Journal of Graph Algorithms and Applications
http://jgaa.info/ vol. 21, no. 5, pp. 915–937 (2017)
DOI: 10.7155/jgaa.00444

Non-aligned Drawings of Planar Graphs

Therese Biedl[1] *Claire Pennarun*[2]

[1]David R. Cheriton School of Computer Science, University of Waterloo, Waterloo, Canada
[2]Univ. Bordeaux, CNRS, LaBRI, UMR 5800, F-33400 Talence, France

Abstract

A *non-aligned* drawing of a graph is a drawing where no two vertices are in the same row or column. Auber et al. showed that not all planar graphs have a non-aligned planar straight-line drawing in the $n \times n$-grid. They also showed that such a drawing exists if up to $n-3$ edges may have a bend.

In this paper, we give algorithms for non-aligned planar drawings that improve on the results by Auber et al. In particular, we give such drawings on an $n \times n$-grid with at most $\frac{2n-5}{3}$ bends, and we study what grid-size can be achieved if we insist on having straight-line drawings.

Submitted:	Reviewed:	Revised:	Accepted:	Final:
November 2016	March 2017	April 2017	June 2017	June 2017

Published: October 2017

Article type: Regular paper
Communicated by: Y. Hu and M. Nöllenburg

Research supported by NSERC. A preliminary version appeared at GD'16.

E-mail addresses: biedl@uwaterloo.ca (Therese Biedl) claire.pennarun@labri.fr (Claire Pennarun)

1 Introduction

At the GD 2015 conference, Auber et al. [3] introduced the concept of *rook-drawings*: These are drawings of a graph on an $n \times n$-grid (i.e., a grid with n rows and n columns) such that no two vertices are in the same row or the same column (thus, if the vertices were rooks on a chessboard, then no vertex could beat any other). They showed that not all planar graphs have a planar straight-line rook-drawing, and then gave a construction of planar rook-drawings with at most $n-3$ bends. From now on, all drawings are required to be planar.

In this paper, we continue the study of rook-drawings. Note that if a graph has no straight-line rook-drawing, then we can relax the restrictions in two possible ways. We could either, as Auber et al. did, allow to use bends for some of the edges, and try to keep the number of bends small. Or we could increase the grid-size and ask what size of grid can be achieved for straight-line drawings in which no two vertices share a row or a column; this type of drawing is known as *non-aligned drawing* [3]. A rook-drawing is then a non-aligned drawing on an $n \times n$-grid.

Existing results: Apart from the paper by Auber et al. [3], non-aligned drawings have arisen in a few other contexts. Alamdari and Biedl [1] showed that every graph that has an inner rectangular drawing also has a non-aligned drawing. These drawings are rectangle-of-influence drawings (in which any axis-aligned rectangle defined by adjacent vertices is empty) and can hence be assumed to be on an $n \times n$-grid. In particular, every 4-connected planar graph with at most $3n-7$ edges therefore has a rook-drawing (see Section 3.1 for details). Di Giacomo et al. [12] created so-called upward-rightward drawings of planar graphs in an $O(n^4) \times O(n^4)$-grid; one can easily see that these are automatically non-aligned drawings (we will give more details and improve their result in Section 2.2.1). Finally, there have been studies about drawing graphs with the opposite goal, namely, creating as many collinear vertices as possible [10].

Our results: In this paper, we show the following (the bounds listed here are upper bounds; see the sections for tighter bounds):

- In Section 2.1, we show that every planar graph has a non-aligned straight-line drawing on an $n^2 \times n^2$-grid. This is achieved by taking any weak barycentric representation (for example, the one by Schnyder [19]), scaling it by a big enough factor, and then moving vertices slightly so that they have distinct coordinates while maintaining a weak barycentric representation.

- We prove in Section 2.2 that every planar graph has a non-aligned straight-line drawing in an $n \times \frac{1}{2}n^3$-grid. This is achieved by creating drawings with the canonical ordering [11] in a standard fashion (similar to [9]). However, we pre-compute all the x-coordinates (and in particular, make them a permutation of $\{1, \ldots, n\}$), and then argue that with the standard construction the slopes do not get too big, and hence the height is quadratic. Modifying the construction a bit, we can also achieve that all y-coordinates are distinct and that the height is cubic.

- In Section 2.3, we show that any so-called nested-triangle graph has a non-aligned straight-line drawing on an $n \times (\frac{4}{3}n - 1)$-grid.
- Finally, in Section 3, we prove that every planar graph has a rook-drawing with at most $\frac{2n-5}{3}$ bends. This is achieved via creating a so-called rectangle-of-influence drawing of a modification of the graph, and arguing that each modification can be undone while adding only one bend. Our bounds are even better for 4-connected planar graphs (see Section 3.1). In particular, every 4-connected planar graph has a rook-drawing with at most one bend (and more generally, the number of bends is no more than the number of so-called filled triangles). Our algorithm for general planar graphs relies on finding a so-called independent-filled-hitting set of edges, which is done in more details in Section 3.3.
- A conclusion and open problems are presented in Section 4.

2 Non-aligned straight-line drawings

In this section, all drawings are required to be straight-line drawings.

2.1 Non-aligned drawings on an $n^2 \times n^2$-grid

We first show how to construct non-aligned drawings on an $n^2 \times n^2$-grid by scaling and perturbing a so-called *weak barycentric representation* (reviewed below).

In the following, a vertex v is assigned to a triplet of non-negative integer coordinates $(p_0(v), p_1(v), p_2(v))$. For two vertices u, v and $i = 0, 1, 2$, we say that $(p_i(u), p_{i+1}(u)) <_{\text{lex}} (p_i(v), p_{i+1}(v))$ if either $p_i(u) < p_i(v)$, or $p_i(u) = p_i(v)$ and $p_{i+1}(u) < p_{i+1}(v)$. Note that in this section, addition on the subscripts is done modulo 3.

Definition 1 (Weak barycentric representation [19]) *A* weak barycentric representation *of a graph G is an injective function $\mathcal{P}$ that maps each $v \in V(G)$ to a point $(p_0(v), p_1(v), p_2(v)) \in \mathbb{N}_0^3$ such that*

- *$p_0(v) + p_1(v) + p_2(v) = c$ for every vertex v, where c is a constant independent of the vertex,*
- *for each edge (u, v) and each vertex $z \neq \{u, v\}$, there is some $k \in \{0, 1, 2\}$ such that $(p_k(u), p_{k+1}(u)) <_{lex} (p_k(z), p_{k+1}(z))$ and $(p_k(v), p_{k+1}(v)) <_{lex} (p_k(z), p_{k+1}(z))$.*

Theorem 1 ([19]) *Every planar graph with n vertices has a weak barycentric representation with $c = n - 1$. Furthermore, $0 \leq p_i(v) \leq n - 2$ for all vertices $v \in V$ and all $i \in \{0, 1, 2\}$.*

Observe that weak barycentric representations are preserved under scaling, i.e., if we have a weak barycentric representation $\mathcal{P}$ (say with constant c), then we can scale all assigned coordinates by the same factor N and obtain another weak barycentric representation (with constant $c \cdot N$). We need to do slightly more, namely scale and "twist", as detailed in the following lemma.

Lemma 1 *Let G be a graph with a weak barycentric representation $\mathcal{P}$ with $\mathcal{P} = \big((p_0(v), p_1(v), p_2(v))_{v\in V}\big)$. Let $N \geq 1 + \max_{v\in V}\{\max_{i=0,1,2} p_i(v)\}$ be a positive integer. Define $\mathcal{P}'$ to be the assignment $p'_i(v) := N \cdot p_i(v) + p_{i+1}(v)$ for $i = 0, 1, 2$. Then $\mathcal{P}'$ is also a weak barycentric representation.*

Proof: We need to check the following properties:

(a) For some constant c we have $p'_1(v) + p'_2(v) + p'_3(v) = c$ for all vertices v.

(b) $\mathcal{P}'$ is injective.

(c) For each edge (u, v) and each vertex $z \neq \{u, v\}$, there is some $k \in \{0, 1, 2\}$ such that $(p'_k(u), p'_{k+1}(u)) <_{\text{lex}} (p'_k(z), p'_{k+1}(z))$ and $(p'_k(v), p'_{k+1}(v)) <_{\text{lex}} (p'_k(z), p'_{k+1}(z))$.

(a) Let c_P be the constant of $\mathcal{P}$. Then for each vertex v, $p'_1(v) + p'_2(v) + p'_3(v) = N\,(p_1(v) + p_2(v) + p_3(v)) + p_1(v) + p_2(v) + p_3(v) = N \cdot c_P + c_P$, which is a constant.

(b) Let $\{u, v\}$ be two vertices of G, $u \neq v$. Since $\mathcal{P}$ is injective, we know that there exists $i \in \{0, 1, 2\}$ such that $p_i(u) \neq p_i(v)$. Without loss of generality, $p_i(u) > p_i(v)$. Since all coordinates p_i are integers, $p_i(u) \geq p_i(v) + 1$. Thus $N \cdot p_i(u) \geq N \cdot p_i(v) + N > N \cdot p_i(v) + p_{i+1}(v) - p_{i+1}(u)$ by $p_{i+1}(v) < N$ and $p_{i+1}(u) \geq 0$. Thus $p'_i(u) > p'_i(v)$ and $\mathcal{P}'$ is injective.

(c) Let (u, v) be an edge of G and $z \neq \{u, v\}$ a vertex of G. Since $\mathcal{P}$ is a weak barycentric representation, there is some $k \in \{0, 1, 2\}$ such that $(p_k(u), p_{k+1}(u)) <_{\text{lex}} (p_k(z), p_{k+1}(z))$ and $(p_k(v), p_{k+1}(v)) <_{\text{lex}} (p_k(z), p_{k+1}(z))$.
We only show the claim for u, and have two cases:

- $p_k(u) < p_k(z)$: As in part (b), then $p'_k(u) < p'_k(z)$.
- $p_k(u) = p_k(z)$: Then $p_{k+1}(u) < p_{k+1}(z)$ and $p'_k(u) = Np_k(u) + p_{k+1}(u) = Np_k(z) + p_{k+1}(u) < Np_k(z) + p_{k+1}(z) = p'_k(z)$.

So either way $p'_k(u) < p'_k(z)$ and hence $(p'_k(u), p'_{k+1}(u)) <_{\text{lex}} (p'_k(z), p'_{k+1}(z))$. □

Applying this to Schnyder's weak barycentric representation, we now have:

Theorem 2 *Every planar graph has a non-aligned straight-line planar drawing on an $(n(n-2)) \times (n(n-2))$-grid.*

Proof: Let $\mathcal{P} = \big((p_0(v), p_1(v), p_2(v))_{v\in V}\big)$ be the weak barycentric representation of Theorem 1; we know that $0 \le p_i(v) \le n-2$ for all v and all i. Now apply Lemma 1 with $N = n-1$ to obtain the weak barycentric representation $\mathcal{P}'$ with $p'_i(v) = (n-1)p_i(v) + p_{i+1}(v)$. Observe that $p'_i(v) \le (n-1)(n-2) + (n-2) = n(n-2)$. Also, $p'_i(v) \ge 1$ since not both $p_i(v)$ and $p_{i+1}(v)$ can be 0. (More precisely, $p_i(v) = 0 = p_{i+1}(v)$ would imply $p_{i+2}(v) = n-1$, contradicting $p_{i+2}(v) \le n-2$.)

As shown by Schnyder [19], mapping each vertex v to point $(p'_0(v), p'_1(v))$ gives a planar straight-line drawing of G. By the above, this drawing has the desired grid-size. It remains to show that it is non-aligned, i.e., for any two vertices u, v and any $i \in \{0, 1\}$, we have $p'_i(u) \ne p'_i(v)$. Assume after possible renaming that $p_i(u) \le p_i(v)$. We have two cases:

- If $p_i(u) < p_i(v)$, then $p_i(u) \le p_i(v) - 1$ since $\mathcal{P}$ assigns integers. Thus $N \cdot p_i(u) \le N \cdot p_i(v) - N < N \cdot p_i(v) - p_{i+1}(u) + p_{i+1}(v)$ since $p_{i+1}(u) < N$ and $p_{i+1}(v) \ge 0$. Therefore $p'_i(u) < p'_i(v)$.
- If $p_i(u) = p_i(v)$, then $p_{i+1}(u) \ne p_{i+1}(v)$ (else the three coordinates of u and v would be the same, which is impossible since $\mathcal{P}$ is an injective function). Then $p'_i(u) = N \cdot p_i(u) + p_{i+1}(u) \ne N \cdot p_i(v) + p_{i+1}(v) = p'_i(v)$.

□

2.2 Non-aligned drawings on an $n \times f(n)$-grid

We now show how to build non-aligned drawings for which the width is the minimum-possible n, and the height is $\approx \frac{1}{2}n^3$. We use the well-known canonical ordering for *triangulated plane graphs*, i.e., graphs for which the planar embedding is fixed and all faces (including the outer-face) are triangles. We hence assume throughout that G is triangulated; we can achieve this by adding edges and delete them in the obtained drawing.

The *canonical ordering* [11] of such a graph is a vertex order $v_1, \dots, v_n$ such that $\{v_1, v_2, v_n\}$ is the outer-face, and for any $3 \le k \le n$, the graph G_k induced by $v_1, \dots, v_k$ is 2-connected. This implies that v_k has at least 2 *predecessors* (i.e., neighbours in G_{k-1}), and its predecessors form an interval on the outer-face of G_{k-1}. We assume (after possible renaming) that v_1 is the neighbour of v_2 found in clockwise order on the outer-face, and enumerate the outer-face of graph G_{k-1} in clockwise order as $c_1, \dots, c_L$ with $c_1 = v_1$ and $c_L = v_2$. Then the predecessors of v_k consist of $c_\ell, \dots, c_r$ for some $1 \le \ell < r \le L$; we call c_ℓ and c_r the *leftmost* and *rightmost* predecessors of v_k (see also Figure 1a). In this section, $x(v)$ and $y(v)$ denote the x- and y-coordinates of a vertex v, respectively.

2.2.1 Distinct x-coordinates

We first give a construction that achieves distinct x-coordinates in $\{1, \dots, n\}$ (but y-coordinates may coincide). Let $v_1, \dots, v_n$ be a canonical ordering. The goal is to build a straight-line planar drawing of the graph G_k induced by $v_1, \ldots, v_k$ using induction on k. The key idea is to define *all* x-coordinates beforehand. Orient the edges of G as follows. Direct (v_1, v_2) as $v_1 \to v_2$. For $k \geq 3$, if c_r is the rightmost predecessor of v_k, then direct all edges from predecessors of v_k towards v_k, with the exception of (v_k, c_r), which is directed $v_k \to c_r$.

By induction on k, one easily shows that the orientation of G_k is acyclic, with unique source v_1 and unique sink v_2, and the outer-face directed $c_1 \to \cdots \to c_L$.[1] Find a topological order $x : V \to \{1, \ldots, n\}$ of the vertices, i.e., if $u \to v$ then $x(u) < x(v)$ (this always exists since the graph is acyclic). We use this topological order as our x-coordinates, and hence have $x(v_1) = 1$ and $x(v_2) = n$. (We thus have two distinct vertex-orderings: one defined by the canonical ordering, which is used to compute y-coordinates, and one defined by the topological ordering derived from the canonical ordering, which directly gives the x-coordinates.)

Now construct a planar straight-line drawing of G_k that respects these x-coordinates by induction on k (see also Figure 1b). Start with v_1 at $(1, 2)$, v_3 at $(x(v_3), 2)$ and v_2 at $(n, 1)$.

For $k \geq 3$, let c_ℓ and c_r be the leftmost and rightmost predecessors of v_{k+1}. Notice that $x(c_\ell) < \cdots < x(c_r)$ due to our orientation, which in particular implies that for any $\ell \leq j \leq r$, the upward ray from c_j intersects no other vertex or edge. Let y^* be the smallest integer value such that any c_j, for $\ell \leq j \leq r$, can "see" the point $p = (x(v_{k+1}), y^*)$ in the sense that that the line segment from c_j to p intersects no other vertices or edges. Such a y^* exists since the upward ray from c_j is empty: by tilting this ray slightly, c_j can also see all sufficiently high points on the vertical line $\{x = x(v_{k+1})\}$. Placing v_{k+1} at $(x(v_{k+1}), y^*)$ hence gives a planar drawing of G_{k+1}.

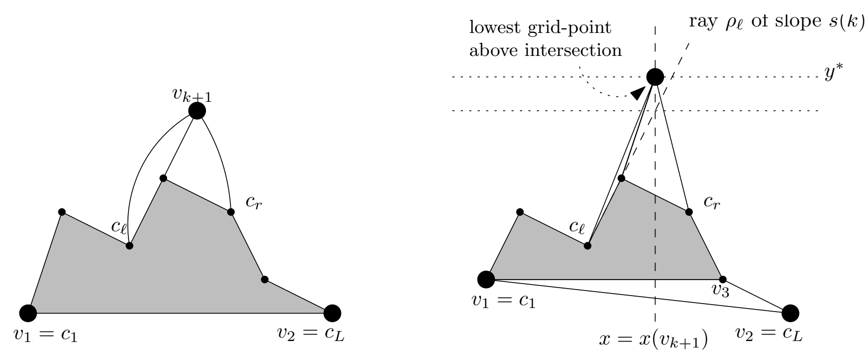

(a) Illustration of a canonical order. (b) Finding a y-coordinate for v_{k+1}.

Figure 1: Drawing algorithm to find distinct x-coordinates in $\{1, \ldots, n\}$.

[1] For readers familiar with how a Schnyder wood T_1, T_2, T_3 can be obtained from a canonical ordering: The orientation is the same as $T_1^{-1} \cup T_2 \cup T_3$, and hence acyclic [19].

To analyze the height of this construction, we bound the slopes of the segments representing the edges.

Lemma 2 *Define $s(k) := k - 3$ for $k \geq 3$. All edges on the outer-face of the constructed drawing of G_k have slope at most $s(k)$ for $k \geq 3$.*

Proof: Clearly this holds for $k = 3$ and $s(3) = 0$: one slope is zero, and the other two are negative. Now assume it holds for some $k \geq 3$, and let $c_\ell, \dots, c_r$ be the predecessors of v_{k+1}.

Consider one predecessor c_j for $\ell \leq j < r$, and define ρ_j to be the ray of slope $s(k)$ starting from c_j. We claim that there are no vertices or edges anywhere in the open region above ray ρ_j. To see this, recall that the upward vertical ray from c_j intersects no vertices or edges. Rotate this vertical ray rightward until for the first time it intersects a vertex or edge; this happens at the latest when it hits c_L. Since we stop at the first intersection, it must necessarily happen when the current ray (call it ρ) intersects a vertex of the outer-face, say c_x. We know $j < x$ since the ray rotated rightward and $x(c_1) < \cdots < x(c_L)$. By induction, every edge on the outerface-path $c_j, c_{j+1}, \dots, c_x$ has slope at most $s(k)$. Therefore ray ρ, whose slope is the weighted average of the slopes of these edges, also has slope at most $s(k)$. This means that either $\rho = \rho_j$ or ρ has rotated past ρ_j; either way the claim holds.

Now consider the point where ray ρ_ℓ intersects the vertical line $\{x = x(v_{k+1})\}$, and set y' to be the smallest integer y-coordinate that is strictly above this intersection. For $j = \ell, \dots, r-1$, we have $x(c_\ell) \leq x(c_j) < x(v_{k+1})$, and therefore point $p = (x(v_{k+1}), y')$ is to the right of c_j and in the open region above ρ_j. Since this region contains no other vertices or edges, therefore c_j can see p.

We claim that point p can also see c_r. This holds because edge (c_{r-1}, c_r) has slope at most $s(k)$, and (due to the chosen edge directions) $x(c_{r-1}) < x(v_{k+1}) < x(c_r)$. Therefore point p is above (c_{r-1}, c_r), and can be connected to both of them without intersection. Also note that line segment (p, c_r) therefore has a smaller slope than (c_{r-1}, c_r), and in particular slope less than $s(k)$.

So point $(x(v_{k+1}), y')$ can see all vertices $c_\ell, \dots, c_r$, and the value of $y^* = y(v_{k+1})$ is no bigger than y'. We already argued that edge (v_{k+1}, c_r) has slope less than $s(k) \leq s(k+1)$, so we only must argue the slope of the unique other new outer-face edge (c_ℓ, v_{k+1}). Since y' is the smallest integer y-coordinate above the point where ρ_ℓ intersects the line $\{x = x(v_{k+1})\}$, we have

$$y^* \leq y' \leq y(c_l) + (x(v_{k+1}) - x(c_l)) \cdot s(k) + 1. \tag{1}$$

By $x(v_{k+1}) - x(c_l) \geq 1$ the slope of (c_l, v_{k+1}) is at most

$$\frac{y^* - y(c_l)}{x(v_{k+1}) - x(c_l)} \leq s(k) + \frac{1}{x(v_{k+1}) - x(c_l)} \leq s(k) + 1 = s(k+1)$$

as desired. □

Vertex v_n has x-coordinate at most $n-1$, and the edge from v_1 to v_n has slope at most $s(n) = n-3$. This shows that the y-coordinate of v_n is at most $2 + (n-2) \cdot (n-3)$. Since triangle $\{v_1, v_2, v_n\}$ bounds the drawing, this gives:

Theorem 3 *Every planar graph has a planar straight-line drawing on an* $n \times (2 + (n-2)(n-3))$*-grid such that all vertices have distinct* x*-coordinates.*

While this theorem per se is not useful for non-aligned drawings, we find it interesting from a didactic point of view: It proves that polynomial coordinates can be achieved for straight-line drawings of planar graphs, and requires for this only the canonical ordering, but neither the properties of Schnyder trees [19] nor the details of how to "shift" that is needed for other methods using the canonical ordering (e.g. [9, 11]). We believe that our bound on the height is much too big, and that the true height is $o(n^2)$ and possibly $O(n)$.

Our theorem also has implications for *upward-rightward drawings* introduced by Di Giacomo et al. [12].[2] These are drawings of directed graphs where any directed edge goes either rightward or upward (or both). Di Giacomo et al. achieve this by creating drawings of the underlying undirected graph where for every vertex v the (closed) 1st and 3rd quadrant relative to v contains no other vertices; in particular such drawings are non-aligned. They create such drawings on an $O(n^4) \times O(n^4)$-grid, by first creating drawings on an $O(n^2) \times O(n^2)$-grid for which all y-coordinates are distinct, and then stretching, scaling and rotating appropriately. If we start with the drawing in Theorem 3 instead, then we can create upward-rightward drawings with smaller area.

Corollary 1 *Let* G *be an embedded planar digraph with* n *vertices. Then* G *admits an upward-rightward straight-line grid drawing* Γ *in area* $2n^3 \times 2n^3$*.*

Proof: Ignore the edge-directions. Obtain a drawing of G from Theorem 3 and rotate it by 90°, so it now resides on an $n^2 \times n$-grid and all y-coordinates are distinct. Now apply Phase 2 and Phase 3 of the algorithm in [12]. Thus, stretch the drawing vertically by a factor w, where $w \leq n^2$ is the width of the drawing. We now have a drawing on an $n^2 \times n^3$-grid. Next scale the drawing by $\sqrt{2}$ and then rotate by 45°. We now have a drawing on (at most) a $2n^3 \times 2n^3$-grid. As argued by Di Giacomo et al. [12], this drawing is a rightward-upward drawing. □

2.2.2 Non-aligned drawings

We now modify the above construction slightly to achieve distinct y-coordinates. Define the exact same x-coordinates and place v_1 and v_2 as before. To place vertex v_{k+1}, let y^* be the smallest y-coordinate such that point $(x(v_{k+1}), y^*)$ can see all predecessors of v_{k+1}, and such that none of $v_1, \dots, v_k$ is in row $\{y = y^*\}$. Clearly this gives a non-aligned drawing. It remains to bound how much this increases the height.

Lemma 3 *Define* $s'(k) := \sum_{i=1}^{k-2} i = \frac{1}{2}(k-1)(k-2)$ *for* $k \geq 3$*. All edges on the outer-face of the constructed non-aligned drawing of* G_k *have slope at most* $s'(k)$ *for* $k \geq 3$*.*

[2] Many thanks to the anonymous reviewer whose comments inspired this result.

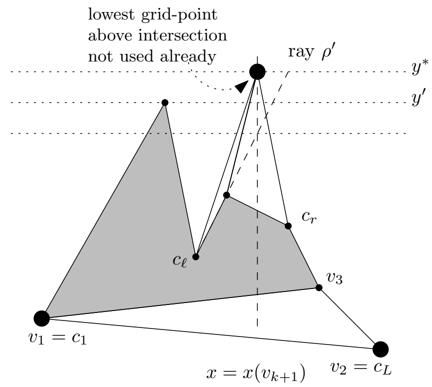

Figure 2: Finding a y-coordinate for v_{k+1} that has not been used by earlier vertices.

Proof: The claim clearly holds for $k = 3$, since v_3 is placed with y-coordinate 3 and therefore (v_1, v_3) has slope at most $1 = s(3)$. Now let $k \geq 3$ and consider the time when adding v_{k+1} with predecessors $c_\ell, \ldots, c_r$, and define ρ' to be the ray of slope $s'(k)$ emanating from c_ℓ. Let y' be the smallest integer coordinate above the intersection of ρ' with the vertical line $\{x = x(v_{k+1})\}$. As in Lemma 2, one argues that $p' = (x(v_{k+1}), y')$ can see all of $c_\ell, \ldots, c_r$.

We may or may not be able to use point p' for v_{k+1}, depending on whether some other vertices were in the row $\{y = y'\}$. Observe that $y' \geq 3$, because $y(c_\ell) \geq 2$ and $s'(k) \geq 1$. Therefore neither v_1 nor v_2 had y-coordinate y', which leaves at most $k-2$ vertices that could be in row y' or higher. In particular

$$y^* \leq y' + (k-2) \leq y(c_\ell) + (x(v_{k+1}) - x(c_\ell)) \cdot s'(k) + 1 + (k-2) \qquad (2)$$

Reformulating as before shows that the slope of (c_ℓ, v_{k+1}) is at most

$$\frac{y^* - y(c_l)}{x(v_{k+1}) - x(c_l)} \leq s'(k) + \frac{k-1}{x(v_{k+1}) - x(c_l)} \leq s'(k) + k - 1 = s'(k+1),$$

which concludes the proof. □

Edge (v_1, v_n) has slope at most $\frac{1}{2}(n-1)(n-2)$. Since $x(v_n) - x(v_1) \leq n-2$ and $y(v_1) = 2$, therefore the height is at most $2 + \frac{1}{2}(n-1)(n-2)^2$.

Theorem 4 *Every planar graph has a non-aligned straight-line drawing in an* $n \times \left(2 + \frac{1}{2}(n-1)(n-2)^2\right)$*-grid.*

Comparing this to Theorem 2, we see that the aspect ratio is much worse, but the area is smaller. We suspect that the method results in a smaller height than the proved upper bound: Equation (2) is generally not tight, and so a smaller slope-bound (implying a smaller height) is likely to hold.

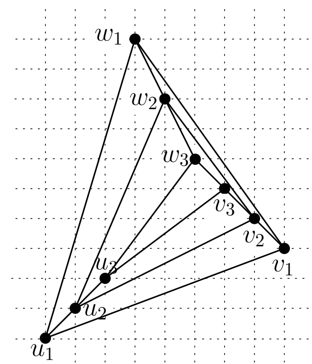

Figure 3: A non-aligned straight-line drawing of a nested-triangle graph with $k = 3$ on an 9×11-grid.

2.3 The special case of nested triangles

We now turn to non-aligned drawings of a special graph class. Define a *nested-triangle graph* G as follows. G has $3k$ vertices for some $k \geq 1$, say $\{u_i, v_i, w_i\}$ for $i = 1, \ldots, k$. Vertices $\{u_i, v_i, w_i\}$ form a triangle (for $i = 1, \ldots, k$). We also have paths $u_1, u_2, \ldots, u_k$ as well as $v_1, v_2, \ldots, v_k$ and $w_1, w_2, \ldots, w_k$. With this the graph is 3-connected; we assume that its outer-face is $\{u_1, v_1, w_1\}$. All interior faces that are not triangles may or may not have a diagonal in them, and there are no restrictions on which diagonal (if any). Nested-triangle graphs are of interest in graph drawing because they are the natural lower-bound graphs for the area of straight-line drawings [13].

Theorem 5 *Any nested-triangle graph with $n = 3k$ vertices has a non-aligned straight-line drawing on an $n \times (\frac{4}{3}n - 1)$-grid.*

Proof: The 4-cycle $\{w_k, v_k, v_{k-1}, w_{k-1}\}$ may or may not have a diagonal in it; after possible exchange of $w_1, \ldots, w_k$ and $v_1, \ldots, v_k$ we assume that there is no edge between v_{k-1} and w_k. For $i = 1, \ldots, k$, place u_i at (i, i), vertex v_i at $(3k + 1 - i, k + i)$, and w_i at $(k + i, 4k + 1 - 2i)$ (see Figure 3). The x- and y-coordinates are all distinct. The x-coordinates range from 1 to n, and the maximal y-coordinate is $4k - 1 = \frac{4}{3}n - 1$. It is easy to check that all interior faces are drawn strictly convex, with the exception of $\{v_k, v_{k-1}, w_{k-1}, w_k\}$ which has a $180°$ angle at v_k, but our choice of naming ensured that there is no edge (v_{k-1}, w_k). Thus any diagonal inside these 4-cycles can be drawn without overlap. Since G is planar, two edges joining vertices of different triangles cannot cross. Thus G is drawn without crossing on an $n \times \left(\frac{4}{3}n - 1\right)$-grid. $\square$

In particular, notice that the octahedron is a nested-triangle graph (for $k = 2$) and this construction gives a non-aligned straight-line drawing on a 6×7-grid. This is clearly optimal since it has no straight-line rook-drawing [3].

We conjecture that this construction gives the minimum-possible height for nested-triangle graphs among all non-aligned straight-line drawings.

3 Rook-drawings with bends

We now construct rook-drawings with bends; as before we do this only for triangulated plane graphs. The main idea is to find rook-drawings with only one bend for 4-connected triangulated plane graphs, then convert any planar graph into a 4-connected triangulated plane graph by subdividing few edges and re-triangulating, and finally argue that the drawing for it, modified suitably, gives a rook-drawing with few bends.

We need a few definitions first. Fix a triangulated plane graph G. A *separating triangle* is a triangle that has vertices both strictly inside and strictly outside the triangle. G is *4-connected* (i.e., cannot be made disconnected by removing 3 vertices) if and only if it has no separating triangle. A *filled triangle* [6] of G is a triangle that has vertices strictly inside. A triangulated plane graph has at least one filled triangle (namely, the outer-face) and every separating triangle is also a filled triangle. We use f_G to denote the number of filled triangles of the graph G.

A *rectangle-of-influence (RI) drawing* is a straight-line drawing such that for any edge (u, v), the minimum axis-aligned rectangle containing u and v is *empty*, i.e. contains no other vertex of the drawing in its relative interior [3]. The following is known:

Theorem 6 ([6]) *Let G be a triangulated 4-connected plane graph and let e be an edge on the outer-face. Then $G - e$ has a planar RI-drawing.*

Moreover, the drawing is non-aligned and on an $n \times n$-grid, the ends of e are at $(1, n)$ and $(n, 1)$, and the other two vertices on the outer-face are at $(2, 2)$ and $(n - 1, n - 1)$.

Figure 4b illustrates such a drawing of a graph. The latter part of the claim is not specifically stated in [6], but can easily be inferred from the construction (see also a simpler exposition in [4]).

RI-drawings are useful because they can be deformed (within limits) without introducing crossings. We say that two drawings Γ and Γ' of a graph have *the same relative coordinates* if for any two vertices v and w, we have $x_\Gamma(v) < x_\Gamma(w)$ if and only if $x_{\Gamma'}(v) < x_{\Gamma'}(w)$, and $y_\Gamma(v) < y_\Gamma(w)$ if and only if $y_{\Gamma'}(v) < y_{\Gamma'}(w)$, where $x_\Gamma(v)$ (resp. $y_\Gamma(v)$) denotes the x-coordinate (resp. y-coordinate) of v in Γ. The following result appears to be folklore; we sketch a proof for completeness.

[3] In the literature there are four kinds of RI-drawings, depending on whether points on the boundary of the rectangle are allowed or not (open vs. closed RI-drawings), and whether an edge (u, v) must exist if $R(u, v)$ is empty (strong vs. weak RI-drawings). The definition here corresponds to open weak RI-drawings.

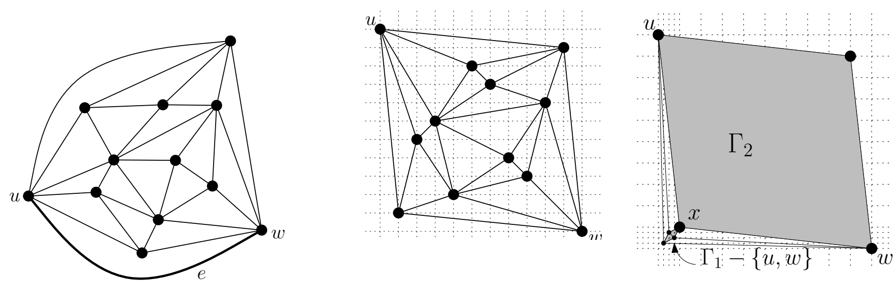

(a) A triangulated plane graph G; e is one of the edges of its outer-face.

(b) An RI-drawing of $G - e$ satisfying the conditions of Theorem 6.

(c) Combining two RI-drawings if all separating triangles contain (u, w).

Figure 4: RI-drawings.

Observation 1 *Let* Γ *be an RI-drawing. If* Γ' *is a straight-line drawing with the same relative coordinates as* Γ*, then* Γ' *is an RI-drawing, and it is planar if and only if* Γ *is.*

Proof: The claim on the RI-drawing was shown by Liotta et al. [17]. It remains to argue planarity. Assume that edge (u, v) crosses edge (w, z) in an RI-drawing. Since all rectangles-of-influence are empty, this happens if and only if (up to renaming) we have $x(w) \le x(u) \le x(v) \le x(z)$ and $y(u) \le y(w) \le y(z) \le y(v)$. This only depends on the relative orders of u, v, w, z, and hence a transformation maintaining relative coordinates also maintains planarity. □

We need a slight strengthening of Theorem 6.

Lemma 4 *Let* G *be a triangulated plane graph, let* $e \in E$ *be an edge on the outer-face, and assume all separating triangles of* G *contain* e*. Then* $G - e$ *has a planar RI-drawing. Moreover, the drawing is non-aligned and on an* $n \times n$*-grid, the ends of* e *are at* $(1, n)$ *and* $(n, 1)$*, and the other two vertices on the outer-face are at* $(2, 2)$ *and* $(n-1, n-1)$*.*

Proof: We proceed by induction on the number of separating triangles of G. In the base case, G is 4-connected and the claim holds by Theorem 6. For the inductive step, assume that $T = \{u, x, w\}$ is a separating triangle. By assumption it contains e, say $e = (u, w)$. Let G_1 be the graph consisting of T and all vertices inside T, and let G_2 be the graph obtained from G by removing all vertices inside T. Apply induction to both graphs. In drawing Γ_2 of $G_2 - e$, vertex x is on the outer-face and hence (after possible reflection) placed at $(2, 2)$. Now insert a (scaled-down) copy of the drawing Γ_1 of G_1, minus vertices u and w, in the square $(1, 2] \times (1, 2]$ (see Figure 4c). Since x was (after possible reflection) in the top-right corner of $\Gamma_1 - \{u, w\}$, the two copies of x can be identified. One easily verifies that this gives an RI-drawing, because within each drawing the relative coordinates are unchanged, and the two drawings have disjoint x-range and y-range except at u and w. Finally, re-assign coordinates to the vertices while keeping relative coordinates intact so that we have an $n \times n$-grid; by Observation 1 this gives a planar RI-drawing. □

We also need a technical lemma on being able to shift vertices in an RI-drawing under a special condition.

Lemma 5 *Let Γ be a planar RI-drawing. Let x be an interior vertex of degree 4 with neighbours u_1, u_2, u_3, u_4 that form a 4-cycle. Assume that none of x, u_1, u_2, u_3, u_4 share a grid-line. Then we can move x to a point on grid-lines of its neighbours and obtain a planar RI-drawing.*

Proof: We assume that the naming is such that u_1, u_2, u_3, u_4 are the neighbours of x in counter-clockwise order. We will also use the notation $R(u, v)$ for the (open) axis-aligned rectangle whose diagonally opposite corners are u and v.

Consider the four quadrants relative to x, using the open sets. Each neighbour of x shares no grid-line with x and hence belongs to some quadrant. We claim that (after suitable renaming) u_i is in quadrant i for $i = 1, 2, 3, 4$. For if any quadrant is empty, then either all four of u_1, u_2, u_3, u_4 are within two consecutive quadrants (in case of which x is outside cycle u_1, u_2, u_3, u_4, violating planarity), or two consecutive vertices of u_1, u_2, u_3, u_4 are in diagonally opposite quadrants (in case of which x is inside their rectangle-of-influence, violating the condition of an RI-drawing). So each quadrant contains at least one of the four vertices, implying that each contains exactly one of them.

Consider the five columns of x, u_1, u_2, u_3, u_4; for ease of description we will denote these columns by $1, \ldots, 5$ (in order from left to right), even though their actual x-coordinates may be different. Likewise let $1, \ldots, 5$ be the five rows of x, u_1, u_2, u_3, u_4. Since x has a neighbour in each quadrant, it must be at $(3, 3)$. The open set $R\big((2, 2), (4, 4)\big)$ contains none of u_1, u_2, u_3, u_4, so the cycle C_u formed by u_1, u_2, u_3, u_4 goes around $R\big((2, 2), (4, 4)\big)$. The only vertex inside C_u is x, which implies that no vertex other than u_3 or u_4 can be at $(2, 2)$ or $(4, 2)$. But not both u_3 and u_4 can be in row 2, so we may assume that point $x' := (2, 2)$ contains no vertex (the other case is similar). Move x to x', which puts it on grid-lines of two of its neighbours.

We claim that we obtain an RI-drawing, and verify the conditions for the four edges (x, u_i) separately:

- Vertex u_3 is on row 1 or 2 and column 1 or 2. But by choice of x' it is not at $(2, 2)$. No matter where it is, rectangle $R(u_3, x)$ has $x' = (2, 2)$ inside or on the boundary, and therefore $R(u_3, x') \subset R(u_3, x)$ is empty since (u_3, x) is an edge in an RI-drawing.
- Vertices u_3 and u_4 are on rows 1 and 2, but not on the same row, so $R(u_3, u_4)$ contains points that are between rows 1 and 2. Further, u_3

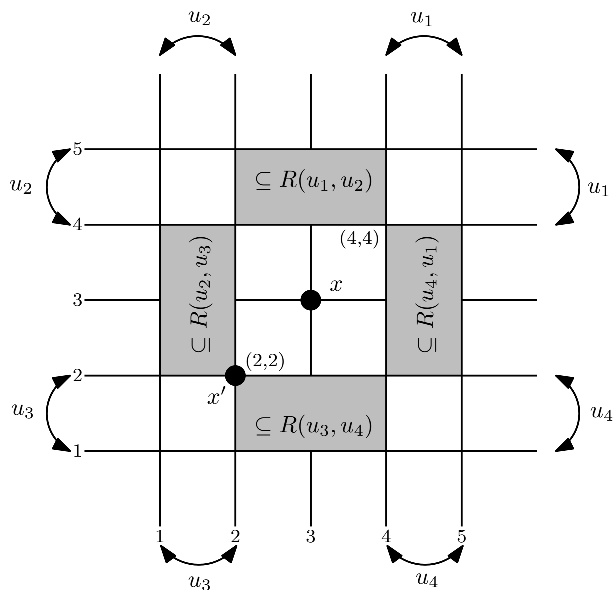

Figure 5: Moving x to two grid-lines among its neighbours.

is on column 2 or left while u_4 is on column 4 or right, so $R(u_3, u_4)$ includes $R\big((2,1),(4,2)\big)$. So the empty rectangle $R(u_3, u_4)$ contains point $(2,2) = x'$ and therefore includes rectangle $R(x', u_4)$.

- Similarly one shows that $R(x', u_2)$ is empty.
- It remains to show that $R(x', u_1)$ is empty, regardless of the position of u_1 within quadrant 1. To do so, split $R(x', u_1)$ into parts and observe that all of them are empty. We already saw that $R_1 := R\big((2,2),(4,4)\big)$ is empty. We also claim that $R_2 := R\big((2,4),(4,5)\big)$ is empty. This holds because u_1 and u_2 are on rows 4 and 5 (but not on the same row) and this rectangle hence is within $R(u_1, u_2)$. Similarly one shows that $R_3 := R\big((4,2),(5,4)\big)$ is empty.

  Notice that $R_1 \cup R_2 \cup R_3$ contains $R(x', u_1)$, unless u_1 is at $(5,5)$. But in the latter case $R_4 := R\big((4,4),(5,5)\big) \subset R(x, u_1)$ is empty. So either way $R(x', u_1)$ is contained within the union of empty rectangles and therefore is empty.

This concludes the proof. □

3.1 4-connected planar graphs

Combining Theorem 6 with Observation 1, we immediately obtain:

Theorem 7 *Let G be a triangulated 4-connected plane graph. Then G has a planar rook-drawing with at most one bend.*

Proof: Fix an arbitrary edge e on the outer-face, and apply Theorem 6 to obtain an RI-rook-drawing Γ of $G-e$. It remains to add in edge $e=(u,v)$. One end u of e is in the top-left corner, and the leftmost column contains no other vertex. The other end v is in the bottom-right corner, and the bottommost row contains no other vertex. We can hence route (u,v) by going vertically from u and horizontally from v, with the bend in the bottom-left corner. □

Corollary 2 *Let G be a 4-connected planar graph. Then G has a rook-drawing with at most one bend, and with no bend if G is not triangulated.*

Proof: If G is triangulated then the result was shown above, so assume G has at least one face of degree 4 or more. Since G is 4-connected, one can add edges to G such that the result G' is triangulated and 4-connected [7]. Pick a face incident to an added edge e as outer-face of G', and apply Theorem 6 to obtain an RI-drawing of $G'-e$. Deleting all edges in $G'-G$ gives the result. □

Since we have only one bend, and the ends of the edge (u,v) that contains it are the top-left and bottom-right corner, we can remove the bend by stretching.

Theorem 8 *Every 4-connected planar graph has a non-aligned planar drawing on an $n \times (n^2-3n+4)$-grid and on a $(2n-2)\times(2n-2)$-grid.*

Proof: Let Γ be the RI-drawing of $G-(u,v)$ with u at $(1,n)$ and v at $(n,1)$. Relocate u to point $(1,n^2-3n+4)$. The resulting drawing is still a planar RI-drawing by Observation 1. Now $y(u)-y(v)=(n-2)(n-1)+1$, hence the line segment from u to v has slope less than $-(n-2)$, and is therefore above point $(n-1,n-1)$ (and with that, also above all other vertices of the drawing). So we can add this edge without violating planarity, and obtain a non-aligned straight-line drawing of G (see Figure 6a).

For the other result, start with the same drawing Γ. Relocate u to $(1,2n-2)$ and v to $(2n-2,1)$. The line segment from u to v has slope -1 and crosses Γ only between points $(n-1,n)$ and $(n-1,n)$, where no points of Γ are located. So we obtain a non-aligned planar straight-line drawing (see Figure 6b). □

3.2 Constructing rook-drawings with few bends

We now explain the construction of a (poly-line) rook-drawing for a triangulated plane graph G with at least 5 vertices. We proceed as follows:

1. Find a small independent-filled-hitting set E_f.

   Here, an *independent-filled-hitting set* is a set of edges E' such that (i) every filled triangle has at least one edge in E' (we say that E' *hits* all filled triangles), and (ii) every face of G has at most one edge in E' (we say that E' is *independent*). We can show the following bound on $|E'|$:

   Lemma 6 *Any triangulated plane graph G of order n has an independent-filled-hitting set of size at most*

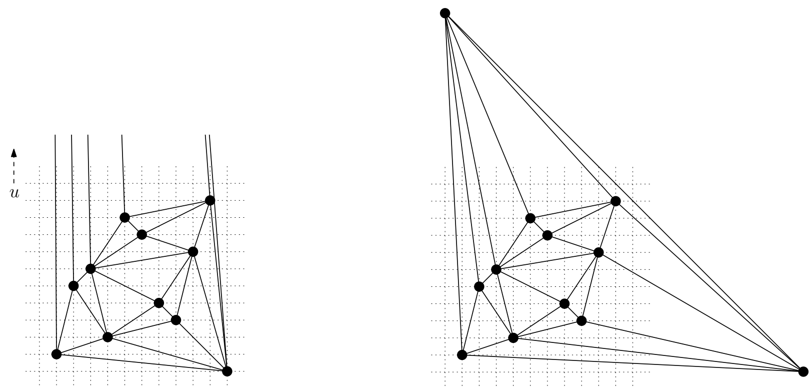

(a) A non-aligned drawing of width n.

(b) A non-aligned drawing on a $(2n-2) \times (2n-2)$-grid.

Figure 6: Transforming Figure 4b into straight-line drawings.

- f_G *(where* f_G *is the number of filled triangles of* G*), and it can be found in* $O(n)$ *time,*
- $\frac{2n-5}{3}$*, and it can be found in* $O((n \log n)^{1.5}\sqrt{\alpha(n,n)})$ *time. Here* α *is the slow-growing inverse Ackermann function.*

The proof of this lemma requires detours into matchings and 4-coloring; to keep the flow of the algorithm we defer it to Section 3.3.

2. Since the outer-face is a filled triangle, there exists one edge $e_o \in E_f$ that belongs to the outer-face. Define $E_s := E_f - \{e_o\}$ and notice that E_s contains no outer-face edges since E_f is independent.

3. As done in some previous papers [16, 8], remove separating triangles by subdividing all edges $e \in E_s$, and re-triangulate by adding edges from the subdivision vertex (see Figure 7). Let V_x be the new set of vertices, and let G_1 be the new graph. Observe that G_1 may still have separating triangles, but all those separating triangles contain e_o since E_f hits all filled triangles.

4. By Lemma 4, $G_1 - e_o$ has a non-aligned RI-drawing Γ where the ends of e_o are at the top-left and bottom-right corner.

5. Transform Γ into drawing Γ' so that the relative orders stay intact, the original vertices (i.e., vertices of G) are on an $n \times n$-grid and the subdivision vertices (i.e., vertices in V_x) are in-between.

   This can be done by enumerating the vertices in x-order, and assigning new x-coordinates in this order, increasing to the next integer for each original vertex and increasing by $\frac{1}{|V_x|+1}$ for each subdivision vertex. Similarly

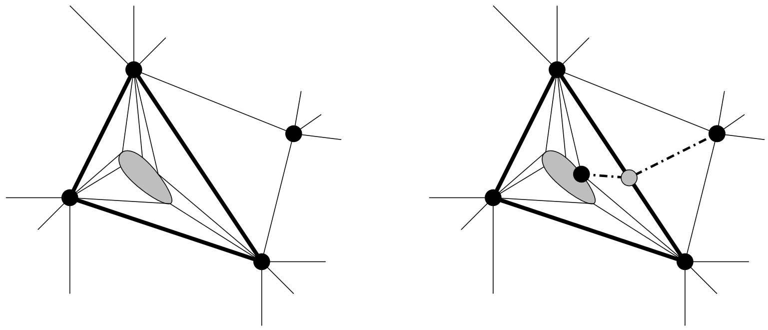

Figure 7: A separating triangle removed by subdividing one of its edges and re-triangulating.

update the y-coordinates (see Figure 8a). Drawing Γ' is still a non-aligned RI-drawing, and the ends of e_o are on the top-left and bottom-right corner.

6. Let e be an edge in E_s with subdivision vertex x_e. Since e is an interior edge of G, x_e is an interior vertex of G_1. Now move x_e to some integer grid-point nearby, which is possible due to Lemma 5. Note that the neighbours of x_e are not in V_x, since E_s is independent. We can thus apply this operation independently to all subdivision-vertices.

7. Now replace each subdivision-vertex x_e by a bend, connected to the ends of e along the corresponding edges from x_e (see Figure 8b). (Sometimes, as is the case in the example, we could also simply delete the bend and draw edge e straight-line.) None of the shifting changed positions for vertices of G, so we now have a rook-drawing of $G - e_o$ with bends. The above shifting of vertices does not affect outer-face vertices, so the ends of e_o are still in the top-left and bottom-right corner. As the final step draw e_o by drawing vertically from one end and horizontally from the other; these segments are not occupied by the rook-drawing.

We added one bend for each edge in E_f. By Lemma 6, we can compute a set E_f with $|E_f| \leq f_G$ and $|E_f| \leq \frac{2n-5}{3}$ (neither bound is necessarily smaller than the other), and hence have:

Theorem 9 *Any planar graph G of order n has a planar rook-drawing with at most b bends, with $b \leq \min\{\frac{2n-5}{3}, f_G\}$.*

Remark that the algorithm proposed to compute such a drawing has polynomial run-time, but is not linear-time due to the time complexity of finding a small independent-filled-hitting set. All other steps (including computing an RI-drawing) can be done in linear time.

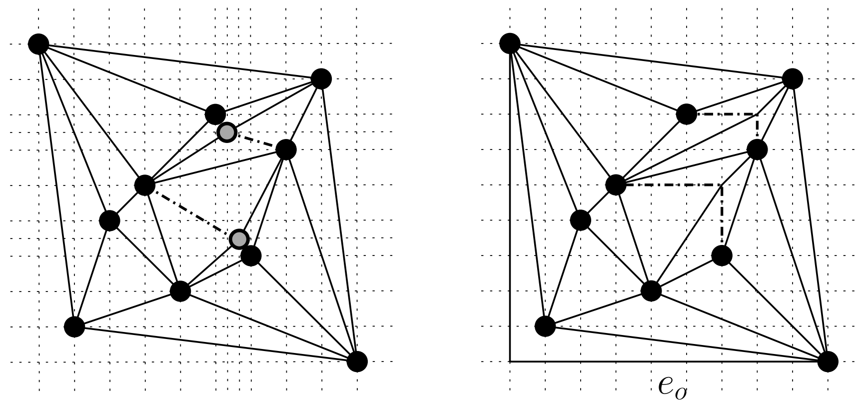

(a) Reorder such that subdivision vertices (grey) are not at grid points.

(b) Subdivision vertices shifted to integer grid-points, and adding e_0.

Figure 8: Creating non-aligned drawings with few bends.

3.3 Independent-filled-hitting sets

It remains to prove Lemma 6. Recall that we want to find a set E_f that *hits* all filled triangles (i.e., contains at least one edge of each filled triangles) and is *independent* (i.e., no face contains two edges of E_f).

Our first result shows how to find a matching of size at most f_G in linear time. The existence of such a matching could easily be proved using the 4-color theorem (see below for more details), but with a different approach we can find it in linear time.

Lemma 7 *Any triangulated plane graph G has an independent-filled-hitting set E_f of size at most f_G. It can be found in linear time.*

Proof: We prove a slightly stronger statement, namely, that we can find such a set E_f and additionally (i) we can prescribe which edge e_o on the outer-face is in E_f, and (ii) every separating triangle has *exactly* one edge in E_f.

We proceed by induction on the number of filled triangles. If G has only one (namely, the outer-face), then use the prescribed edge e_o; this satisfies all claims. Now assume G has multiple filled triangles, and let $T_1, \ldots, T_k$ (for $k \geq 1$) be the maximal separating triangles in the sense that no other separating triangle contains T_i inside. Define G_i (for $i = 1, \ldots, k$) to be the graph consisting of T_i and all vertices inside T_i. Since we chose maximal separating triangles, graphs $G_1, \ldots, G_k$ are disjoint. Let the *skeleton* G^{skel} of G be the graph obtained from G by removing the interior of $T_1, \ldots, T_k$.

Let $(G^{\text{skel}})^*$ be the dual graph of G^{skel}, i.e., it has a vertex for every face of G^{skel} and a dual edge e^* for any edge e that connects the two faces that e is incident to. Since G^{skel} is triangulated, its dual graph is 3-regular and 3-connected, and therefore has a perfect matching M by Petersen's theorem (see e.g. [5]). If the dual edge e_o^* of e_o is not in M, then find an alternating cycle that contains e_o^* and swap matching/non-matching edges so that e_o^* is in M.

For every maximal separating triangle T_i of G, exactly one edge e_i of T_i has its dual edge in matching M, since T_i forms a face in G^{skel}. Find an independent-filled-hitting set $E_f(G_i)$ of G_i that contains e_i recursively. Combine all these independent-filled-hitting sets into one set and add e_o to it (if not already in it); the result is set E_f. Every filled triangle of G is either the outer-face or a filled triangle of one of the subgraphs G_i, so this hits all filled triangles. Also, every filled triangle contains exactly one edge of E_f, so $|E_f|$ is as desired. Finally if f is a face of G, then it is either an inner face of one of the subgraphs G_i or a face of G^{skel}. Either way at most one edge of f is in E_M; therefore E_f is independent.

The time complexity is dominated by splitting the graph into its 4-connected components at all separating triangles, which can be done in linear time [15], and by finding the perfect matching in the dual graph, which can also be done in linear time [5]. □

We now give the second result, which gives a different (and sometimes better) bound at the price of being slower to compute. Recall that the 4-color-theorem for planar graphs states that we can assign colors $\{1, 2, 3, 4\}$ to vertices of G such that no edge has the same color at both endpoints [2]. Define M_1 to be all those edges where the ends are colored $\{1, 2\}$ or $\{3, 4\}$, set M_2 to be all those edges where the ends are colored $\{1, 3\}$ or $\{2, 4\}$ and set M_3 to be all those edges where the ends are colored $\{1, 4\}$ or $\{2, 3\}$. Since every face of G is a triangle and colored with 3 different colors, the following is easy to verify:

Observation 2 *For $i = 1, 2, 3$, edge set M_i contains exactly one edge of each triangle.*

So each M_i is an independent-filled-hitting set. Define E_i to be the set of edges obtained by deleting from M_i all those edges that do not belong to any filled triangle. Clearly E_i is still an independent-filled-hitting set, and since it contains exactly one edge of each filled triangle, its size is also at most f_G. The best of these three disjoint independent-filled-hitting sets contains at most $\frac{2n-5}{3}$ edges, due to the following:

Lemma 8 *Any triangulated plane graph G with $n \geq 4$ has at most $2n - 5$ edges that belong to a filled triangle.*

We would like to mention first that Cardinal et al. [8] gave a very similar result, namely that every triangulated plane graph contains at most $2n-7$ edges that belong to a separating triangle. (This immediately implies that at most $2n-4$ edges belong to a filled triangle.) Their proof (not given in [8], but kindly shared via private communication) is quite different from ours below, and does not seem to adapt easily to filled triangles. Since our proof is quite short, we give it below despite the rather minor improvement.

Proof: For every edge e that belongs to a filled triangle, fix an arbitrary filled triangle T containing e and assign to e the face that is incident to e and inside triangle T. We claim that no face f can have been assigned to two edges e_1 and e_2. Assume for contradiction that it did, so there are two distinct filled triangles T_1 and T_2, both having f inside and with e_i incident to T_i for $i = 1, 2$. Since face f is inside both T_1 and T_2, one of the two triangles (say T_2) is inside the other (say T_1). Since e_1 belongs to f, it is on or inside T_2, but since it is also on T_1 (which contains T_2 inside) it therefore must be on T_2. But then T_2 contains e_1 and e_2, and these two (distinct) edges hence determine the three vertices of T_2. But these three vertices also belong to the triangular face f, and so T_2 is an inner face and hence not a filled triangle by $n \geq 4$: contradiction.

With that, we can assign a unique inner face to every edge of a filled triangle, therefore in total there are at most $2n - 5$ of them. □

We could find the smallest of edge sets E_1, E_2, E_3 by 4-coloring the graph (which can be done in $O(n^2)$ time [18]), but a better approach is the following: Compute the dual graph G^*, and assign weight 1 to an edge e^* if the corresponding edge e in G belongs to a filled triangle; else assign weight 0 to e^*. Now find a minimum-weight perfect matching M in G^*; this can be done in $O((n \log n)^{1.5}\sqrt{\alpha(n,n)})$ time [14] since we have $m \in O(n)$ and maximum weight 1. Deleting from M all edges of weight 0 then gives an independent-filled-hitting set E_f, and it has size at most $\min\{f_G, \frac{2n-5}{3}\}$ since one of the three perfect matching of G^* induced by a 4-coloring would have at most this weight. We can hence conclude:

Corollary 3 *Every planar graph G has an independent-filled-hitting set of size at most $\frac{2n-5}{3}$. It can be found in $O((n \log n)^{1.5}\alpha(n,n))$ time.*

It seems quite plausible that such an independent-filled-hitting set could be computed in linear time, for example by modifying the perfect-matching-algorithm for 3-regular biconnected planar graphs [5] to take into account 0-1-edge-weights, with edges having weight 1 only if they occur in a non-trivial 3-edge-cut. This remains for future work.

Combining Lemma 7 and Corollary 3 gives Lemma 6.

4 Conclusion

In this paper, we continued the work on planar rook-drawings initiated by Auber et al. [3]. We constructed planar rook-drawings with at most $\frac{2n-5}{3}$ bends; the number of bends can also be bounded by the number of filled triangles. We also considered drawings that allow more rows and columns while keeping vertices on distinct rows and columns; we proved that such non-aligned planar straight-line drawings always exist with maximum area $O(n^4)$. As for open problems, the most interesting question is lower bounds. No planar graph is known that needs more than one bend in a planar rook-drawing, and no planar graph is known that needs more than $2n + 1$ grid-lines in a planar non-aligned drawing. The "obvious" approach of taking multiple copies of the octahedron fails because the property of having a rook-drawing is not closed under taking subgraphs: if vertices are added, then they could "use up" extraneous grid-lines in the drawing of a subgraph. We conjecture that the $n \times (\frac{4}{3}n - 1)$-grid achieved for nested-triangle graphs is optimal for planar straight-line non-aligned drawings with width n.

References

[1] S. Alamdari and T. Biedl. Planar open rectangle-of-influence drawings with non-aligned frames. In *Graph Drawing (GD'11)*, volume 7034 of *Lecture Notes in Computer Science*, pages 14–25. Springer-Verlag, 2012. doi:10.1007/978-3-642-25878-7_3.

[2] K. Appel and W. Haken. Every planar map is four colorable. I. Discharging. *Illinois J. Math.*, 21(3):429–490, 1977. doi:10.1090/conm/098.

[3] D. Auber, N. Bonichon, P. Dorbec, and C. Pennarun. Rook-drawings for plane graphs. *Journal of Graph Algorithms and Applications*, 21(1):103–120, 2017. doi:10.7155/jgaa.00402.

[4] T. Biedl and M. Derka. The (3,1)-ordering for 4-connected planar triangulations. *Journal of Graph Algorithms and Applications*, 20(2):347–362, 2016. doi:10.7155/jgaa.00396.

[5] T. Biedl, P. Bose, E. Demaine, and A. Lubiw. Efficient algorithms for Petersen's matching theorem. *J. Algorithms*, 38(1):110–134, 2001. doi:10.1006/jagm.2000.1132.

[6] T. Biedl, A. Bretscher, and H. Meijer. Rectangle of influence drawings of graphs without filled 3-cycles. In *Graph Drawing (GD'99)*, volume 1731 of *Lecture Notes in Computer Science*, pages 359–368. Springer-Verlag, 1999. doi:10.1007/3-540-46648-7_37.

[7] T. Biedl, G. Kant, and M. Kaufmann. On triangulating planar graphs under the four-connectivity constraint. *Algorithmica*, 19(4):427–446, 1997. doi:10.1007/PL00009182.

[8] J. Cardinal, M. Hoffmann, V. Kusters, C. D. Tóth, and M. Wettstein. Arc diagrams, flip distances, and hamiltonian triangulations. *Computational Geometry*, 2017. doi:10.1016/j.comgeo.2017.06.001.

[9] M. Chrobak and G. Kant. Convex grid drawings of 3-connected planar graphs. *Internat. J. Comput. Geom. Appl.*, 7(3):211–223, 1997. doi:10.1142/S0218195997000144.

[10] G. Da Lozzo, V. Dujmović, F. Frati, T. Mchedlidze, and V. Roselli. Drawing planar graphs with many collinear vertices. In *International Symposium on Graph Drawing and Network Visualization*, pages 152–165. Springer, 2016. doi:10.1007/978-3-319-50106-2_13.

[11] H. de Fraysseix, J. Pach, and R. Pollack. How to draw a planar graph on a grid. *Combinatorica*, 10:41–51, 1990. doi:10.1007/BF02122694.

[12] E. Di Giacomo, W. Didimo, M. Kaufmann, G. Liotta, and F. Montecchiani. Upward-rightward planar drawings. In *Information, Intelligence, Systems and Applications (IISA 2014)*, pages 145–150. IEEE, 2014. doi:10.1109/IISA.2014.6878792.

[13] D. Dolev, F. Leighton, and H. Trickey. Planar embedding of planar graphs. *Advances in Computing Research*, 2:147–161, 1984.

[14] H. Gabow and R. Tarjan. Faster scaling agorithms for general graph matching problems. *Journal of the ACM*, 38:815–853, 1991. `doi:10.1145/115234.115366`.

[15] G. Kant. A more compact visibility representation. *Internat. J. Comput. Geom. Appl.*, 7(3):197–210, 1997. `doi:10.1142/S0218195997000132`.

[16] M. Kaufmann and R. Wiese. Embedding vertices at points: Few bends suffice for planar graphs. *Journal of Graph Algorithms and Applications*, 6(1):115–129, 2002. `doi:10.7155/jgaa.00046`.

[17] G. Liotta, A. Lubiw, H. Meijer, and S. Whitesides. The rectangle of influence drawability problem. *Comput. Geom.*, 10(1):1–22, 1998. `doi:10.1016/S0925-7721(97)00018-7`.

[18] N. Robertson, D. Sanders, P. Seymour, and R. Thomas. The four-colour theorem. *J. Combin. Theory Ser. B*, 70(1):2–44, 1997. `doi:10.1006/jctb.1997.1750`.

[19] W. Schnyder. Embedding planar graphs on the grid. In *ACM-SIAM Symposium on Discrete Algorithms (SODA '90)*, pages 138–148, 1990.