



Universal 3-Dimensional Visibility Representations for Graphs

Helmut Alt, Michael Godau, Sue Whitesides

► To cite this version:

Helmut Alt, Michael Godau, Sue Whitesides. Universal 3-Dimensional Visibility Representations for Graphs. RR-2622, INRIA. 1995. inria-00074064

HAL Id: inria-00074064

<https://inria.hal.science/inria-00074064>

Submitted on 24 May 2006

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Universal 3-Dimensional Visibility Representations for Graphs

Helmut Alt Michael Godau , Sue Whitesides

N° 2622

Juillet 1995

PROGRAMME 4



***apport
de recherche***



Universal 3-Dimensional Visibility Representations for Graphs

Helmut Alt* Michael Godau** , Sue Whitesides***

Programme 4 — Robotique, image et vision
Projet Prisme

Rapport de recherche n° 2622 — Juillet 1995 — 18 pages

Abstract: This paper studies 3-dimensional visibility representations of graphs in which objects in 3-d correspond to vertices and vertical visibilities between these objects correspond to edges. We ask which classes of simple objects are *universal*, i.e. powerful enough to represent all graphs. In particular, we show that there is no constant k for which the class of all polygons having k or fewer sides is universal. However, we show by construction that every graph on n vertices can be represented by polygons each having at most $2n$ sides. The construction can be carried out by an $O(n^2)$ algorithm. We also study the universality of classes of simple objects (translates of a single, not necessarily polygonal object) relative to cliques K_n and similarly relative to complete bipartite graphs $K_{n,m}$.

Key-words: Computational geometry, Visibility, Graphs.

(Résumé : *tsvp*)

Research of H. Alt was supported in part by ESPRIT Basic Research Actions 7141 (ALCOMII). Work of S. Whitesides was written while the author was visiting Freie Universität Berlin. Her research supported by NSERC and FCAR grants.

*Institut für Informatik, Freie Universität Berlin, Takustr. 9, 14195 Berlin, Germany. alt@inf.fu-berlin.de.

**Freie Universität Berlin godau@inf.fu-berlin.de

***INRIA, B.P.93, 06902 Sophia-Antipolis cedex (France), and School of Computer Science, McGill University, 3480 University St. #318, Montréal, Québec H3A 2A7 Canada. sue@cs.mcgill.ca

Visibilité en 3 dimensions, une représentation universelle des graphes

Résumé : Cet article étudie la représentation des graphes par des relations de visibilité en trois dimensions. Les objets 3d correspondent aux sommets du graphe et les visibilités verticales aux arêtes. Nous recherchons quelles classes d'objets simples sont *universelles* c'est-à-dire suffisamment puissantes pour représenter tous les graphes. En particulier, nous montrons qu'il n'y a pas de constante k pour laquelle la classe des polygones ayant k côtés ou moins soit universelle. Nous montrons cependant que les graphes à n sommets peuvent être représentés par des polygones à $2n$ côtés. La construction est effectuée en temps $O(n^2)$. Nous étudions également l'universalité de classes d'objets simples (translations d'un même objet, non nécessairement polygonal) en regardant les représentations des cliques K_n et des graphes bipartis $K_{n,m}$.

Mots-clé : Géométrie algorithmique, Visibilité, Graphes.

1 Introduction

This paper considers 3-dimensional visibility representations for graphs. Vertices are represented by 2-dimensional objects floating in 3-d parallel to the x, y -plane (these objects can be swept in the z direction to form thick objects if desired). There is an edge in the graph if, and only if, the objects corresponding to its endpoints can see each other along a thick line of sight parallel to the z -axis. A thick line of sight is a tube of arbitrarily small but positive radius whose ends are contained in the objects. Throughout this paper, we use the term “visibility representation” to refer to this particular model.

The corresponding notion of 2-dimensional visibility has received wide attention due to its applications to such areas as graph drawing, VLSI wire routing, algorithm animation, CASE tools and circuit board layout. See [DETT] for a survey on graph drawing in general; for 2-dimensional visibility representations, see for example [DH], [TT], [KKU], [W].

Exploration of 3-dimensional visibility is still in the early stages. From the point of view of geometric graph theory, it is natural to consider visibility representations of graphs in dimensions higher than 2. From the point of view of visualization of graphs, it is basic to ask whether 3-dimensional representations give useful visualizations. For a 3-dimensional representation to be useful for visualization, it should be powerful enough to represent all graphs, or at least basic kinds of graphs. This motivates us to ask which classes of objects are *universal*, i.e., can give visibility representations for all graphs, or all graphs of a given kind?

The visibility representation considered in this paper has also been studied in [SS] (an abstract was presented at GD’92), in [Rom], and in [FW]. In these papers, the objects representing vertices are axis-aligned rectangles, or disks, and the properties of graphs that can be represented by these objects are studied. By contrast, this paper begins with families of graphs (all graphs, or all graphs of a specific kind), and explores simple ways to represent all graphs in the family.

Section 2 considers which translates of a given, fixed figure are universal for cliques K_n and complete bipartite graphs $K_{m,n}$. Section 3 uses counting arguments based on arrangements to show that no class of polygons having at most some fixed number k of sides is strong enough to represent all graphs. Section 4 shows that every graph on n vertices has a visibility representation by polygons each of which has at most $2n$ sides. These sections also contain additional results not listed here in the introduction.

2 Graphs realizable by translates of a figure

In this section we will investigate which complete and which complete bipartite graphs can be realized as visibility graphs of *translates* of one fixed figure. Here a *figure* is defined as an open bounded set whose boundary is a Jordan-curve. We say that a graph G can be *realized* by a figure F iff G is the visibility graph of translates of F . It will turn out that with many figures arbitrary complete graphs can be realized whereas each figure can only realize a finite number of stars, i.e. complete bipartite graphs of the form $K_{1,n}$.

2.1 Complete graphs

The realization of complete graphs K_n by translates of special figures like squares and disks has been investigated by Fekete, Houle, and Whitesides [FHW] and by Bose et al. [SS]. In [FHW] it was shown that K_7 can be realized by a square, whereas any K_n , $n \geq 8$ cannot. On the other hand, any K_n can be realized by a disk. We will consider more general figures in the following theorem.

First, we need the following definitions:

A curve C is called *strictly convex*, iff for any two points $p, q \in C$ the interior of the line segment \overline{pq} does not intersect C . We say that a figure F has a *local roundness* if there is some open set U such that $U \cap \partial F$ is a strictly convex curve.

Theorem 2.1 a) Any K_n can be realized by any nonconvex polygon.

b) For any convex polygon P there is an $n \in \mathbb{N}$ such that no K_m , $m \geq n$ can be realized by P .

c) To any K_n there is a convex polygon realizing it.

d) Any figure F with a local roundness can realize any K_n .

Proof:

a) We first observe that the figure in Fig.1 can realize any K_n . If P is a nonconvex polygon then it has at least one nonconvex vertex. Arranging copies of P in a neighborhood of this vertex as in Fig. 1 realizes any K_n .

b) (Sketch) Let P_1, \dots, P_k be a sequence of (projections of) translates of a convex

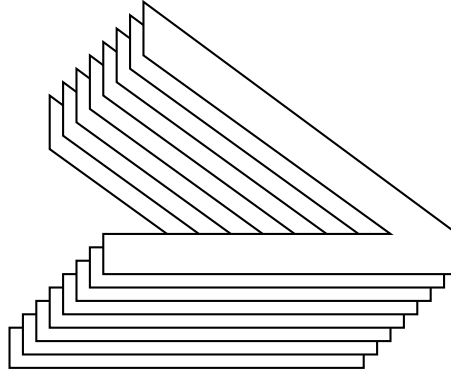


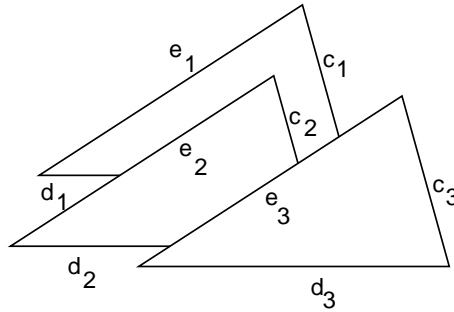
Figure 1: Realization of an arbitrary K_n with a nonconvex polygon

n -gon ordered by increasing z -coordinates, e_1, \dots, e_k the corresponding translates of one edge, and H_i the halfplane bounded by the straight line through e_i which contains P_i , $i = 1, \dots, k$. We define a linear order on e_1, \dots, e_k (more precisely, on the set of lines passing through them) by: $e_i \leq e_j \iff H_i \supseteq H_j$. By geometric considerations it can be shown:

Claim: If P_1, P_2, P_3 are translates of a convex polygon realizing K_3 , then not all sequences e_1, e_2, e_3 of translates of one edge can be monotone in the above order.

For example in Fig. 2 e_1, e_2, e_3 is monotone increasing, d_1, d_2, d_3 monotone decreasing, but c_1, c_2, c_3 is not monotone.

Now let $f(k) = (k - 1)^2 + 1$ for $k \in \mathbb{N}$ and let for $n \in \mathbb{N}$ $N := f^n(3)$ (n -fold iteration of f ; actually $N = 2^{2^n} + 1$). Using an argument from [SS] we will show that K_N cannot be realized by any convex n -gon. Suppose otherwise and let e^1, \dots, e^n be the edges and P_1, \dots, P_N the translates of the n -gon. Since $N = (f^{n-1}(3) - 1)^2 + 1$ by the Theorem of Erdős-Szekeres [ES] the sequence e_1^1, \dots, e_N^1 of translates of edge e^1 has a monotone subsequence of length $f^{n-1}(3)$. Considering the corresponding subsequence of polygons it must have a subsequence of length $f^{n-2}(3)$ where both the e^1 - and e^2 -sequences are monotone. Iterating this process

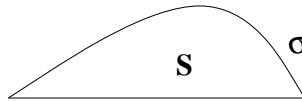
Figure 2: Triangles realizing K_3 .

we would obtain a subsequence of length $f^0(3) := 3$ where all edge-sequences are monotone in contradiction to the claim above.

c) follows from the fact that any K_n can be realized by disks and any disk can be approximated to arbitrary precision by convex polygons.

d) Consider a nondegenerate segment of F' 's boundary that is strictly convex. We can select a suitable subsegment σ with the following property: If l is the straight line through σ 's endpoints then no line perpendicular to l intersects σ in more than one point.

Assume also w.l.o.g. that l is horizontal, so σ looks as in Fig. 3.

Figure 3: Curve segment σ

Let S be the convex figure bounded by σ and the line segment between its endpoints. We will show by an inductive construction:

Claim: For any K_n there exists a realization by n translates S_1, \dots, S_n of S with the following properties:

- i) Let S'_1, \dots, S'_n be the projections of S_1, \dots, S_n into the xy -plane. There exists a horizontal straight line g such that all the horizontal segments of S'_1, \dots, S'_n lie strictly below g .
- ii) There is a visibility for any pair $S_i, S_j, i \neq j$ strictly above g .
- iii) Let s_{ij} be the intersection point of S'_i, S'_j . For $i = 1, \dots, n-1$ some nondegenerate part c_i of S_i 's boundary and some part of its interior are visible from $z = \infty$ in any neighborhood of $s_{i,n}$.
- iv) The z -coordinate of S_i is i for $i = 1, \dots, n$.

The claim is obviously true for $n = 1$.

Suppose now by inductive hypothesis that we positioned S_1, \dots, S_n satisfying the claim. We choose some point p on the boundary of S_n to the right of all $s_{1,n}, \dots, s_{n-1,n}$ as intersection point $s_{n+1,n}$ (see Fig. 4). Now we position S_{n+1} in the plane $z = n+1$ as follows:

First we put it exactly over S_n . Then we move it upwards slightly so that i) is still correct. Then it is moved to the left until it intersects S_n in p (see Fig. 4). The total motion can be made arbitrarily small, in fact, small enough so that iii) is still satisfied. ii) is satisfied by part iii) of the inductive hypothesis since s_{n+1} covers all points $s_{1,n} \dots s_{n-1,n}$.

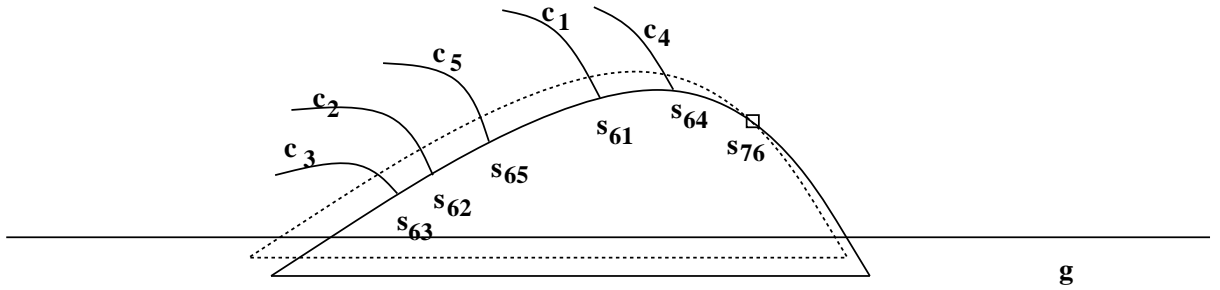


Figure 4: Construction of S_7 .

□

2.2 Complete Bipartite Graphs

[FHW] considers the realization of complete bipartite graphs by unit disks and unit squares. It is shown that $K_{2,3}$ and $K_{3,3}$ can be realized but $K_{j,3}$, $j \geq 4$ cannot. Here we will consider translates of more general convex objects and in particular the realization of stars $K_{1,n}$. In fact, we will show:

- Theorem 2.2**
- a) $K_{1,5}$ but no $K_{1,n}$, $n \geq 6$ can be realized with parallelograms.
 - b) If B is a strictly convex body $K_{1,6}$ but no $K_{1,n}$, $n \geq 7$ can be realized by B .
 - c) To any figure F there exists an $n \in \mathbb{N}$ such that for all $k \in \mathbb{N}$, $k \geq n$ $K_{1,k}$ is not realizable by F .
 - d) To any $K_{n,m}$ there exists a quadrilateral realizing it.

Proof:

a) A realization of $K_{1,5}$ by parallelograms is quite straightforward. $K_{1,n}$ $n \geq 6$ is not possible since one square cannot intersect 5 or more disjoint squares of the same size.

b) (Sketch) Here we use some results from convexity theory obtained by Hadwiger [H] and Grünbaum [G]. In fact, they showed that at most 8 translates of a convex body B in two dimensions can touch B without intersecting it or each other. The number 8 is only achieved by parallelograms, otherwise it is 6. For strictly convex bodies we observe that the tangent rays from B separating two neighboring touching translates all point into different directions and their slopes form a strictly monotone sequence (see Figure 5). From this it is possible to conclude that if one of the translates is removed one can distribute the others so that they still touch B but not each other any more. Then they can be pushed slightly inward B and we have a realization of $K_{1,5}$. Placing another copy on the other side exactly over B gives a realization of $K_{1,6}$.

The impossibility of $K_{1,7}$ is derived with similar arguments from the fact that no 6 translates of B can intersect B without intersecting each other (see [G]).

c) Consider a realization of $K_{1,n}$ and its projection onto the $x - y$ -plane. Then no point of the plane can be covered by the projections of more than three of the figures. Furthermore the figure representing the center of the star must be intersected by all others, so all projections must lie within in a circle whose diameter is at most three

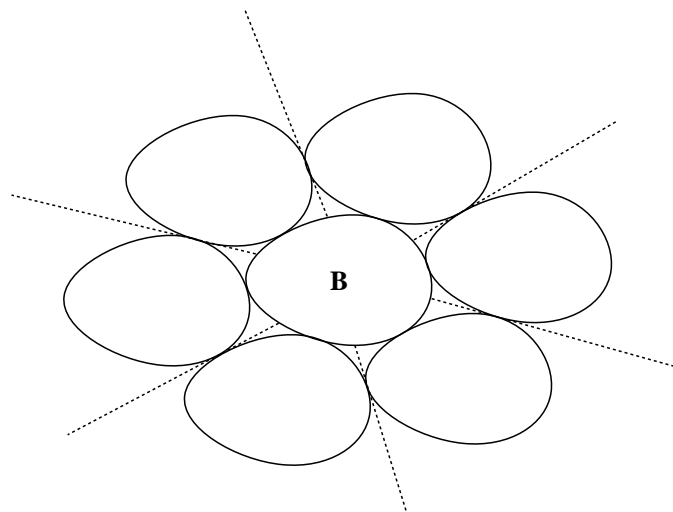


Figure 5: B touched by 6 of its translates.

times the diameter of F . These two properties imply that the number of figures is limited.

d) The construction is shown in Fig. 6.

□

3 An upper bound on the number of graphs representable by k -gons

In this section we will show that not each graph has a visibility representation by k -gons for some fixed $k \in \mathbb{N}$. In fact, we will even see that in order to represent all graphs with n vertices by polygons, some of those must have more than $\lfloor \frac{\alpha n}{\log n} \rfloor$ vertices for some constant $\alpha > 0$.

Definition 3.1 *A graph is said to be k -representable iff there is a visibility representation with (not necessarily convex) simple polygons each having at most k vertices.*

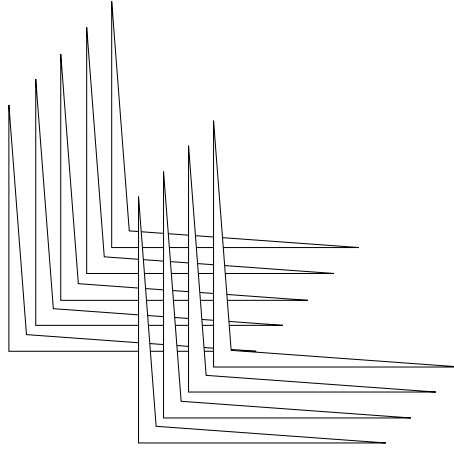


Figure 6: Realization of $K_{4,5}$ by quadrilaterals

The interesting fact that for every k there is a graph which is not k -representable follows from the following theorem.

Theorem 3.1 *There is an $\alpha > 0$ and there are graphs G_2, G_3, G_4, \dots such that G_n has n vertices and is not $\lfloor \frac{\alpha n}{\log n} \rfloor$ -representable.*

The theorem follows quite easily from the following lemma.

Lemma 3.2 *There is a β such that for all n, k there can be at most $2^{\beta nk \log(nk)}$ many graphs with a fixed vertex set $V = \{v_1, \dots, v_n\}$ which are k -representable.*

Proof: We consider an arbitrary k -representable graph $G = (V, E)$ with $V = \{v_1, \dots, v_n\}$. It can easily be shown that if G is k -representable then there exists a representation by polygons P_1, \dots, P_n parallel to the xy -plane with $m_1, \dots, m_n \leq k$ corners and z -coordinates $z_1, \dots, z_n \in \{1, \dots, n\}$.

Consider the projections of all the polygons into the xy -plane. Extend each edge s of each polygon to a line l_s obtaining a family \mathcal{L} of not necessarily distinct straight lines. For each edge s let h_s denote the open halfplane bounded by l_s lying on the same side of s as the polygon to which s belongs. Thus we obtain a family \mathcal{H} of halfplanes. The *arrangement* \mathcal{A} of \mathcal{H} is obtained by considering the line arrangement of \mathcal{L} and attaching to each cell the set of halfplanes in \mathcal{H} containing it.

It can be shown by a geometric argument, that from the triple $(\mathcal{A}, \vec{m}, \vec{z})$ where $\vec{z} := (z_1, \dots, z_n)$ and $\vec{m} := (m_1, \dots, m_n)$ the graph G can be reconstructed uniquely. Consequently, the number of representable graphs cannot be greater than the number of such triples.

Now we estimate the number of different arrangements of $m := m_1 + \dots + m_n < kn$ halfplanes h_0, \dots, h_{m-1} (which are not necessarily distinct). A family of m lines can be turned into a family of n points by dualization and the arrangement of the lines corresponds to the *order type* of the dual points. Alon [A] has shown that the number of different order types of m points is at most $2^{\gamma m \log m}$ for some constant $\gamma > 0$. (The points do not necessarily have to be in general position, not even to be distinct.) So the same bound holds for the number of different line arrangements. The number of different halfplane arrangements is bounded by 2^m times this bound. This is because we obtain a halfplane arrangement from the arrangement of the bounding lines by adding to each line l one bit of information telling which on of the two possible halfplanes is meant.

Altogether there are at most

$$\begin{array}{lll} 2^{\gamma m \log m} 2^m & \text{many different} & \text{halfplane arrangements } \mathcal{A} \\ k^n = 2^{n \log k} & \text{many different} & \vec{m} \\ n^n = 2^{n \log n} & \text{many different} & \vec{z} \end{array}$$

and thus at most $2^{\gamma m \log m + m + n \log k + n \log n}$ many triples of the form above. So this is also a bound on the number of different representable graphs and it is less than $2^{\beta m \log m}$ for an appropriate β . \square

Theorem 3.1 follows from the lemma since Since there are exactly $2^{\binom{n}{2}}$ graphs with vertex set V there are at least $2^{\binom{n}{2}}/n!$ (pairwise nonisomorphic) graphs with n vertices which is greater than $2^{\delta n^2}$ for some $\delta > 0$. Theorem 3.1 follows from this lower bound and Lemma 3.2.

On the other hand, every graph with n vertices is $(2n+1)$ -representable, which will be shown in the next section.

4 The Construction

This section gives a general construction which produces for any graph $G = (V, E)$ a 3-dimensional visibility representation for G . The construction can be carried out in a straight-forward manner by an algorithm that runs in $O(n^2)$ time, where n is the number of vertices of G . Each vertex is represented by a polygon of $O(n)$ sides (the polygons may differ in shape).

If desired, the basic construction can be modified easily and with the same time complexity to produce convex polygonal (or polyhedral) pieces. Furthermore, these pieces can be made to have all vertex angles of at least $\pi/6$. By using the technique of [CDR], it is also possible to implement the algorithm in $O(n^2)$ time with respect to a Turing machine model of computation.

4.1 The Basic Pieces

Let W denote a regular, convex $2n$ -gon centered at the origin O , and let w_1, w_2, \dots, w_{2n} denote the locations of its vertices. We use W to define the basic pieces representing the vertices of G . For this purpose, let X denote a regular, convex n -gon with vertices located at the odd-indexed vertices of W . Imagine adding triangular “tabs” to X to obtain W as follows. Call edge w_{2i-1}, w_{2i+1} of X *tab position* i , and for each i from 1 to n , add a triangle whose vertices are $w_{2i-1}, w_{2i}, w_{2i+1}$ to X at tab position i . W is X together with its tabs (see Fig. 7).

The pieces of our construction are obtained from X in a similar way, except that the tabs may vary in size. The construction may attach to tab position i of X a tab T_i with vertices w_{2i-1}, t_i, w_{2i+1} . Vertex t_i is called the *tab vertex* of T_i . In general, T_i lies inside the corresponding tab on W , with vertex t_i lying on the radial line through O and w_{2i} .

Definition 4.1 Let p_{2i} denote the point of intersection of the radial line through O and w_{2i} with the line through w_{2i-1} and w_{2i+1} . The size s_i of tab T_i is defined by $s_i = nd(t_i, p_{2i})/d(w_{2i}, p_{2i})$.

A tab of full size n has its tab vertex t_i positioned at w_{2i} .

We depth first search G , assigning to each vertex a number i indicating the order in which the search discovers the vertex. The i^{th} vertex discovered is represented by a polygon P_i consisting of a wedge-shaped portion of X with tabs of various sizes adjoined. See Fig. 8.

The bounding wedge of P_i is defined by two radial segments emanating from O , one to w_{2i-1} and the other to $w_{2(i+n_i)+1}$, $n_i \geq 0$. Between these radial segments, X has $1 + n_i$ tab positions. Each piece P_i has a tab of full size n at its lowest indexed tab position, i.e., at position i . Hence P_i has a tab vertex $t_i(P_i) = w_{2i}$. For $i < j \leq i + n_i$, the existence and location of the tab vertex $t_j(P_i)$ of $T_j(P_i)$ depends on the size $s_j(P_i)$ assigned to $T_j(P_i)$.

The idea behind the construction is as follows. Realize a depth first search tree for G by polygonal pieces floating parallel to the x, y -plane. Arrange these pieces so

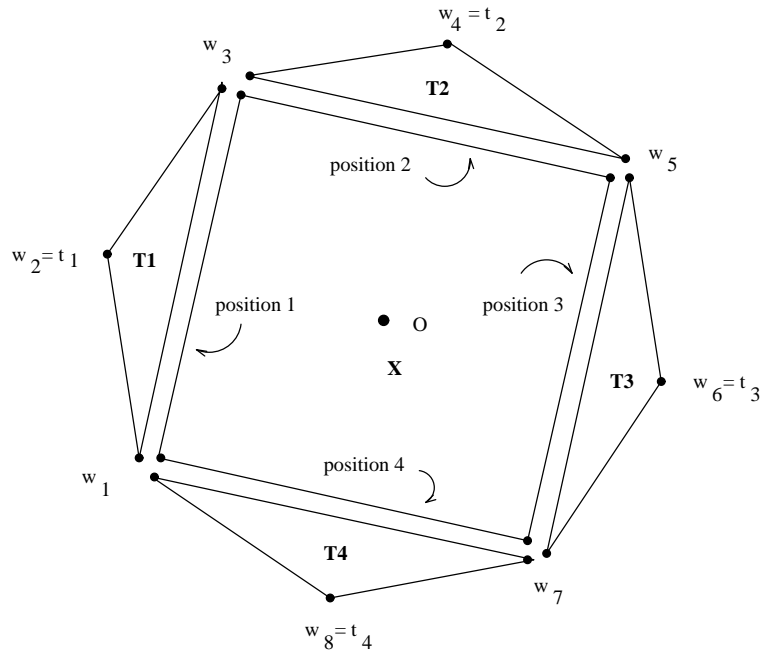
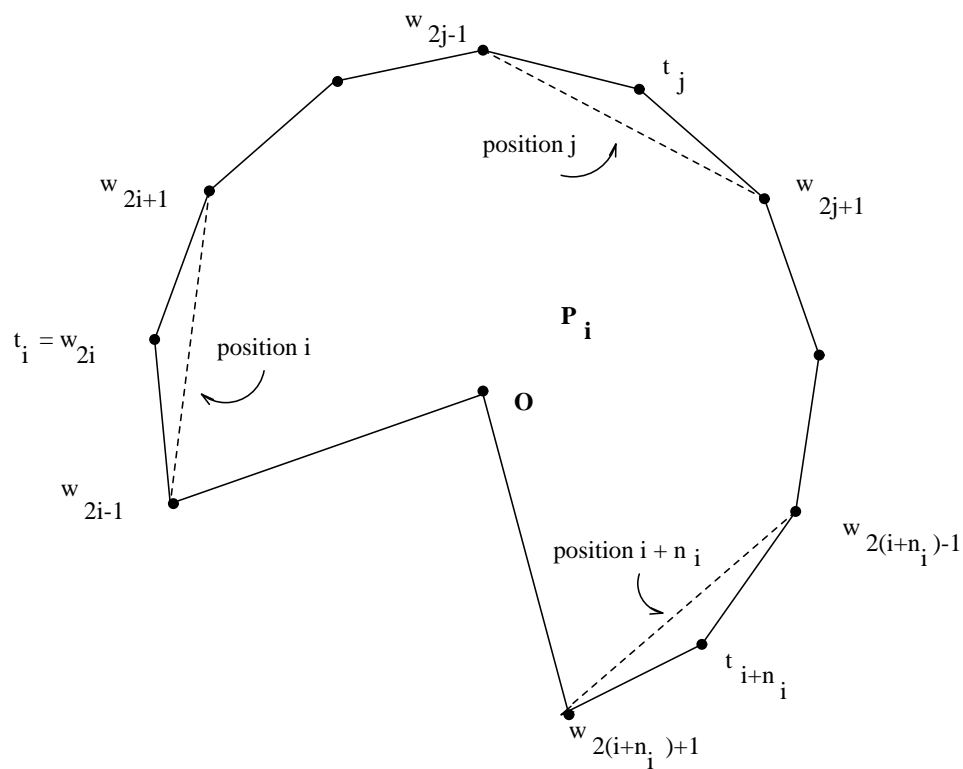


Figure 7: Regular n -gon X for $n = 4$ tabs.

Figure 8: Piece P_i .

that the piece $P(v)$ representing a vertex v lies above the pieces representing vertices in the subtree rooted at v , with the x, y -projection of $P(v)$ containing exactly the projections of the pieces $P(w)$ for which w belongs to the subtree rooted at v . Thus each piece has the possibility of seeing its ancestors and descendants, but nothing else.

Unless G itself is a tree, depth first search discovers back edges, i.e., edges of G that do not appear as tree edges in the depth first search tree. A familiar property of depth first search trees for graphs is that each back edge must connect an ancestor, descendant pair in the tree. The purpose of adding tabs of varying sizes is to control which ancestors and descendants see each other.

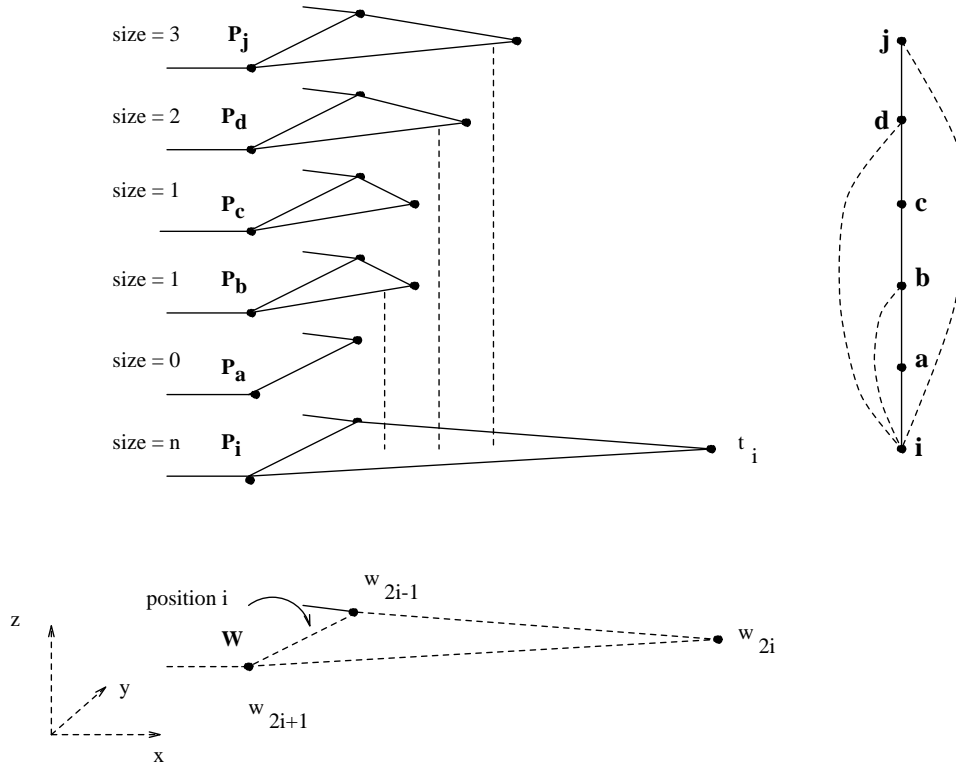


Figure 9: Back edges from i and their inverted staircase of tabs.

Suppose the depth first search tree has a back edge between i and ancestor j of i . Our construction creates a visibility between the tab T_i of full size n in position i on P_i and a tab in position i on P_j . See Fig. 9.

Of course there may be back edges in the tree joining i to k , where k lies on the path from i to its ancestor j . (Consider $k = b, c, d$ in the figures.) In this case, our construction creates a visibility between the tab in position i on P_k and the full sized tab in position i on P_i . Note that the visibility between the tabs in position i on P_k and P_j must be blocked if the graph G contains no edge between j and k . Hence, for example, the tabs in position i on P_b and P_j must be blocked from seeing each other by intervening tabs.

Blocking inappropriate visibilities between tabs is achieved by creating an inverted staircase of tabs above the tab of full size on P_i and the tab in position i on P_j . The tab on P_i has full size n . The tab in position i on the piece immediately above P_i is assigned size 0, as this piece sees P_i in any case. The tab on the next piece above P_i is also assigned size 0 unless there is a back edge from i to the vertex corresponding to this piece; in this case, the tab size is increased to 1. Tab size remains the same or increases with increasing integer z values. In fact, tab size increases precisely when P_i and the piece at the z value in question should be mutually visible. Thus the size of the tab in position i on P_j is equal to the number of back edges of the form i, k , where k lies on the path from i to j (possibly $k = j$).

Theorem 4.1 *Let G be a connected graph. The following assignment of parameters to the piece representing an arbitrary vertex v of G gives a 3-dimensional visibility representation for G :*

- v is assigned its depth first search order i ;
- the index n_i of v is set equal to the number of descendants of v in the depth first search tree;
- the tab $T_i(P_i)$ in position i on P_i is assigned size $s_i(P_i) = n$;
- for $i < j \leq i + n_i$ the size $s_j(P_i)$ of the tab $T_j(P_i)$ on P_i at position j is set equal to the number of nodes on the tree path from j , up to and including i , that receive a back edge from j ; and
- the z coordinate of P_i is set equal to 1 less than the z coordinate of its parent.

Proof:

(Sketch) A well-known property of depth first search ordering is that the descendants of v are numbered with consecutive integers, beginning with $i + 1$. Thus P_i has, in addition to a tab of full size at position i , a tab in position j for $1 < j \leq i + n_i$.

It is easy to check that the pieces are disjoint and that that P_i cannot see any P_k representing a vertex w unless w is either an ancestor or a descendant of v . Clearly, P_i sees its parent (if any) and all of its children.

Let us check that if the depth first search tree has a back edge from v to some ancestor u of v with depth first search number k , then P_i and P_k are mutually visible. P_k has a tab in position i . This tab aligns with the tab of full size in position i on P_i . Furthermore, the tab on P_k has size greater than the intervening tabs in position i , as the number of back edges from i on the path from i to k is at least one greater than the number of back edges on the path from i to k , up to but not including k . Hence P_i and P_k have a line of visibility between their tabs at position i . Thus all back edges are represented.

It can also be checked that no inappropriate visibilities are present.

□

Corollary 4.1 *The construction of Theorem 4.1 can be modified to produce convex pieces, fat pieces, polyhedral pieces, or pieces having any combination of these properties.*

Proof: To produce convex pieces, use a W with sufficiently many vertices ($12n$) that each piece has a vertex angle at O of at most $\Pi/6$. To produce fat pieces, move the vertex at O sufficiently close to the chord through the first and last vertices of P_i shared with W . To produce polyhedral pieces, take the cross product of P_i with a short line segment parallel to the z axis. □

It is straightforward to design an algorithm that, by computing tab sizes efficiently, carries out the construction in $O(n^2)$ time.

References

- [A] N. Alon, "The number of polytopes, configurations and real matroids", *Mathematika* 33 (1986), pp. 62–71
- [SS] P. Bose, H. Everett, S. Fekete, A. Lubiw, H. Meijer, K. Romanik, T. Shermer and S. Whitesides, "On a visibility representation for graphs in three dimensions," in *Snapshots of Computational and Discrete Geometry*, v. 3, eds. D. Avis and P. Bose, McGill University School of Computer Science Technical Report SOCS-94.50, July 1994, pp. 2 - 25

- [CDR] J. Canny, B. Donald and E. K. Ressler, "A rational rotation method for robust geometric algorithms", *Proc. 8th ACM Symp. Comput. Geom.*, 1992, 251 - 260.
- [DH] A. Dean and J. Hutchison, "Rectangle-visibility representations of bipartite graphs," in *Proc. Graph Drawing '94*, Princeton, NJ, 1994. Lecture Notes in Computer Science LNCS v. 894, Springer-Verlag, 1995.
- [ES] P. Erdős, Gy. Szekeres, "A Combinatorial Problem in Geometry", *Compositio Math.* 2, 1935, 463-470.
- [DETT] G. Di Battista, P. Eades, R. Tamassia and I. Tollis, "Algorithms for drawing graphs: an annotated bibliography", *Computational Geometry Theory and Applications*, v. 4, 1994, 235 - 282. Also available from wilma.cs.brown.edu by ftp.
- [FHW] S. Fekete, M. Houle and S. Whitesides, manuscript, 1995.
- [G] B. Grünbaum, "On a Conjecture of H. Hadwiger", *Pacific J. Math.*, 11, 215-219.
- [H] H. Hadwiger, "Über Treffanzahlen bei translationsgleichen Eikörpern", *Arch. Math.*, Vol. VIII, 1957, 212-213.
- [KKU] E. Kranakis, D. Krizanc and J. Urrutia, "On the number of directions in visibility representations of graphs," in *Proc. Graph Drawing '94*, Princeton, NJ, 1994. Lecture Notes in Computer Science LNCS v. 894, Springer-Verlag, 1995, 167 - 176.
- [Rom] K. Romanik, "Directed VR-representable graphs have unbounded dimension," in *Proc. Graph Drawing '94*, Princeton, NJ, 1994. Lecture Notes in Computer Science LNCS v. 894, Springer-Verlag, 1995, 177 - 181.
- [TT] R. Tamassia and I. Tollis, "A unified approach to visibility representations of planar graphs," *Discrete Comput. Geom.* v. 1, 1986, 321 - 341.
- [W] S. Wismath, "Characterizing bar line-of-sight graphs," in *Proc. ACM Symp. on Computational Geometry*, 1985, 147 - 152.



Unité de recherche INRIA Lorraine, Technopôle de Nancy-Brabois, Campus scientifique,
615 rue du Jardin Botanique, BP 101, 54600 VILLERS LÈS NANCY
Unité de recherche INRIA Rennes, Irista, Campus universitaire de Beaulieu, 35042 RENNES Cedex
Unité de recherche INRIA Rhône-Alpes, 46 avenue Félix Viallet, 38031 GRENOBLE Cedex 1
Unité de recherche INRIA Rocquencourt, Domaine de Voluceau, Rocquencourt, BP 105, 78153 LE CHESNAY Cedex
Unité de recherche INRIA Sophia-Antipolis, 2004 route des Lucioles, BP 93, 06902 SOPHIA-ANTIPOLIS Cedex

Éditeur
INRIA, Domaine de Voluceau, Rocquencourt, BP 105, 78153 LE CHESNAY Cedex (France)
ISSN 0249-6399