Small poly-line drawings of series-parallel graphs *

Therese Biedl †

Abstract

In this paper, we study small planar drawings of planar graphs. For arbitrary planar graphs, $\theta(n^2)$ is the established upper and lower bound on the area. It is a long-standing open problem for what graphs smaller area can be achieved. We show here that series-parallel can be drawn in $O(n^{3/2})$ area and outer-planar graphs can be drawn in $O(n \log n)$ area, but partial 3-trees and 2-outer-planar graphs require $\Omega(n^2)$ area. Our drawings are visibility representations, which can be converted to poly-line drawings of asymptotically the same area.

1 Introduction

A planar graph is a graph that can be drawn without crossing. Fáry, Stein and Wagner [Fár48, Ste51, Wag36] proved independently that every planar graph has a drawing such that all edges are drawn as straight-line segments. De Fraysseix, Pach and Pollack [FPP90], and independently Schnyder [Sch90] established that in fact $O(n^2)$ area suffices for a straight-line drawing of an n-vertex planar graph, with vertices placed at grid points. This is asymptotically optimal, since there are planar graphs that need $\Omega(n^2)$ area [FPP88].

A number of other graph drawing models (e.g., poly-line drawings, orthogonal drawings, visibility representations) exist for planar graphs. See Section 2.1 for precise definitions. In all these models, $O(n^2)$ area can be achieved for planar graphs, see for example [Kan96, FKK97, Wis85]. On the other hand, $\Omega(n^2)$ area is needed, even in these models, for the graph in [FPP88].

An important open question in graph drawing [BEG+03] is whether an area of $o(n^2)$ is possible, at least in a weaker drawing model such as poly-line drawings, for subclasses of planar graphs. The only subclass of planar graphs for which such drawings were known are trees, and outer-planar graphs under special restrictions (see below.) We show in this paper that in fact all series-parallel graphs can be drawn in $o(n^2)$ area, and also consider some other subclasses such as outer-planar graphs, planar partial k-trees and k-outer-planar graphs.

*Part of these results appeared in a preliminary version at *Graph Drawing '02* [Bie02]. The author would like to thank participants at the Algorithmic Problem Session at University of Waterloo, and at the Dagstuhl Seminar 05191 on Graph Drawing, for inspiring discussions.

†School of Computer Science, University of Waterloo, Waterloo, ON N2L 3G1, Canada, biedl@uwaterloo.ca. Research supported by NSERC. Part of the work was done while the author was on sabbatical leave at University of Passau.

1.1 Known results

Every **tree** has a straight-line drawing in $O(n \log n)$ area [Shi76], and in $O(n)$ area if the maximum degree is asymptotically smaller than n [GGT96]. See [DBETT98] for many other upper and lower bounds regarding drawings of trees graphs under special restrictions.

Not many drawing results are known for **outer-planar graphs**. It is quite easy (and appears to be folklore) to create straight-line drawings of area $O(nd)$, where d is the diameter of the dual tree of the graph. In another result, any n points in general position can be used for a straight-line drawing of any outer-planar graph [CU96]. However, $O(n^2)$ area is needed to create n grid points in general positions, so this result does not lead to area bounds smaller than the trivial $O(n^2)$. Since the appearance of a preliminary version of this paper [Bie02], Garg and Rusu [GR03] showed that any outer-planar graph with maximum degree Δ has a straight-line drawing of area $O(\Delta n^{1.48})$, so $o(n^2)$ area drawings exist for outer-planar graphs with small maximum degree.

Many drawing results are known for **series-parallel graphs**, see e.g. [BCDB+94, CBTT95, HEQL98]. However, the emphasis here was on displaying the series-parallel structure of the graph, and/or to use the structure to allow for additional constraints. All known algorithms bound the area by $O(n^2)$ area or worse.

No graph drawing results specifically tailored to k**-outer-planar graphs** (for $k \geq 2$), or **planar partial** k**-trees** appear to be known, at least not for 2-dimensional drawings.

While higher-dimensional drawings are not the focus of our paper, we would like to mention briefly that all graph classes considered in this paper can be drawn with linear area in 3D, because they are partial k-trees for constant k; see [DMW05], and also [FLW03] for some earlier results for outer-planar graphs.

We would also like to note that all these graphs have constant-size separators, and hence by a result Leiserson [Lei80] they have a two-dimensional orthogonal point-drawing in $O(n)$ area if the maximum degree is at most 4. However, the drawing need not be planar.

1.2 Our results

In this paper, we study two-dimensional planar poly-line drawings of some subclasses of planar graphs, and provide the following drawing results:

- As our main result, we show that every series-parallel graph has a visibility representation in $O(n^{3/2})$ area if we can choose the planar embedding.
- We can improve this under some circumstances. In particular, a series-parallel graph for which at most f graphs are combined in parallel has a visibility representation in $O(fn \log n)$ area; this is at most $O(\Delta n \log n)$ area.
- This easily implies that every outer-planar graph has a visibility representation in $O(n \log n)$ area.
- For outer-planar graphs, we can maintain the "standard" planar embedding with the same area bound: Every outer-planar graph has an orthogonal box-drawing with area $O(n \log n)$ in which all vertices are drawn on the outer-face.

We also provide the following lower bounds for drawings in various models:

- There are series-parallel graphs that require $\Omega(n^2)$ area in any poly-line drawing that respects the planar embedding.
- There are series-parallel graphs that required $\Omega(n \log n)$ area in any orthogonal box-drawing, even if we can choose the planar embedding.
- There are series-parallel graphs that required $\Omega(n \log n / \log \log n)$ area in any straight-line drawing, even if we can choose the planar embedding.
- There are outer-planar graphs that required $\Omega(n^2)$ area in any poly-line drawing such that the vertices are on the bounding box.
- There are partial 3-trees that require $\Omega(n^2)$ area in any poly-line drawing.
- There are 2-outer-planar graphs that require $\Omega(n^2)$ area in any poly-line drawing.

For algorithms, we restrict our attention to visibility representations and orthogonal drawings, because, as we briefly recall in Section 2, any such drawing can be converted to a poly-line drawing without increasing the area by more than a constant. Hence all our upper bounds (given in Section 3 and Section 4) also hold for poly-line drawings. On the other hand, our lower bounds in Section 5 are mostly for poly-line drawings; this also implies the same asymptotic lower bounds for all other drawing models.

2 Background

2.1 Graphs and graph classes

Let $G = (V, E)$ be a graph with $n = n(G) = |V|$ vertices and $m = m(G) = |E|$ edges. We assume that G is *simple*, i.e., it has no loops and multiple edges. Throughout this paper, we will assume that G is planar, i.e., that G can be drawn without crossing. Such a planar drawing can be characterized by the cyclic order of edges around each vertex. A planar drawing splits the plane into connected pieces; the unbounded piece is called the *outer-face*, all other pieces are called *interior faces.*

An *outer-planar graph* is a planar graph that can be drawn such that all vertices are incident to the outer-face. A *maximal outer-planar graph* is an outer-planar graph to which we cannot add an edge without destroying simplicity or outer-planarity. Such a graph consists of an n-cycle with chords and every interior face is a triangle. When drawing outer-planar graphs, we will generally assume that they are maximal outer-planar, because we can make outer-planar graphs maximal by adding edges, draw the resulting graph, and then delete added edges.

A *2-terminal series-parallel graph with terminals s, t* is a graph defined recursively as follows: (a) An edge (s, t) is a 2-terminal series-parallel graph. (b) If G_i, $i = 1, 2$ are 2-terminal series-parallel graphs with terminals s_i and t_i,, then in a *series composition* we identify t_1 with s_2 to obtain a 2-terminal series-parallel graph with terminals s_1 and t_2. (c) If G_i, $i = 1, \ldots, k$, are 2-terminal series-parallel graphs with terminals s_i and t_i, then in a *parallel composition* we identify $s_1, s_2, \ldots, s_k$ into one terminal s and $t_1, t_2, \ldots, t_k$ into one terminal t to obtain a 2-terminal series-parallel graph with terminals s and t. We assume for a parallel composition that k has been chosen as large as possible, i.e., none of the graphs G_i

used for the parallel composition is itself obtained via a parallel composition. The *fan-out* of a series-parallel graph is the maximum number of subgraphs k used in a parallel composition.

Given a 2-terminal series-parallel graph G, a *subgraph from the composition* is any of the subgraphs $G_1, \ldots, G_k$ that was used to create G, or recursively any subgraph from the composition of $G_1, \ldots, G_k$. Since we never consider any other subgraphs, we will say "subgraphs" instead of "subgraphs from the composition".

A *series-parallel graph*, or *SP-graph* for short, is a graph for which every biconnected component is a 2-terminal series-parallel graph. A *maximal series-parallel graph* is a series-parallel graph to which we cannot add any edge without destroying simplicity or series-parallelness. Similarly as for outer-planar graphs, whenever we speak of a series-parallel graph from now on, we will in fact assume that it is maximal series-parallel, since this makes no difference for asymptotic upper bounds on the area of graph drawings. One can easily show that a maximal series-parallel graph is a 2-terminal series-parallel graph, for which additionally in any parallel composition there exists an edge between the terminals, and in any series composition each subgraph is either an edge or obtained from a parallel composition.

2.2 Drawing models and their relationships

We consider the following drawing models (see also Figure 1):

- *Straight-line drawings*: Vertices are points, all edges are straight-line segments.
- *Poly-line drawings*: Vertices are points, all edges are contiguous sequences of straight-line segments.
- *Orthogonal point-drawings*: A poly-line drawing where all edge segments are horizontal or vertical. Such drawings exist only if the maximum degree is at most 4.[1]
- *Orthogonal box-drawings*: Vertices are boxes,[2] edges are contiguous sequences of horizontal or vertical segments.
- *Visibility representations*: Vertices are boxes, edges are horizontal or vertical segments.

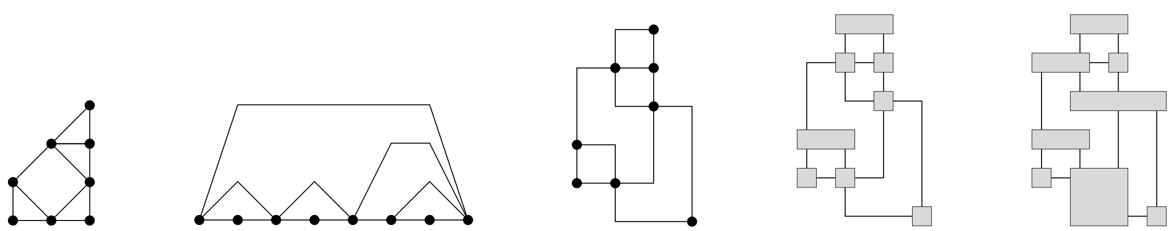

Figure 1: The same graph in a straight-line drawing, a poly-line drawing, an orthogonal point-drawing, an orthogonal box-drawing, and a visibility representation.

For a planar graph, such drawings should be planar, i.e., have no crossing. We also assume that all defining features have integral coordinates; in particular points of vertices

[1] We mention orthogonal point-drawings here for completeness' sake; they will not be studied in this paper, and from now on, whenever we say "orthogonal drawing" we mean "orthogonal box-drawing."

[2] In this paper, the term "box" always refers to an axis-parallel box.

and transition-points (*bends*) in the routes of edges have integral coordinates, and boxes of vertices have integral corner points. We allow boxes to be degenerate, i.e., to be line segments or points.

The *width* of a box is the number of vertical grid lines (*columns*) that are occupied by it. The *height* of a box is the number of horizontal grid lines (*rows*) that are occupied by it. A drawing whose minimum enclosing box has width w and height h is called a $w \times h$-drawing, and has *area* $w \cdot h$.

There are some relationships between these graph drawing models, which we will exploit to obtain the same area bounds for most models. Note that by definition every visibility representation is an orthogonal box-drawing, and every straight-line drawing is a poly-line drawing. Other relationships can be obtained by modifying the drawings, which we explain now.

2.2.1 Orthogonal box-drawings to poly-line drawings

From any orthogonal box-drawing, one can easily obtain a poly-line drawing with asymptotically the same area as follows: Add empty grid-lines until every segment of every edge has length at least 2. Now for every vertex v, create a point for v at an arbitrary grid point inside the box of v. For each incident edge e of v, re-route e to end at this point by placing a new bend (if needed) at the grid point next to where e used to attach to the box of v. (Such a grid point must exist, and is not contained in any other vertex, since e had length at least 2 by assumption.) See Figure 2 for an illustration.

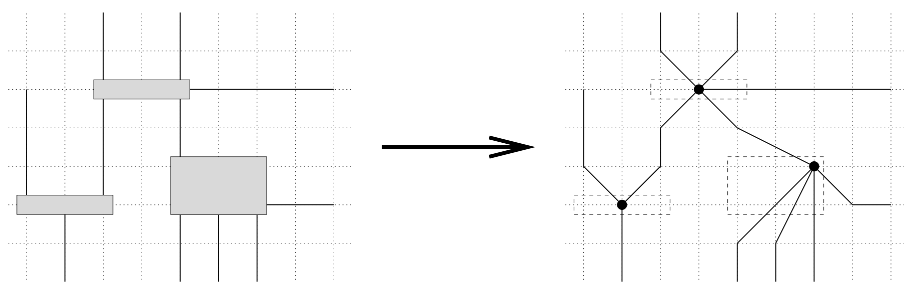

Figure 2: Converting to a point-drawing.

The initial step at most doubles the width and height of the drawing, since by adding a new row/column after every existing row/column, all edges are guaranteed to have length at least 2. Thus, if the orthogonal box-drawing has area A, the new drawing has area at most $4A$.

Lemma 2.1 *If G has an orthogonal box-drawing with area A, then it also has a poly-line drawing with area at most $4A$.*

2.2.2 Flat visibility representations to 1-directional visibility representations

Recall that a visibility representation is an orthogonal box-drawing where edges have no bends, i.e., they are either horizontal or vertical line segments. Visibility representations can be distinguished into different types. A *1-directional visibility representation* uses only vertical lines for edges. A *flat visibility representation* has height 1 for every box. Every 1-directional visibility representation can be made flat, since no vertex has incident horizontal edges. The other direction is also feasible by transforming the drawing.

Assume we are given a flat visibility representation. Replace every grid-line in it by two new grid-lines. If $v_1, \ldots, v_k$ are the vertices in row r from left to right, then place $v_1, \ldots, v_k$ alternatingly in the two rows r_1 and r_2 that replaced r. If a vertex intersected column c before, then let it now intersect both columns that replaced c.

Replace edges as follows. If e is routed vertically in column c, then place it in the right of the two columns that replaced c. If e is routed horizontally, then (since the drawing is flat) it must connect two vertices v_i and v_{i+1} that were placed consecutively in the same row r. Vertex v_i is now in (say) row r_1 and v_{i+1} in row r_2 that replaced r. Extend vertex v_i to the right until it overlaps the leftmost column of v_{i+1}, and route e vertically in this column. See Figure 3.

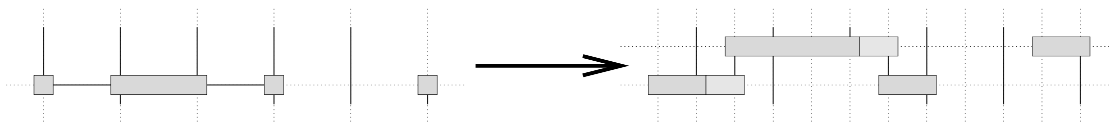

Figure 3: Converting a flat visibility representation to a 1-directional visibility representation.

Assume e is a horizontal edge between v_i and v_{i+1}. Note that no vertical edge has been placed in the columns between v_i and v_{i+1} (otherwise, this would cross edge e in a planar drawing), and there is no other vertex in this space (since v_i and v_{i+1} are consecutive.) Also, the leftmost column of v_{i+1} does not contain a vertical edge (since any such edge is placed in the right of the two columns that replaced it.) So this transformation maintains a visibility representation and converts all horizontal edges into vertical edges, hence gives a 1-directional visibility representation.

Lemma 2.2 *If G has a flat visibility representation with area A, then it also has a 1-directional visibility representation with area at most $4A$.*

Figure 4 summarizes all relationships between drawing models; solid arrows mean that drawings in one model are automatically drawings in the other model as well, whereas dashed arrows mean that drawings in one model can be transformed to drawings in the other model without asymptotic increase in the area.

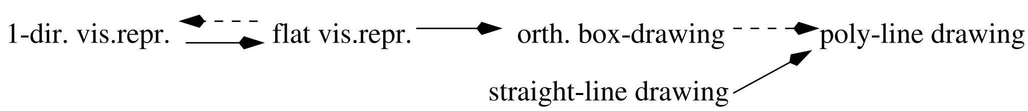

Figure 4: Relationships between drawing models.

From Figure 4, it should be clear why in our constructions we (mostly) consider flat visibility representations: this implies the same asymptotic bounds for all other drawing models except straight-line drawings. On the other hand, poly-line drawings are the best model for lower bounds.

3 Flat visibility representations of series-parallel graphs

In this section, we study how to create small visibility representations of series-parallel graphs, and first give an algorithm depending on some parameter L. If the series-parallel graph has fan-out at most L, i.e., all parallel compositions of the series-parallel graph have at most L subgraphs, then the resulting drawing has an area of $O(Ln \log n)$. In particular, this yields area $O(\Delta n \log n)$ for all series-parallel graphs, by choosing $L = \Delta$. If the fan-out exceeds L, then by choosing L suitably we get an area bound of $O(n^{3/2})$.

3.1 The invariant

Presume from now on that we are given a maximal SP-graph G, which we know to be a 2-terminal series-parallel graph. Our algorithm draws G and recursively all its subgraphs H (recall that this means subgraphs used for compositions to obtain G). To ease putting drawings together, we put constraints on the location of the terminals (see also Figure 5):

- Vertex s is placed in the upper right corner of the bounding box.
- Vertex t is placed in the lower right corner of the bounding box.

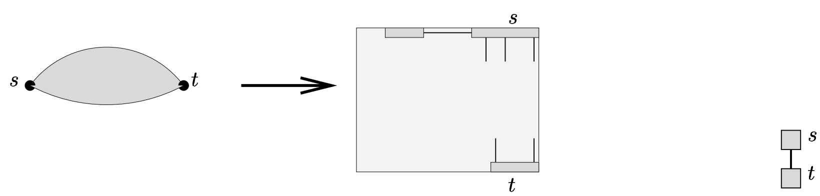

Figure 5: Illustration of the induction hypothesis, and the base case $n = 2$.

With our construction we will also develop a recursive formula for the height of the drawing of H, and denote it by $h(m)$, where H has m edges. More precisely, $h(m)$ is the maximum height of a drawing obtained with our algorithm over all maximal SP-graphs with m edges.[3] In the base case ($m = 1$), simply place s atop t; see Figure 5. The conditions of the induction hypothesis are clearly satisfied, and we have $h(m) = 2$ for $m = 1$.

3.2 Some operations

If our subgraph is more than an edge, then we will obtain a suitable drawing for it by splitting it into further subgraphs, drawing them, and merging their drawings together suitably. Before merging the drawings, we will sometimes have to modify them, and hence study in this sub-section how to modify drawings. Both operations given below can be applied not only to a flat visibility representation, but in fact to any *flat orthogonal drawing* (i.e., an orthogonal drawing for which all vertices have height at most 1.) Neither of them adds bends (the second one may remove bends), hence applying them to a visibility representation will again give a visibility representation.

[3]It is easy to show that maximal SP-graphs have $m = 2n - 3$ edges; we will use m instead of $2n - 3$ to simplify the computations.

The first operation is absolutely straightforward, but the ability to do this operation is one of the main reasons why creating orthogonal drawings is so much easier than straight-line drawings.

Lemma 3.1 *Let $\Gamma(H)$ be a flat orthogonal drawing of H of height $h \geq 2$ that satisfies the invariant. Then for any $h' \geq h$, there exists a flat orthogonal drawing $\Gamma'(H)$ of H of height h' that satisfies the invariant and has the same number of bends as $\Gamma(H)$.*

Proof: Simply insert $h' - h$ empty rows between the top row and the bottom row. All edge segments that intersect a new empty row are necessarily vertical, and hence simply can be extended. □

The second operation is a bit more involved. We say that in a drawing a vertex *spans the top (bottom) row* if its vertex box contains both the top (bottom) left point and the top (bottom) right point of the smallest enclosing box of the drawing.

Lemma 3.2 *Let $\Gamma(H)$ be a flat orthogonal drawing of H of height $h \geq 2$ that satisfies the invariant. Then there exists a flat orthogonal drawing $\Gamma'(H)$ of H of height $h+1$ that satisfies the invariant, has no more bends than $\Gamma(H)$, and vertex s spans the top row.*

Proof: Add a new row above $\Gamma(H)$, and move s into this row, spanning it all the way. The main question is how to re-route the edges incident to s. For all edges that attach at s vertically, we simply extend the edge to reach the new position of s.

Since s has height 1 and contains the top right corner, at most one edge attaches horizontally at s. If there is one, then let z be the other endpoint of the horizontal segment in the top row (z could be a vertex or a bend.) Since s spans the top row in $\Gamma'(H)$, s must be now above z, so we can re-route the edge by continuing vertically upward from z towards s. See also Figure 6. □

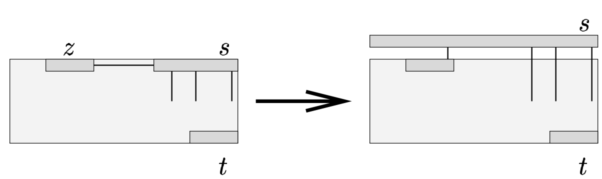

Figure 6: Moving s so that it spans the top row.

We call this approach *releasing terminal s*. With exactly the same approach at the bottom end, we can obtain the following results:

Lemma 3.3 *Let $\Gamma(H)$ be a flat orthogonal drawing of H of height $h \geq 2$ that satisfies the invariant. Then there exists a flat orthogonal drawing $\Gamma'(H)$ of H of height $h+1$ that satisfies the invariant, has no more bends than $\Gamma(H)$, and vertex t spans the bottom row.*

Lemma 3.4 *Let $\Gamma(H)$ be a flat orthogonal drawing of H of height $h \geq 2$ that satisfies the invariant. Then there exists a flat orthogonal drawing $\Gamma'(H)$ of H of height $h+2$ that satisfies the invariant, has no more bends than $\Gamma(H)$, and vertices s and t span the top and bottom rows.*

3.3 Subgraphs from parallel compositions

Now we are ready to work on recursive cases. Assume first that H is a subgraph of G which is obtained in a parallel composition from subgraphs $H_1, \dots, H_k$, $k \geq 2$. As before, we assume that k is as big as possible, i.e., each H_i is an edge or is obtained from a series composition. Let m_i be the number of edges of H_i, and we assume that the naming is such that $m_1 \geq m_2 \geq \dots \geq m_k$.

Recursively obtain drawings of $H_1, \dots, H_k$; the drawing of H_i has height at most $h(m_i)$. Since they are being combined in parallel, each of $H_1, \dots, H_k$ has the same set of terminals, say s and t. For $i = 2, \dots, k$, release both s and t at H_i, thereby adding two units of height to the drawing of this subgraph. Let $h \leq \max\{h(m_1), h(m_2) + 2, \dots, h(m_k) + 2\}$ be the maximum height among the drawings of the subgraphs. Apply Lemma 3.1 to increase the heights of all drawings to exactly h.

Now place the drawing of H_1 leftmost, and the drawings of $H_2, \dots, H_k$ to the right of it. Note that in all drawings (which are all of height h) vertex s is in the top row, and with the exception of H_1, it spans the top row. Similarly vertex t is in the bottom row, and with the exception of H_1, spans the bottom row. Hence we can combine all the drawings of s into one long box of height 1, and similarly combine all drawings of t. This gives a drawing of H that satisfies the invariant, see Figure 7. The height of this drawing is

$$h(m) \leq \max\{h(m_1), h(m_2) + 2, \dots, h(m_k) + 2\} = \max\{h(m_1), h(m_2) + 2\} \qquad (1)$$

since $m_2 \geq m_3 \geq \dots \geq m_k$ and the height-function is monotone.

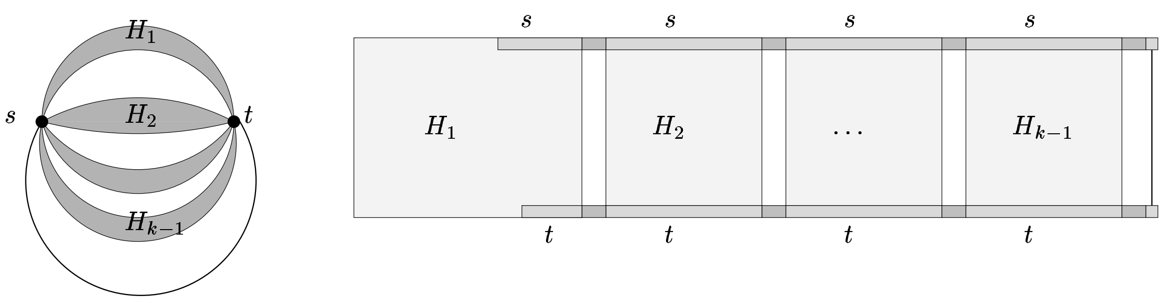

Figure 7: Combining subgraphs in parallel.

3.4 Subgraphs from series compositions

Now we turn to the more difficult case when H is obtained from a series composition, say of graphs H_a and H_b. Assume s and t are the terminals of H and x is the common terminal of H_a and H_b. Since we study maximal SP-graphs, each of H_a and H_b is either an edge or obtained from a parallel composition. We distinguish cases.

Case (S1): One subgraph is an edge. Assume first that H_b is an edge (x, t). In this case we draw H_a recursively, extend the drawing of terminal s to the right, place t in the bottom row, and connect edge (x, t) horizontally. See Figure 8. The case that H_a is an edge is symmetric.

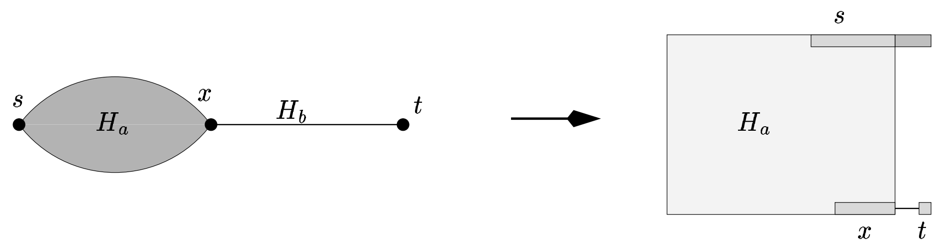

Figure 8: A series composition when one subgraph is an edge.

Case (S2): Both subgraphs are obtained from parallel compositions. So assume from now on that both H_a and H_b contain at least two edges, and hence must be obtained from parallel compositions. Furthermore assume that H_b has no more vertices than H_a (the other case is symmetric.) We break H_b down further, say it was obtained from a parallel composition of subgraphs $H_1, \ldots, H_k$. We assume that subgraphs $H_1, \ldots, H_k$ are sorted by decreasing number of edges, so H_1 has the most edges, and H_k is an edge (x, t) (which exists since the SP-graph is maximal.)

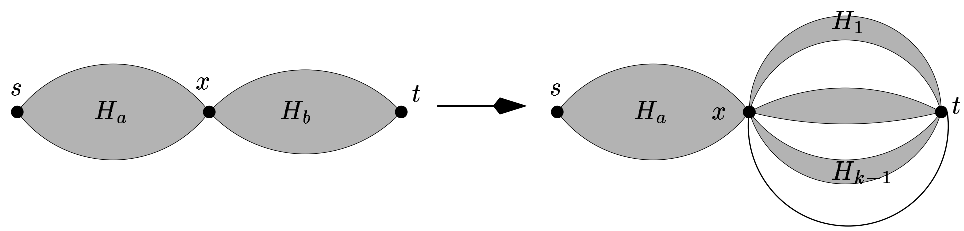

Figure 9: Breaking down subgraph H_b.

We again break down some (but not all) subgraphs farther. Let L be an integer; we will discuss later how to choose L. We break down H_i for all $i < \min\{L, k\}$; we know that H_i is not an edge since we have no multiple edges and H_k is an edge. Since k was as big as possible, H_i is obtained in a series composition of two subgraphs H_i^a and H_i^b with a common terminal y_i. The naming is such that H_i^a has terminals x and y_i, and H_i^b has terminals y_i and t. See also Figure 10. In what follows, for any strings α and β we use m_α^β to denote the number of edges of H_α^β.

Recursively draw each of the subgraphs H_a, H_i^a, H_i^b for $i = 1, \ldots, \min\{k, L\} - 1$,and H_i for $i = L, \ldots, k-1$. Before we can combine these drawings, we need to release some terminals again (recall Lemma 3.2). We proceed as follows:

- The drawing of H_a is unchanged and has height $h(m_a)$.
- For $i = 1, \ldots, \min\{L, k\} - 1$, release terminal x in the drawing of H_i^a, and terminal t in the drawing of H_i^b. The drawings hence have height at most $h(m_i^a) + 1$ and $h(m_i^b) + 1$.
- For $i = L, \ldots, k-1$, release both terminals in the drawing of H_i. The drawing hence has height at most $h(m_i) + 2$.

To explain how we put these drawings together, we distinguish two sub-cases:

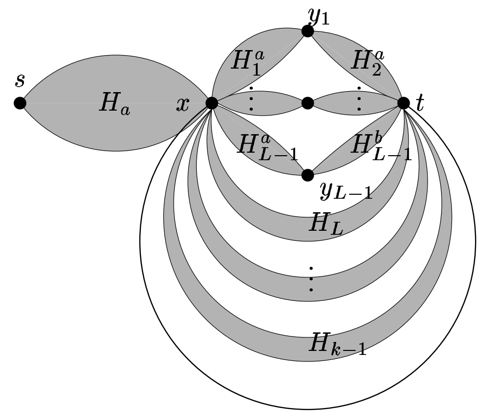

Figure 10: Breaking down some subgraphs further.

Case (S2a): Assume first that $k \leq L$, and consider Figure 11. We place H_a on the left, followed by $H_1^a, H_2^a, \ldots, H_{k-1}^a$. All these graphs share terminal x, which is placed in the bottom row. Note that x spans the drawings of $H_1^a, \ldots, H_{k-1}^a$ by construction, and hence all drawings of x can be combined into one. Now for $i = 1, \ldots, k-1$, flip the drawing of H_i^b horizontally such that terminal t spans the bottom row and terminal y_i occupies the top left corner. We place these flipped drawings in reversed order, i.e., place the drawings as $H_{k-1}^b, H_{k-2}^b, \ldots, H_1^b$ in such a way that their common terminal t is in the bottom row. Since x and t are in the same row, we can draw the edge between them horizontally.

Increase heights (refer to Lemma 3.1) such that H_i^a and H_i^b have the same height; this then places y_i in the same row in both occurrences in H_i^a and H_i^b. For $i < k-1$, increase heights further in both H_i^a and H_i^b such that the row of y_i is at least one row above the top row of the drawing of H_{i+1}^a and H_{i+1}^b; then the two drawings of y_i can be combined into one.

Finally increase the height of H_a such that it is at least one row taller than the drawing of H_1^a and H_1^b. Then we can extend terminal s of H_a to the right until it occupies the top right corner of the whole drawing. We hence obtain a drawing of H that satisfies the invariant.

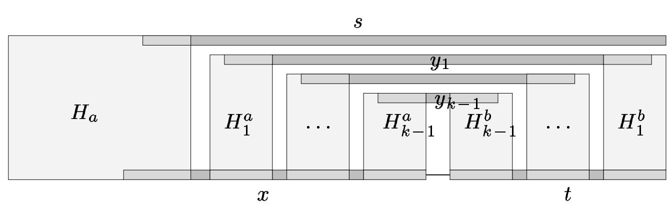

Figure 11: Combining the subgraphs for a series composition. The case $k \leq L$.

Let h_i be the height of the drawing of H_i^a and H_i^b together in the final drawing. Then $h_{k-1} \leq \max\{h(m_{k-1}^a)+1, h(m_{k-1}^b)+1\} \leq h(m_{k-1})+1$, where the last inequality holds since $m_{k-1}^a \leq m_{k-1}$ and $m_{k-1}^b \leq m_{k-1}$ and the height-function is monotone. Also, for $i < k-1$ heights have been increased such that h_i exceeds h_{i+1}, so $h_i \leq \max\{h(m_i)+1, h_{i+1}+1\}$. Using induction, one therefore obtains

$$h_1 \leq \max\{h(m_1)+1, h(m_2)+2, \ldots, h(m_{k-1})+k-1\}$$

The total height is at most $\max\{h(m_a), h_1 + 1\}$, so we have

$$h(m) \leq \max\{h(m_a), h(m_1) + 2, h(m_2) + 3, \ldots, h(m_{k-1}) + k\} \tag{2}$$

Case (S2b): Now we study the case $k > L$, which is similar, but we treat the graphs $H_L, \ldots, H_{k-1}$ differently.[4] Place $H_a, H_1^a, \ldots, H_{L-1}^a, H_{L-1}^b, \ldots, H_1^b$ exactly as before.

Increase heights until $H_L, \ldots, H_{k-1}$ all have the same height h_d, and place them *below* the rectangle of x; this uses $h_d - 1$ additional rows. We may have to add some columns to x if it is not wide enough for the subgraphs. Now terminal t is placed $h_d - 1$ rows below the row of terminal x. To make the two occurrences of t match up, we extend the drawings of $H_{L-1}^b, \ldots, H_1^b$ downwards by $h_d - 1$ lines and draw edge (x, t) vertically. See Figure 12.

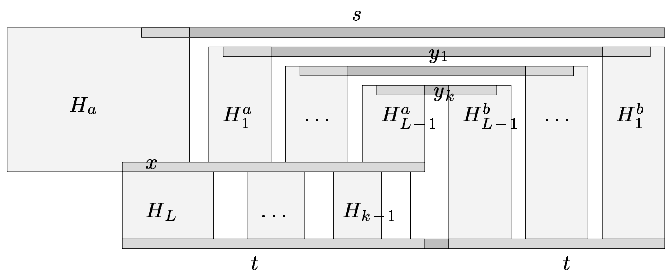

Figure 12: Combining the subgraphs for a series composition. The case $k \geq L$.

To obtain a formula for the resulting height, we hence need to add $h_d - 1$ to the formula of Inequality 2 (after replacing k by L in it.) Since h_d is the maximum height among $H_L, \ldots, H_{k-1}$, and $m_L \geq \ldots \geq m_k$, we have $h_d \leq h(m_L) + 2$ (recall that both terminals were released for $H_L, \ldots, H_{k-1}$), and therefore

$$h(m) \leq \max\{h(m_a), h(m_1) + 2, h(m_2) + 3, \ldots, h(m_{L-1}) + L\} + h(m_L) + 1 \tag{3}$$

3.5 Analysis

Now we show that the above algorithm indeed yields a small area, by evaluating Inequalities 1, 2 and 3 for small upper bounds on $h(m)$. We need a few auxiliary claims:

Claim 1 *For a parallel composition, $m_i \leq m/2$ for $i \geq 2$.*

Proof: We have $m_1 \geq m_i$, and hence $m \geq m_1 + m_i \geq 2m_i$ as desired. □

Claim 2 *In case (S2), we have $m_i \leq m/2$ for $i \geq 1$ and $m_i \leq m_a/2$ for $i \geq 2$.*

Proof: We have $m_a \geq m_b = m_1 + \ldots + m_k$, so $m \geq m_a + m_i \geq 2m_i$. Also, $m_1 \geq m_i$, so $m_a \geq m_1 + m_i \geq 2m_i$ for $i \geq 2$. □

Now we are ready for the main proof:

[4] In the way of motivation: Since the H_i's are sorted by decreasing size, these graphs are very small. If we placed them all as in case (S2a), then each would increase the height by 1, which is too much height increase for a small number of edges.

Lemma 3.5 *For a suitable choice of L, we have $h(m) \le 12\sqrt{m}$.*

Proof: In the base case, $m = 1$ and $h(m) = 2$ and the lemma holds. For a parallel composition, Claim 1 gives $m_i \le m/2$ for $i \ge 2$, and hence by Inequality 1

$$\begin{aligned} h(m) &\le \max\{h(m_1), h(m_2)+2, \ldots, h(m_k)+2\} \\ &\le \max\{h(m), h(m/2)+2\} \le \max\{12\sqrt{m}, 12\sqrt{m/2}+2\} \le 12\sqrt{m}. \end{aligned}$$

In case (S1), we have $h(m) = h(m_a) \le 12\sqrt{m_a} \le 12\sqrt{m}$. In case (S2), we assumed $m_a \ge m_b$. Also, $m_b \ge 3$ (because H_1^a and H_1^b have each an edge, and (x,t) exists), and hence $m \ge 6$. We choose $L = 3\sqrt{m_a} + 1$.[5] Now for case (S2a), we have by Inequality 2

$$\begin{aligned} h(m) &\le \max\{h(m_a), h(m_1)+2, h(m_2)+3, \ldots, h(m_{k-1})+k\} \\ &\le \max\{h(m_a), h(m/2)+L\} \text{ since } m_i \le m/2 \text{ and } k \le L \\ &\le \max\{12\sqrt{m_a}, 12\sqrt{m/2} + 3\sqrt{m_a} + \frac{1}{\sqrt{6}}\sqrt{m}\} \text{ since } L = 3\sqrt{m_a}+1 \text{ and } m \ge 6 \\ &\le \max\{12, (\frac{12}{\sqrt{2}} + 3 + \frac{1}{\sqrt{6}})\}\sqrt{m} \text{ since } m_a \le m \\ &\le 12\sqrt{m} \end{aligned}$$

Finally we consider case (S2b), which is by far the most difficult. Recall that the height in this case is by Inequality 3

$$\begin{aligned} h(m) &\le \max\{h(m_a), h(m_1)+2, h(m_2)+3, \ldots, h(m_{L-1})+L\} + h(m_L) + 1 \\ &\le \max\{h(m_a), h(m_1), h(m_2)+1, \ldots, h(m_{L-1})+L-2\} + h(m_L) + 3 \\ &\le \max\{h(m_a), h(m_a/2)+L-2\} + h(m_L) + 3 \text{ by } m_1 \le m_a \text{ and } m_i \le m_a/2 \text{ for } i \ge 2 \\ &\le 12\sqrt{m_a} + 12\sqrt{m_L} + 3, \end{aligned}$$

where the last inequality holds because $h(m_a/2) + L - 2 \le 12\sqrt{\frac{m_a}{2}} + 3\sqrt{m_a} - 1 \le 12\sqrt{m_a}$.

We now need to show that $12\sqrt{m_a} + 12\sqrt{m_L} + 3 \le 12\sqrt{m}$. At first this seems unintuitive, because m_a can be arbitrarily close to m. But since we are in case (S2b), there are $L \approx \sqrt{m_a}$ subgraphs of H_b, each of which has at least 2 edges. This leaves just enough room for the other terms. More precisely, we have

$$\begin{aligned} (\sqrt{m_a} + \sqrt{m_L} + \frac{1}{4})^2 &= m_a + m_L + \frac{1}{16} + 2\sqrt{m_a}\sqrt{m_L} + \frac{1}{2}\sqrt{m_a} + \frac{1}{2}\sqrt{m_L} \\ &\le m_a + m_L + \frac{1}{16}\sqrt{m_a}m_L + 2\sqrt{m_a}m_L + \frac{1}{2}\sqrt{m_a}m_L + \frac{1}{3}\sqrt{m_a}m_L \text{ by } \sqrt{m_a} \ge \sqrt{3} \ge \tfrac{3}{2} \\ &\le m_a + m_L + 3\sqrt{m_a}m_L = m_a + m_L + (L-1)m_L \text{ by } L = 3\sqrt{m_a}+1 \\ &\le m_a + m_L + m_1 + m_2 + \ldots + m_{L-1} \text{ by } m_i \ge m_L \text{ for } i < L \end{aligned}$$

which is at most m, so $12(\sqrt{m_a} + \sqrt{m_L} + \frac{1}{4}) \le 12\sqrt{m}$ as desired. □

[5] Many thanks to Jason Schattman for helping find small constants that work, using MAPLE.

Theorem 1 *Any series-parallel graph has a flat visibility representation with area $O(n^{3/2})$.*[6]

Proof: By the previous lemma, the height is $O(\sqrt{m}) = O(\sqrt{n})$ by $m = 2n - 3$. To analyze the width, notice that at the most we use one column for each edge. (Each vertex obtains at least one incident vertical edge in the base case, and hence does not contribute additional width.) Hence the width is at most $m \leq 2n - 3$, and the total area is $O(n^{3/2})$. □

We get easier proofs and better bounds if case (S2b) does not happen, i.e., if every parallel composition uses at most L subgraphs, i.e., if the series-parallel graph has small fan-out.

Lemma 3.6 *For a series-parallel graph with fan-out f, we have $h(m) \leq 2 + f \log m$.*

Proof: We proceed by induction on the number of edges. In the base case $h(1) = 2 \leq 2 + f \log m$. In case of a parallel composition, by Inequality 1 we have height

$$\begin{aligned} h(m) &\leq \max\{h(m_1), h(m_2) + 2\} \leq \max\{h(m_1), h(m/2) + 2\} \\ &\leq \max\{2 + f \log m_1, 2 + f \log(m/2) + 2\} \leq 2 + f \log m \end{aligned}$$

since $f \geq 2$. For case (S1), the height is $h(m) = h(m_a) \leq 2 + f \log m_a \leq 2 + f \log m$. In case (S2), we choose $L = f$, and hence always have $k \leq L$ and are in case (S2a). Here, the height is by Inequality 2

$$\begin{aligned} h(m) &\leq \max\{h(m_a), h(m_1) + 2, h(m_2) + 3, \ldots, h(m_{k-1}) + k\} \\ &\leq \max\{h(m_a), h(m/2) + f\} \text{ since } m_i \leq m/2 \text{ and } k \leq f \\ &\leq \max\{2 + f \log m_a, 2 + f \log(m/2) + f\} \leq 2 + f \log m. \end{aligned}$$

□

Arguing as in Theorem 1 about the width, we hence obtain:

Theorem 2 *Any series-parallel graph with fan-out f has a flat visibility representation of area $O(fn \log n)$.*

Note in particular that a series-parallel graph with maximum degree Δ has fan-out at most Δ, so any series-parallel graph has a flat visibility representation of area $O(\Delta n \log n)$.

4 Drawing outer-planar graphs

Every outer-planar graph is a series-parallel graph, and hence by the results of the previous subsection, has a visibility representation of area $O(n^{3/2})$ or $O(fn \log n)$. We now show that outer-planar graphs have fan-out $f \leq 2$, and therefore they have visibility representations of area $O(n \log n)$. Secondly, we show that by modifying the construction, we can obtain orthogonal box-drawings of area $O(n \log n)$ that have all vertices of the graph drawn on the outer-face.

[6] The emphasis in this paper was on achieving an area of $o(n^2)$, not on obtaining particularly small constant. The constant hidden behind O is 48 in our case, but can likely be improved with a more careful analysis, and in particular by considering larger subgraphs for base cases.

4.1 Visibility representation

We first prove the claim on the fan-out of an outer-planar graph by re-proving the well-known fact that any outer-planar graph is series-parallel.

Lemma 4.1 *Let G be a maximal outer-planar graph, embedded such that all vertices are on the outer-face, and let (s,t) be an edge of G on the outer-face. Then G is a series-parallel graph with terminals s,t with fan-out at most 2. Moreover, for every parallel composition one subgraph is an edge.*

Proof: We proceed by induction on the number of vertices; the claim is trivial for $n = 2$ since G is an edge. For $n \geq 3$, consider the interior face F incident to edge (s,t). Since G is maximal outer-planar, F is a triangle; let x be the third vertex on the triangle. Let G_s be the subgraph induced by all vertices between s and x on the outer-face of G, and let G_t be the subgraph induced by all vertices between x and t on the outer-face of G. See also Figure 13.

G_s is outer-planar with edge (s,x) on the outer-face, and G_t is outer-planar with edge (t,x) on the outer-face, so we obtain a series-parallel composition for them with terminals s,x and x,t, respectively. Combine these two compositions in series, and then apply a parallel composition with edge (s,t). This gives a series-parallel composition for G that satisfies all conditions. □

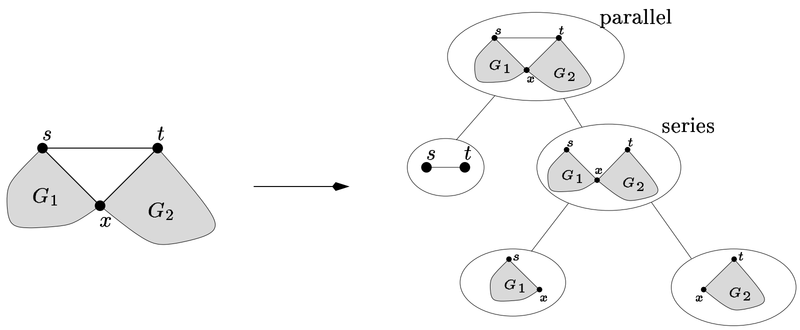

Figure 13: A series-parallel composition for an outer-planar graph.

From Theorem 2, we hence obtain the following for outer-planar graphs:

Theorem 3 *Every outer-planar graph has a visibility representation of area $O(n \log n)$.*

4.2 Orthogonal box-drawings

Next, we study how to modify our construction such that for outer-planar graphs, we can draw all vertices on one face. To do so, we use a series-parallel composition for which the terminals are an edge (s,t) on the outer-face of the graph (Lemma 4.1.) Then each parallel composition contains only one sub-graph that is not an edge. The constructions of Section 3 hence simplify to the ones shown in Figure 14; note in particular that case (S2b) does not happen.

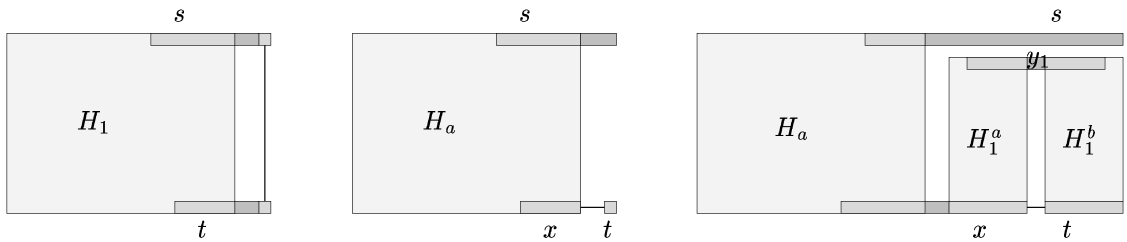

Figure 14: The drawing algorithm applied to outer-planar graphs. Case (P), (S1) and (S2a).

This keeps all vertices on the outer-face except in case (S2a): here edge (x, t) cuts off y_i (and possibly other vertices in H_1^a and H_1^b) from the outer-face. We can avoid this by routing edge (x, t) with two bends, see Figure 15. The clockwise order of edges around x in the drawing then exactly corresponds to the clockwise order of edges around x in the planar embedding that had all vertices on the outer-face; therefore the drawing has all vertices on the outer-face.

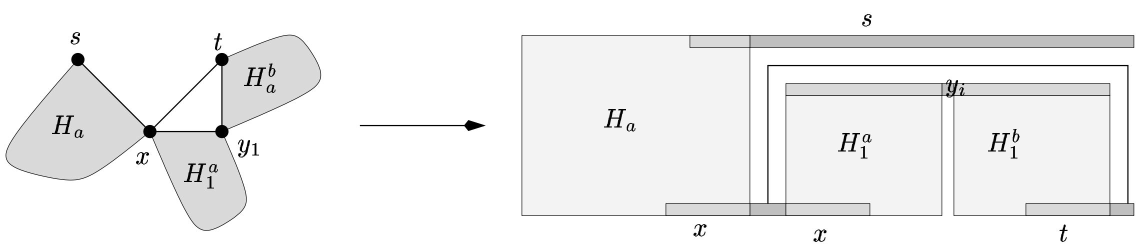

Figure 15: The alternate construction to keep all vertices on the outer-face.

Note that we use one more row to route the edge (x, t), but in exchange, there is no need to free the terminals in H_1^a and H_1^b; the height requirement therefore is unchanged compared to the original construction, and is $O(\log m)$ since the fan-out is 2. The width requirement increases (every edge with two bends now uses two columns instead of none), but still clearly remains $O(m)$.

Theorem 4 *Let G be an outer-planar graph. Then G has an orthogonal box-drawing of area $O(n \log n)$ such that all vertices are on the outer-face. Every edge has at most 2 bends.*

5 Lower bounds

5.1 Series-parallel graphs

Most of the previously given lower bounds for planar drawings (see e.g. [FPP88, MNN01, BB05]) rely on an argument that we call the "stacked cycle argument", which we briefly repeat here because we will modify it later. Assume we have a planar graph G with a fixed planar embedding and outer-face. Furthermore, assume that G has vertex-disjoint cycles $C_1, \ldots, C_k$ such that for any $i > 1$, any path from a vertex in C_i to the outer-face crosses all cycles C_j with $j > i$. We say that G has *k stacked cycles.* Now in any drawing of G that respects the planar embedding and outer-face, each cycle C_i must be drawn "around" all cycles $C_1, \ldots, C_{i-1}$. Since the cycles are vertex-disjoint, this is possible only if the drawing of C_i uses at least one more row on top, one more column on the right, one more row below, and one more column on the left than the drawing that contained $C_1, \ldots, C_{i-1}$. (This argument holds in all drawing models that we study.) Therefore, a graph with k stacked cycles needs at least $2k$ rows and $2k$ columns in any drawing that respects this planar embedding and outer-face.

This stacked cycles argument is usually applied to graphs with a fixed planar embedding and a fixed outer-face. However, we obtain the same asymptotic lower bound even if we allow a choice of the outer-face. For let f be an arbitrary face, and presume that i is such that f lies outside C_i but inside C_{i+1}. Then in a planar drawing where f is the outer-face, cycles $C_1, \ldots, C_i$ are stacked and cycles $C_k, \ldots, C_{i+1}$ are stacked (in reverse order.) Therefore, we have at least $\lceil k/2 \rceil$ stacked cycles and need k rows and k columns in any planar drawing that respects the planar embedding (but not necessarily the outer-face.)

Using the stacked cycles argument, it is easy to come up with $\Omega(n^2)$ lower bounds on the area, simply by constructing graphs that consist of $n/3$ stacked cycles [FPP88], or $\Omega(n)$ stacked cycles for some graph classes that do not allow stacked triangles [MNN01, BB05]. In particular, there is a series-parallel graph that consists of $n/3$ stacked cycles, and hence needs $\Omega(n^2)$ area in any drawing that respects the planar embedding. See Figure 16.

Theorem 5 *There exists a series-planar graph that requires $\Omega(n^2)$ area in any drawing that respects the planar embedding.*

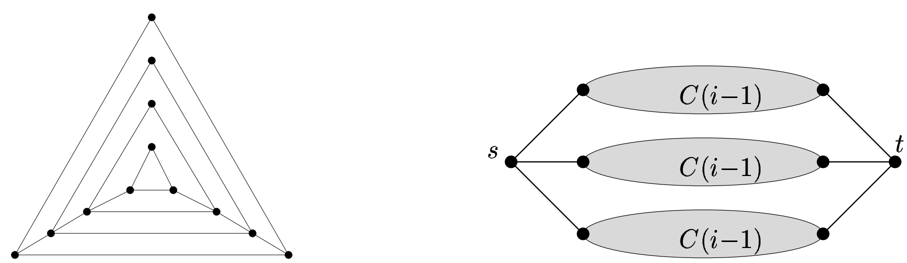

Figure 16: Two lower bounds for series-parallel graphs. The left graph has $n/3$ stacked cycles, and hence needs $\Omega(n^2)$ area in any drawing that respects the planar embedding. The right graph has $\Omega(\log n)$ stacked cycles in any planar embedding.

Note, however, that our graph (in contrary to the other lower bound graphs cited above) does have many different planar embeddings, and using a different embedding one can easily construct drawings of it in area $O(n)$. We will have to work harder to obtain series-parallel graphs with good lower bounds even if we can choose the planar embedding.

To do so, we first define a class of series-parallel graphs $C(i)$ such that in any poly-line drawing, the smaller of width and height is $i \in \Omega(\log n)$. $C(1)$ is a triangle. $C(i)$ consists of three copies of $C(i-1)$ and two more vertices s, t, with each s-terminal of $C(i-1)$ incident to s and each t-terminal of $C(i-1)$ incident to t. See Figure 16.

Note that $C(i)$ essentially has only one embedding, because changing the order of edges incident to s and t does not change what subgraph is attached there. Regardless of the embedding of $C(i)$, there is one cycle (the outer-face), which is around (and entirely disjoint of) one copy of $C(i-1)$. Therefore, by induction $C(i)$ has i stacked cycles. One can easily show by induction that $C(i)$ has $4 \cdot 3^{i-1} - 2$ vertices, and hence $C(i)$ has $\Omega(\log n)$ stacked cycles.

Note that $C(i)$ has fan-out 3, so it proves that our construction (which gives $O(\log n)$ in this case) gives asymptotically optimal height.

Now we use the graphs $C(i)$ to create graphs with good lower bounds on the area for orthogonal box-drawings. Recall that $K_{2,n}$ is the complete bipartite graph with 2 and n vertices, respectively. This is a series-parallel graph, with the two vertices of degree n as terminals. For given parameter N, let $C(\log_3 N) \cup K_{2,N}$ denote the graph obtained from a parallel composition of $C(\log_3 N)$ with $K_{2,N}$, and note that this graph has $n = \theta(N)$ vertices.

Theorem 6 *There exists a series-parallel graph that needs $\Omega(n \log n)$ area in any orthogonal box-drawing. In particular, if the graph is drawn in a $W \times H$-grid with $W \geq H$, then $W \in \Omega(n)$ and $H \in \Omega(\log n)$.*

Proof: Consider an arbitrary orthogonal box-drawing of $C(\log_3 N) \cup K_{2,N}$. Then $H \geq \log_3 N \in \theta(\log n)$, since $C(\log_3 N)$ has $\log_3 N$ stacked cycles. Also, $K_{2,N}$ has a vertex of degree N. The box representing this vertex must have perimeter at least N, so the larger of its width and height must be at least $N/4$, so the drawing of $K_{2,N}$ must have $W \geq N/4 \in \theta(n)$. □

Note that this theorem crucially relies on that vertices are represented by boxes; once we allow poly-lines, a vertex of high degree does not require particularly much area by itself. However, using a different argument on the same graph, we can get non-trivial lower bounds also for straight-line drawings.

Theorem 7 *There exists a series-parallel graph that requires $\Omega(n \log n / \log \log n)$ area in any straight-line drawing.*

Proof: We use again the graph $C(\log_3 N) \cup K_{2,N}$, which has $\theta(N)$ vertices. Assume we have a straight-line drawing of this graph with width W and height H, and $W \geq H$. As before, we have $H \geq \log_3 N$. Now we apply a result on straight-line drawings of $K_{2,N}$, which states that $W \log H \in \Omega(N)$ [BCLO03].

Note that the function $x/\log x$ is minimized when x is as small as possible (presuming $x \geq 2$). Therefore $H/\log H \geq \log_3 N / \log \log_3 N$ (as long as $N \geq 9$), which gives

$$W \cdot H = W \cdot \log H \ \cdot \ H/\log H \geq W \cdot \log H \ \cdot \ \log_3 N/\log\log_3 N \in \Omega(N) \cdot \log_3 N/\log\log_3 N$$

and hence the result. □

5.2 Outer-planar graphs

The lower bound graph given above is not outer-planar. While we suspect that $\Omega(n \log n)$ area is required for some outer-planar graph (in particular the so-called *snowflake graph* in

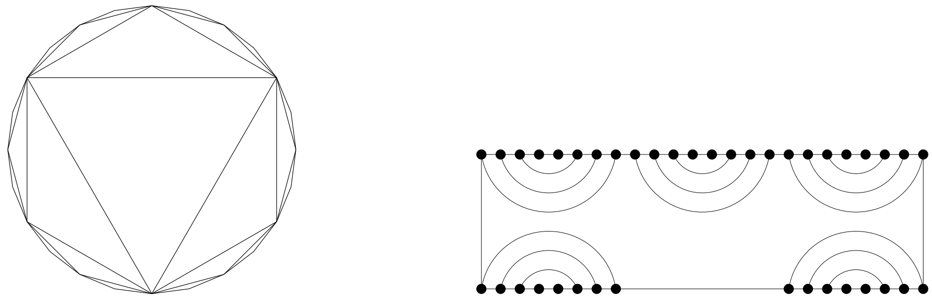

Figure 17: The snowflake graph, and a graph that requires $\Omega(n^2)$ area if vertices are drawn on the bounding box.

Figure 17), we leave this as an open problem, and instead give lower bounds under a stronger drawing requirement.

While in our orthogonal box-drawings of outer-planar graph all vertices are on the outer-face, they are not clearly visible as such (consider in particular vertex y_1 in Figure 15.) Another more natural way of displaying outer-planar graphs would be to require that all vertices are on the boundary of the minimum enclosing rectangle (respectively touch it if they are boxes). It is known that trees may require $\Omega(n \log n)$ are in this model [Ull83, p.83ff]. We show now that for outer-planar graphs, $\Omega(n^2)$ area is required in this model.

Theorem 8 *There exists an outer-planar graph G such that any poly-line drawing Γ of G with all vertices on the boundary of the minimum bounding box has area $\Omega(n^2)$.*

Proof: Let G be the outer-planar graph illustrated in Figure 17. It consists of five groups of $n/5$ vertices, connected with chords between them. Assume we have a drawing Γ such that all vertices are on the boundary of the minimum bounding box $\mathcal{B}$. Of the five groups, hence at least one must be entirely on one side of $\mathcal{B}$, say the top. Duplicate the drawing and flip it upside down. The resulting drawing has asymptotically the same area and contains a multi-graph with $n/2$ stacked cycles, so its area is $\Omega(n^2)$. □

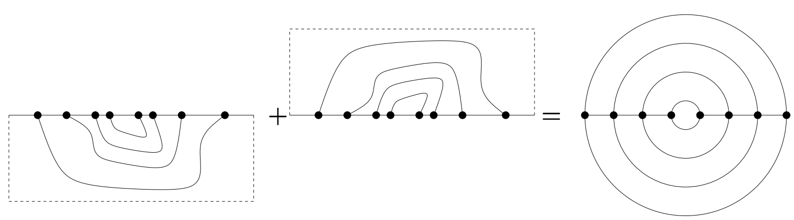

Figure 18: Any poly-line drawing must have $\Omega(n^2)$ area.

We note here without proof that the same graph also yields a lower bound of $\Omega(n^2)$ for other models of "being visible on the outer-face", such as "having a horizontal or vertical segment that reaches to the boundary of the enclosing rectangle", or even "having an escape hatch" (see [Lei80]).

5.3 Other graph classes

A natural question for this paper is what other sub-classes of planar graphs could be studied to see whether $o(n^2)$ area can be achieved. We give here two sub-classes for which this is *not* possible, because we can show a lower bound of $\Omega(n^2)$ area.

5.3.1 k-outer-planar graphs

A k-outer-planar graph is defined as follows. Let G be a graph with a fixed planar embedding. G is called 1-*outer-plane* if all vertices of G are on the outer-face (i.e., if G is outer-planar in this embedding.) G is called k-*outer-plane* if the graph that results from removing all vertices from the outer-face of G is $(k-1)$-outer-plane in the induced embedding. A graph G is called k-*outer-planar* if it is k-outer-plane in some planar embedding.

Clearly, k-outer-planar graphs generalize the concept of outer-planar graphs, and hence for small (constant) k are good candidates for $o(n^2)$ area. Also, by definition we cannot use a stacked cycle argument on them (a k-outer-planar graph has at most k stacked cycles.) Nevertheless, we can show an $\Omega(n^2)$ lower bound on the area even for 2-outer-planar graphs.

In some ways, this is to be expected. For in order to draw a 2-outer-planar graph G, we must draw the outer-planar graph G_1 that results from removing the outer-face of G. Moreover, G_1 must be drawn in such a way that all vertices are accessible to their neighbours in $G - G_1$. But we know from Section 5.2 that such drawings typically require $\Omega(n^2)$ area. This in itself is not a proof (because we would have to make "accessible" more precise), and for the exact proof we instead modify the stacked-cycle argument.

Let G be a graph with a fixed planar embedding, and let $C_1, \ldots, C_k$ be k edge-disjoint cycles in G. Moreover, for $1 \leq i < k$, cycles C_i and C_{i+1} are vertex-disjoint, except that they may have one vertex v_i in common. (We allow that $v_i = v_{i+1}$, i.e., any number of consecutive cycles may have one vertex in common.) We say that $C_1, \ldots, C_k$ are *1-fused stacked cycles* if for $i > 1$, any path from a vertex $v \neq v_i$ in C_i to the outer-face must cross all cycles $C_{i+1}, \ldots, C_k$. Figure 19 illustrates this concept.

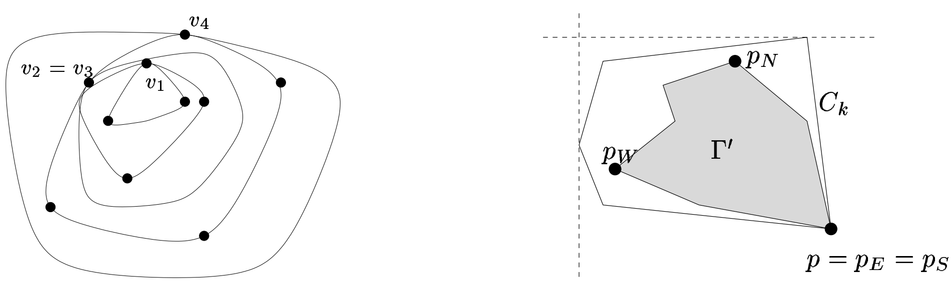

Figure 19: 5 cycles that are 1-fused stacked cycles, and adding a 1-fused cycle around a drawing.

Lemma 5.1 *Let G be a planar graph with a fixed planar embedding and outer-face, and assume G has k 1-fused stacked cycles $C_1, \ldots, C_k$. Then any poly-line drawing of G that respects the planar embedding has width and height at least $k+1$.*

Proof: We proceed by induction on k. Clearly we need width and height 2 to draw the cycle C_1. For $k > 1$, let G' be the subgraph obtained from G by deleting all vertices of C_k except the vertex v_{k-1} that C_k has in common with C_{k-1} (if it exists.) Then G' has the 1-fused stacked cycles $C_1, \ldots, C_{k-1}$, and by induction needs width and height at least k in any poly-line drawing.

Consider an arbitrary poly-line drawing Γ of G, and let Γ' be the induced drawing of G', which has width and height at least k. Consider Figure 19. The drawing of C_k in Γ must stay outside Γ', except at the point p where v_{k-1} is drawn. Let p_N and p_S be points in the top and bottom row of Γ'; by $k \geq 2$ they are distinct. So $p \neq p_N$ or $p \neq p_S$; assume the former. To go around p_N, the drawing of C_k in Γ must reach a point strictly higher than p_N, and hence uses at least one more row above Γ'. Similarly one shows that Γ has at least one more column than Γ'. □

Similarly as for stacked cycles, we get asymptotically the same lower bounds for 1-fused stacked if we can choose the outer-face (with a fixed planar embedding), because at least half of the cycles remain 1-fused stacked cycles.

Now all that remains to do is to create a 2-outer-planar graph that contains $\Omega(n)$ 1-fused stacked cycles. Figure 20 shows such a graph. Note that this graph is 3-connected, hence no other planar embedding is possible. Also note that in fact all stacked cycles of this graph share the same vertex, which is specifically allowed in our definition.

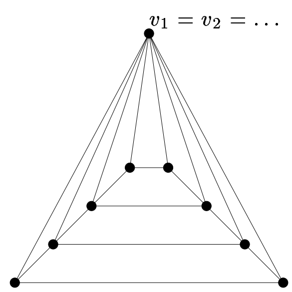

Figure 20: A 2-outer-planar graph that requires $\Omega(n^2)$ area in any poly-line drawing.

Theorem 9 *There exists a 3-connected 2-outer-planar graph that requires $\Omega(n^2)$ area in any poly-line drawing.*

5.3.2 Partial k-trees

Since we only consider partial k-trees for $k = 2, 3$, we will not give the general definition here, but refer the interested reader for example to [Bod97]. Partial 2-trees are exactly the series-parallel graphs. Planar partial 3-trees can be characterized as follows. A planar graph is a *planar 3-tree* if it is obtained recursively as follows: A triangle is a planar 3-tree. If G is a planar 3-tree, and T is a face of G that is a triangle, then the graph obtained from G by adding a new vertex inside T and making it adjacent to all vertices of T is also a planar 3-tree. Any subgraph of a planar 3-tree is a *planar partial 3-tree.*

The graph in Figure 20 can easily be shown to be a partial 3-tree. We give another planar partial 3-tree in Figure 21 (the vertices are labeled in the order in which they are added), which also requires $\Omega(n^2)$ area, since it has $n/3$ stacked cycles and is 3-connected, hence has only one possible planar embedding. This graph has some additional features, hence eliminating other possible candidates for graph classes to be drawn in $o(n^2)$ area. In particular, it is triangulated and has maximum degree 6. Without going into the definition of these terms, we would like mention that this graph also has pathwidth 3, and in fact has proper pathwidth 3 (all these define strict subclasses of each other.)

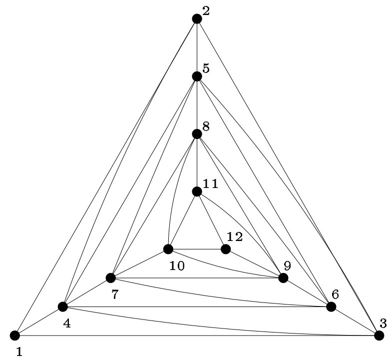

Figure 21: A planar 3-tree that requires $\Omega(n^2)$ area.

Theorem 10 *There exists a planar partial 3-tree that requires $\Omega(n^2)$ area in any planar poly-line drawing. Moreover, this graph is triangulated and has maximum degree 6, and has proper pathwidth 3.*

Since planar partial 3-trees are also partial k-trees for any $k \geq 3$, our lower bounds holds for all partial k-trees with $k \geq 3$, hence destroying the hope that the linear-area layouts in 3D [DMW05] could be replicated in 2D.

6 Conclusion

In this paper, we studied planar poly-line drawings of some subclasses of planar graphs. Using a recursive algorithm, we achieved $O(n^{3/2})$ for series-parallel graphs, and $O(fn \log n)$ area for series-parallel graphs with fan-out f. In particular, this implies $O(\Delta n \log n)$ area for series-parallel graphs with maximum degree Δ, and $O(n \log n)$ area for outer-planar graphs. The drawings we create are in fact flat visibility representations, from which poly-line drawings of the same asymptotic area are easily obtained. For outer-planar graphs, we also gave a variant that achieves orthogonal box-drawings of area $O(n \log n)$ which keeps all vertices on the outer-face.

We also studied lower bounds, both for these two graph classes, and for some other subclasses of planar graphs where we showed that $\Omega(n^2)$ area is required.

Many open problems remain:

- What subclasses of planar graphs have *straight-line drawings* of area $o(n^2)$? Can we achieve $o(n^2)$ for all outer-planar graphs, not only those where the diameter and/or the maximum degree is small? Can we achieve $o(n^2)$ straight-line drawings of series-parallel graphs, at least under some conditions on other graph-parameters?

- For outer-planar graphs, can we achieve visibility representations of area $O(n \log n)$ that keep all vertices on the outer-face? Or better even, that respect any given planar embedding?
- What are correct lower bounds? Does the snowflake graph (or some other outer-planar graph) require $\Omega(n \log n)$ area? Does some series-parallel graph (with large fan-out) require $\omega(n \log n)$ area?

  Of particular interest are also other techniques for proving lower bounds. The technique of stacked cycles cannot be applied to outer-planar graph. We also conjecture that more than $\Omega(\log n)$ stacked cycles are not possible for series-parallel graphs. Either way, to obtain better lower bounds we need different techniques, and the only one that we are aware of only works for $K_{2,n}$ for straight-line drawings [BCLO03]. What are other techniques of proving lower bounds on the area for planar drawings?

References

[BB05] T. Biedl and F. Brandenburg. Drawing planar bipartite graphs with small area. In *Canadian Conference on Computational Geometry (CCCG'05)*, 2005. To appear.

[BCDB+94] P. Bertolazzi, R.F. Cohen, G. Di Battista, R. Tamassia, and I.G. Tollis. How to draw a series-parallel digraph. *Intl. J. Comput. Geom. Appl.*, 4:385–402, 1994.

[BCLO03] T. Biedl, T. Chan, and A. López-Ortiz. Drawing $K_{2,n}$: a lower bound. *Information Processing Letters*, 85/6:303–305, 2003.

[BEG+03] F. Brandenburg, D. Eppstein, M.T. Goodrich, S.G. Kobourov, G. Liotta, and P. Mutzel. Selected open problems in graph drawing. In *Graph Drawing (GD'03)*, volume 2912 of *Lecture Notes in Computer Science*, pages 320–331. Springer-Verlag, 2003.

[Bie02] T. Biedl. Drawing outer-planar graphs in $O(n \log n)$ area. In *Graph Drawing*, volume 2528 of *Lecture Notes in Computer Science*, pages 54–65. Springer, August 2002.

[Bod97] H. L. Bodlaender. Treewidth: algorithmic techniques and results. In *Mathematical foundations of computer science 1997 (Bratislava)*, volume 1295 of *Lecture Notes in Comput. Sci.*, pages 19–36. Springer, Berlin, 1997.

[CBTT95] R. F. Cohen, G. Di Battista, R. Tamassia, and I. G. Tollis. Dynamic graph drawings: Trees, series-parallel digraphs, and planar st-digraphs. *SIAM J. Comput.*, 24(5):970–1001, 1995.

[CU96] N. Castañeda and J. Urrutia. Straight line embeddings of planar graphs on point sets. In *Canadian Conference on Computational Geometry (CCCG '96)*, pages 312–318, 1996.

[DBETT98] G. Di Battista, P. Eades, R. Tamassia, and I. Tollis. *Graph Drawing: Algorithms for Geometric Representations of Graphs.* Prentice-Hall, 1998.

[DMW05] V. Dujmovic, P. Morin, and D.K. Wood. Layout of graphs with bounded treewidth. *SIAM J. on Computing*, 34(3):553–579, 2005.

[Fár48] I. Fáry. On straight line representation of planar graphs. *Acta. Sci. Math. Szeged*, 11:229–233, 1948.

[FKK97] U. Fößmeier, G. Kant, and M. Kaufmann. 2-visibility drawings of planar graphs. In S. North, editor, *Symposium on Graph Drawing, GD 96*, volume 1190 of *Lecture Notes in Computer Science*, pages 155–168. Springer-Verlag, 1997.

[FLW03] S. Felsner, G. Liotta, and S. Wismath. Straight-line drawings on restricted integer grids in two and three dimensions. *Journal of Graph Algorithms and Applications*, 7(4):335–362, 2003.

[FPP88] H. de Fraysseix, J. Pach, and R. Pollack. Small sets supporting fary embeddings of planar graphs. In *Twentieth Annual ACM Symposium on Theory of Computing*, pages 426–433, 1988.

[FPP90] H. de Fraysseix, J. Pach, and R. Pollack. How to draw a planar graph on a grid. *Combinatorica*, 10:41–51, 1990.

[GGT96] A. Garg, M.T. Goodrich, and R. Tamassia. Planar upward tree drawings with optimal area. *International J. Comput. Geom. Appl.*, 6:333–356, 1996.

[GR03] A. Garg and A. Rusu. Area-efficient drawings of outerplanar graphs. In *Graph Drawing '03*, volume 2912 of *Lecture Notes in Computer Science*, pages 129–134. Springer-Verlag, 2003.

[HEQL98] S.-H. Hong, P. Eades, A. Quigley, and S. H. Lee. Drawing algorithms for series-parallel digraphs in two and three dimensions. In *Graph Drawing*, volume 1547 of *Lecture Notes in Computer Science*, pages 198–209, Springer-Verlag, 1998.

[Kan96] G. Kant. Drawing planar graphs using the canonical ordering. *Algorithmica*, 16:4–32, 1996.

[Lei80] C. Leiserson. Area-efficient graph layouts (for VLSI). In *21st IEEE Symposium on Foundations of Computer Science*, pages 270–281, 1980.

[MNN01] K. Miura, T. Nishizeki, and S. Nakano. Grid drawings of 4-connected plane graphs. *Discrete Computational Geometry*, 26:73–87, 2001.

[Sch90] W. Schnyder. Embedding planar graphs on the grid. In *1st Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 138–148, 1990.

[Shi76] Y. Shiloach. *Arrangements of Planar Graphs on the Planar Lattice*. PhD thesis, Weizmann Institute of Science, 1976.

[Ste51] S. Stein. Convex maps. In *Amer. Math. Soc.*, volume 2, pages 464–466, 1951.

[Ull83] J.D. Ullman. *Computational Aspects of VLSI*. Computer Science Press, 1983.

[Wag36] K. Wagner. Bemerkungen zum Vierfarbenproblem. *Jahresbericht der Deutschen Mathematiker-Vereinigung*, 46:26–32, 1936.

[Wis85] S. Wismath. Characterizing bar line-of-sight graphs. In *1st ACM Symposium on Computational Geometry*, pages 147–152, Baltimore, Maryland, USA, 1985.