Journal of Graph Algorithms and Applications
http://www.cs.brown.edu/publications/jgaa/
vol. 2, no. 3, pp. 1–16 (1998)

A Visibility Representation for Graphs in Three Dimensions

Prosenjit Bose
Carleton University, Canada
jit@scs.carleton.ca

Hazel Everett
Université du Québec à Montréal, Canada
everett@info.uqam.ca

Sándor P. Fekete
Universität zu Köln, Germany
sandor@zpr.uni-koeln.de

Michael E. Houle
University of Newcastle, Australia
mike@cs.newcastle.edu.au

Anna Lubiw
University of Waterloo, Canada
alubiw@uwaterloo.ca

Henk Meijer
Queen's University, Canada
henk@qucis.queensu.ca

Kathleen Romanik
University of Maryland, USA
romanik@cfar.umd.edu

Günter Rote
Technische Universität Graz, Austria
rote@opt.math.tu-graz.ac.at

Thomas C. Shermer
Simon Fraser University, Canada
shermer@cs.sfu.ca

Sue Whitesides
McGill University, Canada
sue@cs.mcgill.ca

Christian Zelle
Technische Universität Graz, Austria
zelle@opt.math.tu-graz.ac.at

Abstract

This paper proposes a 3-dimensional visibility representation of graphs $G = (V, E)$ in which vertices are mapped to rectangles floating in $\mathbb{R}^3$ parallel to the x, y-plane, with edges represented by vertical lines of sight. We apply an extension of the Erdős-Szekeres Theorem in a geometric setting to obtain an upper bound of $n = 56$ for the largest representable complete graph K_n. On the other hand, we show by construction that $n \geq 22$. These are the best existing bounds. We also note that planar graphs and complete bipartite graphs $K_{m,n}$ are representable, but that the family of representable graphs is not closed under graph minors.

Communicated by M. T. Goodrich: submitted February 1996; revised February 1997.

Research on this paper was supported by FCAR, NSERC, NSF, and DFG.

1 Introduction

The problem of drawing or representing graphs has been studied extensively in the literature (see the paper of Di Battista et al. [11] for a survey of graph drawing research and its applications). In particular, determining a *visibility representation* of a graph, where the vertices of the graph map to disjoint sets or objects in the plane and the edges are expressed as visibility relations between these sets, has received considerable attention recently due to the large number of applications in areas such as VLSI wire routing, algorithm animation, CASE tools, and circuit board layout. However, visibility representations in 3-dimensions have received less attention. In this paper, we define a 3-dimensional visibility representation and study its properties. This paper combines experimental results with the results of two conference papers [4, 14] and a technical report [5], which have motivated a number of other conference papers such as [2, 7, 15, 17, 21]. Closely related questions have been examined in [10, 19]. This paper introduces the basic concepts and fundamental results in this area.

1.1 Previous results

We begin by reviewing some of the results in 2-dimensional visibility representations. To do this, we first discuss in more detail the most common types of visibility representations that have been studied.

A visibility representation is defined by specifying a class of objects to represent the vertices and the visibility relation between the objects. In two dimensions, common choices for objects are axis-aligned line segments or rectangles. Two visibility relations that have been considered previously are defined as follows:

- Two objects are (*mutually*) *visible* if and only if they can be joined by a line segment that does not intersect any other object. Often, the direction of the line segment is restricted to be axis-parallel.
- Two objects are ϵ-*visible* if and only if they can be joined by a family of parallel line segments such that no segment intersects any other objects, and the union of the segments covers a region of positive area, i.e., positive "width" ϵ. Again, the direction of the line segments is often restricted to be axis-parallel.

For example, in Figure 1, line segments 1 and 2 are visible but not ϵ-visible. Line segments 2 and 3 as well as line segments 1 and 3 are ϵ-visible. However, line segments 1 and 3 are not vertically ϵ-visible. Line segments 2 and 4 are not visible by either of the above notions of visibility. Although these two different types of visibility seem closely related, the restrictions imposed by ϵ-visibility often radically change the class of graphs that admit visibility representations. See [22] for a discussion of the different definitions of visibility and the sensitivity of results to the choice of visibility definition.

Now we can describe some previous results concerning 2-dimensional visibility representations. Wismath [24] and Tamassia and Tollis [22] independently

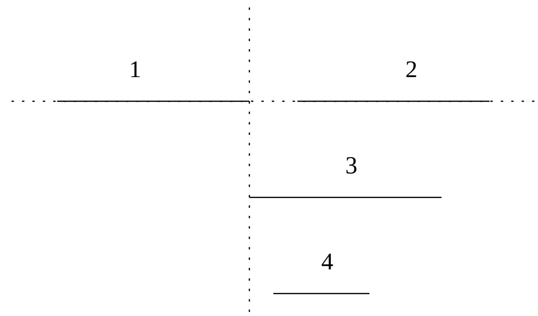

Figure 1: An illustration of visibility relations.

showed that every 2-connected planar graph admits a visibility representation where the vertices are represented by closed disjoint horizontal line segments in the plane and where two vertices are adjacent if and only if their corresponding segments are vertically ϵ-visible. (More precisely, they showed that such a representation exists for a graph if and only if the graph is planar and has an embedding with all cut vertices on the exterior face.) For the same model, Kant et al. [18] studied the visibility representations of trees. When the vertices are represented by disjoint axis-aligned rectangles in the plane and visibility is defined as ϵ-visibility in the horizontal and vertical directions, Wismath [24] showed that every planar graph admits a visibility representation. Dean and Hutchinson [10] proved that K_8 is the largest complete graph that admits a visibility representation in this model. For an overview of the various visibility representations studied, see [11].

1.2 3-Dimensional visibility representation

This paper studies a 3-dimensional visibility representation for graphs $G = (V, E)$. First we define this representation. Consider an arrangement of closed, disjoint rectangles in $\mathbb{R}^3$ such that the planes determined by the rectangles are perpendicular to the z-axis, and the sides of the rectangles are parallel to the x- or y-axes. Two rectangles R_i and R_j are *z-visible* if and only if between the two rectangles there is a closed cylinder C of positive length and radius such that the ends of C are contained in R_i and R_j, the axis of C is parallel to the z-axis, and the intersection of C with any other rectangle in the arrangement is empty. Througout this paper, a given graph $G = (V, E)$ is said to be *representable* if and only if its n vertices can be associated with n disjoint rectangles parallel to the x, y-plane in $\mathbb{R}^3$ such that vertices v_i and v_j are adjacent in G if and only if their corresponding rectangles R_i and R_j are z-visible.

One motivation for considering this type of 3-dimensional representation is the fact that, in addition to all planar graphs, many non-planar graphs can also be represented. Indeed we will exhibit a representation of K_{22}, but prove that complete graphs of size $n \geq 56$ are not representable.

The fact that a clique as large as K_{22} is representable may prove useful for graph visualization purposes. Given an arbitrary graph on $n \leq 22$ vertices, one can take a visibility representation (as defined above) of K_n as defined above and, for each edge $e \in E(G)$, draw in z-axis parallel lines of sight between the rectangles corresponding to the endpoints of e. This gives what is called a *weak* visibility representation, as G cannot be recovered from the positions of the rectangles[1]. However, weak representation is entirely appropriate when, for example, the purpose is to visualize the graph or to produce a VLSI layout.

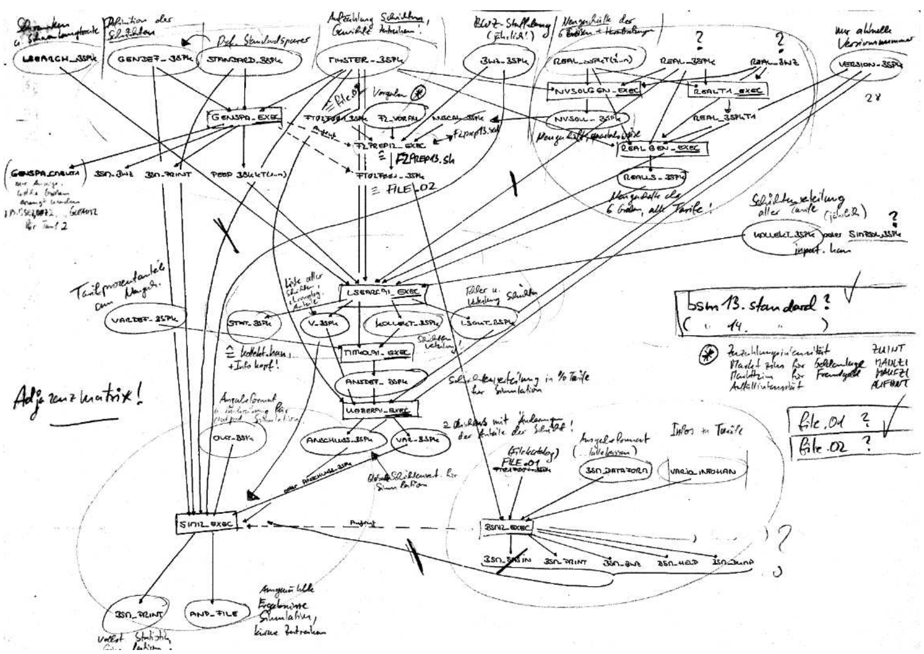

Figure 2: The structure of a software package, as drawn by one of its users.

Figures 2 and 3 suggest that for visualizing diagrams involving graphs whose vertices are associated with text and various other properties, a 3-dimensional visibility representation may prove useful. Figure 2 gives a hand-drawn diagram of a software package. (This package is used for analyzing and forecasting the savings of millions of customers of the German public loan banks. See [3].) Figure 3 gives a view from the top of a 3-dimensional visibility representation; in addition to the adjacency information, the size of a rectangle corresponds approximately to the importance of the files and routines it represents; input files are black, whereas output files are white. Furthermore, this type of representation permits taking advantage of the fact that the given directed graph is

[1] When the graph can be recovered from the geometry of its representation, the representation is sometimes called *strong*.

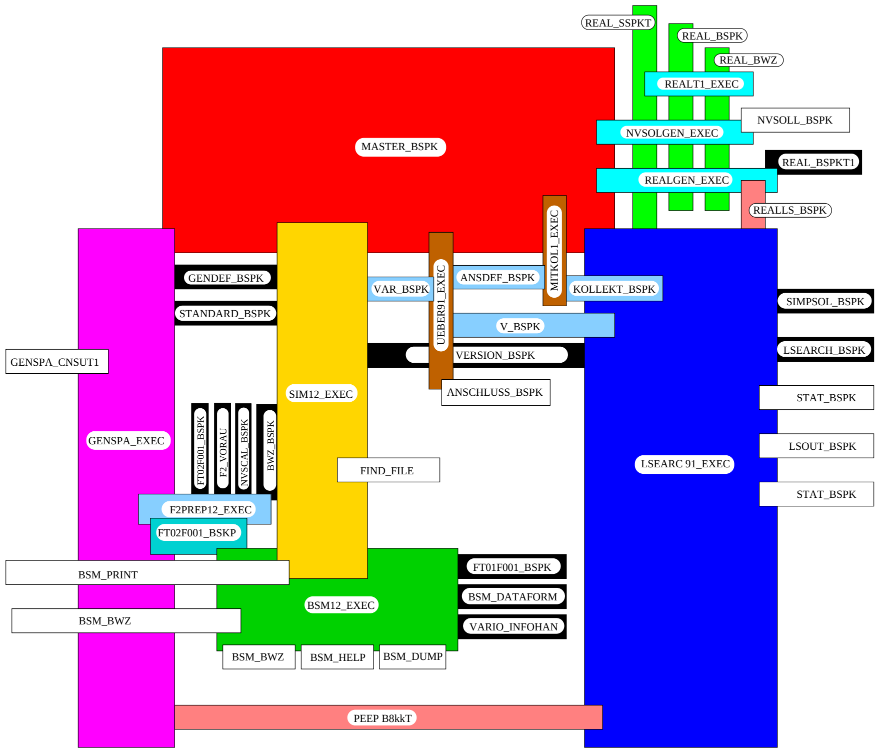

Figure 3: Top view of a representation of the software package.

acyclic: the edges are directed from lower rectangles to higher rectangles.

Alternatively, instead of viewing the representation from the top, on a workstation one could view the rectangles as boxes that have thickness in the z direction, so that text could be written on the sides of the box. The representation could then be examined from multiple viewpoints and could be explored by travelling around inside the picture.

We make no claim that the top view of the 3-dimensional representation shown in Fig. 3 is easier to understand than the handmade drawing of Fig. 2. However, we believe that the interplay between geometry and combinatorics demonstrated by the type of visibility representation defined in this paper makes such representations interesting in their own right. Furthermore, we believe that the possibility of drawing graphs in 3 dimensions, with vertices represented as solid objects such as boxes, is well worth further investigation.

The rest of this paper is organized as follows: In Section 2, we show that K_n does not have a representation for $n \geq 56$. On the other hand, we show in Section 3 that K_n has a representation for values of $n \leq 22$. In Section 4, we note that all planar graphs and all complete bipartite graphs are representable, but that the family of representable graphs is not closed under graph minors.

Section 5 concludes the paper.

2 An upper bound

In this section, we provide an upper bound on the size of the largest clique that can be represented. We show that K_n does not have a representation for any $n > 55$. This result, based on [14], improves the previous best known result of $n > 102$ from [4, 5].

2.1 Sequences of 4-tuples

We consider sequences of n rectangles lying parallel to the x, y-plane in $\mathbb{R}^3$, and ordered by increasing z-coordinate. We call a sequence *valid* if its associated visibility graph is K_n. Consider the projections of all the rectangles in a valid sequence onto the x, y-plane. Because each two must intersect and the objects are axis-aligned rectangles, application of a Helly-type argument shows that the intersection of all the projections must be non-empty [9]. Thus we can choose a common point O (henceforth regarded as the origin) belonging to the interior of each of the projections. To simplify the notation, we do not distinguish between a rectangle and its projection onto the x, y-plane; the meaning will be clear from the context. Without loss of generality, we may assume that all rectangles in a representation have distinct, non-negative integer z-coordinates.

Each rectangle R in a valid sequence can be described in terms of the perpendicular distances from O to each of its sides. Instead of giving the x, y coordinates of the vertices of R, we describe R as a 4-tuple (E_r, N_r, W_r, S_r) whose coordinates give, respectively, the distances from $O \in R$ to the east, north, west and south sides of R.

We can assume without loss of generality that no two rectangles of a valid sequence share the same value on any of the four coordinates E,N,W,S. Hence we can assume that each coordinate value of each of the n rectangles is an integer in the range $[1, n]$ without changing the visibility relationships among the rectangles.

Consider two rectangles $A = (E_a, N_a, W_a, S_a)$ and $B = (E_b, N_b, W_b, S_b)$ in a valid sequence, and denote by $A \cap B$ the intersection of their projections onto the x, y-plane. Then $A \cap B$ contains O, and the coordinates of $A \cap B$ are $E_{A \cap B} = \min\{E_a, E_b\}$, $N_{A \cap B} = \min\{N_a, N_b\}$, $W_{A \cap B} = \min\{W_a, W_b\}$ and $S_{A \cap B} = \min\{S_a, S_b\}$. We say that a corner of $A \cap B$ is *free* if it is not covered by any of the projections of rectangles occuring between A and B in the sequence.

Suppose A and B are rectangles in a valid sequence. Then since O belongs to all the rectangles, at least one of the corners of $A \cap B$ must be free. This is because any rectangle that covers a corner also covers O and hence an entire quadrant of $A \cap B$. Thus if $A \cap B$ had no free corner, it would be covered by the union of the intervening rectangles that cover at least one corner of $A \cap B$.

The northeast corner of $A \cap B$ is not covered by a particular rectangle $R = (E_r, N_r, W_r, S_r)$ between A and B if and only if the boolean expression $E_r <$

$\min\{E_a, E_b\}$ *OR* $N_r < \min\{N_a, N_b\}$ is true. Similar conditions hold for the other three corners.

We summarize: the rectangles A and B can see each other if and only if at least one of the following free corner conditions **FC** holds *simultaneously* for *all* the rectangles R between A and B:

$\mathbf{FC}_{ne}(A, B)$ northeast is free: ($E_r < \min\{E_a, E_b\}$ *OR* $N_r < \min\{N_a, N_b\}$);

$\mathbf{FC}_{nw}(A, B)$ northwest is free: ($N_r < \min\{N_a, N_b\}$ *OR* $W_r < \min\{W_a, W_b\}$);

$\mathbf{FC}_{sw}(A, B)$ southwest is free: ($W_r < \min\{W_a, W_b\}$ *OR* $S_r < \min\{S_a, S_b\}$);

$\mathbf{FC}_{se}(A, B)$ southeast is free: ($S_r < \min\{S_a, S_b\}$ *OR* $E_r < \min\{E_a, E_b\}$).

Now we give a definition that is needed in the following discussions. Given a valid sequence of n rectangles $(E_1, N_1, W_1, S_1), \ldots, (E_n, N_n, W_n, S_n)$, we define the sequences $V_E = (E_1, ...E_n)$, $V_N = (N_1, ...N_n)$, $V_W = (W_1, ...W_n)$, and $V_S = (S_1, ...S_n)$.

2.2 Unimaximal sequences

The next definition and lemma provide a key tool in our analysis.

Definition 2.1 *A sequence $x_1, x_2, \ldots$ of distinct integers is called* unimaximal *if it has exactly one local maximum, i.e., for all i, j, k with $i < j < k$, we have $x_j > \min\{x_i, x_k\}$ for $i < j < k$.*

Lemma 2.2 *For all $m > 1$, in every sequence of $\binom{m}{2} + 1$ distinct integers, there exists at least one unimaximal subsequence of length m. On the other hand, there exists a sequence of $\binom{m}{2}$ distinct integers that has no unimaximal subsequence of length m.*

This result arises from the Erdős-Szekeres Theorem [13], whose pigeon-hole proof was given by [16]. Lemma 2.2 is attributed by F. P. K. Chung [6] to V. Chvátal and J. M. Steele, among others.

Lemma 2.3 *In a representation of K_5 by five rectangles, with no other rectangles present, it is impossible that both sequences V_N and V_S are unimaximal.*

Proof: Suppose both V_N and V_S are unimaximal. Then the relations $N_r > \min\{N_a, N_b\}$ and $S_r > \min\{S_a, S_b\}$ must hold for all rectangles A, B, and R between A and B. Now consider the conditions $\mathbf{FC}_{ne}$, $\mathbf{FC}_{nw}$, $\mathbf{FC}_{sw}$, $\mathbf{FC}_{se}$. For $\mathbf{FC}_{ne}(A, B)$ to be true, it must be the case that $E_r < \min\{E_a, E_b\}$ for all R between A and B. The same is true for $\mathbf{FC}_{se}(A, B)$. Similarly, for $\mathbf{FC}_{nw}(A, B)$ or $\mathbf{FC}_{sw}(A, B)$ to be true, W_r must be less than $\min\{W_a, W_b\}$. Hence the free corner conditions reduce to the following. One of the two possibilities ($W_r < \min\{W_a, W_b\}$) or ($E_r < \min\{E_a, E_b\}$) holds simultaneously for all rectangles R

between A and B. This means that all rectangles A and B can see each other along a line of sight with y-coordinate 0. By intersecting the arrangement of five rectangles with the x, z-plane, we get an arrangement of 5 line segments in this plane that all see each other. This contradicts the fact that only planar graphs can be represented by vertical visibility of horizontal line segments in a plane. (See [22, 24] for results on such representations in the plane.) □

2.3 The bound

Theorem 2.4 *No complete graph K_n has a z-visibility representation for $n \geq 56$.*

Proof: Suppose we had a representation of K_n with $n \geq 56$. Lemma 2.2 implies that V_N has a unimaximal sequence V'_N of length 11. Consider the associated subsequence V'_S of length 11. It follows again from Lemma 2.2 that there is a subsequence V''_S of length 5 that is unimaximal. Remove the rectangles not associated with the subsequence. This destroys no visibility lines, so the five remaining rectangles represent K_5. However, both V''_S and its corresponding subsequence V''_N are unimaximal. This contradicts Lemma 2.3. □

3 A lower bound

In this section we show by construction that the complete graph on 22 vertices is representable. This improves the previous lower bound of 20 in [5].

3.1 The construction

Figure 4 gives the representation of K_{22} discovered by the algorithm. At each stage, a new rectangle is added; each white dot indicates a corner of visibility for a rectangles at a lower level.

The following table shows the coordinates of the 22 rectangles in the notation of the previous section.

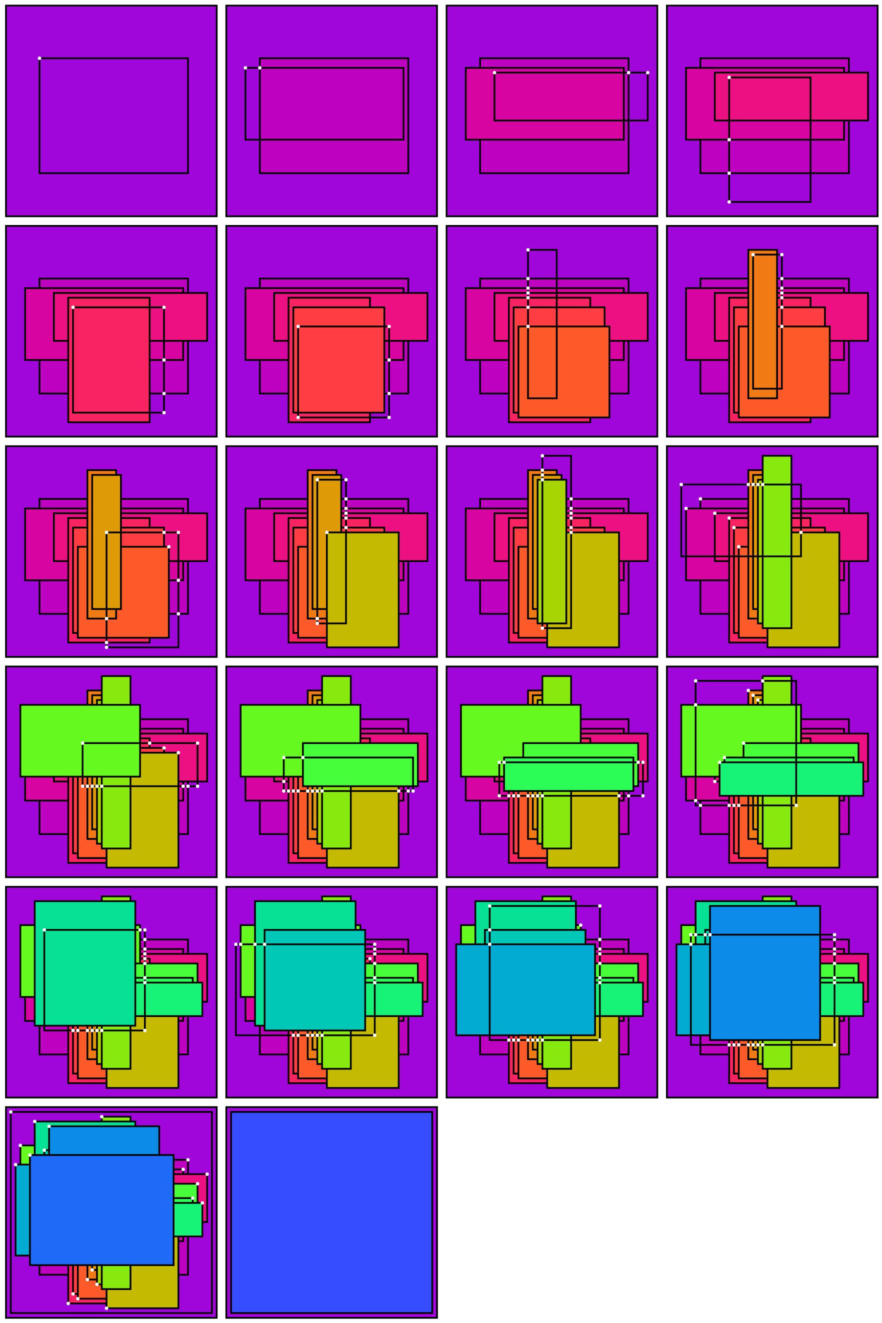

Figure 4: A representation of K_{22}.

Rectangle	North	South	West	East
1	22	22	22	22
2	11	13	15	16
3	9	6	18	15
4	8	2	12	20
5	7	19	9	8
6	5	17	8	11
7	1	18	7	12
8	17	14	5	1
9	16	12	4	2
10	4	20	1	14
11	15	15	3	3
12	19	16	2	4
13	14	1	19	6
14	6	3	6	18
15	3	4	10	17
16	2	5	11	19
17	20	7	16	5
18	13	8	14	7
19	10	9	20	9
20	18	10	13	10
21	12	11	17	13
22	21	21	21	21

3.2 How the representation was found

The representation of K_{22} was found using simulated annealing, a general randomized heuristic approach for finding good solutions for optimization problems. By starting with a randomly-generated candidate solution and applying local modifications, other candidate solutions are generated. The choice of local transformation is guided by an objective function which the procedure attempts to minimize. See [1, 20, 23] for more information regarding simulated annealing techniques.

For a given n, the algorithm tries to find a realization of K_n. As described in the previous section, any configuration can be assumed to consist of a list of 4-tuples of numbers between 0 and $n-1$, where all numbers for the same coordinate are distinct. Two configurations I and J are regarded as *adjacent* if the configuration J is obtained from I by swapping one of the edge coordinates of two rectangles in I. For a configuration I, the set of all configurations adjacent to I shall be called the *neighborhood* of I.

The *objective function* that we want to minimize is defined as follows: for every pair of rectangles A and B in the collection, we assign a score equal to the minimum number of rectangles needed to be removed from the collection in order for A and B to become mutually visible. This score can be determined in a straightforward manner from the free corner conditions **FC** stated in the previous section, in linear time. The objective function is obtained by summing these scores over all pairs of rectangles; this function is zero for any feasible realization of K_n and is positive otherwise.

Simulated annealing generates a random element J in the neighborhood of the current solution I and compares their two objective values. If J is better than I, then J is accepted as the new current solution; otherwise, J is accepted with a probability that depends on the difference in the objective functions and on a parameter T, which is called the *temperature*. The higher the temperature, the more likely it is to accept worse solutions. The initial temperature should be of the same order of magnitude as the average change of objective function for a random swap. The temperature is gradually decreased until it is so low that no changes are accepted and the algorithm gets stuck at a local minimum.

After some experimentation we chose to decrease T by a factor of 0.96 after every 100-th iteration. When spending more than 160 iterations in the same solution, the algorithm decides that it is caught in a local minimum, generates a new random configuration, and restarts. The program stops when the objective function reaches zero.

Initially, we determined the neighboring solution by selecting two rectangles uniformly at random and swapping one of the four edge coordinates, again chosen randomly, in accordance with the general principles of simulated annealing. However, in order to accelerate the algorithm, we changed to a non-uniform selection. We defined a penalty function for each rectangle R in the collection, whose value is based on

1. the number of rectangle pairs whose visibility is blocked by R.
2. the number of rectangles not visible to R.

If these numbers are large, the rectangle R is likely to be important for the objective function, and hence it is favored in the selection. We also considered an enlargement of the neighborhood by defining two configurations to be adjacent if one can be obtained from the other by swapping two, three, or all four edge coordinates between two of the rectangles.

After experimenting with the choices of T and the size of the neighborhood for approximately one month, a z-visibility representation for K_{22} was found. The search for a representation of K_{23} was run more or less continuously on various workstations and personal computers for almost a year, without success.

4 Additional Results

The section presents some simple, additional representability results, most notably, that all planar graphs are representable.

4.1 Representations for planar graphs

In this section, we show that all planar graphs have a representation. There are two main ingredients in the proof. The first is the result due independently to

Wismath [24] and to Tamassia and Tollis [22] that any 2-connected planar graph has what [22] calls an ε-*visibility* representation. (Vertices correspond to closed, disjoint, horizontal line segments in the plane, and two vertices are adjacent in the graph if and only if their corresponding segments are ε-visible in the vertical direction.) The second ingredient is the use of the third dimension to deal with cut vertices. This is similar to an idea of [24] for obtaining a visibility representation for all planar graphs by rectangles in $\mathbb{R}^2$ that have ε-visibility in both the x and y directions. For ease of notation, let $P_{-\infty}$ represent the plane $z = -\infty$ and let P_∞ represent the plane defined by $z = +\infty$.

Theorem 4.1 *Every planar graph admits a z-visibility representation.*

Proof: We will prove, by induction on the number of 2-connected components, the following stronger result:

Claim: Let G be a connected planar graph and let v be a vertex of G. Then G has a z-visibility representation such that each rectangle is visible from P_∞, but only the rectangle representing v is visible from $P_{-\infty}$.

Base case for the induction: G is 2-connected. We say a line segment is a y-segment if it is parallel to the y-axis and lies in the y,z-plane. By the result of Wismath [24] and Tamassia and Tollis [22], G can be represented by (planar) ε-visibility restricted to the z direction of y-segments, with only the segment representing the vertex v visible from below. Place the y,z-plane containing this configuration in 3-space at $x = 0$. Number the segments in order of decreasing z-coordinate. Expand each segment to an x,y-rectangle by pulling it out until its x-length is equal to its number. The rectangles then form a "staircase" shape. Each rectangle is visible from P_∞, and only the rectangle repesenting v is visible from $P_{-\infty}$.

Now assume the result is true for graphs with at most k 2-connected components. Let the number of 2-connected components of G be $k+1$. Let x be a cut vertex of G, and break G at x into two subgraphs $G_1 \cup \{x\}$ and $G_2 \cup \{x\}$. (Vertex x may still be a cut vertex in these subgraphs.) Suppose that v lies in G_1. By induction G_1 has a z-visibility representation with all rectangles visible from P_∞ and only the rectangle representing v visible from $P_{-\infty}$. Identify a rectangular area A of the rectangle corresponding to the vertex x, such that A is visible from P_∞. By induction G_2 has a z-visibility representation with all rectangles visible from P_∞ and only the rectangle representing x visible from $P_{-\infty}$. Scale the representation for G_2, so that it can be placed upward of the rectangular area A in the z-direction. After the representation of G_2 is in place, remove the rectangle corresponding to x from the representation of G_2. The result is a z-visibility representation of G; all rectangles are visible from P_∞, and only the rectangle representing v is visible from $P_{-\infty}$. This completes the proof. $\square$

4.2 Representations for bipartite graphs

As Figure 5 illustrates, every complete bipartite graph $K_{m,n}$ admits a z-visibility representation.

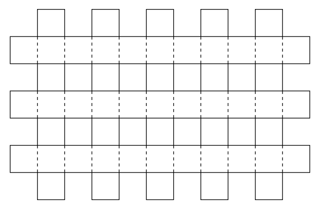

Figure 5: A z-visibility representation of $K_{m,n}$

4.3 Nonclosure under graph minors

The family of representable graphs is not closed under graph minors. To prove this, consider the complete bipartite graph $K_{56,56}$, which has a z-visibility representation as illustrated in Figure 5. When an edge of a complete bipartite graph is contracted, it yields a vertex that is adjacent to all the vertices of the graph. If we contract the 56 edges of a perfect matching in $K_{56,56}$, then we obtain the complete graph K_{56}, which does not have a z-visibility representation by Theorem 2.4. This gives us the following result:

Theorem 4.2 *The class of graphs admitting a z-visibility representation is not closed under graph minors.*

5 Conclusions and open problems

We have shown that for the 3-dimensional z-visibility representation introduced in this paper, all planar graphs have a representation, K_n has a representation for $n \leq 22$, and K_n does not have a representation for $n \geq 56$. Concerning bipartite graphs, we showed that $K_{m,n}$ is representable for all m and n. Finally, we showed that the family of graphs with a representation is not closed under graph minors.

¿From the point of view of pure discrete geometry and combinatorics, the problem of finding the exact upper upper bound for the representability of K_n remains a tantalizing one. The lower bound was found by computation (simulated annealing), and the upper bound was found by pushing the Erdős-Szekeres approach hard, using an extended version of the original theorem. However, from the point of view of graph visualization, it is already interesting that K_{22} has a z-visibility representation, as this means that *all* graphs with at most 22 vertices have a weak representation.

It was shown in [5] that $K_{5,5}$ minus a perfect matching has a representation. We conjecture that $K_{6,6}$ minus a perfect matching does not have a representation. It was also shown in [5] that the complete tripartite graphs $K_{m,n,2}$ and $K_{m,4,3}$ can be represented.

What is the smallest graph that does not have a z-visibility representation?

Acknowledgments

Our study of visibility representations began at Bellairs Research Institute of McGill University during the International Workshop on Visibility Representations organized by S. Whitesides and J. Hutchinson, February 12-19, 1993. We are grateful to the other conference participants Joan Hutchinson, Goos Kant, Marc van Kreveld, Beppe Liotta, J. R., Steve Skiena, Roberto Tamassia, Yanni Tollis, and Godfried Toussaint for their encouragement and comments.

References

[1] E. Aarts, J. Korst. *Simulated Annealing and Boltzmann Machines. A Stochastic Approach to Combinatorial Optimization and Neural Computing.* Wiley, Chichester 1989

[2] H. Alt, M. Godau, and S. Whitesides. Universal 3-dimensional visibility representations for graphs. *Proc. Graph Drawing 95*, Passau 1996. Lecture Notes in Computer Science Vol. 1027, Springer-Verlag, 1996, pp. 8–19.

[3] A. Bachem, S. P. Fekete, B. Knab, R. Schrader, I. Vannahme, I. Weber, R. Wegener, K. Weinbrecht, B. Wichern. Analyse großer Datenmengen und Clusteralgorithmen im Bausparwesen. To appear in *Geld, Finanzwirtschaft, Banken und Versicherungen.* Editors C. Hipp et al., VVW Karlsruhe, 1997.

[4] P. Bose, H. Everett, S. P. Fekete, A. Lubiw, H. Meijer, K. Romanik, T. Shermer, and S. Whitesides. On a visibility representation for graphs in three dimensions. *Proc. Graph Drawing '93*, Paris (Sèvres), 1993, pp. 38–39.

[5] P. Bose, H. Everett, S. P. Fekete, A. Lubiw, H. Meijer, K. Romanik, T. Shermer, and S. Whitesides. On a visibility representation for graphs in three dimensions. Snapshots of Computational and Discrete Geometry, **3**, *eds.* D. Avis and P. Bose, McGill University School of Computer Science Technical Report SOCS-94.50, July 1994, pp. 2–25.

[6] F. R. K. Chung. On unimodal subsequences. *J. Combinatorial Theory*, Series A, **29**, 1980, pp. 267–279.

[7] F. J. Cobos, J. C. Dana, F. Hurtado, A. Marquez, and F. Mateos On a visibility representation of graphs. *Proc. Graph Drawing 95*, Passau

1996. Lecture Notes in Computer Science Vol. 1072, Springer-Verlag, 1996, pp. 152–161.

[8] R. Cohen, P. Eades, T. Lin, and F. Ruskey. Three-dimensional graph drawing. *Proc. Graph Drawing '94*, Princeton NJ, 1994. Lecture Notes in Computer Science LNCS Vol. 894, Springer-Verlag, 1995, pp. 1–11.

[9] L. Danzer, B. Grünbaum, and V. Klee. Helly's theorem and its relatives. *Convexity, AMS Proc. Symp. Pure Math.* **7**, 1963, pp. 101–181.

[10] A. Dean and J. Hutchison. Rectangle visibility representations of bipartite graphs. *Proc. Graph Drawing '94*, Princeton NJ, 1994. Lecture Notes in Computer Science LNCS Vol. 894, Springer-Verlag, 1995, pp. 159–166.

[11] G. Di Battista, P. Eades, R. Tamassia, and I. G. Tollis, Algorithms for automatic graph drawing: an annotated bibliography. *Comput. Geometry: Theory and Applications*, **4**, 1994, pp. 235–282. Also available from wilma.cs.brown.edu by ftp.

[12] G. Di Battista and R. Tamassia. Algorithms for plane representations of acyclic digraphs. *Theoretical Computer Science*, **61**, 1988, pp. 175–198.

[13] P. Erdős and A. Szekeres. A combinatorial problem in geometry. *Compositio Mathematica* **2**, 1935, pp. 463–470.

[14] S. P. Fekete, M. E. Houle, and S. Whitesides. New results on a visibility representation of graphs in 3D. *Proc. Graph Drawing 95*, Passau 1996. Lecture Notes in Computer Science Vol. 1072, Springer-Verlag 1996, pp. 234–241.

[15] S. P. Fekete and H. Meijer. Rectangle and box visibility graphs in 3D. To appear in *Intern. J. Comp. Geom. Appl.*

[16] J. M. Hammersley. A few seedlings of research. *Proc. 6^{th} Berkeley Symp. Math. Stat. Prob.*, U. of California Press, 1972, pp. 345–394.

[17] J. Hutchinson, T. Shermer, and A. Vince. On representations of some thickness-two graphs. *Proc. Graph Drawing 95*, Passau 1996. Lecture Notes in Computer Science Vol. 1072, Springer-Verlag 1996, pp. 324–332.

[18] G. Kant, G. Liotta, R. Tamassia, and I. G. Tollis. Area requirements of visibility representations of trees. *Proceedings of the 5th Canadian Conf. on Comp. Geom.*, Waterloo, Ontario, 1993, pp. 192–197.

[19] E. Kranakis, D. Krizanc, and J. Urrutia. On the number of directions in visibility representations of graphs. *Proc. Graph Drawing '94*, Princeton NJ, 1994, Lecture Notes in Computer Science LNCS Vol. 894, Springer-Verlag, 1995, pp. 167–176.

[20] P. J. M. van Laarhoven, E. H. L. Aarts. *Simulated Annealing. Theory and Applications.* (Reprint with corrections.) Reidel, Dordrecht 1988.

[21] K. Romanik. Directed VR-representable graphs have unbounded dimension. *Proc. Graph Drawing '94*, Princeton, NJ, 1994, Lecture Notes in Computer Science LNCS Vol. 894, Springer-Verlag, 1995, pp. 177–181.

[22] R. Tamassia and I. G. Tollis. A unified approach to visibility representations of planar graphs. *Discrete Comput. Geom.* **1**, 1986, pp. 321–341.

[23] R. V. V. Vidal (ed.) *Applied Simulated Annealing.* Lecture Notes in Economics and Mathematical Systems, Vol. 396, Springer-Verlag, 1993.

[24] S. K. Wismath. Characterizing bar line-of-sight graphs. *Proc. ACM Symp. on Comp. Geometry*, 1985, pp. 147–152.