Algorithmica (2019) 81:2527–2556
https://doi.org/10.1007/s00453-018-00542-9

Universal Slope Sets for 1-Bend Planar Drawings

Patrizio Angelini[1] · Michael A. Bekos[1] · Giuseppe Liotta[2] · Fabrizio Montecchiani[2]

Received: 24 July 2017 / Accepted: 29 December 2018 / Published online: 8 January 2019

Abstract

We prove that every set of $\Delta - 1$ slopes is 1-bend universal for the planar graphs with maximum vertex degree Δ. This means that any planar graph with maximum degree Δ admits a planar drawing with at most one bend per edge and such that the slopes of the segments forming the edges can be chosen in any given set of $\Delta - 1$ slopes. Our result improves over previous literature in three ways: Firstly, it improves the known upper bound of $\frac{3}{2}(\Delta - 1)$ on the 1-bend planar slope number; secondly, the previously known algorithms compute 1-bend planar drawings by using sets of $O(\Delta)$ slopes that may vary depending on the input graph; thirdly, while these algorithms typically minimize the slopes at the expenses of constructing drawings with poor angular resolution, we can compute drawings whose angular resolution is at least $\frac{\pi}{\Delta - 1}$, which is worst-case optimal up to a factor of $\frac{3}{4}$. Our proofs are constructive and give rise to a linear-time drawing algorithm.

Keywords Graph drawing · Slope number · 1-Bend planar drawings

An extended abstract of this paper has been presented at the 33rd International Symposium on Computational Geometry (SoCG 2017). Research supported in part by the project: "Algoritmi e sistemi di analisi visuale di reti complesse e di grandi dimensioni" - Ricerca di Base 2018, Dipartimento di Ingegneria dell'Università degli Studi di Perugia, and by the DFG Grants Ka812/17-1 and Ka812/18-1.

✉ Fabrizio Montecchiani
fabrizio.montecchiani@unipg.it

Patrizio Angelini
angelini@informatik.uni-tuebingen.de

Michael A. Bekos
bekos@informatik.uni-tuebingen.de

Giuseppe Liotta
giuseppe.liotta@unipg.it

1 Wilhelm-Schickhard-Institut für Informatik, Universität Tübingen, Tübingen, Germany

2 Dipartimento di Ingegneria, University of Perugia, Perugia, Italy

1 Introduction

This paper is concerned with planar drawings of graphs such that each edge is a poly-line with few bends, each segment has one of a limited set of possible slopes, and the drawing has a good angular resolution, i.e. it forms large angles between consecutive edges incident to a common vertex. Besides their theoretical interest, visualizations with these properties find applications in software engineering and information visualization (see, e.g., [15,34,49]). For example, degree-4 planar graphs (that is, with maximum degree four) are widely used in database design, where they are typically represented by *orthogonal* drawings, i.e. drawings such that every edge is a polygonal chain of horizontal and vertical segments. Clearly, orthogonal drawings of degree-4 planar graphs are optimal both in terms of angular resolution and in terms of the number of distinct slopes for the edges. A classical result in the graph drawing literature is that every degree-4 planar graph, except for the octahedron, admits a planar orthogonal drawing with at most two bends per edge [5].

It is immediate to see that more than two slopes are needed in any planar drawing of a graph with vertex degree $\Delta \geq 5$. The *k-bend planar slope number* of a graph G with degree Δ is defined as the minimum number of distinct slopes that are sufficient to compute a crossing-free drawing of G with at most k bends per edge. Keszegh et al. [39] generalize the aforementioned technique by Biedl and Kant [5] and prove that for any $\Delta \geq 5$, the 2-bend planar slope number of a degree-Δ planar graph is $\lceil \Delta/2 \rceil$; the construction in their proof has an optimal angular resolution, that is $\frac{2\pi}{\Delta}$.

For the case of drawings with one bend per edge, Keszegh et al. [39] also show an upper bound of 2Δ and a lower bound of $\frac{3}{4}(\Delta - 1)$ on the 1-bend planar slope number, while a recent paper by Knauer and Walczak [40] improves the upper bound to $\frac{3}{2}(\Delta - 1)$. Both these papers use a similar technique: The graph is first realized as a contact representation with T-shapes [13], which is then transformed into a planar drawing where vertices are points and edges are poly-lines with at most one bend. The set of slopes depends on the initial contact representation and may change from graph to graph; also, each slope is either very close to horizontal or very close to vertical, which in general gives rise to a bad angular resolution. Note that Knauer and Walczak [40] also consider subclasses of planar graphs. In particular, they prove that any set of $\lceil \frac{\Delta}{2} \rceil$ slopes can be used to construct 1-bend outerplanar drawings of outerplanar graphs with $\Delta > 2$, and present an upper bound of $\Delta + 1$ and a lower bound of $\frac{2}{3}(\Delta - 1)$ for planar bipartite graphs. In addition, Durocher and Mondal [23] prove that $\frac{2}{3}\Delta$ slopes and Δ slopes always suffice to compute 1-bend planar drawings of 2-trees and of planar 3-trees, respectively.

In this paper, we study the trade-off between number of slopes, angular resolution, and number of bends per edge in a planar drawing of a graph with maximum degree Δ. Our main contribution is a constructive proof that any set of $\Delta - 1$ slopes is universal for 1-bend planar drawings of planar graphs with maximum degree $\Delta \geq 4$. More precisely, we prove the following.

Theorem 1 *For any $\Delta \geq 4$, every set S of $\Delta - 1$ slopes is universal for 1-bend planar drawings of planar graphs with maximum degree Δ. Namely, every such graph has a planar drawing with the following properties: (i) each edge has at most one bend;*

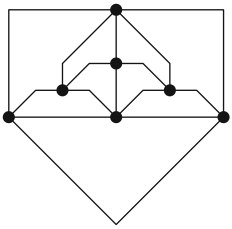

Fig. 1 A 1-bend planar drawing with 4 slopes and angular resolution $\frac{\pi}{4}$ of a graph with $\Delta = 5$.

(ii) each edge segment uses one of the slopes in S. Furthermore, a drawing with these properties can be constructed in linear time.

The first implication of Theorem 1 is an improvement on the upper bound of Knauer and Walczak [40] on the 1-bend planar slope number of planar graphs. In fact, our result, in conjunction with [35], implies that the 1-bend planar slope number of planar graphs with $n \geq 5$ vertices and maximum degree $\Delta \geq 3$ is at most $\Delta - 1$.

Further, by using an *equispaced* set of slopes, in which the minimum angle between any two slopes is $\frac{\pi}{\Delta-1}$, we can also guarantee that the constructed drawings have angular resolution $\frac{\pi}{\Delta-1}$. We formalize this observation in the following corollary of Theorem 1.

Corollary 1 *For any $\Delta \geq 4$, every planar graph with maximum degree Δ has a planar drawing with the following properties: (i) each edge has at most one bend; (ii) each edge segment uses one of the slopes of an equispaced set of $\Delta - 1$ slopes; and (iii) the minimum angle between any two consecutive edge segments incident to a vertex or a bend is at least $\frac{\pi}{\Delta-1}$. Furthermore, a drawing with these properties can be constructed in linear time.*

An implication of the proof given in [39] for the $\frac{3}{4}(\Delta-1)$ lower bound on the 1-bend planar slope number is that a 1-bend planar drawing may require angular resolution smaller than $\frac{4\pi}{3(\Delta-2)}$ (see Corollary 2 in p. 27). Hence, the result in Corollary 1 is worst-case optimal up to a multiplicative factor of at least $\frac{3}{4}$ (as Δ tends to infinity). We remark that our result also improves the best-known upper bound of $\frac{\pi}{4\Delta}$, proved by Duncan and Kobourov [22], on the angular resolution of 1-bend planar drawings (where the number of slopes is not limited). This comes at the cost of increased drawing area, since our algorithm may produce drawings with a non-polynomial area.

We prove Theorem 1 by using an approach that is conceptually different from that of Knauer and Walczak [40]: We do not construct an intermediate representation and then transform it into a 1-bend planar drawing, but we provide an algorithm that directly computes a 1-bend drawing of any planar graph with degree at most Δ on any set of $\Delta - 1$ slopes. Our proof is constructive and it gives rise to a linear-time algorithm, assuming the real RAM model of computation. Figure 1 shows a drawing computed by our algorithm with an equispaced set of slopes, while Figs. 8 and 15 show larger examples with different sets of slopes. The construction for triconnected planar graphs uses a variant of the shifting method of de Fraysseix, Pach, and Pollack [14]; this construction is the building block for the drawing algorithm for biconnected planar graphs, which is based on the SPQR-tree decomposition of the graph into its triconnected components (see, e.g., [15]). Finally, the result is extended to connected graphs by using a block-cutvertex tree decomposition as a guideline to assign subsets of the universal slope set to the different biconnected components of the input graph. If the graph is disconnected, since we use universal sets of slopes, the distinct connected components can be drawn independently.

Related Work The results on the slope number of graphs are mainly classified into two categories based on whether the constructed drawings are required to be planar or not. For a (planar) graph G with maximum degree Δ, the *slope number* (*planar slope number*) is the minimum number of slopes that are sufficient to compute a straight-line (planar) drawing of G. The slope number of non-planar graphs is lower bounded by $\lceil \Delta/2 \rceil$ [50] but it can be arbitrarily large, even when $\Delta = 5$ [1,20]. For $\Delta = 3$ this number is 4 [44], while it is unknown for $\Delta = 4$, to the best of our knowledge. Upper bounds on the slope number are known for complete graphs [50] and outer 1-planar graphs [16] (i.e., graphs that can be drawn in the plane such that each edge is crossed at most once, and all vertices are on the external boundary). Deciding whether a graph has slope number 2 is NP-complete [26]. Concerning poly-line drawings, Knauer and Walczak [40] proved that general graphs with maximum degree Δ have a 1-bend drawing with $\lceil \frac{\Delta}{2} \rceil + 1$ slopes, which improves a previous result by Dujmovic et al. [20].

For a planar graph G with maximum degree Δ, the planar slope number of G is lower bounded by $3\Delta - 6$ and upper bounded by $O(2^{\Delta})$ [39]. Improved upper bounds are known for certain subclasses of planar graphs, e.g., planar graphs with $\Delta \leq 3$ [18,19,36], outerplanar graphs with $\Delta \geq 4$ [41], partial 2-trees [42], and planar partial 3-trees [33]. Note that determining the planar slope number of a graph is hard in the existential theory of the reals [32].

The slope number problem has been studied also in the upward setting, i.e., for straight-line and poly-line drawings of directed graphs such that all edges flow towards a common direction (see, e.g., [4,28]). Upward drawings of ordered sets with straight-line edges and few slopes have been studied by Czyzowicz [10] and by Czyzowicz et al. [11]. Concerning 1-bend drawings, Di Giacomo et al. [17] proved that every series-parallel directed graph has an upward drawing with one bend per edge and Δ slopes. The construction is worst-case optimal in terms of the number of slopes, and it gives rise to drawings with optimal angular resolution $\frac{\pi}{\Delta}$. Also, Czyzowicz et al. [12] studied upward drawings of ordered sets with one bend per edge and few slopes.

Closely related to our problem is the problem of finding *d-linear* drawings of graphs, in which all angles (that are formed either between consecutive segments of an edge or between edge-segments incident to the same vertex) are multiples of $2\pi/d$. Bodlaender and Tel [8] showed that, for $d = 4$, an angular resolution of $2\pi/d$ implies d-linearity and that this is not true for any $d > 4$. Special types of d-linear drawings are the *orthogonal* [5,7,28,48] and the *octilinear* [2,3,45] drawings, for which $d = 2$ and $d = 4$ holds, respectively. As already recalled, Biedl and Kant [5], and independently Liu et al. [43], have shown that any planar graph with $\Delta \leq 4$ (except the octahedron) admits a planar orthogonal drawing with at most two bends per edge. Deciding whether a degree-4 planar graph has an orthogonal drawing with no bends is NP-complete [28], while it is solvable in polynomial time if one bend per edge is allowed [6] (see also the work by Felsner et al. [25]). On the other hand, octilinear drawings have been mainly studied in the context of metro map visualization and map schematization [46,47].

Nöllenburg [45] proved that deciding whether a given embedded planar graph with $\Delta \leq 8$ admits a straight-line planar octilinear drawing is NP-complete. Bekos et al. [2] showed that a planar graph with $\Delta \leq 5$ always admits a planar octilinear drawing with at most one bend per edge and that such drawings do not always exist if $\Delta \geq 6$. Note that in our work we generalize their positive result to any Δ. Later, Bekos et al. [3] studied bounds on the total number of bends of planar octilinear drawings.

Finally, trade-offs between number of bends, angular resolution, and area requirement of planar drawings of graphs with maximum degree Δ are, for example, studied in [9,21,22,24,27,30].

Paper Organization The rest of this paper is organized as follows. Preliminaries are given in Sect. 2. In Sect. 3, we describe a drawing algorithm for triconnected planar graphs. The technique is extended to biconnected and to general planar graphs in Sects. 4 and 5, respectively. Finally, in Sect. 6 we list some open problems.

2 Preliminaries

A graph $G = (V, E)$ containing neither loops nor multiple edges is *simple*. We consider simple graphs, if not otherwise specified. The *degree* of a vertex of G is the number of its neighbors. We say that G has *maximum degree* Δ if it contains a vertex with degree Δ but no vertex with degree larger than Δ. A graph is *connected* if for any pair of vertices there is a path connecting them, and is *k-connected*, $k \geq 1$, if the removal of any set of $k - 1$ vertices leaves it connected. A 2-connected (3-connected) graph is also called *biconnected* (*triconnected*, respectively).

A *drawing* Γ of a graph G maps each vertex of G to a point in the plane and each edge of G to a Jordan arc between its two endpoints. A drawing is *planar* if no two edges cross (except at common endpoints). A planar drawing divides the plane into connected regions, called *faces*. The unbounded one is called *outer face*. A graph is *planar* if it admits a planar drawing. A *planar embedding* of a planar graph is an equivalence class of planar drawings that combinatorially define the same set of faces and outer face.

The *slope* s of a straight line ℓ is the angle that a horizontal line needs to be rotated counter-clockwise in order to make it overlap with ℓ.[1] The slope of an edge-segment is the slope of the line containing it. Let S be a set of slopes sorted in increasing order; assume without loss of generality up to a rotation, that S contains the 0 angle, which we call *horizontal slope*. A *1-bend* planar drawing Γ of graph G *on* S is a planar drawing of G in which every edge is composed of at most two straight-line segments, each of which has a slope that belongs to S. We say that S is *equispaced* if the difference between any two consecutive slopes of S is $\frac{\pi}{|S|}$. For a vertex v in G, each slope $s \in S$ defines two different *rays* that emanate from v and have slope s. If s is the horizontal slope, then these rays are called *horizontal*. Otherwise, the one going upward is a *top* and the other one is a *bottom* ray of v. Consider a 1-bend planar drawing Γ of a graph

[1] Note that, formally, for a straight line ℓ, the angle that a horizontal line needs to be rotated counter-clockwise in order to make it overlap with ℓ is the *angle of incline* of ℓ, while its *slope s* is the tangent of this angle. However, for simplicity reasons, and similarly as in previous papers (see, e.g., [16,19,20]), we refer directly to the angle of incline of ℓ as its slope.

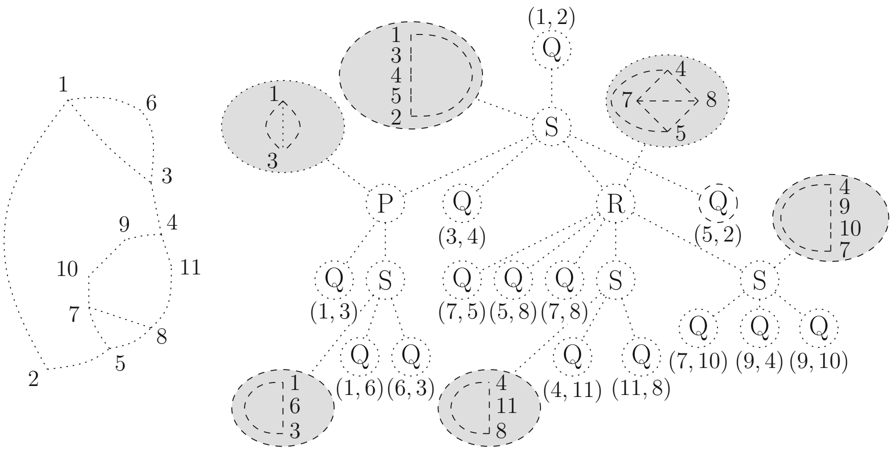

Fig. 2 A biconnected planar graph (left) and its SPQR-tree (right). For S-, P-, and R-nodes, the skeleton is depicted in the gray balloons and the reference edge is dashed; for Q-nodes the corresponding edge is shown

G and a ray r_v emanating from a vertex v of G. We say that r_v is *free* if there is no edge attached to v through r_v. We also say that r_v is *incident* to face f of Γ if r_v is free and the first face encountered when moving from v along r_v is f.

Let Γ be a 1-bend planar drawing of a graph and let e be an edge on the boundary of the outer face of Γ that has a horizontal segment. A *cut at e* is a strictly y-monotone curve that (i) starts at any point of the horizontal segment of e, (ii) ends at any point of a horizontal segment of an edge e' incident to the outer face of Γ, where $e' \neq e$, unless both sides of e are incident to the outer face of Γ, and (iii) crosses only horizontal segments of Γ. A cut at e is called *degenerate* if both sides of e are incident to the outer face of Γ and it does not cross any edge of Γ different from e.

Central in our approach is the canonical order of triconnected planar graphs [14,37]. Let $G = (V, E)$ be a triconnected plane graph and let $\Pi = (P_0, \ldots, P_m)$ be a partition of V 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 of G. For $k = 0, \ldots, m$, let G_k be the subgraph induced by $\cup_{i=0}^{k} P_i$ and denote by C_k the cycle delimiting the outer face of G_k. We say that Π is a *canonical order* of G if for each $k = 1, \ldots, m-1$ the following hold: (i) G_k is biconnected, (ii) all neighbors of P_k in G_{k-1} are on C_{k-1}, (iii) $|P_k| = 1$ or the degree of each vertex of P_k is 2 in G_k, and (iv) all vertices of P_k with $0 \leq k < m$ have at least one neighbor in P_j for some $j > k$. A canonical order of any triconnected planar graph can be computed in linear time [37].

An SPQR-tree $\mathcal{T}$ represents the decomposition of a biconnected graph G into its triconnected components [15]; refer to Fig. 2 for an illustration. Every triconnected component of G is associated with a node μ of $\mathcal{T}$. The triconnected component itself is called the *skeleton* of μ, denoted by G_μ^{skel}. A node μ in $\mathcal{T}$ can be of four different types: (i) μ is an *R-node*, if G_μ^{skel} is a triconnected graph, (ii) a simple cycle of length at least three classifies μ as an *S-node*, (iii) a bundle of at least three parallel edges classifies μ as a *P-node*, (iv) the leaves of $\mathcal{T}$ are *Q-nodes*, whose skeleton consists of two parallel edges. Neither two S- nor two P-nodes are adjacent in $\mathcal{T}$. For each node

ν that is adjacent to μ in $\mathcal{T}$, graph G_μ^{skel} contains a *virtual edge* that corresponds to ν; symmetrically, G_ν^{skel} contains a virtual edge that corresponds to μ. We also say that these two virtual edges correspond to each other. If we assume that $\mathcal{T}$ is rooted at a Q-node ρ, then the virtual edge of G_μ^{skel} that corresponds to its parent is the *reference edge* of μ, and its endpoints are the *poles* of μ. Note that the reference edge always exists, unless $\mu = \rho$. Every subtree $\mathcal{T}_\mu$ rooted at a node μ of $\mathcal{T}$ induces a subgraph G_μ of G called *pertinent*, which is described by $\mathcal{T}_\mu$ in the decomposition. We remark that the SPQR-tree of a biconnected planar graph can be computed in linear time [31].

Finally, the *BC-tree* $\mathcal{B}$ of a connected graph G represents the decomposition of G into its biconnected components. Namely, $\mathcal{B}$ has a *B-node* for each biconnected component of G and a *C-node* for each cutvertex of G. Each B-node is connected to the C-nodes that are part of its biconnected component.

3 Triconnected Planar Graphs

Let G be a triconnected planar graph with maximum degree $\Delta \geq 4$ and let S be any set of $\Delta - 1$ slopes containing the horizontal one (as already observed, this assumption is not restrictive). We consider the vertices of G according to a canonical order $\Pi = (P_0, \ldots, P_m)$. At each step $k = 1, \ldots, m$, we consider the planar graph G_k^- obtained by removing edge (v_1, v_2) from G_k. Let C_k^- be the path from v_1 to v_2 obtained by removing (v_1, v_2) from C_k. We seek to construct a 1-bend planar drawing of G_k^- on S satisfying the following invariants.

I.1 For any two vertices $u \in P_i$ and $v \in P_j$ with $0 \leq i \leq j \leq k$, one of the following hold:

(a) If $i = j$ or $j = 1$, then u and v have the same y-coordinate; also, if edge (u, v) exists, then it is drawn as a horizontal segment.

(b) Otherwise (i.e., $i < j$ and $j \neq 1$), vertex u is strictly below v; also, if edge (u, v) exists, then its bend-point (if any) is strictly above u, and below or at the same y-coordinate as v.

I.2 Every edge on C_k^- has a horizontal segment. Also, for any point p of this segment there is a strictly y-monotone curve on the outer face of the drawing that starts at p and ends at any point strictly above the drawing (i.e., all its interior points are above p and it does not cross the drawing).

I.3 Each vertex v on C_k^- has at least as many free top rays incident to the outer face of G_k^- as the number of its neighbors in $G \setminus G_k$.

We briefly discuss some direct implications of Invariants I.1–I.3. Invariant I.1 implies that no part of the drawing lies below vertices v_1 and v_2, which are horizontally aligned. In particular, Invariant I.1a implies that any edge between two vertices in the same path P_i, with $i \geq 1$, is drawn as a horizontal segment. Invariant I.2 implies that the drawing of C_k^- does not "spiralize". Finally, note that if Invariant I.2 holds for one point of the horizontal segment of an edge of C_k^-, then it also holds for any point of this segment.

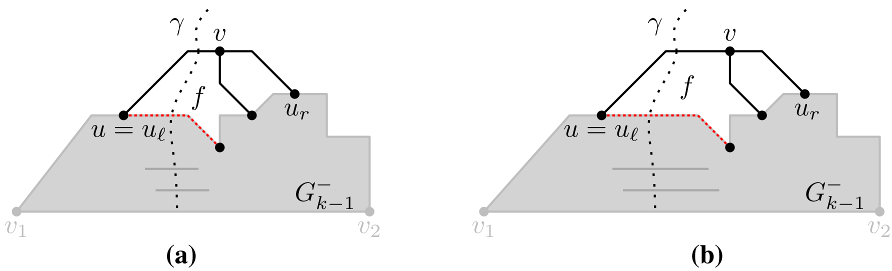

Fig. 3 Illustrations for Lemma 1; edge (w, z) in the proof is drawn dotted-red (Color figure online)

Once a 1-bend planar drawing on S of G^-_m satisfying Invariants I.1–I.3 has been constructed, a 1-bend planar drawing on S of $G = G^-_m \cup \{(v_1, v_2)\}$ can be obtained by drawing edge (v_1, v_2) as a polyline composed of two straight-line segments, one attaching at the first clockwise bottom ray of v_1 and the other one at the first anti-clockwise bottom ray of v_2. Note that, since $\Delta \geq 4$, there are at least three slopes in S, and thus these two rays cross at a point below v_1 and v_2 with finite coordinates. By Invariant I.1, adding this edge does not introduce any crossing.

In the following lemma we show another important property of any 1-bend planar drawing on S satisfying Invariants I.1–I.3.

Lemma 1 *Let Γ_k be a 1-bend planar drawing on S of G^-_k satisfying Invariants I.1–I.3 and let σ be any positive number. For any edge (u, v) of C^-_k such that u precedes v along path C^-_k, there exists a 1-bend planar drawing Γ'_k on S of G^-_k, satisfying Invariants I.1–I.3, such that: (i) the horizontal distance between u and v is increased by σ; (ii) the horizontal distance between any two other vertices that are consecutive along C^-_k is the same as in Γ_k.*

Proof We first show that there exists a cut of Γ_k at (u, v) that separates the subpath of C^-_k connecting v_1 to u from the subpath of C^-_k connecting v to v_2. We use this cut to construct Γ'_k as a copy of Γ_k in which all the horizontal segments that are crossed by the cut are elongated by σ. For an illustration refer to Fig. 3.

The existence of the cut is proved by induction on k. In the base case, $k = 1$ holds. In this case, each edge (u, v) in C^-_1 is drawn as a straight-line segment by Invariant I.1a, and thus any point of this edge can be used to define a degenerate cut. For the inductive case $k \geq 2$, assume that a cut exists for all edges of C^-_ν, with $1 \leq \nu < k$. Consider now an edge (u, v) that belongs to C^-_k. If (u, v) also belongs to C^-_{k-1}, then the existence of the cut is guaranteed by the induction hypothesis. Otherwise, at least one between u and v belongs to P_k, say v (the other case is symmetric). By Invariant I.2, edge (u, v) has a horizontal segment. Let $u \in P_i$ and $v \in P_k$, for some $i \leq k$. We claim that the horizontal segment of (u, v) is incident to v: If $i = k$, then (u, v) is drawn as a single horizontal segment by Invariant I.1a; otherwise, the bend-point of (u, v) cannot be at the same y-coordinate as u by Invariant I.1b, and thus it must be at the same y-coordinate as v, as otherwise there would be no horizontal segment.

Let f be the internal face of Γ_k to which edge (u, v) is incident to; this face is uniquely defined since G_k is biconnected and (u, v) is incident to the outer face. By the canonical order, there exists at least an edge (w, z) incident to f that belongs to G^-_{k-1} but not to G^-_k. In particular, (w, z) belongs to C^-_{k-1} and thus it contains a horizontal segment by Invariant I.2. Also, by Invariant I.1b, both w and z are strictly below v, and thus the horizontal segment of (w, z) is strictly below the horizontal segment of (u, v). Let p and p' be any two points that belong to the horizontal segments of (u, v) and (w, z), respectively. We have that p lies strictly above p'; also, by Invariant I.2, f is below p; finally, by Invariant I.1a, all vertices of P_k have the same y-coordinate. It follows that we can join p and p' with a y-monotone curve that lies in the interior of f. By induction, there is a cut starting at any point of the horizontal segment of (w, z), and in particular at p', that separates the subpath of C^-_{k-1} connecting v_1 to w from the subpath of C^-_{k-1} connecting z to v_2. We can join this cut with the y-monotone curve described above, obtaining the desired cut. In particular, recall that C^-_k is obtained by replacing a subpath π of C^-_{k-1} containing w and z and with end-vertices denoted by u_ℓ and u_r, with a subpath starting at u_ℓ, visiting all vertices of P_k, and finishing at u_r. (Note that u_ℓ or u_r may coincide with u; see e.g. Fig. 3). Thus, the obtained cut at (u, v) starting from point p separates the subpath of C^-_k connecting v_1 to u from the subpath of C^-_k connecting v to v_2, as desired.

We now describe how to obtain a drawing Γ'_k of G^-_k satisfying all the required properties. Consider the y-monotone curve γ that is obtained by concatenating the cut at (u, v) starting at point p with the y-monotone curve that starts at p and ends at any point above Γ_k, which exists by Invariant I.2. Let L and R be the two sets of vertices separated by γ. All the vertices in L and all the edges between any two of them are drawn in Γ'_k as in Γ_k; all the vertices in R and all the edges between any two of them are drawn in Γ'_k as in Γ_k, after a translation to the right by σ. Finally, for each edge that is crossed by the cut, its horizontal segment is elongated by σ; also, if this edge contains a segment that is not horizontal, then this segment is either drawn in Γ'_k as in Γ_k, if it is to the left of the cut, or translated by σ, if it is to the right of the cut.

We prove that Γ'_k satisfies all properties of the lemma. First, Γ'_k is a 1-bend drawing of G^-_k on S since Γ_k is, and since all the edge-segments have the same slope in Γ'_k as in Γ_k. Also, Γ'_k is planar, since Γ'_k is obtained from Γ_k by elongating only the horizontal segments intersected by curve γ, which is y-monotone and spans the whole vertical extension of Γ_k. The fact that the horizontal distances between consecutive vertices of C^-_k are the required ones descends from the fact that L contains all the vertices in the path of C^-_k from v_1 to u, while R contains all the vertices in the path of C^-_k from v to v_2. We finally prove that Γ'_k satisfies the three invariants. Namely, Invariant I.1 holds since the y-coordinates of the vertices and of the bend-points, as well as the slope of the edge-segments, have not been changed, while Invariants I.2–I.3 hold since all the edges are attached to their incident vertices in Γ'_k using the same rays as in Γ_k. □

Invariant I.3 guarantees that every vertex on C^-_k has enough free top rays incident to the outer face to attach all its incident edges following it in the canonical order. The next lemma exploits Lemma 1 to show that these rays can be always used to actually draw these edges.

Lemma 2 *Let Γ_k be a 1-bend planar drawing on S of G^-_k satisfying Invariants I.1–I.3. Let u be any vertex of C^-_k, and let r_u be any free top ray of u that is incident to the*

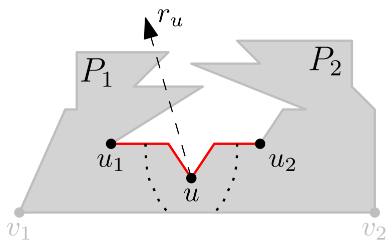

Fig. 4 Illustration for Lemma 2

outer face of G_k^- in Γ_k. Then, there exists a 1-bend planar drawing Γ'_k on S of G_k^-, satisfying Invariants I.1–I.3, in which r_u does not cross any edge of Γ'_k.

Proof Since r_u is a top ray of u incident to the outer face of Γ_k^-, if r_u crosses some edges of G_k^-, then it also crosses some edges of C_k^-. So, we can focus on removing the crossings with the edges of C_k^-.

Let P_1 be the path of C_k^- between v_1 and u, and let P_2 be the path of C_k^- between u and v_2. Also, let u_1 and u_2 be the neighbors of u in P_1 and P_2, respectively. Refer to Fig. 4. By Lemma 1, we can elongate (u, u_1) to eliminate all crossings between r_u and edges of P_1 without introducing any new crossings between r_u and edges of P_2. Analogously, we can elongate (u, u_2) to eliminate all crossings between r_u and edges of P_2 without introducing any new crossings between r_u and edges of P_1. The obtained drawing Γ'_k satisfies all the requirements of the lemma. The statement follows. □

We now describe our algorithm. First, we draw $P_0 = \{v_1, v_2\}$ and $P_1 = \{v_3, \ldots, v_j\}$ of partition Π such that $v_1, v_3, \ldots, v_j, v_2$, and the path connecting them, lie along a horizontal line, in this order (recall that edge (v_1, v_2) is not drawn at this stage). Invariants I.1a and I.2 clearly hold. Invariant I.3 follows from the fact that S contains $\Delta - 2$ top rays and all vertices drawn so far (including v_1 and v_2) have at most $\Delta - 2$ neighbors later in the canonical order.

We now describe how to add path P_k, for some $k > 1$, to a drawing Γ_{k-1} satisfying Invariants I.1–I.3, in such a way that the resulting drawing Γ_k of G_k^- is a 1-bend planar drawing on S satisfying Invariants I.1–I.3. We distinguish two cases, based on whether P_k is a chain or a singleton.

Suppose first that P_k is a chain, say $\{v_i, v_{i+1}, \ldots, v_j\}$; refer to Fig. 5a. Let u_ℓ and u_r be the neighbors of v_i and v_j in C_{k-1}^-, respectively. By Invariant I.3, each of u_ℓ and u_r has at least one free top ray that is incident to the outer face of Γ_{k-1}; among them, we denote by $\tau_a(u_\ell)$ the first one in anti-clockwise order for u_ℓ, and by $\tau_c(u_r)$ the first one in clockwise order for u_r. By Lemma 2, we can assume that $\tau_a(u_\ell)$ and $\tau_c(u_r)$ do not cross any edge in Γ_{k-1}. Consider any horizontal line that lies completely above Γ_{k-1}. By possibly applying Lemma 1 at any edge between u_ℓ and u_r we may assume that the crossing point of $\tau_a(u_\ell)$ is to the left of the crossing point of $\tau_c(u_r)$ along this horizontal line. Let h be the horizontal line segment between these two crossing points. We proceed by placing all the vertices $v_i, v_{i+1}, \ldots, v_j$ of P_k on interior points of h, in this left-to-right order. Finally, we draw edge (u_ℓ, v_i) with a segment along h and the other one along $\tau_a(u_\ell)$; we draw edge (u_r, v_j) with a segment along h and the other one along $\tau_c(u_r)$, and we draw every edge (v_q, v_{q+1}), with $q = i, \ldots, j-1$, with a unique segment along h.

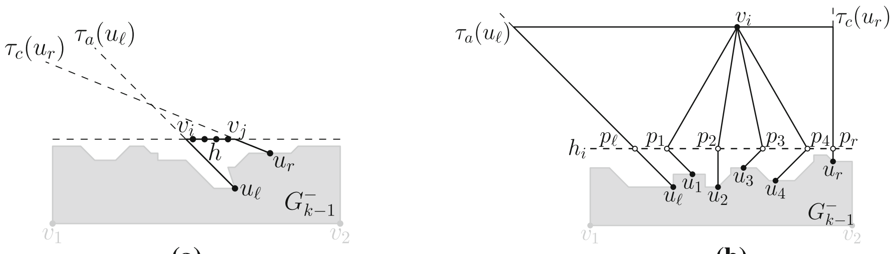

Fig. 5 Illustration of the cases in which P_k is: **a** a chain, and **b** a singleton of degree δ_i in G^-_k

By construction, Γ_k is a planar drawing on S and all the vertices of P_k lie above u_ℓ and u_r. Further, the bend-points of edges (u_ℓ, v_1) and (u_r, v_j) lie at the same y-coordinate as v_1 and v_j. Hence, Invariant I.1b is satisfied. Since all the edges of P_k are drawn as horizontal segments, Invariant I.1a is also satisfied. For Invariant I.2, note that every edge that is drawn at this step has a segment along h, which is horizontal and lies above Γ_{k-1}. Hence, the edges in $C^-_k \setminus C^-_{k-1}$ satisfy Invariant I.2. Consider an edge e that is in $C^-_k \cap C^-_{k-1}$. By Invariant I.2, there is a y-monotone curve starting at any point of the horizontal segment of e and lying in the outer face of Γ_{k-1}. Note that this curve may now be crossed by some edge in $C^-_k \setminus C^-_{k-1}$ of Γ_k. In this case, however, by following the drawing of the crossed edges, one can easily modify this curve so that it remains y-monotone and lies in the outer face of Γ_k. Hence, Invariant I.2 is also maintained. Finally, Invariant I.3 is satisfied since we attached edges (u_ℓ, v_i) and (u_r, v_j) at vertices u_ℓ and u_r using the first anti-clockwise free top ray of u_ℓ and the first clockwise free top ray of u_r among those that are incident to the outer face, respectively. Thus, we reduced only by one the number of free top rays incident to the outer face for u_ℓ and u_r. For the other vertices of P_k, the invariant is satisfied since their $\Delta - 2$ top rays are free and incident to the outer face. This concludes our description when P_k is a chain.

Suppose now that P_k is a singleton, say $\{v_i\}$, of degree $\delta_i \leq \Delta$ in G^-_k. This also includes the case in which $k = m$, that is, P_k is the last path of Π. If $\delta_i = 2$, then v_i is placed as in the case of a chain. So, we may assume in the following that $\delta_i \geq 3$. Let $u_\ell, u_1, u_2, \ldots, u_{\delta_i-2}, u_r$ be the neighbors of v_i as they appear along C^-_{k-1}.

By Invariant I.3, each neighbor of v_i in C^-_{k-1} has at least one free top ray that is incident to the outer face of Γ_{k-1}; among them, we denote by $\tau_a(u_\ell)$ the first one in anti-clockwise order for u_ℓ and by $\tau_c(u_r)$ the first one in clockwise order for u_r, while for each vertex u_q, with $q = 1, \ldots, \delta_i - 2$, we denote by $\tau(u_q)$ any of these rays arbitrarily. Refer to Fig. 5b. By Lemma 2, we can assume that these rays do not cross any edge in Γ_{k-1}.

Consider any horizontal line h_i lying above all vertices of Γ_{k-1}. Rays $\tau_a(u_\ell)$, $\tau(u_1), \ldots, \tau(u_{\delta_i-2})$, $\tau_c(u_r)$ cross h_i; however, the corresponding intersection points $p_\ell, p_1, \ldots, p_{\delta_i-2}, p_r$ may not appear in this left-to-right order along h_i; see Fig. 6a. To guarantee this property, we perform a sequence of stretchings of Γ_{k-1} by repeatedly applying Lemma 1. First, if p_ℓ is not the leftmost of these intersection points, then let σ be the distance between p_ℓ and the leftmost intersection point. We apply Lemma 1 at any edge between u_ℓ and u_1 along C^-_{k-1} to stretch Γ_{k-1} so that all the vertices in the

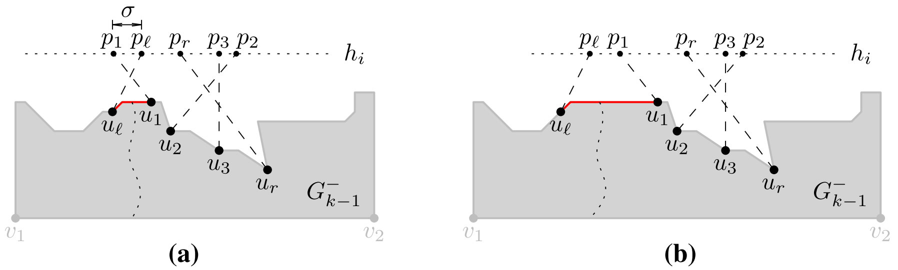

Fig. 6 **a** Intersection points $p_\ell, p_1, \ldots, p_{\delta_i-2}, p_r$ appear in a wrong order along h_i. **b** Applying Lemma 1 to make p_ℓ be the leftmost intersection point

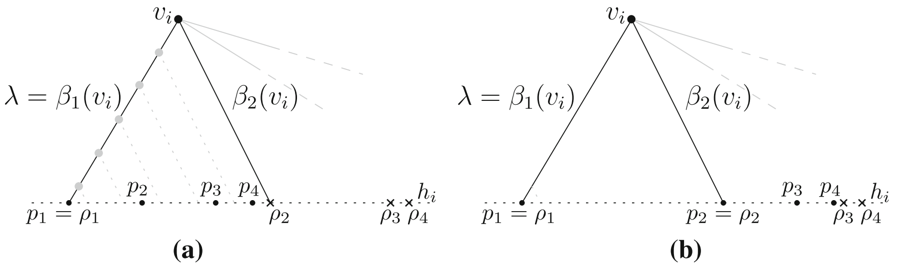

Fig. 7 Placement of a singleton v_i. **a** Moving v_i upwards along λ increases the distance among any two points ρ_q and ρ_{q+1}. For readability, ρ_3 and ρ_4 are not drawn in their correct position, which is more to the right. **b** After the stretching, p_2 coincides with ρ_2, and $p_3, \ldots, p_{\delta_i-2}$ lie between ρ_2 and ρ_3

path of C_k^- from u_1 to v_2 are moved to the right by a quantity σ' slightly larger than σ. This implies that p_ℓ is not moved, while all the other intersection points are moved to the right by a quantity σ', and thus they all lie to the right of p_ℓ in the new drawing; see Fig. 6b. Analogously, we can move p_1 to the left of every other intersection point, except for p_ℓ, by applying Lemma 1 at any edge between u_1 and u_2 along C_{k-1}^-. Thus, by applying this stretching at most $\delta_i - 1$ times, we obtain that in Γ_{k-1} the intersection points appear in the correct left-to-right order along h_i.

We now describe how to place v_i. Let $\beta_1(v_i), \ldots, \beta_{\delta_i-2}(v_i)$ be any set of $\delta_i - 2 \leq \Delta - 2$ consecutive bottom rays of v_i. Observe that, if we place v_i at any point above h_i, rays $\beta_1(v_i), \beta_2(v_i), \ldots, \beta_{\delta_i-2}(v_i)$ intersect h_i in this left-to-right order. Let $\rho_1, \ldots, \rho_{\delta_i-2}$ be the corresponding intersection points. The goal is to place v_i so that each ρ_q, for each $q = 1, \ldots, \delta_i - 2$, coincides with p_q. To do so, consider the line λ passing through p_1 with the same slope as $\beta_1(v_i)$. Refer to Fig. 7a. Note that placing v_i along λ and above h_i results in ρ_1 to coincide with p_1. Also note that, while moving v_i upwards along λ, the distance $d(\rho_q, \rho_{q+1})$ between any two consecutive points ρ_q and ρ_{q+1}, with $q = 1, \ldots, \delta_i - 3$, increases. Thus, we can place v_i along λ in such a way that $d(\rho_q, \rho_{q+1}) > d(p_1, p_{\delta_i-2})$, for each $q = 1, \ldots, \delta_i - 3$. This implies that all points $p_2, \ldots, p_{\delta_i-2}$ lie strictly between ρ_1 and ρ_2.

Then, we apply Lemma 1 at any edge between u_1 and u_2 along C_{k-1}^- to stretch Γ_{k-1} so that all the vertices in the path of C_k^- from u_2 to v_2 are moved to the right by a quantity $d(p_2, \rho_2)$. Refer to Fig. 7b. In this way, u_1 is not moved, and thus p_1

still coincides with ρ_1; also, p_2 is moved to the right to coincide with ρ_2; finally, since $d(\rho_2, \rho_3) > d(p_1, p_{\delta_i - 2}) > d(p_2, p_{\delta_i - 2})$, all points $p_3, \ldots, p_{\delta_i - 2}$ lie strictly between ρ_2 and ρ_3. By repeating this transformation for all points $p_3, \ldots, p_{\delta_i - 2}$, we guarantee that each ρ_q, with $q = 1, \ldots, \delta_i - 2$, coincides with p_q. We draw each edge (v_i, u_q), with $q = 1, \ldots, \delta_i - 2$, with one segment along $\tau(u_q)$ and one along $\beta_q(v_i)$; see Fig. 5b.

It remains to draw edges (v_i, u_ℓ) and (v_i, u_r), which must have a horizontal segment, in order to satisfy Invariant I.2. By possibly applying Lemma 1 at any edge between u_ℓ and u_1 along C^-_{k-1} to stretch Γ_{k-1}, we can guarantee that $\tau_a(u_\ell)$ crosses the horizontal line through v_i to the left of v_i. Similarly, by possibly applying Lemma 1 at any edge between $u_{\delta_i - 2}$ and u_r, we can guarantee that $\tau_c(u_r)$ crosses the horizontal line through v_i to the right of v_i. We draw edge (v_i, u_ℓ) with one segment along $\tau_a(u_\ell)$ and one along the horizontal line through v_i, and we draw edge (v_i, u_r) with one segment along $\tau_c(u_r)$ and one along the horizontal line through v_i. A drawing Γ_k produced in this step of the algorithm is illustrated in Fig. 5b.

The fact that Γ_k is a 1-bend planar drawing on S follows by the construction. We show that it satisfies Invariants I.1b, I.2 and I.3. For vertices v_i, u_ℓ, and u_r, for the edges (v_i, u_ℓ) and (v_i, u_r), and for the edges of $C^-_k \setminus C^-_{k-1}$, this can be proved as in the case in which P_k is a chain. For the other vertices and edges, note that $u_1, \ldots, u_{\delta_i - 2}$ do not have neighbors in $G \setminus G_k$ and do not belong to C^-_k, and thus do not need to satisfy Invariants I.2 and I.3. On the other hand, the fact that the edges connecting these vertices to v_i satisfy Invariant I.1b follows from the fact that their bend-points are along h_i, which lies above all the vertices of Γ_{k-1} and below v_i. This concludes our description for the case in which P_k is a singleton.

From the above discussion, it follows that the algorithm described in this section produces a 1-bend planar drawing on S of any planar graph with maximum degree Δ, where S is any set of $\Delta - 1$ slopes. We formalize this result in the following theorem, where we also prove that the algorithm can be performed in linear time. We refer to Fig. 8 for an example of a drawing constructed by our algorithm.

Theorem 2 *For any $\Delta \geq 4$, every set S of $\Delta - 1$ slopes is universal for 1-bend planar drawings of triconnected planar graphs with maximum degree Δ. Also, for any such graph on n vertices, a 1-bend planar drawing on S can be computed in $O(n)$ time.*

Proof Let G be any triconnected planar graph with maximum degree $\Delta \geq 4$. Apply the algorithm described above to produce a 1-bend planar drawing of G on S. The correctness has been proven throughout the section.

We now prove the time complexity. As already mentioned, computing the canonical order Π of G takes linear time [37]. Hence, our algorithm can be easily implemented in quadratic time. In fact, when a path P_k of Π is added, we first apply Lemma 2 twice, once for $\tau_a(u_\ell)$ and once for $\tau_c(u_r)$; each application uses Lemma 1 twice with a suitable stretching amount σ, which can be computed in constant time as follows. Let s be the slope of S such that $a = \min\{s, \pi - s\}$ is minimum over all slopes in S (note that s is the "flattest" slope between the one that follows the horizontal slope in clockwise order in S and the corresponding one in anti-clockwise order). Let y_{max} be the largest y-coordinate of the current drawing. Consider the intersection point p between the horizontal line at y_{max} and the half-line starting at the origin and having

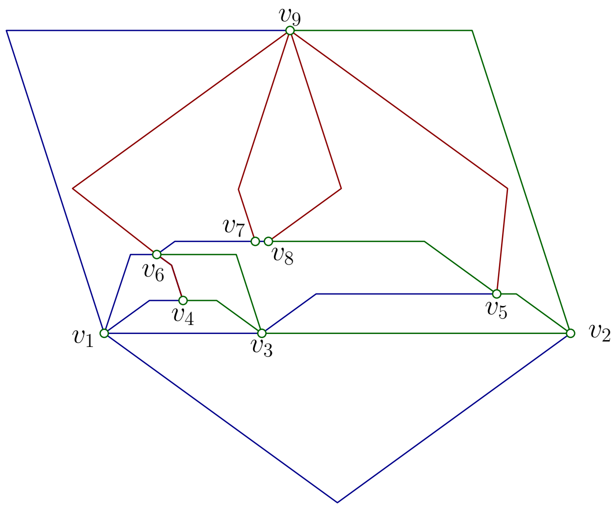

Fig. 8 A 1-bend drawing with 5 slopes of a triconnected planar graph with maximum degree 6 constructed by applying the algorithm of Theorem 2

slope s. Note that no segment of the drawing can have a horizontal projection larger than the absolute value of the x-coordinate of point p. Then we can choose σ equal to twice this quantity, which guarantees that the considered ray will not cross any segment of the drawing anymore after the stretching.

After having applied Lemma 2 twice as explained above, we apply Lemma 1 once if P_k is a chain, or $O(\delta_i)$ times if P_k is a singleton v_i of degree $\delta_i \leq \Delta$. In the latter case, since $\sum_{i=1}^{n} \delta_i = O(n)$, the total number of applications of the lemma over all singletons is $O(n)$. In each application of Lemma 1, the stretching amount σ can be computed in constant time as described in the algorithm. The total quadratic time comes from the fact that a straightforward application of Lemma 1 may require linear time.

To improve the time complexity of our algorithm to linear we seek to use the shifting method of Kant [35]. However, as the y-coordinates of the vertices are not necessarily consecutive, this method is not directly applicable. On the other hand, the y-coordinates of the vertices that have been placed at some step of our algorithm do not change in later steps. As noted by Bekos et al. [2], one can exploit this property so to allow the usage of the shifting method (even in the case of non-consecutive y-coordinates) in order to perform all applications of Lemma 1 in total linear time. This approach relies on two main observations. The first one is that each chain is not imposing any restriction on the height of the drawing, i.e., one may assume w.l.o.g. that the vertices of each chain are placed one unit above the drawing constructed so far. This can be guaranteed, e.g., by appropriately stretching the drawing horizontally. The second observation is that, in order to determine the y-coordinate of a singleton, one needs to compute the x-distances between its neighbors in the drawing constructed so far. This inevitably must be done by computing the exact x-distances between the neighbors of this singleton (which means that one has to "switch" from relative to exact x-distances for these vertices). The key observation is that the computation of exact x-distances will not involve the same set of vertices more than once, which guarantees that the total time of our algorithm will stay linear. □

We conclude this section by observing that every planar graph G with maximum degree Δ, even if not triconnected, can be triangulated by adding edges so that the resulting triangulated planar graph G' has maximum degree at most $\lceil \frac{3}{2}\Delta \rceil + 11$, as shown in [38]. Together with Theorem 2, it follows that every set S of $\lceil \frac{3}{2}\Delta \rceil + 10$ slopes is universal for 1-bend planar drawings of planar graphs, which matches the upper bound on the 1-bend planar slope number by Knauer and Walczak [40] up to a small additive factor. In the next sections we improve this bound to $\Delta - 1$ with a more sophisticated argument, hence extending Theorem 2 to all planar graphs.

4 Biconnected Planar Graphs

In this section we describe how to extend Theorem 2 to biconnected planar graphs. Note that variants of the canonical order for biconnected planar graphs exist (see, e.g., [29,30]) and one may wonder whether these variants can be used to prove a bound of $\Delta - 1$ slopes also for biconnected planar graphs. These variants however allow vertices having only one predecessor, which may create issues for our construction. More precisely, let v be a vertex with one predecessor and $\Delta - 1$ successors. Since we have $\Delta - 1$ slopes in total, and thus $\Delta - 2$ top rays, we must use a horizontal segment incident to v when representing an edge that connects v to one of its successors in the order, but this would violate Invariant I.1b. To overcome this problem, we present a technique based on the SPQR-tree data structure (see Sect. 2).

The idea is to traverse the SPQR-tree of the input biconnected planar graph G bottom-up and to construct for each visited node a special drawing of its pertinent graph (except for its two poles) inside an axis-aligned rectangle. In particular, besides being a 1-bend planar drawing on S, this drawing must have additional properties concerning the placement of the neighbors of the two poles and the slope of the edge-segments incident to them. In the following we give a formal definition and describe how to compute this drawing for each type of node of the SPQR-tree.

Let $\mathcal{T}$ be the SPQR-tree of G rooted at an arbitrary Q-node ρ. Let μ be a node of $\mathcal{T}$ with poles s_μ and t_μ. Let G_μ be the pertinent graph of μ. Let $\overline{G}_\mu$ be the graph obtained from G_μ as follows: (i) Remove edge (s_μ, t_μ), if it exists; (ii) subdivide each edge incident to s_μ (to t_μ) with a dummy vertex, which is called a *pin of* s_μ (is called a *pin of* t_μ); and (iii) remove s_μ and t_μ, and their incident edges. Note that, if μ is a Q-node other than the root ρ, then $\overline{G}_\mu$ is the empty graph. We denote by $\delta(s_\mu, \mu)$ and $\delta(t_\mu, \mu)$ the degree of s_μ and t_μ in G_μ, respectively; note that the number of pins of s_μ (of t_μ) is $\delta(s_\mu, \mu) - 1$ (is $\delta(t_\mu, \mu) - 1$), if edge (s_μ, t_μ) exists in G, otherwise it is $\delta(s_\mu, \mu)$ (it is $\delta(t_\mu, \mu)$).

The goal is to construct a 1-bend planar drawing of $\overline{G}_\mu$ on S, which lies inside an axis-aligned rectangle, called *chip* of μ and denoted by C_μ, that satisfies the following invariant properties (see Fig. 9a):

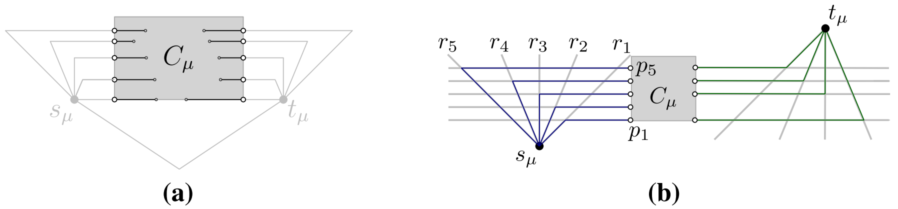

Fig. 9 **a** Illustration of a stretchable drawing of a node μ inside a chip C_μ. The poles s_μ and t_μ (in gray) do not belong to the drawing. **b** Illustration for Lemma 3

P.1: All pins of s_μ lie on one vertical side of C_μ and all pins of t_μ lie on the opposite side.
P.2: For each pin of either s_μ or t_μ, the (unique) edge incident to the pin is horizontal.
P.3: There exist pins on the bottom-left and on the bottom-right corners of C_μ.

A drawing of $\overline{G}_\mu$ that satisfies Properties P.1–P.3 is such that we can increase its width without changing the slope of its segments by elongating the horizontal edges incident to the pins on the left or right side of the chip; for this reason, we say that any of these drawings is *horizontally-stretchable* (or *stretchable*, for short). Note that a stretchable drawing remains stretchable after any uniform scaling and any translation.

Before giving the details of the algorithm, we describe a subroutine that we will often use to add the poles of a node μ to a stretchable drawing of $\overline{G}_\mu$ and draw the edges incident to them.

Lemma 3 *Let $u \in \{s_\mu, t_\mu\}$ be a pole of a node $\mu \in \mathscr{T}$ and let $u_1, \ldots, u_q$ be the neighbors of u in $\overline{G}_\mu$. Suppose that there exists a set of q consecutive free rays of u and a stretchabledrawing of $\overline{G}_\mu$ such that the elongation of the edge incident to each pin of u intersects all these rays. Then, it is possible to draw edges $(u, u_1), \ldots, (u, u_q)$, each with two straight-line segments whose slopes are in S, without introducing any crossing.*

Proof Refer to Fig. 9b. Let $p_1, \ldots, p_q$ be the pins of u, where p_i is adjacent to u_i ($i = 1, \ldots, q$). First note that, since $p_1, \ldots, p_q$ are all on the same side of C_μ, the elongations of their incident edges intersect the q free rays of u in the same order; we name the rays as $r_1, \ldots, r_q$ according to this order. Also note that, since the elongations of the edges incident to all the pins intersect all of $r_1, \ldots, r_q$, the elongation of the edge incident to either p_1 or p_q separates u from all the other pins. We assume without loss of generality that the elongation of the edge incident to p_1 separates u from $p_2, \ldots, p_q$, as in Fig. 9b. We then place each pin p_i, with $1 \leq i \leq q$, on the intersection point between the elongation of its incident edge and r_i, and draw edge (u, u_i) as a poly-line with a single bend at p_i. The fact that no crossing is introduced directly follows from the construction. This concludes the proof of the lemma. □

We now describe the algorithm. At each step of the bottom-up traversal of $\mathscr{T}$, we consider a node $\mu \in \mathscr{T}$ with children $\nu_1, \ldots, \nu_h$, and we construct a stretchable drawing of $\overline{G}_\mu$ starting from the stretchable drawings of $\overline{G}_{\nu_1}, \ldots, \overline{G}_{\nu_h}$ that have already been constructed. In the following, we distinguish four different cases, based on whether μ is a Q-, a P-, an S-, or an R-node.

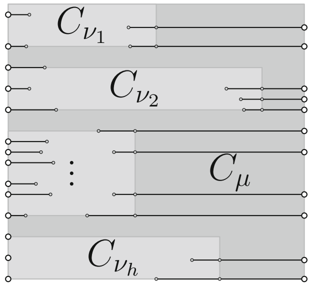

Fig. 10 Construction of a stretchable drawing of $\overline{G}_\mu$ inside C_μ, when μ is a P-node

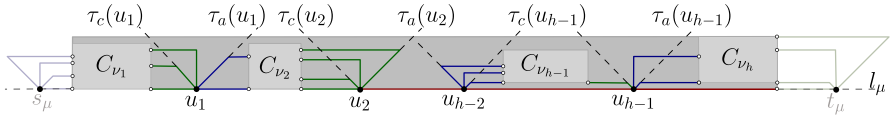

Fig. 11 Construction of a stretchable drawing of $\overline{G}_\mu$ when μ is an S-node

Suppose that μ is a Q-node If μ is not the root ρ of $\mathscr{T}$, we do not do anything, since $\overline{G}_\mu$ is the empty graph; edge (s_μ, t_μ) of G corresponding to μ will be drawn when visiting either the parent ξ of μ, if ξ is not a P-node, or the parent of ξ. In the case in which $\mu = \rho$, we observe that it has only one child ν_1. In particular, $\overline{G}_\mu = \overline{G}_{\nu_1}$, and thus the stretchable drawing of $\overline{G}_{\nu_1}$ is also a stretchable drawing of $\overline{G}_\mu$. Vertices s_μ and t_μ, and their incident edges, will be added after the traversal of $\mathscr{T}$.

Suppose that μ is a P-node; refer to Fig. 10. Since $\overline{G}_\mu$ does not contain the edge between the poles s_μ and t_μ of μ, for simplicity we assume that none of $\nu_1, \dots, \nu_h$ is a Q-node. We create a chip C_μ whose height is larger than the sum of the heights of the chips $C_{\nu_1}, \dots, C_{\nu_h}$ and whose width is larger than the maximum width of any of them. Then, we place the stretchable drawings of $\overline{G}_{\nu_1}, \dots, \overline{G}_{\nu_h}$ so that the chips $C_{\nu_1}, \dots, C_{\nu_h}$ lie inside C_μ without overlapping with each other, their left sides lie along the left side of C_μ, and the bottom side of C_{ν_h} lies along the bottom side of C_μ. Finally, we elongate the edges incident to the pins on the right sides of $C_{\nu_1}, \dots, C_{\nu_h}$ till reaching the right side of C_μ. The resulting drawing is stretchable since each of the drawings of $\overline{G}_{\nu_1}, \dots, \overline{G}_{\nu_h}$ is stretchable. In particular, Property P. 3 holds for C_μ since it holds for C_{ν_h}.

Suppose that μ is an S-node; refer to Fig. 11. Let $u_1, \dots, u_{h-1}$ be the internal vertices of the path of virtual edges between s_μ and t_μ that is obtained by removing the virtual edge (s_μ, t_μ) from the skeleton G_μ^{skel} of μ. We proceed as follows.

First, we place vertices $u_1, \dots, u_{h-1}$ in this order along a horizontal line l_μ. For $i = 1, \dots, h-1$, let $\tau_c(u_i)$ and $\tau_a(u_i)$ be the first top rays of u_i in clockwise and in anti-clockwise order, respectively. Then, for each child ν_i of μ, with $i = 1, \dots, h$, we place the chip C_{ν_i} as follows. For $i = 2, \dots, h-1$, we possibly apply a uniform scaling-down of C_{ν_i} and then we place it in such a way that its left side is to the right of u_{i-1}, its right side is to the left of u_i, it does not cross $\tau_a(u_{i-1})$ and $\tau_c(u_i)$, and either its bottom side lies on line l_μ (if edge $(u_{i-1}, u_i) \notin G$; see C_{ν_2} in Fig. 11), or it lies slightly above it (otherwise; see $C_{\nu_{h-1}}$ in Fig. 11). Analogously, we place C_{ν_1} in such a way that its right side is to the left of u_1, it does not cross $\tau_c(u_1)$, and either its bottom side lies on line l_μ (if edge $(s_\mu, u_1) \notin G$, as in Fig. 11), or it lies slightly above it (otherwise). Finally, we place C_{ν_h} in such a way that its left side is to the right of u_{h-1}, it does not cross $\tau_a(u_{h-1})$, and either its bottom side lies on line l_μ (if edge $(t_\mu, u_{h-1}) \notin G$), or it lies slightly above it (otherwise, as in Fig. 11).

We now draw all the edges incident to each vertex u_i, with $i = 1, \ldots, h-1$. Namely, if edge $(u_{i-1}, u_i) \in G$, for $i = 2, \ldots, h-1$, then it can be drawn as a horizontal segment, by construction (see (u_{h-2}, u_{h-1}) in Fig. 11). Otherwise, u_i can be connected with a horizontal segment to its neighbor in $\overline{G}_{\nu_i}$ corresponding to the pin on the bottom-right corner of C_{ν_i}, which exists by Property P. 3 (see u_1 and u_2 in Fig. 11). In both cases, one of these edges is attached at a horizontal ray of u_i. Analogously, one of the edges connecting u_i to its neighbors in $\overline{G}_{\nu_{i+1}} \cup \{u_{i+1}\}$ can be attached at the other horizontal ray of u_i. Thus, it is possible to draw the remaining $\delta(u_i, \nu_i) + \delta(u_i, \nu_{i+1}) - 2 \leq \Delta - 2$ edges incident to u_i by attaching them at (a subset of) the $\Delta - 2$ top rays of u_i, by applying Lemma 3. In fact, since the chips of ν_i and ν_{i+1} lie to the left and to the right of u_i, respectively, and do not cross $\tau_c(u_i)$ and $\tau_a(u_i)$, the elongations of the edges incident to the pins of u_i intersect all the top rays of u_i, hence satisfying the preconditions to apply the lemma.

Finally, we construct chip C_μ as the smallest axis-aligned rectangle enclosing the current drawing. We now show that it is possible to place the pins of poles s_μ and t_μ along the two vertical sides of C_μ so to satisfy Properties P.1–P.3.

Suppose first that ν_1 is not a Q-node. Then, the left side of C_μ contains the left side of C_{ν_1}, and thus all the pins of s_μ that correspond to pins of pole s_{ν_1} of ν_1 are already on the left side of C_μ. Note that, if (s_μ, u_1) does not belong to G, then these are all the pins of s_μ; also, in this case, the bottom-left corner of C_μ coincides with the bottom-left corner of C_{ν_1}, which contains a pin due to Property P.3. Hence, Properties P.1–P.3 are satisfied (see Fig. 11). On the other hand, if (s_μ, u_1) belongs to G, then C_{ν_1} lies above l_μ, and so it is possible to place the pin of s_μ corresponding to (s_μ, u_1) on the bottom-left corner of C_μ, and connect it with a horizontal segment to u_1 (see the corresponding case for C_{ν_h} in Fig. 11); thus, Properties P.1–P.3 are satisfied also in this case. Suppose now that ν_1 is a Q-node; note that in this case edge (s_μ, u_1) belongs to G, and the pin corresponding to this edge is the only pin of s_μ. Also note that, by construction, vertex u_1 lies on the bottom-left corner of C_μ. In order to avoid overlapping between u_1 and the pin corresponding to (s_μ, u_1), we slightly enlarge C_μ to the left and place the pin on its bottom-left corner, hence satisfying Properties P.1–P.3.

A symmetric argument applies to the pins of t_μ, which implies that the constructed drawing of $\overline{G}_\mu$ is stretchable. This concludes the case in which μ is an S-node.

Suppose that μ is an R-node In order to compute a stretchable drawing of $\overline{G}_\mu$, we first construct a 1-bend planar drawing on S of the whole pertinent graph G_μ of μ, including its poles s_μ and t_μ, and then we remove these poles and their incident edges; we finally define a chip C_μ and place the pins on its two vertical sides so to satisfy Properties P.1–P.3.

To compute the drawing of G_μ, we exploit the fact that the skeleton G_μ^{skel} of μ is triconnected. Hence, we can use the algorithm described in Sect. 3 as a main tool for drawing G_μ, with suitable modifications to take into account the fact that each virtual edge (u, v) of G_μ^{skel} actually corresponds to a whole subgraph, namely the pertinent graph G_ν of the child ν of μ with poles $s_\nu = u$ and $t_\nu = v$. Thus, when a virtual edge (u, v) is considered in the canonical order, we add the stretchable drawing of $\overline{G}_\nu$, together with the edge (u, v), if it exists; this enforces additional requirements for our drawing algorithm.

The first obvious requirement is that the edges connecting u and v to their neighbors in G_ν will occupy $\delta(u, \nu)$ consecutive rays of u and $\delta(v, \nu)$ consecutive rays of v, and not just a single ray for each of them, as in the triconnected case. Note that, however, reserving the correct amount of rays of u and v is not always sufficient to add the chip C_ν of $\overline{G}_\nu$. In fact, we need to ensure that there exists a placement for C_ν such that the elongations of the horizontal edge-segments incident to the pins of u (of v) intersect all the reserved rays of u (of v), hence satisfying the preconditions to apply Lemma 3. As an additional difficulty, this has to be done while considering that the edge (u, v), which does not belong to $\overline{G}_\nu$, may belong to the original graph. This requires two slightly different constructions, depending on the existence of this edge, in order to guarantee that some rays of u and v do not remain unused.

From a high-level point of view, for the virtual edges that would be drawn with a horizontal segment in the triconnected case (all the edges of a chain, and the first and last edges of a singleton), we handle these requirements by using a construction similar to the one of the case in which μ is an S-node. For the edges that do not have any horizontal segment (the internal edges of a singleton), instead, we need a more complicated construction.

We now describe in detail the algorithm, which is based on considering the vertices of G_μ according to a canonical order $\Pi = (P_0, \ldots, P_m)$ of G_μ^{skel}, in which $v_1 = s_\mu$ and $v_2 = t_\mu$. For $k = 1, \ldots, m$, we denote by H_k the graph that is the union of the pertinent graphs of the virtual edges of G_μ^{skel} connecting vertices in $P_0, P_1, \ldots, P_k$, except for the reference virtual edge (v_1, v_2) of μ. Note that the reference virtual edge (v_1, v_2) of μ represents the rest of the graph with respect to G_μ, and thus we have that $H_m = G_\mu$.

For each graph H_k, the algorithm constructs a 1-bend planar drawing on S satisfying a modified version of Invariants I.1–I.3. We use C_k to denote the outer face of H_k and we let C_k^- be the path from v_1 to v_2 containing the vertices in P_k.

M.1 Let (u, v) be a virtual edge of G_μ^{skel}, such that $u \in P_i$ and $v \in P_j$ $(i \leq j)$. Let ν be the node of $\mathcal{T}$ associated with (u, v).

(a) If $i = j$ or $j = 1$, then u and v have the same y-coordinate; also, if edge (u, v) exists, then it is drawn as a horizontal segment. Finally, the drawing of G_ν is inside a rectangle with u and v along its bottom side (possibly at the corners).

(b) Otherwise (i.e., $i < j$ and $j \neq 1$), vertex u is strictly below v; also, if edge (u, v) exists, then its bend-point (if any) is strictly above u and below or at the same y-coordinate as v. Finally, the drawing of G_ν is inside a rectangle with u along its bottom side and v along one of its other sides.

M.2 Let $(u, v) \neq (v_1, v_2)$ be a virtual edge of G_μ^{skel} with both u and v on C_k^-. Let ν be the node of $\mathcal{T}$ associated with (u, v). Each edge of G_ν incident to u or v (possibly including (u, v)) has a horizontal segment. Also, consider the edge incident to

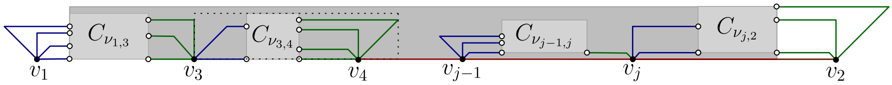

Fig. 12 Construction of a drawing of $P_0 \cup P_1$ satisfying Invariants M.1–M.3

u that is incident to C_k^-. Then, for any point p of its horizontal segment there is a strictly y-monotone curve on the outer face of the drawing that starts at p and ends at any point strictly above the drawing. The same holds for the edge incident to v that is incident to C_k^-.

M.3 Each vertex of G_μ^{skel} on C_k^- has at least as many free top rays incident to the outer face of H_k as the number of its neighbors in G_μ that have not been drawn yet.

Note that Invariant M.1 is similar to Invariant I.1. In particular, it ensures that v_1 and v_2 are the bottommost vertices in the drawing and are horizontally aligned. Invariant M.2 corresponds to Invariant I.2, as it ensures that we can still apply Lemma 1 and Lemma 2. Finally, Invariant M.3 is the natural extension of Invariant I.3 to take into account our previous observation that a virtual edge of G_μ^{skel} corresponds to a whole subgraph of G_μ.

At the first step, we draw $P_0 = \{v_1, v_2\}$ and $P_1 = \{v_3, \ldots, v_j\}$. Consider the path of virtual edges $(v_1, v_3), (v_3, v_4), \ldots, (v_j, v_2)$, and let $\nu_{1,3}, \nu_{3,4}, \ldots, \nu_{j,2}$ be the corresponding children of μ. We consider this path as the skeleton of an S-node μ with poles $s_\mu = v_1$ and $t_\mu = v_2$, and we apply the same algorithm as described in the previous subsection; refer to Fig. 12. Namely, we place the chips of $\nu_{1,3}, \nu_{3,4}, \ldots, \nu_{j,2}$ inside a rectangle, which we call the chip C_{v_1,v_2} of $P_0 \cup P_1$ (with a slight abuse of notation).

Note that, by construction, C_{v_1,v_2} has pins on its bottom-left and on its bottom-right corners. We then place v_1 and v_2 with the same y-coordinate as the bottom side of C_{v_1,v_2}, with v_1 to the left and v_2 to the right of C_{v_1,v_2}. This ensures that we can draw one of the edges incident to v_1 horizontal, and the remaining $\delta(v_1, \nu_{1,3}) - 1$ by applying Lemma 3. We draw v_2 symmetrically.

We denote the resulting drawing by Γ_1, and prove that it satisfies the three invariants. First, observe that for every virtual edge (v_i, v_{i+1}), with $i = 3, \ldots, j-1$, vertices v_i and v_{i+1} lie on the two bottom side of the rectangle enclosing the drawing of $G_{\nu_{i,i+1}}$; see the dotted rectangle in Fig. 12 for $G_{\nu_{3,4}}$. Note that this rectangle degenerates to a horizontal segment if $\nu_{i,i+1}$ is a Q-node. Also, the same property holds for virtual edges (v_1, v_3) and (v_j, v_2). Hence, Invariant M.1a is satisfied. Since the existence of the y-monotone curves required by Invariant M.2 is guaranteed by construction, we can prove that Invariant M.2 is satisfied in the same way we proved that Invariant P.4 holds for S-nodes. Finally, for Invariant M.3, observe that each vertex v_i, with $i = 4, \ldots, v_{j-1}$, has consumed $\delta(v_i, \nu_{i-1,i}) + \delta(v_i, \nu_{i,i+1}) - 2$ top rays out of the $\Delta - 2$ available ones, since exactly two of its incident edges are attached at its horizontal rays. Thus, v_i has $\Delta - \delta(v_i, \nu_{i-1,i}) - \delta(v_i, \nu_{i,i+1})$ free top rays incident to the outer face, which can be used to attach the at most $\Delta - \delta(v_i, \nu_{i-1,i}) - \delta(v_i, \nu_{i,i+1})$ remaining neighbors it may have in G_μ, and so it satisfies the invariant. Analogous arguments hold for v_3 and v_j, which have consumed $\delta(v_3, \nu_{1,3}) + \delta(v_3, \nu_{3,4}) - 2$ and $\delta(v_j, \nu_{j-1,j}) + \delta(v_j, \nu_{j,2}) - 2$ top rays, respectively. As for v_1 and v_2, they have consumed $\delta(v_1, \nu_{1,3}) - 1$ and

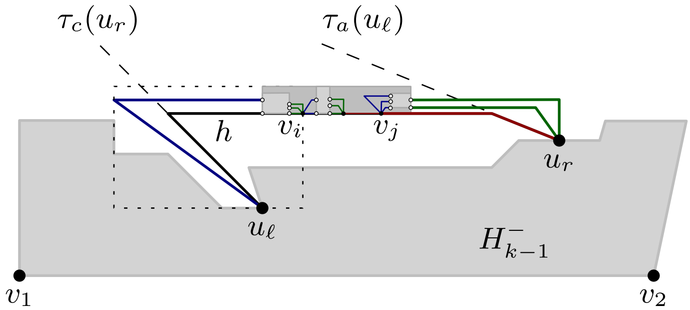

Fig. 13 Addition of a chain P_k to drawing Γ_{k-1}, resulting in a drawing Γ_k satisfying Invariants M.1–M.3

$\delta(v_2, \nu_{j,2}) - 1$ top rays, respectively, since each of them has only one edge attached to a horizontal ray; thus, they have $\Delta - \delta(v_1, \nu_{1,3}) - 1$ and $\Delta - \delta(v_2, \nu_{j,2}) - 1$ free top rays incident to the outer face, respectively. However, since both v_1 and v_2 have at least a neighbor in $G_\mu \setminus H_1$, their remaining neighbors in G_μ are at most $\Delta - \delta(v_1, \nu_{1,3}) - 1$ and $\Delta - \delta(v_2, \nu_{j,2}) - 1$, respectively, and so they also satisfy Invariant M.3. Note that no other vertices of H_1 have further neighbors in G_μ.

We now describe how to add path P_k, for some $k > 1$, to the current drawing Γ_{k-1} in the two cases in which P_k is a chain or a singleton, so that the resulting drawing Γ_k satisfies the three invariants.

Suppose that P_k is a chain, say $\{v_i, v_{i+1}, \dots, v_j\}$. Let u_ℓ and u_r be the neighbors of v_i and v_j in G^{skel}_μ that lie on C^-_k, respectively. Denote by ν_ℓ and ν_r the children of μ corresponding to virtual edges (u_ℓ, v_i) and (v_j, u_r), respectively, and by $\nu_i, \dots, \nu_{j-1}$ the children of μ corresponding to virtual edges $(v_i, v_{i+1}), \dots, (v_{j-1}, v_j)$, respectively.

We define rays $\tau_a(u_\ell)$ and $\tau_c(u_r)$, and the horizontal segment h between them, as in the triconnected case. However, in this case, we possibly perform an additional stretching of the drawing, by applying Lemma 1 at an edge between u_ℓ and u_r along C^-_k, in order to ensure that there exists at least an internal point of h whose x-coordinate lies between those of u_ℓ and u_r; see Fig. 13. Due to Lemma 2, we can assume that $\tau_a(u_\ell)$ and the $\delta(u_\ell, \nu_\ell) - 1$ top rays of u_ℓ following it in anti-clockwise order do not cross any edge of Γ_{k-1}, and the same for $\tau_c(u_r)$ and the $\delta(u_r, \nu_r) - 1$ top rays of u_r following it in clockwise order. Note that, by Invariant M.3, all these rays are free and incident to the outer face of the drawing.

Then, as in the previous step in which we considered P_0 and P_1 of Π, we use the algorithm for the case in which μ is an S-node to construct a drawing of the subgraph composed of vertices $v_i, \dots, v_j$ and of the pertinent graphs of $\nu_\ell, \nu_i, \dots, \nu_{j-1}$ and ν_r, inside a rectangle C_{u_ℓ,u_r} having pins on its bottom-left and on its bottom-right corners. After possibly performing a uniform scaling-down of C_{u_ℓ,u_r}, we place it so that: (i) its bottom side lies on h; (ii) its left side is to the right of u_ℓ; (iii) its right side is to the left of u_r; and (iv) it does not cross $\tau_a(u_\ell)$ and $\tau_c(u_r)$; see Fig. 13. Note that conditions (ii) and (iii) can be met due to the previous application of Lemma 1, which ensures that there exists at least an internal point of h whose x-coordinate lies between those of u_ℓ and u_r. We will use these two conditions to guarantee that virtual edges (u_ℓ, v_i) and (v_j, u_r) satisfy Invariant M.1. Finally, we draw the $\delta(u_\ell, \nu_\ell)$ edges between u_ℓ and its neighbors in $\overline{G}_{\nu_\ell} \cup \{v_i\}$, and the $\delta(u_r, \nu_r)$ edges between

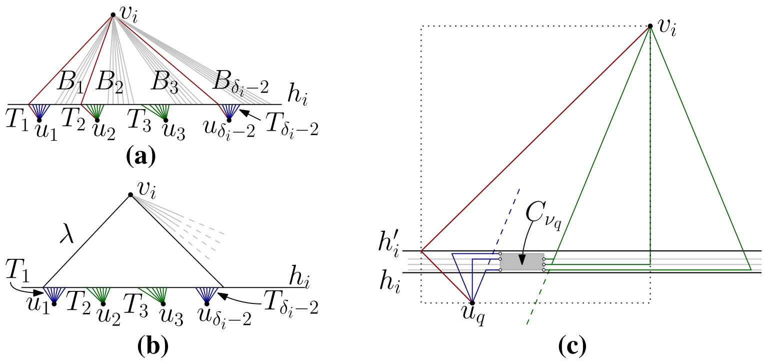

Fig. 14 Illustrations for the case in which P_k is a singleton, when μ is an R-node. **a** An ordering of the intersection points of sets B_q and T_q, with $q = 1, \dots, \delta_i - 2$, respecting the BT-ordering condition. **b** Placement of v_i along line λ. **c** Addition of chip C_{ν_q} to the drawing

u_r and its neighbors in $\overline{G}_{\nu_r} \cup \{v_j\}$, by applying Lemma 3, whose preconditions are satisfied.

The fact that the constructed drawing satisfies the three invariants can be proved similarly as in the step in which we considered P_0 and P_1. In particular, the existence of the y-monotone curves required by Invariant M.2 is guaranteed by the fact that h lies above Γ_{k-1}. Hence, as anticipated above, the only additional argument we need is to show that (u_ℓ, v_i) and (v_j, u_r) satisfy Invariant M.1b. Due to condition (ii), to the construction of C_{u_ℓ,u_r}, and to the fact that h lies above u_ℓ, we have that the pertinent graph of ν_ℓ lies inside a rectangle having u_ℓ along the bottom side and v_i along the right side; see the dotted rectangle in Fig. 13. Analogously, we can prove that the pertinent graph of ν_r lies inside a rectangle having u_r along the bottom side and v_i along the left side. This concludes the analysis of the case in which P_k is a chain.

Suppose finally that P_k is a singleton $\{v_i\}$, and let $u_\ell, u_1, \dots, u_{\delta_i-2}, u_r$ be the neighbors of v_i in G^{skel}_μ that lie (in this order) along C^-_{k-1}, with $\delta_i \leq \Delta$. As in the triconnected case, we assume that $\delta_i \geq 3$. Let $\nu_\ell, \nu_1, \dots, \nu_{\delta_i-2}, \nu_r$ be the children of μ corresponding to the virtual edges connecting v_i to these vertices.

For each $q = 1, \dots, \delta_i - 2$, we select any set T_q of consecutive $\delta(u_q, \nu_q)$ free top rays of u_q incident to the outer face, which exist by Invariant M.3, and a set B_q of consecutive $\delta(v_i, \nu_q)$ bottom rays of v_i; refer to Fig. 14a. Sets $B_1, \dots, B_{\delta_i-2}$ are selected in such a way that all the rays in B_q precede all the rays in B_{q+1} in anti-clockwise order. Since $\delta(v_i, \nu_\ell) + \delta(v_i, \nu_r) \geq 2$, vertex v_i has enough bottom rays for sets $B_1, \dots, B_{\delta_i-2}$. We also define sets T_ℓ and T_r as composed of the first $\delta(u_\ell, \nu_\ell)$ free top rays of u_ℓ in anti-clockwise order and of the first $\delta(u_r, \nu_r)$ free top rays of u_r in clockwise order, respectively, which exist by Invariant M.3.

We then select a horizontal line h_i lying above every vertex in Γ_{k-1}. As in the triconnected case, after possibly applying Lemma 1 at most $\Delta - 1$ times, we can assume that all the rays in sets $T_\ell, T_1, \dots, T_{\delta_i-2}, T_r$ intersect h_i in the correct order. Namely, when moving along h_i from left to right, we encounter all the intersections with the rays in T_ℓ, then all those with the rays in T_1, and so on. On the other hand, this property is guaranteed by construction for the rays in $B_1, \dots, B_{\delta_i-2}$. This defines two total left-to-right orders $\mathcal{O}_T$ and $\mathcal{O}_B$ of the intersection points of $T_\ell, T_1, \dots, T_{\delta_i-2}, T_r$

and of $B_1, \ldots, B_{\delta_i-2}$ along h_i, respectively. To simplify the description, we extend these orders to the rays in $T_\ell, T_1, \ldots, T_{\delta_i-2}, T_r$ and in $B_1, \ldots, B_{\delta_i-2}$, respectively.

Our goal is to merge the two sets of intersection points, while respecting $\mathcal{O}_T$ and $\mathcal{O}_B$, in such a way that for each $q = 1, \ldots, \delta_i - 2$ the following condition, called *BT-ordering condition*, holds. If edge (v_i, u_q) belongs to H, then the first intersection point of T_q in $\mathcal{O}_T$ coincides with the first intersection point of B_q in $\mathcal{O}_B$, and the second intersection point of B_q in $\mathcal{O}_B$ is to the right of the last intersection point of T_q in $\mathcal{O}_T$; see T_1 and B_1 in Fig. 14a. Otherwise, that is $(v_i, u_q) \notin H$, the first intersection point of B_q in $\mathcal{O}_B$ is to the right of the last intersection point of T_q in $\mathcal{O}_T$; see T_3 and B_3 in Fig. 14a. In both cases, the intersection points of T_q and B_q are to the left of those of T_{q+1} and B_{q+1}.

To achieve this goal, we perform a procedure analogous to the one described in the triconnected case of our algorithm to make points $p_1, \ldots, p_{\delta_i-2}$ coincide with points $\rho_1, \ldots, \rho_{\delta_i-2}$. Namely, we consider a half-line λ whose slope is the one of the first ray in B_1 such that, if edge (v_1, u_1) belongs to H, then λ starts at the first intersection point of T_1 in $\mathcal{O}_T$, otherwise, it starts at any point between the last intersection point of T_1 and the first intersection point of T_2 in $\mathcal{O}_T$. Then, we place v_i along λ, far enough from h_i so that the distance between any two consecutive intersection points in $\mathcal{O}_B$ is larger than the distance between the first and the last intersection points in $\mathcal{O}_T$; see Fig. 14b. Finally, we apply Lemma 1 at most $\delta_i - 3$ times to move the intersection points of sets $T_2, \ldots, T_{\delta_i-2}$, one by one, so to meet the BT-ordering condition; see Fig. 14a.

Once the BT-ordering condition is met for each $q = 1, \ldots, \delta_i - 2$, we consider another horizontal line h'_i lying slightly above h_i such that its intersections with the rays in $T_\ell, T_1, \ldots, T_{\delta_i-2}, T_r$ and $B_1, \ldots, B_{\delta_i-2}$ appear along it in the same order as along h_i; refer to Fig. 14c. For each $q = 1, \ldots, \delta_i - 2$, we place the chip C_{ν_q} of ν_q, after possibly scaling it down uniformly, in the interior of the region delimited by these two lines, by the last ray in T_q, and by a ray in B_q (either the second or the first, depending on whether $(v_i, u_q) \in G_\mu$ or not).

We draw the edges incident to v_i and u_q, for each $q = 1, \ldots, \delta_i - 2$, as follows. If edge (v_i, u_q) belongs to G_μ, then we draw it with one segment along the first ray in T_q and one along the first ray in B_q (see Fig. 14c). For the other edges we apply Lemma 3 twice, whose preconditions are satisfied due to the placement of C_{ν_q} (see the blue and green edges in Fig. 14c).

We conclude by drawing the edges connecting v_i, u_ℓ, and vertices in $\overline{G}_{\nu_\ell}$; the edges connecting v_i, u_r and vertices in $\overline{G}_{\nu_r}$ are drawn symmetrically. First, after possibly applying Lemma 1 at the edge incident to u_ℓ that belongs to the path of C^-_{k-1} from u_ℓ to u_r, we assume that v_i is to the right of u_ℓ, and that the last ray of T_ℓ intersects the horizontal line through v_i at a point p_i that is to the left of v_i. After possibly scaling down uniformly the chip C_{ν_ℓ} of ν_ℓ, we place it so that: (i) its left side is to the right of both p_i and u_ℓ; (ii) its right side is to the left of v_i; (iii) it does not cross the first top ray of v_i in clockwise order; and (iv) its bottom side lies either above the horizontal line through v_i, if edge $(u_\ell, v_i) \in G_\mu$, or along it, otherwise. Then, if $(u_\ell, v_i) \in G_\mu$, we draw it with one segment along the last ray of T_ℓ and the other one along the horizontal line through v_i. Otherwise, we draw as a horizontal segment the edge connecting v_i to its neighbor in $\overline{G}_{\nu_\ell}$ corresponding to the pin on the bottom-right corner of C_{ν_ℓ}. We finally apply Lemma 3 twice, to draw the edges from u_ℓ and from v_i to their neighbors in $\overline{G}_{v_\ell}$.

We prove that Γ_k satisfies Invariants M.1b, M.2 and M.3. For vertices u_ℓ, v_i, and u_r, and for the virtual edges (u_ℓ, v_i), and (v_i, u_r), the same arguments as in the case in which P_k is a chain hold. On the other hand, vertices $u_1, \ldots, u_{\delta_i-2}$ do not belong to C_k^-, and so these vertices and the virtual edges connecting them to v_i do not need to satisfy Invariants M.2 and M.3, but only Invariant M.1b. To see that this is the case, observe that for each virtual edge (u_q, v_i), with $q = 1, \ldots, \delta_i - 2$, we have that u_q is at the bottom side and v_i is at the top side of the rectangle enclosing G_{v_q}, by construction; see the dotted rectangle in Fig. 14c.

Once the last path P_m of Π has been added, we have a drawing Γ_m of $H_m = G_\mu$ satisfying Invariants M.1–M.3. We construct a stretchable drawing of $\overline{G}_\mu$ starting from Γ_m, as follows. We initialize the chip C_μ of μ as the smallest axis-aligned rectangle enclosing Γ_m. By Invariant M.1a, vertices v_1 and v_2 lie on the bottom side of C_μ. Also, by Invariant M.2, all the edges connecting either v_1 or v_2 to one of their neighbors attach to this neighbor with a horizontal segment (this is due to the fact that v_1 and v_2 belong to C_k^-, for all $k = 1, \ldots, m$). We thus remove v_1 and v_2 (and their incident edges) from Γ_m, we elongate the horizontal segments incident to their neighbors till reaching the vertical sides of C_μ, and we place pins at their ends. The fact that the obtained drawing satisfies Properties P.1–P.3 follows from the observation that v_1 and v_2 were on the bottom side of C_μ, and from the fact that they were both using a horizontal ray to attach to one of their neighbors. This concludes the case in which μ is an R-node.

Once the bottom-up traversal of $\mathscr{T}$ has been completed, after visiting the root ρ of $\mathscr{T}$, we have a stretchable drawing of $\overline{G}_\rho$ inside a chip C_ρ, which we extend to a drawing of G as follows. Refer to Fig. 9a. We place the poles s_ρ and t_ρ of ρ at the same y-coordinate as the bottom side of C_ρ, with s_ρ to its left and t_ρ to its right, so that C_ρ does not cross any of the rays of s_ρ and of t_ρ. Then, we draw edge (s_ρ, t_ρ) with one segment along the first bottom ray in clockwise order of s_ρ and the other one along the first bottom ray in anti-clockwise order of t_ρ. Note that, since $\Delta \geq 4$, these rays intersect at a point with finite coordinates. Also, we draw the edges connecting s_ρ and t_ρ to the vertices corresponding to the pins on the bottom corners of C_ρ as horizontal segments. Finally, we draw all the remaining edges incident to s_ρ and t_ρ by applying Lemma 3 twice. We refer to Fig. 15 for an example of a drawing constructed by our algorithm. The following theorem summarizes the discussion in this section.

Theorem 3 *For any $\Delta \geq 4$, every set S of $\Delta - 1$ slopes is universal for 1-bend planar drawings of biconnected planar graphs with maximum degree Δ. Also, for any such graph on n vertices, a 1-bend planar drawing on S can be computed in $O(n)$ time.*

Proof Let G be any biconnected planar graph with maximum degree Δ. Apply the algorithm described above to produce a 1-bend planar drawing of G on S. The correctness has been proved throughout the section.

For the time complexity, first observe that the SPQR-tree $\mathscr{T}$ of G can be computed in linear time [31]. Also, for each node $\mu \in \mathscr{T}$, we can compute a stretchable drawing of $\overline{G}_\mu$ in time linear in the size of G_μ^{skel} assuming that for each chip we only store the

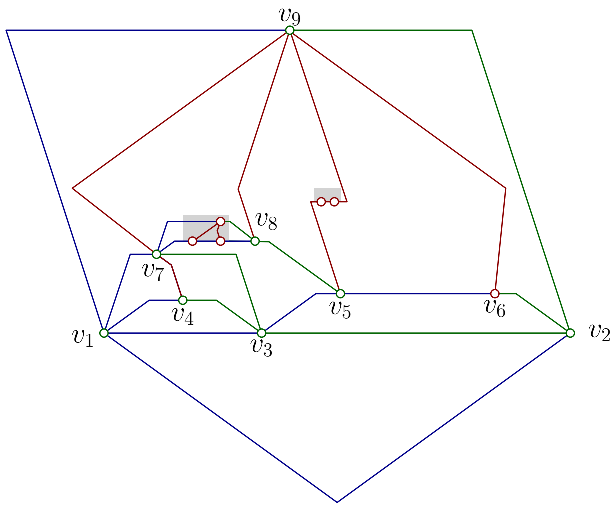

Fig. 15 A 1-bend drawing with 5 slopes of a biconnected planar graph with maximum degree 6 constructed by applying the algorithm of Theorem 3. The SPQR-tree of this graph contains a Q-node for each of its edges, two R-nodes, and one S-node. The two shaded in gray rectagles illustrate one R-node and one S-node

coordinates of two opposite corners. Final coordinates can then be assigned by traversing the SPQR-tree top-down. In particular, notice that the linear-time complexity for R-nodes can be proved analogously as for the algorithm for triconnected graphs; see Theorem 2. Since the total size over all the skeletons of the nodes of $\mathscr{T}$ is linear in the size of G, the time complexity of the algorithm is linear. □

5 General Planar Graphs

Let G be a connected planar graph with maximum degree Δ and let $\mathscr{B}$ be its BC-tree. Let β be a B-node whose parent in $\mathscr{B}$ is the C-node γ. Let (γ, ξ) be an edge of β, and consider the SPQR-tree of β rooted at the Q-node ρ corresponding to edge (γ, ξ). We call $\overline{G}_\beta$ the graph $\overline{G}_\rho$, as defined in Sect. 4.

For each B-node β of $\mathscr{B}$, we compute a stretchable drawing $\Gamma(\beta)$ of $\overline{G}_\beta$, satisfying an additional property other than P.1–P.3. Consider any vertex c of β different from γ that is a cut-vertex in G, and let δ_c be the number of neighbors of c in G that do not belong to β.

P.4 There exists a set of δ_c consecutive bottom rays of c that are not used in $\Gamma(\beta)$.

To satisfy the above property, we slightly modify the algorithm of Theorem 3, as follows. Recall that cut-vertex c is drawn in $\Gamma(\beta)$ by this algorithm when considering a node μ of the SPQR-tree of β such that μ is either an S-node or an R-node and c is an internal vertex of the corresponding skeleton. If μ is an S-node, all the bottom rays of c are free by construction. If μ is an R-node, c is either a singleton or part of a chain in the canonical order used to draw the pertinent graph of μ. When c is part of a chain, its bottom rays are free again by construction. When c is a singleton, note that

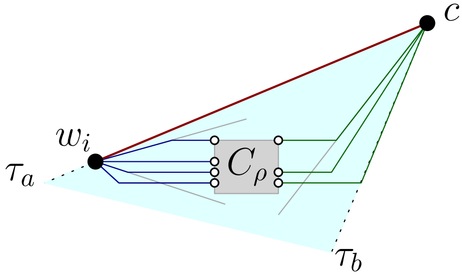

Fig. 16 Placement of the chip C_ρ for a child ζ_i of β

the algorithm reserves a set of consecutive bottom rays B_q (as in Fig. 14a) for each virtual edge connecting c with one of its predecessors in the canonical order. Hence, it is enough to reserve one more set B^* with cardinality δ_c. The argument that c has enough bottom rays is implied by the fact that c uses both its horizontal rays in $\Gamma(\beta)$, it has $\Delta - 2$ bottom rays, and it has maximum degree Δ in G.

We now show how to combine the drawings of all the B-nodes of $\mathscr{B}$. Let β be a B-node of $\mathscr{B}$. For each C-node c that is a child of β in $\mathscr{B}$, consider all its children $\zeta_1, \ldots, \zeta_q$, with $q \leq \Delta - 2$. Note that for each of these blocks ζ_i, with $i = 1, \ldots, q$, the drawing $\Gamma(\zeta_i)$ of $\overline{G}_{\zeta_i}$ inside a chip C_{ζ_i} has been computed by rooting the SPQR-tree of ζ_i at a Q-node ρ_i corresponding to an edge incident to c. This implies that c does not belong to $\overline{G}_{\zeta_i}$. We now show how to add C_{ζ_i} and the pole w_i of ρ_i different from c to the drawing $\Gamma(\beta)$ of $\overline{G}_\beta$, and how to draw the edges connecting c and w_i to their neighbors in $\overline{G}_{\zeta_i}$.

Let $\delta(c, \zeta_i)$ be the degree of c in ζ_i. Since $\sum_{i=1}^{q} \delta(c, \zeta_i) = \delta_c$, we can insert these drawings into $\Gamma(\beta)$ using the δ_c free bottom rays of c, which are guaranteed to exist by Property P.1, as follows. Let τ_a and τ_b be the a-th and b-th rays within the δ_c free bottom rays of c in anti-clockwise order, where $a = 1 + \sum_{j=1}^{i-1} \delta(c, \zeta_j)$ and $b = \sum_{j=1}^{i} \delta(c, \zeta_j)$. These two rays define a cone in which we place the chip C_ρ, refer to Fig. 16. Let τ'_a be the top ray of w_i with the same slope as τ_a. Suppose that w_i has t top rays following τ'_a in clockwise order. After possibly performing a scaling-down of C_ρ, we place it such that: (a) C_ρ lies in the intersection of the cone delimited by τ_a and by the bottom ray of c following τ_a in anti-clockwise order, and of the cone delimited by the first top ray of w_i in anti-clockwise order and by the first bottom ray of w_i in clockwise order; (b) the $(t+1)$-th pin of w_i from top to bottom has the same y-coordinate as w_i. These two properties guarantee that all the edges that connect c and w_i to vertices in $\overline{G}_{\zeta_i}$ can be drawn with one bend and with slopes in S. In particular, note that edge (c, w_i) exists, since it corresponds to the Q-node ρ_i; this edge is drawn as a single segment along τ_a. Also, the edge incident to w_i and to the vertex corresponding to the $(t+1)$-th pin of w_i from top to bottom is drawn as a horizontal segment. Further, the edges connecting c to its neighbors in $\overline{G}_{\zeta_i}$ are drawn by applying Lemma 3. Finally, the edges connecting w_i to its neighbors in $\overline{G}_{\zeta_i}$ are drawn by applying twice Lemma 3, once for those above the $(t+1)$-th pin of w_i, and once for those below it.

Since the modification to the algorithm in Theorem 3 that is applied to satisfy Property P.1 does not alter its linear-time complexity, and since the procedure to

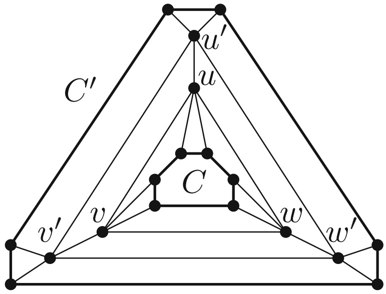

Fig. 17 Illustration for the proof of Corollary 2

merge all the computed drawing can also be implemented to run in linear time with respect to the size of $\mathscr{B}$, the overall procedure takes linear time. Thus, the algorithm described in this section extends the result of Theorem 3 to every connected graph.

In order to obtain a complete proof of Theorem 1, we need to extend this result even to disconnected graphs. This is however immediately implied by the fact that we construct universal slope sets, and thus the same set of slopes can be used to draw each connected component independently.

We conclude this section by proving the following corollary of Theorem 1 that proves that planar graphs of degree at most Δ admit 1-bend planar drawings whose angular resolution is worst-case optimal up to a multiplicative factor of at least $\frac{3}{4}$ (as Δ tends to infinity).

Corollary 2 *A planar graph with maximum degree $\Delta \geq 3$ admits a planar drawing with at most one bend per edge and angular resolution at least $\frac{\pi}{\Delta-1}$. Also there exist planar graphs with maximum degree Δ whose planar drawings with at most one bend per edge all have angular resolution strictly less than $\frac{4\pi}{3(\Delta-2)}$.*

Proof For $\Delta \geq 4$, the first part of the statement is a consequence of Theorem 1 when considering an equispaced set of $\Delta - 1$ slopes. When $\Delta = 3$, the statement follows from a work by Kant [35] who proved that an orthogonal planar drawing with at most one bend per edge can always be constructed.

The second part of the statement is an almost straightforward consequence of the technique by Keszegh et al. [39] used to prove a $\frac{3}{4(\Delta-1)}$ lower bound on the 1-bend planar slope number. For every integer $\Delta \geq 3$, consider the graph G_Δ of Fig. 17 (which shows the case when $\Delta = 5$). In any planar embedding of G_Δ, either the cycle C (bold in Fig. 17) is in the interior of the 3-cycle with vertices u, v, w or the cycle C' (bold in Fig. 17) is in the interior of the 3-cycle with vertices u', v', w'. Assume that C is inside the 3-cycle induced by vertices u, v, w (the proof in the other case is analogous). In every 1-bend planar drawing of G, the 3-cycle with vertices u, v, w is represented as a k-gon Ψ with $3 \leq k \leq 6$. It follows that the sum of the internal angles of Ψ at vertices u, v, w is strictly less than 4π and, therefore, at least one of these three angles–say the one at u–is strictly less than $\frac{4\pi}{3}$. Since the number of edges connecting u to C is $\Delta - 3$, it follows that the minimum angle between any two such edges is strictly less than $\frac{4\pi}{3(\Delta-2)}$. □

6 Conclusions and Open Problems

In this paper, we improved the best-known upper bound of Knauer and Walczak [40] on the 1-bend planar slope number from $\frac{3}{2}(\Delta - 1)$ to $\Delta - 1$, for $\Delta \geq 4$. We obtained this as a corollary of a stronger result, namely, that any set of $\Delta - 1$ slopes is universal for 1-bend planar drawings of planar graphs with maximum degree $\Delta \geq 4$. By using an equispaced set of slopes, the angular resolution of our drawings is at least $\frac{\pi}{\Delta - 1}$.

A side-result of our work is the following. For $\Delta = 4$, our algorithm guarantees that planar graphs with maximum degree 4 admit 1-bend planar drawings on a set of slopes $\{0, \frac{\pi}{3}, \frac{2\pi}{3}\}$, while previously it was known that such graphs can be embedded with one bend per edge on a set of slopes $\{0, \frac{\pi}{4}, \frac{\pi}{2}, \frac{3\pi}{4}\}$ [2] and with two bends per edge on a set of slopes $\{0, \pi\}$ [5].

Our work raises several open problems.

- Reduce the gap between the $\frac{3}{4}(\Delta - 1)$ lower bound and the $\Delta - 1$ upper bound on the 1-bend planar slope number.
- Our drawings may have super-polynomial area. Is this unavoidable for 1-bend planar drawings with few slopes (and good angular resolution)?
- Study the straight-line case (e.g., for degree-4 graphs). Note that the stretching operation might be difficult in this setting.

Acknowledgements This work started at the 19th *Korean Workshop on Computational Geometry*. We wish to thank the organizers and the participants for creating a pleasant and stimulating atmosphere and in particular Fabian Lipp and Boris Klemz for useful discussions. We also gratefully thank the anonymous reviewers for their useful comments.

References

1. Barát, J., Matoušek, J., Wood, D.R.: Bounded-degree graphs have arbitrarily large geometric thickness. Electron. J. Comb. **13**(1), 3 (2006)
2. Bekos, M.A., Gronemann, M., Kaufmann, M., Krug, R.: Planar octilinear drawings with one bend per edge. J. Graph Algorithms Appl. **19**(2), 657–680 (2015)
3. Bekos, M.A., Kaufmann, M., Krug, R.: On the total number of bends for planar octilinear drawings. J. Graph Algorithms Appl. **21**(4), 709–730 (2017)
4. Bertolazzi, P., Di Battista, G., Mannino, C., Tamassia, R.: Optimal upward planarity testing of single-source digraphs. SIAM J. Comput. **27**(1), 132–169 (1998)
5. Biedl, T.C., Kant, G.: A better heuristic for orthogonal graph drawings. Comput. Geom. **9**(3), 159–180 (1998)
6. Bläsius, T., Krug, M., Rutter, I., Wagner, D.: Orthogonal graph drawing with flexibility constraints. Algorithmica **68**(4), 859–885 (2014)
7. Bläsius, T., Lehmann, S., Rutter, I.: Orthogonal graph drawing with inflexible edges. Comput. Geom. **55**, 26–40 (2016)
8. Bodlaender, H.L., Tel, G.: A note on rectilinearity and angular resolution. J. Graph Algorithms Appl. **8**, 89–94 (2004)
9. Bonichon, N., Saëc, B. L., Mosbah, M.: Optimal area algorithm for planar polyline drawings. In: Kucera, L. (ed.) WG, vol. 2573 of LNCS, pp. 35–46. Springer, Berlin (2002)
10. Czyzowicz, J.: Lattice diagrams with few slopes. J. Comb. Theory Ser. A **56**(1), 96–108 (1991)
11. Czyzowicz, J., Pelc, A., Rival, I.: Drawing orders with few slopes. Discrete Math. **82**(3), 233–250 (1990)
12. Czyzowicz, J., Pelc, A., Rival, I., Urrutia, J.: Crooked diagrams with few slopes. Order **7**(2), 133–143 (1990)

13. de Fraysseix, H., de Mendez, P.O., Rosenstiehl, P.: On triangle contact graphs. Comb. Probab. Comput. **3**, 233–246 (1994)
14. de Fraysseix, H., Pach, J., Pollack, R.: How to draw a planar graph on a grid. Combinatorica **10**(1), 41–51 (1990)
15. Di Battista, G., Eades, P., Tamassia, R., Tollis, I.G.: Graph Drawing: Algorithms for the Visualization of Graphs. Prentice-Hall, Englewood Cliffs (1999)
16. Di Giacomo, E., Liotta, G., Montecchiani, F.: Drawing outer 1-planar graphs with few slopes. J. Graph Algorithms Appl. **19**(2), 707–741 (2015)
17. Di Giacomo, E., Liotta, G., Montecchiani, F.: 1-bend upward planar drawings of SP-digraphs. In: Hu, Y., Nöllenburg, M. (eds.), Graph Drawing, vol. 9801 of LNCS, pp. 123–130. Springer, Berlin (2016)
18. Di Giacomo, E., Liotta, G., Montecchiani, F.: Drawing subcubic planar graphs with four slopes and optimal angular resolution. Theor. Comput. Sci. **714**, 51–73 (2018)
19. Dujmović, V., Eppstein, D., Suderman, M., Wood, D.R.: Drawings of planar graphs with few slopes and segments. Comput. Geom. **38**(3), 194–212 (2007)
20. Dujmović, V., Suderman, M., Wood, D.R.: Graph drawings with few slopes. Comput. Geom. **38**(3), 181–193 (2007)
21. Duncan, C.A., Eppstein, D., Goodrich, M.T., Kobourov, S.G., Nöllenburg, M.: Drawing trees with perfect angular resolution and polynomial area. Discrete Comput. Geom. **49**(2), 157–182 (2013)
22. Duncan, C.A., Kobourov, S.G.: Polar coordinate drawing of planar graphs with good angular resolution. J. Graph Algorithms Appl. **7**(4), 311–333 (2003)
23. Durocher, S., Mondal, D.: On balanced +-contact representations. In: Graph Drawing, vol. 8242 of LNCS, pp. 143–154. Springer, Berlin (2013)
24. Durocher, S., Mondal, D.: Trade-offs in planar polyline drawings. In: Duncan, C.A., Symvonis, A. (eds.) Graph Drawing, vol. 8871 of LNCS, pp. 306–318. Springer, Berlin (2014)
25. Felsner, S., Kaufmann, M., Valtr, P.: Bend-optimal orthogonal graph drawing in the general position model. Comput. Geom. **47**(3), 460–468 (2014)
26. Formann, M., Hagerup, T., Haralambides, J., Kaufmann, M., Leighton, F.T., Symvonis, A., Welzl, E., Woeginger, G.J.: Drawing graphs in the plane with high resolution. SIAM J. Comput. **22**(5), 1035–1052 (1993)
27. Garg, A., Tamassia, R.: Planar drawings and angular resolution: algorithms and bounds (extended abstract). In: van Leeuwen, J. (ed.) ESA, vol. 855 of LNCS, pp. 12–23. Springer, Berlin (1994)
28. Garg, A., Tamassia, R.: On the computational complexity of upward and rectilinear planarity testing. SIAM J. Comput. **31**(2), 601–625 (2001)
29. Gronemann, M.: Bitonic st-orderings of biconnected planar graphs. In: Graph Drawing, vol. 8871 of LNCS, pp. 162–173. Springer, Berlin (2014)
30. Gutwenger, C., Mutzel, P.: Planar polyline drawings with good angular resolution. In: Graph Drawing, vol. 1547 of LNCS, pp. 167–182. Springer, Berlin (1998)
31. Gutwenger, C., Mutzel, P.: A linear time implementation of SPQR-trees. In: Marks, J. (ed), Graph Drawing, vol. 1984 of LNCS, pp. 77–90. Springer, Berlin (2000)
32. Hoffmann, U.: On the complexity of the planar slope number problem. J. Graph Algorithms Appl. **21**(2), 183–193 (2017)
33. Jelínek, V., Jelínková, E., Kratochvíl, J., Lidický, B., Tesar, M., Vyskocil, T.: The planar slope number of planar partial 3-trees of bounded degree. Graphs Comb. **29**(4), 981–1005 (2013)
34. Jünger, M., Mutzel, P. (eds.): Graph Drawing Software. Springer, Berlin (2004)
35. Kant, G.: Drawing planar graphs using the lmc-ordering (extended abstract). In: FOCS, pp. 101–110. IEEE Computer Society (1992)
36. Kant, G.: Hexagonal grid drawings. In: Mayr, E.W. (ed.) WG, vol. 657 of LNCS, pp. 263–276. Springer, Berlin (1992)
37. Kant, G.: Drawing planar graphs using the canonical ordering. Algorithmica **16**(1), 4–32 (1996)
38. Kant, G., Bodlaender, H.L.: Triangulating planar graphs while minimizing the maximum degree. Inf. Comput. **135**(1), 1–14 (1997)
39. Keszegh, B., Pach, J., Pálvölgyi, D.: Drawing planar graphs of bounded degree with few slopes. SIAM J. Discrete Math. **27**(2), 1171–1183 (2013)
40. Knauer, K., Walczak, B.: Graph drawings with one bend and few slopes. In: Kranakis, E., Navarro, G., Chávez, E. (eds.), LATIN, vol. 9644 of LNCS, pp. 549–561. Springer, Berlin (2016)
41. Knauer, K.B., Micek, P., Walczak, B.: Outerplanar graph drawings with few slopes. Comput. Geom. **47**(5), 614–624 (2014)

42. Lenhart, W., Liotta, G., Mondal, D., Nishat, R.I.: Planar and plane slope number of partial 2-trees. In: Wismath, S.K., Wolff, A. (eds.) Graph Drawing, vol. 8242 of LNCS, pp. 412–423. Springer, Berlin (2013)
43. Liu, Y., Morgana, A., Simeone, B.: A linear algorithm for 2-bend embeddings of planar graphs in the two-dimensional grid. Discrete Appl. Math. **81**(1–3), 69–91 (1998)
44. Mukkamala, P., Pálvölgyi, D.: Drawing cubic graphs with the four basic slopes. In: van Kreveld, M.J., Speckmann, B. (eds.) Graph Drawing, vol. 7034 of LNCS, pp. 254–265. Springer, Berlin (2011)
45. Nöllenburg, M.: Automated drawings of metro maps. Technical Report 2005-25, Fakultät für Informatik, Universität Karlsruhe (2005)
46. Nöllenburg, M., Wolff, A.: Drawing and labeling high-quality metro maps by mixed-integer programming. IEEE Trans. Vis. Comput. Graph. **17**(5), 626–641 (2011)
47. Stott, J.M., Rodgers, P., Martinez-Ovando, J.C., Walker, S.G.: Automatic metro map layout using multicriteria optimization. IEEE Trans. Vis. Comput. Graph. **17**(1), 101–114 (2011)
48. Tamassia, R.: On embedding a graph in the grid with the minimum number of bends. SIAM J. Comput. **16**(3), 421–444 (1987)
49. Tamassia, R. (ed.) Handbook on Graph Drawing and Visualization. CRC Press, Amsterdam (2013)
50. Wade, G.A., Chu, J.: Drawability of complete graphs using a minimal slope set. Comput. J. **37**(2), 139–142 (1994)