Drawing Plane Triangulations with Few Segments

Stephane Durocher* Debajyoti Mondal†

Abstract

Dujmović, Eppstein, Suderman, and Wood showed that every 3-connected plane graph G with n vertices admits a straight-line drawing with at most $2.5n-3$ segments, which is also the best known upper bound when restricted to plane triangulations. On the other hand, they showed that there exist triangulations requiring $2n-6$ segments. In this paper we show that every plane triangulation admits a straight-line drawing with at most $(7n-2\Delta_0-10)/3 \le 2.33n$ segments, where Δ_0 is the number of cyclic faces in the minimum realizer of G. If the input triangulation is 4-connected, then our algorithm computes a drawing with at most $(9n-9)/4 \le 2.25n$ segments. For general plane graphs with n vertices and m edges, our algorithm requires at most $(16n-3m-28)/3 \le 5.33n-m$ segments, which is smaller than $2.5n-3$ for all $m \ge 2.84n$.

1 Introduction

A *plane graph* is a fixed combinatorial embedding of a planar graph. Given a plane graph G, a *straight-line drawing* of G in $\mathbb{R}^2$ maps each vertex of G to a point, and each edge of G to a straight line segment such that no two edges intersect except possibly at their common endpoints. A *segment* in a straight-line drawing Γ is a maximal path such that all the vertices on the path are collinear in Γ. A *k-segment drawing* is a straight-line drawing with at most k segments. A k-segment drawing of a plane graph G is called a *minimum-segment drawing* if G does not admit any straight-line drawing with fewer than k segments. Figure 1(a) illustrates a plane graph G. Figures 1(b) and (c) are two straight-line drawings of G with 10 and 8 segments, respectively. The drawing of Figure 1(c) is a minimum segment drawing of G.

Straight-line drawings are preferable since the use of bends makes it difficult to follow the edges in the drawing. Drawings with few segments further enhance this straightness aesthetic. Besides, a k-segment drawing corresponds to an edge decomposition of the underlying graph into k induced paths. Although the problem of computing a drawing with minimum number of segments is NP-hard for arrangement graphs [4], drawings with minimum number of segments have been achieved for trees [3], plane 2-trees with maximum degree three [6], and cubic plane graphs [5]. Dujmović et al. [3] proved tight upper and lower bounds on the number of segments for several classes of plane graphs such as outerplane graphs (n segments), plane k-trees ($2n$ segments) with $k \in \{2,3\}$, and a nearly tight upper bound of $5n/2$ segments for 3-connected plane graphs. As a natural open problem they asked to determine the minimum constant c such that every plane graph with n vertices admits a straight-line drawing with at most $cn+O(1)$ segments. They also examined planar graphs, i.e., when the embedding of the input graph is not given.

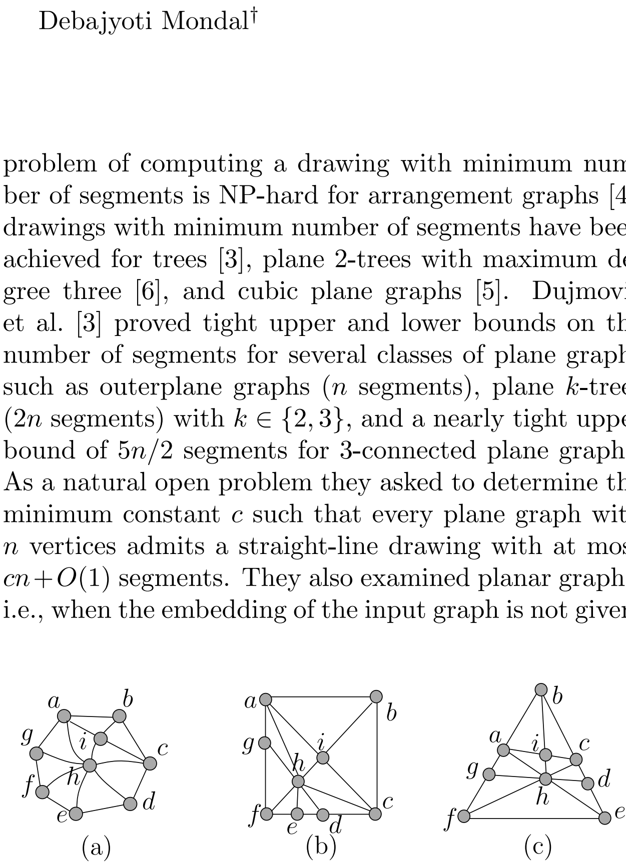

Figure 1: (a) A plane graph G. (b–c) Two straight-line drawings of G.

In this paper we only examine plane graphs, i.e., the combinatorial embedding of the input graph is given as an input, and the output drawing must respect the given embedding. Table 1 summarizes the best known upper and lower bounds on the number of segments for different classes of plane graphs.

Graph Class	L.B.	U.B.	Ref.
Trees	$\lambda/2$	$\lambda/2$	[4]
Maximal outerplane graphs	n	n	[4]
Plane 2-tree (max-degree 3)	$2n$	$2n$	[6]
Plane 2- and 3-trees	$2n$	$2n$	[3]
3-connected cubic plane graphs	$n/2$	$n/2$	[5]
3-connected plane graphs	$2n$	$5n/2$	[3]
This Paper			
3-connected triangulations	$2n$	$7n/3$	Th. 4
4-connected triangulations	$2n$	$9n/4$	Th. 5

Table 1: Upper and lower bounds on the number of segments, ignoring additive constants. Here λ is the number of vertices of odd degree.

*Department of Computer Science, University of Manitoba, Canada. Work of the author is supported in part by the Natural Sciences and Engineering Research Council of Canada (NSERC), durocher@cs.umanitoba.ca

†Department of Computer Science, University of Manitoba, Canada. Work of the author is supported in part by a University of Manitoba Graduate Fellowship. jyoti@cs.umanitoba.ca

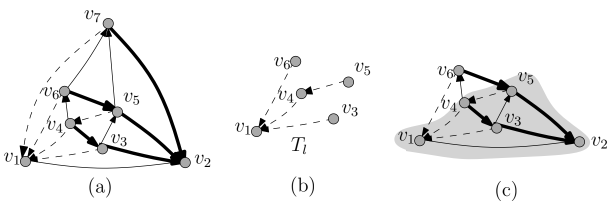

Figure 2: (a) A plane triangulation G with a canonical ordering of its vertices. The associated realizer is a minimum realizer, where the l-, r- and m- edges are shown in dashed, bold-solid, and thin-solid edges respectively. The only cyclic face is v_4, v_6, v_5, v_4, which is oriented clockwise. (b) T_l. (c) Illustration for canonical ordering, when $k+1=6$.

2 Preliminaries

In this section we introduce some preliminary definitions and results.

2.1 Canonical Ordering and Schnyder Realizer

Let G be a connected plane graph. G is called k-*connected*, where $k > 1$, if removal of fewer than k vertices does not disconnect the graph. G is called *triangulated* if and only if each of its faces (including the outer face) is a cycle of length three. G is *internally triangulated* if each of its inner faces is a cycle of length three. Let G be an n-vertex triangulated plane graph, and let v_1, v_2 and v_n be the outer vertices of G in clockwise order. Let $\sigma = (v_1, v_2, ..., v_n)$ be an ordering of all vertices of G. By G_k, $3 \leq k \leq n$, we denote the subgraph of G induced by $v_1 \cup v_2 \cup ... \cup v_k$. By P_k we denote the path (while walking clockwise) on the outer face of G_k that starts at v_1 and ends at v_2. We call σ a *canonical ordering* of G with respect to the outer edge (v_1, v_2) if for each k, $3 \leq k \leq n$, the following conditions are satisfied [2].

(a) G_k is 2-connected and internally triangulated.

(b) If $k+1 \leq n$, then v_{k+1} is an outer vertex of G_{k+1} and the neighbors of v_{k+1} in G_k appears consecutively on P_k.

For some k, where $3 \leq k \leq n$, let P_k be the path $w_1(=v_1), \ldots, w_l, v_k(=w_{l+1}), w_r, \ldots, w_t(=v_2)$. We call the edges (w_l, v_k) and (v_k, w_r) the *l-edge* and the *r-edge* of v_k, respectively. The other edges incident to v_k in G_k are called the *m-edges* of v_k. For example, in Figure 2(c), the edges $(v_1, v_6), (v_6, v_5)$, and (v_4, v_6) are the l-, r- and m-edges of v_6, respectively. Let E_m be the set of all m-edges in G. Then the graph T_m induced by the edges in E_m is a tree with root v_n. Similarly, the graph T_l induced by all l-edges except (v_1, v_n) is a tree rooted at v_1 (Figure 2(b)), and the graph T_r induced by all r-edges except (v_2, v_n) is a tree rooted at v_2. These three trees form the *Schnyder realizer* [7] of G. A *minimum realizer* is a Schnyder realizer with all the cyclic inner faces oriented clockwise. The number of cyclic inner faces in a minimum realizer is denoted by Δ_0 [9]. Each of T_l, T_r and T_m corresponds to a canonical ordering of G, and hence called a *canonical ordering tree* of G.

2.2 Monotone Chains, Rays and Visibility

Let p be a point in $\mathbb{R}^2$. We denote the x and y-coordinates of p by p_x and p_y, respectively. Let $b_1, b_2, \ldots, b_k$ be a strictly x-monotone polygonal chain C. For each i, where $0 < i < k$, an edge (b_i, b_{i+1}) is called a left (respectively, right) edge if $b_{i_y} < b_{i+1_y}$ (respectively, $b_{i_y} > b_{i+1_y}$). Let Γ be a straight-line drawing of a plane graph G. A *segment* in Γ is a maximal path of G whose vertices are collinear in Γ. A segment is a *left or right segment* if it contains a left or right edge, respectively. The *tip of a left or right segment s* is the vertex on s with the highest y-coordinate. The *tip of an edge e* is the tip of the segment that contains e. Two points p and q are *visible* to each other with respect to Γ if they does not intersect Γ at any point except possibly at p and q. By l_{pq} we denote the line through p and q. We denote the slope of l_{pq} by $slope(p, q)$. A set of rays is *divergent* if no two rays in the set are parallel, and no two rays intersect (except possibly at their common origin).

Lemma 1 *Let $a, b_1, b_2, \ldots, b_k, c(= b_{k+1})$ be a strictly x-monotone polygonal chain C. Let p be a point above C such that the segments ap and cp does not intersect C except at a and c. If the slopes of the left edges of C are smaller than $slope(a, p)$, and the slope of the right edges of C are greater than $slope(p, c)$, then every b_i is visible from p (e.g., Figure 3(a)).*

Proof. Suppose for a contradiction that some vertex b_i, where $1 \leq i \leq k$, is not visible to p. Without loss of generality assume that b_i is in the left half-plane of the vertical line through p. Since C is strictly x-monotone, no left edge of C can block the visibility between p and b_i. Hence let (b_j, b_{j+1}) be the right edge that blocks the visibility, where $i < j \leq k$, as shown in Figure 3(b).

If the slope of the line l_{b_i,b_j} is smaller than $slope(a, p)$, then it is also smaller than $slope(p, b_i)$, and hence b_i must be visible to p. We may thus assume that $slope(b_i, b_j)$ is as large as $slope(a, p)$, which implies that (b_i, b_j) cannot be an edge in C. Consider now the path $P = (b_i, b_{i+1}, \ldots, b_j)$. Observe that the edge (b_i, b_{i+1}) must lie to the right half plane of l_{b_i,b_j}. On the other hand, the edge (b_{j-1}, b_j) must lie to the left half plane of l_{b_i,b_j}. Hence there exists some edge e in P that crosses l_{b_i,b_j}. Since P is strictly x-monotone, e must be a left edge. Furthermore, since e crosses l_{b_i,b_j}, we have $slope(e) > slope(b_i, b_j) \geq slope(a, p)$, which is a contradiction. $\square$

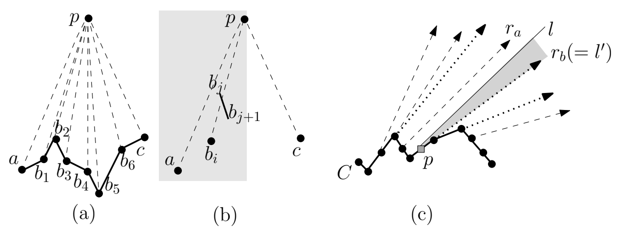

Figure 3: (a–b) Illustration for the proof of Lemma 1. (c) Illustration for the proof of Lemma 2, where the sets R_1 and R_2 are shown in dotted and dashed lines, respectively.

Lemma 2 *Let C be a strictly x-monotone polygonal chain, and let R_1 be the set of rays obtained by the extending each left segment of C above C. Let R_2 be another set of rays each with origin on C, directed above C and with slope less than $90°$. Assume that the rays in $R_1 \cup R_2$ are divergent, and none of the rays intersect C except at their origin. Given a point p on C, one can find a ray r with origin p such that the rays in $R_1 \cup R_2 \cup \{r\}$ are divergent, and r does not intersect C except at p.*

Proof. We prove the lemma by constructing the ray r. Let r_a be the last ray that we encounter while walking on C from left to right before we visit p, as shown in Figure 3(c). If there are several candidates for r_a, i.e., all with the same origin on C, then we choose the last ray above C in the clockwise order around the origin. If we do not encounter any ray before visiting p, then there is no left edge before p, and we choose r_a as a vertical ray directed upward starting at the leftmost point on C. We now draw a line l parallel to r_a at p.

Similarly, find the last ray r_b that we encounter while walking on C from right to left until we visit p (here p itself could be the origin of r_b). If there are several candidates for r_b, i.e., all with the same origin on C, then we choose the first ray above C in the clockwise order around the origin. If we do not encounter any ray before visiting p, then there is no left edge before p, and we choose r_b as a horizontal ray directed to the right starting at the rightmost point on C. We now draw a line l' parallel to r_b at p.

We now construct the ray r with origin p and slope $(slope(l') + slope(l))/2$. Since l' and l are divergent, the set $R_1 \cup R_2 \cup \{r\}$ is divergent. Since C is strictly x-monotone and no ray in $R_1 \cup R_2$ intersects C except at its origin, all the points in the region bounded by l and l' above C is visible to p. Hence r cannot intersect C except at p. □

3 Drawing Plane Triangulations

In this section we prove that every n-vertex plane triangulation G admits a straight-line drawing with at most $(7n - 2\Delta_0 - 10)/3$ segments.

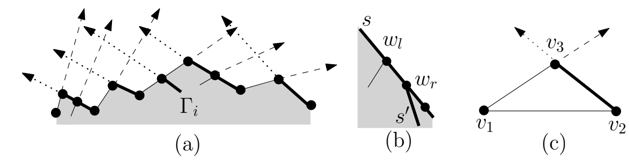

Figure 4: (a) Illustrating the invariants for some Γ_i. The path P_i is the upper envelope of the shaded region. The edges of T_l and T_r are shown in thin and bold solid lines, respectively. The set Q_l and Q_r are shown in dashed and dotted lines, respectively. (b) Illustration for Case 1. (c) Drawing of G_3.

Let $\sigma = (v_1, v_2, \ldots, v_n)$ be a canonical ordering of the vertices of G, which corresponds to the minimum realizer of G. We first construct a drawing of G using σ, and then bound the number of segments in the constructed drawing.

3.1 Algorithm FewSegDraw

We first draw the edge (v_1, v_2) using a horizontal straight line segment. We now complete the drawing of G by adding the vertices $v_3, v_4, \ldots, v_n$ incrementally. Let Γ_i be the drawing of G_i. At each addition, Γ_i will maintain the following invariants, as shown in Figure 4(a).

1. The drawing of P_i in Γ_i is strictly x-monotone.

2. Let (u, v) be an edge in G_i, where u appears before v in some P_j, where $3 \le j \le i$. If (u, v) is an l-edge, then $u_x < v_x$ and $u_y < v_y$, i.e., (u, v) is a left edge, in Γ_i. If (u, v) is an r-edge, then $u_x < v_x$ and $u_y > v_y$, i.e., (u, v) is a right edge, in Γ_i.

3. Let Q_l be the set of rays obtained by shooting for every right segment s that has an end point on the outer face of Γ_i, an upward ray with origin $tip(s)$ and slope $slope(s)$. Then any two rays of Q_l are divergent. Analogously, we define a set of rays Q_r for the right segments, which must be divergent.

4. No ray in $Q_l \cup Q_r$ intersects Γ_i except at its origin. Any two rays $r \in Q_l$ and $r' \in Q_r$ intersect if and only if the origin of r precedes the origin of r' on P_i.

We now add v_3 such that Γ_i is an isosceles triangle with $\angle v_3v_1v_2 = \angle v_3v_2v_1$, as shown in Figure 4(c). Since (v_1, v_3) and (v_3, v_2) are the only l- and r- edges in Γ_3, Invariants 1–4 are straightforward to verify. Assume that the invariants hold for the additions of v_i, where $i < n$, and let Γ_i be the drawing of G_i that respects Invariants 1–4. We now show how to add v_{i+1} to Γ_i

such that the constructed drawing Γ_{i+1} respects all the invariants.

We call a vertex $w \in P_i$ a *peak vertex* if all of its neighbors have smaller y-coordinates than w_y in Γ_i. The distinction between 'tip' and 'peak' is important, i.e., a vertex w may be a tip of some left (respectively, right) segment, but w is not a peak unless it is also a tip of some right (respectively, left) segment.

Let $w_l, w_{l+1}, \ldots, w_r$ be the neighbors of v_{i+1} in G_i. Note that (w_l, v_{l+1}) and (v_{l+1}, w_r) are the l- and r-edges of v_{l+1}, respectively. We now consider the following three cases. For convenience we assume that v_1 and v_2 are the tips of some left and right segments, respectively, such that the cases when $v_1(= w_l)$ or $v_2(= w_r)$ are handled by Case 2.

Case 1 (w_l is a tip of some left segment and w_r is a tip of some right segment): We claim that the segment containing w_l is different than the segment containing w_r. Otherwise, without loss of generality assume that both lie on some right segment s. By definition, $w_{ly} > w_{ry}$. Hence w_r cannot be a tip of s. If w_r is a tip of some right segment s' other than s, as shown in Figure 4(b), then Invariant 2 will imply that w_r is a child of two different vertices in T_r, which is a contradiction that T_r is a tree. Note that we can use the above argument to claim that w_r cannot be an internal vertex of a right segment, and similarly, w_l cannot be an internal vertex of a left segment. Figure 5(a–d) depict the remaining four scenarios.

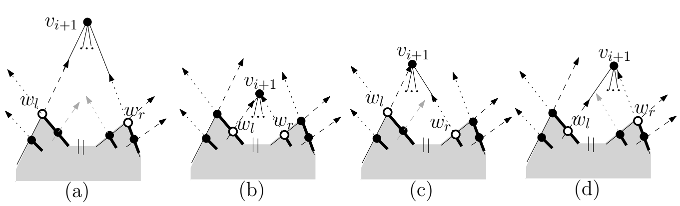

Figure 5: (a–d) Illustration for different drawings in Case 1.

Observe now that by Invariants 3–4, the ray in Q_l emanating from w_l intersects the ray in Q_r emanating from w_r, and none of these rays intersect Γ_i. Let c be the intersection point of these two rays. We place v_{i+1} at c and draw the edges (v_{i+1}, w), where $w \in \{w_l, w_{l+1}, \ldots, w_r\}$. We claim that the drawing of the m-edges does not create any edge crossing, as follows. By invariant 3, all the right edges in the path $w_l, w_{l+1}, \ldots, w_r$ have slope larger than $slope(v_{i+1}, w_r)$. Similarly, all the left edges have slope smaller than $slope(v_{i+1}, w_l)$. Hence by Lemma 1, the drawing of the m-edges does not create any edge crossings.

Case 2 (w_l is a tip of some right segment and w_r is a tip of some left segment): If Case 1 is also satisfied, i.e., if w_l and w_r both are peaks in Γ_i, then we add v_{i+1} as in Case 1. Otherwise, at most one of w_l and w_r are peaks.

Case 2A. If none of w_l and w_r are peaks, then we construct two rays r_1 and r_2 starting from w_l and w_r, respectively, such the slope of r_1 is $slope(w_l, w_{l+1})+\epsilon_1$, and the slope of r_2 is $slope(w_{r-1}, w_r) - \epsilon_2$. Here ϵ_1 and ϵ_2 are two constants such that the sets Q_l and Q_r respect Invariant 3. Lemma 2 guarantees the existence of such constants. Figure 6(a) illustrates such a scenario. We then place the vertex v_{i+1} at the intersection point of r_1 and r_2, and draw its l-, r- and m-edges. By Lemma 1, the drawing of these edges does not create any edge crossing.

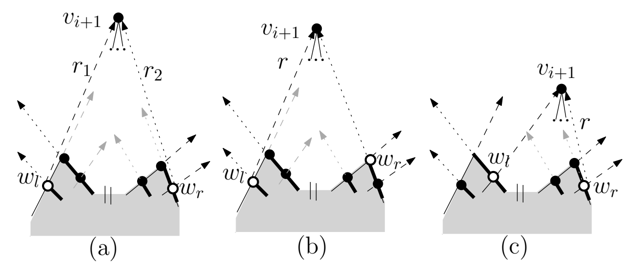

Figure 6: (a–b) Illustration for Case 2. (c) Illustration for Case 3.

Case 2B. If exactly one of w_l and w_r is a peak, then without loss of generality assume that w_r is a peak vertex. We then construct a ray r starting from w_l with $slope(w_l, w_{l+1}) + \epsilon$ such that the rays of $Q_l \cup \{r\}$ are divergent and maintain Invariant 3. Lemma 2 guarantees the existence of such a constant ϵ. Figure 6(b) illustrates this scenario. We then place the vertex v_{i+1} at the intersection point of r and the ray in Q_r emanating from w_r. Finally, we draw the l-, r- and m-edges of v_{i+1}. Lemma 1 ensures that the drawing of these edges does not create any edge crossing.

Case 3 (w_l and w_r both are tips of the same types of segments): Consider first the case when at least one of w_l and w_r is a peak. If both are peaks, then we follow Case 1. Otherwise, exactly one of them is a peak. If w_l is a peak, then we insert v_{i+1} following either Case 1 or Case 2B depending on whether w_r is a tip of some right or left segment. Similarly, if w_r is a peak, then we insert v_{i+1} following either Case 1 or Case 2B depending on whether w_l is a tip of some left or right segment. Finally, if none of w_l and w_r is a peak, without loss of generality assume that both w_l and w_r are tips of some left segments.

In such a scenario we construct a ray r starting from w_r with $slope(w_r, w_{r-1})-\epsilon$ such that the rays of $Q_r \cup \{r\}$ are divergent and maintain Invariant 3. Lemma 2 guarantees the existence of such a constant ϵ. Figure 6(c) illustrates this scenario. We then place the vertex v_{i+1}

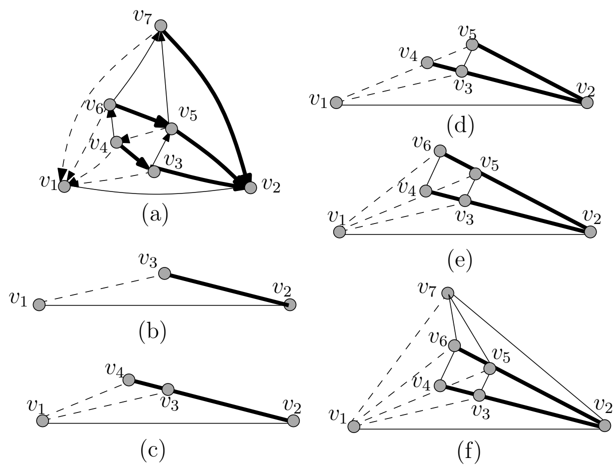

Figure 7: A plane graph and the incremental construction of its drawing.

at the intersection point of r and the ray in Q_l emanating from w_l. Finally, we draw the l-, r- and m-edges of v_{i+1}. Lemma 1 ensures that the drawing of these edges does not create any edge crossing.

This completes the description of our drawing algorithm. Figure 7 illustrates a drawing computed by our algorithm.

Γ_{i+1} **respects Invariants 1–4:** According to our construction, v_{i+1} is a peak in Γ_{i+1} such that $w_{lx} < v_{i+1x} < w_{rx}$. Hence Invariants 1–2 hold for Γ_{i+1} in all the three cases. We now consider Invariants 3–4. Since Case 1 does not increase the number of rays, it is straightforward to check that Γ_{i+1} respects these invariants. On the other hand, Cases 2–3 create new rays. Note that these new rays have been constructed according to Lemma 2, which ensures that for any new ray $r \in Q_l$ (respectively, $r' \in Q_r$), the set $r \cup Q_l$ (respectively, $r' \cup Q_r$) is divergent. Since all the rays have origin on P_{i+1}, it is straightforward to observe that the ray r intersects all the other rays that belong to Q_r and appear after r while visiting P_{i+1} from left to right. The rays emanating from $w_{l+1}, \ldots, w_{r-1}$ in Γ_i disappears in Γ_{i+1}. Hence no ray in Q_l and Q_r in Γ_{i+1} intersects Γ_{i+1} except at its origin.

3.2 Computing the Upper Bound

Let $\Gamma = \Gamma_n$ be the drawing of G computed using the above drawing algorithm. Let T_l, T_r, T_m be the Schnyder realizer that corresponds to σ. We now claim that the drawing has $leaf(T_l) + leaf(T_r) + n$ segments.

Lemma 3 *Let G be plane triangulation. Then Algorithm* FewSegDraw *computes a drawing Γ of G with $leaf(T_l) + leaf(T_r) + n$ segments, where T_l and T_r are a pair of trees in a Schnyder realizer of G.*

Proof. The idea is to show that the drawings of T_l and T_r has $leaf(T_l)$ and $leaf(T_r)$ segments in Γ, respectively. Since $G \setminus (T_l \cup T_r)$ has n edges, the claim follows.

Let Γ'_i, where $3 \le i \le n$, be the drawing obtained from Γ_i by deleting the edges of T_m. While adding v_i, the algorithm adds one edge of T_l (i.e., the l-edge of v_{i+1}), and one edge of T_r (i.e., the r-edge of v_{i+1}) to Γ'_{i-1}. Case 1 does not create any new segment. A new segment in the drawing of T_l and T_r can appear only in Cases 2–3. Whenever the algorithm creates a new segment above P_{i-1}, it ensures that the corresponding vertex w on P_{i-1} is an internal vertex of some left or right segment in Γ'_{i-1}. For example, see Figure 6.

We claim that any segment that starts at some nonleaf vertex of T_l, ends at some leaf of T_l, which will imply that the drawing of T_l has at most $leaf(T_l)$ segments. Suppose for a contradiction that there exists a left segment s that starts at some nonleaf vertex w of T_l and ends at some nonleaf vertex w' of T_l. If w' is not internal to any other segment in Γ'_n, then it is a leaf, and the claim holds. Otherwise, let w' be an internal vertex of some segment s'. If the segment s' is a right segment, then the property that w' is an end point of s will imply that w' is a leaf of T_l, which is a contradiction. The remaining scenario, where s' is a left segment, implies that w' is a child of two different parents, which contradicts that T_l is a tree. Similarly, the drawing of T_r has at most $leaf(T_r)$ segments. □

In a minimum Schnyder realizer T_l, T_r, T_m of G, we have $leaf(T_l) + leaf(T_r) + leaf(T_m) = 2n - 5 - \Delta_0$ [1], where $0 \le \Delta_0 \le \lfloor (n-1)/2 \rfloor$. Note that the tree with the largest number of leaves must have at least $(2n - 5 - \Delta_0)/3$ leaves. Hence the remaining two trees have at most $2(2n - 5 - \Delta_0)/3 \le (4n - 2\Delta_0 - 10)/3$ leaves. Using Lemma 3 we obtain the following theorem.

Theorem 4 *Let G be an n-vertex plane triangulation. Then G admits a drawing with at most $(7n - 2\Delta_0 - 10)/3$ segments.*

3.3 Constraints and Generalizations

We can improve the upper bound of $7n/3 - O(1)$ segments for triangulations to $9n/4 - O(1)$ segments under 4-connectivity constraint, as follows.

For any Schnyder realizer, $leaf(T_l) + leaf(T_r) + leaf(T_m) = 2n - 5 - \Delta$ [1], where Δ is the number of cyclic faces. Zhang and He [9] showed that for 4-connected triangulations, there exists a canonical ordering tree with at most $(n+1)/2$ leaves. Without loss of generality assume that $leaf(T_l) \le (n+1)/2$. Then $leaf(T_r) + leaf(T_m) \le 2n - 5 - leaf(T_l)$. Hence either T_r or T_m has at most $(2n - 5 - leaf(T_l))/2$ leaves. Without loss of generality assume that $leaf(T_m) \le (2n - 5 - leaf(T_l))/2$. Therefore, $leaf(T_m) + leaf(T_l) \le$

$(2n-5)/2 - leaf(T_l)/2 + leaf(T_l) = (2n-5)/2 + leaf(T_l)/2$. Since $leaf(T_l) \leq (n+1)/2$, we have $leaf(T_m) + leaf(T_l) \leq (5n-9)/4$. In summary, there exists a Schnyder realizer such that two of its trees has at most $(5n-9)/4$ leaves. Using Lemma 3 we obtain the following theorem.

Theorem 5 *Let G be an n-vertex 4-connected plane triangulation. Then G admits a drawing with at most $(9n-9)/4$ segments.*

It is straightforward to use our algorithm to draw general plane graphs: Given a plane graph G, we first triangulate the graph, then draw the triangulation with $(7n-10)/3$ segments using Theorem 4, and finally remove the added edges. Note that removal of edges may increase the number of segments in the drawing. Since removal of one edge from any segment of some straight-line drawing can increase the number of segments by at most one, the over all increase in the number of segments is at most the total number of edges removed. Since an n-vertex triangulation has exactly $m = 3n-6$ edges, the drawing we obtain can have at most $(7n-10)/3 + (3n-6-m) = (16n-3m-28)/3$ segments.

Theorem 6 *Let G be a plane graph with n vertices and m edges. Then G admits a straight-line drawing with at most $(16n-3m-28)/3$ segments.*

Dujmović et al. [3] gave an algorithm to draw n-vertex m-edge 3-connected plane graphs with at most $\min\{m - n/2 + \alpha - 3, m - \alpha\}$ segments, where the parameter α lies in the interval $[0, 3n-6-m]$, giving an upper bound of $2.5n$ segments. Theorem 6 gives a better upper bound when the graph is dense, i.e., when $m \geq 2.84n$.

4 Conclusion

In this paper we have given an algorithm to draw any n-vertex plane triangulation with at most $7n/3$ segments, which improves to $9n/4$ when the input triangulation is 4-connected. Since the realizers we use can be computed in linear time [9], our algorithm runs in linear time.

Dujmović et al. [3] showed that the lower bounds on the number of segments for the general plane triangulations and 4-connected plane triangulations are $2n-2$ and $2n-6$ (Figure 8), respectively. A natural open question is to reduce the gap between the lower and upper bounds.

Another limitation of the drawings we compute is the rational coordinates for vertex positions, which may be exponential. Thus it would be interesting to examine the area requirement of these drawings, where the vertices are restricted to integer grid points.

Since a k-segment drawing is an arrangement of a set of k straight line segments, an interesting generalization would be to represent planar graphs as arrangement of other objects such as circles, ellipses and lower order splines. Recently, Schulz [8] has presented such a generalization considering circular arcs.

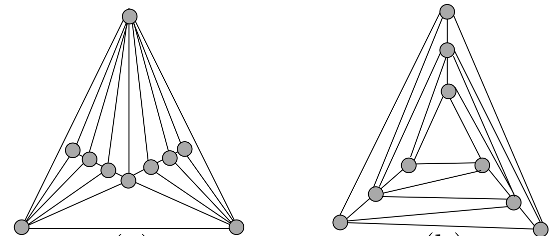

Figure 8: (a) Illustration for lower bounds for general plane triangulations. (b) A 'Nested triangle graph', which is also a 4-connected triangulation. Such a graph requires $2n-6$ segments even when the embedding is not fixed [3].

References

[1] N. Bonichon, B. L. Saëc, and M. Mosbah. Wagner's theorem on realizers. In *Colloquium on Automata, Languages and Programming (ICALP 2002)*, volume 2380 of *LNCS*, pages 1043–1053. Springer, 2002.

[2] H. De Fraysseix, J. Pach, and R. Pollack. How to draw a planar graph on a grid. *Combinatorica*, 10(1):41–51, 1990.

[3] V. Dujmović, D. Eppstein, M. Suderman, and D. R. Wood. Drawings of planar graphs with few slopes and segments. *Computational Geometry*, 38(3):194–212, 2007.

[4] S. Durocher, D. Mondal, R. I. Nishat, and S. Whitesides. A note on minimum-segment drawings of planar graphs. *Journal of Graph Algorithms and Applications*, 17(3):301–328, 2013.

[5] D. Mondal, R. I. Nishat, S. Biswas, and M. S. Rahman. Minimum-segment convex drawings of 3-connected cubic plane graphs. *Journal of Combinatorial Optimization*, 25(3):460–480, 2013.

[6] M. A. H. Samee, M. J. Alam, M. A. Adnan, and M. S. Rahman. Minimum segment drawings of series-parallel graphs with the maximum degree three. In *Graph Drawing*, volume 5417 of *LNCS*, pages 408–419. Springer, 2008.

[7] W. Schnyder. Embedding planar graphs on the grid. In *ACM-SIAM Symposium on Discrete Algorithms (SODA 1990)*, pages 138–148. ACM, January 22–24 1990.

[8] A. Schulz. Drawing graphs with few arcs. In *Workshop on Graph-Theoretic Concepts in Computer Science (WG 2013)*, volume 8165 of *LNCS*, pages 406–417. Springer, 2013.

[9] H. Zhang and X. He. Canonical ordering trees and their applications in graph drawing. *Discrete & Computational Geometry*, 33(2):321–344, February 2005.