

A 2-Approximation for the Height of Maximal Outerplanar Graph Drawings

by

Philippe Demontigny

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Master of Mathematics
in
Computer Science

Waterloo, Ontario, Canada, 2016

© Philippe Demontigny 2016

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

In this thesis, we study drawings of maximal outerplanar graphs that place vertices on integer coordinates. We introduce a new class of graphs, called umbrellas, and a new method of splitting maximal outerplanar graphs into systems of umbrellas. By doing so, we generate a new graph parameter, called the umbrella depth (ud), that can be used to approximate the optimal height of a drawing of a maximal outerplanar graph. We show that for any maximal outerplanar graph G , we can create a flat visibility representation of G with height at most $2 \cdot ud(G) + 1$. This drawing can be transformed into a straight-line drawing of the same height. We then prove that the height of any drawing of G is at least $ud(G) + 1$, which makes our result a 2-approximation for the optimal height. The best previously known approximation algorithm gave a 4-approximation. In addition, we provide an algorithm for finding the umbrella depth of G in linear time. Lastly, we compare the umbrella depth to other graph parameters such as the pathwidth and the rooted pathwidth, which have been used in the past for outerplanar graph drawing algorithms.

Acknowledgments

I would like to thank my advisor Therese Biedl for her expert guidance and infinite patience, my colleagues in the school of computer science for the good times and support, and my family for always being there for me.

Table of Contents

List of Figures	vii
1 Introduction	1
1.1 Background	1
1.2 Existing Results for Outerplanar Graphs	2
1.3 Overview of Thesis	4
2 Preliminaries	5
3 New Graph Parameters	8
3.1 Umbrellas and Umbrella Systems	8
3.2 Free vs. Rooted Umbrella Depth	11
4 From Umbrella Systems to Flat Visibility Representations	15
4.1 Drawing the Root Umbrella	15
4.2 Drawing an Umbrella System	20
5 From Drawings to Umbrella Systems	22
5.1 Left-free and Right-free Edges	22
5.2 Lower Bound	25
6 Finding the Umbrella Depth	29

7	Comparison with Pathwidth	37
7.1	Definitions	37
7.2	Umbrella Depth Inequalities	39
7.3	Comparison of Bounds	41
7.4	Tightness of Bounds	42
8	Conclusions and Future Work	48
	References	50

List of Figures

1.1	A straight-line drawing including the dual tree (left, dashed edges) and a shaded outerplanar path, and a flat visibility representation (right) of a maximal outerplanar graph. Each drawing has height 4.	3
3.1	An umbrella as the union of three outerplanar paths (left) and an example of an umbrella with the handle shaded (right).	9
3.2	A graph with root-edge (u, v) , root umbrella U , anchor edge (x, y) , and hanging subgraph S	10
3.3	A rooted umbrella system of depth 2 with root-edge (u, v) and root umbrella shaded. Edge (x, y) is the root-edge of a hanging umbrella system of depth 1.	10
3.4	An umbrella with shaded subgraph $S = S_P \cup S_1$, where S_P is the lightly shaded region and S_1 is the darker region.	12
3.5	An example from the proof of Lemma 3.6 where $(x, y) \in U_0$ and $S_{x,y}$ is part of the handle of U_0	13
3.6	An example from the proof of Lemma 3.7 where the shaded umbrella is U_0 , the striped region is $U_0 \cap U_0^*$, and the unshaded regions are part of U_0^*	14
4.1	An umbrella U_0 for which all vertices are neighbors of u or v (left) and a construction for U_0 on two layers where edge (u, v) spans the top layer (right).	16
4.2	An example standard embedding of the handle P of an umbrella (left) and a drawing of P that satisfies Claim 4.2 (right).	17
4.3	Releasing edge (u, v) in a flat visibility representation so that (u, v) spans the top layer.	18

4.4	A flat visibility representation of a root umbrella on three layers.	19
4.5	Inserting the drawings of hanging subgraphs into the flat visibility representation from Figure 4.4.	21
5.1	A flat visibility representation in which vertex w has a left escape path, r_1 has a right escape path, (ℓ_1, ℓ_2) is a left-free edge, and (r_1, r_2) is a right-free edge.	23
5.2	Expanding Γ' as part of the proof of Lemma 5.3	25
5.3	The standard embedding of the root-umbrella U_0 and one of its hanging subgraphs S . The dotted path represents the dividing path Π_1 while the dashed path represents Π_2	26
5.4	Extracting dividing paths from a flat visibility representation. The dotted path represents Π_1 while the dashed path represents Π_2	27
6.1	A fan umbrella of b (left) and a partial umbrella of b (right).	30
6.2	The labeling used for the formulas in the dynamic programming algorithm.	30
6.3	A handle umbrella $U_{a,b}^h$ with cap (a, b) . The shaded region is $U_{a,b}^h \cap U_{b,c}^h$, and the dashed lines indicate possible anchor edges.	31
6.4	A fan umbrella $U_{b,a}^f$ of b . The shaded region is $U_{b,a}^f \cap U_{b,c}^f$, and the dashed lines indicate possible anchor edges.	32
6.5	A partial umbrella $U_{b,a}^p$ of b . The shaded region is $U_{b,a}^p \cap U_{b,c}^p$, and the dashed lines indicate possible anchor edges.	34
6.6	An umbrella as the union of a fan umbrella of a (shaded) and a partial umbrella of b (unshaded). The striped region belongs to both the fan and partial umbrellas, and the dashed lines indicate possible anchor edges. . .	35
7.1	A root umbrella with the dual trees for the handle and both fans included. S_1^*, S_2^* , and S^* indicate components of $T - P^*$, where T is the dual tree of the entire graph.	40
7.2	A drawing of G_1 from the proof of Lemma 7.11. The dotted lines indicate anchor edges for components of G_H when $H > 1$	42
7.3	A drawing of G_1 from the proof of Lemma 7.12. The dotted lines indicate anchor edges for components of G_H when $H > 1$	44

7.4	A drawing of G_1 from the proof of Lemma 7.13 with the dual tree included. The dotted lines indicate anchor edges for components of G_H when $H > 1$.	45
7.5	A flat visibility representation of the graph from Figure 7.4 created using the algorithm from Chapter 4.	46
7.6	Drawing the dual of a main path (top), merging the recursively created subgraph (middle), and releasing edge (u, v) (bottom) as part of the algorithm from [5] (Theorem 1.6) for creating a flat visibility representation of the graph from Figure 7.4.	47

Chapter 1

Introduction

1.1 Background

Graph drawing is the art of creating a picture of a graph that is both functional and visually appealing. Although there are many possible objectives for such a drawing, a common goal is to keep the area small without making the graph hard to see. One way to accomplish this is by minimizing the area while restricting the placement of vertices to integer coordinates. By doing so, one can create drawings in which the representations of vertices and edges are not overly crowded or small. Such an approach has many applications, including data visualization [9], DNA mapping [25], and circuit layout [20].

Of particular interest are drawings of planar graphs, which are graphs that can be drawn without edge crossings. Planar graphs are very popular, and have been studied for many years. (See, for example, [9] and the references therein.) Numerous results that bound the total area of planar graph drawings have been found (which are summarized in [11]), but there are still many questions left unsolved.

It is known that finding the minimum area drawing for a given planar graph is NP-hard [19]. However, for some subclasses of planar graphs, such as planar 3-trees, it is possible to find a minimum area drawing in polynomial time [21]. For general planar graphs, it has been known for a while that a straight-line drawing can always be found that uses an $O(n) \times O(n)$ grid [16, 22], where n is the number of vertices. It is also known that there are certain planar graphs that require an $\Omega(n) \times \Omega(n)$ grid [15].

When visualizing any kind of graph in a software application, such as a program for visualizing computer networks, it is important that the information be easy for a user to see

and process. If a drawing doesn't fit on a single screen, then the user would need to scroll in order to see all of the information. Scrolling in one direction is easy and intuitive, but scrolling in multiple directions (up-down and left-right, for example), can quickly become overwhelming. Thus it is often better to focus on minimizing one dimension of a 2D drawing at a time in order to avoid this situation.

In this thesis, we focus on minimizing the height of a planar graph drawing, but note that minimizing the width is equivalent after rotation. So far, it not known whether or not finding the minimum height is NP-hard for general planar graphs. The closest result is an NP-hardness proof by Heath and Rosenberg [18] for so-called proper drawings in which the y-coordinates of the endpoints for every edge are exactly one unit apart. We also know that when given the height H , testing whether a drawing of height H exists for a particular graph is fixed parameter tractable in H [12].

It is also known that any graph with a planar drawing of height H has pathwidth at most H [13]. This makes the pathwidth a useful parameter for approximating the height of a planar graph drawing. Indeed, when one considers only trees, upper and lower bounds based on the pathwidth have been found for a variety of different drawing styles (including proper, straight-line, upward, order-preserving, etc.). A selection of such results can be found in [23, 7, 2]. For the specific case of straight-line drawings of trees, a linear-time algorithm for finding the minimum height was discovered later [1].

1.2 Existing Results for Outerplanar Graphs

This thesis introduces drawing algorithms for a subclass of planar graphs, the so-called *maximal outerplanar graphs*. In this section, we review a number of previous results for the area and height of outerplanar graph drawings. One of the first results in this field is the following by Biedl, which establishes an upper bound for flat visibility representations, a drawing style that we will use often in this thesis. Figure 1.1 illustrates a flat visibility representation as well as a straight-line drawing of an outerplanar graph.

Theorem 1.1. [3] *Every outerplanar graph with n vertices has a flat visibility representation in a $(\frac{3}{2}n - 2) \times (3 \log n - 1)$ grid.*

For straight-line drawings, one of the first results for the area was based on the *degree* of an outerplanar graph, which is the maximum number of edges incident to a single vertex in the graph.

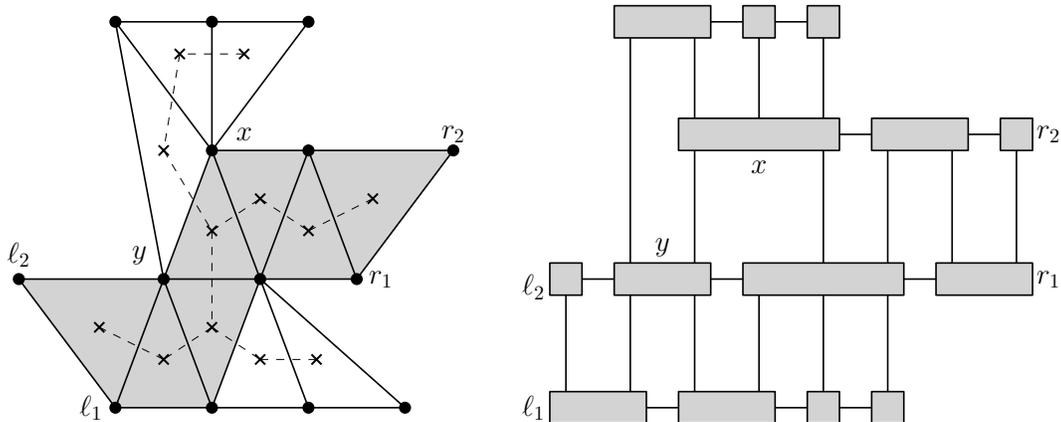


Figure 1.1: A straight-line drawing including the dual tree (left, dashed edges) and a shaded outerplanar path, and a flat visibility representation (right) of a maximal outerplanar graph. Each drawing has height 4.

Theorem 1.2. [17] *Every outerplanar graph with n vertices and degree d admits a planar straight-line drawing with area $O(dn^{1.48})$.*

Theorem 1.2 was further improved by Frati.

Theorem 1.3. [14] *Every outerplanar graph with n vertices and degree d admits a planar straight-line drawing with area $O(dn \log n)$.*

A few years later, the first sub-quadratic upper bound for the area of straight-line drawings was found by Di Battista and Frati.

Theorem 1.4. [10] *Every outerplanar graph with n vertices admits a planar straight-line drawing with area $O(n^{1.48})$.*

Lastly, we have the following result, which gives a linear upper bound for a particular subclass of outerplanar graphs.

Theorem 1.5. [10] *Every balanced outerplanar graph with n vertices admits a planar straight-line drawing for which both the height and the width are $O(\sqrt{n})$.*

The results above all establish bounds on the area, which is the product of the width and height of a drawing. If we only care about one dimension, say the height of a drawing, then for flat visibility representations we have an $O(\log n)$ bound [3], and for straight-line

drawings we have an $O(d \log n)$ bound [14]. These bounds can be derived from the results for the area in Theorems 1.1 and 1.3, respectively.

Lastly, Biedl improved her result from Theorem 1.1 by approximating the optimal height using bounds based on the pathwidth of the dual tree. This result holds for all maximal outerplanar graphs, and even more broadly for all outerplanar graphs that are 2-connected. Recall that if the solution to a minimization problem is H , then a k -approximation algorithm finds a solution that is always less than or equal to $k \cdot H$. The following result is therefore a 4-approximation for the optimal height.

Theorem 1.6. [5] *Every 2-connected outerplanar graph G has a flat visibility representation with height $4pw(G) - 3$, where $pw(G)$ is the pathwidth of G .*

Theorem 1.6 was the primary motivation for this thesis. We were interested in finding a better approximation for the height of maximal outerplanar graphs using a parameter other than the pathwidth.

1.3 Overview of Thesis

After giving preliminaries in Chapter 2, we introduce a new parameter for maximal outerplanar graphs called the umbrella depth. Details on the umbrella depth and what it represents can be found in Chapter 3. Then, in Chapter 4, we show that any outerplanar graph G has a flat visibility representation of height at most $2ud(G) + 1$, where $ud(G)$ is the umbrella depth of G . In Chapter 5, we show that the optimal height for any drawing of G is at least $ud(G) + 1$. This proves that our result is a 2-approximation for the optimal height, which must fall in the range $[ud(G) + 1, 2ud(G) + 1]$.

Our algorithm to compute the visibility representation assumes that the umbrella depth of G is known. In Chapter 6, we provide an algorithm for finding the umbrella depth in $O(n)$ time. In Chapter 7, we compare the umbrella depth to the pathwidth and rooted pathwidth, which have been used in previous papers to establish bounds on the optimal height for drawings of a maximal outerplanar graph. We show that our height-bounds are never worse than the bounds from those papers except for a small additive term. Lastly, in Chapter 8, we discuss possibilities for future research and other problems that remain open.

Chapter 2

Preliminaries

Let $G = (V, E)$ be a simple graph with n vertices and m edges. A special class of graphs are the *planar graphs*, which are graphs that admit a straight-line drawing without edge crossings. In this thesis, all graphs will in fact be *outerplanar*, which means they have a *standard planar embedding* in which all vertices are in the *outer face*, the infinite connected region outside the drawing. By contrast, any finite region enclosed by edges in an outerplanar graph is called an *interior face*, which is denoted by the vertices and edges that are adjacent to it. A *cutting edge* of a graph G is an edge that, when its ends are removed, splits G into multiple disjoint subgraphs. In the left side of Figure 1.1, the edge (x, y) is a cutting edge, while (ℓ_2, y) is not. All other edges in G are referred to as *non-cutting edges*.

We say a graph G is *maximal outerplanar* if adding any edge to it makes it no longer a simple outerplanar graph. In this thesis, we are concerned only with maximal outerplanar graphs that have at least 3 vertices, which are always 2-connected and in which all interior faces are triangles. An example of a maximal outerplanar graph can be found on the left side of Figure 1.1. A cutting edge in a maximal outerplanar graph can be seen to be the same as an edge that borders two interior faces of G . If (u, v) is a cutting edge in a general graph G , then $G - \{u, v\}$ splits into k connected components S_1, S_2, \dots, S_k . Define a *cut-component* of (u, v) to be $S_i \cup (u, v)$ for any $i \in [1, k]$. A simple (but frequently used) fact is that any cutting edge in an outerplanar graph has exactly two cut-components.

The *dual tree* of a maximal outerplanar graph G has a vertex for each interior face of G , and edges between vertices if their corresponding faces in G share an edge. Note that this definition is different from the so-called *dual graph* of G , which includes the outer face of G as a vertex. In this thesis, we will not make use of the dual graph. Since G is assumed to be maximal, and all faces are triangles, it follows that the maximum degree of the dual

tree of G is 3. Figure 1.1 includes an example of a dual tree for a maximal outerplanar graph as well.

Let an *outerplanar path* be any maximal outerplanar graph whose dual tree is a path. We will refer to any path between the vertices themselves as a *vertex path* to differentiate them from the outerplanar variety. We say the *endpoints* of an outerplanar path P are the vertices of degree 2 in P . Note that in any outerplanar path where $n > 3$, there will be exactly two such vertices. If $n = 3$, then G is a triangle, and all three vertices are endpoints by definition. We say that the four incident edges to an endpoint of P are the *end-edges* of P . In the left side of Figure 1.1, the shaded subgraph is an outerplanar path with endpoints ℓ_1 and r_1 , and end-edges (ℓ_1, ℓ_2) and (r_1, r_2) .

If edges (ℓ_1, ℓ_2) and (r_1, r_2) are distinct non-cutting edges in a maximal outerplanar graph G , then each of them is adjacent to a single inner face in the standard embedding of G . For any two such edges, we define the outerplanar path *between* (ℓ_1, ℓ_2) and (r_1, r_2) as the path whose dual connects the inner faces adjacent to (ℓ_1, ℓ_2) and (r_1, r_2) in G . Such a definition makes (ℓ_1, ℓ_2) and (r_1, r_2) end-edges of the resulting path. This idea is summarized in the following observation.

Observation 2.1. *For any two non-cutting edges (ℓ_1, ℓ_2) and (r_1, r_2) in a maximal outerplanar graph G , there exists an outerplanar path P in G that has (ℓ_1, ℓ_2) and (r_1, r_2) as end-edges.*

A *drawing* of a graph consists of a point or an axis-aligned box for every vertex, and a curve for every edge that intersects each of the points/boxes of its endpoints once. Such a drawing is *planar* if none of the points, boxes, or curves intersect unless the corresponding elements do in the original graph. Note that a planar drawing need not reflect a graph's standard planar embedding. In this thesis, whenever we discuss a drawing of a graph, we are referring to one that is planar. Of primary interest are *flat visibility representations*, in which vertices are represented by horizontal line segments, and edges are vertical or horizontal straight-line segments. For convenience, we will use boxes to represent vertices in flat visibility representations. We will also consider *straight-line drawings*, in which vertices are represented by points and edges are line segments between points. An example of both can be found in Figure 1.1.

In either type of drawing, we require that the vertex points or the ends of vertex segments are placed at points with integer y-coordinates. We call such a drawing a *layered graph drawing*, where each *layer* is the horizontal line defined by a single y-value. Two vertices are said to be in the same layer if they have the same y-coordinate. Note that in a layered drawing, vertex points and the ends of vertex segments do not need to be placed on

integer x-coordinates. The *height* of a layered graph drawing is the total number of layers in the drawing, including layers that contain no vertices. In any layered graph drawing, we can define a left-to-right ordering of vertices based on their x-coordinates. In a flat visibility representation, we can also define a left-to-right ordering of the vertical edges in the same manner.

In our thesis, we create flat visibility representations of maximal outerplanar graphs. However, some of the previous results concerning outerplanar graphs create straight-line drawings instead. If the objective is to minimize the height of the drawing, then this distinction is unimportant because of the following.

Theorem 2.2. *[6] Any planar straight-line drawing can be transformed into a planar flat visibility representation of the same height that preserves y-coordinates and left-to-right orders.*

The reverse direction is also possible.

Theorem 2.3. *[6] Any flat visibility representation can be transformed into a planar straight-line drawing of the same height that preserves y-coordinates and left-to-right orders.*

Chapter 3

New Graph Parameters

3.1 Umbrellas and Umbrella Systems

In this section, we introduce a special class of outerplanar graphs called *umbrellas*, and a method of splitting maximal outerplanar graphs into systems of umbrellas. These systems are the key to achieving the main results of this thesis, which are presented in later sections.

Definition 3.1. *Let an umbrella U be a maximal outerplanar graph that can be split into three outerplanar paths P, F_1 , and F_2 such that:*

1. *P is an outerplanar path with two end-edges (u, v) and (x, y) that are non-cutting edges of U .*
2. *F_1 contains only u and neighbors of u , while F_2 contains only v and neighbors of v .*
3. *In the standard embedding of U , the paths P, F_1 , and F_2 have no faces in common.*

We refer to the edge (u, v) as the cap of U , P as the handle, and F_1 and F_2 as the fans. See Figure 3.1 for an example.

Note that in any umbrella, each non-empty fan shares a single edge with the handle P . This edge is a cutting edge that is adjacent to one face of P , and one face of the fan in the standard embedding. In this thesis, we will use a consistent ordering for the vertices in the fans. Let a_1, a_2, \dots, a_ℓ be the ℓ neighbors of u in fan F_1 , labeled such that the outer face of F_1 in the standard embedding is $u, a_1, a_2, \dots, a_\ell$ with (u, a_ℓ) as the single edge shared

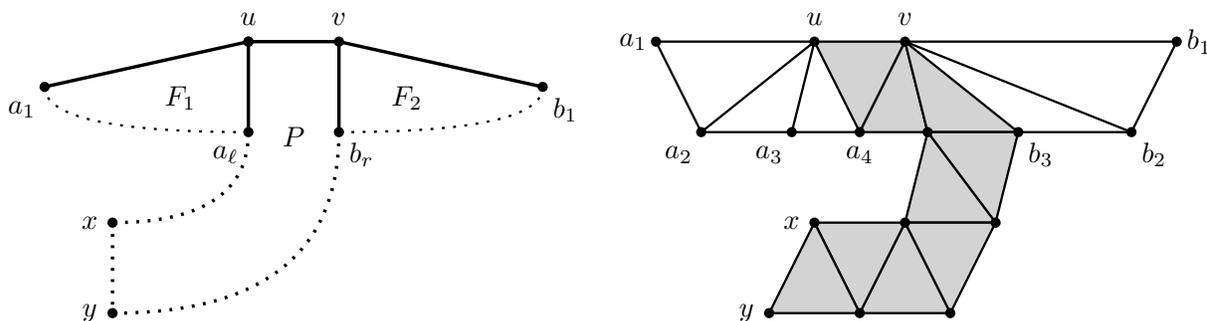


Figure 3.1: An umbrella as the union of three outerplanar paths (left) and an example of an umbrella with the handle shaded (right).

between P and F_1 . We define the r neighbors b_1, b_2, \dots, b_r of v in fan F_2 similarly, with (v, b_r) as the common edge of P and F_2 . See Figure 3.1 for an example of this labeling.

We now introduce a special type of cutting edge and cut-component, both of which will be used to partition maximal outerplanar graphs in many of the results in this thesis. This partitioning depends on the location of a given *root-edge*, which is any non-cutting edge.

Definition 3.2. *Given a maximal outerplanar graph G with root-edge (u, v) and a maximal outerplanar subgraph U of G that contains (u, v) , an anchor edge (or just anchor) of U is any cutting edge of G that belongs to U but is not a cutting edge of U . For any such anchor edge, the cut-component that does not contain (u, v) is called a hanging subgraph of U .*

See also Figure 3.2. Given any maximal outerplanar graph G and a root-edge (u, v) of G , we can partition G into a collection of umbrellas in the following manner.

Definition 3.3. *Given an outerplanar graph G with non-cutting root-edge (u, v) , a rooted umbrella system \mathcal{U} on G is a collection of umbrellas that satisfy the following:*

1. *There exists one umbrella $U_0 \in G$, called the root umbrella, that contains all neighbors of u and v in G and has cap (u, v) .*
2. *If S_1, \dots, S_k are the hanging subgraphs of U_0 for some $k \geq 0$, then $\mathcal{U} = \{U_0\} \cup \mathcal{U}_1 \cup \dots \cup \mathcal{U}_k$, where \mathcal{U}_i is a rooted umbrella system on S_i whose root-edge is the anchor edge of U_0 that has S_i as its hanging subgraph. We call \mathcal{U}_i a hanging umbrella system of U_0 .*

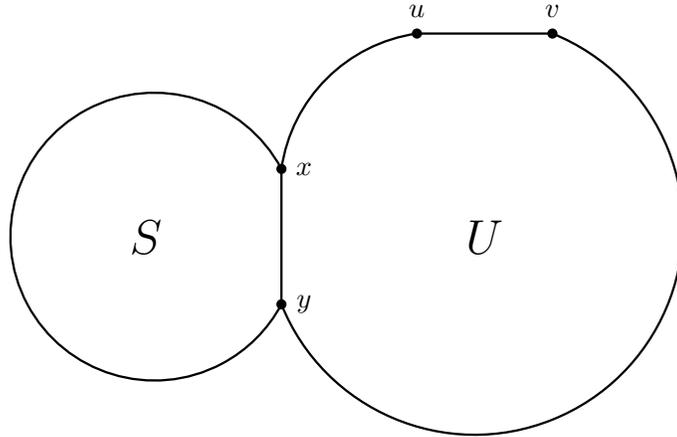


Figure 3.2: A graph with root-edge (u, v) , root umbrella U , anchor edge (x, y) , and hanging subgraph S .

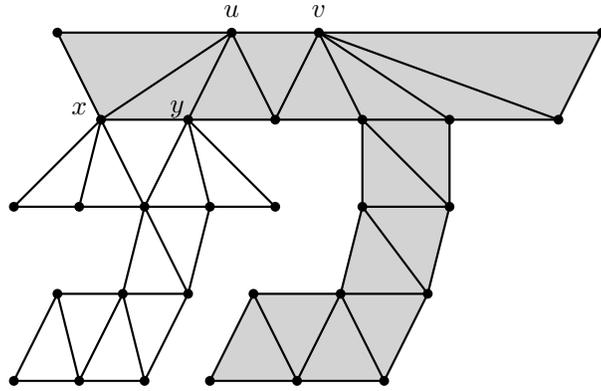


Figure 3.3: A rooted umbrella system of depth 2 with root-edge (u, v) and root umbrella shaded. Edge (x, y) is the root-edge of a hanging umbrella system of depth 1.

See also Figure 3.3. Given a rooted umbrella system \mathcal{U} with root umbrella U_0 , if $\mathcal{U} = \{U_0\}$, let the *depth* $d(\mathcal{U}) = 1$. Otherwise, if $\mathcal{U} = \{U_0\} \cup \mathcal{U}_1 \cup \dots \cup \mathcal{U}_k$, where $\mathcal{U}_1, \dots, \mathcal{U}_k$ are hanging umbrella systems of U_0 , the depth is

$$d(\mathcal{U}) = 1 + \max [d(\mathcal{U}_1), d(\mathcal{U}_2), \dots, d(\mathcal{U}_k)].$$

For any maximal outerplanar graph G and non-cutting edge $(u, v) \in G$, we say that the *rooted umbrella depth* with respect to (u, v) (denoted $ud^{rooted}(G, u, v)$, or just $ud(G)$ if the root-edge is clear from the context) is the minimum depth of all possible rooted umbrella systems on G with root-edge (u, v) . Note that normally G is required to have at least 3 vertices, but it will be convenient to define $ud(G) = 0$ if G is a single edge.

The following lemma will be helpful later.

Lemma 3.4. *Given a rooted umbrella system \mathcal{U} on a maximal outerplanar graph G with root-edge (u, v) , no hanging subgraph of the root umbrella U_0 contains u or v in its anchor.*

Proof. Assume that some hanging subgraph S of U_0 had anchor edge (u, x) , for some vertex $x \in G$. By definition, (u, x) is a cutting edge of G with S as a cut-component. By the properties of cut-components, u has at least one neighbor in S that is not in U_0 . This contradicts the fact that all neighbors of u must be part of the root umbrella U_0 . It follows that (u, x) cannot be the anchor of S , as desired. The argument is similar for vertex v . \square

3.2 Free vs. Rooted Umbrella Depth

For any rooted umbrella system on a maximal outerplanar graph G , the root-edge must be given. One can also consider a *free umbrella system* on G in which the root-edge for the entire graph can be any non-cutting edge of G . Let the *free umbrella depth* of G (denoted $ud^{free}(G)$) be the minimum depth of any free umbrella system on G . We will argue that the difference between the free umbrella depth and the rooted umbrella depth for a given root-edge is at most one. This difference is therefore small enough that we can ignore it for practical purposes, and in the remainder of this thesis, we will use the term *umbrella depth* in place of the rooted umbrella depth unless otherwise noted.

To prove this claim, we first make the following observation.

Lemma 3.5. *If U is an umbrella with cap (u, v) and S is a subgraph of U that is maximal outerplanar with at least 3 vertices, then S is an umbrella.*

Proof. Recall that an umbrella is the union of three outerplanar paths P , F_1 , and F_2 , where P is the handle, and F_1 and F_2 are the fans such that $u \in F_1$ and $v \in F_2$. If S is a subgraph of P , F_1 , or F_2 , then it is an outerplanar path, and therefore an umbrella by definition. Now let S_P , S_1 , and S_2 be the portions of S inside P , F_1 , and F_2 , respectively. Then there are three cases to consider.

1. If S_P , S_1 , and S_2 are all non-empty, then S must contain (u, v) , and is therefore an umbrella with cap (u, v) .
2. If only S_P and S_2 are non-empty, and $(u, v) \in S$, then S is an umbrella with cap (u, v) by definition. If $(u, v) \notin S$, then S must contain vertex v , since v is in both

S_P and S_2 . In this case, by using the end-edge of S_P that is incident to v and not shared with S_2 as the cap, we again see that S is an umbrella. See Figure 3.4 for an example in which edge (v, w) is the cap of S .

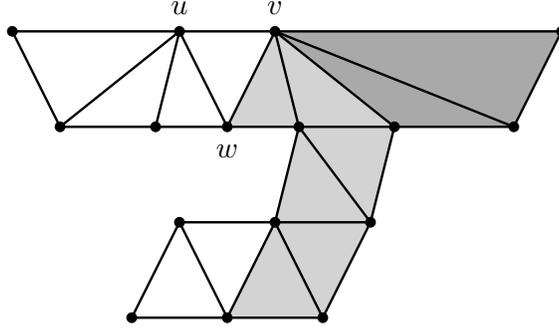


Figure 3.4: An umbrella with shaded subgraph $S = S_P \cup S_1$, where S_P is the lightly shaded region and S_1 is the darker region.

3. If S_P and S_1 are non-empty, then the argument is the same as case 2 above, except that S must contain u if $(u, v) \notin S$.

Note that S_1 and S_2 cannot be the only non-empty subgraphs, as then S would not be a connected graph. It follows that any subgraph of U is an umbrella, as desired. \square

Using Lemma 3.5, we can prove the following lemma, which states that the rooted umbrella depth of G cannot increase if we consider a subgraph of G .

Lemma 3.6. *If G is a maximal outerplanar graph with root-edge (u, v) , and (x, y) is a cutting edge of G , then*

$$ud^{rooted}(S_{x,y}, x, y) \leq ud^{rooted}(G, u, v)$$

where $S_{x,y}$ is the cut-component of (x, y) that does not contain edge (u, v) .

Note that Lemma 3.6 is not trivial because the root umbrella of any rooted umbrella system on G must include all neighbors of the root-edge, and therefore a change of root-edge may trigger changes for one or more umbrellas in hanging subgraphs.

Proof. We proceed by induction on the rooted umbrella depth of G . For the base case, let $H = ud(G) = 1$, which makes G an umbrella with cap (u, v) . By Lemma 3.5, $S_{x,y}$ is an umbrella, and therefore has umbrella depth 1, as desired.

For the inductive step, there are two cases to consider. In the first case, (x, y) is part of a hanging subgraph of U_0 (and possibly an anchor edge of U_0). This hanging subgraph by definition has umbrella depth at most $H - 1$. It follows by induction that the umbrella depth of $S_{x,y}$ is also at most $H - 1$.

In the second case, (x, y) does not belong to a hanging subgraph of U_0 , and is therefore a cutting edge of U_0 . As in Figure 3.5, consider the rooted umbrella system \mathcal{U}' on $S_{x,y}$ defined as follows. Let U_1 be the union of $(S_{x,y} \cap U_0)$ and the remaining neighbors of x and y in $S_{x,y}$. Since (x, y) is not an anchor edge of U_0 , $(S_{x,y} \cap U_0)$ is a maximal outerplanar graph with at least 3 vertices and a subgraph of U_0 . Lemma 3.5 tells us that it is also an umbrella, which implies that U_1 is an umbrella as well. We aim to make U_1 the root umbrella of a rooted umbrella system of depth at most H on $S_{x,y}$.

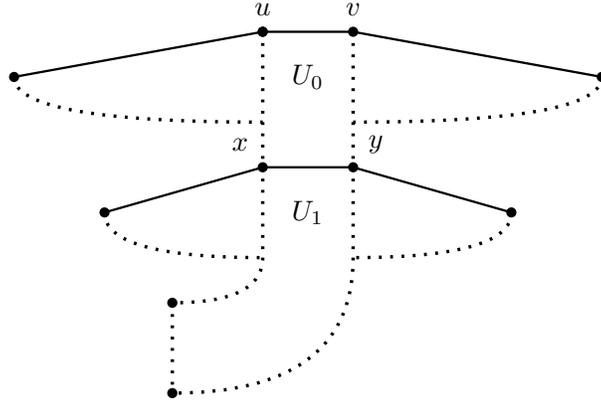


Figure 3.5: An example from the proof of Lemma 3.6 where $(x, y) \in U_0$ and $S_{x,y}$ is part of the handle of U_0 .

Consider any hanging subgraph S' of U_1 with anchor edge (x', y') . Since $(S_{x,y} \cap U_0) \subseteq U_1$, (x', y') is part of a hanging subgraph S of U_0 . Any hanging subgraph of U_0 has rooted umbrella depth at most $H - 1$, so by induction applied to S' and S with edge (x', y') , the subgraph S' has rooted umbrella depth at most $H - 1$, and therefore the rooted umbrella depth of $S_{x,y}$ is at most H , as desired. \square

We now proceed with the main result of this section, which relates the rooted umbrella depth to the free umbrella depth of a maximal outerplanar graph G .

Lemma 3.7. *Given a maximal outerplanar graph G , we have the following relationship between the free umbrella depth of G and the rooted umbrella depth of G :*

$$ud^{free}(G) = \min_{(u,v)} (ud^{rooted}(G, u, v)) \leq \max_{(u,v)} (ud^{rooted}(G, u, v)) \leq ud^{free}(G) + 1$$

where the minimum and maximum are taken over all non-cutting edges (u, v) of G .

Proof. Recall that the free umbrella depth of G is by definition the minimum rooted umbrella depth of G from any root-edge in G . The second inequality is obvious, so we will focus on the claim that $\max_{(u,v)} ud^{rooted}(G, u, v) \leq ud^{free}(G) + 1$. Let \mathcal{U} be a rooted umbrella system on G with depth $H = ud^{free}(G)$. Let U_0^* be the root umbrella of \mathcal{U} , and let (u^*, v^*) be the cap of U_0^* . We will show that there exists a rooted umbrella system on G for any non-cutting edge $(u, v) \neq (u^*, v^*)$ whose depth is at most $H + 1$.

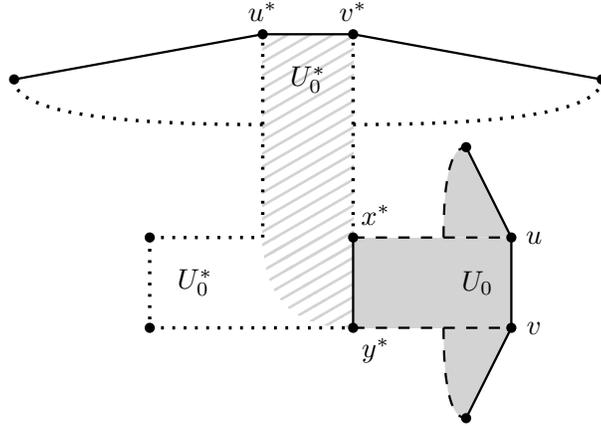


Figure 3.6: An example from the proof of Lemma 3.7 where the shaded umbrella is U_0 , the striped region is $U_0 \cap U_0^*$, and the unshaded regions are part of U_0^* .

Recall that root-edges are always non-cutting edges by definition. Therefore, by Observation 2.1, there exists an outerplanar path with (u, v) and (u^*, v^*) as end-edges. Let P be that path, and let U_0 be the umbrella defined by the union of P and all neighbors of u and v in G (see Figure 3.6). Let S be any hanging subgraph of U_0 , and let (x, y) be the anchor of S . Then (x, y) is a cutting edge of G , and its cut-component S does not contain U_0 , and in particular does not contain (u^*, v^*) . Therefore, by Lemma 3.6, its rooted umbrella depth is at most H . Therefore U_0 is the root umbrella of a rooted umbrella system of depth at most $H + 1$, as desired. \square

Chapter 4

From Umbrella Systems to Flat Visibility Representations

In this chapter, we show how to create a flat visibility representation given a maximal outerplanar graph G and a rooted umbrella system on G . The drawings we create will not be standard drawings of G (i.e., with all vertices on the outer face), as we will allow drawings of hanging subgraphs to be rotated and placed inside an inner face of the root umbrella. We will also only consider flat visibility representations in this section, as Theorem 2.3 can be used to transform any flat visibility representation into a straight-line drawing of the same height.

4.1 Drawing the Root Umbrella

In this section, we create a flat visibility representation for the root umbrella of a rooted umbrella system. In such a drawing, we would like the root-edge (u, v) to *span* the top layer, which means that u touches the top left corner of the drawing, and v touches the top right corner, or vice versa (see for example Figure 4.1).

Crucial to our construction is the following result, which will be used both for the base case and the induction step of the drawing for a rooted umbrella system in the following section.

Lemma 4.1. *Let U_0 be the root umbrella of a rooted umbrella system with root-edge (u, v) . Then there exists a flat visibility representation Γ of U_0 on three layers such that*

1. (u, v) spans the top layer of Γ .
2. Any anchor edge of U_0 is drawn horizontally in the middle or bottom layers.

The remainder of this section is dedicated to the proof of Lemma 4.1. Recall that the umbrella U_0 is the union of the handle P and two fans, where edge (u, v) is the cap of U_0 and, by definition, an end-edge of the outerplanar path P . Assume that u is an endpoint of P (the construction is similar, but flipped horizontally, if v is the endpoint). Let F_A be the fan of U_0 that contains u , and F_B the fan that contains v . As before, let a_1, a_2, \dots, a_ℓ be the ℓ neighbors of u in F_A , and let b_1, b_2, \dots, b_r be the r neighbors of v in F_B , with a_ℓ as the single vertex from F_A in P , and b_r is the single vertex from F_B in P . Now, either all vertices in P other than u and v are neighbors of u or v , or not. We will discuss each case in turn below.

Assume first that all vertices in U_0 are either u , v , or neighbors of the two. In this case, we can create a drawing of U_0 on two layers in which (u, v) spans the top layer (to have 3 layers, assume there is an empty layer in the middle). Let f_1, f_2, \dots, f_k be the k inner faces in the standard embedding of U_0 such that f_i shares an edge with f_{i-1} and f_{i+1} for all $i \in [2, k-1]$. Using two layers, assign to each face a square in order from left to right. Since edge (u, v) is a non-cutting edge of G , at most one face of U_0 can be adjacent to both u and v . Let f_i be that face, and assign u a flat box that extends from f_i to cover the entire top side of all faces from f_1 to f_{i-1} . Do the same for vertex v so that it covers the entire top side of all faces from f_{i+1} to f_k . Doing so defines a unique placement in the bottom layer for all other vertices in U_0 . See Figure 4.1 for an illustration of this construction. One can easily verify all conditions since (v, a_1) and (v, b_1) cannot be anchor edges by Lemma 3.4.

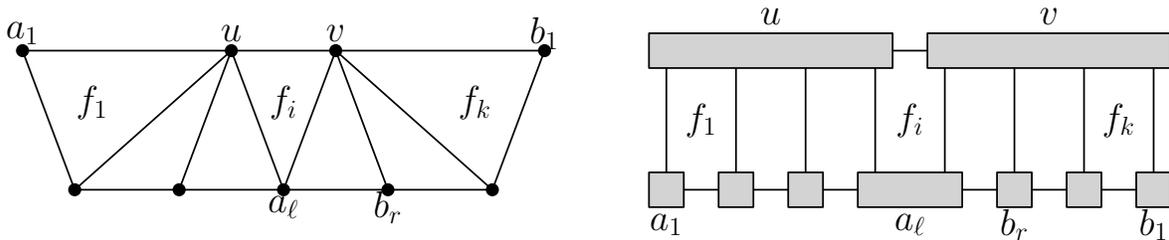


Figure 4.1: An umbrella U_0 for which all vertices are neighbors of u or v (left) and a construction for U_0 on two layers where edge (u, v) spans the top layer (right).

Now assume that some vertex of U_0 is neither u , v , or adjacent to them. This implies that there is some endpoint x of P that is not adjacent to u or v . By the definition of an

umbrella, at least one end-edge incident to x is not a cutting edge, and therefore not an anchor edge of U_0 . Let (x, y) be that edge. We explain how to draw the handle P in the following claim.

Claim 4.2. *There exists a flat visibility representation Γ of P on two layers that meets the following conditions (see also Figure 4.2).*

1. Edge (u, a_ℓ) is the vertical edge farthest to the left among all vertical edges in Γ .
2. Edge (x, y) is the vertical edge farthest to the right among all vertical edges in Γ .
3. Edge (u, v) is the horizontal edge that is farthest to the left in the top layer.
4. Edge (v, b_r) is a horizontal edge in the top layer.

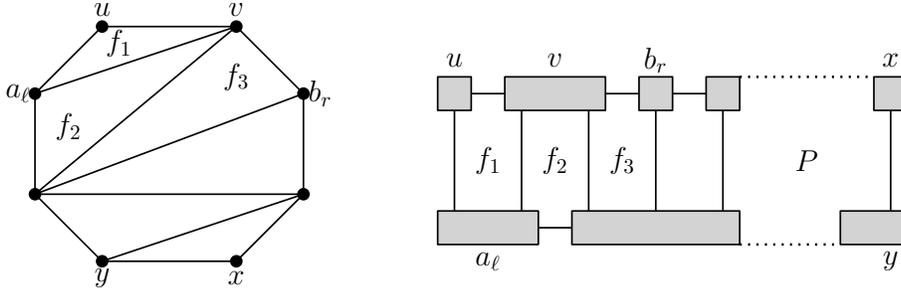


Figure 4.2: An example standard embedding of the handle P of an umbrella (left) and a drawing of P that satisfies Claim 4.2 (right).

Proof. Let f_1, f_2, \dots, f_k be the k inner faces in the standard embedding of P such that f_1 is incident to (u, v) , f_k is incident to (x, y) , and f_i shares an edge with f_{i-1} and f_{i+1} for all $i \in [2, k - 1]$. Using two layers, assign to each face a square in order from left to right. Now assign the vertices in P to boxes as follows.

- Assign u to the top left corner. Since u is an endpoint of P , it has degree 2 and gets drawn as a point (shown as a small square in Figure 4.3).
- Assign a_ℓ (which is a neighbor of u) to the bottom-left corner, and expand its segment horizontally to touch all squares of all faces that a_ℓ is incident to. With this, condition 1 is satisfied.

- If the clockwise (or counterclockwise) order of vertices around the outer face of P is $a_\ell, u, v, b_r, \dots, x, y, \dots$, then assign x to the top right corner. Otherwise, the order of vertices will be $a_\ell, u, v, b_r, \dots, y, x, \dots$. In this situation, assign x to the bottom right corner. For either case, place y on the opposite layer of x so that edge (x, y) is on the right side of f_k , thus satisfying condition 2. Vertex x is an endpoint of P , and is therefore drawn as a point in the same manner as vertex u . Note that $x \neq a_\ell$, since it is not adjacent to u , and therefore this does not contradict the earlier placement of a_ℓ . Vertex y is expanded horizontally to touch all squares of all faces that it is incident to. In the situation where $y = a_\ell$, its box occupies the entire bottom layer.
- Draw all other vertices as flat boxes that touch all squares of faces that the vertex is incident to. The choice between layers is done so that the order along the outer face of P in the standard embedding is respected. In particular, v will be to the right of u in the top layer, and b_r will be to the right of v . If $y = b_r$, then $u, v, y = b_r$ fill the top layer. We know that $x \neq b_r$, since x is not adjacent to v , and so the placement of v and b_r does not contradict the placement of x . With this, conditions 3 and 4 are satisfied.

□

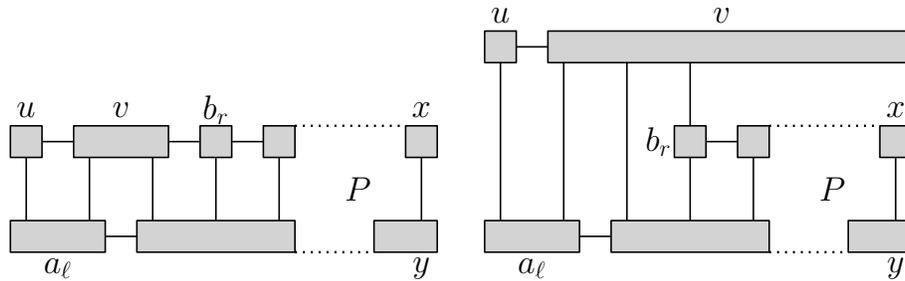


Figure 4.3: Releasing edge (u, v) in a flat visibility representation so that (u, v) spans the top layer.

Claim 4.2 gives us a drawing of the handle P with end-edge (u, v) in the top layer. We would like to have (u, v) span the top layer. In any drawing where (u, v) is in the top layer but does not span the top layer, we can *release* (u, v) as in [5] by adding a layer to the drawing and moving (u, v) so that it spans the new layer. To do this, first place (u, v) in the new layer, then, if u is to the left of v , extend u so it reaches the top left corner of Γ , and extend v so it reaches the top right corner. Do the opposite if u is to the right of v . For any neighbor w of u , w was connected to u by either a vertical or a horizontal line before u

was moved. In the former case, simply extend the existing line so it reaches u in the new top layer. In the latter case, replace the horizontal line with a vertical line connecting w and u . This is possible because Γ is a flat visibility representation, and therefore w and u must be in the same layer if they are connected by a horizontal edge. In a similar manner, reconnect v with all of its neighbors. See Figure 4.3 for an illustration of releasing an edge.

To complete the drawing of U_0 , we must add the vertices $a_1, \dots, a_{\ell-1}$ and b_1, \dots, b_{r-1} , which form fans with (u, a_ℓ) and (v, b_r) , respectively. By construction, u and a_ℓ are the leftmost vertices in their respective layers, and b_r has no immediate neighbor to its left after releasing edge (u, v) . Thus we can draw $a_1, \dots, a_{\ell-1}$ in order to the left of a_ℓ in the bottom layer of Γ , and b_1, \dots, b_{r-1} to the left of b_r in the middle layer. Expanding the box of u so that its left end aligns vertically with the left end of a_1 completes the construction of U_0 . See Figure 4.4 for an example of the final drawing.

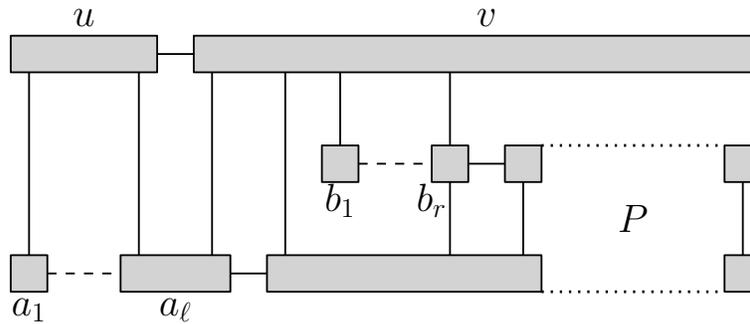


Figure 4.4: A flat visibility representation of a root umbrella on three layers.

We now argue that all anchor edges are horizontal. In the drawing of Claim 4.2, all vertical edges other than (x, y) are either cutting edges of P , or incident to u . Releasing (u, v) adds more vertical edges, but all of them are incident to u or v . Likewise, all vertical edges added when inserting the fans are incident to u or v .

Recall from Lemma 3.4 that no anchor edge (a, b) of U_0 can contain u or v . Also, (a, b) cannot be a cutting edge of U_0 since any cutting edge in a maximal outerplanar graph has at most two cut-components. Finally, $(a, b) \neq (x, y)$, since an anchor edge by definition is a cutting edge of G , and (x, y) was chosen to be a non-cutting edge. Therefore, any anchor edge of U_0 is drawn horizontally in the bottom two layers. This finished the proof of Lemma 4.1.

4.2 Drawing an Umbrella System

We now state the main result of this chapter. Its proof provides the algorithm for constructing a drawing of any maximal outerplanar graph.

Lemma 4.3. *Given a rooted umbrella system \mathcal{U} of depth H on a maximal outerplanar graph G , there exists a flat visibility representation Γ of G with height $2H + 1$ such that the root-edge spans the topmost layer of Γ .*

Proof. We prove this lemma by induction on the depth H of \mathcal{U} . For the base case, let $H = 1$. Here our rooted umbrella system consists of the single umbrella U_0 . By Lemma 4.1, we can draw U_0 on 3 layers, as desired.

For the inductive step, assume that our rooted umbrella system has depth H , with U_0 as the root umbrella. Let Γ_0 be the flat visibility representation for U_0 on three layers created with Lemma 4.1. Thus any anchor edge (a, b) in Γ_0 is drawn as a horizontal edge in the bottom two layers of Γ_0 .

Now add $2H - 2$ layers to Γ_0 between the middle and bottom layers. More precisely, if there are k hanging subgraphs S_1, S_2, \dots, S_k , then it suffices to add $\max_{1 \leq i \leq k} (\text{height of } S_i) - 1$ layers to Γ_0 . Each of the hanging umbrella systems of U_0 has depth at most $H - 1$, so by induction the hanging subgraphs can be drawn using at most $2H - 1$ layers with the anchor edge (a, b) spanning the top layer. Let Γ_1 be one such drawing of one such subgraph. For the merge step, we distinguish cases by the layer containing (a, b) .

1. If (a, b) is in the bottom layer of U_0 , then we can rotate (and reflect, if necessary) Γ_1 so that (a, b) is in the bottom layer and the left-to-right order of a and b in Γ_1 is the same as their left-to-right order in Γ_0 . This updated drawing of Γ_1 can then be inserted in the space between (a, b) in Γ_0 . This fits because Γ_1 has height at most $2H - 1$, and in the insertion process we can re-use the layer spanned by (a, b) .
2. If (a, b) is in the middle layer of U_0 , then we can reflect Γ_1 (if necessary) so that (a, b) has the same left-to-right order in Γ_1 as in Γ_0 . This updated drawing of Γ_1 can then be inserted in the space between (a, b) in Γ_0 .

Inserting all hanging subgraph drawings through one of the cases above completes the drawing. See Figure 4.5 for an example with inserted drawings highlighted. Since we added $2H - 2$ layers to a drawing of height 3, the total height of the final drawing is $2H + 1$, as desired. \square

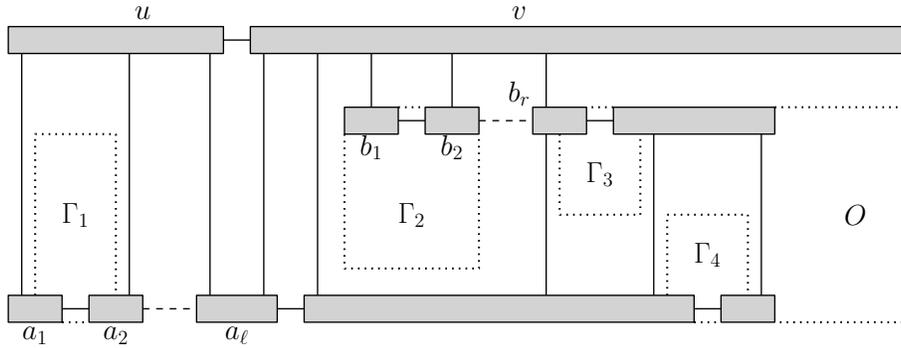


Figure 4.5: Inserting the drawings of hanging subgraphs into the flat visibility representation from Figure 4.4.

The following theorem summarizes the results of our construction, and provides a new upper bound for the optimal height of a maximal outerplanar graph. This bound will be compared to previous results in Chapter 7.

Theorem 4.4. *Any maximal outerplanar graph G has a planar flat visibility representation of height at most $2ud(G) + 1$.*

Chapter 5

From Drawings to Umbrella Systems

In this chapter, we will show that the height of a planar flat visibility representation of a maximal outerplanar graph can be lower-bounded by the depth of a rooted umbrella system. We will focus only on flat visibility representations in this section, as Theorems 2.2 and 2.3 can be used to convert from straight-line drawings to flat visibility representations and vice versa. We make no assumptions about the embedding of the graph induced by the visibility representation. In particular, we do not assume that the embedding must be the same as the standard embedding.

5.1 Left-free and Right-free Edges

We begin with the introduction of a few definitions and lemmas that will be needed in the lower bound argument. Let Γ be a flat visibility representation of a maximal outerplanar graph G , and let B_Γ be a minimum-height bounding box of Γ . A vertex $v \in G$ has a *left escape path* in Γ if there exists a polyline inside B_Γ from v to a point on the left side of B_Γ that is vertex-disjoint from Γ except at v , and for which all bends are on layers. We say that (ℓ_1, ℓ_2) is a *left-free edge* of Γ if it is vertical, and for every intersection point of (ℓ_1, ℓ_2) with a layer, the layer is empty to the left of that point. In particular, this implies that there is a left escape path from this intersection point by walking along the layer. Let *right escape paths* and *right-free edges* be defined symmetrically. See Figure 5.1 for an example.

Lastly, let a *dividing path* P in Γ be any polyline from the left side of the bounding box B_Γ to the right, and for which all bends are on layers. A dividing path P in B_Γ divides it into two disjoint regions, the *top* region which contains the top layer, and the *bottom*

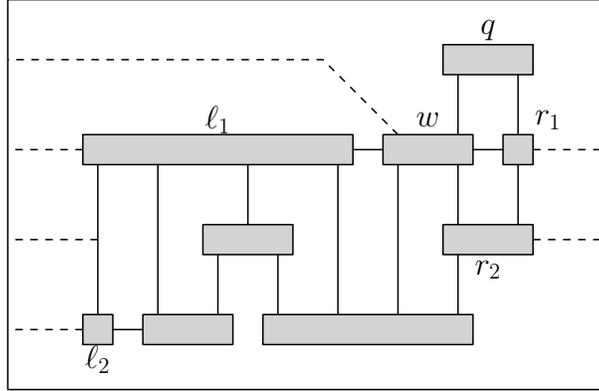


Figure 5.1: A flat visibility representation in which vertex w has a left escape path, r_1 has a right escape path, (ℓ_1, ℓ_2) is a left-free edge, and (r_1, r_2) is a right-free edge.

region that contains the bottom layer. We say that a subgraph of G that is vertex disjoint from P is *above* P if it is in the top region, and *below* P if it is in the bottom region. An example of a dividing path is the union of edge (ℓ_1, w) , edge (w, r_1) , the left escape path from ℓ_1 , and the right escape path from r_1 in Figure 5.1. In this example, vertex q would be above the dividing path, and r_2 would be below it. If P consists of the entire top layer in B_Γ , then the top region is empty. Similarly, the bottom region is empty if P consists of the entire bottom layer. We also have the following lemma.

Lemma 5.1. *Given a visibility representation Γ of a graph G with height H and dividing path P , any subgraph S of G that is vertex disjoint from P uses at most $H - 1$ layers in Γ .*

Proof. This follows from the definition of a dividing path. If the subgraph S is above the dividing path P , then S cannot touch the bottom layer of the drawing Γ without intersecting P . Similarly if S is below P , then S cannot intersect the top layer of Γ without intersecting P . \square

We now introduce the following lemma concerning the existence of left-free and right-free edges in flat visibility representations of maximal outerplanar graphs.

Lemma 5.2. *Any flat visibility representation Γ of a maximal outerplanar graph G with $n \geq 3$ vertices has at least one left-free edge and at least one right-free edge.*

Proof. We only provide the proof for the left-free edge, as the right-free case is symmetrical. Since G is maximal, and $n \geq 3$, it contains cycles, and therefore the height of Γ is at least

2. Furthermore, since Γ is flat, there must be at least one vertical edge in Γ . Consider the vertical edge (v_1, v_2) that is farthest to the left, breaking ties arbitrarily. We claim that (v_1, v_2) is a left-free edge in Γ . To see why, assume to the contrary that layer i is non-empty to the left of the intersection point ρ of (v_1, v_2) and layer i . This implies that there is either a vertical edge to the left of (v_1, v_2) that crosses layer i , or there are one or more vertices to the left of ρ on layer i . The former case is a contradiction to the choice of (v_1, v_2) . For the latter case, let v_ℓ be the leftmost vertex on layer i . Since all maximal outerplanar graphs are 2-connected and $n \geq 3$ by assumption, v_ℓ must have at least two neighbors. Furthermore, since v_ℓ is the leftmost vertex in its layer and the drawing is flat, at least one of its neighbors must lie on a different layer, and the edge to it must be vertical. But then there is a vertical edge in Γ that is farther to the left than (v_1, v_2) , which is a contradiction. It follows that (v_1, v_2) must be a left-free edge in the flat visibility representation Γ . \square

For the proof of the lower bound, we will want to create handles for umbrellas using outerplanar paths that have left-free and right-free edges as end-edges. This requires that the left-free and right-free edges are non-cutting edges, which isn't always the case in drawings of maximal outerplanar graphs. This motivates the following lemma, which allows us to convert any flat visibility representation of an outerplanar graph to another that contains non-cutting left-free and right-free edges.

Lemma 5.3. *Let Γ be a flat visibility representation of a maximal outerplanar graph G . We have the following.*

1. *Let (r_1, r_2) be a right-free edge of Γ , and let v_r be a vertex that has a right escape path. Then there exists a drawing Γ' in which v_r has a right escape path, (r_1, r_2) is a right-free edge, and there exists at least one left-free edge that is not a cutting edge of G .*
2. *Let (ℓ_1, ℓ_2) be a left-free edge of Γ , and let v_ℓ be a vertex that has a left escape path. Then there exists a drawing Γ' in which v_ℓ has a left escape path, (ℓ_1, ℓ_2) is a left-free edge, and there exists at least one right-free edge that is not a cutting edge of G .*

In either case, the height of Γ and y -coordinates of all vertices in Γ are unchanged in Γ' .

Proof. We prove the claim by induction on the number of vertices n in the maximal outerplanar graph G . We will show only the first claim, since the other is symmetric. For the base case, let $n = 3$, and note that G is a triangle without cutting edges. Therefore, by Lemma 5.2, Γ contains a left-free edge which is not a cutting edge.

For the induction step, let (ℓ_1, ℓ_2) be a left-free edge of Γ , which exists by Lemma 5.2. If (ℓ_1, ℓ_2) is not a cutting edge of G , we are done. Otherwise, (ℓ_1, ℓ_2) is a cutting edge of G . Let A and B be the cut-components of (ℓ_1, ℓ_2) such that $v_r \in A$.

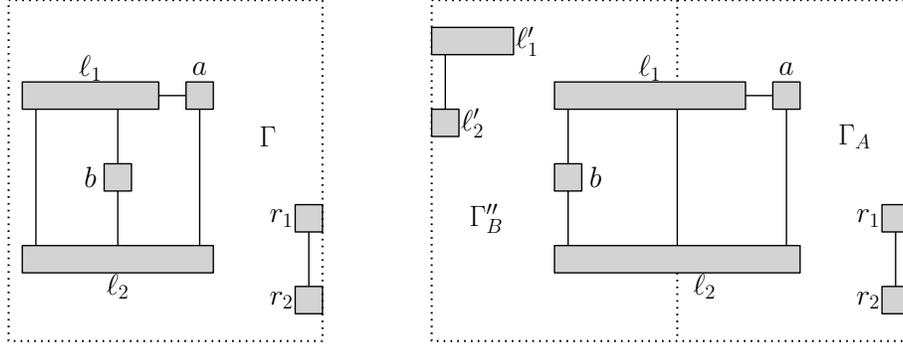


Figure 5.2: Expanding Γ' as part of the proof of Lemma 5.3

Let Γ_A be the drawing of A induced by Γ , and let Γ_B be the drawing of B induced by Γ . Note that (ℓ_1, ℓ_2) is a left-free edge for both Γ_A and Γ_B , and the height of Γ_A and Γ_B cannot exceed the height of Γ . Now let Γ'_B be the drawing of Γ_B reflected horizontally so (ℓ_1, ℓ_2) is now a right-free edge of Γ'_B . Observe that Γ'_B has at least one fewer vertex than Γ' , since B is a cut-component. By induction, we can create a drawing Γ''_B from Γ'_B in which (ℓ_1, ℓ_2) is a right-free edge and there is a left-free edge (ℓ'_1, ℓ'_2) that is not a cutting edge of B . Since the y-coordinates of ℓ_1 and ℓ_2 are the same in both Γ_A and Γ''_B , we can create a new drawing that places Γ''_B to the left of Γ_A and extends ℓ_1 and ℓ_2 to join the two copies. This is possible since (ℓ_1, ℓ_2) is left-free in Γ_A and right-free in Γ''_B . The drawing Γ_A is unchanged, so v_r will have the same right escape path in Γ' as in Γ , and Γ' will have right-free edge (r_1, r_2) and left-free edge (ℓ'_1, ℓ'_2) , as desired. See Figure 5.2 for an illustration of this drawing. \square

5.2 Lower Bound

We now prove the lower bound for the optimal height of drawings of maximal outerplanar graphs that meet certain conditions.

Lemma 5.4. *Let Γ be a flat visibility representation of a maximal outerplanar graph G with height H , and let (u, v) be a non-cutting edge of G . If there exists an escape path from u or v in Γ , then G has a rooted umbrella system with root-edge (u, v) and depth at most $H - 1$.*

The remainder of this section is dedicated to the proof of Lemma 5.4. We will proceed by induction on H .

Assume without loss of generality that there exists a right escape path from v (all other cases are symmetric). Using Lemma 5.3, we can modify Γ without increasing the height so that v has a right escape path, and there is a left-free edge in Γ that is a non-cutting edge of G . Let (ℓ_1, ℓ_2) be that edge, with the left escape path from ℓ_1 touching the bounding box of Γ above the left escape path from ℓ_2 . Define the outerplanar path P as the outerplanar path between (ℓ_1, ℓ_2) and (u, v) , which exists by Observation 2.1. Let U_0 be the union of P , the neighbors of u , and the neighbors of v . By definition, U_0 is an umbrella. Define Γ_U to be the drawing of U_0 induced by Γ , and Γ_U^* to be the standard planar embedding of U_0 in which all vertices are on the outer face. See also Figures 5.3 and 5.4.

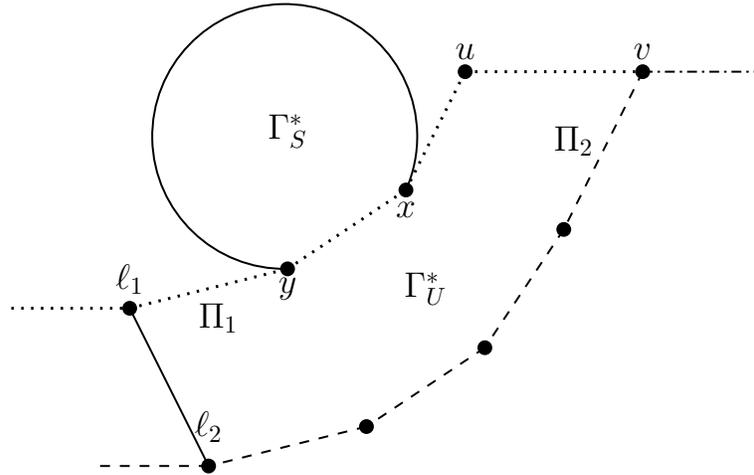


Figure 5.3: The standard embedding of the root-umbrella U_0 and one of its hanging subgraphs S . The dotted path represents the dividing path Π_1 while the dashed path represents Π_2 .

Now we define two dividing paths using U_0 as follows. Let P_1 be the vertex path in U_0 that starts from ℓ_1 , and continues along the outer face of Γ_U^* in the opposite direction of ℓ_2 until it reaches v . Similarly, let P_2 be the path that starts from ℓ_2 , and continues along the outer face of Γ_U^* in the opposite direction of ℓ_1 until it reaches v . Let Π_1 be the dividing path that contains P_1 , the escape path from v , and the escape path from ℓ_1 . Similarly, let Π_2 be the dividing path that contains P_2 , the escape path from v , and the escape path from ℓ_2 . Note that by definition, all vertices of $\Pi_1 - \{v\}$ are above Π_2 , and all vertices of $\Pi_2 - \{v\}$ are below Π_1 in Γ . In Figure 5.4, Π_1 is the path marked by a dotted line, while Π_2 is the path marked by a dashed line.

We now proceed with the induction. For the base case, let $H = 2$. We claim that in

this case, $G = U_0$, which gives the desired umbrella depth of 1. To see why, assume to the contrary that there exists a hanging subgraph S of U_0 with anchor edge (x, y) . Recall that a hanging subgraph of U_0 is any maximal outerplanar subgraph of G with at least 3 vertices that intersects U_0 only at the anchor edge (x, y) . We also know that the anchor edge must be a cutting edge of G . This rules out (ℓ_1, ℓ_2) as an anchor edge, as it was chosen to be a non-cutting edge of G . Furthermore, by Lemma 3.4, the anchor edge cannot contain vertex u or v . Thus S must be vertex-disjoint from one of Π_1 and Π_2 , and therefore by Lemma 5.1 it must be drawn with height 1. However, no maximal outerplanar graph with $n \geq 3$ can be drawn in one layer. It follows that S cannot exist, and thus $G = U_0$, as desired.

For the induction step, we once again let S be a hanging subgraph of U_0 with anchor (x, y) . As before, S is disjoint from Π_1 or Π_2 , and hence drawn with height at most $H - 1$. To apply induction, we must show that there exists a valid escape path from x or y that occupies the same $H - 1$ layers as S . The following cases cover all possible locations for (x, y) , because (x, y) is neither incident to v nor a cutting edge of U_0 .

Case 1: (x, y) belongs to $\Pi_1 - \{v\}$. After a possible renaming of x and y , we may assume that y is the vertex in (x, y) that is closer to vertex ℓ_1 in Π_1 . Let Γ_S be the drawing of S induced by Γ . Since Γ_S intersects $\Pi_1 - \{v\}$ at (x, y) , it is above the dividing path Π_2 . The subpath of Π_1 from y to ℓ_1 is also above Π_2 , and in combination with the left escape path from ℓ_1 is an escape path for Γ_S in the top $H - 1$ layers of Γ (see Figure 5.4). By induction, S has a rooted umbrella system of depth at most $H - 2$ for which (x, y) is the cap of the root umbrella.

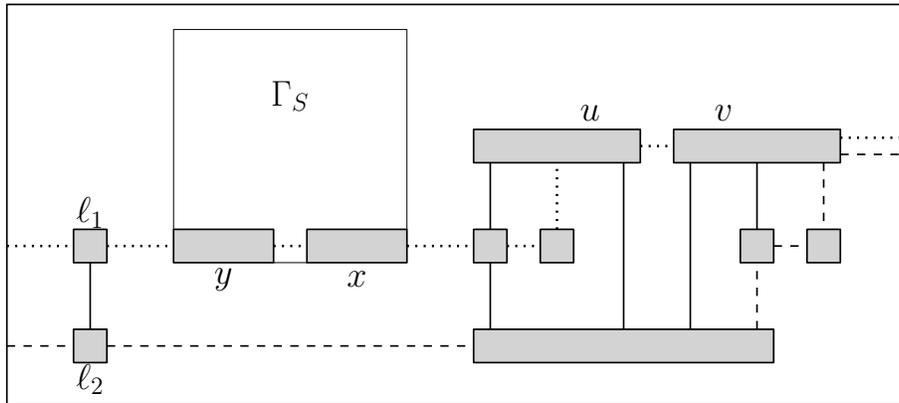


Figure 5.4: Extracting dividing paths from a flat visibility representation. The dotted path represents Π_1 while the dashed path represents Π_2 .

Case 2: (x, y) belongs to $\Pi_2 - \{v\}$. The argument is similar to Case 1, except we use Π_1 as our dividing path, with $\Pi_2 - \{v\}$ below it.

The cases above show that we can define a rooted umbrella system of depth at most $H - 2$ on any hanging subgraph S of U_0 . It follows that U_0 is the root umbrella of a rooted umbrella system with depth at most $H - 1$, as desired. This ends the proof of Lemma 5.4.

The following theorem summarizes the lower bound argument.

Theorem 5.5. *Given any flat visibility representation Γ of a maximal outerplanar graph G with height H , we can create a rooted umbrella system for G with depth at most $H - 1$.*

Proof. Using Lemma 5.3, we can convert Γ into a drawing Γ' in which some edge (u, v) is a non-cutting right-free edge. This implies that there is a right escape path from v , and it follows from Lemma 5.4 that we can define a rooted umbrella system on G with root-edge (u, v) and depth $H - 1$, as desired. \square

Recall that a 2-approximation algorithm for a minimization problem with optimal solution H is an algorithm that finds a solution that is always less than or equal to $2H$. By Theorem 4.4, we know that any maximal outerplanar graph G can be drawn with height at most $2ud(G) + 1$. Furthermore, Theorem 5.5 gives us a lower bound of $ud(G) + 1$ for the optimal height. It follows that our algorithm produces a drawing that is always less than twice the optimal height, and is therefore a 2-approximation by definition. This significantly improves on the result of [5], which was only a 4-approximation.

Chapter 6

Finding the Umbrella Depth

The algorithm in Chapter 4 requires that we have a rooted umbrella system of small depth. In this chapter, we introduce a dynamic programming algorithm for finding the rooted umbrella depth of a maximal outerplanar graph G with root-edge (u, v) . As always, we assume that the root-edge is a non-cutting edge of G . The runtime of our algorithm is linear in the number of edges m in G . Since $m = 2n - 3$ in any maximal outerplanar graph, this is an $O(n)$ algorithm. For each edge, we calculate the umbrella depth of a cut-component with that edge as the root-edge. The goal is to find the rooted umbrella depth of the entire graph with respect to the root-edge (u, v) .

Our algorithm performs a tree traversal on the rooted dual tree T of G , whose root is the vertex that corresponds to the inner face of G containing the root-edge (u, v) . For each cutting edge (a, b) in G , let $S_{a,b}$ be the cut-component of (a, b) that does not contain (u, v) . We use $ud(a, b)$ as a shorthand for $ud^{rooted}(S_{a,b}, a, b)$, and note that $ud(u, v) = ud(G)$.

We now introduce a collection of special umbrella types, each of which will be important for finding the umbrella depth $ud(a, b)$.

1. A *handle umbrella* with cap (a, b) is an outerplanar path P with end-edge (a, b) .
2. A *partial umbrella* of a is an umbrella with cap (a, b) in which b has exactly two neighbors.
3. A *partial umbrella* of b is an umbrella with cap (a, b) in which a has exactly two neighbors.
4. A *fan umbrella* of a is an umbrella with cap (a, b) in which all vertices other than a are neighbors of a .

5. A *fan umbrella* of b is an umbrella with cap (a, b) in which all vertices other than b are neighbors of b .

Informally speaking, a partial umbrella is an umbrella where one of the fans is missing, a handle umbrella is missing both fans, and a fan umbrella is missing one fan, and has a handle with only a single inner face. See Figure 6.1 for some examples of these umbrellas.

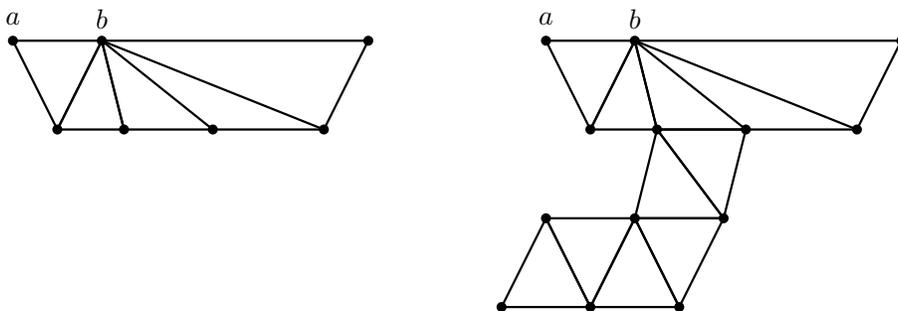


Figure 6.1: A fan umbrella of b (left) and a partial umbrella of b (right).

The umbrella depth $ud(a, b)$ can be derived from the handle umbrella depth, the partial umbrella depth, and the fan umbrella depth of edge (a, b) , which we will define in turn below. First we introduce some notation. Let c be the neighbor of a and b in $S_{a,b}$, which must exist since we only study maximal outerplanar graphs with at least 3 vertices. Now let d be the neighbor of both a and c that is not b , and let e be the neighbor of b and c that is not a . Note that d and e need not exist. This situation is illustrated in Figure 6.2. Using this notation, we discuss each intermediate value below.

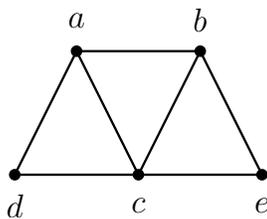


Figure 6.2: The labeling used for the formulas in the dynamic programming algorithm.

Definition 6.1. The handle umbrella depth of (a, b) , denoted $ud^{handle}(a, b)$, is defined as

$$ud^{handle}(a, b) = \min_{U_{a,b}^h} \left(\max_{S \subseteq S_{a,b} - U_{a,b}^h} ud(S) \right)$$

where $U_{a,b}^h$ is a handle umbrella with cap (a, b) , and S is a hanging subgraph of $U_{a,b}^h$ in $S_{a,b}$.

The following lemma describes how to calculate the handle umbrella depth using the notation from Figure 6.2.

Lemma 6.2. *The handle umbrella depth of (a, b) can be calculated as follows:*

$$ud^{handle}(a, b) = \min \left(\max [ud^{handle}(a, c), ud(b, c)], \max [ud^{handle}(b, c), ud(a, c)] \right).$$

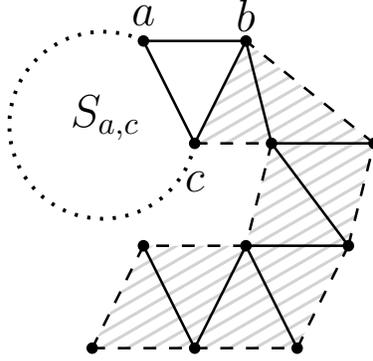


Figure 6.3: A handle umbrella $U_{a,b}^h$ with cap (a, b) . The shaded region is $U_{a,b}^h \cap U_{b,c}^h$, and the dashed lines indicate possible anchor edges.

Proof. Let $U_{a,b}^h$ be the handle umbrella in $S_{a,b}$ with cap (a, b) that achieves the minimum in Definition 6.1. This gives us the following equation for the handle umbrella depth

$$ud^{handle}(a, b) = \max_{S \subseteq S_{a,b} - U_{a,b}^h} ud(S).$$

By the definition of a handle umbrella, at least one of the edges (a, c) and (b, c) is not a cutting edge of $U_{a,b}^h$. Assume without loss of generality that this edge is (a, c) . Then $S_{a,c}$ either does not exist (if (a, c) is not a cutting edge of $S_{a,c}$) or $S_{a,c}$ is a hanging subgraph of $U_{a,b}^h$. In either case, the above equation is equivalent to the following

$$ud^{handle}(a, b) = \max \left(ud(a, c), \max_{S' \subseteq S_{b,c} - U_{b,c}^h} ud(S') \right)$$

where $U_{b,c}^h$ is the subgraph of $U_{a,b}^h$ in $S_{b,c}$. This implies that

$$ud^{handle}(a, b) \geq \max \left(ud(a, c), \min_{U_{b,c}^h} \left[\max_{S' \subseteq S_{b,c} - U_{b,c}^h} ud(S') \right] \right)$$

where the minimum goes over all possible handle umbrellas of $S_{b,c}$. It follows by Definition 6.1 that

$$ud^{handle}(a, b) \geq \max \left(ud(a, c), ud^{handle}(b, c) \right).$$

Equality is shown easily. If $U_{b,c}^h$ is the handle umbrella of $S_{b,c}$ that minimizes $ud^{handle}(b, c)$, then we can extend $U_{b,c}^h$ by adding a in order to obtain a handle umbrella of $S_{a,b}$ for which all hanging subgraphs are either $S_{a,c}$, or a hanging subgraph of $U_{b,c}^h$ (see also Figure 6.3). When (b, c) is an anchor edge of $U_{a,b}^h$, the situation is symmetric, and we get the following equality

$$ud^{handle}(a, b) = \max \left(ud(b, c), ud^{handle}(a, c) \right).$$

Combining the results from both cases gives us the desired formula. □

The next two values are the fan umbrella depths of both a and b , denoted $ud^{fan}(a; b)$ and $ud^{fan}(b; a)$, respectively. Note that we separate the a and b with a semi-colon here because the order of the vertices is important.

Definition 6.3. *The fan umbrella depth of a is defined as*

$$ud^{fan}(a; b) = \max_{\substack{S \subseteq S_{a,b} - U_{a,b}^f \\ b \notin S}} ud(S)$$

where $U_{a,b}^f$ is the unique fan umbrella with cap (a, b) that contains all neighbors of a in $S_{a,b}$, and S is any hanging subgraph of $U_{a,b}^f$ in $S_{a,b}$ that does not contain vertex b . The fan umbrella depth of b is defined symmetrically, and denoted $ud^{fan}(b; a)$.

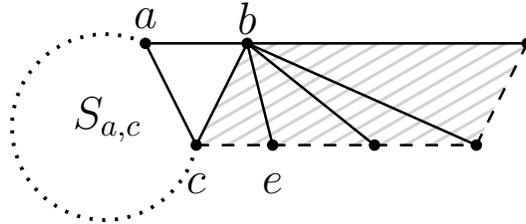


Figure 6.4: A fan umbrella $U_{b,a}^f$ of b . The shaded region is $U_{b,a}^f \cap U_{b,c}^f$, and the dashed lines indicate possible anchor edges.

The following lemma describes how to calculate the fan umbrella depth using the notation from Figure 6.2.

Lemma 6.4. *The fan umbrella depth of a is 0 if there are no neighbors of a and c other than b in $S_{a,b}$. Otherwise, if vertex $d \neq b$ is adjacent to a and c , we have*

$$ud^{fan}(a; b) = \max [ud^{fan}(a; c), ud(d, c)].$$

The fan umbrella depth of b is 0 if there are no neighbors of b and c other than a in $S_{a,b}$. Otherwise, if vertex $e \neq a$ is adjacent to b and c , we have

$$ud^{fan}(b; a) = \max [ud^{fan}(b; c), ud(c, e)].$$

Proof. We will consider only the fan umbrella depth for b , as the argument for a is similar. Let $U_{b,a}^f$ be the fan umbrella for b in $S_{a,b}$, and let $U_{b,c}^f$ be the fan umbrella for b in $S_{b,c}$. Then all neighbors of b in $U_{b,a}^f$ are shared with $U_{b,c}^f$ except for a , and thus $U_{b,c}^f \subset U_{b,a}^f$ (see also Figure 6.4). By Definition 6.3, the fan umbrella depth of $U_{b,c}^f$ is the maximum of the umbrella depth for all hanging subgraphs in $U_{b,a}^f$ except for those with anchor edge (c, e) or (a, c) . The desired formula follows. \square

Lastly, we need to find the *partial umbrella depth* of both a and b , denoted $ud^{partial}(a; b)$ and $ud^{partial}(b; a)$, respectively.

Definition 6.5. *The partial umbrella depth of a is defined as*

$$ud^{partial}(a; b) = \min \left(\max_{\substack{S \subseteq S_{a,b} - U_{a,b}^p \\ b \notin S}} ud(S) \right)$$

where $U_{a,b}^p$ is a partial umbrella of a with cap (a, b) that contains all neighbors of a in $S_{a,b}$, and S is any hanging subgraph of $U_{a,b}^p$ in $S_{a,b}$ that does not contain vertex b . The partial umbrella depth for b is defined symmetrically, and denoted $ud^{partial}(b; a)$.

The following lemma describes how to calculate the partial umbrella depth using the notation from Figure 6.2.

Lemma 6.6. *The partial umbrella depth of a is 0 if there are no neighbors of a and c other than b in $S_{a,b}$. Otherwise, if vertex $d \neq b$ is adjacent to a and c , we have*

$$ud^{partial}(a; b) = \min \left(\max [ud^{partial}(a; c), ud(d, c)], \max [ud^{fan}(a; c), ud^{handle}(d, c)] \right).$$

The partial umbrella depth of b is 0 if there are no neighbors of b and c other than a in $S_{a,b}$. Otherwise, if vertex $e \neq a$ is adjacent to b and c , we have

$$ud^{partial}(b; a) = \min \left(\max [ud^{partial}(b; c), ud(c, e)], \max [ud^{fan}(b; c), ud^{handle}(c, e)] \right).$$

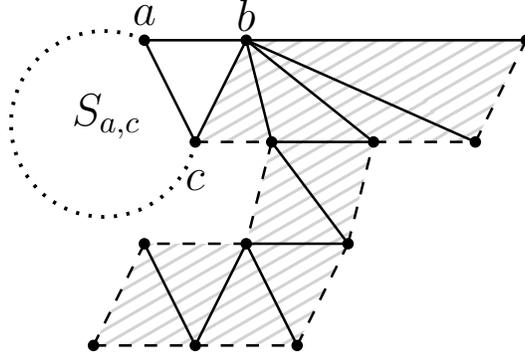


Figure 6.5: A partial umbrella $U_{b,a}^p$ of b . The shaded region is $U_{b,a}^p \cap U_{b,c}^p$, and the dashed lines indicate possible anchor edges.

The proof of this lemma is very similar to that of Lemma 6.2, and is left to the reader. Note the similarities between the situation illustrated in Figures 6.3 and 6.5. Lastly, using the fan and partial umbrella depths together, we can calculate the umbrella depth of $S_{a,b}$ as follows.

Lemma 6.7. *The umbrella depth of (a, b) can be calculated as follows:*

$$ud(a, b) = 1 + \min \left(\max [ud^{partial}(a; b), ud^{fan}(b; a)], \max [ud^{partial}(b; a), ud^{fan}(a; b)] \right).$$

Proof. Let U_0 be the root umbrella of a rooted umbrella system that achieves the umbrella depth $ud(a, b)$. Let P be the handle of U_0 and let F_A and F_B be the fans of U_0 such that $a \in F_A$ and $b \in F_B$. Note that P must have at least one of the edges (a, c) and (b, c) as a non-cutting edge. Assume without loss of generality that this edge is (a, c) , which implies that $(a, c) \in F_A$. We can define a partial umbrella $U_{a,b}^p$ of b and a fan umbrella $U_{b,a}^f$ of a which are subgraphs of U_0 such that $U_0 = U_{a,b}^p \cup U_{b,a}^f$. Both $U_{a,b}^p$ and $U_{b,a}^f$ include the inner face $\{a, b, c\}$ (see also Figure 6.6).

By Definition 6.5, we know that the partial umbrella depth of $U_{a,b}^p$ is equal to the maximum umbrella depth for any hanging subgraph of $U_{a,b}^p$ that does not have vertex b in its anchor edge. By Definition 6.3, we know that the fan umbrella depth of $U_{b,a}^f$ is equal to the maximum umbrella depth for any hanging subgraph of $U_{b,a}^f$ that does not have vertex a in its anchor edge. The hanging subgraphs of $U_{a,b}^p$ and $U_{b,a}^f$ coincide exactly with the hanging subgraphs of U_0 . Hence

$$ud(S_{a,b}) = 1 + \max \left(\max_{\substack{S \subseteq S_{a,b} - U_{a,b}^f \\ b \notin S}} ud(S), \max_{\substack{S \subseteq S_{a,b} - U_{a,b}^p \\ a \notin S}} ud(S) \right).$$

It then follows from Definitions 6.3 and 6.5 that

$$ud(a, b) \geq 1 + \max \left(ud^{fan}(a, b), ud^{partial}(b, a) \right).$$

Equality is then easily proven as in Lemma 6.2. A similar argument can be made when (b, c) is an anchor edge of P , and the desired formula follows. \square

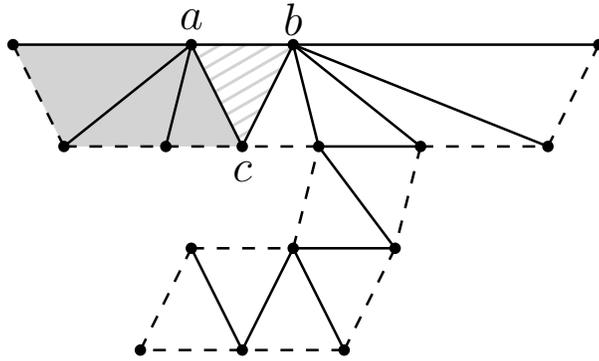


Figure 6.6: An umbrella as the union of a fan umbrella of a (shaded) and a partial umbrella of b (unshaded). The striped region belongs to both the fan and partial umbrellas, and the dashed lines indicate possible anchor edges.

We now discuss our dynamic programming algorithm, which is a standard bottom-up traversal in a tree. Given a maximal outerplanar graph G with dual tree T and root-edge (u, v) , let X be the set of all non-cutting edges of G with the exception of edge (u, v) . For each edge $(a, b) \in X$, set the handle umbrella depth, the fan umbrella depths, and the partial umbrella depths to zero. Note that the faces of G whose corresponding vertex in T is a leaf will have two such non-cutting edges and one cutting edge.

Now for each dynamic programming step, consider any face F in the standard embedding of G for which the umbrella depth has been computed for all but one edge. Such a face is guaranteed to exist at each step since the dual T of G is a tree. Let (a, b) be the edge on the face for which the umbrella depth has not been computed yet. Since we are traversing T from the leaves to the root vertex, we can use the formulas above to calculate all six values, including the umbrella depth, for (a, b) . Repeat this process until the umbrella depth for the root-edge (u, v) is found.

For every edge (a, b) in G , our algorithm computes a total of six values. Each of these values has a closed-form expression that can be computed in $O(1)$ time using a lookup

table. It follows that the runtime is linear in the number of edges in G , as desired. This result is summarized in the following theorem.

Theorem 6.8. *Given a non-cutting edge (u, v) , there exists an $O(n)$ algorithm for finding the rooted umbrella depth of a maximal outerplanar graph G with n vertices and root-edge (u, v) .*

Note that our algorithm computes the rooted umbrella depth for G , since the root-edge (u, v) must be given. One way to instead find the free umbrella depth is to repeat the process described above for every choice of root-edge in G . This would give an $O(n^2)$ algorithm for finding the free umbrella depth. One could likely compute the free umbrella depth in $O(n)$ time by initializing $ud(a, b)$ at all leaves of the (unrooted) dual tree, and then updating at the face where the resulting umbrella depth is minimized. However, by Lemma 3.7, the free umbrella depth is at most one more than the rooted umbrella depth, and therefore it does not seem worth the minor improvement to pursue this line of research.

Chapter 7

Comparison with Pathwidth

In this section, we will compare our results concerning the height of layered drawings of maximal outerplanar graphs to the current state of the art. The most recent bounds on the height come from Theorem 1.6, which establishes a 4-approximation based on the pathwidth of the dual tree in 2-connected outerplanar graphs. We therefore want to compare the umbrella depth of a maximal outerplanar graph to the pathwidth of the dual tree. We will do the same for the so-called rooted pathwidth, which has also been used a parameter for drawings of maximal outerplanar graphs.

7.1 Definitions

We say that a graph G has *pathwidth* at most k if there exists an ordering of the vertices v_1, v_2, \dots, v_n of G such that, for any $j \geq k$, there are at most k vertices in $\{v_1, v_2, \dots, v_j\}$ that are adjacent to vertices in $\{v_{j+1}, v_{j+2}, \dots, v_n\}$. By this definition, a graph with pathwidth 0 has no edge, while a graph with pathwidth 1 is a *caterpillar*, i.e., a tree in which deleting all leaves results in a path. For general trees, the pathwidth can also be described recursively using the idea of *main paths* that was introduced by Suderman [23].

Lemma 7.1. (Based on [23]) *The pathwidth of a tree T , denoted $pw(T)$, satisfies the following.*

1. $pw(T) = 0$ if T is a single vertex.
2. $pw(T) = 1 + \min_P(\max_{T'} pw(T'))$, where P is a path in T and T' is a component of $T - P$.

Any path where the minimum is achieved is called a main path.

Restating Lemma 7.1, for any tree T with pathwidth ρ and main path P , all components of $T - P$ have pathwidth at most $\rho - 1$. We may further assume that every main path ends at leaves of T , otherwise it could be extended to leaves of T while remaining a main path. When considering maximal outerplanar graphs, we will only be interested in the pathwidth of the dual tree, as opposed the graphs themselves.

The algorithm in [5] and the upper bound in Theorem 1.6 are based on the pathwidth. This algorithm splits a graph G at the outerplanar path P whose dual is a main path, recursively draws each connected component in $G - P$, then merges them into a drawing of P . This process was the primary motivation for the definition of an umbrella. We wanted to extend the concept of an outerplanar path while maintaining a subgraph small enough to be drawn on 3 layers.

In this section, we will also consider the *rooted pathwidth*, which was first defined by Biedl [7]. A *rooted tree* is a tree with a single vertex r called the *root*. In a rooted tree, a *root-to-leaf path* is a path in T from its root r to a leaf of T . We now have the following definition for the rooted pathwidth, which is very similar to the characterization in Lemma 7.1.

Definition 7.2. *The rooted pathwidth of a rooted tree T , denoted $rpw(T)$, is defined as follows.*

1. $rpw(T) = 1$ if T is a root-to-leaf path.
2. $rpw(T) = 1 + \min_P(\max_{T'} rpw(T'))$, where P is a root-to-leaf path in T , and T' is a component of $T - P$.

Any path where the minimum is achieved in is called a rooted main path.

There is a strong relationship between the umbrella depth of a maximal outerplanar graph and the rooted pathwidth of its dual tree. Note that the dual of an umbrella is a root-to-leaf path plus two paths for the fans. As such, the umbrella depth is never more than the rooted pathwidth of the dual tree. This will be proved in more detail in Lemma 7.5. We will also need the following result from [7], which relates the pathwidth to the rooted pathwidth in a rooted tree.

Lemma 7.3. [7] *For any rooted tree T , we have*

$$pw(T) \leq rpw(T) \leq 2pw(T) + 1.$$

One can show (Biedl, private communication) that the bound on the height of outerplanar graph drawings in Theorem 1.6 can be described as follows.

Theorem 7.4. (Based on [5]) *Let G be a 2-connected outerplanar graph. For any choice of root-vertex in the dual tree T of G , there exists a flat visibility representation of G with height at most $2rpw(T)$.*

7.2 Umbrella Depth Inequalities

We now compare our graph parameter (umbrella depth) to the graph parameters used in Theorems 1.6 and 7.4 (pathwidth and rooted pathwidth, respectively). We begin with the rooted pathwidth in Lemma 7.5, and afterwards discuss the pathwidth in Lemma 7.6.

Lemma 7.5. *In any maximal outerplanar graph G with dual tree T , we have*

$$\frac{rpw(T)}{2} \leq ud(G) \leq rpw(T).$$

More precisely, for every non-cutting edge (u, v) of G , if we root T at the inner face adjacent to (u, v) , then

$$\frac{rpw(T)}{2} \leq ud^{rooted}(G; u, v) \leq rpw(T).$$

Proof. As before, we use the shorthand $ud(G)$ to represent $ud^{rooted}(G; u, v)$. We will show by induction on $ud(G)$ that $\frac{rpw(T)}{2} \leq ud(G)$, which is equivalent to $rpw(T) \leq 2 \cdot ud(G)$. Recall that an umbrella is the union of three outerplanar paths P , F_1 , and F_2 , where P is the handle, and F_1 and F_2 are the fans.

For the base case, let $H = 1$ be the umbrella depth of G , which makes G an umbrella. We show that the dual tree T of G has rooted pathwidth at most 2 when the dual tree P of the handle (which is a vertex path) is chosen as the first root-to-leaf path. In this case, the dual trees of F_1 and F_2 are the only components of $T - P$, and they each have rooted pathwidth at most 1 since they are paths with leaves as their roots.

For the induction step, let \mathcal{U} be a rooted umbrella system on G with depth H , and let U_0 be the root umbrella of \mathcal{U} with cap (u, v) , handle P , and fans F_1 and F_2 . Let P^* be the dual tree of P , and F_i^* for $i = 1, 2$ be the dual tree of F_i . Let S_1^* be the component of $T - P^*$ that contains F_1^* , and S_2^* be the component of $T - P^*$ that contains F_2^* . See Figure 7.1 for an example of this labeling. Now let P_1^* be the root-to-leaf path created by extending

F_1^* to a leaf of S_1^* . Recall that by definition and Lemma 3.6 any hanging subgraph of U_0 with its anchor edge in F_1 has umbrella depth at most $H - 1$. Since $F_1^* \subseteq P_1^*$, it follows by induction that the rooted pathwidth of any component of $S_1^* - P_1^*$ is at most $2H - 2$.

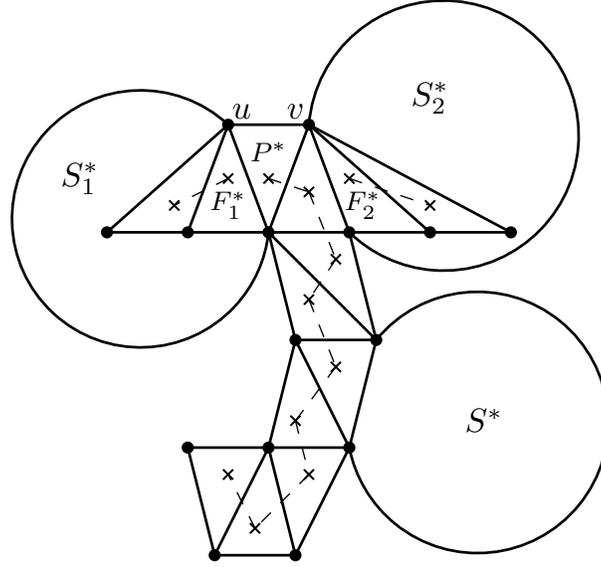


Figure 7.1: A root umbrella with the dual trees for the handle and both fans included. S_1^* , S_2^* , and S^* indicate components of $T - P^*$, where T is the dual tree of the entire graph.

Therefore, the rooted pathwidth of S_1^* is at most $2H - 1$. Similarly, one can show that the rooted pathwidth of S_2^* is at most $2H - 1$. Any other component S^* of $T - P^*$ corresponds to the dual of a hanging subgraph of U_0 , which has umbrella depth at most $H - 1$. By induction, the rooted pathwidth of S^* is at most $2H - 2$. It follows that the rooted pathwidth of T is at most $2H$, and the left inequality holds, as desired.

We will now show that $ud(G) \leq rpw(T)$ through induction on the rooted pathwidth. For the base case, let $H = 1$, which implies that T consists of a single root-to-leaf path. Thus G is an outerplanar path with (u, v) as an end-edge, and it follows by definition that G is an umbrella with cap (u, v) that has umbrella depth 1, as desired.

For the inductive step, let P^* be a rooted main path of T . By definition, P^* is a path from the face containing (u, v) to a leaf of T . We can define a root umbrella U_0 for G with cap (u, v) and the outerplanar path whose dual is P^* as the handle, plus all other neighbors of u and v in G . Now let S be any hanging subgraph of U_0 . Then the dual S^* of S is part of a subtree of $T - P^*$ that is rooted at the face adjacent to the anchor edge of S . Thus $rpw(S^*) \leq rpw(T) - 1 = H - 1$. By induction, we have $ud(S) \leq rpw(S^*) \leq H - 1$,

and it follows that the umbrella depth of G is at most H , as desired. \square

The following bounds for the pathwidth follow directly from Lemmas 7.3 and 7.5.

Corollary 7.6. *In any maximal outerplanar graph G with dual tree T , we have*

$$\frac{pw(T)}{2} \leq ud(G) \leq 2 \cdot pw(T) + 1.$$

7.3 Comparison of Bounds

In this section, we prove that the bounds introduced in this thesis are as good as the current state of the art (up to a constant term). Using Lemmas 7.5 and 7.6, we can compare our results directly to previous results that use the pathwidth and rooted pathwidth. First, for the rooted pathwidth, we have the following corollary of Theorem 4.4 and Lemma 7.5.

Corollary 7.7. *For any maximal outerplanar graph G with dual tree T , the construction in Chapter 4 gives a flat visibility representation of height at most $2rpw(T) + 1$.*

The bound established in Corollary 7.7 matches the one in Theorem 7.4, except for a constant term. A similar result for the pathwidth follows from Theorem 4.4 and Lemma 7.6, which match Theorem 1.6 up to a ‘+6’ term.

Corollary 7.8. *For any maximal outerplanar graph G with dual tree T , the construction in Chapter 4 gives a flat visibility representation of height at most $4pw(T) + 3$.*

We can also compare our result to the construction from [3], where the height is based on the number of vertices n in a maximal outerplanar graph. To this end, we need the following result.

Lemma 7.9. *[7] Any rooted tree T has $rpw(T) \leq \log(n + 1)$.*

Combining Lemma 7.9 with Corollary 7.7 gives us the following, which proves that our bounds are better than those established in Theorem 1.1 for all $n > 9$.

Corollary 7.10. *For any maximal outerplanar graph G with dual tree T , the construction in Chapter 4 gives a flat visibility representation of height at most $2 \log(n + 1) + 1$.*

In summary, we can say that our bounds are better than the bound of [3] (Theorem 1.1) for all $n > 9$, and that they match, up to a small constant term, the bound of [5] (Theorems 1.6 and 7.4). In the next section, we will see an example of a graph where our construction is strictly better than the one from [5].

7.4 Tightness of Bounds

In this section, we show that the bounds in Lemmas 7.5 and 7.6 are tight. We will accomplish this through the following three lemmas, each of which introduce a recursively defined family of maximal outerplanar graphs for which one or more of the bounds are tight.

Lemma 7.11. *For all $H \geq 1$, there exists a maximal outerplanar graph G_H with dual tree T_H for which the rooted umbrella depth of G_H is H for some choice of root-edge, and the rooted pathwidth of T_H is H if T_H is rooted at the face adjacent to the root-edge of G_H .*

Proof. For $H = 1$, let G_1 be defined as follows (see also Figure 7.2).

1. Start with a single edge (u, v) , which will be the root-edge of G_1 .
2. Add a vertex b that is adjacent to both u and v .
3. Add a vertex a that is adjacent to both u and b .
4. Add a vertex c that is adjacent to both a and b .

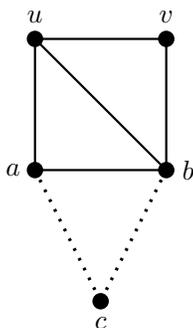


Figure 7.2: A drawing of G_1 from the proof of Lemma 7.11. The dotted lines indicate anchor edges for components of G_H when $H > 1$.

G_1 is an outerplanar path, and its dual tree, rooted at the face adjacent to (u, v) , is a root-to-leaf path. Therefore G_1 has umbrella depth 1, and T_1 has rooted pathwidth 1.

For G_H when $H > 1$, attach two copies of G_{H-1} to a copy of G_1 for such that (a, c) and (b, c) are the root-edges of the copies. Let P^* be any root-to-leaf path in T_H . Since there are two copies of G_{H-1} in G_H , one component of $T_H - P^*$ must contain a copy of

T_{H-1} . Since the rooted pathwidth of T_{H-1} is $H - 1$, it follows that the rooted pathwidth of T_H is at least H . One can easily show equality using the copy of T_1 that contains the root vertex of T_H (extended to a leaf) as the rooted main path.

Now we show that the $ud(G_H) \geq H$, which proves the claim since $ud(G_H) \leq rpw(T_H) = H$ by Lemma 7.5. Clearly, this holds for G_1 , which is itself an umbrella. Now consider an arbitrary umbrella system of G_H with root-edge (u, v) . The root-umbrella U_0 can have at most one of the edges (a, c) and (b, c) as a cutting edge. The other edge is therefore an anchor edge, and thus a copy of G_{H-1} is a hanging subgraph of U_0 . Since $ud(G_{H-1}) \geq H - 1$ by induction, this implies that $ud(G_H) \geq H$, as desired. \square

Lemma 7.11 defines a family of graphs G_H with dual trees T_H for which $rpw(T_H) = ud(G_H)$, and therefore the upper bound in Lemma 7.5 is tight. The following lemma defines a different family for which $ud(G_H) = 2 \cdot pw(T_H)$.

Lemma 7.12. *For all $H \geq 1$, there exists a maximal outerplanar graph G_H with dual tree T_H for which the pathwidth of T_H is H , while the rooted umbrella depth of G_H is $2H$ for some choice of root-edge.*

Proof. For $H = 1$, let G_1 be defined as follows (see also Figure 7.3).

1. Start with a single edge (u, v) , which will be the root-edge of G_1 .
2. Add a vertex x that is adjacent to both u and v .
3. Add a vertex y_ℓ that is adjacent to both u and w .
4. Add a vertex a_ℓ that is adjacent to both y_ℓ and w .
5. Add a vertex b_ℓ that is adjacent to both y_ℓ and a_ℓ .
6. Add a vertex c_ℓ that is adjacent to both a_ℓ and b_ℓ .
7. Add a vertex y_r that is adjacent to both v and x .
8. Add a vertex a_r that is adjacent to both x and y_r .
9. Add a vertex b_r that is adjacent to both a_r and y_r .
10. Add a vertex c_r that is adjacent to both a_r and b_r .

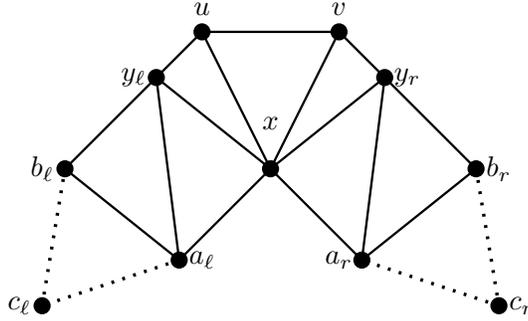


Figure 7.3: A drawing of G_1 from the proof of Lemma 7.12. The dotted lines indicate anchor edges for components of G_H when $H > 1$.

G_1 is an outerplanar path, and therefore T_1 has pathwidth one. For the umbrella depth, note that any handle P of an umbrella U_0 with cap (u, v) can only contain one of the endpoints of the outerplanar path G_1 . This endpoint is not in either fan of U_0 by construction. Thus the umbrella depth of G_1 cannot be 1. One can verify that it has umbrella depth 2.

For G_H when $H > 1$, start with a copy of the outerplanar path G_1 , then use each of its four end-edges as anchors, and attach four copies of G_{H-1} as hanging subgraphs. Clearly, T_H has pathwidth at most H by using the copy of T_1 that contains the root-edge of T_H as a main path. Equality can be shown since the pathwidth of T_{H-1} is $H - 1$, and T_H contains four copies of T_{H-1} .

For the umbrella depth of G_H , consider any umbrella U_0 with cap (u, v) , and let P_0 be its handle. (x, y_l) and (x, y_r) cannot both be cutting edges of U_0 , so let one of them, say (x, y_r) , be the anchor of a hanging subgraph S of U_0 that contains two copies of G_{H-1} as hanging subgraphs. Now consider any root umbrella U_1 of S with cap (x, y_r) . U_1 can have at most one of (a_r, c_r) and (b_r, c_r) as a cutting edge, with the other being the anchor of a hanging subgraph of U_1 . This subgraph contains a copy of G_{H-1} , and therefore has umbrella depth at least $2H - 2$ by induction. It follows that the umbrella depth of G_H is at least $2H$. Equality is shown easily by covering G_1 with two umbrellas and recursing. \square

Lemma 7.12 proves that the upper bound in Lemma 7.6 is almost tight, leaving only an $O(1)$ gap. In our final lemma, we define a family of graphs G_H with dual trees T_H for which $pw(T_H) = 2 \cdot ud(G_H)$ and $rpw(T_H) = 2 \cdot ud(G_H)$, thus proving that the lower bounds in Lemma 7.5 and Lemma 7.6 are tight as well.

Lemma 7.13. *For all $H \geq 1$, there exists a maximal outerplanar graph G_H with dual tree T_H for which the rooted umbrella depth of G_H is H for some choice of root-edge, while the*

pathwidth and rooted pathwidth of T_H are $2H$, if T_H is rooted at the face adjacent to the root-edge of G_H .

Proof. For $H = 1$, let G_1 be defined as follows (see also Figure 7.4).

1. Start with a single edge (u, v) , which will be the root-edge of G_1 .
2. Let a_1, a_2, a_3, a_4 be four other neighbors of u labeled in counterclockwise order around u with a_4 also adjacent to v .
3. Let b_1, b_2, b_3, b_4 be four other neighbors of v labeled in clockwise order around v with b_4 also adjacent to a_4 .
4. Lastly, let c_1, c_2, c_3 be four other neighbors of b_4 labeled in counterclockwise order around b_4 with c_1 also adjacent to a_4 .

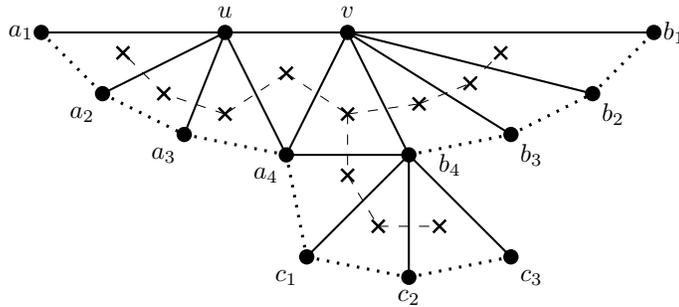


Figure 7.4: A drawing of G_1 from the proof of Lemma 7.13 with the dual tree included. The dotted lines indicate anchor edges for components of G_H when $H > 1$.

It is easy to see that G_1 is an umbrella, with the outerplanar path P between (u, v) and (b_4, c_3) as the handle, and the remaining inner faces as part of the two fans F_1 and F_2 . Furthermore, the dual tree of G_1 is not a caterpillar, so it cannot have pathwidth 1. One can verify that it has pathwidth 2, using the dual of the handle P as a main path.

For G_H when $H > 1$, start with a copy of G_1 , then attach nine copies of G_{H-1} such that each of them has one of the dotted edges in Figure 7.4 as its root-edge. This construction makes G_1 the root umbrella of a rooted umbrella system of depth H on G_H , so $ud(G_H) \leq H$.

We will show that $pw(T_H) \geq 2H$ by induction on H . For the base case when $H = 1$, we have already shown that the pathwidth of G_1 is 2. For the inductive step, consider any main path P_1^* of T_H for $H > 1$. T_H contains three copies of a path of length 3, each with

three copies of T_{H-1} attached, and P_1^* can contain at most two of these paths. Therefore, there must be at least one component S_1^* of $T - P_1^*$ that contains three copies of T_{H-1} . For any main path P_2^* of S_1^* , there must be at least one component S_2^* of $S_1^* - P_2^*$ that contains a copy of T_{H-1} . By induction, the pathwidth of T_{H-1} is at least $2H - 2$, and therefore the pathwidth of S_2^* is at least $2H - 2$. This makes the pathwidth of S_1^* at least $2H - 1$, and the pathwidth of T_H at least $2H$. One can show that equality holds by splitting the copy of T_1 that contains the root of T_H into two paths and recursing in all attached copies of T_{H-1} . Thus $pw(G_H) = 2H$, and $ud(G_H) \geq pw(G_H)/2 = H$, and it follows that $ud(G_H) = H$, as desired.

For the rooted pathwidth, we have $rpw(T_H) \geq pw(T_H) \geq 2H$ by Lemma 7.3. This is easily shown to be tight using the dual of the handle P in the copy of G_1 that contains the root-edge as the rooted main path of G_H . In each component S of $G_H - P$ that contains a fan of G_1 , let the dual of the fan be the rooted main path of S . Recursing on the remaining components gives the desired result. \square

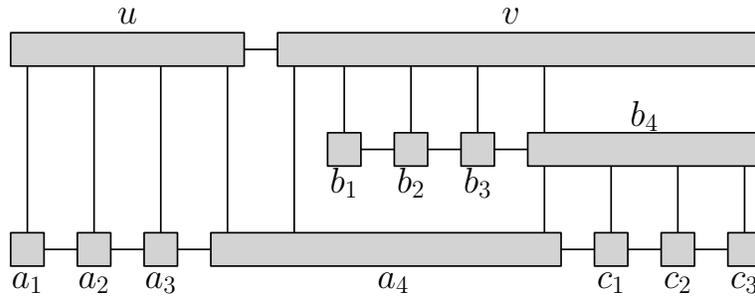


Figure 7.5: A flat visibility representation of the graph from Figure 7.4 created using the algorithm from Chapter 4.

Lemma 7.13 proves that the umbrella depth of the graph G_H in Figure 7.4 is strictly less than the pathwidth and the rooted pathwidth. One can also verify that the drawing produced by our algorithm has smaller height than the drawing from [5]. To see why, we include drawings of G_1 from both algorithms in Figure 7.5 and Figure 7.6. While we do not explain exactly how the algorithm from [5] works, we include some intermediate steps in Figure 7.6 to hint at how the final drawing was obtained.

Our drawing requires a total of 3 layers, while the drawing from [5] requires 4 layers. More generally, one can show that our drawing of G_H for $H > 1$ has height $2H + 1$ while the drawing from [5] has height $3H - 1$. Thus our algorithm is strictly better than the algorithm from [5] on some graphs.

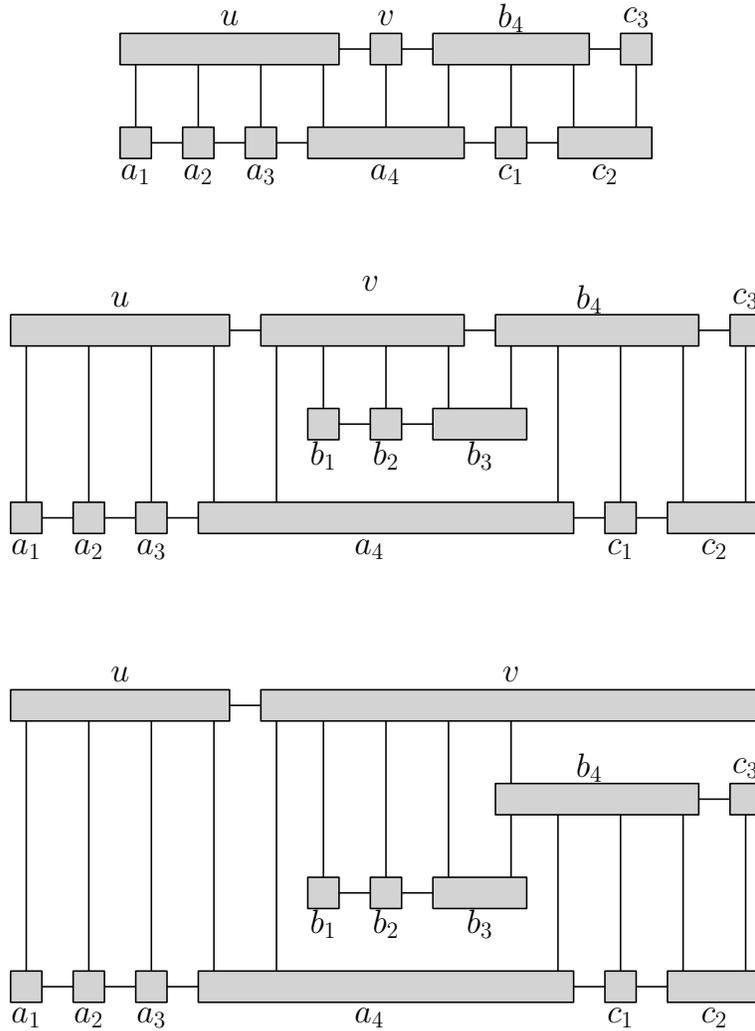


Figure 7.6: Drawing the dual of a main path (top), merging the recursively created subgraph (middle), and releasing edge (u, v) (bottom) as part of the algorithm from [5] (Theorem 1.6) for creating a flat visibility representation of the graph from Figure 7.4.

Chapter 8

Conclusions and Future Work

In this thesis we presented an algorithm for drawing maximal outerplanar graphs that is a 2-approximation for the optimal height. To this end, we introduced the umbrella depth as a new graph parameter for bounding the optimal height. Previous approximation algorithms were based on the pathwidth and rooted pathwidth of the dual tree. When compared to these approximation algorithms, we found that our bounds are never worse, and that there exist graphs for which our algorithm produces a flat visibility representation with height smaller than the drawing described in [5] (Theorem 1.6). Lastly, we showed that for all $n > 9$, our bound is better than the $O(\log n)$ bound established in [3] (Theorem 1.1).

There are still a number of problems that remain open.

- Our result only holds for maximal outerplanar graphs. Can the algorithm be modified so that it works for all outerplanar graphs?
- The algorithm from Chapter 4 creates a drawing that does not place all vertices on the outer face. Can we create an algorithm that minimizes or approximates the optimal height when the planar embedding is fixed?
- What is the width achieved by the algorithm from Chapter 4? Any visibility representation can be modified without changing the height so that the width is at most $m + n$, where m is the number of edges and n is the number of vertices [6]. Thus the width is $O(n)$, but what is the constant?
- Is it possible to determine the optimal height for maximal outerplanar graphs in polynomial time, or can we at least achieve a smaller approximation factor? Since the

pathwidth of outerplanar graphs can be approximated efficiently [8], the algorithm to find the optimal height in [12] becomes faster. Does this give us a pseudo-polynomial algorithm for the optimal height?

- The algorithm in Chapter 4 can be generalized to any system of subgraphs in which each subgraph can be drawn on three layers with (u, v) spanning the top layer. This includes umbrellas, but can be generalized to special umbrellas where the “handle” is an outerplanar path for which the cap (u, v) is not necessarily an end-edge. Perhaps this could even work for a maximal outerplanar graph whose dual is a tripod (i.e., a subdivision of $K_{1,3}$)? Could this be used to find better approximation factors? The bottleneck here is proving better lower bounds.
- Finally, are there approximation algorithms for the height or the area of drawings for other, more general planar graph classes? Of particular interest are planar 3-trees, which are graphs that, like maximal outerplanar graphs, naturally feature a tree-like description.

References

- [1] Muhammad Jawaherul Alam, Md. Abul Hassan Samee, Mashfiqui Rabbi, and Md. Saidur Rahman. Minimum-layer upward drawings of trees. *J. Graph Algorithms Appl.*, 14(2):245–267, 2010.
- [2] Johannes Batzill and Therese C. Biedl. Order-preserving drawings of trees with approximately optimal height (and small width). *CoRR*, abs/1606.02233, 2016.
- [3] Therese Biedl. Drawing outer-planar graphs in $O(n \log n)$ area. In *Graph Drawing, 10th International Symposium, GD 2002, Irvine, CA, USA, August 26-28, 2002*, pages 54–65, 2002.
- [4] Therese Biedl. Small drawings of outerplanar graphs, series-parallel graphs, and other planar graphs. *Discrete & Computational Geometry*, 45(1):141–160, 2011.
- [5] Therese Biedl. A 4-approximation for the height of drawing 2-connected outer-planar graphs. In *Approximation and Online Algorithms - 10th International Workshop, WAOA 2012, Ljubljana, Slovenia, September 13-14, 2012*, pages 272–285, 2012.
- [6] Therese Biedl. Height-preserving transformations of planar graph drawings. In *Graph Drawing - 22nd International Symposium, GD 2014, Würzburg, Germany, September 24-26, 2014*, pages 380–391, 2014.
- [7] Therese Biedl. Optimum-width upward drawings of trees I: rooted pathwidth. *CoRR*, abs/1502.02753, 2015.
- [8] Hans L. Bodlaender and Fedor V. Fomin. Approximation of pathwidth of outerplanar graphs. *J. Algorithms*, 43(2):190–200, 2002.
- [9] Giuseppe Di Battista, Peter Eades, Roberto Tamassia, and Ioannis G. Tollis. *Graph Drawing: Algorithms for the Visualization of Graphs*. Prentice-Hall, 1999.

- [10] Giuseppe Di Battista and Fabrizio Frati. Small area drawings of outerplanar graphs. *Algorithmica*, 54(1):25–53, 2009.
- [11] Giuseppe Di Battista and Fabrizio Frati. A survey on small-area planar graph drawing. *CoRR*, abs/1410.1006, 2014.
- [12] Vida Dujmovic, Michael R. Fellows, Matthew Kitching, Giuseppe Liotta, Catherine McCartin, Naomi Nishimura, Prabhakar Ragde, Frances A. Rosamond, Sue Whitesides, and David R. Wood. On the parameterized complexity of layered graph drawing. *Algorithmica*, 52(2):267–292, 2008.
- [13] Stefan Felsner, Giuseppe Liotta, and Stephen K. Wismath. Straight-line drawings on restricted integer grids in two and three dimensions. *J. Graph Algorithms Appl.*, 7(4):363–398, 2003.
- [14] Fabrizio Frati. Straight-line drawings of outerplanar graphs in $O(dn \log n)$ area. *Comput. Geom.*, 45(9):524–533, 2012.
- [15] Hubert de Fraysseix, János Pach, and Richard Pollack. Small sets supporting fary embeddings of planar graphs. In *Proceedings of the 20th Annual ACM Symposium on Theory of Computing, May 2-4, 1988, Chicago, Illinois, USA*, pages 426–433, 1988.
- [16] Hubert de Fraysseix, János Pach, and Richard Pollack. How to draw a planar graph on a grid. *Combinatorica*, 10(1):41–51, 1990.
- [17] Ashim Garg and Adrian Rusu. Area-efficient planar straight-line drawings of outerplanar graphs. *Discrete Applied Mathematics*, 155(9):1116–1140, 2007.
- [18] Lenwood S. Heath and Arnold L. Rosenberg. Laying out graphs using queues. *SIAM J. Comput.*, 21(5):927–958, 1992.
- [19] Marcus Krug and Dorothea Wagner. Minimizing the area for planar straight-line grid drawings. In *Graph Drawing, 15th International Symposium, GD 2007, Sydney, Australia, September 24-26, 2007*, pages 207–212, 2007.
- [20] Thomas Lengauer. *Combinatorial algorithms for integrated circuit layout*. Applicable theory in computer science. Teubner, 1990.
- [21] Debajyoti Mondal, Rahnuna Islam Nishat, Md. Saidur Rahman, and Muhammad Jawaherul Alam. Minimum-area drawings of plane 3-trees. *J. Graph Algorithms Appl.*, 15(2):177–204, 2011.

- [22] Walter Schnyder. Embedding planar graphs on the grid. In *Proceedings of the First Annual ACM-SIAM Symposium on Discrete Algorithms, 22-24 January 1990, San Francisco, California.*, pages 138–148, 1990.
- [23] Matthew Suderman. Pathwidth and layered drawings of trees. *Int. J. Comput. Geometry Appl.*, 14(3):203–225, 2004.
- [24] Kozo Sugiyama, Shojiro Tagawa, and Mitsuhiro Toda. Methods for visual understanding of hierarchical system structures. *IEEE Trans. Systems, Man, and Cybernetics*, 11(2):109–125, 1981.
- [25] Michael S Waterman and Jerrold R Griggs. Interval graphs and maps of DNA. *Bulletin of Mathematical Biology*, 48(2):189–195, 1986.