Universal 3-Dimensional Visibility Representations for Graphs

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

B 95-14
November 1995

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 $K_{n,m}$.

*`alt@inf.fu-berlin.de`, Institut für Informatik, FU Berlin, Takustr. 9, 14195 Berlin, Germany.
This research was supported by the ESPRIT Basic Research Action No. 7141, Project ALCOM II.
**`godau@inf.fu-berlin.de`, Freie Universität Berlin, Germany
***`sue@cs.mcgill.ca`, School of Computer Science, McGill University, 3480 University St. #318, Montréal, Québec H3A 2A7 Canada. Written while the author was visiting INRIA-Sophia Antipolis and Freie Universität Berlin. Research supported by NSERC and FCAR grants.

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 xy-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 [BEF+] (an abstract of some of its results was presented at GD'92), in [Rom], and in [FHW]. 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 if and only if G is the visibility graph of translates of F. It will turn out, for example, that there are many figures that can realize all complete graphs. On the other hand, no figure can realize more than 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. [BEF+]. In [FHW] it is shown that K_7 can be realized by a square, whereas no K_n, $n \geq 8$, can be realized. 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* if and only if 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. A figure bounded by a strictly convex curve is a *strictly convex figure.*

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

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

c) For 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) Let $P_1, \ldots, P_k$ be a sequence of (projections of) translates of a convex n-gon ordered by increasing z-coordinate, let $e_1, \ldots, e_k$ be the corresponding translates of one edge, and H_i the halfplane bounded by the straight line through e_i which contains P_i, $i = 1, \ldots, k$. We define a linear order on $e_1, \ldots, e_k$ (more precisely, on the set of lines passing through them) by: $e_i \leq e_j \iff H_i \supseteq H_j$. First, we will show:
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 is monotone decreasing, but c_1, c_2, c_3 is not monotone.

To prove the claim, consider a point p (in the $xy-$plane) where P_1 and P_3 see each other. Then p lies outside (the projection of) P_2 and therefore there exists an edge c_2 of P_2 so that the straight line g through c_2 separates p from P_2. Let c_1, c_3

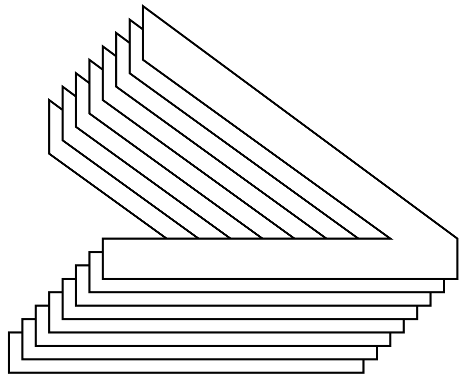

Figure 1: Realization of an arbitrary K_n with a nonconvex polygon

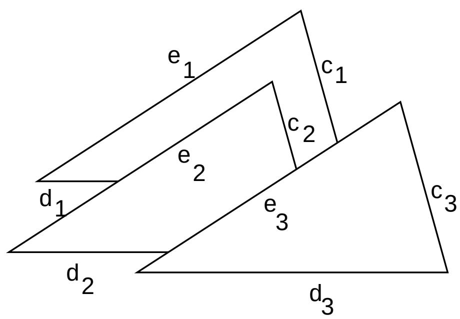

Figure 2: Triangles realizing K_3.

be the edges of P_1, P_3, respectively, corresponding to c_2. Assume a line parallel to g is being moved towards the scene from the outside. It will first meet P_1 and P_3 before it meets P_2 (or vice versa). Consequently, the order in which edges c_1, c_2, c_3 are met is not monotone.

For $n, k \in I\!N$, let $f(k) := (k-1)^2+1$ and let $N := f^n(3)$ (i.e., the n-fold iteration of $f(k)$ evaluated at $k := 3$; actually $N := 2^{2^n}+1$). Using an argument from [BEF+] we will show that K_N cannot be realized by any convex n-gon. Suppose otherwise and let $e^1, \ldots, e^n$ be the edges of the n-gon and $P_1, \ldots, P_N$ the translates of the n-gon. Since $N = (f^{n-1}(3) - 1)^2 + 1$, by the theorem of Erdös-Szekeres [ES] the sequence $e_1^1, \ldots, e_N^1$ of translates of edge e^1 has a monotone subsequence of length $f^{n-1}(3)$. The corresponding subsequence of polygons must have a subsequence of length $f^{n-2}(3)$ where both the e^1- and e^2-sequences are monotone. Iterating this process yields a subsequence of length $f^0(3) := 3$ where all edge-sequences are monotone in contradiction to the claim above. N can be reduced from doubly exponential to exponential in n using properties of edge colorings in graphs [F].

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

d) Consider a nondegenerate segment of the boundary of F that is strictly convex. We can select a suitable subsegment σ with the following property: if l is the straight line through the endpoints of σ, then no line perpendicular to l intersects σ in more than one point.
Assume also without limitation of generality that l is horizontal, so σ looks as in Fig. 3.

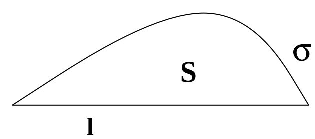

Figure 3: Curve segment σ

Let S be the closed 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 of S by n translates $S_1, \ldots, S_n$ with the following properties:

i) Let $S'_1, \ldots, S'_n$ be the projections of $S_1, \ldots, S_n$ into the xy-plane, and let $\sigma'_1, \ldots, \sigma'_n$ and $l'_1, \ldots, l'_n$ denote the pieces of the boundaries of these projections that arise from σ and l. There exists a horizontal line g such that all the $l'_1, \ldots, l'_n$ lie strictly below g.

ii) Any pair S_i, S_j, $i \neq j$, see each other along a line of sight that intersects the xy-plane strictly above g.

iii) For $1 \leq i < n$, the boundary pieces σ'_i and σ'_n have exactly one common intersection point above g. Let s_{in} denote this point, and let $D_{in}(\epsilon)$ denote the closed disk of positive radius ϵ centered at s_{in}. Consider the set $D_{in}(\epsilon) \cap S'_i \backslash S'_n$. For all sufficiently small $\epsilon > 0$, all points in S_i with x, y-projections in this set see upward to $z = \infty$.

iv) For $i = 1, \ldots, n$ the z-coordinate of S_i is i.

The claim is obviously true for $n = 1$.
Suppose now by inductive hypothesis that we positioned $S_1, \ldots, S_n$ satisfying the claim. We choose some point p on the boundary of S_n to the right of all $s_{1,n}, \ldots, s_{n-1,n}$ 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 (i.e. in positive y-direction) slightly so that i) is still correct. Then it is moved to the left until it intersects S_n at p (see Fig. 4). The total motion can be made arbitrarily small, in fact, small enough so that iii) is satisfied with n replaced by $n + 1$ and points s_{in}

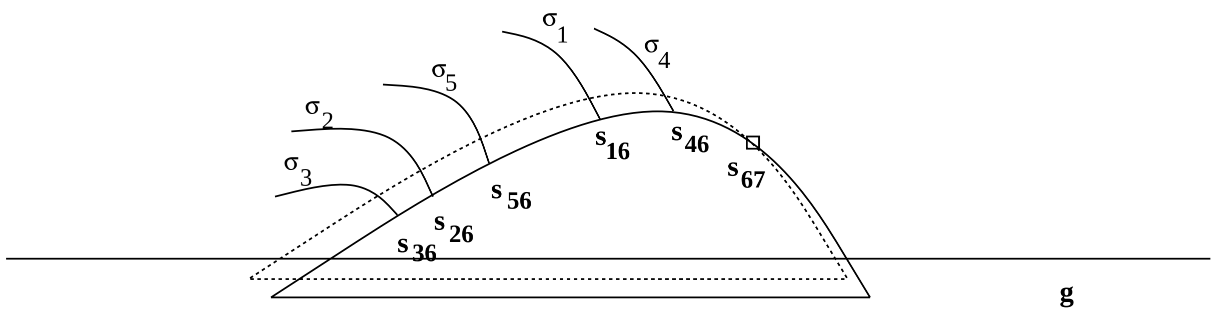

Figure 4: Construction of $S_{n+1} = S_7$.

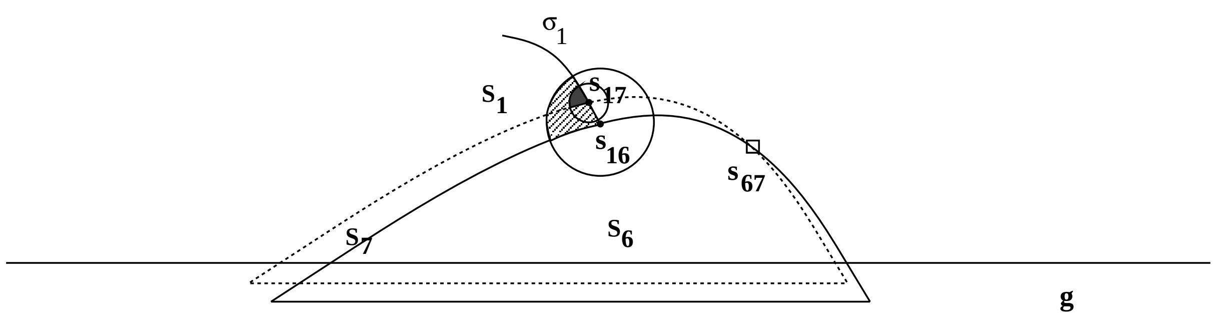

Figure 5: Visible parts of S_1 in neighborhoods of s_{16} and s_{17}.

replaced by points $s_{i,n+1}$ (see Fig. 4). Item ii) is satisfied by part iii) of the inductive hypothesis since S_{n+1} covers all points $s_{1,n} \dots s_{n-1,n}$.

□

2.2 Complete Bipartite Graphs

[BEF+] 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 claimed that $K_{j,3}$, $j \geq 4$ 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 then $K_{1,6}$ but no $K_{1,n}$, $n \geq 7$, can be realized by B.

c) For any figure F there exists an $n \in \mathbb{N}$ such that for all $k \in \mathbb{N}$ with $k \geq n$ $K_{1,k}$ is not realizable by F.

d) For any $K_{n,m}$ there exists a quadrilateral realizing it.

For the proof of the theorem we need the following lemma.

Lemma 2.1 *Let A be a strictly convex body and let A_1, A_2 translates of A such that A, A_1, A_2 pairwise touch each other (i.e., the boundaries intersect but not the interiors). Then for any sufficiently small $\varepsilon > 0$ A_2 can be translated by a vector t of length ε such that $A_2 + t$ still touches A but is disjoint from A_1.*

Proof: Assume without loss of generality that the origin $0 \in A$ and let $A_i = A + t_i, i = 1, 2$, so t_1, t_2 are reference points within A_1, A_2 corresponding to 0 within A. Define A' by the Minkowski sum $A' := A \oplus (-A)$ and define $A'_i := A' + t_i, i = 1, 2$. Then A', A'_1, A'_2 are also strictly convex. The fact that two of these figures, say A, A_1, touch is equivalent to the fact that the reference points $0, t_1$ lie on the boundaries $\partial A'_1, \partial A'$, respectively. So altogether we have the situation illustrated in Fig. 6. Because of their strict convexity the curves $\partial A'$ and $\partial A'_1$ intersect properly in t_2, so

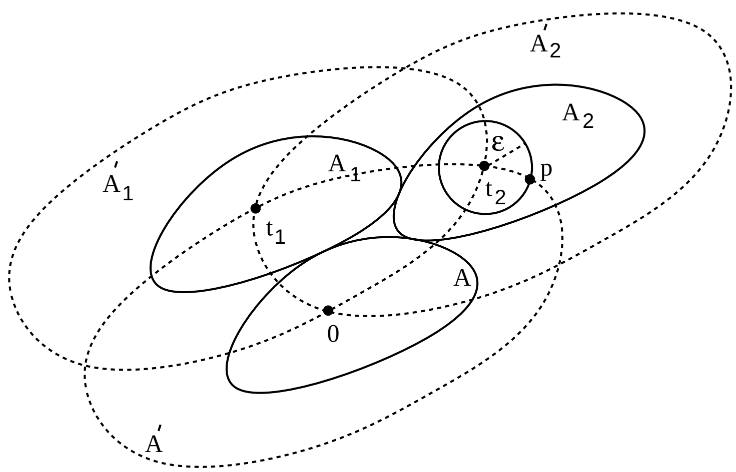

Figure 6: Three translates touching each other

any sufficiently small ε-circle around t_2 has an intersection point p with $\partial A' \setminus A'_1$. A translation of A_2 by $t = p - t_2$ then has the desired properties. □

Proof of Theorem 2.2:

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

b) 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 (see Fig. 7). Suppose one of the 6 outer translates is removed. Then we can apply Lemma 2.1 to one of the neighboring ones and move it away from its neighbor that is touching it. Repeating this process, we can adjust the five outer translates so that each still touches the

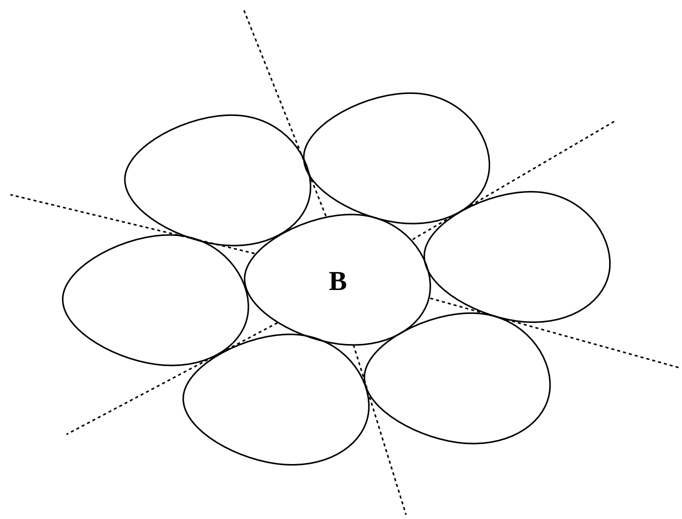

Figure 7: B touched by 6 of its translates.

inner one but no two outer ones touch or intersect each other. Clearly, it is then possible to push each of them slightly inward so that all properly intersect the inner one still without touching each other. Placing the five outer translates at, say, $z = 0$, the inner one at $z = 1$, and another one exactly above it at $z = 2$ realizes $K_{1,6}$.

To show the impossibility of $K_{1,7}$ we assume without loss of generality that the object B is closed. Suppose $K_{1,7}$, could be realized and let A be (the projection of) the copy of B realizing the central vertex. Then at most one of the other vertices can be realized by a translate of B having exactly the same projection. Otherwise, since the translate representing the central vertex would be covered from both sides by two other translates, any additional translate would either fail to see the translate for the central vertex or would see at least two translates. So there are (at least) six vertices whose representations have projections $A_1, ..., A_6$ different from A, but intersecting A. For $i = 1, ..., 6$, let $t_i \neq 0$ be the translation vector such that $A_i = A + t_i$. Further let $\lambda_i > 0$ be the unique positive number such that $C_i := A + \lambda_i t_i$ just touches A in one point.

Claim: $C_i \cap C_j = \emptyset$ for $i \neq j$.

In fact, we will show that there is a straight line separating C_i from C_j. Let $B_i := A_i \setminus A$ for all i. Then the interiors of $B_1, ..., B_6$ do not intersect. Even their convex hulls do not intersect, as easily can be seen. So for $i \neq j$ there is a straight line l separating B_i from B_j (see Fig. 8). Furthermore l must intersect the interior of A. Since l does not intersect the curve γ in Fig. 8 it cannot intersect C_i. Likewise it cannot intersect C_j, so it separates C_i and C_j.

By the claim we would have $C_1, ..., C_6$ all touching A but not two touching each, which is not possible by the results of Hadwiger and Grünbaum.

c) Consider a realization of $K_{1,n}$ and its projection into the $xy-$plane. Then no

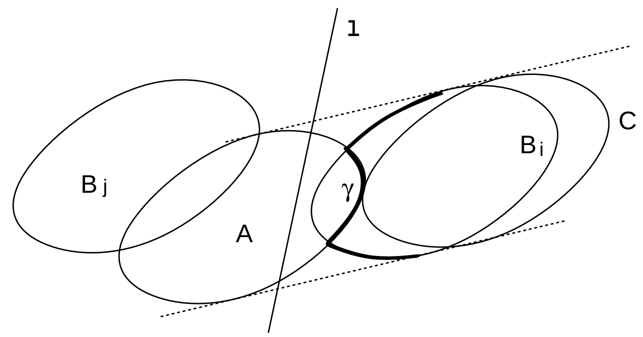

Figure 8: A line separating C_i and C_j.

point of the plane can be covered by the projections of more than three of the figures. Furthermore the projection of the figure representing the center of the star must be intersected by the projections of all the other figures, so all projections must lie within a circle whose diameter is at most three times the diameter of F. These two properties imply that the number of figures is limited by an area argument.

d) The construction is shown in Fig. 9.

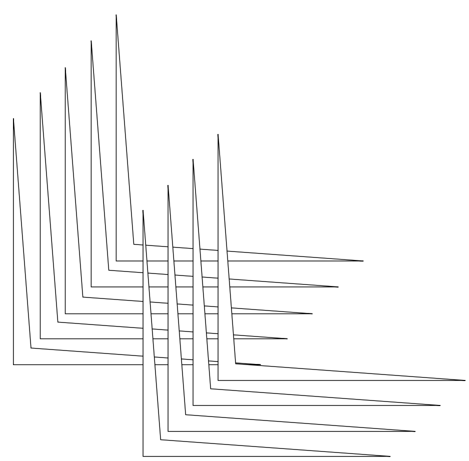

Figure 9: Realization of $K_{4,5}$ by quadrilaterals

□

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

In this section we will show that there is no fixed $k \in I\!N$ such that every graph has a visibility representation by k-gons. In fact, we will even see that there is a constant

$\alpha > 0$ such that in order to represent all graphs with n vertices by polygons, some of those polygons must have more than $\lfloor \frac{\alpha n}{\log n} \rfloor$ vertices.

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

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

Theorem 3.1 *There is an* $\alpha > 0$ *and there are graphs* $G_2, G_3, G_4, ..., G_n, ...$ *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* $V = \{v_1, ..., v_n\}$ *that are* k*-representable.*

Proof: We consider an arbitrary k-representable graph $G = (V, E)$ with $V = \{v_1, ..., v_n\}$. Obviously, if G is k-representable then there exists a representation by polygons $P_1, ..., P_n$ parallel to the xy-plane with at most k edges each. Without loss of generality we can assume that P_i has z-coordinate i for $i = 1, ..., n$.

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 at most $m := nk$ not necessarily distinct straight lines. Each edge s and, thus, each line l_s can be oriented by the convention that the polygon lies, say, left of s. Now, G can be uniquely identified by the information in the following items.

1. the *arrangement* of the lines in $\mathcal{L}$.

2. Each polygon $P_i, i = 1, ..., n$, is identified by the description of a counterclockwise tour around its boundary. In particular, the starting point s is given by a line $l \in \mathcal{L}$ containing it and by a number $n_0 \leq m$ meaning that s is the n_0^{th} intersection point when traversing l in the direction of its orientation. Then a sequence of at most k numbers $n_1, ..., n_r \in \{1, ..., m\}$ is given, meaning that the tour starts at s, goes straight on l for n_1 intersections, then turns into the oriented line crossing there, goes straight for n_2 intersections, etc. Clearly, this describes a tour within the arrangement.

Clearly, the information in the above items uniquely identifies the pairwise intersections of the projections of the polygons into the xy-plane. This together with the convention that P_i has z-coordinate equal to i makes it possible to determine all visibilities, and hence G itself.

It remains to count the number of different possibilities for the data in the above items:

1. As is well known (see [A]), the number of different arrangements of m oriented straight lines is at most $2^{\beta_1 m \log m}$ for some constant $\beta_1 > 0$.

2. For each polygon there are m possibilities for the starting line l, and at most m possibilities for each number $n_0, ..., n_r$, $r \leq k$. So the number of possibilities per polygon is bounded by m^{k+2}. Altogether, the number of possibilities is at most $m^{(k+2)n}$, which is at most $2^{\beta_2 m \log m}$ for some constant $\beta_2 > 0$.

Multiplying the upper bounds in 1 and 2 gives the desired total upper bound of $2^{\beta m \log m}$, where $\beta = \beta_1 + \beta_2$. □

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 more 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 $w_1, w_2, \ldots w_{2n}$ 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. 10).

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}.

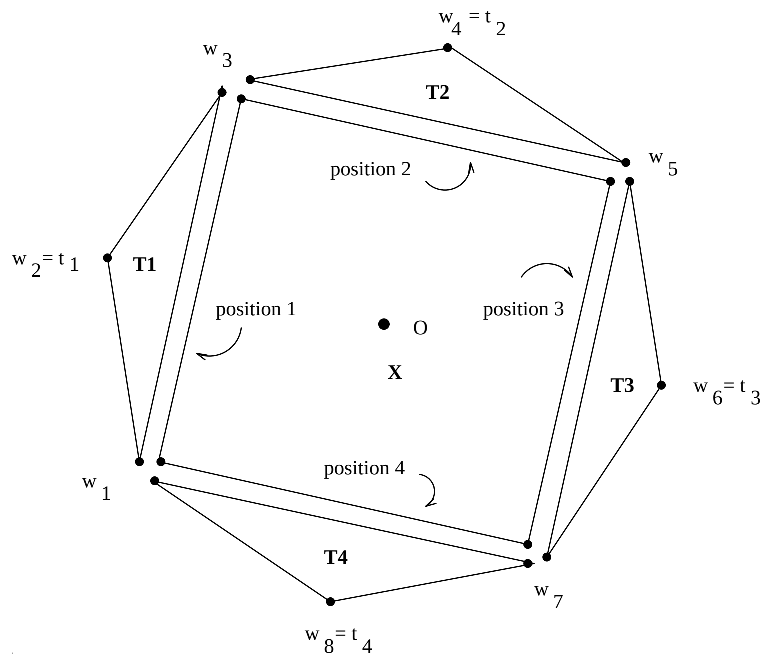

Figure 10: Regular n-gon X for $n = 4$ tabs.

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. 11.

The bounding wedge of P_i is defined by two radial segments emanating from O, one to w_{2i-1} and the other to $w_{2(i+n_i)+1}$, for some $n_i \geq 0$ to be determined. 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 tab $T_j(P_i)$ depends on the size $s_j(P_i)$ assigned to tab $T_j(P_i)$.

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 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

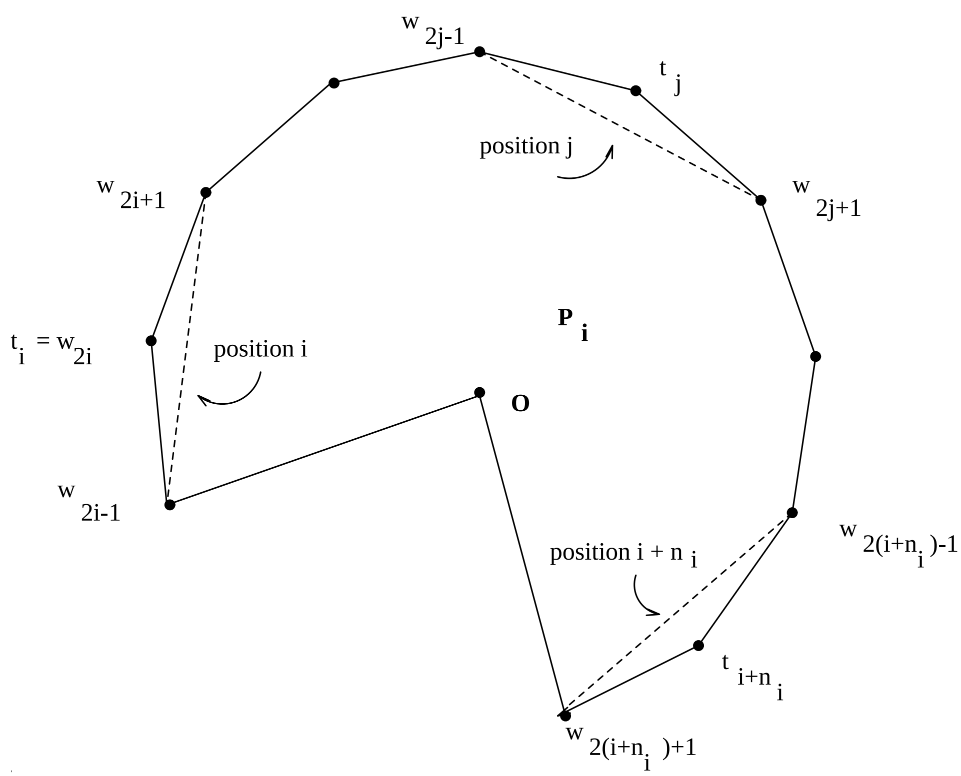

Figure 11: Piece P_i.

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.

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. 12.

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 in position i on P_i and the tab in position i on P_j. The tab in position 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

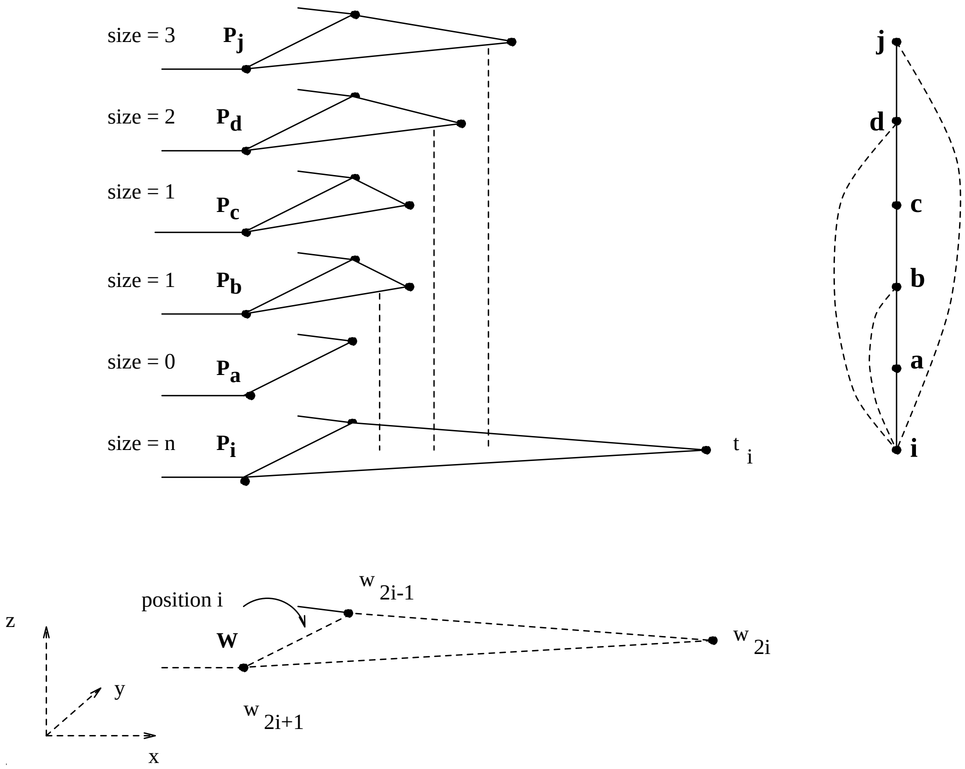

Figure 12: Back edges from i and their inverted staircase of tabs.

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 $k = j$).

Lemma 4.2 *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:

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 (possibly of size 0) in position j for $1 < j \leq i + n_i$.

It is easy to check that the pieces have disjoint interiors and that P_i representing a vertex v cannot see any P_k representing a vertex w unless w is either an ancestor or a descendant of v. (Note that if two pieces have the same parent, they are assigned the same z-coordinate and may share an edge. However, the pieces can be perturbed slightly to make all the pieces disjoint.) 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, where v is numbered i, to some ancestor u of v, where u is numbered 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.

Now we check that no inappropriate visibilites are present. Clearly pieces corresponding to vertices in disjoint subtrees do not even overlap in projection, so no visibilities occur between pieces that are not ancestor-descendant pairs. Now consider a vertex u, numbered k, and a vertex v, numbered i, where k is an ancestor of i but not the parent of i. Suppose there is no edge $(u, v) \in G$ but that pieces P_i and P_k are mutually visible. Clearly any visibility line must pass through some tab $T_j(P_i)$ on P_i and some corresponding tab $T_j(P_k)$ on P_k.

Suppose first that $j = i$. Of course tab $T_i(P_i)$ has full size. Because there is no back edge from i to k, and because k is not the parent of i, tab $T_i(P_k)$ has the same size (possibly 0) as the tab T_i of the piece immediately below P_k on the path of pieces between P_i and P_k. This piece blocks visibility between $T_i(i)$ and $T_i(k)$.

Now suppose that $j > i$. Then the tab T_j of the piece immediately above P_i in the path of pieces between P_i and P_k has size equal to or greater than the size of $T_j(P_i)$. Hence the tabs in position j on P_i and P_k are not visible to one another.

This completes the proof that no inappropriate visibilities occur, and hence the proof of the lemma. □

Now we can state the main result of this section.

Theorem 4.1 *Every graph on n vertices is $2n$-representable. Furthermore, a representation can be constructed in $O(n^2)$ time.*

Proof: If G is connected, the statement holds by Lemma 4.2. If G is not connected, a representation can be obtained by representing each connected component and then translating these representations so that their projections do not overlap.

It is straightforward to design an algorithm that runs in $O(n^2)$ time for carrying out the construction of Lemma 4.2. This can be done by modifying the usual depth-first search algorithm to compute the description of P_i at the time the search returns from i to the parent of i.

To facilitate the computation of P_i, a list B_i is maintained that records the number j of any vertex for which (j, i) is a back edge to i. When search of the subtree rooted at i has been completed, the value of n_i is set to the number of the most recently discovered vertex. The tab size of $T_i(P_i)$ is set to n. Then the remaining sizes for tabs on P_i are initialized to 0. The tab sizes of tabs on the children of P_i are copied to the sizes of the tabs in the same positions on P_i. Finally, the list B_i is processed. For each $j \in B_i$, the tab size for the tab in position j on P_i is increased by 1. The z-coordinate of P_i can be determined when i is first labeled, as it is equal to 1 less than the z-coordinate of the parent of P_i. Hence the computation of the description of P_i can be completed when the search is about to return from i to its parent. Each tab on P_i is computed in constant time. □

We can generalize our results as follows

Corollary 4.1 *The construction of Lemma 4.2 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. □

Acknowledgements

We would like to thank Stefan Felsner and Emo Welzl for helpful discussions and hints concerning this research.

References

[A] N. Alon, "The number of polytopes, configurations and real matroids", *Mathematica 33* (1986), pp. 62–71

[BEF+] 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.

[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.

[DH] A. Dean and J. Hutchinson, "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. Erdös, Gy. Szekeres, "A Combinatorial Problem in Geometry", *Compositio Math. 2*, 1935, 463-470.

[F] S. Felsner, personal communication, 1995.

[FHW] S. Fekete, M. Houle and S. Whitesides, "New results on a visibility representation of graphs in 3D," *in* Proc. Symposium on Graph Drawing (GD '95), Passau, Germany, 1995, Springer-Verlag LNCS series.

[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.