Upward Planar Drawing of Single Source Acyclic Digraphs* (Extended Abstract)

Michael D. Hutton†
University of Waterloo

Anna Lubiw‡
University of Waterloo

Abstract

An upward plane drawing of a directed acyclic graph is a plane drawing of the graph in which each directed edge is represented as a curve monotone increasing in the vertical direction. Thomassen [14] has given a non-algorithmic, graph-theoretic characterization of those directed graphs with a single source that admit an upward plane drawing. We present an efficient algorithm to test whether a given single-source acyclic digraph has an upward plane drawing and, if so, to find a representation of one such drawing.

The algorithm decomposes the graph into biconnected and triconnected components, and defines conditions for merging the components into an upward plane drawing of the original graph. To handle the triconnected components we provide a linear algorithm to test whether a given plane drawing admits an upward plane drawing with the same faces and outer face, which also gives a simpler, algorithmic proof of Thomassen's result. The entire testing algorithm (for general single-source directed acyclic graphs) operates in $O(n^2)$ time and $O(n)$ space.

1 Introduction

There are a wide range of results dealing with drawing, representing, or testing planarity of graphs. Fáry [4] showed that every planar graph can be drawn in the plane using only straight line segments for the edges. Tutte [15] showed that every 3-connected planar graph admits a convex straight-line drawing, where the facial cycles other than the unbounded face are all convex polygons. The first linear time algorithm for testing planarity of a graph was given by Hopcroft and Tarjan [6].

An *upward plane drawing* of a digraph is a plane drawing such that each directed arc is represented as a curve monotone increasing in the y-direction. In particular the graph must be a directed acyclic graph (DAG). A digraph is *upward planar* if it has an upward plane drawing. Consider the digraphs in Figure 1. By convention, the edges in the diagrams in this paper are directed upward unless specifically stated otherwise, and direction arrows are omitted unless necessary. The digraph on the left is upward planar: an upward plane drawing is given. The digraph on the right is not upward planar—though it is planar, since placing v inside the face f would eliminate crossings, at the cost of producing a downward edge. Kelly [10] and

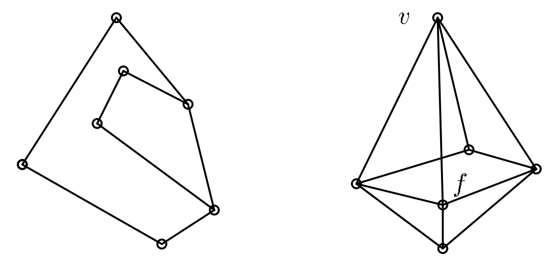

Figure 1: Upward planar and non-upward planar graphs.

Kelly and Rival [11] have shown that for every upward plane drawing there exists a *straight-line* upward plane drawing with the same faces and outer face, in which every edge is represented as a straight line segment. This is an analogue of Fáry's result for general planar graphs. The problem of recognizing upward planar digraphs is not known to be in P, nor known to be NP-hard. For the case of single-source single-sink digraphs there is a polynomial time recognition algorithm provided by Platt's result [12] that such a graph is upward planar iff the graph with a source-to-sink edge added is planar. An algorithm to find an upward plane drawing of such a graph was given DiBattista and Tamassia [2].

In this paper we will solve these problems for single-source digraphs. For the most part we will be concerned only with constructing an upward planar *representation*—enough combinatorial information to specify an upward plane drawing without giving actual numerical coordinates for the vertices. This notion will

*Supported in part by NSERC.

†Currently mdhutton@theory.toronto.edu.

‡Currently alubiw@water.waterloo.edu.

be made precise in Section 3. We will remark on the extension to a drawing algorithm in the Conclusions. Our main result is an $O(n^2)$ algorithm to test whether a given single-source digraph is upward planar, and if so, to give an upward planar representation for it. This result is based on a graph-theoretic result of Thomassen [14, Theorem 5.1]:

THEOREM 1.1. (THOMASSEN) *Let Γ be a plane drawing of a single-source digraph G. Then there exists an upward plane drawing Γ' strongly equivalent to (i.e. having the same faces and outer face as) Γ if and only if the source α of G is on the outer face of Γ, and for every cycle Σ in Γ, Σ has a vertex β which is not the tail of any directed edge inside or on Σ.*

The necessity of Thomassen's condition is clear: for a graph G with upward plane drawing Γ', and for any cycle Σ of Γ', the vertex of Σ with highest y-coordinate cannot be the tail of an edge of Σ, nor the tail of an edge whose head is inside Σ.

Thomassen notes that a 3-connected graph has a unique planar embedding (up to the choice of the outer face) and concludes that his theorem provides a "good characterization" of 3-connected upward planar graphs (i.e. puts the class of 3-connected upward planar graphs in NP intersect co-NP). An efficient algorithm is not given however, nor does Thomassen address the issue of non-3-connected graphs.

The problem thus decomposes into two main issues. The first is to describe Thomassen's result algorithmically; we do this in Section 4 with a linear time algorithm, which provides an alternative proof of his theorem. The second issue is to isolate the triconnected components of the input graph, and determine how to put the "pieces" back together after the embedding of each is complete. This more complex issue is treated in Section 5. The combined testing/embedding algorithm is left out due to space constraints; a full version is available in the first author's Masters Thesis [8], or in [9].

The algorithm for splitting the input into triconnected components and merging the embeddings of each operates in $O(n^2)$ time. Since a triconnected graph is uniquely embeddable in the plane up to the choice of the outer face, and the number of possible external faces of a planar graph is linear by Euler's formula, the overall time to test a given triconnected component is also $O(n^2)$, so the entire algorithm is quadratic.

2 Preliminaries

In addition to the definitions below we will use standard terminology and notation of Bondy and Murty [1].

A digraph G is *connected* if there exists an undirected path between any two vertices. For S a set of vertices, $G\backslash S$ denotes G with the vertices in S and all edges incident to vertices in S removed. If S contains a single vertex v we will use the notation $G\backslash v$ rather than $G\backslash\{v\}$. G is k-connected if the removal of at least k vertices is required to *disconnect* the graph. By Menger's Theorem [1] G is k-connected if and only if there exist k vertex-disjoint undirected paths between any two vertices. A set of vertices whose removal disconnects the graph is a *cut-set*. The terms *cut vertex* and *separation pair* apply to cut-sets of size one and two respectively. A graph which has no cut vertex is *biconnected* (2-connected). A graph with no separation pair is *triconnected* (3-connected). For G with cut vertex v, a *component* of G with respect to v is formed from a connected component H of $G\backslash v$ by adding to H the vertex v and all edges between v and H. For G with separation pair $\{u, v\}$, a *component* of G with respect to $\{u, v\}$ is formed from a connected component H of $G\backslash\{u, v\}$ by adding to H the vertices u, v and all edges between u, v and vertices of H. The edge (u, v), if it exists, forms a component by itself. An algorithm for finding triconnected components[1] in linear time is given in Hopcroft and Tarjan [7]. A related concept is that of *graph union.* We define $G_1 \cup G_2$, for components with "shared" vertices to be the *inclusive* union of all vertices and edges. That is, for v in both G_1 and G_2, the vertex v in $G_1 \cup G_2$ is adjacent to edges in each of the subgraphs G_1 and G_2.

Contracting an edge $e = (u, v)$ in G results in a graph, denoted G/e, with the edge e removed, and vertices u and v *identified.* Inserting new vertices within edges of G generates a *subdivision* of G. A *directed subdivision* of a digraph results from repeatedly adding a new vertex w to divide an edge (u, v) into (u, w) and (w, v). G_1 and G_2 are *homeomorphic* if both are subdivisions of some other graph. G is planar if and only if every subdivision of G is planar [1].

In a directed graph, the *in-degree* of a vertex v is the number of edges directed towards v, denoted $deg^- v$. Analogously the *out-degree* ($deg^+ v$) of v is the number of edges directed away from v. A vertex of in-degree 0 is a *source* in G, and a vertex of out-degree 0 is a *sink.*

Adopting some poset notation: we will write $u < v$ is there is a directed path $u \overset{*}{\to} v$. Vertices u and v are *comparable* if $u < v$ or $v < u$, and incomparable otherwise. If (u, v) is an edge of a digraph then u *dominates* v.

3 A Combinatorial View of Upward Planarity

As discussed by Edmonds and others (see [5]) a connected graph G is planar iff it has a *planar representa-*

[1]Note that Hopcroft and Tarjan's "components" include an extra (u, v) edge.

tion: a cyclic ordering of edges around each vertex such that the resulting set of *faces* F satisfies $2 = |F| - |E| + |V|$ (Euler's formula). A *face* is a cyclically ordered sequence of edges and vertices $v_0, e_0, v_1, e_1, \ldots, v_{k-1}, e_{k-1}$, where $k \geq 3$, such that for any $i = 0, \ldots, k-1$ the edges e_{i-1} (subscipt addition modulo k) and e_i are incident with the vertex v_i and consecutive in the cyclic edge ordering for v_i.

One method of combinatorially specifying an upward planar drawing is provided by the result of (independently) DiBattista and Tamassia [2], and Kelly [10] that a DAG G is upward planar iff edges can be added to obtain a *planar s-t graph*, defined to be a DAG which has a single source s, a single sink t, contains the edge (s, t), and is planar. DiBattista and Tamassia give an algorithm using $O(n \log n)$ arithmetic steps to find an upward plane drawing of a planar *s-t* graph. We will find it useful to have a slightly different notion for our special case:

Definition 3.1. *An* upward planar representation *of a single source DAG $G = (V, E)$ consists of a planar representation together with: a designated* outer *face; and a vertex ordering $1, 2, \ldots, n$ such that vertex 1 is on the outer face, and for each $i = 2, \ldots, n$ vertex i is a sink in G_i and is on the outer face of G_i. Here G_i is the subgraph induced on vertices $1, \ldots i$, and inherits its planar representation and outer face from those of G.*

Proposition 3.1. *A single-source acyclic digraph is upward planar iff it has an upward planar representation.*

Proof. ($\Rightarrow$) An upward plane drawing provides a planar representation and a distinguished outer face. Order the vertices by increasing y coordinate. It is easy to verify that the required conditions hold.

($\Leftarrow$) Let G be the digraph, and Γ be an upward planar representation of G. We will show how to add edges to augment G to a planar *s-t* graph G' and augment Γ to a planar representation of G'. Then by the result of DiBattista and Tamassia [2], G' has an upward plane drawing corresponding to the representation Γ'. Thus G has an upward plane drawing whose representation is Γ.

For each face f of Γ, let v_f be the vertex of f maximum in the ordering. Add edges (u, v_f) for each vertex $u \neq v_f$ in face f for which such an edge does not already exist. Call the result G'. Clearly G' is acyclic. G' is upward planar—we must just augment the planar representation for G to obtain a planar representation for G', and use the same vertex ordering. G' has a single source s, the vertex numbered 1, and the vertex t, numbered n, is a sink. The edge (s, t) has been added in G'. In order to prove that t is the only sink, we will prove that any other vertex v has some edge leaving it: If v is on the outer face of Γ then the edge (v, t) has been added to G'. Otherwise consider the minimum i such that v is not on the outer face of G_i. Note that $i > v$. Then v is in a face f containing vertex i, and the edge (v, i) has been added to G'. ∎

Note that this proof provides a simple linear-time algorithm to convert an upward planar representation of G to the set of edges which should be added to G to produce a planar *s-t* graph.

Two plane drawings are *equivalent* if they have the same representation—i.e. the same faces. Two plane drawings are *strongly equivalent* if they have the same representation and the same outer face.

In the remainder of this section we give some operations which preserve upward planarity. The first operation contracts an edge connected to a vertex of in- (out-) degree 1. The second attaches one upward planar graph to another at a single vertex. The third attaches an upward planar graph in place of an edge of another upward planar graph. The last splits a vertex into two vertices. These results can be proved using Proposition 3.1. Proofs can be found in [8] or [9].

Lemma 3.1. *Let G be a DAG and v, dominated by u, be a vertex of G with in-degree 1. Then, $G/(u, v)$ is upward planar if G is upward planar.*

Note that the same result holds for G and edge (u, v) with $deg^+ u = 1$ by symmetry.

Lemma 3.2. *Let G be an upward planar digraph with a vertex u, and let H be a digraph which has an upward planar representation with a source u' on the outer face. Let G' be the graph formed by identifying u and u' in $G \cup H$. Then G' is upward planar.*

Lemma 3.3. *Let G be an upward planar digraph with an edge (u, v), and H be a digraph which has an upward planar representation with a source u' and a sink v' on the outer face. Let G' be the graph formed by removing the (u, v) edge of G and adding H, identifying vertex u with u' and vertex v with v'. Then G' is upward planar.*

Lemma 3.4. *Let G be a DAG which has an upward planar representation where the cyclic edge order about vertex v is $e_1, \ldots, e_k$. Let G' be the DAG formed by splitting v into two vertices: v' incident with edges $e_1, \ldots, e_i$, and v'' incident with edges $e_{i+1}, \ldots, e_k$. Then G' is upward planar.*

4 Strongly-Equivalent Upward Planarity

Consider the following question: Given a single-source acyclic digraph G and a plane drawing Γ of G, does G admit an upward plane drawing strongly equivalent to Γ? Expressing this in terms of representations: Given a planar representation of G with some distinguished outer face, can we augment this to an upward planar representation of G—i.e. can we find a vertex ordering of G satisfying Definition 3.1?

We will rework Thomassen's condition in the form of a linear time algorithm to answer this question.

Define a *violating cycle* of G with respect to Γ to be a cycle Σ such that every vertex of Σ is the tail of an edge inside or on Σ. Our algorithm will find either a violating cycle of G—evidence that G does not have an upward plane drawing strongly equivalent to Γ—or a vertex ordering satisfying Definition 3.1—evidence (by Proposition 3.1) that G has an upward plane drawing strongly equivalent to Γ. The correctness proof for the algorithm will provide a new proof of Thomassen's theorem.

The algorithm is recursive, and the proof that it works is by induction. If there is a sink v on the outer face of Γ, give it the number n, and recurse on $G\backslash v$ with respect to the induced plane drawing $\Gamma\backslash v$. By induction we will find an upward planar representation of $G\backslash v$ augmenting $\Gamma\backslash v$, or a violating cycle for $G\backslash v$. In the first case we get the required ordering for G; and in the second case we get a violating cycle for G.

It remains to deal with the case when the outer face of Γ has no sink. We claim that in this case G has a violating cycle: If the outer face of Γ is a cycle then it is a violating cycle. If the outer face is a walk, then follow it starting at s, and let v be the first vertex which repeats. Vertex v must be a cut vertex. Consider the segment of the walk from v to v. If this segment contains only one other vertex, say u, then u is a sink, contradiction. Otherwise we obtain a cycle C from v to v. The two edges incident with v must be directed away from v, and thus C is a violating cycle.

5 Separation into Tri-Connected Components

The algorithm of Section 4 tests for upward planarity of a single-source DAG G starting from a given planar representation and outer face of G. In principle, we could apply this test to all planar representations of G, but this would take exponential time. In order to avoid this, we will decompose the graph into biconnected and then into triconnected components. Each triconnected component has a unique planar representation (see [1]), and only a linear number of possible outer faces. We can thus test upward planarity of the triconnected components in quadratic time using the algorithm of Section 4. Since we will perform the splitting and merging of triconnected components in quadratic time, the total time will then be quadratic.

To decompose G into biconnected components we use:

Lemma 5.1. *A DAG G with a single source s and a cut vertex v is upward planar iff each of the k components H_i of G (with respect to v) is upward planar.*

Proof. If G is upward planar then so are its subgraphs the H_i's. For the converse, note that if $v \neq s$ then v is the unique source in all but one of the H_i's; and if $v = s$ then v is the unique source in each H_i. Apply Lemma 3.2. ■

Dividing G into triconnected components is more complicated, because the cut-set vertices impose restrictive structure on the merged graph. In the biconnected case, it is sufficient to simply test each component separately, since biconnected components do not interact in the combined drawing; this is not the case for triconnected components, as illustrated by the two examples in Figure 2. (Recall our convention that direction arrow-heads are assumed to be "upward" unless otherwise specified.) In (a), the union of the graphs is upward planar, but adding the edge (u, v) to each makes the second component non-upward-planar. In (b), the graph is non-upward-planar, but each of the components is upward planar with (u, v) added.

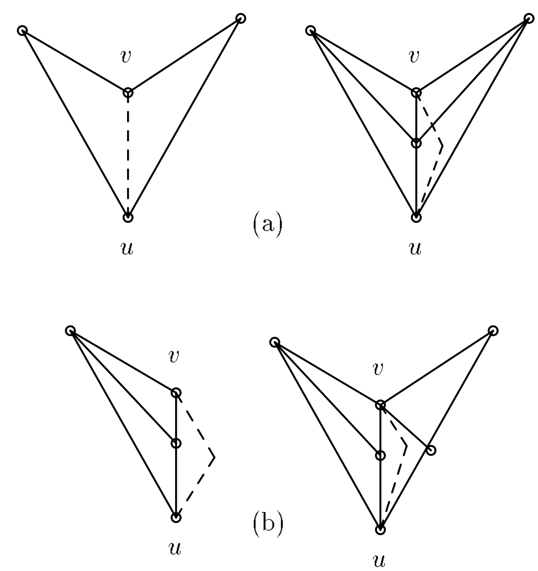

Figure 2: Added complication of 2-vertex cut-sets.

We will find it convenient, particularly for the case where the source s is in a separation pair, to split the graph into exactly two pieces at separation pairs.

There are two main issues. Firstly, we must identify which component will be the "outer" component, because this imposes restrictions on the other ("inner" component) to adapt to its facial structure (in order to be injected within a face). It will always be true that the inner component will have more restrictions upon its embedding, because it must fit within the prescribed face. Specifically, a list of vertices will be required to be on the outer face of any embedding to retain planarity in the merge. Secondly, we must be able to properly represent the facial structures of the two components to ensure that the recursively calculated embeddings can be merged without destroying upward planarity.

Our general subproblem instance consists of a biconnected graph G, and a set of vertices $X = \{x_i\} \subseteq V(G)$ which will be required to be on the outer face of any planar embedding of G. G is broken up into two components at a cut-set $\{u, v\}$, and recursive calls made. We will give conditions based on the type of cut-set involved as to whether upward plane drawings of the two components can be put back together into an upward plane drawing of the whole. These conditions are broken into three cases: where u and v are incomparable; where u and v are comparable with $u \neq s$ (i.e. $s < u < v$); and where u, v are comparable with $u = s$.

The conditions prescribed will be in the form of *markers* added to each component to represent the shape of the other component in the decomposition. If the graph were undirected, it would be sufficient to add a single edge between the cut vertices in each component, because the only requirement would be that the vertices share a face. We would also not need to require any vertices to be on the outer face, because any face can be made the outer face. This is not true for upward planarity. The type of markers needed will depend on the particular graph.

The markers are necessary for three reasons: firstly, to ensure that the original graph is upward planar iff the two components (with markers) are upward planar; secondly, to maintain biconnectedness; and thirdly, to maintain the single source property. The markers we will be interested in are shown in Figure 3.

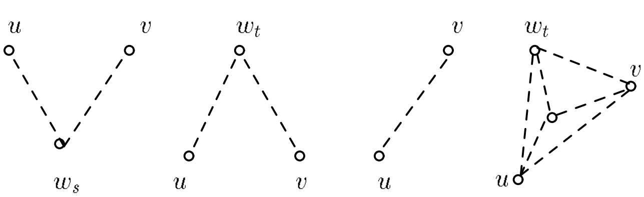

(a) M_s (b) M_t (c) M_{uv} (d) M_{uvt}

Figure 3: Marker Graphs.

An important note to make at this time is that the markers, except for M_{uv}, are subgraphs attached at only two vertices, which means that $\{u, v\}$ will still constitute a cut-set. For the purposes of determining cut-sets, and making recursive calls, the markers should be treated as distinguished edges—a single edge labelled to indicate its role. As long as the type of marker is identified, the algorithm can continue to treat the vertices of attachment as source, sink or neither, as appropriate for the particular operation.

Due to space constraints, we have not included the proofs of the first and third cases, when u and v are incomparable, and when $u = s$ respectively. The second case, $u < v$ and $u \neq s$ is indicative of the type of proof required, so we attempt to provide some detail. For full proofs see [8] or [9].

5.1 Cut-set $\{u, v\}$; u and v are incomparable.

THEOREM 5.1. *Let G be a biconnected directed acyclic graph with a single source s and let $X = \{x_i\} \subseteq V(G)$ be a set of vertices. Let $\{u, v\}$ be a separation pair of G, with u and v incomparable. Let S be the connected component of G with respect to $\{u, v\}$ containing s, and H be the union of all other components. Then, G admits an upward plane drawing with all vertices of X on the outer face if and only if*

(i) $S' = S \cup M_t$ admits an upward plane drawing with all vertices of X in S on the outer face, and w_t on the outer face if some $x \in X$ is contained in H.

(ii) $H' = H \cup M_s$ admits an upward plane drawing with all vertices of X in H on the outer face.

5.2 Cut-set $\{u, v\}$, where $u < v$, $u \neq s$.

Here we consider any other vertex cut-sets not involving the source s. We divide the graph at a vertex cut $\{u, v\}$ into two subgraphs—the *source component* S (the one component which contains the source s), and the union of the remaining components H. Note that v can be a source in S, as long as there is a u, v path in H.

In this section we will give the full proof. First we need some preliminary results:

PROPOSITION 5.1. *If G is a connected DAG with exactly two sources u and v, then there exists some w_t such that two vertex disjoint (except at w_t) directed paths $u \xrightarrow{+} w_t$ and $v \xrightarrow{+} w_t$ exist in G.*

Proof. Let G be such a DAG and let P be an undirected path from u to v. Note that every x in P is comparable with either u or v, otherwise G has more than two sources. Follow P from u to the first node x (following y on P) incomparable with u (in G). Then x is comparable with v and (x, y) is an edge in G (otherwise $u < x$), so y is also comparable with v. Taking the first common vertex in the paths $u \xrightarrow{+} y$ and $v \xrightarrow{+} y$ gives w_t. ∎

The following results show the existence of lower bounds and upper bounds (in the partial order corresponding to G) under certain conditions. This allows us to prove the necessity conditions in Theorem 5.2 (to come).

LEMMA 5.2. *If G is a biconnected DAG with a single source s, and u and v are incomparable vertices in G, then there exists some w_s such that two vertex disjoint (except at w_s) directed paths $w_s \xrightarrow{+} u$ and $w_s \xrightarrow{+} v$ exist in G. If $\{u,v\}$ is a cut-set in G, then there also exists some w_t such that two vertex disjoint (except at w_t) directed paths $u \xrightarrow{+} w_t$ and $v \xrightarrow{+} w_t$ exist in G.*

Proof. Since G is a single source digraph, there exist directed paths from s to u and s to v in G. Taking the last common vertex in these paths gives w_s.

For the existence of w_t, let u and v be an incomparable separation pair of G. Since $\{u,v\}$ cuts G into at least two connected components, any non-source component H has u and v as its (exactly) two sources, and the result follows from Proposition 5.1. ■

LEMMA 5.3. *If G is a biconnected DAG with a single source s and cut-set $\{u,v\}$, where $u < v$ in G and $u \neq s$, then in any non-source component H of G with respect to $\{u,v\}$, where $deg^+ v > 0$, there exists some w_t such that $u \xrightarrow{+} w_t$ and $v \xrightarrow{+} w_t$ are vertex disjoint directed paths in H.*

Proof. No vertex other than u and v can be a source in H, otherwise G has more than one source. u is always a source in H, otherwise G contains a directed cycle. If v is also a source, then we are done by Proposition 5.1.

If v is not a source, let $w \in H$ be a vertex dominated by v. G is biconnected, so there are two vertex disjoint $u \xrightarrow{+} w$ undirected paths in G. But u and v are cut-vertices in G, so at least one of the paths P lies completely within H and does not contain v (as w is in H and the only exit points from H are u and v). Every x on P is comparable with either u or v, or else G has more than one source. Find the last vertex y on P which has a $u \xrightarrow{+} y$ path (in G) without v. If $y = w$, then we are done. Otherwise, the vertex x following y on P has any $u \xrightarrow{+} x$ path necessarily going through v. Then there exist directed paths $v \xrightarrow{+} x$, $u \xrightarrow{+} x$ with the latter not containing v so the last common vertex on these paths provides a w_t. ■

We are now ready to proceed with the statement of the main result of the decomposition.

THEOREM 5.2. *Let G be a biconnected directed acyclic graph with a single source s, and let $X = \{x_i\} \subseteq V(G)$ be a set of vertices. Let $\{u,v\}$ be a separation pair of G with $u < v$ in G and $u \neq s$. Let S be the source component of G with respect to $\{u,v\}$ and H be the union of all other components. Then, G admits an upward plane drawing with all vertices of X on the outer face if and only if*

(i) $S' = (S \cup H\text{-marker})$ admits an upward plane drawing with all vertices of X in S on the outer face and w_t (if it exists, otherwise the edge (u,v)) on the outer face if some $x \in X$ is contained in H.

(ii) $H' = (H \cup S\text{-marker})$ admits an upward plane drawing with w_t (if it exists, otherwise the edge (u,v)) and all vertices of X in H on the outer face.

where

$$H\text{-}marker = \begin{cases} M_t & \text{if } v \text{ is a source in } H \\ M_{uv} & \text{if } v \text{ is a sink in } H \\ M_{uvt} & \text{otherwise.} \end{cases}$$

and

$$S\text{-}marker = \begin{cases} M_t & \text{if } v \text{ is a source in } S \\ M_{uv} & \text{otherwise.} \end{cases}$$

Proof. (Necessity) Suppose G admits an upward plane drawing with all $x_i \in X$ on the outer face.

(Necessity of condition (i)): If v is a source in H, then there exists some w_t in H and vertex disjoint paths $u \xrightarrow{+} w_t$ and $v \xrightarrow{+} w_t$ by Proposition 5.1; so $S' = S \cup M_t$ is homeomorphic to a subgraph of G and is upward planar. If v is a sink in H, then u is the single source of H, as only u and v are possible sources. Thus, in H, there is a path $u \xrightarrow{+} v$, so $S' = S \cup M_{uv}$ is homeomorphic to a subgraph of G and is upward planar. If v is neither a source nor a sink in H then, by Lemma 5.3, there is also some $w_t > v$ and disjoint directed paths $u \xrightarrow{+} w_t$ and $v \xrightarrow{+} w_t$ in G. Since v is a non-source in H, there is also a $u \xrightarrow{+} v$ path in H. This path crosses the $u \xrightarrow{+} w_t$ path at some latest vertex z on that path, so $S \cup (u \xrightarrow{*} z) \cup (z \xrightarrow{+} v) \cup (z \xrightarrow{+} w_t) \cup (v \xrightarrow{+} w_t)$ is a subgraph of G and hence upward planar. Note that these four paths are disjoint. Since z has in-degree one we can contract the $u \xrightarrow{*} z$ path to u without destroying upward planarity, by Lemma 3.1, so $S \cup \{(u,v),(u,w_t),(v,w_t)\}$ has an upward planar subdivision and is upward planar itself. No other vertices of this graph can lie inside the u, v, w_t triangle, as $s < u$ (and hence below/outside the triangle) and there are no other (u,v) components in S' (as we chose S to be the *single* source component), so the extra edges and vertex for M_{uvt} can be added without destroying planarity[2].

(Necessity of condition (ii)): If v is a source in S, then, by Proposition 5.1, there are vertex disjoint paths $s \xrightarrow{+} w_t$ and $v \xrightarrow{+} w_t$ in S. There must be an $s \xrightarrow{+} u$ path in S, otherwise there is either a second source (u is a source in H, so it cannot also be a source in S) or a cycle in G ($u < v$ in G, so there can be no $v \xrightarrow{+} u$ directed path in S). Let z be the last vertex common to paths

[2] The point of adding these edges is to fix the face in S for the sufficiency conditions.

$s \xrightarrow{+} u$ and $s \xrightarrow{+} w_t$. Then, $H \cup \{(z,u),(z,w_t),(v,w_t)\}$ is homeomorphic to a subgraph of G and is upward planar. Since $deg^- u = 1$ (in this graph), the edge (z,u) can be contracted without destroying upward planarity, by Lemma 3.1, and $H' = H \cup M_t$ is upward planar.

Otherwise (v a non-source), if $u < v$ in S, then $H' = H \cup M_{uv}$ is homeomorphic to a subgraph of G and, hence, is upward planar. If u and v are incomparable in S, then they share a greatest lower bound w_s, by Lemma 5.2, and $H \cup \{(w_s,u),(w_s,v)\}$ is upward planar. Again, $deg^- u = 1$ in H, so the (w_s,u) edge can be contracted to give $H' = H \cup M_{uv}$.

The necessity of the outer facial conditions on the x_i's can also be shown.

(Sufficiency) Suppose S' and H' admit upward plane drawings meeting the requirements (i) and (ii).

Case 1: v is a source in H: If v is a source in H it cannot at the same time be a source in S, as $u < v$ in either S or H. Thus $H' = H \cup (u,v)$ is upward planar with single source u. Using Lemma 3.2, add H' (with u and v renamed as u' and v') to S', identifying u' with w_t. We can do this so that edges (v,w_t) and (w_t,v') are consecutive in the cyclic order about w_t. Using Lemma 3.4, split w_t by making these two edges incident with a new vertex u_1 and the remaining edges incident with a new vertex u_2. Now v and u_1 have in-degree 1, so use Lemma 3.1 to contract their in-edges, thus identifying v and v'. Vertex u_2 has in-degree 1 so contract (u,u_2). The result is the graph G, and thus G is upward planar. See Figure 4(a).

Case 2: v is a sink in H: If v is a non-source in S, then $H' = H \cup (u,v)$ is upward planar with u and v on the outer face by assumption. If v is a source in S, then $H' = H \cup M_t$ is upward planar with w_t on the outer face. In either case H is upward planar with source u and sink v on the outer face. By Lemma 3.3 we can add H to S' in place of the (u,v) edge in S', and the result, G, is upward planar. The two possibilities are illustrated in Figure 4 (b) and (c).

Case 3: v is a non-source/sink in H: Suppose v is a source in S. Then $H' = H \cup M_t$ is upward planar with the sink w_t on the outer face. Using Lemma 3.3, add H' (renaming u and v to u' and v' respectively) to S' in place of the edge (u,v), identifying u' with u and w_t with v. Throw away the edge (u,w_t) and the remaining marker edges of S'. Vertex v now has in-degree 1 so the edge (v',v) can be contracted by Lemma 3.1, and the result, G, is upward planar. See Figure 4(d).

Suppose then that v is a non-source in S. Consider a plane representation of S' and throw away the marker edges, save for $(u,w_t),(v,w_t),(u,v)$, which then form a face. $H' = H \cup (u,v)$ is upward planar with u and v on the outer face. Let z be some sink on the outer face,

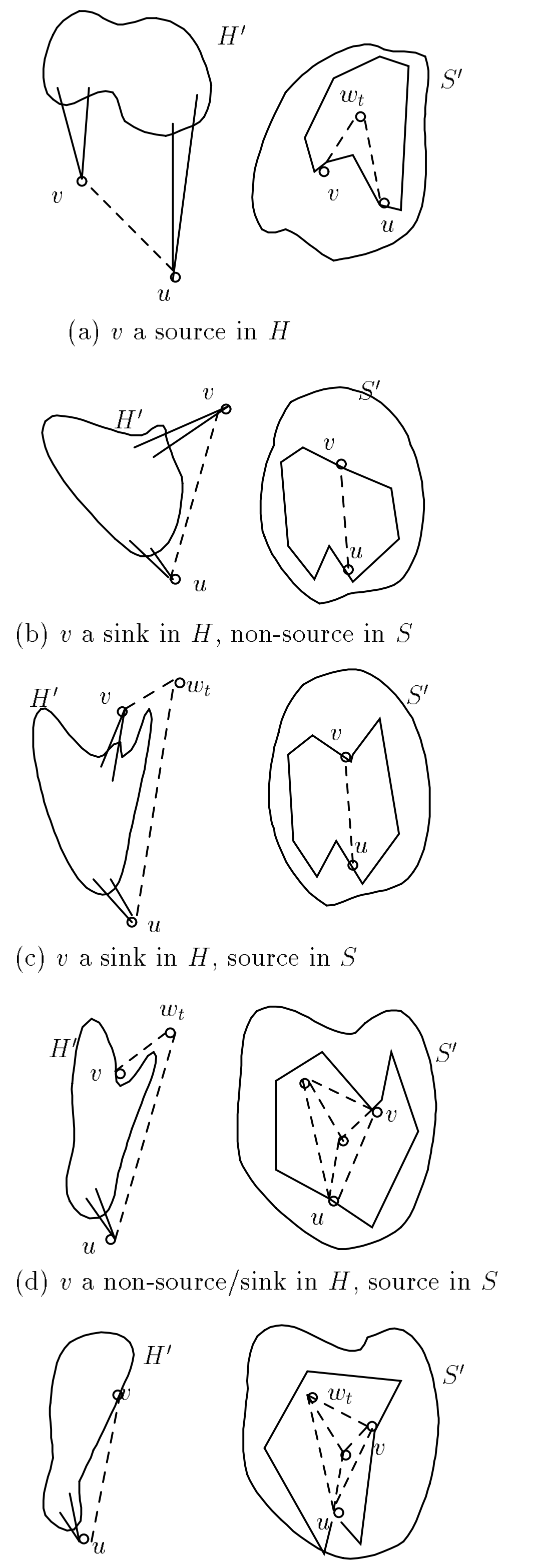

(a) v a source in H

(b) v a sink in H, non-source in S

(c) v a sink in H, source in S

(d) v a non-source/sink in H, source in S

(e) v a non-source/sink in H, non-source in S

Figure 4: Merging S and H; cut-set $\{u,v\}$ and $u < v$.

and add the edge (v, z) to obtain H'', upward planar with u, v, z on the outer face. Using Lemma 3.3, add H'' (with u and v renamed to u' and v') to S' in place of the edge (u, w_t), identifying u' with u and z with w_t. Do this so that v' and v share the face of edges $(u, v'), (v', z), (u, v), (v, z)$. Clearly we can now identify the vertices v and v'. We obtain an upward planar graph containing G as a subgraph. See Figure 4(e).

It can be shown that the conditions on the x_i's, when they arise, are also sufficient. ■

5.3 Cut-set $\{s, v\}$.

As mentioned in the introduction to this chapter, it is important to be able to distinguish the "inner" and "outer" components. The inner component will be embedded in a face of the outer one, and thus the inner component will have to have its marker on its outer face since this marker is a proxy for the outer component. If we have to check each component as a potential inner component, we must recursively solve two subproblems for each component, and an exponential time blowup results.

Until now, the outer component has been uniquely identified as the *source component*, since that component cannot lie within an internal face of any other component. If we have a cut-set of the form $\{s, v\}$ where s is the source, then we lose this restriction, so we handle it instead by requiring one of the components, E, to be 3-connected so that deciding if it can be the inner face does not require recursive calls. To decide if E can be the inner face we need to test if it is upward planar with s, v on the outer face. This can be done in linear time using the algorithm of Section 4. If G has only cut-sets of the form $\{s, v\}$, then, for at least one such cut-set, one of the components will be triconnected. Given the list of cut-sets we can find such a cut-set and such a component in linear time using depth-first search.

We capture these ideas in terms of two theorems. One is applicable if the triconnected component can be the inner component, and one if it cannot. Note that in the statement of these theorems, we continue to use u (redundant since $u = s$) for consistency with previous usage.

THEOREM 5.3. *Let G be a biconnected DAG with a single source s, and let $X = \{x_i\} \subseteq V(G)$ be a set of vertices. Let $\{u, v\}$ be a separation pair of G where $u = s$, E be a 3-connected component of G with respect to $\{u, v\}$, and F be the union of all other components of G with respect to $\{u, v\}$. If E admits an upward plane drawing with u and v on the outer face, then G admits an upward plane drawing with all vertices of X on the outer face if and only if*

(i) $F' = (F \cup E\text{-marker})$ admits an upward plane drawing with all vertices of X in F on the outer face, and w_t (if it exists, otherwise the edge (u, v)) also on the outer face if some $x \in X$ contained in E.

(ii) $E' = (E \cup F\text{-marker})$ admits an upward plane drawing with w_t (if it exists, otherwise the edge (u, v)) and all vertices of X in E on the outer face,

where

$$E\text{-marker} = \begin{cases} M_t & \text{if } v \text{ is a source in } E \\ M_{uv} & \text{if } v \text{ is a sink in } E \\ M_{uvt} & \text{otherwise.} \end{cases}$$

and

$$F\text{-marker} = \begin{cases} M_t & \text{if } v \text{ is a source in } F \\ M_{uv} & \text{otherwise.} \end{cases}$$

Proof. Similar to that of Theorem 5.2. ■

THEOREM 5.4. *Let G be a biconnected DAG with a single source s and let $X = \{x_i\} \subseteq V(G)$ be a set of vertices. Let $\{u, v\}$ be a separation pair of G where $u = s$, E be a 3-connected component of G with respect to $\{u, v\}$, and F be the union of all other components of G with respect to $\{u, v\}$. If E does not admit an upward plane drawing with u and v on the outer face, then G admits an upward plane drawing with all vertices of X on the outer face if and only if*

(i) There is no $x \in X$ contained in F.

(ii) $F' = (F \cup E\text{-marker})$ admits an upward plane drawing with w_t (if it exists, otherwise the edge (u, v)) also on the outer face if some $x \in X$ is contained in E.

(iii) $E' = (E \cup F\text{-marker})$ admits an upward plane drawing with all $x \in X$ on the outer face,

where

$$E\text{-marker} = \begin{cases} M_t & \text{if } v \text{ is a source in } E \\ M_{uv} & \text{if } v \text{ is a sink in } E \\ M_{uv} & \text{otherwise.} \end{cases}$$

and

$$F\text{-marker} = \begin{cases} M_t & \text{if } v \text{ is a source in } F \\ M_{uv} & \text{if } v \text{ is a sink in } F \\ M_{uvt} & \text{otherwise.} \end{cases}$$

Proof. (outline) Since E has no upward plane drawing with s and v both on the outer face, the only way G could be upward planar is if F can be embedded within a face of E. Thus, the outer face of G is fixed as being some face of the drawing of E' not containing v. It remains to ensure that there is some embedding of F which will fit the structural constraints of the shape of a face shared by s and v in the drawing of E. These are exactly the conditions previously required by E for embedding within the drawing of F. The remainder of the proof does not rely on the triconnectedness of either component, and is similar to the proof of Theorem 5.3.

6 Conclusions and Further Work

We have given a linear time algorithm to test whether a given single-source digraph has an upward plane drawing strongly equivalent to a given plane drawing, and give a representation for this drawing if it exists. We have used this result to outline an efficient $O(n^2)$ algorithm to test upward planarity of a single-source digraph.

A lower bound for the single-source upward planarity problem is not known, although we believe that it may be possible to perform the entire test in subquadratic (perhaps linear) time. An obvious extension of this work would be to find such an algorithm or prove a lower bound.

This paper has concentrated on the issues of efficiently testing for an upward plane drawing and outputting an abstract representation of such a drawing. Using the representation it is easy to add edges to the graph to get a planar s-t graph. Then an upward plane drawing can be obtained from the algorithm of DiBattista and Tamassia in $O(n \log n)$ arithmetic steps [2]. However, DiBattista, Tamassia and Tollis have shown that there exist upward planar graphs which require an exponential sized integer grid [3] so the algorithm is actually output sensitive, and hence exponential in the worst case. It would be interesting to characterize some classes of digraphs which permit upward plane drawings on a polynomially sized grid. Guaranteeing minimum area in all cases is, however, NP-hard [13].

The more general problem of testing upward planarity of an arbitrary acyclic digraph is open. The only known characterization is that any such graph is a subgraph of a planar s-t graph [2].

References

[1] J. A. Bondy and U. S. R. Murty. *Graph Theory with Applications.* MacMillian Co. New York, 1976.

[2] G. DiBattista and R. Tamassia. *Algorithms for plane representations of acyclic digraphs.* Theoretical Computer Science 61, pp. 175-178, 1988.

[3] G. DiBattista, R. Tamassia and I. G. Tollis. *Area requirement and symmetry display in drawing graphs.* Proc. ACM Symposium on Computational Geometry, pp. 51-60, 1989.

[4] I. Fáry. *On straight line representations of planar graphs.* Acta. Sci. Math. Szeged 11, pp. 229-233, 1948.

[5] I. S. Filotti, G. L. Miller, J. Reif. *On determining the genus of a graph in $O(v^{O(g)})$ steps,* Proc. 11th ACM Symp. on Theory of Computing, 1979, pp. 27-37.

[6] J. Hopcroft and R. Tarjan. *Efficient planarity testing.* J.ACM 21,4, pp. 549-568, 1974.

[7] J. Hopcroft and R. Tarjan. *Dividing a graph into triconnected components.* SIAM J. Comput. 2, pp. 135-158, 1972.

[8] M. D. Hutton. *Upward planar drawing of single source acyclic digraphs.* Masters Thesis, University of Waterloo, Sept. 1990.

[9] M. D. Hutton and A. Lubiw. *Upward planar drawing of single source acyclic digraphs.* In preparation.

[10] D. Kelly. *Fundamentals of planar ordered sets.* Discrete Math 63, pp. 197-216, 1987.

[11] D. Kelly and I. Rival. *Planar lattices.* Can. J. Math., Vol 27, No. 3, pp. 636-665, 1975.

[12] C. R. Platt. *Planar lattices and planar graphs.* J. Comb. Theory (B) 21, pp. 30-39, 1976.

[13] R. Tamassia and P. Eades. *Algorithms for drawing graphs: an annotated bibliography.* Brown University TR CS-89-09, 1989.

[14] C. Thomassen. *Planar acyclic oriented graphs.* Order 5, pp. 349-361, 1989.

[15] W. T. Tutte. *Convex representations of graphs.* Proc. London Math. Soc. Vol.10, pp. 304-320, 1960.