Discrete Comput Geom 1:321-341 (1986)

Discrete & Computational Geometry
© 1986 Springer-Verlag New York Inc.

A Unified Approach to Visibility Representations of Planar Graphs*

Roberto Tamassia[1]† and Ioannis G. Tollis[2]‡

[1] Coordinated Science Laboratory and Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA

[2] Coordinated Science Laboratory and Department of Computer Science, University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA

Abstract. We study *visibility representations* of graphs, which are constructed by mapping vertices to horizontal segments, and edges to vertical segments that intersect only adjacent vertex-segments. Every graph that admits this representation must be planar. We consider three types of visibility representations, and we give complete characterizations of the classes of graphs that admit them. Furthermore, we present linear time algorithms for testing the existence of and constructing visibility representations of planar graphs. Many applications of our results can be found in VLSI layout.

1. Introduction

Several layout compaction strategies for VLSI are based on the concept of *visibility* between parallel segments [12] where we say that two parallel segments of a given set are *visible* if they can be joined by a segment orthogonal to them, which does not intersect any other segment. In this paper, we study *visibility representations* of graphs, which are constructed by mapping vertices to horizontal segments, and edges to vertical segments drawn between visible vertex-segments. It is easy to see that a graph that admits such a representation must be planar.

---

* This research was partially supported by the Joint Services Electronics Program under Contract N00014-84-C-0149. Roberto Tamassia was supported in part by a Fulbright grant.

† On leave from the Dipartmento di Informatica e Sistemistica, Universita' di Roma, "La Sapienza", Via Buonarroti 12, 00185 Roma, Italy.

‡ Portions of this paper were written while this author was visting the Tektronix Computer Research Laboratory in Beaverton, Oregon, USA.

Various visibility representations have been considered in the literature, where vertices are represented either by horizontal intervals or by horizontal segments. An interval may or may not contain one or both of its endpoints. Otten and van Wijk [10] gave an algorithm for constructing a representation of a 2-connected planar graph such that vertices are represented by horizontal segments and edges by vertical segments having only points in common with the pair of horizontal segments corresponding to the vertices they connect (Fig. 1(b)). In the following, this representation will be referred to as *weak-visibility representation* (*w-visibility representation*). The algorithm of Otten and van Wijk can be implemented to run in linear time, though they give no time bound. Duchet *et al.* [2] independently proved that every planar graph admits a w-visibility representation.

Melnikov [9] suggested the problem of characterizing the graphs whose vertices can be represented by horizontal intervals in the plane such that two vertices are adjacent if and only if their associated intervals are visible (Fig. 1(c)). From the result of Duchet *et al.*, it follows that every maximal planar graph admits a representation of the latter type, which will be called *ε-visibility representation.* Thomassen [16] extended this by showing that all 3-connected planar graphs admit an ε-visibility representation. Note that the ε-visibility representation differs from the w-visibility representation because: (a) vertices are represented by intervals and not only by segments, and (b) visible vertex-intervals always correspond to adjacent vertices.

Another problem that naturallly arises in this context is the following: characterize the class of graphs whose vertices can be represented by horizontal segments such that two vertices are adjacent if and only if their corresponding segments are visible (see Fig. 1(d)). Such a representation will be called *strong-visibility representation* (*s-visibility representation*) and it differs from the w-visibility representation because it *requires* that visible vertex-segments correspond to adjacent vertices. It also differs from the ε-visibility representation because the

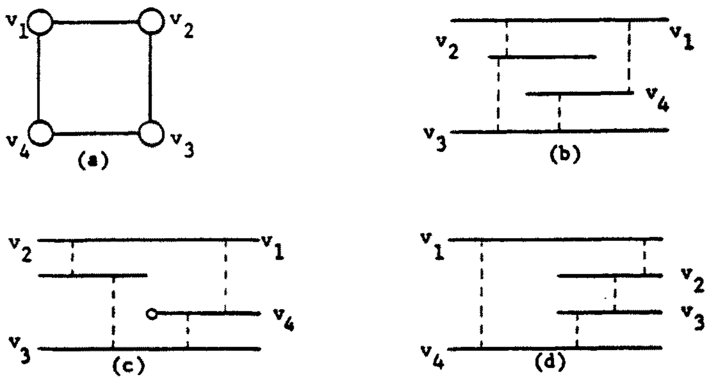

Fig. 1. The three visibility representations: (a) a cycle of length 4; (b) w-visibility representation; (c) ε-visibility representation; (d) s-visibility representation.

vertices are always represented by segments. Luccio *et al.* [8] gave a partial solution to the above problem by requiring that the endpoints of all the horizontal segments have distinct x-coordinates. Namely, they defined a new family of graphs, called *ipo-triangular graphs* (graphs that can be transformed into planar multigraphs with all triangular internal faces, by successive duplications of existing edges), and proved that a graph admits an s-visibility representation with the above restriction if and only if it is ipo-triangular. Notice that the restriction on the x-coordinates of vertex-segments is essential to their characterization. Consider for example any cycle of length greater than three (see Fig. 1).

The main contributions of this paper are:

(1) We unify and extend the results of Otten and van Wijk and of Duchet *et al.* on the w-visibility representation. First, we propose a linear time algorithm for constructing a w-visibility representation for a 2-connected planar graph. This algorithm is a variant of the Otten-van Wijk algorithm.[1] Next, we extend our algorithm such that it constructs a w-visibility representation of any planar graph without increasing the time complexity.
(2) We present a complete solution of Melnikov's problem by showing that a graph admits an ε-visibility representation if and only if it is planar and there is a planar embedding for it such that all cutpoints appear on the boundary of the same face. We also give two linear time algorithms, one for testing the above condition, and the other for constructing an ε-visibility representation.
(3) Finally, we give a complete characterization of the class of graphs that admit an s-visibility representation, and we show how to construct one efficiently in the case of maximal planar graphs and 4-connected planar graphs.

Another application of our results in the field of VLSI layout is to the problem of *minimal-node-cost* planar embedding. This problem has been considered by Storer [13] and consists of finding an embedding of a graph in the rectilinear grid where the total number of bends along edges is minimum. The technique described in this paper can be used as the core of a linear time heuristic algorithm [15] for this problem which yields better performance guarantees than the heuristics given by Storer.

The rest of this paper is organized as follows. Section 2 contains complete definitions of the above visibility representations, and basic properties of them. Section 3 is concerned with the w-visibility representation. In Section 4, we present the results on the ε-visibility representation. Section 5 deals with the s-visibility representation. Finally, in Section 6 we present a summary of our results and discuss open problems for further research on the subject.

[1] After the submission of this paper, we became aware that Rosenstiehl and Tarjan [11] independently proposed another variant of the Otten-van Wijk algorithm and considered the "interlocking" layout of the dual graph.

2. Preliminaries and Definitions

The basic graph theoretic definitions can be found in many textbooks [1], [3]. Here, we recall some terminology on connectivity properties of graphs. A *cut-point* of a graph is a vertex whose removal disconnects the graph. A *separation pair* is a pair of vertices whose removal disconnects the graph. A graph is said to be *k-connected* if it cannot be disconnected by the removal of less than k vertices. Clearly, if a graph is k-connected, then it is also l-connected for all $l \leq k$. A *block* of a graph G is a maximal 2-connected subgraph of G. The *block-cutpoint tree* of G is a tree whose vertices are the cutpoints and the blocks of G, and whose edges connect each cutpoint to the blocks that contain it.

Let S be a set of horizontal nonoverlapping segments in the plane. Two segments s, s' of S are said to be *visible* if they can be joined by a vertical segment not intersecting any other segment of S. Furthermore, s and s' are called *ε-visible* if they can be joined by a vertical band of nonzero width that does not intersect any other segment of S. This is equivalent to saying that s and s' can be joined by two distinct vertical segments not intersecting any other segment of S.

Definition 1. A *w-visibility representation* for a graph $G = (V, E)$ is a mapping of vertices of G into nonoverlapping horizontal segments (called *vertex-segments*) and of edges of G into vertical segments (called *edge-segments*) such that, for each edge $(u, v) \in E$, the associated edge-segment has its endpoints on the vertex-segments corresponding to u *and* v, and it does not cross any other vertex-segment.

In order to study the visibility representations in a unified way, we give a definition of ε-visibility representations using segments instead of intervals.

Definition 2. An *ε-visibility representation* for a graph G is a w-visibility representation with the additional property that two vertex-segments are ε-visible if and only if the corresponding vertices of G are adjacent.

Now, we show that our definition is equivalent to the one of Melnikov with respect to the class of graphs that admit an ε-visibility representation. First, from a Melnikov ε-visibility representation of a graph G, we can obtain an ε-visibility representation of G by closing all the intervals. Similarly, from an ε-visibility representation, we can derive a Melnikov ε-visibility representation by transforming each segment into an interval, removing its right endpoint.

Definition 3. An *s-visibility representation* for a graph G is a w-visibility representation with additional property that two vertex-segments are visible if and only if the corresponding vertices of G are adjacent.

If a graph admits any of the three aforementioned visibility representations, then it is planar, since a planar embedding of it can be immediately obtained from the visibility representation by shrinking each vertex-segment into a point.

A *face* of a visibility representation Γ is a maximal region of the plane such that, for every two points x and y in it, there is a Jordan curve from x to y which does not intersect any segment of Γ. Let C_w, C_ε, and C_s be the classes of graphs which admit a w-visibility representation, ε-visibility representation, and s-visibility representation, respectively. Clearly, if $G \in C_w$, then G is a spanning subgraph of some graph $H \in C_\varepsilon$, and furthermore H is a spanning subgraph of another graph $N \in C_s$. As we will see in the following, the three classes of graphs defined above are hierarchically related, i.e.: C_s is properly included in C_ε, and C_ε is properly included in C_w.

In the remaining part of this section, we present some preliminary results that will be used later. A *PERT-digraph* $D = (V, A)$ is an acyclic digraph with exactly one source, s, and one sink, t. We usually associate a positive length with each arc of D. A well-known problem on PERT-digraphs is the following: for each vertex v of D, find the length of the longest path from s to v. This quantity will be denoted by $\alpha(v)$. The *critical path method* solves this problem in $O(|A|)$ time [3].

An *st-numbering* for a graph $G = (V, E)$, where s and t are two distinct vertices of G, is a one-to-one mapping ξ: $V \to \{1, 2, \ldots, |V|\}$, such that $\xi(s) = 1$, $\xi(t) = |V|$, and each vertex $v \neq s, t$ has two adjacent vertices u, w for which $\xi(u) < \xi(v) < \xi(w)$. Given an st-numbering ξ for a graph $G = (V, E)$, we construct a digraph $D = (V, A)$ by orienting every edge from the lowest numbered vertex to the highest one. Namely, $[u, v] \in A$ if and only if $(u, v) \in E$ and $\xi(u) < \xi(v)$. The digraph D, which is *induced* by ξ, is clearly acyclic and has exactly one source, s, and one sink, t, i.e., it is a PERT-digraph. Conversely, any topological sorting of the vertices of a PERT-digraph is an st-numbering for the underlying undirected graph.

Lempel *et al.* [7] showed that for every 2-connected graph and every edge (s, t), there exists an st-numbering. A linear time algorithm for finding it has been presented by Even and Tarjan [4].

Fact 1. Every directed path of D visits vertices in increasing order.

Proof. Otherwise, there would be an arc $[w, v]$ with $\xi(v) < \xi(w)$, which contradicts the definition of D. □

Fact 2. For every vertex v of D there exists a simple directed path P from s to t containing v.

Proof. Let P be any maximal path containing v. Let u and w be the first and last vertices of P, respectively. Then u is a source, and w is a sink. Hence, we have $u = s$ and $w = t$. □

Let D be a 2-connected planar digraph, induced by some st-numbering, and $\hat{D}$ any planar embedding of D. For any vertex v of D we define $\deg^+(v)$ and $\deg^-(v)$ to be the number of arcs outgoing from v and incoming to v, respectively. Furthermore, we denote with $l(f)$ and $h(f)$ the lowest and highest numbered vertices on the boundary of a face f of $\hat{D}$.

Lemma 1. *Each face f of $\hat{D}$ consists of two directed paths from $l(f)$ to $h(f)$.*

Proof. Let f be a face of $\hat{D}$ for which the lemma is not true. Then there exists an arc $[w, u]$ on the boundary of f directed from $h(f)$ to $l(f)$. From Fact 2, there are directed paths P_1 from u to t and P_2 from s to w (Fig. 2). From Facts 1 and 2, and the planarity of D, these two paths must intersect at a common vertex x. But then $\hat{D}$ (D) has a cycle which consists of the arc $[w, u]$, the subpath of P_1 from u to x, and the subpath of P_2 from x to w. This contradicts the acyclicity of D. □

Lemma 2. *All outgoing (ingoing) arcs of any vertex v of $\hat{D}$ appear consecutively around v.*

Proof. The lemma holds trivially for the vertices s and t. Let v be any other vertex, and suppose, for a contradiction, that there are arcs $[v, w_0]$, $[w_1, v]$, $[v, w_2]$, and $[w_3, v]$, appearing in clockwise order around v (Fig. 3). From Fact 2, there are directed paths P_0 and P_2 from w_0 and w_2 to t, respectively. Similarly, there are directed paths P_1 and P_2 from s to w_1 and w_3, respectively. But then one of P_2 and P_0 must intersect either P_1 or P_3 at a common vertex x. This implies that $\hat{D}$ (D) has a cycle, which contradicts the acyclicity of D. □

Lemma 3. *Every vertex $v \in V - \{s, t\}$ is the lowest numbered vertex for* $\deg^+(v) - 1$ *faces and the highest numbered vertex for* $\deg^-(v) - 1$ *faces. s is the lowest numbered vertex for* $\deg^+(s)$ *faces, and t is the highest numbered vertex for* $\deg^-(v)$ *faces.*

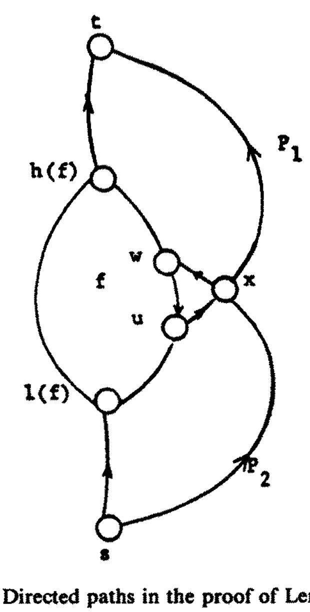

Fig. 2. Directed paths in the proof of Lemma 1.

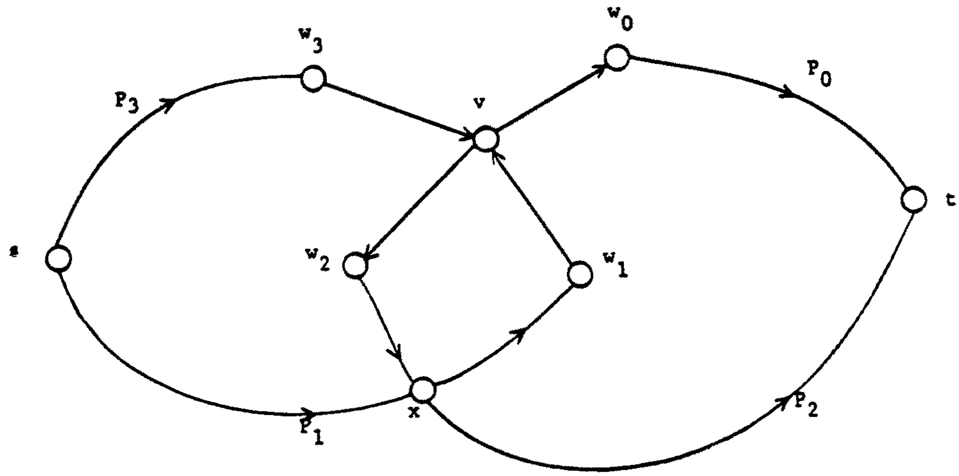

Fig. 3. Directed paths in the proof of Lemma 2.

Proof. Since D is 2-connected, every vertex v is in the boundary of deg(v) distinct faces. By Lemma 2, all incoming and outgoing arcs incident to v appear consecutively around v. Therefore, if $v \neq s, t$, there are $\deg^+(v)-1$ faces around v that contain two directed paths originating from v (Fig. 4).Hence, from Lemma 1, v is the lowest numbered vertex for $\deg^+(v)-1$ faces. A similar argument applies to the outgoing arcs. The different result for s and t is due to the fact that all their incident arcs have the same orientation. □

3. Weak-Visibility Representation

First, we describe a linear time algorithm for constructing a w-visibility representation of a 2-connected planar graph $G=(V, E)$. Next, we extend this algorithm in order to construct a w-visibility representation of any planar graph.

For the sake of simplicity, we will use the same notation for the vertex-segments of the visibility representations and their corresponding vertices in the graph. The same will be done for the edge-segments and their corresponding edges.

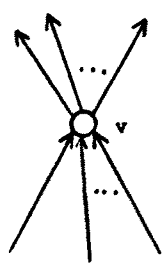

Fig. 4. Faces around vertex v.

Algorithm *W-VISIBILITY*
Input: A 2-connected planar graph $G=(V, E)$.
Output: A w-visibility representation for G such that each vertex- and edge-segment has endpoints with integer coordinates.

1. Select an edge $(s, t) \in E$.
2. Compute an st-numbering for G. Let D be the directed graph induced by the st-numbering.
3. Find a planar representation $\hat{D}$ of D such that the arc $[s, t]$ is on the external face and the rest of D lies on the right side of $[s, t]$. Use $\hat{D}$ to construct a new digraph D^* as follows:
   3.1. Vertices of D^* are the faces of $\hat{D}$.
   3.2. There is an arc $[f, g]$ in D^* if face f shares an arc $a=[v, w]$, distinct from $[s, t]$, with face g and a is positively oriented with respect to f, i.e., face f is on the left side of a, when a is traversed from the tail to the head.

   Note that D^* is a 2-connected planar PERT-digraph, with source s^* (the internal face containing arc $[s, t]$), and sink t^* (the external face). If we ignore arc directions, we observe that D^* is the dual graph of $\hat{D}$ without the dual edge of (s, t).
4. Apply the critical path method to D^* with all arc-lengths equal to 2. This gives the function $\alpha(f)$ for each vertex f of D^*.
5. Construct the w-visibility representation as follows:
   5.1. Use the st-numbering computed in step 2 to assign y-coordinates to horizontal vertex-segments.
   5.2. Set the x-coordinate of arc $[s, t]$ equal to -1.
   5.3. For any other arc a of $\hat{D}$, set the x-coordinate of the corresponding vertical edge-segment equal to an integer j, with $\alpha(f)<j<\alpha(g)$, where f and g are the faces of $\hat{D}$ sharing a in their contour.
   5.4. Set the y-coordinates of the endpoints of each edge-segment equal to the ones of the connected vertex-segments.
   5.5. Set the x-coordinates of the left and right endpoint of each vertex-segment equal to the minimum and maximum x-coordinates of their incident arcs, respectively. If a vertex-segment v is incident to exactly two edge-segments with the same x-coordinate, x_v, then set the x-coordinates of the endpoints of v to x_v-1 and x_v, respectively.

An example of the construction performed by the algorithm *W-VISIBILITY* is given in Fig. 5. Figure 5(a) shows a planar embedding $\hat{D}$ along with the corresponding D^*. Vertices of $\hat{D}$ and D^* are represented by white and black circles, respectively. The white vertices are numbered according to the st-numbering. For each black vertex f_i the value of $\alpha(f_i)$ is shown in parentheses. Figure 5(b) illustrates the w-visibility representation produced by the algorithm. We will use the following lemma to prove the correctness of the algorithm.

Lemma 4. *For any two vertices f and g of D^*, either there is a directed path of D^* between them, or there is a directed path of D from* $\min\{h(f), h(g)\}$ *to* $\max\{l(f), l(g)\}$.

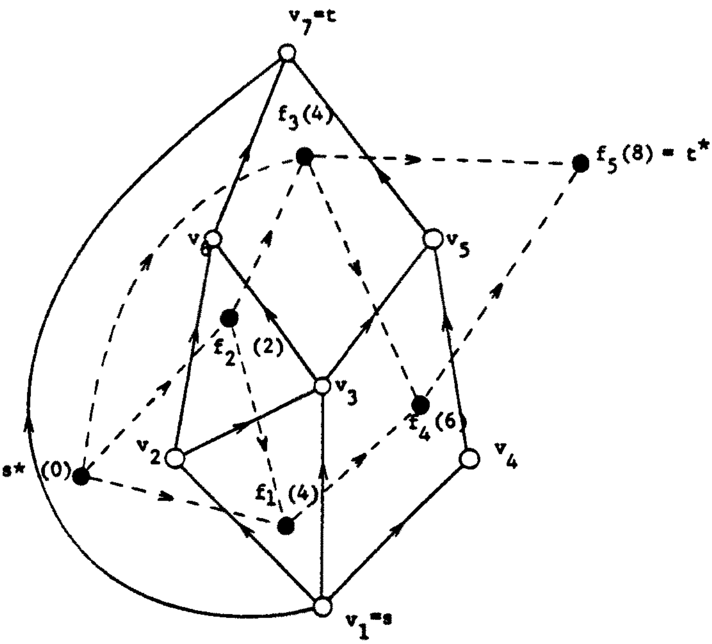

(a)

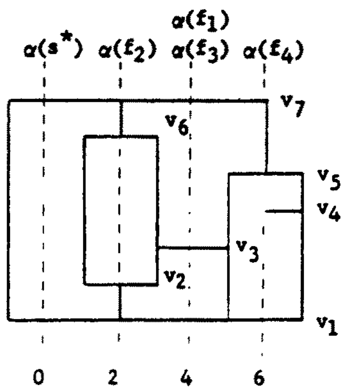

(b)

Fig. 5. Running example for algorithm *W-VISIBILITY*. (a) Directed graphs $\hat{D}$ and D^* derived from a graph G; (b) w-visibility representation for G.

Proof. Assume without loss of generality that $h(f) < l(g)$. A path from vertex v of $\hat{D}$ that always takes the leftmost outgoing arc (i.e., the first outgoing arc in the clockwise order around the vertex) will be called a *leftmost path* from v. A *rightmost path* is defined similarly. Let P_1 and P_2 be the leftmost and rightmost paths of $\hat{D}$ from $h(f)$ to t. Similarly, let P_3 and P_4 be the corresponding paths for $l(g)$. If there is a directed path of $\hat{D}$ from $h(f)$ to $l(g)$, we are done. Otherwise, either P_2 crosses P_3 (at a common vertex), or P_1 crosses P_4. For simplicity, we will consider only the first case. Let x be the first vertex at which P_2 and P_3 intersect (Fig. 6). Clearly, from Lemma 2, every arc incident to any vertex in path P_2 from the right side of P_2 is incoming. The same happens for the arcs incident to P_3 from the left. Because of the construction of D^*, there is a directed path in D^* from f to g. □

Theorem 1. *The algorithm W-VISIBILITY correctly computes a* w*-visibility representation of G.*

Proof. Since each horizontal segment has a distinct y-coordinate, no two horizontal segments intersect. Because of Lemma 1 and the assignment of y-coordinates to the horizontal segments, each face f of the w-visibility representation is a horizontally convex rectilinear polygon, i.e., the intersection of every horizontal line with f is either empty or consists of only one segment (see Fig. 5(b)). Furthermore, from steps 4 and 5, the vertical line with abscissa $\alpha(f)$ separates the two paths on the sides of f. Hence, it is impossible for two distinct edges of a face f to overlap in the w-visibility representation. Finally, considering Lemma 4, we can conclude that no two faces of $\hat{D}$ intersect in the representation constructed by the algorithm, except for the common edges. Therefore, the algorithm computes a correct w-visibility representation of the graph G. □

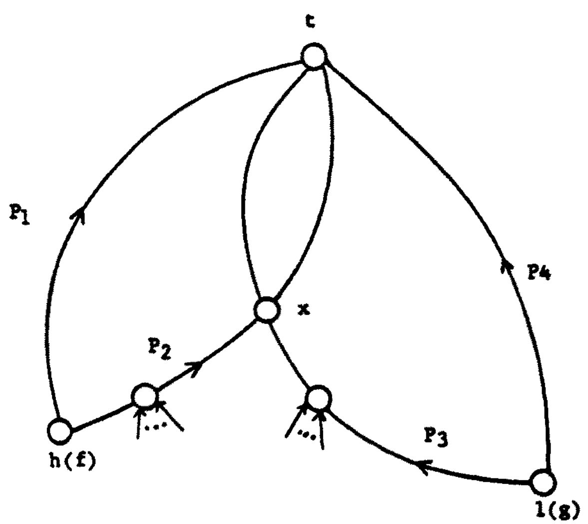

Fig. 6. Directed paths in the proof of Lemma 4.

We now discuss the time complexity of the algorithm *W-VISIBILITY*. Step 1 takes constant time. Using Even and Tarjan's algorithm [4] step 2 can be performed in time $O(|V|+|E|)$. By suitably modifying Hopcroft and Tarjan's planarity testing algorithm [6] step 3 take $O(|V|)$ time. The critical path method of step 4 has complexity $O(|E|)$. Step 5 takes time $O(|V|+|E|)$. Because of the planarity of G, $|E|=O(|V|)$. We thus have:

Theorem 2. *The overall time complexity of algorithm W-VISIBILITY is* $O(|V|)$. □

Notice that a more compact w-visibility representation can be obtained by applying the critical path method also to the graph G and by using the value $\alpha(v)$ as the y-coordinate of the vertex-segment v, for each vertex v in V. This modification does not affect the asymptotic computing time. See also [11].

The algorithm *W-VISIBILITY* can be extended to work for a 1-connected graph as shown below.

Algorithm *W-VISIBILITY2*
Input: A planar graph G.
Output: A w-visibility representation for G.
1. Find the blocks $B_1, \ldots, B_m$ of the graph G. Let $T := \{B_1, \ldots, B_m\}$ and $S := \emptyset$.
2. Construct a w-visibility representation for B_1;
   $T := T-\{B_1\}$;
   $S := S\cup\{B_1\}$;
3. **while** $T \neq \emptyset$ **do**
   let $B_{c_1}, \ldots, B_{c_k}$ be all the blocks of T which have a cutpoint c in common with some blocks in S, i.e. $(\bigcap_{i=1}^{k} B_{c_i}) \cap S=\{c\}$;
   find a w-visibility representation for each B_{c_i} using algorithm *W-VISIBILITY*, where in step 1 c is chosen to be the source vertex s;
   scale down the above representations in such a way that they all fit on the top of the vertex-segment corresponding to c in the w-visibility representation already constructed for S;

$$T := T - \bigcup_{i=1}^{k} \{B_{c_i}\};$$

$$S := S \cup \left(\bigcup_{i=1}^{k} \{B_{c_i}\}\right);$$

endwhile

We can summarize the results of this section in the following theorem.

Theorem 3. *A graph admits a* w-*visibility representation if and only if it is planar. Furthermore, a* w-*visibility representation for a planar graph can be constructed in linear time.* □

4. ε-Visibility Representation

In this section we present a complete characterization of the class of graphs that admit an ε-visibility representation. Moreover, we give linear time algorithms for testing the existence of and for constructing an ε-visibility representation of a planar graph. The following lemma provides a necessary condition for the existence of an ε-visibility representation.

Lemma 5. *If the graph G admits an ε-visibility representation, then there exists a planar embedding $\hat{G}$ of G such that all cutpoints appear on the boundary of the external face.*

Proof. Let Γ be an ε-visibility representation for G. Construct $\hat{G}$ by shrinking every vertex-segment of Γ into a point, and bending the edge-segments in order to maintain the adjacencies (Fig. 7). Suppose, for a contradiction, that there is a cutpoint c that does not appear on the boundary of the external face. Then there are blocks $B_0, B_1, \dots, B_m$ in G such that the embedding of $B_1, \dots, B_m$ in $\hat{G}$ lies entirely inside an internal face f of the embedding of B_0, and every path from a vertex of any B_i, $i=1,\dots,m$, to the rest of the graph $\bar{G}=G-\bigcup_{i=1}^{m} B_i$ passes through c (Fig. 8). Let σ be a segment of some B_j, $j\neq 0$, such that σ is distinct from c and is either the topmost or the bottommost segment of all B_i, $i=1,\dots,m$. Since face f is internal, segment σ is visible by some segment τ of $\bar{G}$ distinct from c. Hence, there is an edge (σ, τ) connecting $\bar{G}$ and B_j, which contradicts the fact that c is a cutpoint. □

The algorithm *W-VISIBILITY* described in the previous section can be extended in order to construct an ε-visibility representation for any 2-connected planar graph G (Fig. 9).

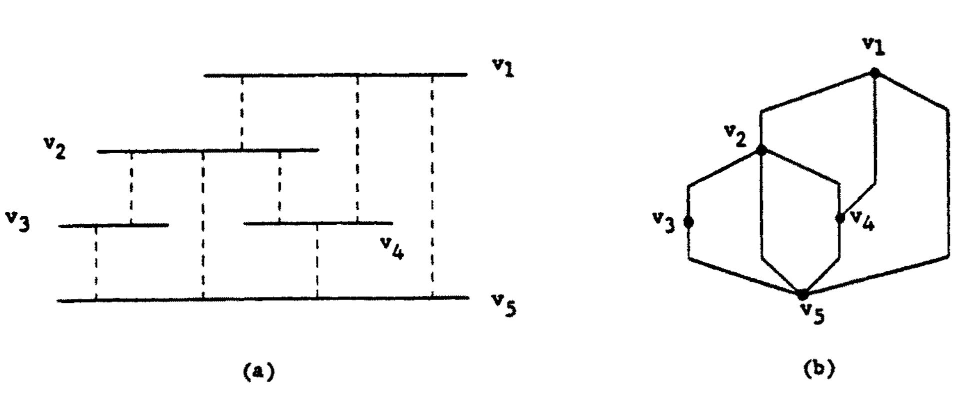

Fig. 7. (a) An ε-visibility representation. (b) A planar embedding constructed from the representation in (a).

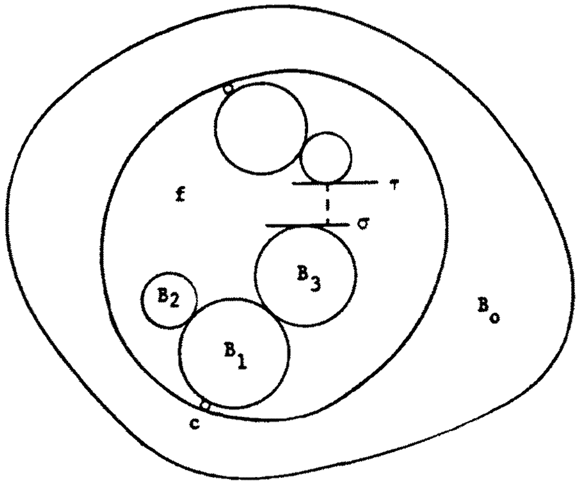

Fig. 8. Arrangement of blocks in the proof of Lemma 5.

Algorithm *ε-VISIBILITY*
Input: A 2-connected planar graph G.
Output: An ε-visibility representation for G such that each vertex- and edge-segment has endpoints with integer coordinates.
1. Compute a w-visibility representation Γ for G using algorithm *W-VISIBILITY.*
2. **for** each internal face f of Γ **do begin**
   2.1. let Λ and Ψ be the sets of vertex segments on the left and right side of f, excluding $l(f)$ and $h(f)$, respectively;
   2.2. **for** each $\lambda \in \Lambda$ **do**
   extend λ moving its right endpoint to the abscissa $\alpha(f)$;
   2.3. **for** each $\psi \in \Psi$ **do**
   extend ψ moving its left endpoint to the abscissa $\alpha(f)$;

end

The correctness of the algorithm stems from the following considerations:

(1) Each vertex-segment has a distinct y-coordinate.
(2) For each internal face f, segments of Λ and Ψ lie on the left and right side of the vertical line with abscissa $\alpha(f)$.
(3) For each internal face f without the edge $(l(f), h(f))$, vertices $l(f)$ and $h(f)$ are no longer ε-visible inside f.

From Theorem 2, step 1 takes $O(|V|)$ time. In step 2 each vertex-segment v is considered at most $\deg(v)$ times, once for every internal face in which it appears. Hence, step 2 has complexity $O(\sum_{v\in V} \deg(v)) = O(|E|) = O(|V|)$. From the above discussion we have:

Theorem 4. *Algorithm ε-VISIBILITY correctly computes an ε-visibility representation of a 2-connected planar graph $G=(V, E)$ in time $O(|V|)$.* □

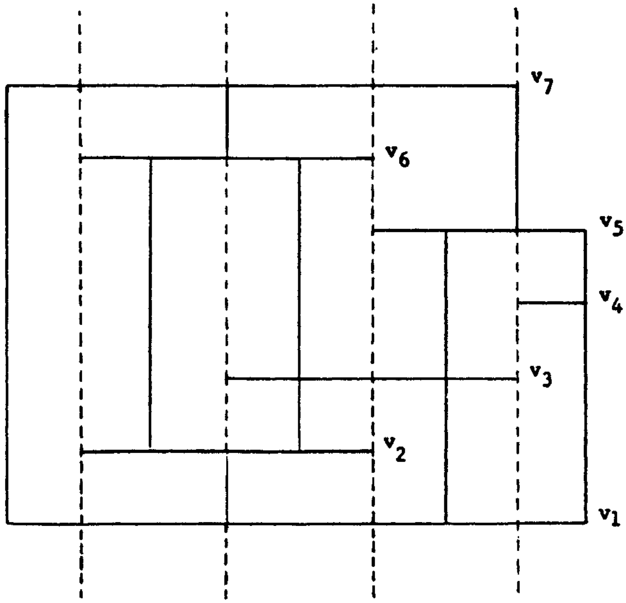

Fig. 9. The ε-visibility representation constructed from the w-visibility representation of Fig. 5(b).

Since every 2-connected planar graph admits an ε-visibility representation, one might question whether the necessary condition given in Lemma 5 for the existence of this representation is also sufficient. The answer is affirmative.

Lemma 6. *Let $\hat{G}$ be a planar embedding of a separable graph $G = (V, E)$ such that every cutpoint of G appears on the external face of $\hat{G}$. Then G admits an ε-visibility representation that can be constructed in time $O(|V|)$.*

Proof. Let B_i, $i = 1, \dots, k$, be the blocks of G that have only one cutpoint c_i, $i = 1, \dots, k$, in common with the rest of G, i.e., the B_i's are the leaves of the block-cutpoint tree of G. Let v_i be a vertex of B_i distinct from c_i, appearing on the external face of $\hat{G}$, $i = 1, \dots, k$. We construct the graph G' from G by adding a new vertex x and connecting it to all the vertices v_i, $i = 1, \dots, k$. G' is 2-connected and planar. Hence, from Theorem 4, it admits an ε-visibility representation. In particular, consider the one, Γ', produced by algorithm ε-*VISIBILITY* when choosing vertex x as the topmost vertex-segment. By removing x from Γ', we obtain an ε-visibility representation for G. The above transformation can clearly be performed in linear time. □

Note. For every boundary circuit C of $\hat{G}$, there exists another planar embedding $\tilde{G}$, of the same graph G, which has the same boundary circuits, but C is external in $\tilde{G}$. Therefore, Lemma 6 still holds if the cutpoints of G lie all in some internal face of $\hat{G}$.

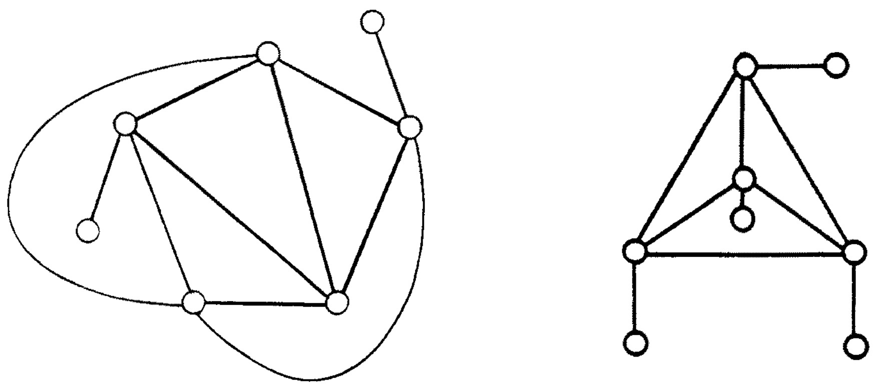

(a) (b)

Fig. 10. Examples of planar graphs that do not admit an ε-visibility representation.

Figure 10 shows two examples of planar graphs that do not admit an ε-visibility representation. Note that the graph in Fig. 10(a) is the planar graph with the minimum number of vertices that does not admit an ε-visibility representation. From Lemma 5 and Lemma 6 we obtain a complete characterization of the class of graphs that admit an ε-visibility representation.

Theorem 5.[2] *A graph G admits an ε-visibility representation if and only if there is a planar embedding $\hat{G}$ for G such that all cutpoints of G appear on the boundary of the same face.* □

The following equivalent characterization may be conveniently used in order to test in linear time whether a graph G admits an ε-visibility representation.

Corollary 1. *Let G' be the graph obtained from G by adding a new vertex x and connecting it to all cutpoints of G. Then G admits an ε-visibility representation if and only if G' is planar.* □

5. Strong-Visibility Representation

In this section we present necessary and sufficient conditions for the existence of an s-visibility representation. We also give efficient algorithms for the construction of this representation in the case of maximal planar graphs and 4-connected planar graphs.

From the results of Section 3, one can immediately derive that:

Theorem 6. *Every maximal planar graph $G=(V,E)$ admits an s-visibility representation that can be computed in time $O(|V|)$.* □

[2] This result has been independently discovered by the authors [14] and by Wismath [18].

Furthermore, one could use an argument similar to the proof of Lemma 5 to prove the following result.

Lemma 7. *If the graph G admits an* s-*visibility representation, then there exists a planar embedding $\hat{G}$ of G such that all cutpoints appear on the boundary of the external face.* □

However, the above necessary condition is not always sufficient to guarantee the existence of an s-visibility representation. In fact, there are 2-connected graphs that do not admit an s-visibility representation. Consider for example the graph $K_{2,4}$ shown in Fig. 11. The reason for this is given in the next theorem.

Theorem 7. *Let G be a 2-connected planar graph that has a separation pair of nonadjacent vertices v and w. If the removal of v and w separates G in at least four components, then G does not admit an* s-*visibility representation.*

Proof. Let $C_1, \ldots, C_k$, $k \geq 4$, be the connected components of G with respect to the separation pair v, w. Suppose, for a contradiction, that G admits an s-visibility representation Γ. We consider two cases for the vertex-segments v and w:

Case 1. There is a vertical band β such that v and w lie on opposite sides of β.

Any component C_i must have some vertex-segment that intersects a vertical line λ inside band β (Fig. 12(a)). Hence, there must be at least $k-1$ edge-segments between the C_i's, which is a contradiction.

Case 2. Otherwise.

Now let β be the vertical band of the plane consisting of the vertical lines crossing both v and w, and λ_1 and λ_2 be the leftmost and rightmost lines, respectively. We define ρ_i, $i = 1, 2$, as the ray of λ_i with origin an endpoint of v and w, and not intersecting the other segment. Clearly, exactly one component occupies the part of β that lies between v and w, because v and w must not be visible. Furthermore, any other component must have a vertex-segment intersecting either one of ρ_1 and ρ_2 (Fig. 12(b)). Since there are at least four components, there are at least two components which have a vertex-segment intersecting the same ray. Therefore, there exist at least two visible vertex-segments which belong

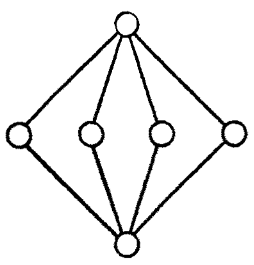

Fig. 11. Example of a planar biconnected graph that does not admit an s-visibility representation.

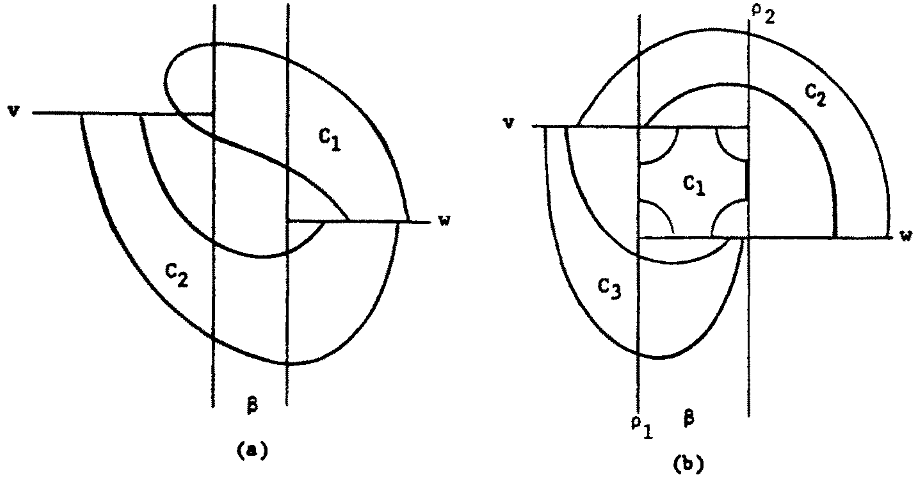

Fig. 12. Connected components with respect to the separation pair v, w in the proof of Theorem 7.

to distinct components. In other words, there is at least one edge between two vertices of distinct components, which is again a contradiction. □

In the rest of this section we present a complete characterization of the class of graphs that admit an s-visibility representation. Moreover, we show that all 4-connected planar graphs admit an s-visibility representation which can be computed in $O(|V|^3)$ time.

Recall that each face of an ε-visibility representation consists of two chains of vertex-segments and edge-segments between its topmost and bottommost vertex-segments. The s-visibility representation imposes further restrictions on the shape of the internal faces, i.e., for each internal face of an s-visibility representation, there is an edge-segment connecting the topmost and bottommost vertex-segments (see Fig. 1(d)).

Let D be the digraph induced by some st-numbering ξ on the 2-connected planar graph G. We say that ξ is a *strong st-numbering* if there is a planar embedding $\hat{D}$ of D such that s and t appear on the boundary of the external face, and for every internal face f of $\hat{D}$, the vertices $l(f)$ and $h(f)$ are joined by the arc $[l(f), h(f)]$.

Theorem 8. *A 2-connected graph G admits an* s*-visibility representation with bottommost vertex-segment s and topmost vertex-segment t if and only if there is a strong st-numbering for G.*

Proof.

Only If. Let Γ be an s-visibility representation for G. We can assume without loss of generality that each vertex-segment of Γ has a distinct y-coordinate. From the previous discussion on the shape of faces in a s-visibility representation, it is easy to see that a strong st-numbering can be obtained by assigning numbers from 1 to $|V|$ to vertices, according to the vertical ordering of the corresponding vertex-segments.

If: Let ξ be a strong *st*-numbering for G. If s and t are not adjacent, we add a new edge (s, t). For the resulting graph, ξ is still a strong *st*-numbering. We then apply the algorithm *ε-VISIBILITY*, where we replace step 2.2. with step 2.2′ shown below, using the *st*-numbering ξ and the associated planar embedding $\hat{D}$.

2.2′. **for** each $\lambda \in \Lambda$ **do if** f contains the arc $[l(f), h(f)]$
then extend λ moving its right endpoint to the abscissa $\alpha(f)$
else extend λ moving its right endpoint to the abscissa $\alpha(f)-\frac{1}{2}$;

Finally, if the edge (s, t), is not in G, we remove the corresponding edge-segment, and cut the vertex-segments s and t at the abscissa $\alpha(s^*)=0$. The result of this construction is an s-visibility representation for G. □

An *st-extension* of G is a 2-connected planar graph G' obtained from G by adding two new vertices s and t, the edge (s, t) and edges connecting s and t to G. Combining Lemma 7 and Theorem 8 we have a complete characterization of the class of graphs that admit an s-visibility representation.

Corollary 2. *A graph G admits an* s*-visibility representation if and only if there exists an st-extension G' of G that admits a strong st-numbering.* □

Now, we give some results that show a connection between Hamiltonian paths and s-visibility representations.

Theorem 9. *Let G be a 2-connected planar graph, and Γ a planar embedding of G. If there is a Hamiltonian path between two vertices s and t of G that lie on the boundary of the external face, then G admits an* s*-visibility representation.*

Proof. Let $P=(s=v_1, v_2, \ldots, v_n=t)$ be a Hamiltonian path between s and t. We claim that $\xi(v_i)=i$, for all $i=1, \ldots, n$, is a strong *st*-numbering for G.

Clearly, ξ is an *st*-numbering. Modify the embedding Γ so that P becomes a straight line and each internal face lies either on the left side or on the right side of P (Fig. 13). This can be done without modifying the face boundaries. Suppose, for a contradiction, that there exists an internal face f that does not contain the edge $(l(f), h(f))$. Then, by Lemma 1, there are two distinct paths from $l(f)$ to $h(f)$ that must have at least one vertex in common with the subpath of P from $l(f)$ to $h(f)$. Since these two paths must lie on the same side of P, we obtain a contradiction with the planarity of G. Therefore, we conclude from Theorem 8 that G admits an s-visibility representation. □

Corollary 3. *Let G be a 1-connected planar graph, and Γ a planar embedding of G. If there is a Hamiltonian path between two vertices s and t of G that lie on the boundary of the external face, then G admits an* s*-visibility representation.*

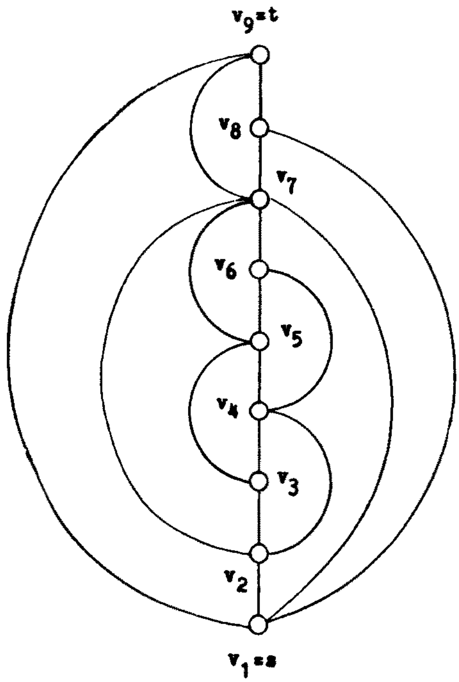

Fig. 13. Modification of the embedding Γ in the proof of Theorem 9.

Proof. We need only consider the case in which G is not 2-connected. Transform G into a 2-connected graph G' by adding the edge (s, t). From Theorem 9, G' admits an s-visibility representation. In particular, consider the representation obtained applying the construction of Theorem 8. Then we can remove from it the edge-segment corresponding to (s, t) and cut the vertex-segments corresponding to s and t as discussed in the proof of Theorem 8. The result is an s-visibility representation for G. □

Every 4-connected planar graph has a Hamiltonian cycle [17] which can be computed in time $O(|V|^3)$ [5]. We thus have:

Corollary 4. *Every 4-connected planar graph $G = (V, E)$ admits an s-visibility representation which can be computed in time $O(|V|^3)$.* □

6. Conclusions

We have derived new results on visibility representations of graphs, where vertices are represented by horizontal segments, and edges by vertical segments joining adjacent vertices. Specifically we have presented:

(1) A linear time algorithm for constructing a w-visibility representation of a planar graph.

(2) A complete characterization of the class of graphs that admits an ε-visibility representation, and a linear time algorithm for deciding whether a given graph admits one.
(3) A linear time algorithm for constructing ε-visibility representations.
(4) A complete characterization of the class of graphs that admits an s-visibility representation.
(5) Efficient algorithms for constructing s-visibility representations for maximal planar graphs and 4-connected planar graphs.

Although now all the three classes of graphs have been completely characterized, there are still open problems on the s-visibility representation:

(1) Is the necessary condition of Theorem 7 also sufficient for the existence of an s-visibility representation?
(2) Is there an efficient algorithm for deciding whether a given graph admits an s-visibility representation?
(3) Is there an efficient algorithm for constructing s-visibility representations in the general case?

Acknowledgments

We wish to thank Franco Preparata for suggesting the connection between Hamiltonian paths and visibility, and Doug West for useful discussions.

References

1. J. A. Bondy and U. S. R. Murty, *Graph Theory with Applications*, North Holland, Amsterdam, 1976.
2. P. Duchet, Y. Hamidoune, M. Las Vergnas, and H. Meyniel, Representing a planar graph by vertical lines joining different levels, *Discrete Math.* **46** (1983), 319-321.
3. S. Even, *Graph Algorithms*, Computer Science Press, Rockville, MD, 1979.
4. S. Even and R. E. Tarjan, Computing an st-numbering, *Theoret. Comput. Sci.* **2**, (1976), 339-344.
5. D. Gouyou-Beauchamps, The Hamiltonian circuit problem is polynomial for 4-connected planar graphs, *SIAM J. Comput.* **11** (1982), 529-539.
6. J. Hopcroft and R. Tarjan, Efficient planarity testing, *J. Assoc. Comput. Mach.* **21** (1974), 549-568.
7. A. Lempel, S. Even, and I. Cederbaum, An algorithm for planarity testing of graphs, in *Theory of Graphs* (Proceedings of an International Symposium, Rome, July 1966), (P. Rosenstiehl, ed.), 215-232, Gordon and Breach, New York, 1967.
8. F. Luccio, S. Mazzone, and C. K. Wong, A note on visibility graphs, Manuscript, Pisa, 1983.
9. L. A. Melnikov, *Problem at the Sixth Hungarian Colloquium on Combinatorics*, Eger, 1981.
10. R. H. J. M. Otten and J. G. van Wijk, Graph representations in interactive layout design, *Proceedings IEEE International Symposium on Circuits and Systems*, 914-918, New York, 1978.
11. P. Rosenstiehl and R. E. Tarjan, Rectilinear planar layouts of planar graphs and bipolar orientations, *Discrete Comput. Geom.*, to appear.
12. M. Schlag, F. Luccio, P. Maestrini, D. T. Lee, and C. K. Wong, A visibility problem in VLSI layout compaction, in *Advances in Computing Research*, Vol. 2 (F. P. Preparata, ed.), 259-282, JAI Press Inc., Greenwich, CT, 1985.
13. J. A. Storer, On minimal node-cost planar embeddings, *Networks* **14** (1984), 181-212.
14. R. Tamassia and I. G. Tollis, Plane Representations of Graphs and Visibility Between Parallel Segments, Technical Report ACT-57, Coordinated Science Laboratory, University of Ilinois at Urbana-Champaign, Urbana, IL, 1985.

15. R. Tamassia and I. G. Tollis, A Provably Good Linear Algorithm for Embedding Graphs in the Rectilinear Grid, Technical Report ACT-64, Coordinated Science Laboratory, University of Illinois at Urbana-Champaign, Urbana, IL, 1985.
16. C. Thomassen, Plane representations of graphs, in *Progress in Graph Theory* (J. A. Bondy and U. S. R. Murty, eds.), 43-69, Academic Press, New York, 1984.
17. W. T. Tutte, A theorem on planar graphs, *Trans. Amer. Math. Soc.* **82** (1956), 99-116.
18. S. K. Wismath, Characterizing bar line-of-sight graphs, *Proceedings of the ACM Symposium on Computational Geometry*, 147-152, Baltimore, Maryland, 1985.

Received July 9, 1985, and in revised form May 8, 1986.