increasing.

In related work, there is an algorithm to morph a convex drawing to another convex drawing of the same graph while preserving planarity and convexity [4]. Such morphs are convexity-increasing by default, but do not address our problem since our initial drawing is not convex.

Many previous morphing algorithms find “piece-wise linear” morphs, where the morph is composed of discrete steps and each step moves vertices along straight lines. A morph is called *linear* if each vertex moves along a straight line at constant speed; different vertices are allowed to move at different speeds, and some may remain stationary. A linear morph is completely specified by the initial and final drawings. If, in addition, all the lines along which vertices move are parallel, then the morph is called *unidirectional* [2].

Alamdari et al. [2] gave an algorithm with runtime $O(n^3)$ that takes as input two planar straight-line drawings of a graph on n vertices with the same combinatorial embedding, i.e., the drawings have the same outer face and for each face (every boundary component) has the same cyclic ordering of edges. It then constructs a planar morph between the two drawings that consists of a sequence of $O(n)$ unidirectional morphs.

Our contribution. In this paper, we give the first algorithm that convexifies a given straight-line planar drawing Γ via a planar convexity-increasing morph. The only requirement is that the plane graph G represented by Γ admits a strictly convex drawing. This is the case if and only if G is internally 3-connected; see Section 2.1 for the definition and related discussions.

In fact, we achieve the following stronger property—our morphs are composed of a linear number of *horizontal* and *vertical* morphs. A *horizontal* morph moves all vertices at constant speeds along horizontal lines; a *vertical* morph is defined analogously. These are special cases of unidirectional morphs. See Fig. 1 for an illustration.

Orthogonality is a very desirable and well-studied criterion for graph drawing [17], in part because there is evidence that the human visual cortex comprehends orthogonal lines more easily [5,31,34]. Similarly, it seems natural that orthogonal *motion* should be easier to comprehend, though this criterion has not been explored in previous work.

Our main result is summarized in the following theorem.

Theorem 1. *Let Γ be a planar straight-line drawing of an internally 3-connected graph G on n vertices. Then Γ can be morphed to a strictly convex drawing via a sequence of at most $3.5n + 2$ convexity-increasing planar morphing steps each of which is either horizontal or vertical.*

In the special cases that G is 3-connected or Γ has a convex outer face, the upper bound on the number of morphing steps can be improved to $1.5n + 2$ or $\max\{2, r + 1\}$, respectively, where r denotes the number of internal reflex angles.

Furthermore, there is an $O(n^{1+\omega/2})$ time algorithm to find the sequence of morphs, where ω is the matrix multiplication exponent.

The run time is $O(n^{2.5})$ with Gaussian elimination, improved to $O(n^{2.1865})$ using the current fastest matrix multiplication method with $\omega \approx 2.3728639$ [27]. Our model of computation is the real-RAM—we do not have a polynomial bound on the bit-complexity of the coordinates of the vertices in the sequence of drawings that specify the morph. However, previous morphing algorithms had no such bounds either.

Our algorithm has another advantage in terms of visualization over previous morphing algorithms such as the one by Alamdari et al. [2]. These algorithms tend to “almost contract” vertices, which destroys the user’s “mental map” of the graph. We do not use contractions, and therefore expect our morphs to be useful for visualizations.

A main ingredient of our proof is a result of Hong and Nagamochi [24] that gives conditions (and an algorithm) for re-drawing a planar straight-line drawing to have convex faces, while preserving the y -coordinates of the vertices (“level planar drawings of hierarchical-st plane graphs,” in their terminology). Angelini et al. [4] strengthened Hong and Nagamochi’s result to strictly convex faces. We give a new proof of the strengthened result using Tutte’s graph drawing algorithm. Thereby, we improve the runtime of Hong and Nagamochi’s result from $O(n^2)$ to $O(n^{\omega/2})$. This is of independent interest, as Hong and Nagamochi’s technique serves as a building block in several other morphing algorithms [2,4,14]. In particular, our improvement also speeds up the run-time of the morphing algorithm of Alamdari et al. [2] from $O(n^3)$ to $O(n^{1+\omega/2})$:

Theorem 2 (Theorem 1.1 in [2] with an improved runtime). *Given a planar graph G on n vertices and two straight-line planar drawings of G with the same combinatorial embedding, there is a planar morph between the two drawings that consists of $O(n)$ unidirectional morphs. Furthermore, the morph can be found in time $O(n^{1+\omega/2})$.*

In Appendix B we describe the algorithm by Alamdari et al. and justify the improved runtime due to our version of Hong and Nagamochi’s result.

Theorem 1 guarantees the existence of a convexity-increasing morph to a strictly convex drawing where the morph is composed of $O(n)$ horizontal/vertical morphs. This is optimal in the worst case. In fact we show something stronger:

Theorem 3. *For any $n \geq 3$, there exists a drawing of an internally 3-connected graph on n vertices for which any convexifying planar morph composed of a sequence of linear morphing steps requires $\Omega(n)$ steps.*

Organization. Our paper is structured as follows: We begin with preliminaries in Section 2. The proof of Theorem 1 is presented in Section 3, and the proof of the improved running time of Theorem 2 is given in Section 4. The lower bound on the number of morphs, namely Theorem 3, is shown in Section 5. Finally, a discussion of the size of the grid needed for the intermediate drawings of our morph can be found in Section 6. We conclude with open problems in Section 7.

2. Preliminaries

In this section, we introduce the concepts we will use. We formally define convex drawings and internally 3-connected graphs in Section 2.1 as well as y -monotone drawings in Section 2.2. We proceed by stating several useful properties of unidirectional morphs in Section 2.3. Finally, we address the concept of finding convex drawings in Section 2.4.

2.1. Convex drawings and internal 3-connectivity

Given a planar straight-line drawing Γ of a graph, its *angles* are formed by pairs of consecutive edges around a face, with the angle measured inside the face. An *internal angle* is an angle of an inner face. We say an angle is *reflex* if it exceeds π , *convex* if it is at most π , and *strictly convex* if it is less than π . A drawing Γ is *convex* if the boundary of every face is a convex polygon, i.e., angles of the inner faces are convex and angles of the outer face are reflex or of size π . The drawing is *strictly convex* if the boundary of every face is a strictly convex polygon.




Fig. 2. (a) An external separation pair $\{u, v\}$ and its two components C and C' . (b–c) In these cases $\{u, v\}$ is a non-external separation pair because the component marked with a jagged arrow has no vertex on the outer face. In these cases there is no convex drawing of G .

Conditions for the existence of convex drawings. Throughout, we assume that our input is a drawing of a graph that admits a strictly convex drawing with the same combinatorial embedding. Necessary and sufficient conditions for the existence of a strictly convex drawing were given by Tutte [40], Thomassen [39], and Hong and Nagamochi [24]. These conditions can be tested in linear time by the algorithm of Chiba et al. [9].

Such conditions are usually stated for a fixed convex drawing of the outer face, but the conditions become simpler when, as in our case, the drawing of the outer face may be freely chosen—in particular, may be chosen to have no 3 consecutive collinear vertices. Internal vertices of degree 2 can also be dealt with directly: In a convex drawing, an internal vertex of degree 2 must be drawn as a point in the interior of the straight line segment formed by its two incident edges. This has two implications. Firstly, a graph with an internal vertex of degree 2 has no strictly convex drawing. Secondly, for a convex drawing we may eliminate every internal degree 2 vertex by repeatedly replacing a path of two edges by a single edge. However, if this produces multiple edges, then there exists no convex drawing.

With these observations, the necessary and sufficient conditions for the existence of a (strictly) convex drawing become quite simple to state. A plane graph G is *internally 3-connected* if the graph is 2-connected and any separation pair $\{u, v\}$ is *external*, meaning that u and v lie on the outer face and that every connected component of $(G - u - v)$ contains a vertex of the outer face of G . Observe that the two neighbors of an internal vertex of degree 2 form a separation pair that is not external. The results of Tutte [40], Thomassen [39], and Hong and Nagamochi [24] become:

Lemma 4. *Let G be a plane graph with outer face C . Then*

1. G has a strictly convex drawing with outer face C if and only if G is internally 3-connected.
2. G has a convex drawing with outer face C if and only if repeatedly eliminating internal vertices of degree 2 produces a graph that has no multiple edges and is internally 3-connected.

Note that a separation pair which is not external can have both of its vertices on the outer face, see Fig. 2(b, c). For this reason, we refer to a separation pair which is not external as *non-external*.¹

Structure of internally 3-connected graphs. There are multiple well-known equivalent definitions of internal 3-connectivity. Each of them provides a different perspective on the concept and it will be convenient to be able to refer to all of them. Hence, we state the following characterization:

Lemma 5. *Let G be a plane 2-connected graph and let f_o denote its outer face. The following statements are equivalent:*

- (I1) G is internally 3-connected.
- (I2) Inserting a new vertex v in f_o and adding edges between v and all vertices of f_o results in a 3-connected graph.
- (I3) From each internal vertex w of G there exist three paths to f_o that are pairwise disjoint except for the common vertex w .

Proof. (I1) obviously implies (I2), which in turn implies (I3) by Menger's theorem. It remains to show that (I3) implies (I1). So suppose that from each internal vertex w there exist three paths to f_o that are disjoint except for the common vertex w . It follows that for each pair of vertices p, q each connected component of $(G - p - q)$ has at least one vertex on f_o , as otherwise there can be at most two disjoint paths from w to f_o . Moreover, both p and q belong to f_o : Assume otherwise. Since G is 2-connected, f_o is a simple cycle and, hence, all vertices of $(f_o - p - q)$ belong to the same connected component of $(G - p - q)$. However, by assumption from each internal vertex w there exist at least one path to f_o in $(G - p - q)$ and, so, $(G - p - q)$ is connected; a contradiction. \square

¹ instead of using the more canonical, but misleading term *internal*.

The following statement gives a characterization of external separation pairs and describes the structure of internally 3-connected graphs. For an illustration see Fig. 2(a).

Observation 6. *Let H be a plane 2-connected graph and let $\{u, v\}$ be a separation pair of H . Then, $\{u, v\}$ is external if and only if all of the following conditions hold:*

- (E1) *Vertices u and v belong to the outer face of H .*
- (E2) *The outer face of H decomposes into two internally disjoint paths (u, w_1, \dots, w_j, v) and $(v, w'_1, \dots, w'_\ell, u)$ each with at least 3 vertices, i.e. $j \geq 1$ and $\ell \geq 1$.*
- (E3) *Vertices w_1, \dots, w_j belong to a connected component C of $(H - u - v)$.*
- (E4) *Vertices w'_1, \dots, w'_ℓ belong to a connected component C' of $(H - u - v)$.*
- (E5) *The graph $(H - u - v)$ has no connected component other than C, C' .*
- (E6) *The components C and C' are distinct.*

Proof. If the six conditions hold, then clearly $\{u, v\}$ is an external separation pair. On other hand, if $\{u, v\}$ is an external separation pair, u and v belong to the outer face of H (E1). Since H is 2-connected, its outer face is a *simple* cycle. Further, the removal of u and v splits the graph into at least two connected components each of which has a vertex that belongs to the outer face of H . Hence, the removal of u and v decomposes the outer face into two internally disjoint paths (u, w_1, \dots, w_j, v) and $(v, w'_1, \dots, w'_\ell, u)$ each with at least 3 vertices (E2). Since (w_1, \dots, w_j) is a path in $(H - u - v)$, its vertices belong to a connected component C (E3). Similarly, (w'_1, \dots, w'_ℓ) is a path in $(H - u - v)$, its vertices belong to a connected component C' (E4). Since these two paths together with u and v cover the entire outer face of H , there can not be any more components (E5). Finally, since there are at least two components, C and C' have to be distinct (E6). \square

2.2. y -Monotone drawings

A face of a planar graph drawing is *y-monotone* if the boundary of the face consists of two y -monotone paths. A path is *y-monotone* if the y -coordinates along the curve realizing the path are strictly increasing. These definitions apply to general planar graph drawings, not just straight-line drawings. We note that the directed graphs that have drawings with y -monotone faces are the *st-planar* graphs, which are well-studied [16].

We say a vertex v is a *local minimum* (*local maximum*) of face f in a drawing Γ if the neighbors of v in f lie above v (below v , respectively). A *local extremum* refers to a local minimum or a local maximum. Note that a face f is y -monotone if and only if it has exactly one local maximum and exactly one local minimum. Alternatively, a face is y -monotone if it has no reflex local extremum.

2.3. Linear and unidirectional morphs

A linear morph is completely specified by the initial and the final drawing. To denote the linear morph from a drawing Γ_1 to a drawing Γ_2 , we use the notation $\langle \Gamma_1, \Gamma_2 \rangle$. Restricting to linear morphs is a sensible way to discretize morphs—essentially, it asks for the vertex trajectories to be piece-wise linear. At first glance, the restriction to unidirectional morphs seems arbitrary and restrictive. However, as discovered by Alamdari et al. [2], it is easier to prove the existence of unidirectional morphs. Also, unidirectional morphs have many nice properties, as we explain in this section. Suppose we do a horizontal morph. Then every vertex keeps its y -coordinate. Alamdari et al. [2] gave conditions on the initial and final drawing that guarantee that the horizontal morph between them is planar:

Lemma 7. [2, in the proof of Lemma 13] *If Γ and Γ' are two planar straight-line drawings of the same graph such that every line parallel to the x -axis crosses the same ordered sequence of edges and vertices in both drawings, then the linear morph from Γ to Γ' is planar.*

Observe that the conditions of the lemma imply that every vertex is at the same y -coordinate in Γ and Γ' so the linear morph between them is horizontal. Also note that the lemma generalizes in the obvious way to any direction, not just the direction of the x -axis. We note several useful consequences of Lemma 7.

Lemma 8. *Let $\Gamma_1, \Gamma_2, \Gamma_3$ be three planar straight-line drawings where the linear morphs $\langle \Gamma_1, \Gamma_2 \rangle$ and $\langle \Gamma_2, \Gamma_3 \rangle$ are horizontal and planar. Then the linear morph $\langle \Gamma_1, \Gamma_3 \rangle$ is a horizontal planar morph.*

Proof. The morphs $\langle \Gamma_1, \Gamma_2 \rangle$ and $\langle \Gamma_2, \Gamma_3 \rangle$ are horizontal and planar, so every line parallel to the x -axis crosses the same ordered sequence of edges and vertices in Γ_1 and Γ_3 . Then by Lemma 7 the morph $\langle \Gamma_1, \Gamma_3 \rangle$ is horizontal and planar. \square

Lemma 9. *Let Γ_1, Γ_2 be two planar straight-line drawings such that $\langle \Gamma_1, \Gamma_2 \rangle$ is a horizontal morph. Then the convexity status of each angle changes at most once in the morph $\langle \Gamma_1, \Gamma_2 \rangle$, i.e., an angle cannot change more than once between reflex and convex or vice versa. If additionally, every convex internal angle of Γ_1 is also convex in Γ_2 then the morph $\langle \Gamma_1, \Gamma_2 \rangle$ is convexity-increasing.*

Proof. This result is a generalization of [4, Lemma 7], and both are proved using basic properties of unidirectional morphs from [2]. Consider an angle formed by points a, b, c . If a and c both lie above b , or both lie below b , then the angle maintains its convexity status (convex or reflex) during any horizontal morph. So suppose that the ordering of the points by y -coordinate is a, b, c . Suppose that the clockwise angle abc is convex at two time points t and t' with $t < t'$ during the horizontal morph. If we add the edge ac we obtain a triangle and the horizontal line through b crosses ac and b in the same order at both time points. Thus, the morph between t and t' is planar by Lemma 7, so the angle is convex at all times between t and t' .

We prove the second statement by contrapositive. Suppose the horizontal morph $\langle \Gamma_1, \Gamma_2 \rangle$ is not convexity-increasing. Then some internal angle changes from convex to reflex during the morph. By the above statement, the angle must be convex in Γ_1 and reflex in Γ_2 . \square

Alamdari et al. gave the following further condition that implies the hypothesis of Lemma 7. We emphasize that the statement applies to planar graph drawings in general, that is, edges are not required to be straight-line.

Observation 10. [2, in the proof of Lemma 13] *Let Γ be a planar graph drawing of a graph G in which all faces are y -monotone and let Γ' be another planar drawing of G that has the same combinatorial embedding, the same y -coordinates of vertices, and y -monotone edges. Then every line parallel to the x -axis crosses the same ordered sequence of edges and vertices in both drawings.*

Note that the property of y -monotone faces is a necessary condition—imagine a rectangle with a stalactite from the top and a stalagmite from the bottom such that a horizontal line through the middle of the rectangle intersects first the stalactite and then the stalagmite. This order of intersection can be switched without changing the combinatorial embedding nor the y -coordinates of the vertices.

The final ingredient we need in order to make use of the above lemmas is a way to redraw a graph to preserve the combinatorial embedding and the y -coordinates of the vertices, while improving convexity. We will again follow Alamdari et al. [2] and make use of a result of Hong and Nagamochi which is described in the next section.

2.4. Redrawing with convex faces while preserving y -coordinates

We build upon an $O(n^2)$ time algorithm due to Hong and Nagamochi [24] that redraws a given drawing with y -monotone faces such that all faces become convex while preserving the y -coordinates of the vertices. Angelini et al. [4] extended the result to strictly convex faces by perturbing vertices to avoid angles of π . They did not analyze the run-time. Both [24] and [4] expressed their results in terms of *level planar drawings of hierarchical-st plane graphs*. Their original statements and further explanations can be found in Appendix A.

In Section 4, we give a new proof of Hong and Nagamochi's result using Tutte's graph drawing method, which finds the vertex coordinates by solving a linear system. Tutte's original linear system gives rise to a symmetric matrix that can be solved quickly using the nested dissection method of Lipton et al. [29]. For our case, the matrix is not symmetric, and we need a recent generalization of nested dissection due to Alon and Yuster [3]. Our approach results in an improved running time of $O(n^{1.5})$ without, and $O(n^{1.1865})$ with, fast matrix multiplication:

Lemma 11. *Let Γ be a planar drawing of an internally 3-connected graph G such that every face is y -monotone. Let C be a strictly convex straight-line drawing of the outer face of G such that every vertex of C has the same y -coordinate as in Γ . Then there is a strictly convex straight-line drawing Γ' of G that has C as the outer face and such that every vertex of Γ' has the same y -coordinate as in Γ . Furthermore, Γ' can be found in time $O(n^{\omega/2})$, where ω is the matrix multiplication exponent.*

3. Computing convexity-increasing morphs

In this section we prove Theorem 1. In fact, we show multiple variants of Theorem 1, starting with a highly specialized version and proceeding to more and more general ones, which use the more specialized cases as building blocks.

3.1. A simple case: morphing y -monotone drawings

To give some intuition about our general proof strategy, we first consider an easy case where the outer face C of Γ is strictly convex and all faces are y -monotone. Then we can immediately apply Lemma 11 with the outer face fixed to obtain a new straight-line strictly convex drawing Γ' with all vertices at the same y -coordinates. By Observation 10, every line parallel to the x -axis crosses the same ordered sequence of edges and vertices in Γ and in Γ' . Then by Lemma 7 the morph from Γ to Γ' is planar. Also it is a horizontal morph. Thus we have a morph from Γ to a strictly convex drawing Γ' by way of a single horizontal morph. Furthermore, the morph is convexity-increasing by Lemma 9 since every internal angle is convex in Γ and Γ' .




Fig. 3. (a) A face that is not y -monotone. The dashed edges inside the face are added by Lemma 12. (b) The face after application of Lemma 13.

3.2. Morphing drawings with a convex outer face

We next consider the case of a planar straight-line drawing Γ in which the outer face is convex (but not necessarily strictly convex) and the inner faces are not necessarily y -monotone. Additionally, assume that Γ has no horizontal edge; we will later show how to ensure this.

Overview. As in Section 3.1, we begin by performing a horizontal morph. Observe that such a morph preserves the local extrema and does not change their convexity status. Thus the only reflex angles that can be made convex via a horizontal morph are the h -reflex angles, where an angle of inner face f is called h -reflex if it is reflex and occurs at a vertex that has one neighbor in f above and the other below—equivalently, the angle is reflex and is not a local extremum of f . We will show that a single horizontal morph suffices to convexify all the internal angles that are not local extrema. The plan is to then conceptually “turn the paper” by 90° and perform a vertical morph to make any v -reflex angle convex, where an angle of inner face f is called a v -reflex angle if it is reflex and occurs at a vertex that has one neighbor in f to the left and the other to the right. By continuing to alternate between the horizontal and the vertical direction, we eventually end up with a convex drawing and, thus, prove Theorem 1 for the case of a convex outer face.

A horizontal morphing step. To find the desired horizontal morph we will apply Lemma 11. Therefore, we must first augment Γ to a drawing with y -monotone faces by inserting y -monotone edges which are not necessarily straight-line. For an example see Fig. 3. This is a standard operation in upward planar (or “monotone”) drawing [16, Lemma 4.1] [32, Lemma 3.1], but we need the stronger property that the new edges are only incident to local extrema; otherwise we would relinquish control of convexity at that vertex. We will use:

Lemma 12. *Let Γ be straight-line planar drawing of an internally 3-connected graph. Then Γ can be augmented by adding edges to obtain a drawing Γ' such that each additional edge is a y -monotone curve joining two local extrema of some face in Γ , every inner face is y -monotone and the augmented graph is internally 3-connected. Furthermore, the additional edges can be found in $O(n \log n)$ time.*

Proof. Recall that a face is y -monotone if and only if it has no reflex local extrema. Our proof is by induction on the number of reflex local extrema in all inner faces of the drawing. If there are none, then all inner faces are y -monotone. Otherwise, consider an inner face f that has a local extremum u . For an illustration consider Fig. 3.

We want to find a local extremum v below u such that we can insert a y -monotone curve from v to u within face f . To construct the curve go vertically downwards from u to the first point p_u on the boundary of f , and then follow a y -monotone chain of f 's boundary downwards from p_u to a local minimum v . Adding the edge (u, v) divides f into two faces, and decreases the total number of local extrema. We note that Pach and Tóth [32] used a similar idea to triangulate with monotone curves, although their curves stopped at the first vertex on f 's boundary and we must continue to the first local minimum. Since we will only insert inner edges, the graph remains internally 3-connected.

To complete the proof we briefly describe how the set of augmenting edges can be found in time $O(n \log n)$. We deal with the local reflex minima; the maxima can be dealt with in a second phase. Find a trapezoidization of the drawing [15] in $O(n \log n)$ time. This gives the point p_u for each local reflex minimum u in face f . We can preprocess the graph in time $O(n)$ to find, for each edge e in face f , the local minimum v that is reached by following a y -monotone chain downward from e in f . This gives the set of augmenting edges. We must still find the cyclic order of augmenting edges incident with each vertex v . We separate into those that arrive at v from the left and those that arrive from the right. Within each of these sets, we sort by the y -coordinate of points p_u . \square

This observation allows us to prove the following:

Lemma 13. *Let Γ be a straight-line planar drawing of an internally 3-connected graph with a convex outer face and no horizontal edge. There exists a horizontal planar morph to a straight-line drawing Γ' with a strictly convex outer face and every internal angle*




Fig. 4. (a) A face that is not y -monotone. (b) The face after application of Lemma 13. There is a vertical edge and the single reflex vertex is not v -reflex. (c) After applying a horizontal shear transformation, the reflex vertex is v -reflex and there are no vertical edges.

that is not a local extremum is strictly convex in Γ' . Furthermore, the morph $\langle \Gamma, \Gamma' \rangle$ is convexity-increasing, and can be found in time $O(n^{\omega/2})$, where ω is the matrix multiplication exponent.

Proof. Use Lemma 12 to augment Γ with a set of edges A such that $\Gamma \cup A$ is a planar drawing in which all faces are y -monotone, and any edge of A goes between two local extrema in some inner face. This takes $O(n \log n)$ time. Let C be the outer face of Γ . Create a new drawing C' of C such that C' is strictly convex and preserves the y -coordinates of vertices.

By Lemma 11 with the outer face C' we obtain (in time $O(n^{\omega/2})$) a new straight-line strictly convex drawing $\Gamma' \cup A'$ with all vertices at the same y -coordinates as in Γ . (Here A' is a set of straight-line edges corresponding to A .) By Observation 10 every line parallel to the x -axis crosses the same ordered sequence of edges and vertices in $\Gamma \cup A$ and in $\Gamma' \cup A'$. Then by Lemma 7 the morph from Γ to Γ' is a planar horizontal morph.

Any internal angle of Γ that is not a local extremum has no edge of A incident to it, and thus becomes strictly convex in Γ' . Any internal angle of Γ that is a local extremum maintains its convex/reflex status in Γ' . Thus by Lemma 9 the morph is convexity-increasing. The run-time to find the morph (i.e., to find Γ') is $O(n^{\omega/2})$. \square

Making progress. Lemma 13 generalizes to directions other than the horizontal direction: for any direction d that is not parallel to an edge of Γ , there exists a convexity-increasing unidirectional (with respect to d) morph that convexifies all internal angles that are not extreme in the direction orthogonal to d . Thus, if we do not insist on a sequence of *horizontal and vertical* morphs, we immediately obtain a proof of Theorem 1 for the case of a convex outer face, since, for each reflex angle, we can choose a direction d that convexifies it.

In order to prove the stronger result that uses only two orthogonal directions, we need to alternate between the horizontal and the vertical direction. Note that after one application of Lemma 13, we do not necessarily obtain a drawing which contains a v -reflex vertex, see Fig. 4(a, b). In order to ensure that our algorithm makes progress after every step, we prove a strengthened version of Lemma 13, which ensures that there is at least one h -reflex or v -reflex vertex available after each step.

Lemma 14. Let Γ be a straight-line planar drawing of an internally 3-connected graph with a convex outer face and no horizontal edge. There exists a horizontal planar morph to a straight-line drawing Γ'' such that

- (i) the outer face of Γ'' is strictly convex,
- (ii) every internal angle that is not a local extremum is convex in Γ'' ,
- (iii) Γ'' has no vertical edge, and
- (iv) if Γ'' is not convex, then it has at least one v -reflex angle.

Furthermore, the morph is convexity-increasing, and can be found in time $O(n^{\omega/2})$.

Proof. We first apply Lemma 13 to obtain a morph from Γ to a drawing Γ' that satisfies (i) and (ii). If Γ' satisfies all the requirements, we are done. Otherwise we will achieve the properties (iii) and (iv) by shearing the drawing Γ' . Eliminating vertical edges via a horizontal shear is easy, so we concentrate on the requirement (iv) about v -reflex angles. Suppose Γ' is not convex and has no v -reflex angle. Consider any reflex angle of Γ' , say at vertex u in the inner face f . By property (ii), u must be a local extremum in Γ' ; otherwise it would be convex. We will apply a horizontal shear transformation to create a drawing Γ'' in which the angle at u becomes v -reflex, i.e., in which the x -coordinate of u is between the x -coordinates of its two neighbors in f . Furthermore, the shear should eliminate all vertical edges, e.g., see Fig. 4. The shear can be found in linear time.

Since shearing is an affine transformation, Γ'' has the same convex/reflex angles as Γ' . Thus Γ'' satisfies all the properties. The linear morph $\langle \Gamma', \Gamma'' \rangle$ is a planar horizontal morph that preserves the convex/reflex status of every angle. By Lemma 8 the morph $\langle \Gamma, \Gamma'' \rangle$ is a horizontal planar morph. By Lemma 9 it is convexity-increasing. The morph can be found in time $O(n^{\omega/2})$. \square

We are now ready to prove Theorem 1 for the case of a convex outer face.

Proof of Theorem 1 for the case of a convex outer face. If the given drawing Γ has a horizontal edge and/or if Γ has internal reflex angles such that none of them is h-reflex, then we use one vertical shear as in the proof of Lemma 14 to remedy this. Then, in the special case that there is no internal reflex angle, but there are angles of degree exactly 180° , we may apply Lemma 14 once to obtain the desired strictly convex drawing. Otherwise, there are internal reflex angles; and we apply Lemma 14 alternately in the horizontal and vertical directions until the drawing is strictly convex. In each step there is at least one h-reflex or v-reflex angle that becomes convex. Thus, the number of morphing steps is at most $\max\{2, r + 1\} \leq n$, where r is the number of inner reflex angles in Γ . The resulting total run-time is $O(n^{1+\omega/2})$. \square

3.3. Morphing drawings of 3-connected graphs

In this section we prove the case of Theorem 1 where the outer face of Γ is not convex. However, we will assume that the given graph G is 3-connected (instead of only *internally* 3-connected).

Overview. On a high level, our approach works as follows: We first augment the outer face of Γ with edges from its convex hull to obtain a drawing of an augmented graph with a convex outer face. We then apply the results from Section 3.2 to morph to a strictly convex drawing and then remove the extra edges on the outer face one-by-one. After each removal of an edge, we morph to a strictly convex drawing of the reduced graph using at most three horizontal or vertical morphs.

Augmenting the outer face. Compute the convex hull of Γ . Any segment of the convex hull that does not correspond to an edge of G becomes a new edge that we add to G . Let A denote the new edges and $G \cup A$ denote the augmented graph with straight-line planar drawing $\Gamma \cup \Gamma_A$. Note that adding edges maintains 3-connectivity. Each edge $e \in A$ is part of the boundary of an inner face f_e of $\Gamma \cup \Gamma_A$. We call f_e the *pocket* of e . We apply the result of Section 3.2 to obtain a strictly convex drawing of $G \cup A$, see Fig. 5(a). Note that the techniques used in that section give us a drawing with no horizontal or vertical edges.

We remark that this step is the reason why we limit ourselves to 3-connected graphs in this subsection: Adding the convex hull edges in a drawing of an *internally* 3-connected graph may create non-external separation pairs, see Fig. 6(a). This would prevent us from using the algorithm from Section 3.2 as this algorithm uses Lemma 11 which requires the input graph to be internally 3-connected.

Popping a pocket outward. In this final step, we describe a way to remove an edge of A and “pop” out the vertices of its pocket so that they become part of the convex hull. Lemma 11 serves once again as an important subroutine. We make ample use of the fact that we may freely specify the desired subdrawing of the outer face after each application of Lemma 11, as long as we maintain either the x -coordinates or the y -coordinates of all vertices.

We remark that the following lemma applies to *internally* 3-connected graphs, not only 3-connected graphs. We plan to use it in the following section, in which we prove Theorem 1 in its general form. In fact, the final algorithm will use the entire procedure described in this section as a subroutine (we will augment the internally 3-connected graph such that adding the convex hull edges does not create non-external separation pairs).

Lemma 15. *Let Γ be a strictly convex drawing of an internally 3-connected graph G without vertical edges and let e be an edge on the outer face. If $G - e$ is internally 3-connected, then $\Gamma - e$ can be morphed to a strictly convex drawing of $G - e$ without vertical edges via at most three convexity-increasing morphs, each of which is horizontal or vertical. Furthermore, the morphs can be found in time $O(n^{\omega/2})$.*

Proof. Our morph will be specified by a sequence of drawings, $\Gamma, \Gamma_1, \Gamma_2, \Gamma_3$, where the first and the last morph are vertical and the second morph, which we can sometimes skip, is horizontal.

Let $e = (u, v)$. We first perform a vertical morph from Γ to a strictly convex drawing Γ_1 in which vertex u is top-most or bottom-most and which does not contain vertical or horizontal edges. Since Γ is strictly convex and has no vertical edges, it is also x -monotone. Therefore, the desired drawing Γ_1 can be found by choosing some strictly convex drawing of the outer face in which u is extreme while maintaining the x -coordinates of all vertices, and then using one application of Lemma 11 (for vertical morphs). Additionally, we may need to apply a vertical shearing transformation in order to get rid of horizontal edges. This is easily done while still guaranteeing that u is extreme in the y -direction. Analogous to Section 3.1, by combining Observation 10, Lemma 7 and Lemma 9 we conclude that the horizontal morph $\langle \Gamma, \Gamma_1 \rangle$ is planar and convexity-increasing.

For the remainder of the proof, assume without loss of generality that u is the top-most vertex and that v lies to the right of u in Γ_1 . The other cases are symmetric. Let p_{uv} denote the path from u to v in $f_e - e$. We distinguish two cases depending on the shape of f_e in Γ_1 .

Case 1: The path p_{uv} is x -monotone in Γ_1 , see Fig. 5(b). In this case we can skip the second step of the morph sequence. We will remove e and compute a vertical morph from $\Gamma_1 - e$ to a strictly convex drawing Γ_3 of $G - e$ without vertical edges. Once again, this can be done by combining Lemma 11 (for vertical morphs), Observation 10, Lemma 7 and Lemma 9




Fig. 5. (a) Schematic of the convex drawing of $G \cup A$. Graph G is depicted in gray, edges of A are dashed, and the pockets are white. (b)–(c) Cases 1 and 2 for Lemma 15, where faint gray arrows indicate explicit placements on the convex hull.

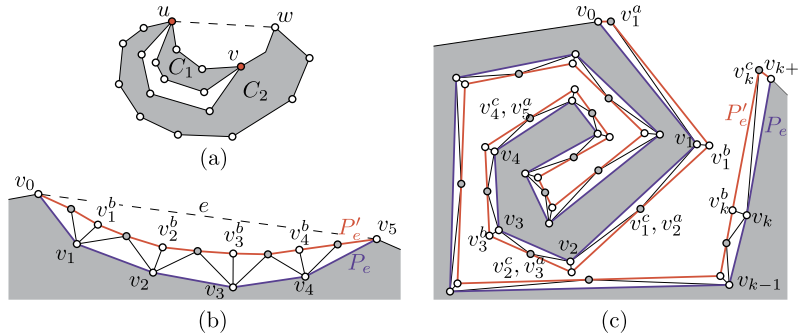


Fig. 6. (a) An internally 3-connected graph with an external separation pair $\{u, v\}$. Adding the convex hull edge uw turns $\{u, v\}$ into a non-external separation pair both since v becomes an internal vertex and since C_1 no longer has a vertex on the outer face. (b) Schematic drawing of the path P'_e created for the pocket defined by the convex hull edge $e = \{v_0, v_5\}$. (c) Geometrically, we embed P'_e very close to P_e . This is possible regardless of the shape of the pocket.

as long as we can specify a strictly convex drawing of the outer face of $\Gamma_1 - e$ in which the x -coordinates match those of Γ_1 . It suffices to compute a suitable new reflex chain for p_{uv} , see Fig. 5(b).

Case 2: The path p_{uv} is not x -monotone. In this case we will compute a horizontal morph from Γ_1 to a strictly convex drawing Γ_2 in which u and v are the unique left-most and the unique right-most vertices and which does not contain a vertical edge. Once again, this can be done by combining Lemma 11, a horizontal shearing transformation, Observation 10, Lemma 7 and Lemma 9 as long as we can specify a strictly convex drawing of the outer face of Γ_1 in which u and v are the left-most and right-most vertices and the y -coordinates match those of Γ_1 . This is possible because u is top-most, see Fig. 5(c).

In the drawing Γ_2 the pocket f_e is convex with extreme points u and v so the path p_{uv} is x -monotone and, hence, by Case 1, there is a vertical morph from $\Gamma_2 - e$ to a strictly convex drawing Γ_3 of $G - e$. \square

Observe that each application of Lemma 15 increases the number of vertices of G on the convex hull. Thus, the proof of Theorem 1 follows by induction on the number of convex hull vertices. Let ρ denote the number of pockets of the strictly convex drawing of $G \cup A$ obtained by applying the algorithm from Section 3.2. We observe that $\rho \leq n/2$, since each pocket can be associated with two private vertices of G : the clockwise first of the two convex hull vertices defining the pocket and its clockwise successor, which is private to the pocket.

We use at most three horizontal and vertical morphs to pop out a pocket. Hence, the number of morphs needed to deal with all the pockets can be bounded by $2\rho + 1$ by observing that each application of Lemma 15 involves a vertical-horizontal-vertical morph sequence, and Lemma 8 allows us to compress two consecutive vertical morphs into one. Additionally, obtaining the strictly convex drawing of $G \cup A$ requires at most $\max\{2, r + 1\}$ morphing steps, where r is the number of internal reflex angles in the drawing $\Gamma \cup \Gamma_A$. This number can be bounded by $\max\{2, r + 1\} \leq n - \rho + 1$ since each pocket can be associated with a private vertex of the convex hull, that cannot have an internal reflex angle. Hence, the total number of morphs is at most $n - \rho + 1 + 2\rho + 1 = n + \rho + 2 \leq 1.5n + 2$, where the last inequality uses the fact that $\rho \leq n/2$. The run time of the algorithm is $O(n^{1+\omega/2})$.

3.4. Morphing drawings of internally 3-connected graphs

So far, if the given drawing does not have a convex outer face, we have restricted our attention to the class of 3-connected graphs. The only reason why the approach described in Section 3.3 is not able to handle the case of internally 3-connected graphs is that the addition of the convex hull edges A may create non-external separation pairs, see Fig. 6(a). As a result,

the augmented drawing $\Gamma_G \cup \Gamma_A$ is no longer a valid input for Lemma 11. In this section we extend our algorithm such that it is able to convexify drawings of internally 3-connected graphs and, thus, we conclude the proof of Theorem 1 in its general form.

Overview. Let Γ be a straight-line planar drawing of an internally 3-connected graph $G = (V, E)$. As a first step, we augment the outer face of Γ by adding new edges and vertices near each pocket. The goal of this step is to ensure that we can add convex hull edges without introducing non-external separation pairs. We then apply the algorithm from Section 3.3, which results in a drawing in which all the new vertices appear on the strictly convex outer face. Finally, we remove the new vertices one-by-one; gradually turning the strictly convex drawing of the augmented graph into a strictly convex drawing of G .

In slightly more detail, our method involves the following steps:

Step 1: Augmenting the outer face. We will augment the graph G to G' , and then augment the drawing Γ to Γ' , which involves geometric arguments. The goal is to ensure that G' is internally 3-connected and has the additional property that adding convex hull edges to Γ' does not introduce non-external separation pairs, as, for example, in Fig. 6(a). We will show that bad cases only arise when a pocket has a vertex that is part of a separating pair. Thus, our method will be to add an extra “buffer” layer of vertices to each pocket boundary, while ensuring that each buffer vertex is not part of a separating pair.

Step 2: Convexifying the augmented drawing. The goal of this step is to find a convexity-increasing morph from the drawing Γ' of the augmented graph G' to a strictly convex drawing of G' . In principle, the idea for accomplishing this task is very simple: we just apply the algorithm described in Section 3.3 to Γ' . The challenging aspect is, that this algorithm uses Lemma 11 as a subroutine and, thus, we need to ensure internal 3-connectivity of the input graph.

Step 3: Removing the additional vertices. At this point we have a strictly convex drawing of G' . We must now “reverse” the augmentation process, removing vertices of G' to get back to G . After each vertex is removed we will morph to obtain a strictly convex drawing. As in Step 2, the graph must remain internally 3-connected at each step. We will therefore treat the augmentation process of Step 1 as an iterative process, adding vertices one-by-one and ensuring that the graph is internally 3-connected at each step.

We now give further details on these three steps.

Step 1: Augmenting the outer face. We compute the convex hull of Γ . Let e be a convex hull edge with $e \notin E$. The following steps are illustrated in Fig. 6(b). Let $P_e = (v_0, v_1, \dots, v_{k+1})$ be the unique path on the outer face of Γ such that $P_e + e$ is a cycle with $G - P_e$ in its exterior. We introduce $2k + 1$ new vertices that form a path P'_e connecting v_0 and v_{k+1} :

$$P'_e = (v_0, v_1^a, v_1^b, v_1^c, v_2^a, v_2^b, v_2^c, v_3^a, \dots, v_{k-1}^c, v_k^a, v_k^b, v_k^c, v_{k+1})$$

Note that every inner vertex v_i of P_e gets a “private copy” v_i^b in P'_e . Two consecutive copies v_i^b and v_{i+1}^b are connected via another vertex which is equipped with two labels $v_i^c = v_{i+1}^a$, which will simplify the notation later on. Additionally, we add the edges $\{v_i, v_i^a\}$, $\{v_i, v_i^b\}$ and $\{v_i, v_i^c\}$ for $i = 1, \dots, k$.

Geometrically, the new path P'_e is embedded in a planar fashion very close to P_e , see Fig. 6(c). This can be accomplished regardless of the shape of P_e : Assume that $k \geq 2$ and let ε be the smallest distance between any pair of disjoint edges on the cycle $P_e + e$. For $i = 1, \dots, k$, we place v_i^b on the angular bisector of the outer angle at v_i such that its distance to v_i is smaller than $\varepsilon/2$. If $i > 1$, the vertex v_i^a is placed in the center of the line-segment $v_{i-1}^b v_i^b$. Similarly, if $i < k$, the vertex v_i^c is placed in the center of the line-segment $v_i^b v_{i+1}^b$.

We make sure that during this procedure no vertex is placed in the exterior of the convex hull of Γ ; with the following exceptions: the vertices v_1^a and v_k^c play a special role and are placed close to v_0 and v_{k+1} , respectively, such that they do appear on the convex hull of the augmented drawing.

Note that for the special case of small pockets with $k = 1$, the value ε would not be well-defined, as there are no disjoint edges on $P_e + e$. However, in this case it is easy to directly compute an embedding of P'_e with the desired properties (i.e. planarity; and only v_1^a and v_k^c appear on the convex hull).

We repeat the process for all convex hull edges $e \notin E$ of Γ and use $G' = (V', E')$ and Γ' to denote the resulting plane graph and drawing respectively. The total run-time for this step is dominated by the time to compute the values ε , which can be done in $O(n^2)$ total time. We remark that this step could be implemented more efficiently, for example by computing a medial axis [10]. However, since our total runtime is $\omega(n^2)$, we refrain from stating the details.

We will now prove that G' is internally 3-connected. Keeping in mind our plan for Step 3, we will add the vertices of G' one-by-one, showing that each addition preserves the property of being internally 3-connected. We begin with two basic operations that preserve internal 3-connectivity.

Lemma 16. *Let H be an internally 3-connected graph with an edge (a, b) on the outer face. Construct H' by adding a new vertex x in the outer face connected to a and b . Then H' is internally 3-connected.*

Proof. By internal 3-connectivity, Property (I3) of Lemma 5 holds for H . Thus, it is clear that H' also satisfies Property (I3) and, hence, H' is internally 3-connected. \square

Lemma 17. *Let H be an internally 3-connected graph with two consecutive edges (a, b) and (b, c) on the outer face. Construct H' by adding a new vertex x in the outer face connected to a, b and c . Then H' is internally 3-connected.*

Proof. By internal 3-connectivity, Property (I3) of Lemma 5 holds for H . It suffices to show that Property (I3) also holds for H' . Let f_0 and f'_0 denote the outer faces of H and H' , respectively. Clearly, b has three paths to f'_0 that are disjoint except for b . So let $v \neq b$ be some internal vertex of H' and note that v is also internal in H . Hence, by Property (I3) of H , v has three paths to f_0 that are disjoint except for v . At most one of these paths does not end at f'_0 , namely if its endpoint on f_0 is b . However, appending the edge $\{b, x\}$ to this path yields the desired three paths from v to f'_0 that are disjoint except for v . Hence, Property (I3) holds for H' . \square

With these operations in hand, we can show that G' is internally 3-connected, and—more strongly—that we can build G' by adding one vertex at a time, preserving internal 3-connectivity. Let $V' = V \cup V^b \cup V^{ac}$ where V^b is the set of all vertices whose upper index is b , and where V^{ac} is the set of the remaining vertices (whose upper index is a and/or c).

Lemma 18. *Starting with G and adding the vertices of V^{ac} one-by-one in any order and then the vertices of V^b one-by-one in any order produces an internally 3-connected graph at each step.*

Proof. The addition of each vertex of V^{ac} maintains internal 3-connectivity by Lemma 16. The addition of each vertex of V^b maintains internal 3-connectivity by Lemma 17. \square

Step 2: Convexifying the augmented drawing. As mentioned above, the plan is to apply the algorithm described in Section 3.3 to Γ' . That algorithm adds the convex hull edges of Γ' and then iteratively removes these edges while performing some morphing steps before and after each removal. Each morphing step requires one application of Lemma 11. Therefore, in order to prove the correctness of Step 2, we need to ensure that before and after each removal of a convex hull edge, the input graph is internally 3-connected.

We begin by observing that the new vertices of G' are not part of separation pairs. Then we show that adding a convex hull edge is safe when none of the vertices of its pocket are in separating pairs.

Observation 19. *No vertex of $V' \setminus V$ is in a separating pair of G' .*

Proof. By construction, removing any one of these vertices can not introduce a cut vertex. \square

We now give one more operation that preserves internal 3-connectivity.

Lemma 20. *Let H be an internally 3-connected graph, with vertices a and b on the outer face. Let P be one of the paths from a to b along the outer face. Assume that no vertex of P is part of a separating pair in H . Let H' be the result of adding the edge (a, b) embedded such that P becomes internal. Then H' is internally 3-connected.*

Proof. Suppose that H' has a separation pair $\{u, v\}$. This pair is also separating in H . Further, since H is internally 3-connected, $\{u, v\}$ is an external separation pair of H . Since both u and v do not belong to P , Observation 6 implies that a and b belong to the same component of $(H - u - v)$. Thus, it is easy to verify that adding the edge between a and b maintains the six conditions of Observation 6 (in particular, Property (E2) holds as a and b remain on the outer face). Therefore, $\{u, v\}$ is an external separation pair of H' . \square

Recall that the construction of Γ' ensures that all convex hull edges e' which do not correspond to edges of G' have the form $e' = \{v_1^a, v_k^c\}$, where v_1^a, v_k^c are the first and last internal vertex of one of the paths P_e' . By Observation 19 the interior vertices of P_e' form a path P on the outer face which does not contain any vertices that are part of a separation pair. Further, adding the edge e' encloses P in an internal face. Thus, we obtain:

Corollary 21. *Let A denote the set of convex hull edges of Γ' which do not correspond to edges of G' . Then, for any $S \subseteq A$ the plane graph $G' + S$ is internally 3-connected.*




Fig. 7. (a) In Step 3, we iteratively remove the vertices v_i^b causing their counterparts v_i to become part of the outer face. (b) The regions corresponding to the Cases 1–4.

Step 3: Removing the additional vertices. At this point we have a strictly convex drawing of G' and want to convert it to a strictly convex drawing of G . We will remove the vertices of V^b iteratively; one-by-one. After each such removal, we will perform up to two morphing steps in order to recover a strictly convex drawing of the reduced graph. Recall that by Lemma 18, all the intermediary graphs are internally 3-connected and, thus, they are valid inputs for Lemma 11.

Lemma 22. Let $B \subseteq V^b$ and let Γ'_B be a strictly convex drawing of $G' - B$ without vertical or horizontal edges. Further, let $v_i^b \in V^b \setminus B$. Then, there is a convexity-increasing morph from $\Gamma'_B - v_i^b$ to a strictly convex drawing Γ'' of $G' - B - v_i^b$ without vertical or horizontal edges. Moreover, there is such a morph which consists of a sequence of up to 2 horizontal / vertical morphs. The morphing sequence can be found in $O(n^{\omega/2})$ time.

Proof. Without loss of generality, we may assume that v_i^c is located to the bottom-right of v_i^a and that v_i^b is located to the right of the oriented line $\overrightarrow{v_i^c v_i^a}$, see Fig. 7(a). We distinguish four cases regarding the position of the vertex v_i , for an illustration see Fig. 7(b).

Case 1: v_i is located to the right of the oriented line $\overrightarrow{v_i^c v_i^a}$. We observe that the subdrawing $\Gamma'_B - v_i^b$ is already a convex drawing of $G' - B - v_i^b$, so there is nothing to show.

Case 2: $y(v_i^a) > y(v_i) > y(v_i^c)$; and we are not in Case 1. Let C be a strictly convex drawing of the outer face of $G' - B - v_i^b$, such that every vertex in C has the same y -coordinate as in $\Gamma'_B - v_i^b$ (we can easily find such a drawing C by adding v_i to the convex hull of $\Gamma'_B - v_i^b$ in a strictly convex fashion). Then, Lemma 11 applied to $\Gamma'_B - v_i^b$ and C (potentially followed by a horizontal shearing transformation in order to remove vertical edges), yields the desired drawing Γ'' . Analogous to Section 3.1, by combining Observation 10, Lemma 7 and Lemma 9 we conclude that the horizontal morph $\langle \Gamma'_B - v_i^b, \Gamma'' \rangle$ is planar and convexity-increasing.

Case 3: $x(v_i^a) < x(v_i) < x(v_i^c)$; and we are not in Case 1 or Case 2. Let C be a strictly convex drawing of the outer face of $G' - B - v_i^b$, such that every vertex in C has the same x -coordinate as in $\Gamma'_B - v_i^b$ (we can easily find such a drawing C by adding v_i to the convex hull of $\Gamma'_B - v_i^b$ in a strictly convex fashion). Then, Lemma 11 (for vertical morphs) (potentially followed by a vertical shearing transformation in order to remove horizontal edges) applied to $\Gamma'_B - v_i^b$ and C yields the desired drawing Γ'' . Analogous to Section 3.1, by combining Observation 10, Lemma 7 and Lemma 9 we conclude that the vertical morph $\langle \Gamma'_B - v_i^b, \Gamma'' \rangle$ is planar and convexity-increasing.

Case 4: we are not in Case 1, Case 2 or Case 3. We reduce to Case 2 or Case 3: using a shearing transformation along the x -axis (or along the y -axis), we obtain a drawing Γ''_B of $G' - B$ satisfying the preconditions of Case 3 (or Case 2). Analogous to Section 3.1, by combining Observation 10, Lemma 7 and Lemma 9 we conclude that the linear morph $\langle \Gamma'_B, \Gamma''_B \rangle$ is planar and convexity-increasing. \square

Starting with the drawing Γ' of G' and iterating Lemma 22, we obtain a strictly convex drawing of $G' - V^b$. By construction, we can simply remove all the vertices of V^{ac} to obtain a strictly convex drawing of G .

Observation 23. Let Γ'_{ac} be a strictly convex drawing of $G' - V^b$. Then, $\Gamma'_{ac} - V^{ac}$ is a strictly convex drawing of G .

We summarize:

Proof of Theorem 1. We analyze the three steps of the algorithm individually.

Step 1: We begin by augmenting G and Γ to G' and Γ' . As discussed in the corresponding section, this can be done in $O(n^2)$ time.

Step 2: Next, we apply the algorithm from Section 3.3 to G' and Γ' . This algorithm was designed for 3-connected graphs. However, note that Lemma 15 applies to internally 3-connected graphs as well). Therefore, the algorithm also works for internally 3-connected graphs as long as adding convex hull edges and then successively removing them never creates a plane graph which is not internally 3-connected. This is the case by Corollary 21. Thereby, we obtain in $O(n'^{1+\omega/2}) \subseteq O(n^{1+\omega/2})$ time a convexity-increasing morph from Γ' to a strictly convex drawing of G' , where n' is the number of vertices of the augmented graph G' . This morph is also convexity-increasing with respect to the subdrawing of G as every internal angle of Γ is also internal in Γ' .

As discussed in the last paragraph of Section 3.3, the upper bound on the number of morphing steps guaranteed by the algorithm is $\max\{2, r' + 1\} + 2\rho' + 1$, which can be bounded by $1.5n' + 2$. Here, ρ' denotes the number of pockets of Γ' . In fact, the bound can be improved to $1.5n + 2 < 1.5n' + 2$ by observing that $r = r'$ and $\rho = \rho'$, where n and ρ denote the number of vertices and pockets, respectively, of the original drawing Γ . The latter equality is obvious. For the former equality, observe that each vertex $v_i^a \in V^{ac}$, $2 \leq i \leq k$ has an angle of π and, hence, it has no reflex angle. Further, the construction of Γ' ensures that an outer angle at a vertex $v_i^b \in V^b$ is reflex if and only if the outer angle at the corresponding vertex v_i is reflex in Γ . Other angles at v_i^b can not be reflex. Moreover, if the outer angle at v_i is reflex in Γ then v_i has no reflex angle in Γ' . Consequently, we can charge the inner reflex angles of the vertices v_i^b to their counterparts v_i . Finally, the vertices $v_0, v_1^a, v_k^c, v_{k+1}$ of each pocket P_e' belong to the convex hull of Γ' and, thus, they do not have any internal reflex angles. The convexity status of the remaining angles is untouched and, hence, $r = r'$ as claimed. Altogether, we obtain the improved bound

$$\max\{2, r' + 1\} + 2\rho' + 1 = \max\{2, r + 1\} + 2\rho + 1 \leq 1.5n + 2,$$

where the last inequality was already discussed in the last paragraph of Section 3.3.

Step 3: Finally, we iteratively apply Lemma 22 to the strictly convex drawing of G' that was obtained in the previous step. Each application increases the number of vertices of G on the convex hull. Thus, by induction we arrive at a strictly convex drawing of G . Each of the intermediary morph steps is convexity-increasing with respect to the respective augmented graph. Once again, since every internal angle of Γ remains internal in (all) the augmented graph(s), we have that the morphing sequence is convexity-increasing for G as well. The number of morphing steps is bounded by $2n$ and the time required to obtain the entire sequence sums up to $O(n'^{1+\omega/2}) \subseteq O(n^{1+\omega/2})$.

Summing up, we end up with $3.5n + 2$ morphing steps and a runtime of $O(n^{1+\omega/2})$. \square

4. Using Tutte's algorithm to find convex drawings preserving y-coordinates

In this section we prove Lemma 11 using Tutte's graph drawing algorithm. This reduces the problem to solving a linear system. Applying a generalized method due to Alon and Yuster [3] for solving a linear system whose non-zero's in the matrix correspond to a planar graph, we obtain an algorithm that produces a straight-line strictly convex redrawing preserving y-coordinates and that runs in $O(n^{\omega/2})$ time, where ω is the matrix multiplication exponent.

In his paper, "How to Draw a Graph," [41] Tutte showed that any 3-connected planar graph $G = (V, E)$ with a fixed convex drawing C of its outer face has a convex drawing with outer face C that can be obtained by solving a system of linear equations. For each $v \in V$ let the variables (x_v, y_v) represent the coordinates of vertex v . Let V_I be the internal vertices of G and let V_B be the vertices of the outer face. For each vertex $v \in V_B$ let (x_v^b, y_v^b) be its (fixed) coordinates in C . Let d_v be the degree of vertex v . Consider the system of equations:

$$\begin{aligned} \forall u \in V_I \quad (x_u, y_u) &= \sum_{(u,v) \in E} \frac{1}{d_u} (x_v, y_v), \\ \forall u \in V_B \quad (x_u, y_u) &= (x_u^b, y_u^b). \end{aligned} \tag{1}$$

Tutte proved that this system of equations has a unique solution and that the solution gives a convex drawing of G with outer face C . In fact, the drawing is even strictly convex if C is strictly convex.

Tutte's Theorem was originally stated for 3-connected graphs. However, it is well known that Tutte's proof also applies to the more general class of internally 3-connected graphs since it only uses Property (I3) of Lemma 5. For the special case of strictly convex drawings of the outer face, the generalization from 3-connected to internally 3-connected graphs is easy to prove: Suppose graph G has an external separation pair (u, v) . We make use of Observation 6. Vertices u and v lie on the outer face, and their removal separates the graph into two connected components D and D' . In the strictly convex drawing C of the outer face, a line segment joining u to v will separate C into two strictly convex subpolygons since by Property (E2) of Observation 6 the vertices u and v are not consecutive on C . By induction, Tutte's algorithm will draw each of the two components with strictly convex faces in the appropriate subpolygon, and—in case (u, v) is not an edge of the

graph—the face between D and D' will also be strictly convex. It is not necessary to apply Tutte’s algorithm separately to the two components—one system of equations will do.

Tutte’s theorem can be generalized to more general “barycenter” weights other than $1/d_u$. Assign a weight $w_{u,v} > 0$ to each ordered pair (u, v) with $(u, v) \in E$ such that for each u it holds that $\sum_v w_{u,v} = 1$. We emphasize that $w_{u,v}$ may be different from $w_{v,u}$. Consider the system of equations:

$$\begin{aligned} \forall u \in V_I \quad (x_u, y_u) &= \sum_{(u,v) \in E} w_{u,v}(x_v, y_v), \\ \forall u \in V_B \quad (x_u, y_u) &= (x_u^b, y_u^b). \end{aligned} \tag{2}$$

This system also has a unique solution that gives a convex drawing of G with outer face C , and a strictly convex drawing of G if C is strictly convex. This generalization was first stated by Floater in 1997 [19] for triangulations and one year later [20] for general 3-connected planar graphs, though the result is not stated as a theorem in either case. Floater proved that the constraint matrix is non-singular, and, for the rest, said that Tutte’s proof² carries over.

An explicit statement that the linear system (2) has a unique solution that gives a strictly convex drawing of G if C is strictly convex is due to Gortler, Gotsman, Thurston in 2006 [22, Theorem 4.1]. They give a new proof using “one-forms”.

We can now give an alternate proof of Hong and Nagamochi’s result:

Proof of Lemma 11. We must show that there is a strictly convex drawing of G with outer face C that preserves the y -coordinates of the vertices from drawing Γ . Our idea is to do this in two steps, first choosing the barycenter weights to force vertices to lie at the required y -coordinates, and then solving system (2) with these barycenter weights to determine the x -coordinates.

For the first step, we solve the following system separately for each $u \in V_I$:

$$y_u = \sum_{(u,v) \in E} w_{u,v}y_v, \quad 1 = \sum_v w_{u,v} \tag{3}$$

Here the y ’s are the known values from Γ and the $w_{u,v}$ ’s are variables. There are two equations and $d_u > 2$ variables, so the system is under-determined and can easily be solved: Because Γ has y -monotone faces, every vertex $u \in V_I$ has neighbors below and above. Let N_u^+ be the neighbors of u that lie above u in Γ . Let $d_u^+ = |N_u^+|$. Similarly, let N_u^- be the neighbors of u that lie below u in Γ and let $d_u^- = |N_u^-|$. Compute the average y -coordinate of u ’s neighbors above and below:

$$y_u^+ = \sum_{v \in N_u^+} \frac{1}{d_u^+} y_v \quad y_u^- = \sum_{v \in N_u^-} \frac{1}{d_u^-} y_v$$

Observe that y_u lies between y_u^+ and y_u^- . Thus we can find a value t_u , $0 < t_u < 1$, such that

$$\begin{aligned} y_u &= t_u y_u^+ + (1 - t_u) y_u^- \\ &= \sum_{v \in N_u^+} \frac{t_u}{d_u^+} y_v + \sum_{v \in N_u^-} \frac{1 - t_u}{d_u^-} y_v \end{aligned}$$

Therefore, setting $w_{u,v} = \frac{t_u}{d_u^+}$ for $v \in N_u^+$ and $w_{u,v} = \frac{1-t_u}{d_u^-}$ for $v \in N_u^-$, yields a solution to (3). Observe that $w_{u,v} > 0$ for all $(u, v) \in E$.

Given values $w_{u,v} > 0$ satisfying the constraints (3) for all $u \in V_I$, we then solve equations (2) to find values for the x_u ’s. By Tutte’s generalized result, this provides a strictly convex drawing of G with outer face C while preserving y -coordinates.

It remains to discuss how to obtain the claimed run-time. Recall that we assume a real RAM model of computation—in particular, each arithmetic operation takes unit time. Observe that solving the system (3) to find the appropriate weights $w_{u,v}$ based on the y -coordinates takes linear time. The significant aspect is solving Tutte’s generalized system of equations (2).

Tutte’s method gives rise to a linear system $Ax = b$ where A is a matrix with a row and column for each vertex, and where the non-zeros in the matrix correspond to edges in the planar graph. In more detail, Tutte’s original linear system (1) can be re-written as

$$\forall u \in V_I \quad d_u(x_u, y_u) = \sum_{(u,v) \in E} (x_v, y_v),$$

² Colin de Verdiere et al. [12] point out that Tutte’s original proof is complicated because Tutte is also re-proving Kuratowski’s theorem, and they recommend the simpler proof by Richter-Gebert [37].




Fig. 8. The drawings (a) $\Gamma^-(7)$, (b) $\Gamma^\Delta(7)$, and (c) $\Gamma^\Delta(14)$.

so the non-zeros in the constraint matrix, A , consist of entries $-d_u$ down the main diagonal, and $a_{u,v} = a_{v,u} = 1$ if (u, v) is an edge. The equations for vertices in V_B give an extra part of the constraint matrix that is just an identity matrix. The crucial property is that the constraint matrix is symmetric. In fact, symmetry holds more generally if we choose weights or “stresses” $s_{u,v} = s_{v,u}$ and define $w_{u,v} = s_{u,v} / \sum_{(u,z) \in E} s_{u,z}$. Tutte’s original theorem is the special case where $s_{u,w} = 1$ for all (u, v) .

When the constraint matrix A is symmetric with non-zero’s corresponding to the edges of a planar graph, the system $Ax = b$ can be solved in $O(n^{\omega/2})$ arithmetic operations using the generalized nested dissection method of Lipton, Rose, Tarjan [29,30]. The fact that this applies to Tutte’s original algorithm is well-known in graph drawing, see for example [11, 36]. However, nested dissection does not apply when the matrix A is not symmetric, so in particular, it does not apply to the linear system (2).³ For the more general case of an asymmetric matrix A we need the following result of Alon and Yuster from 2013 [3]. They consider a linear system $Ax = b$ where A has a row and a column for each vertex of an associated graph G and there is an edge (u, v) in G if and only if $a_{u,v} \neq 0$ or $a_{v,u} \neq 0$ (the diagonal entries of A play no role in the definition of G).

Theorem 24 (Theorem 1.1 in [3], specialized to \mathbb{Q} and to planar graphs). *Let $A \in \mathbb{Q}^{n \times n}$ be a nonsingular matrix and let $b \in \mathbb{Q}^n$. If the graph associated with A is planar, then $Ax = b$ can be solved in $O(n^{\omega/2})$ time.*

They assume that each arithmetic operation takes unit time, i.e., that algorithms are measured in terms of their algebraic complexity. Moreover, the matrix is assumed to be given in an implicit form. They also assumed that $\omega > 2$; otherwise the run-time is in $O(n \log n)$.

In terms of practicality, they note that although the fastest matrix multiplication algorithms are only theoretical, using naive matrix multiplication gives a run time of $O(n^{1.5})$ and all ingredients of the algorithm become practically implementable, and Strassen’s algorithm is sensible for larger n and gives a run time of $O(n^{1.41})$.

Note that in our case the matrix A is non-singular because of Tutte’s generalized result. This completes the proof of Lemma 11. \square

5. Lower bound on the number of morphing steps

In this section, we prove Theorem 3, that is, we show a linear lower bound on the number of required morphing steps. The proof builds on the following result by Alamdari et al. [2].

Theorem 25 ([2]). *There exist two straight-line planar drawings $\Gamma^\Delta(n')$ and $\Gamma^-(n')$ of a path with n' vertices such that any planar morph between them which is composed of a sequence of linear morphing steps requires $\Omega(n')$ steps.*

In the drawing $\Gamma^-(n')$ the n' vertices $a_1, \dots, a_{n'}$ are placed on a horizontal line with a_i to the left of a_{i+1} for $1 \leq i < n'$, see Fig. 8(a). In the drawing $\Gamma^\Delta(n')$ the path forms a spiral, see Fig. 8(b). More precisely, let e_i denote the edge $\{a_i, a_{i+1}\}$. Then for each i with $i \bmod 3 \equiv 1$, the edge e_i is horizontal and a_i is to the left of a_{i+1} . For each i with $i \bmod 3 \equiv 2$, the edge e_i is parallel to the line $y = \tan(2\pi/3)x$ and a_i is to the right of a_{i+1} . Finally, for each i with $i \bmod 3 \equiv 0$, the edge e_i is parallel to the line $y = \tan(-2\pi/3)x$ and a_i is to the right of a_{i+1} .

In order to prove Theorem 3, we present a drawing $\Gamma^\Delta(n)$ of a cycle on n vertices which contains a subdrawing of $\Gamma^\Delta(n')$ for some $n' \in \Theta(n)$, see Fig. 8(c). The existence of a convexifying planar morph for $\Gamma^\Delta(n)$ with $o(n)$ linear morphing steps would imply the existence of a planar morph $\Gamma^\Delta(n')$ and $\Gamma^-(n')$ with $o(n')$ linear morphing steps, contradicting Theorem 25.

Proof of Theorem 3. Let $n' = \lfloor n/2 \rfloor$. Let $\Gamma^\Delta(n)$ be some planar straight-line drawing of the cycle $C = (a_1, \dots, a_{n'}, b_1, \dots, b_{n-n'}, a_{n-n'}, \dots, a_1)$ such that the induced subdrawing of the path $P = (a_1, \dots, a_{n'})$ is $\Gamma^\Delta(n')$. The exact realization of the path $(b_1, \dots, b_{n-n'})$ is irrelevant for the purposes of this proof. We give an example in Fig. 8(c).

³ Although we mistakenly claimed this in a preliminary version [26] of this paper.




Fig. 9. Illustration of the drawings Γ_k for small values of k . Any redrawing of Γ_k that preserves the y -coordinates and has convex inner faces needs an exponential increase in width.

Assume for a contradiction that there exists a morph \mathcal{M} composed of sequence of $o(n)$ linear morphing steps which convexifies $\Gamma^\Delta(n)$. Restricting the morph \mathcal{M} to the path P transforms $\Gamma^\Delta(n')$ into a reflex chain Γ_P on the boundary of a strictly convex polygon. It is easy to find $O(1)$ additional morphing steps that transform Γ_P into the drawing $\Gamma^-(n')$; for example we can intermediately aim for coordinates of a_1 and a_n which are extreme in some direction as in the proof of Lemma 15 (in fact, our situation here is much simpler, as we are not restricted to horizontal and vertical morphs anymore). Extending \mathcal{M} by these additional morphs yields planar morph with $o(n) \subseteq o(n')$ linear morphing steps that transforms $\Gamma^\Delta(n')$ into $\Gamma^-(n')$. This is a contradiction to Theorem 25. \square

6. Lower bound on grid size

Every 3-connected planar graph can be drawn with convex faces on a $n \times n$ grid [18] or with strictly convex faces on a $O(n^2) \times O(n^2)$ grid [6]. It is desirable to find morphs in which the intermediate drawings lie on a polynomial-sized grid, i.e., the coordinates of each vertex have a logarithmic number of bits. In this section, we show that this is not achievable with our current approach and, more generally, with any approach that uses Hong and Nagamochi’s redrawing technique. To do so, we design a family of drawings to show that a single horizontal morph to a convex drawing may unavoidably blow up the width of the drawing from $O(n)$ to $\Omega(n!)$. Thus, there is no hope of restricting to a polynomial-sized grid.

A *grid-drawing* of a planar graph G is a straight-line planar drawing of G in which all vertices are placed at integer coordinates. The *width* of a grid-drawing Γ is the length of a horizontal side of the smallest bounding box of Γ .

Lemma 26. *There exists a family of grid-drawings $(\Gamma_k)_k$ of internally 3-connected graphs G_k on n_k vertices such that the width of Γ_k is $w(\Gamma_k) = 2k$ and any grid-drawing of G_k in which the y -coordinates match those of Γ_k and in which every inner face is convex has at least width $w_c(\Gamma_k) \geq 4^{k-1}(2k - 2)! \in \Omega(w(\Gamma_k)!) = \Omega(n_k!)$.*

The drawing Γ_1 of G_1 is depicted in Fig. 9. The drawing Γ_{k+1} of G_{k+1} is obtained from Γ_k by introducing four new y -coordinates (called *levels*) and six new vertices: Introduce a cycle on six vertices as the new outer face and four new internal edges as shown in Fig. 9.

Note that G_k has exactly two vertices of degree two on the outer face. It follows directly from the construction that G_k has $6k + 1$ vertices, $4k$ levels, and admits a grid-drawing of width at most $2k$. It is easy to see that if G_{k-1} has a drawing with convex inner faces that preserves the levels of Γ_{k-1} , then the analogous statement holds for G_k ; we can simply shift the degree-2 vertices in G_{k-1} and the new vertices on levels 2 and $4k - 1$ far enough outwards.

We now argue that any convex drawing necessarily blows up the width with the following geometric observation. A similar idea was used by Lin and Eades [28] for the construction of hierarchical drawings where every straight-line drawing has a large width.

Observation 27. *Consider a grid-drawing of a path (a_1, b_1, c, b_2, a_2) such that c belongs to level k ; b_1 is left of b_2 on level $k + 1$; and a_1 is left of a_2 on level $k + j$ where $j \geq 2$. Let D denote the distance of b_1 and b_2 . If a_1 is not to the right of the oriented line $\overrightarrow{cb_1}$ and a_2 is not to the left of the oriented line $\overrightarrow{cb_2}$, then the distance of a_1 and a_2 is at least jD .*

For an illustration of Observation 27 consider Fig. 10. We use it to prove Lemma 26.

Proof of Lemma 26. Since the third level of Γ_1 contains three vertices, we have $w_c(\Gamma_1) \geq 2$. For $k \geq 2$, we prove the claim by induction with the stronger induction hypothesis that in any level-preserving grid-drawing Γ'_k of G_k with convex inner faces, the two vertices on the second level have distance at least $w_c(\Gamma_k)$; and that $w_c(\Gamma_k) \geq 4^{k-1}(2k - 2)!$.




Fig. 10. Illustration of Observation 27.

For the base case, consider any level-preserving redrawing Γ'_2 of Γ_2 with convex inner faces. Note that in any grid-drawing the three vertices on level 5 need a width of at least 2. Applying Observation 27 to the outermost vertices on levels 6, 5, 2 in Γ'_2 (highlighted by the red path in Fig. 9), it follows that in Γ'_2 the distance of the two vertices on level 2 is at least 8 and, thus, also $w_c(\Gamma_2) \geq 8$.

For the induction step with $k \geq 3$, consider a level-preserving redrawing Γ'_k of Γ_k in which all inner faces are convex. We apply Observation 27 twice in Γ'_k . By the induction hypothesis, the distance of the vertices on level 4 in Γ'_k is at least $w_c(\Gamma_{k-1})$ since Γ'_k contains a level-preserving redrawing of Γ_{k-1} in which all inner faces are convex. Applying Observation 27 to the outermost vertices on levels 3, 4, and $4k-3$ in Γ'_k (highlighted by the blue path in Fig. 9) shows that the vertices on level $4k-3$ in Γ'_k have a distance of at least $(4k-6) \cdot w_c(\Gamma_{k-1})$. Applying Observation 27 to the outermost vertices on levels 2, $4k-3$, and $4k-2$ (highlighted by the red path in Fig. 9) shows that the vertices on level 2 in Γ'_k have distance at least $(4k-4) \cdot (4k-6)w_c(\Gamma_{k-1}) =: \delta$. Since $w_c(\Gamma_{k-1}) \geq 4^{k-2}(2k-4)!$, a simple calculation yields that $\delta = 4^{k-1}(2k-2)!$ and, thus, the distance of the vertices in level 2 in Γ'_k is at least $w_c(\Gamma_k) \geq \delta = 4^{k-1}(2k-2)!$. \square

Note that each graph G_k resulting from this construction is internally 3-connected by Lemma 4 since it has a convex drawing. However, G_k is not 3-connected. We can obtain the same result for 3-connected graphs simply by contracting the two outer edges incident to the vertex on level $4k$. Note that these contractions remove the two degree-2 vertices on the outer face of G_k .

7. Conclusions

We have shown how to morph any straight-line planar drawing of an internally 3-connected graph to a strictly convex drawing while preserving planarity and increasing convexity throughout the morph. Moreover, our morph is composed of a linear number of horizontal and vertical steps, which is asymptotically optimal. The following questions are open:

1. Recall that during a convexity-increasing morph, the set of *internal* convex angles never decreases. We conjecture that every straight-line planar drawing of a (internally) 3-connected graph admits a convexity-increasing morph to a strictly convex drawing such that during the morph the set of *external* reflex angles also never decreases.
2. Our algorithm for finding a convexity-increasing morph to a convex drawing can be executed in $O(n^{1+\omega/2})$ time on a Real-RAM. Our new version of Lemma 11 also speeds up the run-time of the algorithm of Alamdari et al. [2] for morphing between two given drawings from $O(n^3)$ to $O(n^{1+\omega/2})$. For both these problems, it would be interesting to find even more efficient algorithms, or to establish non-trivial lower bounds on the run-time.
3. A main open question is to design piece-wise linear morphs with a polynomial bound on the bit complexity of the intermediate drawings. This would be a step towards having intermediate drawings that lie on a polynomial-sized grid, i.e. with a logarithmic number of bits for each vertex's coordinates. This is open both for our problem of morphing to a convex drawing and for the problem of morphing between two given planar straight-line drawings.
4. We have introduced the idea of using horizontal and vertical morphs. It would be interesting to further explore their visual quality or to explore what can be accomplished with this restriction. Horizontal and vertical morphs were very recently used by Da Lozzo et al. [14] for a special case of morphing between two upward planar drawings.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

We thank André Schulz for helpful discussions on generalizations of Tutte's algorithm. This work was begun at Dagstuhl workshop 17072, "Applications of Topology to the Analysis of 1-Dimensional Objects." We thank Dagstuhl, the organizers, and the other participants for a stimulating workshop. In particular, we thank Carola Wenk and Regina Rotmann for joining some of our discussions, and Irina Kostitsyna for contributing many valuable ideas.

Appendix A. Previous versions of Lemma 11

In this appendix we give more details on the versions of Lemma 11 proved by Hong and Nagamochi [24] and Angelini et al. [4]. First we introduce some terminology from [24].

A *hierarchical graph* is a graph with vertices assigned to *layers* which are horizontal lines. A *level drawing* of a hierarchical graph maps each vertex to a point on its layer line, and each edge to a y -monotone curve. Each edge can be directed upwards, so there is an underlying directed graph, and the standard notions of sink and source. A *hierarchical-st plane graph* is a hierarchical graph that has a level drawing that is planar and that has only one source and one sink. This is equivalent to our definition that all faces are y -monotone.

Lemma 28 (Hong and Nagamochi [24]). *Let G be an internally 3-connected hierarchical-st plane graph. Then a drawing of the outer face of G on a strictly convex polygon C can be extended in $O(n^2)$ time to a convex drawing of G by choosing an x -coordinate for each internal vertex.*

Implicit in the above statement is that the polygon C and the convex drawing of G respect the layers of the hierarchical graph. Hong and Nagamochi also gave necessary and sufficient conditions to deal with the case where the drawing of C need not be strictly convex.

The drawings generated by Lemma 28 are, in general, not strictly convex since the corresponding recursive algorithm draws certain paths of the graph by placing all its vertices on a common line. Angelini et al. [4, Theorem 6] describe an extension of the algorithm that perturbs the vertex coordinates in order to avoid angles of size π and, thereby, strengthen the above lemma to obtain *strictly* convex drawings. They do not analyze the runtime of their extension.

Appendix B. Morphing algorithm by Alamdari et al.

In this section we give some details to justify Theorem 2. The morphing algorithm of Alamdari et al. [2] depended on the basic result that it is possible to morph a triangulation so that a given quadrilateral formed by two adjacent triangles becomes convex. This problem—called “Quadrilateral Convexification”—was solved in Section 6 of the paper by means of a single unidirectional morph achieved by applying Hong and Nagamochi’s algorithm to redraw the triangulation so that the quadrilateral is convex. The run-time is $O(n^2)$ because of Hong and Nagamochi’s algorithm. This is improved to $O(n^{\omega/2})$ by our Lemma 11.

In fact, Quadrilateral Convexification, applied $O(n)$ times, was the bottleneck that gave a run-time of $O(n^3)$ for the algorithm of Alamdari et al. [2], as we now justify briefly. Their algorithm consisted of a preliminary step to find a compatible triangulation, thereby reducing the general problem to the case of triangulations. Then the main part of the algorithm iterated a contraction step that reduced the number of vertices by 1 in each step. Finally, there was a “clean-up” step that replaced contraction by “almost-contraction.” The preliminary compatible triangulation step (Theorem 4.1) involved $O(n)$ calls to Quadrilateral Convexification. As noted at the end of the proof of Theorem 4.1, this dominated the other work. Thus our Lemma 11 improves the run-time of the preliminary step from $O(n^3)$ to $O(n^{1+\omega/2})$. The main iterative algorithm (Theorem 5.1) consisted of $O(n)$ iterations, and each iteration involved a constant number of applications of Quadrilateral Convexification. Thus the original run-time of $O(n^3)$ is reduced to $O(n^{1+\omega/2})$. Finally, the clean-up step (Theorem 3.3) did not use Quadrilateral Convexification and ran in time $O(n)$ per iteration, for a total of $O(n^2)$.

References

- [1] O. Aichholzer, G. Aloupis, E.D. Demaine, M.L. Demaine, V. Dujmovic, F. Hurtado, A. Lubiw, G. Rote, A. Schulz, D.L. Souvaine, A. Winslow, Convexifying polygons without losing visibilities, in: Canadian Conference on Computational Geometry, CCCG, 2011.
- [2] S. Alamdari, P. Angelini, F. Barrera-Cruz, T.M. Chan, G. Da Lozzo, G. Di Battista, F. Frati, P. Haxell, A. Lubiw, M. Patrignani, V. Roselli, S. Singla, B.T. Wilkinson, How to morph planar graph drawings, SIAM J. Comput. 46 (2) (2017), 29 pages.
- [3] N. Alon, R. Yuster, Matrix sparsification and nested dissection over arbitrary fields, J. ACM 60 (4) (2013) 25.
- [4] P. Angelini, G. Da Lozzo, F. Frati, A. Lubiw, M. Patrignani, V. Roselli, Optimal morphs of convex drawings, in: L. Arge, J. Pach (Eds.), Proceedings of the 31st International Symposium on Computational Geometry, SoCG 2015, in: Leibniz International Proceedings in Informatics (LIPIcs), vol. 34, Dagstuhl, Germany, 2015, pp. 126–140.
- [5] S. Appelle, Perception and discrimination as a function of stimulus orientation: the “oblique effect” in man and animals, Psychol. Bull. 78 (4) (1972) 266.
- [6] I. Bárány, G. Rote, Strictly convex drawings of planar graphs, Doc. Math. 11 (2006) 369–391.
- [7] S. Cairns, Deformations of plane rectilinear complexes, Am. Math. Mon. 51 (5) (1944) 247–252.
- [8] J.H. Cantarella, E.D. Demaine, H.N. Iben, J.F. O’Brien, An energy-driven approach to linkage unfolding, in: Proceedings of the 20th Annual Symposium on Computational Geometry, SoCG, ACM, 2004, pp. 134–143.
- [9] N. Chiba, K. Onoguchi, T. Nishizeki, Drawing plane graphs nicely, Acta Inform. 22 (2) (1985) 187–201.
- [10] F.Y.L. Chin, J. Snoeyink, C.A. Wang, Finding the medial axis of a simple polygon in linear time, Discrete Comput. Geom. 21 (3) (1999) 405–420.
- [11] M. Chrobak, M.T. Goodrich, R. Tamassia, Convex drawings of graphs in two and three dimensions (preliminary version), in: S. Whitesides (Ed.), Proceedings of the Twelfth Annual Symposium on Computational Geometry, Philadelphia, PA, USA, PA, USA, May 24–26, 1996, ACM, 1996, pp. 319–328.
- [12] É. Colin De Verdière, M. Pocchiola, G. Vegter, Tutte’s barycenter method applied to isotopies, Comput. Geom. 26 (1) (2003) 81–97.
- [13] R. Connelly, E.D. Demaine, G. Rote, Straightening polygonal arcs and convexifying polygonal cycles, Discrete Comput. Geom. 30 (2003) 205–239.
- [14] G. Da Lozzo, G. Di Battista, F. Frati, M. Patrignani, V. Roselli, Upward planar morphs, in: International Symposium on Graph Drawing and Network Visualization, GD, Springer, 2018, pp. 92–105.

- [15] M. De Berg, M. Van Kreveld, M. Overmars, O. Cheong, *Computational Geometry: Algorithms and Applications*, 3rd edition, Springer, 2008.
- [16] G. Di Battista, R. Tamassia, Algorithms for plane representations of acyclic digraphs, *Theor. Comput. Sci.* 61 (2–3) (1988) 175–198.
- [17] M. Eiglsperger, S.P. Fekete, G.W. Klau, Orthogonal graph drawing, in: *Drawing Graphs*, Springer, 2001, pp. 121–171.
- [18] S. Felsner, Convex drawings of planar graphs and the order dimension of 3-polytopes, *Order* 18 (1) (2001) 19–37.
- [19] M.S. Floater, Parameterization and smooth approximation of surface triangulations, *Comput. Aided Geom. Des.* 14 (3) (1997) 231–250.
- [20] M.S. Floater, Parametric tilings and scattered data approximation, *Int. J. Shape Model.* 4 (03n04) (1998) 165–182.
- [21] M.S. Floater, C. Gotsman, How to morph tilings injectively, *J. Comput. Appl. Math.* 101 (1–2) (1999) 117–129.
- [22] S.J. Gortler, C. Gotsman, D. Thurston, Discrete one-forms on meshes and applications to 3D mesh parameterization, *Comput. Aided Geom. Des.* 23 (2) (2006) 83–112.
- [23] C. Gotsman, V. Surazhsky, Guaranteed intersection-free polygon morphing, *Comput. Graph.* 25 (1) (2001) 67–75.
- [24] S.-H. Hong, H. Nagamochi, Convex drawings of hierarchical planar graphs and clustered planar graphs, *J. Discret. Algorithms* 8 (3) (2010) 282–295.
- [25] H.N. Iben, J.F. O'Brien, E.D. Demaine, Refolding planar polygons, *Discrete Comput. Geom.* 41 (3) (2009) 444–460.
- [26] L. Kleist, B. Klemz, A. Lubiw, L. Schlipf, F. Staals, D. Strash, Convexity-increasing morphs of planar graphs, in: *Graph-Theoretic Concepts in Computer Science - Proceedings of the 44th International Workshop, WG 2018*, in: *Lecture Notes in Computer Science*, vol. 11159, Springer, 2018, pp. 318–330.
- [27] F. Le Gall, Powers of tensors and fast matrix multiplication, in: *Proceedings of the 39th International Symposium on Symbolic and Algebraic Computation*, ACM, 2014, pp. 296–303.
- [28] X. Lin, P. Eades, Towards area requirements for drawing hierarchically planar graphs, *Theor. Comput. Sci.* 292 (3) (2003) 679–695.
- [29] R.J. Lipton, D.J. Rose, R.E. Tarjan, Generalized nested dissection, *SIAM J. Numer. Anal.* 16 (2) (1979) 346–358.
- [30] R.J. Lipton, R.E. Tarjan, Applications of a planar separator theorem, *SIAM J. Comput.* 9 (3) (1980) 615–627.
- [31] K. Marriott, H. Purchase, M. Wybrow, C. Goncu, Memorability of visual features in network diagrams, *IEEE Trans. Vis. Comput. Graph.* 18 (12) (2012) 2477–2485.
- [32] J. Pach, G. Tóth, Monotone drawings of planar graphs, *J. Graph Theory* 46 (1) (2004) 39–47.
- [33] H.C. Purchase, E. Hoggan, C. Görg, How important is the “mental map”? – An empirical investigation of a dynamic graph layout algorithm, in: *International Symposium on Graph Drawing*, Springer, 2006, pp. 184–195.
- [34] H.C. Purchase, C. Pilcher, B. Plimmer, Graph drawing aesthetics—created by users, not algorithms, *IEEE Trans. Vis. Comput. Graph.* 18 (1) (2012) 81–92.
- [35] S. Rahman, Convex graph drawing, in: M.-Y. Kao (Ed.), *Encyclopedia of Algorithms*, Springer, Berlin, Heidelberg, 2015, pp. 1–7.
- [36] A. Ribó Mor, G. Rote, A. Schulz, Small grid embeddings of 3-polytopes, *Discrete Comput. Geom.* 45 (1) (2011) 65–87.
- [37] J. Richter-Gebert, Realization spaces of polytopes, in: *Lecture Notes in Mathematics*, Springer-Verlag, 1996.
- [38] C. Thomassen, Deformations of plane graphs, *J. Comb. Theory, Ser. B* 34 (3) (1983) 244–257.
- [39] C. Thomassen, Plane representations of graphs, in: J.A. Bondy, U.S.R. Murty (Eds.), *Progress in Graph Theory*, Academic Press, 1984, pp. 43–69.
- [40] W.T. Tutte, Convex representations of graphs, *Proc. Lond. Math. Soc.* s3–10 (1) (1960) 304–320.
- [41] W.T. Tutte, How to draw a graph, *Proc. Lond. Math. Soc.* 3 (1) (1963) 743–767.