



Strictly-Convex Drawings of 3-Connected Planar Graphs

Michael A. Bekos¹, Martin Gronemann², Fabrizio Montecchiani³,
and Antonios Symvonis⁴

¹ Department of Mathematics, University of Ioannina, Ioannina, Greece
`bekos@uoi.gr`

² Algorithms and Complexity Group, TU Wien, Vienna, Austria
`mgronemann@ac.tuwien.ac.at`

³ Department of Engineering, University of Perugia, Perugia, Italy
`fabrizio.montecchiani@unipg.it`

⁴ School of Applied Mathematical & Physical Sciences, NTUA, Athens, Greece
`symvonis@math.ntua.gr`

Abstract. Strictly-convex straight-line drawings of 3-connected planar graphs in small area form a classical research topic in Graph Drawing. Currently, the best-known area bound for such drawings is $O(n^2) \times O(n^2)$, as shown by Bárány and Rote by means of a sophisticated technique based on perturbing (non-strictly) convex drawings. Unfortunately, the hidden constants in such area bound are in the 10^4 order.

We present a new and easy-to-implement technique that yields strictly-convex straight-line planar drawings of 3-connected planar graphs on an integer grid of size $2(n-1) \times (5n^3 - 4n^2)$.

Keywords: Strictly-convex drawings · Area bounds · Planar graphs

1 Introduction

Drawing planar graphs is a fundamental topic in Graph Drawing with several important contributions over the last few decades [8, 14, 17, 21, 24]. One of the most influential is due to de Fraysseix, Pach and Pollack [8], who back in 1988 showed that every n -vertex planar graph admits a straight-line planar drawing on a $(2n-4) \times (n-2)$ grid, which can be computed in $O(n)$ time [7]. Since then, several improvements on the size of the underlying grid have been proposed in the literature [12, 15, 17, 18, 21, 25]. The best-known upper bound is $(n-2) \times (n-2)$ by Chrobak and Kant [6] and by Schnyder [21], who propose two conceptually different approaches to derive this bound. The former is an incremental drawing algorithm inspired by [8], while the latter is based on a face counting technique.

Straight-line drawings of planar graphs have also been extensively studied by requiring convexity [22], that is, the boundary of every face must be a convex polygon. Such drawings are called *convex* and always exist for 3-connected

Research of FM partially supported by Dip. Ingegneria, Univ. of Perugia, grant RICBA21LG “Algoritmi, modelli e sistemi per la rappresentazione visuale di reti”.




Fig. 1. A strictly-convex drawing of a 3-connected planar graph on 7 vertices.

planar graphs [23,24]. Again the aim is to keep the size of the underlying grid as small as possible; see [10] for a survey. Early results date back to Schnyder and Trotter [22], Chrobak and Kant [6], Di Battista et al. [11] and Felsner [12]. The latter guarantees the existence of a convex drawing of a 3-connected planar graph on a $(f - 1) \times (f - 1)$ integer grid, where $f = O(n)$ is the number of faces.

Note that in a convex drawing three vertices on the boundary of a face can be collinear. If this is not allowed, then the corresponding drawings are called *strictly-convex*. Since an n -vertex cycle cannot be drawn strictly-convex on a grid of size $o(n^3)$ [5], it follows that strictly-convex drawings are more demanding in terms of required area. As an adaptation of the standard incremental drawing algorithms or the face-counting methods is rather difficult, the only approach that has been exploited so far to obtain strictly-convex drawings is to perturb convex drawings. This idea was pioneered by Chrobak, Goodrich and Tamassia [5], who claimed (without giving details) that every 3-connected planar graph admits a strictly-convex drawing on an $O(n^3) \times O(n^3)$ grid. The area bound was improved to $O(n^{7/3}) \times O(n^{7/3})$ by Rote [19] and to $O(n^2) \times O(n^2)$ by Bárány and Rote [1], which is currently the best-known asymptotic upper bound. However, as the authors mention “the constants hidden in the O -notation are on the order of 100 for the width and on the order of 10,000 for the height. This is far too much for applications where one wants to draw graphs on a computer screen” [1].

Our Contribution. We continue the research on strictly-convex drawings of 3-connected planar graphs. Our contribution is a new technique that computes strictly-convex drawings of 3-connected n -vertex planar graphs on an integer grid of size $2(n - 1) \times (5n^3 - 4n^2)$, as outlined in the following theorem. Although the asymptotic area bound is the same as the one in [1], the multiplicative constants are significantly smaller. Also, the proposed technique is elegant and can be readily implemented to run in linear time. On the other hand, the aspect-ratio of the produced drawings is quadratic rather than constant.

Theorem 1. *Every 3-connected planar graph with n vertices admits a strictly-convex planar straight-line drawing on an integer grid of size $2(n - 1) \times (5n^3 - 4n^2)$. Also, the drawing can be computed in $O(n)$ time.*

Our technique starts with a convex drawing computed by Kant’s algorithm [17]. We rely on properties of such a drawing to show that shifting vertices upwards




Fig. 2. In both drawings, (x, y) is above s , while (x', y') is below s .

by using a strictly-increasing and strictly-convex function preserves planarity; a property of independent interest. Also, the obtained planar drawing is convex and collinear vertices in a face, if any, are horizontally aligned. For such vertices, a second shifting yields an internally strictly-convex drawing. A suitable augmentation guarantees that the outer face is also strictly-convex.

Paper Structure. Section 2 contains basic definitions and tools. In Sect. 3, we introduce properties of Kant’s algorithm that we leverage in our technique. Section 4 describes our algorithm. Section 5 concludes the paper with a brief discussion and open problems. For space reasons, some proofs are omitted (the corresponding statements are marked with \star) and can be found in [4].

2 Preliminaries

Basic Definitions. Let $f : \mathbb{R} \rightarrow \mathbb{R}$ be a function. If $f(a) < f(b)$ for every pair $a, b \in \mathbb{R}$ with $a < b$, then f is *strictly-increasing*. Function f is *strictly-convex* if for all $t, 0 < t < 1$, and all $a, b \in \mathbb{R}$, it holds $f(ta + (1-t)b) < tf(a) + (1-t)f(b)$. Consider three points $(x_1, y_1), (x, y), (x_2, y_2)$, with $x_1 < x < x_2$ and let s be the line-segment connecting (x_1, y_1) and (x_2, y_2) . We say that (x, y) is *above* (resp., *below*) s if the slope of s is smaller (resp., larger) than the slope of the line-segment connecting (x_1, y_1) and (x, y) ; see Fig. 2. The next lemma easily follows from the chordal slope lemma [20].

Lemma 1 (\star). *Let $(x_1, y_1), (x, y)$ and (x_2, y_2) be three collinear points with $x_1 < x < x_2$ that are not horizontally aligned. If $f : \mathbb{R} \rightarrow \mathbb{R}$ is a strictly-convex function, then point $(x, f(y))$ is below the line-segment with endpoints $(x_1, f(y_1))$ and $(x_2, f(y_2))$.*

Drawings and Embeddings. We assume familiarity with basic graph drawing concepts [9]. In particular, a *plane graph* is a graph with a prescribed planar embedding. Unless otherwise specified, we consider drawings that are straight-line, planar and whose vertices are on an integer grid. A drawing is *convex* (*strictly-convex*) if the boundary of each face is a convex (strictly-convex) polygon. Similarly, a drawing is *internally convex* (*internally strictly-convex*) if the boundary of each inner face is a convex (strictly-convex) polygon. Given a drawing Γ of a graph G , denote by (x_u, y_u) the coordinates of vertex u in Γ . For two vertices u and v in Γ , we denote by Δ_{uv} the interior of the right triangle whose

corners are u, v and the intersection of the vertical line through the vertex having the lowest y -coordinate with the horizontal line through the vertex having the highest y -coordinate (among u, v). For example, in Fig. 3a, the Δ_{uv} triangle of the endpoints of each edge (u, v) is striped.

Canonical Order. Let G be a 3-connected plane graph with n vertices. Let $\delta = (P_0, \dots, P_m)$ be a partition of the vertices of G into paths, such that $P_0 = \{v_1, v_2\}$, $P_m = \{v_n\}$, and edges (v_1, v_2) and (v_1, v_n) exist and belong to the outer face. For $k = 0, \dots, m$, let G_k be the subgraph induced by $\cup_{i=0}^k P_i$. Let C_k be the *contour* of G_k defined as follows: If $k = 0$, then C_0 is the edge (v_1, v_2) , while if $k > 0$, then C_k is the path from v_1 to v_2 obtained by removing (v_1, v_2) from the cycle delimiting the outer face of G_k . Partition δ is a *canonical order* [17] of G if for each $k = 1, \dots, m-1$ the following conditions hold: (i) G_k is biconnected and internally 3-connected, (ii) all neighbors of P_k in G_{k-1} are on C_{k-1} , (iii) either P_k is a *singleton* (i.e., $|P_k| = 1$), or P_k is a *chain* (i.e., $|P_k| > 1$) and the degree of each vertex of P_k is 2 in G_k , (iv) all vertices of P_k with $0 \leq k < m$ have at least one neighbor in P_j for some $j > k$. For example, a canonical order for the graph of Fig. 1 is $P_0 = \{v_1, v_2\}$, $P_1 = \{v_3, v_4\}$, $P_2 = \{v_5, v_6\}$ and $P_3 = \{v_7\}$. A canonical order of G can be computed in $O(n)$ time [17].

Kant's Algorithm. Kant [17] describes an incremental drawing algorithm that, in linear time, computes a convex straight-line planar drawing Γ of an n -vertex plane graph G on an integer grid of size $(2n - 4) \times (n - 2)$. The drawing Γ has the same planar embedding as the input graph G . The algorithm is based on a canonical order δ of G and works as follows: Initially, vertices v_1 and v_2 of P_0 are placed at points $(0, 0)$ and $(1, 0)$, respectively. For $k = 1, \dots, m$, assume that a convex drawing Γ_{k-1} of G_{k-1} has been constructed in which the edges of contour C_{k-1} are drawn with slopes 0 and ± 1 (*contour condition*; see Fig. 3a). Let (w_1, \dots, w_p) be the vertices of C_{k-1} from left to right in Γ_{k-1} , where $w_1 = v_1$ and $w_p = v_2$. Each vertex v in G_{k-1} has been associated with a *shift-set* $S(v)$, such that Γ_{k-1} is *stretchable*, that is, for each $i = 1, \dots, p$ the result of shifting $S(w_i), \dots, S(w_p)$ by one (or more) units to the right is a convex drawing of G_{k-1} . Let $P_k = \{z_1, \dots, z_p\}$ be the next path in δ . Let w_ℓ and w_r be the leftmost and rightmost neighbors of P_k on C_{k-1} in Γ_{k-1} , where $1 \leq \ell < r \leq p$. To introduce P_k and to avoid edge-overlaps, the algorithm first identifies two so-called critical vertices $w_{\ell'}$ and $w_{r'}$ with $\ell \leq \ell', r' \leq r$ and then shifts (i) by one unit to the right each vertex in $\cup_{i=\ell'}^p S(w_i)$ and then (ii) by one unit to the right each vertex in $\cup_{i=r'}^p S(w_i)$. Then, z_1 is placed at intersection of the line of slope $+1$ through w_ℓ with the line through of slope -1 point w_r ; see Fig. 3b. If P_k is a chain, then for $i = 2, \dots, p$, vertex z_i is placed one unit to the right of z_{i-1} by shifting each vertex in $\cup_{i=r'}^p S(w_i)$ one unit to the right. Finally, the shift-sets of the vertices of P_k are defined accordingly to ensure that Γ_k is stretchable.

3 Properties of Kant's Algorithm

We provide properties of drawings computed by Kant's algorithm that we leverage in the next section; some of these properties are indirectly mentioned also




Fig. 3. Introducing a singleton $P_k = \{z_1\}$ in Γ_{k-1} in the algorithm by Kant [17].

in [16]. To ease the presentation, we first introduce a 4-coloring for the edges of G similar to the one by Schnyder [13,21]. We color edge (v_1, v_2) of G_0 black. Given a 4-coloring for G_{k-1} with $k = 1, \dots, m$, we extend it for G_k as follows (see Figs. 1 and 3a). We first color the edges of G_k that do not belong to G_{k-1} and are on contour C_k . Namely, the first such edge encountered in a clockwise walk of C_k from v_1 to v_2 is blue, the last one is green and all remaining ones (that is, those having both endpoints in P_k when P_k is a chain) are black. The remaining edges of G_k not in G_{k-1} are red (i.e., those that are incident to P_k and are not part of C_k ; this case only arises if P_k is a singleton by Condition (iii) of the canonical order), which implies that C_k has no red edges.

Since a shift to introduce a path of δ in the incremental construction of Γ can only decrease the slope of a blue edge, increase the slope of a green edge, while the black and the red edges maintain their slope [17], we have that:

- the slope of each blue edge ranges in $(0, 1]$,
- the slope of each black edge is 0,
- the slope of each green edge ranges in $[-1, 0)$, and
- the slope of each red edge ranges in the complement of $[-1, 1]$.

Since each inner face in Γ is formed when a path of δ is introduced during the incremental construction, part of it belongs to the contour, while its remaining part is formed by the introduced path, which gives rise to the following property.

Property 1. *Let x be the leftmost vertex of an inner face g in Γ (in case of more than one such vertices select the bottommost one). A counterclockwise walk of g starting from x consists of the following boundary parts (see Fig. 4):*

- i. a (possibly empty) strictly descendant path of green edges,
- ii. at most one black edge,
- iii. a (possibly empty) strictly ascendant path of blue edges,
- iv. a green or red edge,
- v. a (possibly empty) horizontal path of black edges, and
- vi. a blue or red edge.




Fig. 4. Illustration for the shape of a face.

Boundary parts (iv)–(vi) (dotted in Fig. 4) are introduced when a path is added during the incremental construction of Γ (which implies that at least one of boundary parts (i)–(iii) is part of the contour and thus is non empty in g). So, boundary parts (iv)–(vi) cannot simultaneously contain black and red edges (by the edge-coloring and Condition (iii) of canonical order).

Property 2 (\star). *Each vertex w of a path P_i , with $0 < i \leq m$, has at least two incident edges (a, w) and (b, w) , such that $y_a \leq y_w$, $y_b \leq y_w$, and $x_a < x_w < x_b$.*

Property 3 (\star). *Every face in Γ has at most one edge drawn vertical.*

Property 4 (\star). *Let u , v and w be three consecutive vertices encountered in this order in a counterclockwise walk along the boundary of an inner face of Γ . If they are collinear and the line through them has zero slope, then they are part of a chain. If they are collinear and the line through them has positive (negative) slope, then $y_u < y_v < y_w$ ($y_u > y_v > y_w$). If they are not collinear, then $v \notin \Delta_{uw}$.*

4 Algorithm Description

We now describe our approach to compute a strictly-convex drawing of a 3-connected plane graph, assuming that its outer face has at most 5 vertices (such a face always exists). We start with Sect. 4.1, in which we describe the properties of what we call lifting functions and liftable drawings. A key property is that applying a non-affine transformation to a liftable drawing by means of a lifting function preserves planarity. As this tool might be of independent interest, we state it as general as possible. In Sect. 4.2, we prove that drawings computed by Kant’s algorithm are indeed liftable and that a transformation via a liftable function makes them internally strictly-convex except for possible horizontally-aligned paths. Up to this point, it was not needed to choose a particular lifting function; in Sect. 4.3 we unveil our choice. We also design a second transformation targeted to faces containing paths of horizontally-aligned vertices. The output of this step is an internally strictly-convex drawing. The last step of the algorithm is described in Sect. 4.4, namely a simple preprocessing in which the outer face of the input graph, which by our assumption has at most 5 vertices, is suitably augmented with dummy vertices, whose removal from the computed drawing guarantees that all faces (including the outer one) are strictly-convex.

4.1 Lifting Functions

Given a drawing Γ of a graph G and a function $f : \mathbb{R} \mapsto \mathbb{R}$, we refer to the drawing Γ_f obtained by applying the transformation $(x_u, y_u) \mapsto (x_u, f(y_u))$ to Γ as the *transformed drawing* of Γ with respect to f . In view of Theorem 2 below, we focus on lifting functions and liftable drawings (see Definitions 1 and 2).

Definition 1. A function $f : \mathbb{R} \mapsto \mathbb{R}$ is *lifting* if and only if (i) f is strictly-convex and strictly-increasing; (ii) $f(r) \geq r, \forall r \in \mathbb{R}$; (iii) $r \in \mathbb{N} \Rightarrow f(r) \in \mathbb{N}$.

The next property follows directly from Definition 1.

Observation 1. Let Γ be a drawing of a graph G . Given a lifting function f , three vertices are horizontally (vertically) aligned in Γ if and only if they are horizontally (vertically) aligned in the transformed drawing Γ_f .

Definition 2. A planar straight-line grid drawing of a graph is called *liftable* if for every edge (u, v) there is no vertex of G in Δ_{uv} .

Theorem 2. Let Γ be a planar straight-line grid drawing of a plane graph G . If Γ is liftable, then its transformed drawing Γ_f with respect to a lifting function f is a planar straight-line grid drawing of G with the same planar embedding as Γ .

Proof. Condition (iii) of Definition 1 trivially implies that Γ_f is a grid drawing. We next prove that Γ and Γ_f have the same planar embedding. (Note that, if Γ is not liftable, Γ_f is not necessarily planar.) Since, by Condition (ii) of Definition 1, Γ_f is obtained from Γ by shifting vertices upwards, the existence in Γ_f of an edge crossing or of a vertex having a circular order of its incident edges different than the one in Γ , implies that there exist a vertex w and an edge (u, v) in G , such that w is below (above) (u, v) in Γ and above (below) (u, v) in Γ_f .

We next argue that the situation described above is not possible. Consider a vertex w and an edge (u, v) of G and let s and s' be the line-segments representing (u, v) in Γ and Γ_f . Clearly, it suffices to consider the case in which $x_u \leq x_w \leq x_v$. Let p and p' be the vertical projection of w on s and s' , respectively. Also, let y_p and $y_{p'}$ be the y -coordinates of p in Γ and of p' in Γ_f , respectively. Suppose $y_u \leq y_v$; the case in which $y_u > y_v$ is symmetric.

Firstly, consider the case in which w is above s , i.e., $y_w > y_p$. Since Γ is liftable, vertex w does not belong to Δ_{uv} . Hence, $y_u \leq y_p \leq y_v \leq y_w$. Since f is strictly-increasing, it follows $f(y_u) \leq f(y_p) \leq f(y_v) \leq f(y_w)$. However, $f(y_w) > y_{p'}$ implies that w is above s' , as desired. Secondly, consider the case in which w is below s in Γ . Here, we distinguish three cases: $x_w = x_u$, $x_w = x_v$ and $x_u < x_w < x_v$. In the first case, we have $y_w < y_p = y_u$ and $y_{p'} = f(y_u)$. Since f is strictly-increasing, it holds $f(y_w) < f(y_u) = y_{p'}$, i.e., w is below s' , as desired. The second case is analogous. For the third case, we know that $y_w < y_p$. Since p lies on s , by Lemma 1, $f(y_p) < y_{p'}$ holds. Also, since f is strictly-increasing, it follows $f(y_w) < f(y_p)$. Thus, $f(y_w) < y_{p'}$ holds, i.e., w is below s' , as desired. \square

4.2 Application to Kant's Drawings

We now show that applying a lifting function to a drawing computed by Kant's algorithm (see Sect. 3) yields a drawing with several important properties.

Lemma 2. *Let Γ be a drawing of a 3-connected plane graph G computed by Kant's algorithm. Drawing Γ is liftable.*

Proof. Consider an edge (u, v) of G and w.l.o.g. assume $y_u < y_v$ in Γ . We prove that there is no vertex w in Δ_{uv} . This is obvious when $x_u = x_v$, since $\Delta_{uv} = \emptyset$. Hence, either $x_u < x_v$ or $x_u > x_v$. Consider the former case; the latter can be treated symmetrically. Suppose for a contradiction that there exists at least one vertex (other than u and v) in Δ_{uv} . Let w be the rightmost vertex out of those in Δ_{uv} . Since w is in Δ_{uv} we know $y_u < y_w < y_v$. Since $y_u \neq y_v$, it follows that (u, v) is not the edge (v_1, v_2) of P_0 , which, in turn, implies that vertex w belongs to a path P_i with $i > 0$, since $y_u < y_w$. Hence, by Property 2, w has at least two incident edges (a, w) and (b, w) , such that $y_a \leq y_w, y_b \leq y_w$, and $x_a < x_w < x_b$. Since b is to the right of w , the way we selected w implies that b does not belong to Δ_{uv} ; consequently, (w, b) crosses (u, v) , contradicting the planarity of Γ . \square

By combining Theorem 2 and Lemma 2, we conclude the following.

Theorem 3 (\star). *Given a 3-connected plane graph G and a lifting function f , let Γ_f be the transformed drawing of a drawing Γ of G computed by Kant's algorithm. Then, Γ_f is internally-convex and planar with the same embedding as Γ . Also, if two consecutive edges of an inner face of Γ_f form an angle π inside this face, then these edges are horizontal.*

Proof sketch. Since, by Lemma 2, Γ is liftable, by Theorem 2 Γ_f has the same planar embedding as Γ . Consider a counterclockwise walk along the boundary of an inner face g in Γ and let u, v and w be three consecutive vertices along this walk. Let α and α' be the angle at v formed by the edges (u, v) and (v, w) inside g in Γ and in Γ_f , respectively. Since Γ is convex, $\alpha \leq \pi$. We claim that $\alpha' \leq \pi$ and that if $\alpha' = \pi$, then (u, v) and (v, w) are horizontally aligned in Γ_f . We prove the claim when u, v and w are collinear in Γ . By Property 3, vertices u, v and w are not vertically aligned in Γ . If they are horizontally aligned in Γ , then by Observation 1 they are horizontally aligned also in Γ_f , as desired. Suppose now the line ℓ through u, v and w in Γ is either of positive or of negative slope, which by Property 4 implies that either $y_u < y_v < y_w$ or $y_u > y_v > y_w$ holds, respectively. Then, by Lemma 1, it follows that v is below the line-segment connecting u and w in Γ_f , hence $\alpha' < \pi$. \square

4.3 Putting Everything Together

We are now ready to put all pieces together. Let G be a 3-connected plane graph with n vertices. Without loss of generality, we can assume that the outer face of G contains at most 5 vertices (which will be useful in the next subsection), since




Fig. 5. (a) Illustration of the procedure of shifting the vertices of a chain upwards, and (b-c) cases that arise in the proof of Theorem 4.

such a face always exists. Let Γ be a convex drawing of G computed by Kant’s algorithm and let $f : \mathbb{R} \rightarrow \mathbb{R}$ be the function $f(y) = 5(n - 2)^2y + y^2$. Clearly, f is a lifting function. Hence, by Theorem 3, in the transformed drawing Γ_f , each inner face g that is not strictly-convex contains at least three horizontally-aligned vertices. By property 4, these vertices are part of a chain in the canonical order δ . Hence, by Condition (iv) of canonical order, each of the vertices of this chain has at least one neighbor placed above it. By definition of f it follows that each of these neighbors is positioned at least $5(n - 2)^2$ units above the chain in Γ_f . We exploit this property to turn Γ_f into an internally strictly-convex drawing by shifting all vertices of each chain upwards while keeping Γ_f planar.

To this end, let $P_k = \{z_1, \dots, z_p\}$ with $p \geq 2$ be a chain in the canonical order δ used to construct Γ . For $i = 1, \dots, p$, we shift vertex z_i of P_k by $(x_{z_i} - x_{z_1})(x_{z_p} - x_{z_i})$ units upwards; see Fig. 5a. It follows that if λ is the total width of P_k in Γ (and thus also in Γ_f), then each vertex in P_k is shifted by at most $\lambda^2/4$ units of length upwards, which is in turn at most $(n - 2)^2$, since the total width of Γ is at most $2n - 4$, and therefore $\lambda \leq 2n - 4$. Also, note that only the *internal vertices* of P_k are shifted (if any), i.e., only the vertices z_i with $2 \leq i \leq p - 1$. Let $\widehat{\Gamma}_f$ be the drawing obtained from Γ_f by applying the aforementioned procedure to each chain, which we call the *curved drawing* of Γ_f . Clearly, $\widehat{\Gamma}_f$ is a grid drawing, we can prove that it is planar and internally strictly-convex.

Theorem 4 (\star). *Given a 3-connected n -vertex plane graph G and the lifting function $f : \mathbb{R} \mapsto \mathbb{R}$ with $f(y) = 5(n - 2)^2y + y^2$, let Γ_f be the transformed drawing of a drawing Γ of G computed by Kant’s algorithm. Then, the curved drawing $\widehat{\Gamma}_f$ of Γ_f is an internally strictly-convex grid drawing of G with the same planar embedding as Γ .*

Proof sketch. Concerning the planarity of $\widehat{\Gamma}_f$, consider any internal vertex of a chain and its neighbors. At high-level, we have that the envelope through such neighbors is a polygon whose boundary is formed by a left and a right path that are y -monotone (see Fig. 5b). Since the vertex remains below its lowest successor in the canonical order (see Fig. 5c), no edge crossing can be introduced. Thus, $\widehat{\Gamma}_f$ is planar with the same embedding as Γ_f (and thus as Γ , by Theorem 3).




Fig. 6. Illustration for the proof of Theorem 4; W denotes the width of Γ_f .

We next prove that $\widehat{\Gamma}_f$ is internally strictly-convex. Let g be an inner face in $\widehat{\Gamma}_f$. Recall that only internal chain-vertices are shifted in the transition from Γ_f to $\widehat{\Gamma}_f$. Hence, if g does not contain internal chain-vertices, then it is strictly-convex by Theorem 3. Thus, we may assume that g contains at least one such vertex z . By Property 1, z is either in the topmost chain of black edges of g , or it is one of the (at most two) bottommost vertices of g . Consider the former, as the latter is similar. Let $P_k = \{z_1, \dots, z_p\}$ be the chain containing z . We argue that any angle inside g incident to these vertices is smaller than π . By construction, this is the case for vertices z_2, \dots, z_{p-1} . Hence, it remains to consider the angles at z_1 and z_p . Since the two cases are symmetric, consider the angle at z_1 . Let w_ℓ be the neighbor of z_1 along C_{k-1} , i.e., w_ℓ is the vertex preceding z_1 in a clockwise walk of g starting from z_1 . We will prove that the slope of (w_ℓ, z_1) is strictly greater than the one of (z_1, z_2) , hence the angle at z_1 is less than π .

Refer to Fig. 6. By the way the vertices of P_k are shifted in the transition from Γ_f to $\widehat{\Gamma}_f$, it follows that the maximum of the slope of (z_1, z_2) is $2(n-2) - 1$ (i.e., achieved when P_k is of maximum x -length in $\widehat{\Gamma}_f$ and the x -distance of z_1 and z_2 in $\widehat{\Gamma}_f$ is 1). We next argue for the slope of the edge (w_ℓ, z_1) . Recall that vertex z_1 is not an interior vertex of P_k , which implies that it has not been shifted in the transition from Γ_f to $\widehat{\Gamma}_f$. The same, however, does not necessarily hold for w_ℓ . As a matter of fact, this vertex may be part of a chain, i.e., when g does not contain boundary part (i) but boundary part (ii) of Property 1. This implies that it may have been shifted upwards by at most $(n-2)^2$ units in the transition from Γ_f to $\widehat{\Gamma}_f$. The minimum of the slope of (w_ℓ, z_1) in Γ_f is achieved, when (w_ℓ, z_1) is of maximum x -length in Γ_f and of minimum y -length. Since the former is at most $2(n-2) - 1$, while the latter is at least $5(n-2)^2$, it follows that the minimum of the slope of (w_ℓ, z_1) is potentially $\frac{5(n-2)^2}{2(n-2)-1}$ in Γ_f . Since in the transition from Γ_f to $\widehat{\Gamma}_f$ vertex w_ℓ may be shifted by at most $(n-2)^2$ units, it follows that the slope of (w_ℓ, z_1) may reduce further to $\frac{5(n-2)^2 - (n-2)^2}{2(n-2)-1}$, which is its minimum value. Therefore, the slope of the edge (w_ℓ, z_1) is strictly greater than the one of (z_1, z_2) , since the following trivially holds:

$$\frac{5(n-2)^2 - (n-2)^2}{2(n-2) - 1} > 2(n-2) - 1 \iff 2(n-2) > 2(n-2) - 1. \quad \square$$




Fig. 7. Treating the outer face when its degree is five (a–b) and four (c–d).

4.4 Outer Face and Final Analysis

To complete the description of our algorithm, it remains to guarantee that the outer face of the computed drawings is strictly-convex. To this aim, we slightly augment the input graph G and suitably choose the canonical order to give as input to Kant’s algorithm. Consider a planar embedding of G and let v_1, v_2, \dots, v_h be the vertices on the outer face (see Fig. 7a); recall that we have assumed $h \leq 5$. If $h = 3$, then the boundary of the outer face is a triangle and hence strictly-convex. So, assume $4 \leq h \leq 5$. To ease the presentation, we let $h = 5$ (see Figs. 7a and 7b), as the case $h = 4$ is simpler (see Figs. 7c and 7d). We proceed by adding two vertices v_1^* and v_2^* in the outer face of G and edges $(v_1^*, v_2^*), (v_1^*, v_5), (v_1^*, v_1), (v_1^*, v_2), (v_2^*, v_3), (v_2^*, v_4),$ and (v_2^*, v_5) . The resulting graph G^* is still planar and 3-connected. In particular, its outer face is a 3-cycle formed by v_1^*, v_2^*, v_5 . We compute a canonical order δ^* of G^* with $P_0 = (v_1^*, v_2^*)$ and $P_m = \{v_5\}$. The key observation is that the second set of δ^* is the chain $P_1 = \{v_2, v_3\}$, since it forms the inner face f^* of G^* with (v_1^*, v_2^*) on its boundary.

Next, we apply the algorithm supporting Theorem 4 to G^* using the aforementioned canonical order δ^* and obtain a drawing of it that is internally strictly-convex. We next prove that the removal of v_1^* and v_2^* from this drawing yields a drawing of G that is strictly-convex. By Theorem 4 and by our augmentation, it suffices to guarantee that the outer face of the obtained drawing is strictly-convex. Consider first the inner angle at v_5 of the polygon bounding the outer face; this angle is strictly less than π , because v_5 is the topmost vertex of the drawing (and no other vertex is horizontally aligned with it). A similar argument applies for the angles at v_1 and v_4 ; in particular, after the removal of v_1^* and v_2^* , vertices v_1 and v_4 are the leftmost and the rightmost neighbors of v_5 , respectively, and therefore they are the leftmost and the rightmost vertices in the drawing, respectively. Concerning v_2 and v_3 , they are horizontally aligned and, after the removal of v_1^* and v_2^* , they are the bottommost vertices of the drawing. Thus, their angles are also strictly less than π completing the proof of our claim. To conclude the proof of Theorem 1, it remains to discuss the area required by the drawing obtained as above and the time complexity to compute it.

Area Bound. The drawing Γ computed by Kant’s algorithm for G^* fits on an integer grid of size $(2n^* - 4) \times (n^* - 2)$, where $n^* = n + 2$ (G^* has two more vertices than G). The transformed drawing Γ_f of Γ by means of the lifting

function $f : \mathbb{R} \mapsto \mathbb{R}$ with $f(y) = 5(n^* - 2)^2 y + y^2$ has the same width as Γ , while the vertices v_1^* and v_2^* have y -coordinate 0 in Γ_f . On the other hand, vertex v_5 has y -coordinate $5(n^* - 2)^2(n^* - 2) + (n^* - 2)^2$, which is also the height of Γ_f . Since no vertex of the outer face of Γ_f is further shifted upwards, the curved drawing $\widehat{\Gamma}_f$ of Γ_f has the same width and height as Γ_f . After removing v_1^* and v_2^* , the width of the final drawing of G is at least two units less than the one of $\widehat{\Gamma}_f$, while its height is at least $5(n^* - 2)^2$ units less. Since $n^* = n + 2$, the final drawing lies on a grid of size $((2(n + 2) - 4) - 2) \times (5((n + 2) - 2)^3 - 4((n + 2) - 2)^2) = 2(n - 1) \times (5n^3 - 4n^2)$.

Time Complexity. Each step of our algorithm can be implemented in $O(n)$ time: (i) finding a planar embedding of G with a face of degree at most 5, (ii) computing G^* , a canonical order of it, and applying Kant's algorithm to G^* , (iii) computing the transformed drawing with respect to our lifting function f and updating the position of the internal chain-vertices. This completes the proof of Theorem 1.

5 Conclusions and Open Problems

We have provided a linear-time algorithm that computes a strictly-convex drawing of a 3-connected planar graph on an integer grid of size $2(n - 1) \times (5n^3 - 4n^2)$. Compared to the previously best-known upper bound for such drawings [1], we largely improve the multiplicative constants by means of an arguably simpler algorithm, which therefore has the potential to be of practical use. Along the way, we proved tools that can be of independent interest (see in particular Theorem 2). Some problems that stem from our research are the following:

- Can we achieve a similar area bound together with a constant aspect ratio?
- Is $\Omega(n^4)$ a lower bound for the area requirement of strictly-convex drawings?
- Can we compute strictly-convex drawings in small area with good edge-vertex resolution [2, 3]?

References

1. Bárány, I., Rote, G.: Strictly convex drawings of planar graphs. *Doc. Math.* **11**, 369–391 (2006)
2. Bekos, M.A., Gronemann, M., Montecchiani, F., Pálvölgyi, D., Symvonis, A., Theoharous, L.: Grid drawings of graphs with constant edge-vertex resolution. *Comput. Geom.* **98**, 101789 (2021)
3. Bekos, M.A., Gronemann, M., Montecchiani, F., Symvonis, A.: Convex grid drawings of planar graphs with constant edge-vertex resolution. In: *IWOCA 2022*. LNCS, vol. 13270, pp. 157–171. Springer (2022)
4. Bekos, M.A., Gronemann, M., Montecchiani, F., Symvonis, A.: Strictly-convex drawings of 3-connected planar graphs. *CoRR* abs/2208.13388 (2022)
5. Chrobak, M., Goodrich, M.T., Tamassia, R.: Convex drawings of graphs in two and three dimensions (preliminary version). In: Whitesides, S. (ed.) *SoCG*, pp. 319–328. ACM (1996)

6. Chrobak, M., Kant, G.: Convex grid drawings of 3-connected planar graphs. *Int. J. Comput. Geom. Appl.* **7**(3), 211–223 (1997)
7. Chrobak, M., Payne, T.H.: A linear-time algorithm for drawing a planar graph on a grid. *Inf. Process. Lett.* **54**(4), 241–246 (1995)
8. de Fraysseix, H., Pach, J., Pollack, R.: How to draw a planar graph on a grid. *Combinatorica* **10**(1), 41–51 (1990)
9. Di Battista, G., Eades, P., Tamassia, R., Tollis, I.G.: *Graph Drawing: Algorithms for the Visualization of Graphs*. Prentice-Hall (1999)
10. Di Battista, G., Frati, F.: Drawing trees, outerplanar graphs, series-parallel graphs, and planar graphs in a small area. In: Pach, J. (ed.) *Thirty Essays on Geometric Graph Theory*, pp. 121–165. Springer, New York (2013)
11. Di Battista, G., Tamassia, R., Vismara, L.: Output-sensitive reporting of disjoint paths. *Algorithmica* **23**(4), 302–340 (1999)
12. Felsner, S.: Convex drawings of planar graphs and the order dimension of 3-polytopes. *Order* **18**(1), 19–37 (2001)
13. Felsner, S.: *Geometric Graphs and Arrangements*. Advanced Lectures in Mathematics, Vieweg (2004)
14. Fáry, I.: On straight lines representation of planar graphs. *Acta Sci. Math. (Szeged)* **11**, 229–233 (1948)
15. He, X.: Grid embedding of 4-connected plane graphs. *Discret. Comput. Geom.* **17**(3), 339–358 (1997)
16. Kant, G.: Hexagonal grid drawings. In: Mayr, E.W. (ed.) *WG 1992*. LNCS, vol. 657, pp. 263–276. Springer, Heidelberg (1993). https://doi.org/10.1007/3-540-56402-0_53
17. Kant, G.: Drawing planar graphs using the canonical ordering. *Algorithmica* **16**(1), 4–32 (1996)
18. Miura, K., Nakano, S., Nishizeki, T.: Grid drawings of 4-connected plane graphs. *Discret. Comput. Geom.* **26**(1), 73–87 (2001)
19. Rote, G.: Strictly convex drawings of planar graphs. In: *SODA*, pp. 728–734. SIAM (2005)
20. Royden, H.L., Fitzpatrick, P.M.: *Real analysis*. NY, fourth. edn, Pearson modern classic, Pearson, New York (2010)
21. Schnyder, W.: Embedding planar graphs on the grid. In: *SODA*, pp. 138–148. SIAM (1990)
22. Schnyder, W., Trotter, W.T.: Convex embeddings of 3-connected plane graphs. *Abstracts AMS* **13**, 502 (1992)
23. Thomassen, C.: A refinement of Kuratowski’s theorem. *J. Comb. Theory, Ser. B* **D**(3), 245–253 (1984)
24. Tutte, W.T.: How to draw a graph. *Proc. London Math. Soc.* **13**, 743–768 (1963)
25. Zhang, H., He, X.: Compact visibility representation and straight-line grid embedding of plane graphs. In: Dehne, F., Sack, J.-R., Smid, M. (eds.) *WADS 2003*. LNCS, vol. 2748, pp. 493–504. Springer, Heidelberg (2003). https://doi.org/10.1007/978-3-540-45078-8_43