Trade-Offs in Planar Polyline Drawings

Stephane Durocher* and Debajyoti Mondal**

Department of Computer Science, University of Manitoba, Canada
{durocher,jyoti}@cs.umanitoba.ca

Abstract. Angular resolution, area and the number of bends are some important aesthetic criteria of a polyline drawing. Although trade-offs among these criteria have been examined over the past decades, many of these trade-offs are still not known to be optimal. In this paper we give a new technique to compute polyline drawings for planar triangulations. Our algorithm is simple and intuitive, yet implies significant improvement over the known results. We present the first smooth trade-off between the area and angular resolution for 2-bend polyline drawings of any given planar graph. Specifically, for any given n-vertex triangulation, our algorithm computes a drawing with angular resolution $r/d(v)$ at each vertex v, and area $f(n, r)$, for any $r \in (0, 1]$, where $d(v)$ denotes the degree at v. For $r < 0.389$ or $r > 0.5$, $f(n, r)$ is less than the drawing area required by previous algorithms; $f(n, r)$ ranges from $7.12n^2$ when $r \leq 0.3$ to $32.12n^2$ when $r = 1$.

1 Introduction

Polyline drawing is a classic style of drawing planar graphs, which has a wide range of applications in the area of software visualization [8,18] and layout of circuit diagrams [7]. Given an n-vertex planar graph G, a *polyline drawing* Γ of G maps each vertex to a distinct point in $\mathbb{R}^2$, and each edge to a simple polygonal chain between its endpoints such that no two edges intersect except possibly at their common end point. Γ is a *k-bend polyline drawing* if the number of line segments per edge is bounded by at most $k + 1$, i.e., each edge contains at most k bend points. Consequently, a k-bend polyline drawing can be considered as a $(k + \lambda)$-bend drawing for any $\lambda > 0$. Figures 1(a) and (b) illustrate a plane graph G and a 2-bend polyline drawing of G, respectively.

Researchers have examined the theoretical aspects of planar polyline drawings over a long time [2,4,9,10,13,17,20]. *Area* (i.e., the size of the smallest integer grid containing the drawing), *angular resolution* (i.e., the smallest angle formed at any vertex), number of bends per edge, edge separation and bend resolution are some examples of such aesthetic criteria. Even after decades of research effort, finding the optimal trade-off between the number of total bends and area still

* Work of the author is supported in part by the Natural Sciences and Engineering Research Council of Canada (NSERC).

** Work of the author is supported in part by a University of Manitoba Graduate Fellowship.

C. Duncan and A. Symvonis (Eds.): GD 2014, LNCS 8871, pp. 306–318, 2014.

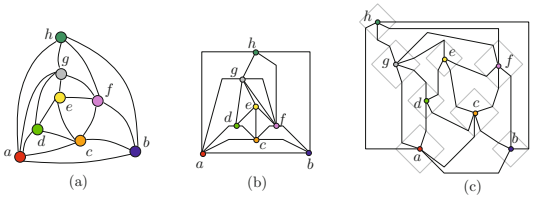

Fig. 1. (a) A planar graph G. (b)–(c) Two polyline drawings of G

seems to be an elusive goal. For example, every planar triangulation with n vertices admits a straight-line drawing (i.e., a 0-bend polyline drawing) in $O(n^2)$ area [9]. Several improvements on the constant hidden in $O(.)$ notation have been achieved [2,4,9,17], and the best known bound is $8n^2/9 = 0.89n^2$ [4]. Better upper bounds, i.e., $4n^2/9 < 0.45n^2$, can be attained if we use 1-bend polyline drawings, which takes at most $2n/3$ bends in total [20]. Although these drawings require small area, the compactness comes at the expense of bad angular resolution, i.e., $\Omega(1/n)$. Garg and Tamassia [19] showed that there exists planar graphs such that any of its straight-line drawing with angular resolution $\Omega(1/\rho)$ requires at least $\Omega(c^{\rho n})$ area, where $c > 1$, which suggests that drawings with angular resolution $\Omega(1/\Delta)$ and polynomial area may exist only if we allow the edges to have bends.

Allowing bends helps both to reduce area and to improve angular resolution, e.g., given an n-vertex planar graph with maximum degree Δ, one can construct a 3-bend polyline drawing with $2/\Delta$ radians of resolution and $3n^2$ area [13]. The angular resolution can be improved to $\Omega(1/d(v))$ radians (for each vertex v) with an expense of higher area [10,12], which also helps to reduce the number of bends per edge. Table 1 presents a brief summary of the related results.

Table 1. Angular resolution, area and total bends in k-bend polyline drawings, where $\alpha \in [1/4, 1/2]$, and $\beta \in [1/3, 1]$.

Graph Class	Area	Resolution	k-Bends	T. Bends	Reference
Maximal Planar	$7n^2/8$	$\Omega(1/n^2)$	0	0	[4]
Maximal Planar	$9n^2/2$	$\Omega(1/n)$	0	0	[16]
Maximal Planar	$12.5n^2$	$0.5/d(v)$	1	$3n$	[10]
Maximal Planar	$450n^2$	$1/d(v)$	1	$3n$	[5]
Maximal Planar	$4n^2/9$	$\Omega(1/n^2)$	1	$2n/3$	[20]
Maximal Planar	$200n^2$	$1/d(v)$	2	$6n$	[12]
Maximal Planar	$\mathbf{(6\alpha + 8/3)^2 n^2}$	$\frac{\alpha}{(d(v)(\alpha^2+1/4)}$	**2**	**5.5n**	**Theorem 2**
Maximal Planar	$\mathbf{(6\beta + 2/3)^2 n^2}$	$\frac{\beta}{(d(v)(\beta^2+1)}$	**2**	**5.5n**	**Theorem 3**
3-connected Planar	$6n^2$	$2/\Delta$	3	$5n - 15$	[15]
General Planar	$3n^2$	$2/\Delta$	3	$5n - 15$	[13]

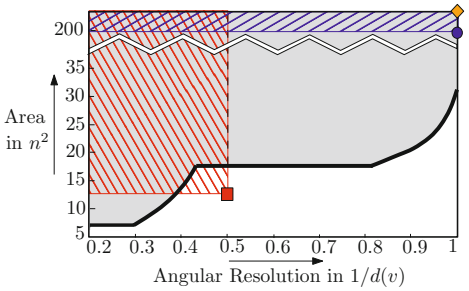

Fig. 2. Trade-off between angular resolution and area for 2-bend polyline drawings, where the bold line denotes the trade-off established in this paper. The square, circle and diamond denote the reference [10], [12] and [5], respectively.

Early polyline drawing algorithms were developed as a generalization of orthogonal drawings [1]. Before Duncan and Kobourov's algorithm [10], all the polyline drawing techniques with good angular resolution and $O(n^2)$ area were based on the idea of assigning an empty square surrounding each vertex (e.g., Figure 1(c)), which forced the constant factor in the $O(.)$ notation to be very large. The algorithm of Duncan and Kobourov [10] finds a drawing with smaller area, but loses the square-emptiness property around the vertices, as well as decreasing the angular resolution by a factor of 2. Observe that two solutions in a multi-objective optimization are comparable if and only if one of them dominates the other with respect to every optimization criteria. Hence although the drawing of [10] has smaller area than that of [5] (see Table 1), it is not an improvement over [5] because of its lower angular resolution.

Contributions. In this paper we examine the trade-offs between the angular resolution and area for 2-bend polyline drawings of planar triangulations. Figure 2 illustrates the solution space dominated by our algorithm in gray, which dominates all the previous 2-bend polyline drawing algorithms except Duncan and Kobourov's algorithm [10], which dominates our algorithm along a small interval of X-axis. Even under the model where each vertex v is surrounded by an empty square of size $d(v) \times d(v)$, we can construct a 2-bend polyline drawing with angular resolution $1/\Delta$ and area $32.12n^2$, where the best known bounds can achieve an $\Omega(1/d(v))$ angular resolution with an area at least $200n^2$ [5,12,14], or an $1/\Delta$ angular resolution with 3 bends per edge [13].

2 Technical Background

Let G be a *plane graph*, i.e., a planar graph with a fixed combinatorial embedding and a specified outerface. If every face of G including (respectively, excluding) the outer face is a cycle of length three, then G is called *triangulated* (respectively, *internally triangulated*). Let G be an n-vertex triangulated plane graph, where v_1, v_2 and v_n are the outer vertices of G in clockwise order, and

let $\sigma = (v_1, v_2, ..., v_n)$ be an ordering of all the vertices of G. Then G_k, where $2 \leq k \leq n$, is the subgraph of G induced by $v_1 \cup v_2 \cup ... \cup v_k$, and P_k is the path (while walking clockwise) on the outer face of G_k that starts at v_1 and ends at v_2. The vertex-ordering σ is called a *canonical ordering* [9] with respect to the outer edge (v_1, v_2) if for each k, $3 \leq k \leq n$, the following conditions are satisfied: (a) G_k is 2-connected and internally triangulated. (b) If $k \leq n$, then v_k is an outer vertex of G_k and the neighbors of v_k in G_{k-1} appears consecutively on P_{k-1}. Figures 3(a)–(b) illustrate an example.

For some j, where $3 \leq j \leq n$, let P_j 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_j) and (v_j, w_r) the *l-edge* and the *r-edge* of v_j, respectively. The other edges incident to v_j in G_j are called the *m-edges* of v_j. For example, in Figure 3(c), the edges (v_6, v_4), (v_6, v_5), and (v_3, v_6) are the l-, r- and m-edges of v_6, respectively. By $d_l(v), d_r(v)$ and $d_m(v)$ we denote the number of l, r and m-edges that are incoming to v, e.g., $d_l(v_6) = 0, d_r(v_6) = 1$ and $d_m(v_6) = 1$.

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 3(d)), 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* [17] of G. A Schnyder realizer is called a *minimum realizer* if all the cyclic inner faces are oriented clockwise. By Δ_0 we denote the number of cyclic inner faces in the minimum realizer [21]. If $\{T_l, T_r, T_m\}$ is a minimum Schnyder realizer of G, then we have $\texttt{leaf}(T_l) + \texttt{leaf}(T_r) + \texttt{leaf}(T_m) = 2n - 5 - \Delta_0$ [3]. Hence we can observe the following property.

Remark 1. *Let $\{T_l, T_r, T_m\}$ be a minimum Schnyder realizer of an n-vertex triangulation. Then* $\min\{\texttt{leaf}(T_l) + \texttt{leaf}(T_r), \texttt{leaf}(T_l) + \texttt{leaf}(T_m), \texttt{leaf}(T_r) + \texttt{leaf}(T_m)\} \leq (4n - 2\Delta_0 - 10)/3$.

A non-root vertex in T_l is called a *primary vertex of T_l* if it is the first child of its parent in the clockwise order. Similarly, a non-root vertex in T_r is a *primary vertex of T_r* if it is the first child of its parent in the anticlockwise order. We now have the following lemma, whose proof is omitted due to space constraints.

Lemma 1. *Let n_l and n_r be the nonprimary vertices in T_l and T_r, respectively. Then* $n_l + n_r \leq \texttt{leaf}(T_l) + \texttt{leaf}(T_r)$.

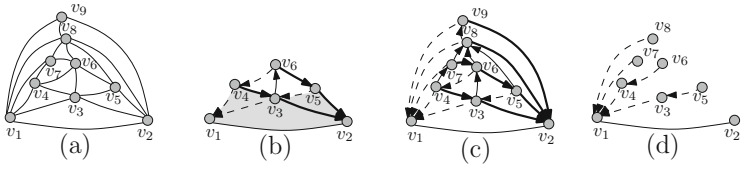

Fig. 3. (a) A canonical ordering of a plane triangulation G. (b) G_6. (c) The l-, r- and m- edges are shown in dashed, bold-solid, and thin-solid edges respectively. (d) T_l.

In a *plus-contact representation* of G, each vertex of G is represented as an axis-aligned plus shape (i.e., a shape consisting of two intersecting line segments) such that two plus shapes touch if and only if their corresponding vertices are adjacent in G [11]. Let Γ be a plus contact representation, and let v be any vertex in Γ. Then by $\mathtt{P}(v)$ we denote the plus-shape that corresponds to v in Γ. By the *center* $\mathtt{C}(v)$ of $\mathtt{P}(v)$, we denote the intersection point of the vertical and horizontal straight line segments of $\mathtt{P}(v)$. The four straight line segments that start at $\mathtt{C}(v)$ and extend to the left, right, above and below $\mathtt{C}(v)$ are the *left, right, up and down hands* of v, which we denote by $\mathtt{L}(v), \mathtt{R}(v), \mathtt{U}(v)$ and $\mathtt{D}(v)$, respectively. A *j-shift operation* on Γ with respect to an infinite horizontal line (respectively, vertical line) ℓ is performed as follows: Remove all the edges that are lying completely above (respectively, to the right of) ℓ. Increase the y-coordinate (respectively, x-coordinate) of every vertex lying above (respectively, to the right of) ℓ by j units. Draw the edges that were removed using the new vertex positions. Extend the edges intersected by ℓ upwards (respectively, to the right) until they reach to their other endpoint.

3 Polyline Drawing

Let G be an n-vertex maximal planar graph. We construct the drawing of G in three phases. In the first phase we construct a plus-contact representation of $G \setminus T_m$ on a rectangular grid. In the next phase we expand the drawing by inserting dummy grid lines, and in the third phase we use these grid lines to draw the edges of T_m, and route the l- and r-edges avoiding degeneracy.

Phase 1 (Plus-Contact): Let $\sigma = (v_1, v_2, \ldots, v_n)$ be a canonical ordering of G and let $\{T_l, T_r, T_m\}$ be the corresponding Schnyder realizer. Let Γ_k, where $2 \leq k \leq n$, be the drawing of all the edges of G_k except the m-edges. We first construct the drawing Γ_2 for G_2, as follows. Place $\mathtt{C}(v_1)$ and $\mathtt{C}(v_2)$ at coordinates $(1, 2)$ and $(2, 1)$, respectively. Then the horizontal and vertical unit-segments to the left and below $(1, 2)$ correspond to $\mathtt{L}(v_1)$ and $\mathtt{D}(v_1)$, respectively. Similarly, the horizontal and vertical unit-segments to the left and below $(2, 1)$ correspond to $\mathtt{L}(v_2)$ and $\mathtt{D}(v_2)$, respectively, as illustrated in Figure 4(b). We now insert the vertices in the canonical ordering maintaining the following invariants. While inserting a new vertex, we only draw the l and r-edges.

- $\mathcal{I}_1$. The upper envelope of Γ_i is x-monotone, where the upper envelope is determined by the left and down hands of the vertices in P_i.
- $\mathcal{I}_2$. The ray with slope $+1$ starting at any outer vertex of Γ_i can be extended towards infinity avoiding any edge crossing.
- $\mathcal{I}_3$. Every l-edge starts as a left hand of some plus shape and ends either at a center or at a down hand of some other plus shape.
- $\mathcal{I}_4$. Every r-edge starts as a down hand of some plus shape and ends either at a center or at a left hand of some other plus shape.

Since the upper envelope of G_2 forms a staircase, and does not contain any l- or r-edge, it is straightforward to verify the invariants for Γ_2. We now assume

that invariants $\mathcal{I}_1$–$\mathcal{I}_4$ hold for $G_2, G_3, \ldots, G_{k-1}$, where $k-1 < n$, and consider the insertion of vertex v_k.

Insertion of v_k: Let $w_l, w_{l+1}, \ldots, w_{r-1}, w_r$ be the neighbors of v_k on P_{k-1}. Consider an infinite horizontal line ℓ_h that lies in between the horizontal grid line determined by $\mathtt{L}(w_l)$ and the horizontal grid line immediately below $\mathtt{L}(w_l)$. Similarly, let ℓ_v be an infinite vertical line that lies in between the vertical grid line determined by $\mathtt{D}(w_r)$ and the vertical grid line immediately to the left of $\mathtt{D}(w_r)$. We now add v_k considering the following cases. The case when $k = n$ is special, which is handled by Case 4.

Case 1 (v_k is a nonprimary vertex in both T_l and T_r): We first perform a 1-shift with respect to ℓ_h. This increases the number of horizontal lines by 1 and ensures that $\mathtt{D}(w_l)$ contains at least 1 grid point p that does not contain any vertex or contact point. Similarly, we perform a 1-shift with respect to ℓ_v, which increases the number of vertical lines by 1 and ensures that $\mathtt{L}(w_r)$ contains at least 1 grid point q that does not contain any vertex or contact point. We now consider the horizontal ray r_p that starts at p. Since the upper envelope of Γ_{k-1} is x monotone and p does not contain any vertex or contact point, r_p does not intersect Γ_{k-1} except at p. Similarly, we define a vertical ray r_q that starts at q, which does not intersect Γ_{k-1} except at q. We now place v_k at the intersection point of r_p and r_q, and draw the edges (v_k, w_l) and (v_k, w_r). Since r_p and r_q do not intersect Γ_{k-1} except at p and q, respectively, drawing of these edges does not introduce any crossing. Figure 4(c) illustrates such a scenario. We omit the proof that Γ_k respects the invariants $\mathcal{I}_1$–$\mathcal{I}_4$ due to space constraints.

Case 2 (v_k is a primary vertex in T_l but a nonprimary vertex in T_r): In this case we perform a 1-shift with respect to ℓ_v, which increases the number of vertical lines by 1 and ensures that $\mathtt{L}(w_r)$ contains at least 1 grid point q that does not contain any vertex or contact point. Assume that $p = \mathtt{C}(w_l)$. We now consider the horizontal ray r_p that starts at p. Since the upper envelope of Γ_{k-1} is x monotone and p does not contain any vertex or contact point, r_p does not intersect Γ_{k-1} except at p. Similarly, we define a vertical ray r_q starting at q, which does not intersect Γ_{k-1} except at q. We now place v_k at the intersection point of r_p and r_q, and draw the edges (v_k, w_l) and (v_k, w_r). Figure 4(e) illustrates such a scenario.

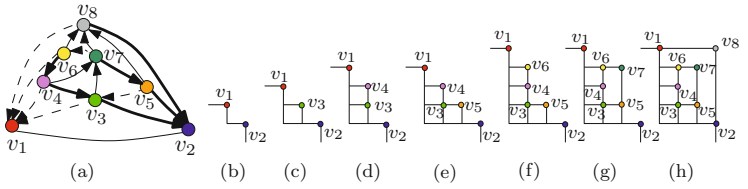

Fig. 4. (a) A plane graph G and a minimum Schnyder realizer of G. (b)–(h) Illustration for the drawing of $G \setminus T_m$.

Case 3 (v_k is a nonprimary vertex in T_l but a primary vertex in T_r): This case is symmetric to Case 2, i.e., we perform a 1-shift with respect to ℓ_h to obtain a new grid point p on $\mathtt{D}(w_l)$ and assume that $q = \mathtt{C}(w_r)$.
Case 4 (v_k is a primary vertex in both T_l and T_r): In this case we do not perform any shift, and assume that $p = \mathtt{C}(w_l)$ and $q = \mathtt{C}(w_r)$.

We now have the following lemma whose proof is omitted due to space constraints.

Lemma 2. *Γ_n is a drawing on a $(W + 2) \times (H + 2)$ grid, where $W + H \leq \mathtt{leaf}(T_l) + \mathtt{leaf}(T_r)$.*

Phase 2 (Expansion): For any plus-contact representation on an integer grid, we define a *free grid line* as a grid line that does not contain any vertex-center or contact points. We refer the reader to Figure 5.

Consider the horizontal grid lines from top to bottom. For every horizontal grid line ℓ containing at least one vertex of Γ, we now perform two $\lfloor d(v)/2 \rfloor$-shifts, where v is the vertex with the largest degree over all the vertices on ℓ. Let ℓ_h (respectively, ℓ'_h) be an infinite horizontal line that lies in between the horizontal grid line ℓ and the horizontal grid line immediately below (respectively, above) ℓ. Perform a $\lfloor d(v)/2 \rfloor$-shift with respect to ℓ_h, and then a $\lfloor d(v)/2 \rfloor$-shift with respect to ℓ'_h. Observe that for each vertex w on ℓ, we now have a set of $\lfloor d(v)/2 \rfloor$ free grid lines above w and a set of $\lfloor d(v)/2 \rfloor$ free grid lines below w. We consider a corresponding set S_w that consists of these $2\lfloor d(v)/2 \rfloor$ free grid lines along with the line ℓ. Furthermore, we assume that the grid lines of S_w are ordered in the increasing order of y-coordinates. Figure 5(b) illustrates S_{v_4}.

Similarly, we consider the vertical grid lines from right to left, and for every vertical grid line ℓ' containing at least one vertex of Γ, we perform two $\lfloor d(v)/2 \rfloor$-shifts to the left and right side of ℓ', where v is the vertex with the largest degree over all the vertices on ℓ'. We consider a corresponding set S'_w that contains these $2\lfloor d(v)/2 \rfloor$ free vertical grid lines along with the line ℓ', where the lines are ordered in the decreasing order of x-coordinates. Let the resulting drawing be Γ'_n, as shown in Figure 5(c). The following property is a straightforward consequence of the Expansion phase.

Remark 2. *For every vertex v in Γ'_n, the point $\mathtt{C}(v)$ lies at the center of an integer grid A_v of size $(2\lfloor d(v)/2 \rfloor + 1) \times (2\lfloor d(v)/2 \rfloor + 1)$. The grid A_v does not contain any vertex, contact point, or edge of Γ' except the four hands of v. Furthermore, for any other vertex $u(\neq v)$, the grids A_u and A_v are disjoint, i.e., they do not share any common grid point.*

Phase 3 (Edge Routing): For each vertex in canonical order, we first route the incoming m-edges incident to v_k, as follows. Recall that the m-edges start at the vertices $w_{l+1}, \ldots, w_{r-1}$ and ends at v_k.

By the construction of Γ'_n, the vertices $w_{l+1}, \ldots, w_{r-1}$ lie below S_{v_k} and to the left of S'_{v_k}. Hence all the boundary grid points of A_{v_k}, which lie below S_{v_k} and to the left of S'_{v_k}, are visible from the top-right corner c_{w_j} of A_{w_j}, for all $l + 1 \leq j \leq r - 1$. Assume that $z = \lceil d_m(v)/2 \rceil$. Let M be the monotone chain

determined by the last line of S_w and first line of S'_w, where $w \in \{w_{l+1}, \dots w_{r-1}\}$. Figure 5(d) illustrates M with a dotted line. For each $w \in \{w_{l+1}, \dots w_z\}$, we now route the m-edge incident to w through the top-right corner c_w upto M, and then to a distinct grid point on the leftmost boundary of A_{v_k} below $\mathtt{L}(v_k)$. Observe that $\lceil d_m(v_k)/2 \rceil \leq d_m(v_k)/2 + 1 \leq (d(v_k) - 3)/2 + 1 \leq (d(v_k) - 1)/2$. Since $(d(v_k) - 1)/2$ is at most $\lfloor d(v_k)/2 \rfloor$ (irrespective of the parity of $d(v_k)$), the grid points on the leftmost boundary of A_{v_k} below $\mathtt{L}(v_k)$ are sufficient to route all the m-edges incident to $\{w_{l+1}, \dots w_z\}$. Similarly, for each $w \in \{w_{z+1}, \dots w_{r-1}\}$, we now route the m-edge incident to w through the top-right corner c_w upto M, and then to a distinct grid point to the left of $\mathtt{D}(v_k)$ on the bottommost boundary of A_{v_k}. Since $\lfloor d_m(v_k)/2 \rfloor \leq \lfloor d(v_k)/2 \rfloor - 1$ (irrespective of the parity of $d(v_k)$), we have sufficient number of boundary points to route all the m-edges incident to $\{w_{z+1}, \dots w_{r-1}\}$.

The l- and r-edges of Γ'_n contain edge overlapping on the left and down hands. From the Expansion phase it is straightforward to observe that the l-edges that are incoming to some vertex v in Γ'_n, are incident to $\mathtt{D}(v)$, and properly intersects the first half of the S'_v. Let ℓ be the nearest vertical grid line to the right of S'_v, and remove the parts of these l-edges that lie in between $\mathtt{D}(v)$ and ℓ (except for the l-edge incident to $\mathtt{C}(v)$). Since all these l-edges lie below S_v, the points where these l-edges are incident to ℓ can see all the grid points on the rightmost boundary of A_v and on the right-half of the bottommost boundary of A_v. Consequently, we can route the l-edges to $\mathtt{C}(v)$ through these boundary grid points, which removes the edge overlaps on $\mathtt{D}(v)$. Figure 5(e) illustrates such a scenario. Symmetrically, we can remove the degeneracy of r-edges on $\mathtt{L}(v)$. Remark 2 and the property that the lines in S_v and S'_v do not contain any vertex except v ensure that the above modifications do not introduce any edge crossing. Let the resulting drawing be Γ'', which is a planar polyline drawing of G, e.g., see Figure 5(e).

Area: By Lemma 2, the area before the Expansion phase was $(W+2)\times(H+2)$. For each i, where $1 \leq i \leq W+2$, the Expansion phase increases the width of the drawing by $2\lfloor d(u_i)/2 \rfloor$, where u_i is the vertex with the largest degree on the ith column. Hence the total increase is at most $(\sum_{i=1}^{W+2} d(u_i)) - 3(n - W - 2) \leq (6n - 12) - 3(n - W - 2) = 3n + 3W - 6$. Similarly, the increase in

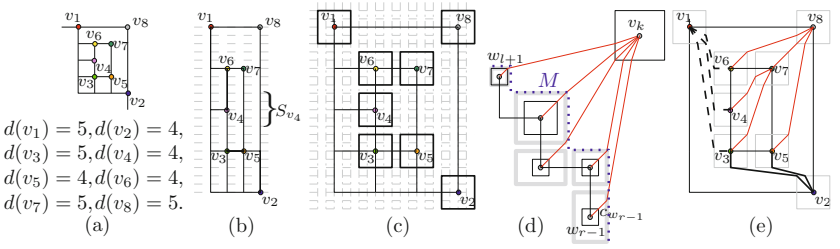

Fig. 5. Illustration for (a) Γ_n, (b) S_{v_k}, and (c) Γ'_n, where the grid A_v, for each vertex v, is shown in black squares. (d) Illustration for M. Note that A_ws are bounded by gray rectangles determined by S_w and S'_w. (e) Γ''.

height is at most $3n + 3H - 6$. Hence Γ'' is a drawing on an integer grid of size $(3n+4W-4)\times(3n+4H-4)$. Since $W+H \leq (4n-2\Delta_0-10)/3$ (see Remark 1), the area can be at most $(3n+4(2n-\Delta_0-5)/3)^2 = ((17n-4\Delta_0-20)/3)^2 \leq 32.12n^2$.

Bends per Edge: If (v, v') is an l-edge or r-edge in Γ_G, which starts at v and ends at v', then the edge has at most 2 bends: one before entering $A_{v'}$, and another at the boundary of $A_{v'}$. If (v, v') is an m-edge, then it contains one bend on M, and another bend on the boundary of $A_{v'}$. The l-and r-edges that connect a primary vertex to its parent, do not contain any bend. Since $\Delta_0 < n/2$ and $\texttt{leaf}(T_m) < n$, the drawing has at most $6n-\texttt{leaf}(T_l)-\texttt{leaf}(T_r)\leq 11n/2$ bends.

Angular Resolution: To compute the angular resolution, observe that the smallest possible angle θ at v is realized by a pair of consecutive integer grid points on the boundary of A_v where one of them is the corner of A_v, e.g., see Figure 6(a). Since A_v is a grid of size $(2\lfloor d(v)/2 \rfloor + 1) \times (2\lfloor d(v)/2 \rfloor + 1)$, the length of the line segment l connecting the center to any corner is $\sqrt{2}\lfloor d(v)/2 \rfloor$. Hence we have $\theta = \arctan\left(\frac{1/\sqrt{2}}{(\sqrt{2}\lfloor d(v)/2 \rfloor - 1/\sqrt{2}}\right) > 1/d(v)$, by the MacLaurin series expansion of arctan [12]. Observe that any edge e that intersects some grid A_v, where v does not correspond to any end vertex of e, must be an m-edge. We can avoid any such intersection by choosing for each vertex u, a rectangular grid of size $(2\lfloor d(u')/2 \rfloor + 1) \times (2\lfloor d(u'')/2 \rfloor + 1)$ (instead of A_u), where u' (respectively, u'') is the vertex with the largest degree over all the vertices on the horizontal (respectively, vertical) line through u. For example, see the gray rectangles in Figure 5(d). However, the angular resolution increases to $1/\Delta$.

Theorem 1. *Every n-vertex maximal planar graph admits a 2-bend polyline drawing Γ with angular resolution at least $1/d(v)$ for each vertex v, and area at most $(3n + 4W - 4) \times (3n + 4H - 4)$, where $W + H \leq (4n - 2\Delta_0 - 10)/3$. Within the same area, we can assign each vertex v in Γ a bounding box of size $(2\lfloor d(v)/2 \rfloor + 1) \times (2\lfloor d(v)/2 \rfloor + 1)$ that only intersect with the edges incident to v, but the angular resolution increases to $1/\Delta$.*

4 Trade-Offs between Angular Resolution and Area

In this section we show that one can significantly improve the area with an small expense of angular resolution. We consider the following two scenarios.

Angular Resolution is $\gamma/d(v)$, where $\gamma \in [0.8, 1]$: Observe that the bottom-left quadrants of A_v (with respect to the center $\texttt{C}(v)$) has at most $2\lfloor d(v)/2 \rfloor - 1 \geq d_m(v)$ boundary points, which are sufficient to route the m-edges, and sometimes necessary. However, the boundary points that are available to route the l-edges (similarly, r-edges) are significantly more than necessary, e.g., the number of boundary points to route the l-edges is $3\lfloor d(v)/2 \rfloor - 2$ (lying on the bottom-right quadrants and on the right-boundary of A_v). Hence assigning a grid of size $(\lfloor d(v)/2 \rfloor + 1 + \lceil d(v)/4 \rceil) \times (\lfloor d(v)/2 \rfloor + 1 + \lceil d(v)/4 \rceil)$ to each vertex v would be sufficient for routing the edges.

Observe that for each vertex v, the increase in width is at most $(\lfloor d(v)/2 \rfloor + \lceil d(v)/4 \rceil) \leq (3d(v)/4 + 1)$. Since one column may contain multiple vertices, and

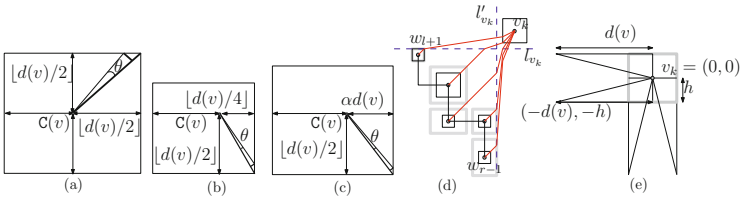

Fig. 6. Illustration for angular resolution

the degree of each vertex is at least three, we are overcounting the increase for $(n-W-2)$ vertices. The amount of over computation for each such vertex v' is at least $\lfloor 3d(v')/4 \rfloor + 1 \geq 3$. Consequently, the total increase in the width in the Expansion phase is now bounded by $(\sum_{i=1}^{W+2}(3d(v_i)/4+1)) - 3(n-W-2) \leq 3n/2+4W-1$. Similarly, the increase in height is at most $3n/2+4H+1$. Since $W+H \leq (4n-2\Delta_0-10)/3$, the area can be at most $(3n/2+5(2n-\Delta_0-5)/3+5)^2 \leq 23.37n^2$. The number of bends remains the same, but the minimum angle θ is now at least $0.8/d(v)$, which is now determined by two consecutive points along the bottom-right corner, as shown in Figure 6(b).

We can parametrize the grid size with a parameter α, i.e., consider the grid assigned to v as $(\lfloor d(v)/2 \rfloor + 1 + \alpha d(v)) \times (\lfloor d(v)/2 \rfloor + 1 + \alpha d(v))$, where $\alpha \geq 1/4$. Then the increase in width is at most $(\sum_{i=1}^{W+2}((\alpha+1/2)d(v_i)+1)) - 3(n-W-2) \leq (6(\alpha+1/2)n - 3n + 4W + 8) \leq (6\alpha n + 4W + 8)$. Similarly, the increase in height is at most $(6\alpha n + 4H + 8)$, respectively. Hence the area is at most $(6\alpha n + 4(W+H)/2 + 10)^2 \leq (6\alpha n + 8n/3 + 10)^2 \approx (6\alpha + 8/3)^2 n^2$. The angular resolution is at least $\frac{\alpha/\sqrt{\alpha^2+1/4}}{d(v)\sqrt{\alpha^2+1/4}} > \frac{\alpha}{d(v)(\alpha^2+1/4)}$, as illustrated in Figure 6(c).

Theorem 2. *Every n-vertex maximal planar graph admits a 2-bend polyline drawing with angular resolution $\frac{\alpha}{d(v)(\alpha^2+1/4)}$ for each vertex v, and area $(6\alpha n + 4W + 10) \times (6\alpha n + 4H + 10)$. Here $\alpha \in [1/4, 1/2]$, and $W+H \leq (4n-2\Delta_0-10)/3$.*

Angular Resolution is $\gamma/d(v)$, where $\gamma \in [0.3, 0.5]$: Recall that the new grid lines in the Expansion phase are inserted such that each vertex v has $h = \beta_v d(v)$ free grid lines, where $\beta_v \geq 1/d(v)$, in each of the four sides (above, below, left, right) around v, i.e., $\mathtt{C}(v)$ is at the center of a free integer grid A_v of size $h \times h$. As in the Expansion phase, let S_v be the ordered set of horizontal free grid lines along with the horizontal line through v, and let S'_v be the ordered set of vertical free grid lines along with the vertical line through v. We now show that these free grids are sufficient for routing the l-, r- and m-edges.

Routing m-edges: Let l_{v_k} and l'_{v_k} be the grid lines that are immediately below and to the left of S_{v_k} and S'_{v_k}, respectively. For each $w \in \{w_{l+1}, \ldots, w_{r-1}\}$, we now extend a line segment with slope $+1$ from $\mathtt{C}(w)$ until we hit either l_{v_k} or l'_{v_k}. Let $B = \{b(w_{l+1}), \ldots, b(w_{r-1})\}$ be the set of points on l_{v_k} and l'_{v_k} reached by these extensions. We now extend these extensions further to reach $\mathtt{C}(v_k)$, as follows:

- If the number of points of B that lie on l_{v_k} is z, where $z \leq h$, then we route the extensions of l_{v_k} through z consecutive grid points lying on the left side of A_{v_k} immediately below $\mathtt{L}(v_k)$. We then route the extensions on l'_{v_k} through the next consecutive grid points along the same vertical line. Since there are at most $d_m(v_k)$ m-edges, we need at most $d(v)$ consecutive grid points below $\mathtt{L}(v_k)$. Figure 6(d) illustrates such a scenario, where $h = 2$.
- If the number of points of B that lie on l'_{v_k} is at most z', where $z' \leq h$, then the drawing is symmetric to the case when $z < h$.
- Otherwise, both l_{v_k} and l'_{v_k} contains more than h extensions. In this case $\min\{z, z'\} > h$, and hence $\max\{z, z'\} \leq d_m(v) - h$. We first extend the extensions on l_{v_k} to the grid points that lie consecutively to the left of A_v (on the first line of S_{v_k}). We then extend the extensions on l'_{v_k} to the grid points that lie consecutively below of A_v (on the last line of S'_{v_k}). Finally, we connect all these new extensions directly to $\mathtt{C}(v_k)$. Note that the maximum horizontal (respectively, vertical) distance between $\mathtt{C}(v)$ and a bend point on l_{v_k} (respectively, l'_{v_k}) is at most $(d_m(v) - h) + h \leq d(v)$.

Routing *l*-edges: Let $u_1, u_2, \ldots, u_q$ be the vertices in top-to-bottom order that are incident to $\mathtt{D}(v_k)$ by incoming l-edges. Let ℓ be the nearest vertical grid line to the right of S'_v, and remove the parts of these l-edges that lie in between $\mathtt{D}(v_k)$ and ℓ (except for the l-edge incident to $\mathtt{C}(v_k)$). We then connect these extensions to the q consecutive grid points on the first line of S'_{v_k} that lie immediately below the top-right corner of A_v. Finally, we connect all these new extensions directly to $\mathtt{C}(v_k)$.

Routing *r*-edges: This scenario is symmetric for routing l-edges.

Angular Resolution and Area: In all the cases, the smallest angle θ at any vertex v is equal to the angle determined by the points $(-d(v), -h)$ and $(-d(v)+1, -h)$ at $\mathtt{C}(v) = (0, 0)$, as illustrated in Figure 6(e). Here the angular resolution is at least $\frac{\beta_v}{d(v)(1+\beta_v^2)}$, where $1/d(v) \leq \beta_v \leq 1$, and the area is $(6\beta + 2/3)^2 n^2$. We omit the details due to space constraints.

Theorem 3. *Every n-vertex maximal planar graph admits a 2-bend polyline drawing with angular resolution $\frac{\beta}{d(v)(1+\beta^2)}$ for each vertex v, and area $(6n\beta + W + 2) \times (6n\beta + H + 2)$. Here $\beta \in [1/3, 1]$, and $W + H \leq (4n - 2\Delta_0 - 10)/3$.*

5 Conclusion

In this paper we have given the first smooth trade-off between the area and angular resolution for 2-bend polyline drawings of any given planar graph. Our algorithm dominates all the previous 2-bend polyline drawing algorithms except Duncan and Kobourov's algorithm [10], which uses 1-bend per edge and dominates our algorithm when the angular resolution is in the interval $[0.38/d(v), 0.5/d(v)]$. Similar to the previously known polyline drawing algorithms, one can implement our algorithm using standard techniques [6] such that the drawings are computed in linear time.

A natural open question is whether Duncan and Kobourov's algorithm could be modified (allowing 2-bends per edge) to achieve a better trade-off. Finding tight lower bounds would also be very interesting. Finally, we hope that the results in this paper will encourage the study of smooth trade-offs among different aesthetic criteria for other styles of drawing graphs.

References

1. Biedl, T.C., Kaufmann, M.: Area-efficient static and incremental graph drawings. In: Burkard, R.E., Woeginger, G.J. (eds.) ESA 1997. LNCS, vol. 1284, pp. 37–52. Springer, Heidelberg (1997)
2. Bonichon, N., Felsner, S., Mosbah, M.: Convex drawings of 3-connected plane graphs. Algorithmica 47(4), 399–420 (2007)
3. Bonichon, N., Le Saëc, B., Mosbah, M.: Wagner's theorem on realizers. In: Widmayer, P., Triguero, F., Morales, R., Hennessy, M., Eidenbenz, S., Conejo, R. (eds.) ICALP 2002. LNCS, vol. 2380, pp. 1043–1053. Springer, Heidelberg (2002)
4. Brandenburg, F.J.: Drawing planar graphs on $\frac{8}{9}n^2$ area. Electronic Notes in Discrete Mathematics 31, 37–40 (2008)
5. Cheng, C.C., Duncan, C.A., Goodrich, M.T., Kobourov, S.G.: Drawing planar graphs with circular arcs. Discrete & Computational Geometry 25(3), 405–418 (2001)
6. Chrobak, M., Payne, T.: A linear-time algorithm for drawing planar graphs. Information Processing Letters 54, 241–246 (1995)
7. CircuitLogix, `https://www.circuitlogix.com/` (accessed June 03, 2014)
8. ConceptDraw: `http://www.conceptdraw.com/` (accessed June 03, 2014)
9. De Fraysseix, H., Pach, J., Pollack, R.: How to draw a planar graph on a grid. Combinatorica 10(1), 41–51 (1990)
10. Duncan, C.A., Kobourov, S.G.: Polar coordinate drawing of planar graphs with good angular resolution. Journal of Graph Algorithms and Applications 7(4), 311–333 (2003)
11. Durocher, S., Mondal, D.: On balanced +-contact representations. In: Proceedings of GD. In: Wismath, S., Wolff, A. (eds.) GD 2013. LNCS, vol. 8242, pp. 143–154. Springer, Heidelberg (2013)
12. Goodrich, M.T., Wagner, C.G.: A framework for drawing planar graphs with curves and polylines. Journal of Algorithms 37(2), 399–421 (2000)
13. Gutwenger, C., Mutzel, P.: Planar polyline drawings with good angular resolution. In: Whitesides, S.H. (ed.) GD 1998. LNCS, vol. 1547, pp. 167–182. Springer, Heidelberg (1999)
14. Hong, S.H., Mader, M.: Generalizing the shift method for rectangular shaped vertices with visibility constraints. In: Tollis, I.G., Patrignani, M. (eds.) GD 2008. LNCS, vol. 5417, pp. 278–283. Springer, Heidelberg (2009)
15. Kant, G.: Drawing planar graphs using the canonical ordering. Algorithmica 16(1), 4–32 (1996)
16. Kurowski, M.: Planar straight-line drawing in an $o(n) \times o(n)$ grid with angular resolution $\omega(1/n)$. In: Vojtáš, P., Bieliková, M., Charron-Bost, B., Sýkora, O. (eds.) SOFSEM 2005. LNCS, vol. 3381, pp. 250–258. Springer, Heidelberg (2005)
17. Schnyder, W.: Embedding planar graphs on the grid. In: Proceedings of ACM-SIAM SODA, January 22-24, pp. 138–148. ACM (1990)

18. SmartDraw Software, LLC, http://www.smartdraw.com/ (accessed June 03, 2014)
19. Tamassia, R., Di Battista, G., Batini, C.: Automatic graph drawing and readability of diagrams. IEEE Transactions on Systems, Man and Cybernetics 18(1), 61–79 (1988)
20. Zhang, H.: Planar polyline drawings via graph transformations. Algorithmica 57(2), 381–397 (2010)
21. Zhang, H., He, X.: Canonical ordering trees and their applications in graph drawing. Discrete & Computational Geometry 33(2), 321–344 (2005)