On a Visibility Representation for Graphs in Three Dimensions

Prosenjit Bose[*], Hazel Everett[†], Sándor Fekete[‡], Anna Lubiw[§]
Henk Meijer[¶], Kathleen Romanik[*], Tom Shermer[†] and Sue Whitesides[*]

Dedicated to Godfried Toussaint

Abstract

Visibility representations of graphs map vertices to sets in Euclidean space and express edges as visibility relations between these sets. Application areas such as VLSI wire routing and circuit board layout have stimulated research on visibility representations where the sets belong to R^2. Here, motivated by the emerging research area of graph drawing, we study a 3-dimensional visibility representation.

In this representation, each vertex of the graph maps to a closed rectangle in R^3 such that the rectangles are disjoint, 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. Edges are expressed by the following visibility relation. Two rectangles R_i and R_j are considered visible provided that there exists a closed cylinder C of non-zero 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 C does not intersect any other rectangle.

If a graph can be represented in this way, we say that the graph is *VR-representable*. Our main results are as follows. All planar graphs are VR-representable, as are many non-planar graphs. In particular, K_n is VR-representable for values of $n \leq 20$. However, K_n is not VR-representable for $n \geq 103$. Concerning bipartite graphs, we show that $K_{m,n}$ is VR-representable for all m and n, and provide bounds on the representability of complete bipartite graphs minus a perfect matching. We also provide bounds on the representability of tripartite graphs. Furthermore, we show that the family of VR-representable graphs is not closed under graph minors. Finally, we consider variant VR-representations, where the rectangles are replaced by squares and discs, and provide bounds on the complete graphs and complete bipartite graphs that admit such representations.

1 Introduction

The problem of drawing or establishing different representations of a given graph has been studied extensively in the literature (see the paper of Di Battista et al.[2] for a survey of the research and applications of graph drawing). In particular, determining

[*] McGill University, Canada; jit@muff.cs.mcgill.ca, romanik@opus.cs.mcgill.ca, sue@cs.mcgill.ca
[†] Université du Québec à Montréal, Canada; hazel@math.uqam.ca
[‡] Universität zu Köln, Köln, Germany; sandor@MI.Uni-Koeln.DE
[§] University of Waterloo, Canada; alubiw@maytag.uwaterloo.ca
[¶] Queen's University, Canada; henk@qucis.queensu.ca
[†] Simon Fraser University, Canada; shermer@cs.sfu.ca

a *visibility representation* of a graph, where the vertices of the graph map to sets or objects in Euclidean space and the edges are expressed as visibility relations between these sets, has received considerable attention recently due to the large number of applications. Application areas such as VLSI wire routing, algorithm animation, CASE tools, and circuit board layout have stimulated research on visibility representations where the sets belong to R^2.

The notion of *visibility* plays a key role in determining the types of graphs that admit visibility representations. Intuitively, we say that two objects are *visible* if they can be joined by a segment which does not intersect any other object. Two objects are *ϵ-visible* if they can be joined by a non-zero thickness band that does not intersect any other object. 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, and line segments 2 and 4 are not visible at all. Although these two different types of visibility seem closely related, the restrictions imposed by ϵ-visibility often radically changes the class of graphs that admit visibility representations (see e.g. [6]). Visibility can also be restricted to a particular direction. For example, two objects may be deemed visible provided that the line segment joining them is vertical. In which case, line segments 1 and 3 are no longer ϵ-visible but only visible.

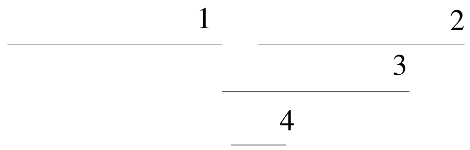

Figure 1: An illustration of visibility relations.

The different sets or objects used to represent the vertices also play an important role in determining the types of graphs that admit visibility representations. In the plane, the two most common objects studied have been line segments and rectangles. For example, Wismath [6] and Tamassia and Tollis [5] independently 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 two vertices are adjacent provided that their corresponding segments are ϵ-visible with respect to vertical visibility. Visibility representations of trees, where the vertices are represented by horizontal line segments, have been studied by Kant et al.[4]

When the vertices are represented by disjoint isothetic rectangles in the plane and visibility is restricted to ϵ-visibility in the horizontal or vertical directions, Wismath [6] proved that every planar graph admits such a visibility representation. Dean and Hutchinson [1] prove that K_8 is the largest complete graph that admits a visibility representation in this model. For a comprehensive overview of the various visibility representations studied, we refer the reader to [2].

Motivated by the emerging research area of graph drawing, we study a 3-dimensional visibility representation. Consider an arrangement of closed, disjoint rectangles in R^3 such that the planes determined by the rectangles are perpendicular to the z (vertical) axis, and the sides of the rectangles are parallel to the x or y (horizontal) axes. Two rectangles R_i and R_j are ϵ-visible restricted to the vertical direction if between the two rectangles there is a closed cylinder C of non-zero 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. When such a line of sight exists, we say that the two rectangles are $z - visible$. Such an arrangement is a *Visibility Representation by Isothetic Rectangles in 3-Space* of a graph $G = (V, E)$, denoted by $VR(G)$. A graph admits such a representation provided that the following hold:

- there exists a 1-1 onto correspondence between the rectangles and the vertices, and
- vertices v_i and v_j are adjacent if, and only if, their corresponding rectangles R_i and R_j are z-visible.

If a graph can be represented in this way, we say that the graph is *VR-representable*. We show that all planar graphs are VR-representable, as are many non-planar graphs. In particular, K_n is VR-representable for values of $n \leq 20$. However, K_n is not VR-representable for $n \geq 103$. Concerning bipartite graphs, we show that $K_{m,n}$ is VR-representable for all m and n, and provide bounds on the representability of complete bipartite graphs minus a perfect matching. We also provide bounds on the representability of tripartite graphs. Furthermore, we show that the family of VR-representable graphs is not closed under graph minors. Finally, we consider variants of VR-representations, where the rectangles are replaced by squares and discs and provide bounds on the complete graphs and complete bipartite graphs that admit such representations.

2 Notation and preliminaries

We begin by defining some of the conventions used in this paper. Unless stated otherwise, ϵ-visibility will be restricted to the z–direction, the planes containing rectangles will be orthogonal to the z-axis and the edges of the rectangles will be parallel to the x or y directions. For ease of reference, most figures will be drawn as their projection into the $z = 0$ plane. Each rectangle, if labeled, will be labeled with its z–coordinate. We will use the terms *up* and *down* to refer to increasing and decreasing z–coordinate, the terms *above* and *below* to refer to increasing and decreasing y–coordinate, and the terms *right* and *left* to refer to increasing and decreasing x–coordinate. Similarly, *upmost* and *downmost* refer to z–coordinate, *top* and *bottom* refer to y–coordinate, and *rightmost* and *leftmost* refer to x–coordinate.

We use R_i to denote a rectangle with z–coordinate i. Given the projection of two rectangles R_i and R_j, we will write $R_i <_l R_j$ if the left edge of R_i is to the left of the left edge of R_j, and $R_i <_r R_j$ if the right edge of R_i is to the left of the right edge of R_j (see Figure 2). Similarly, we will say that $R_i >_t R_j$ if the top edge of R_i is above the top edge of R_j, and $R_i >_b R_j$ if the bottom edge of R_i is above the bottom of R_j. The relations $<_l$, $<_r$, $<_t$ and $<_b$ are partial orders on the set of vertices of a given VR-representation.

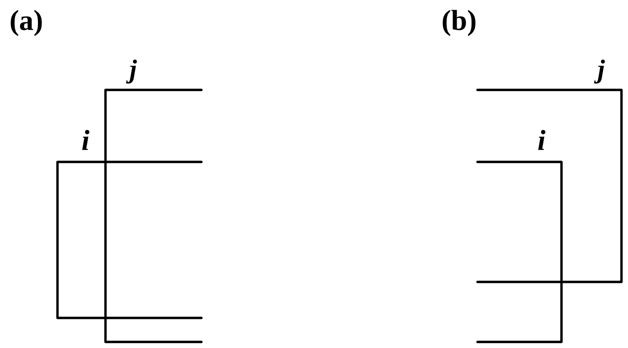

Figure 2: Examples of $R_i <_l R_j$ and $R_i <_r R_j$

3 VR-representability of planar graphs

In this section, we show that all planar graphs are VR-representable. There are two main ingredients in the proof. The first is the result due independently to Wismath [6] and to Tamassia and Tollis [5] that any 2-connected planar graph has what [5] 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 [6] for obtaining a visibility representation for all planar graphs by rectangles in R^2 that have ε-visibility in both the x and y directions. For ease of exposition, let $\mathcal{U}$ represent the plane at $z = +\infty$ and $\mathcal{D}$ represent $z = -\infty$.

Theorem 3.1 *Every planar graph is VR-representable.*

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 is VR-representable in such a way that each rectangle is visible from $\mathcal{U}$, and only the rectangle v is visible from $\mathcal{D}$.

Basis 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 on the y-z plane. By the result of Wismath [6] and Tamassia and Tollis [5], G can be represented by ε-visibility restricted to the z direction of y-segments, with only the segment representing the vertex v visible from below. Place this configuration 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 $\mathcal{U}$, and only the rectangle v is visible from $\mathcal{D}$.

Assume the result is true for graphs with at most k 2-connected components. Let G have $k+1$ 2-connected components. Let x be a cut vertex of G, and break G at x into two subgraphs G_1 and G_2. (Vertex x may still be a cut vertex in these subgraphs.) Suppose that v lies in G_1. By induction G_1 is VR-representable, with all rectangles visible from $\mathcal{U}$ and only the rectangle v visible from $\mathcal{D}$. Identify a rectangular area A of the rectangle X, corresponding to the vertex x, that is visible from $\mathcal{U}$. By induction G_2 is VR-representable with all rectangles visible from $\mathcal{U}$ and only the rectangle x visible from $\mathcal{D}$. Scale the representation for G_2, delete the rectangle corresponding to x, and place the remaining configuration over-top of the rectangular area A. The result is a VR-representation of G; all rectangles are visible from $\mathcal{U}$, and only rectangle v is visible from $\mathcal{D}$. This completes the proof. ∎

4 VR-representability of complete graphs

In this section, we present the construction of a VR-representation of the complete graph on 20 vertices. The construction is based on two main building blocks. Each of the blocks has some special properties that are exploited in order to achieve the construction. For any VR-representation of a complete graph, each rectangle must lie on a different plane; thus, the upmost rectangle U and downmost D can be extended indefinitely without interfering with the visibility of the remaining rectangles.

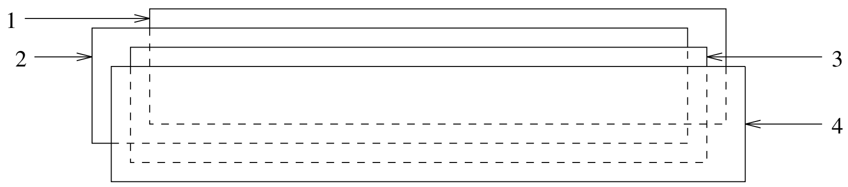

Figure 3: K_4

The first building block is an arrangement of 4 rectangles as depicted in Figure 3. Not only is the visibility graph of this configuration complete, but $R_1 >_t R_2 >_t R_3 >_t R_4$, $R_1 >_b R_2 >_b R_3 >_b R_4$, and all 4 rectangles see D and U.

The other building block consists of five rectangles arranged as in Figure 4. Notice that the visibility graph of this configuration is also complete and we have that $R_1 >_t R_2 >_t R_3 >_t R_4 >_t R_5$, $R_1 >_r R_2 >_r R_3 >_r R_4 >_r R_5$, and all 5 rectangles see D and U.

A K_9 can be built by placing these two blocks one upon the other as in Figure 5. The right edges of the K_5 and the bottom edges of the K_4 provide the visibility needed to achieve a K_9. Note that this arrangement of 9 rectangles, whose visibility graph is complete, satisfies the relation $R_1 >_t R_2 >_t R_3 \cdots >_t R_9$, and all see D and U.

A K_{18} is constructed as a composite of two K_9 configurations. One K_9 is placed as in

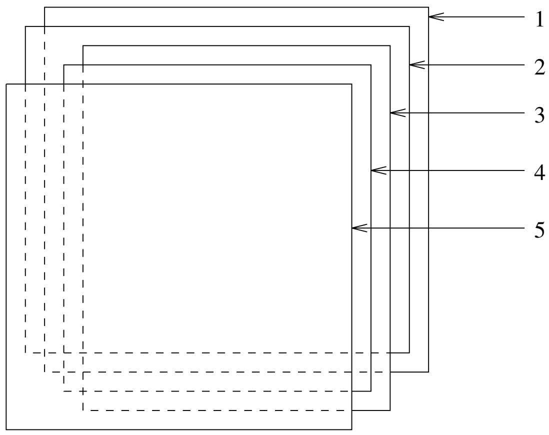

Figure 4: K_5

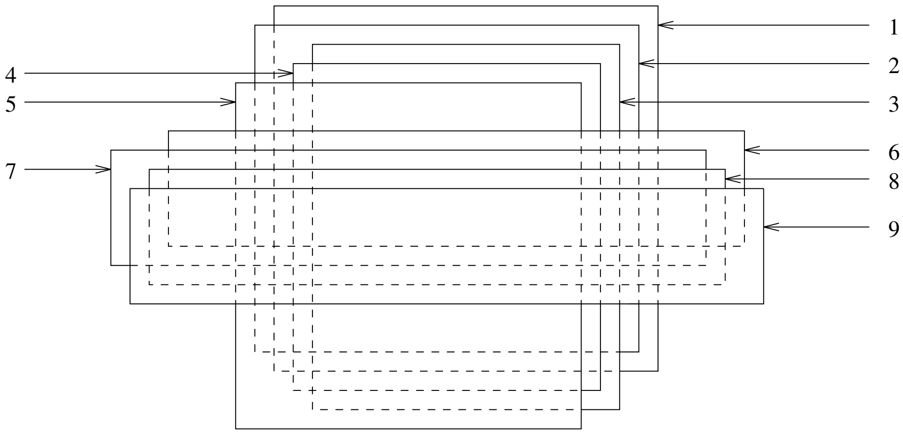

Figure 5: K_9

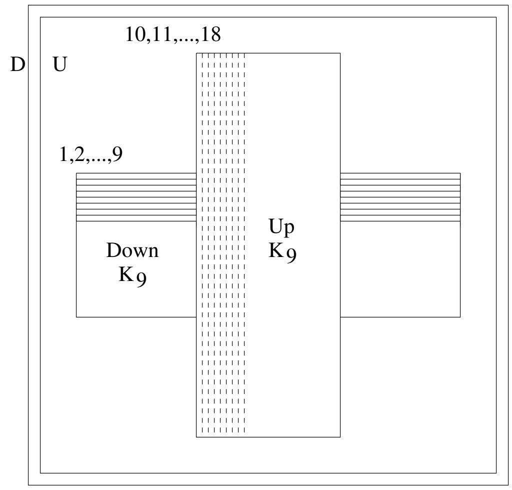

Figure 6: K_{20}

Figure 5, referred to as the down K_9. The other K_9 is place upon the down K_9. The up K_9 is an exact copy of the down one, but flipped over and rotated 90° so that the relation $R_1 <_l R_2 <_l R_3 \cdots <_l R_9$ holds. The top edges of the down K_9 and the left edges of the up K_9 provide the visibility needed to achieve K_{18}. Notice that the K_{18} is constructed such that all rectangles see D and U. Therefore, by adding D and U, we have a VR-representation of K_{20} (see Figure 6). Hence, we have the following theorem.

Theorem 4.1 *The complete graph on 20 vertices has a VR-representation.*

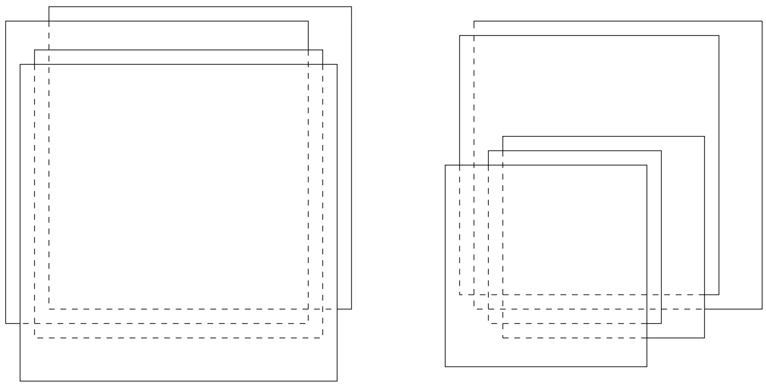

Figure 7: Square building blocks.

Notice that all the rectangles in Figure 3 and Figure 4 can be enlarged into squares without removing any of the properties of either building block. Figure 7 shows the two main building blocks of squares that can be used to form a complete visibility graph of size 20. The construction is the same as the K_{20} construction for rectangles. Therefore, we have the following corollary.

Corollary 4.2 *The complete graph on 20 vertices has a visibility representation by isothetic squares in 3-space.*

It is interesting to note that although rectangles are more flexible objects than squares, we are unable to capitalize on this flexibility as our lower bounds for VR-representations of complete graphs for both these objects is the same.

5 Non-VR-representability of large complete graphs

In this section, we provide an upper bound on the size of the largest clique admitting a VR-representation. We show that K_n is not VR-representable for any $n > 102$. We will make extensive use of the following theorem of Erdös and Szekeres[3]:

Theorem 5.1 *Given any two positive integers j and k, any sequence of more than jk distinct integers has a (not necessarily contiguous) ascending subsequence of length $j+1$ or descending subsequence of length $k+1$.*

Consider an arrangement $\mathcal{R}$ of rectangles that has a clique for its visibility graph. The upmost rectangle U, and the downmost D, can be extended indefinitely without interfering with the visibility amongst the other rectangles. Thus, we may assume that U and D have projections on the $z=0$ plane that enclose all of the other rectangles of $\mathcal{R}$. Let $\mathcal{R}'$ be $\mathcal{R} \setminus \{U, D\}$. The collection $\mathcal{R}'$ must itself form a clique, and each of its rectangles must be able to see upwards past all other elements of $\mathcal{R}'$ (in order to see U) and downwards past all other elements of $\mathcal{R}'$ (in order to see D).

We begin by proving two key lemmas.

Lemma 5.2 *Let R_i, R_j, R_k, and R_l be four rectangles of $\mathcal{R}'$ with $i<j<k<l$. It cannot be the case that $R_i <_l R_j <_l R_k <_l R_l$ and $R_l <_r R_k <_r R_j <_r R_i$.*

Proof: As R_j must see the downmost rectangle D, either the bottom edge of R_j lies below the bottom edge of R_i, or the top edge of R_j lies above the top edge of R_i, or both. Note, however that it is impossible for both conditions to hold since then R_k will not see R_i. As the two conditions are symmetric, we will assume without loss of generality that the bottom edge of R_j lies below the bottom edge of R_i (see Figure 8a).

Similarly, R_k must see D. Thus, either the top edge of R_k lies above the top edge of R_i, or the bottom edge of R_k lies below the bottom edge of R_j, or both. If the top edge of R_k lies above the top edge of R_i (or both conditions hold), then, because R_k must see R_j, the bottom edge of R_k must be below the top edge of R_j (see Figure 8b). In this situation, R_k and R_j will block R_l from seeing R_i. If the bottom edge of R_k lies below the bottom

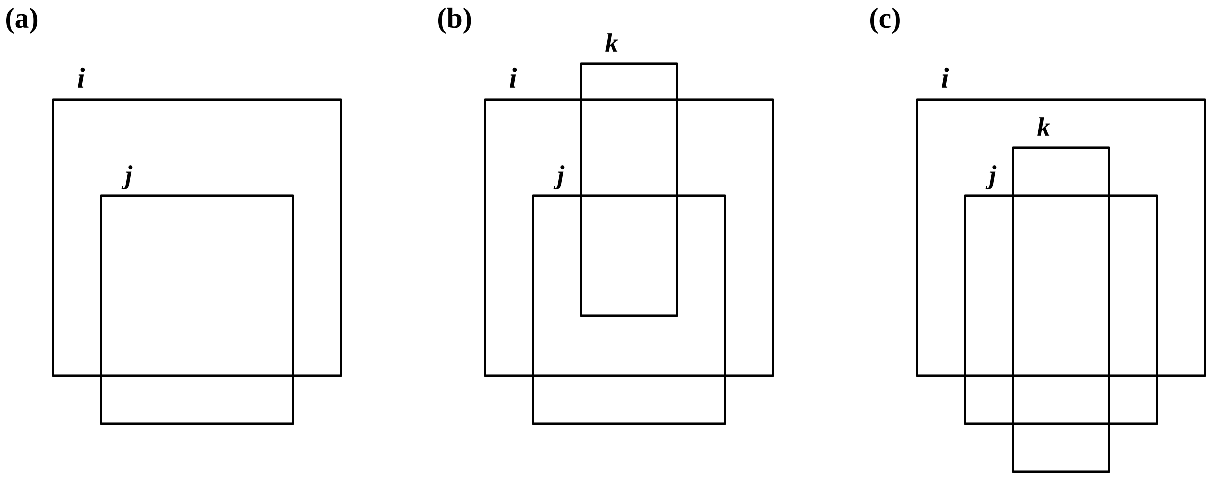

Figure 8: Trying to lay out $R_i <_l R_j <_l R_k <_l R_l$ and $R_l <_r R_k <_r R_j <_r R_i$

edge of R_j, then, because R_k must see R_i, the top edge of R_k lies between the top edges of R_i and R_j (see Figure 8c). In this case, R_l will not see R_j. Thus, there is no possible layout of the situation, and the lemma holds. ∎

Lemma 5.3 *Let R_i, R_j, R_k, R_l, and R_m be five rectangles of $\mathcal{R}'$ with $i < j < k < l < m$. It cannot be the case that $R_i <_l R_j <_l R_k <_l R_l <_l R_m$ and $R_i <_r R_j <_r R_k <_r R_l <_r R_m$.*

Proof: We consider visibility amongst $\{U, D, R_i, R_j, R_k, R_l, R_m\}$. All of the left edges of these rectangles are visible from the large upmost rectangle U; this means that we may without loss of generality assume that R_m is the rectangle that has the topmost and the bottommost edges (it can be extended without blocking any visibilities amongst $\{U, D, R_i, R_j, R_k, R_l, R_m\}$). Similarly, the right edges of these rectangles are visible from D, and we may therefore assume that R_i extends as far above and below as R_m. In order to see R_m, the rectangles R_j, R_k, and R_l must all be visible from the upward direction not only on the left, but also somewhere to the right of the left edge of R_m (which is somewhere to the right of the left edge of R_l). Similarly, in order to see R_i, they must be visible from the downward direction from somewhere to the left of the right edge of R_i (or R_j).

We consider cases based on the layout of R_j and R_k. In order to see R_i, the rectangle R_k must extend either above or below R_j, or both. If R_k extends both above and below R_j, then R_j will not see R_m (see Figure 9a). Suppose that R_k extends only below R_j (or, symmetrically, only above R_k). Then, for R_l to see R_i, R_l must extend above R_j or below R_k or both. If R_l extends above R_j (or both above R_j and below R_k), then, in order to see R_k, the bottom edge of R_l must lie below the top edge of R_k (see Figure 9b). In this case, R_j is blocked from seeing R_m by R_k and R_l. If R_l extends below R_k, then its top edge must lie above the top edge of R_k for it to see R_j (see Figure 9c). Thus, R_l blocks

R_k from seeing R_m.

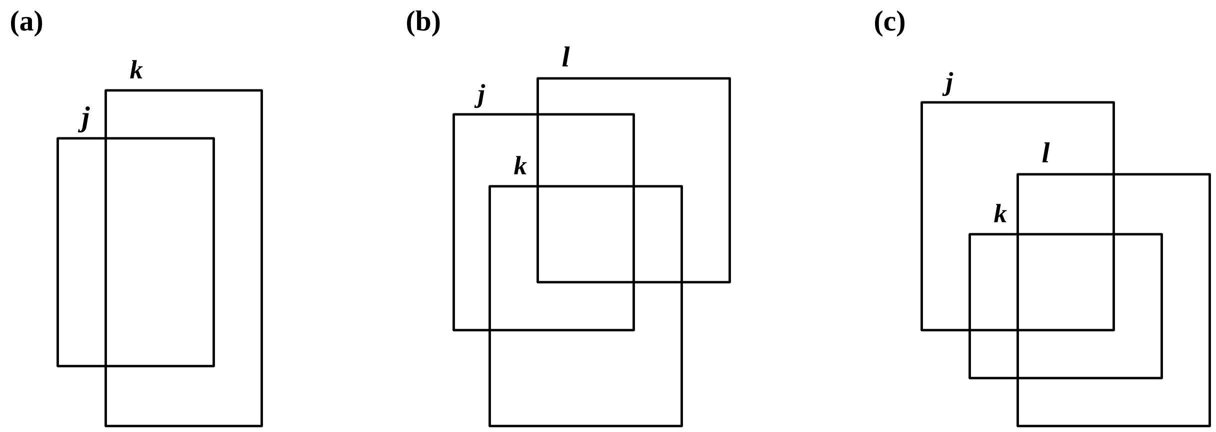

Figure 9: Trying to lay out $R_i <_l R_j <_l R_k <_l R_l <_l R_m$ and $R_i <_r R_j <_r R_k <_r R_l <_r R_m$

As every case has led to a contradiction, the lemma is proved. ■

We can now show an upper bound of 102 on the size of $\mathcal{R}$:

Theorem 5.4 *No clique of size greater than 102 admits a VR-representation.*

Proof: Suppose that there were some collection $\mathcal{R}$ of 103 (or more) rectangles that represent a clique. We remove U and D and concentrate on $\mathcal{R}'$ which has at least 101 elements.

List the z–coordinates of the rectangles in $\mathcal{R}'$ in increasing order in the order $<_l$. By Theorem 5.1 (with $j = k = 10$), there must be an ascending or descending subsequence of length 11. By symmetry (of z–coordinates), we may assume that this sequence is ascending. This sequence corresponds to a collection $\mathcal{T}$ of rectangles that have the same order in both z–coordinate and $<_l$.

Let T_1 be the upmost rectangle in $\mathcal{T}$. We consider visibility amongst $\mathcal{T} \cup \{U, D\}$. The rectangle T_1 could interfere only with visibility between the rectangle U and the remainder of $\mathcal{T}$; however, every element of $\mathcal{T}$ sees U along its left edge. Thus, we can assume that T_1 has the topmost, rightmost, and bottommost edges in $\mathcal{T}$. Let $\mathcal{T}'$ be $\mathcal{T} \setminus \{T_1\}$; $\mathcal{T}'$ has at least ten elements.

List the x–coordinates of the right edges of the rectangles of $\mathcal{T}'$ in increasing order in the order $<_l$. By Theorem 5.1 (with $j = k = 3$), there is an ascending or descending subsequence of length four. A descending subsequence of length four corresponds to rectangles in the situation prohibited by Lemma 5.2. An ascending subsequence of length four may be augmented with T_1 (as its last element) to create an ascending subsequence of length five. This corresponds to rectangles in the situation prohibited by Lemma 5.3. As neither case is permitted, our assumption that $\mathcal{R}$ has more than 102 elements is false. ■

There is a large gap between the upper and lower bounds. While we have constructed an explicit VR-representation for K_{20}, and hence for K_n, $n \leq 20$, we have also shown that K_n has no VR-representation for $n \geq 103$. We conjecture that $n = 20$ is close to the correct bound. The upper bound presented in this section also holds for squares, however, we believe that the correct upper bound for squares is even lower than that for rectangles.

6 VR-representability of bipartite graphs

In this section, we are concerned with VR-representations of complete bipartite graphs and graphs that are formed by removing a perfect matching from a complete bipartite graph. We denote the former type of graphs by $K_{m,n}$, where m and n are the sizes of each partition, and the latter type of graphs by $K_{j,j} - M$, where j is the size of both of the partitions; note that both partitions must have the same size for there to be a perfect matching. The graphs $K_{4,4} - M$ and $K_{5,5} - M$ are shown in Figure 10. When we speak of the *directed* $K_{j,j} - M$, we mean the normal $K_{j,j} - M$ with all edges directed from one of the two partitions to the other.

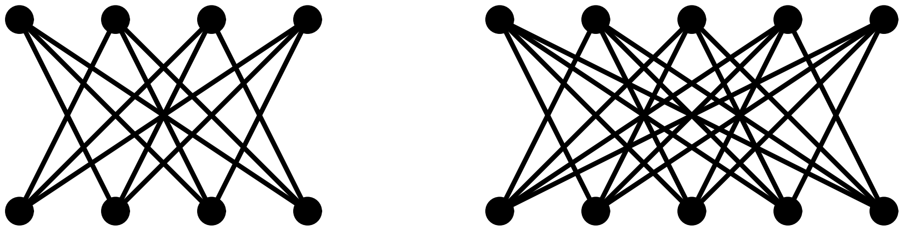

Figure 10: The graphs $K_{4,4} - M$ and $K_{5,5} - M$

We say that a set of rectangles represents a directed graph if it represents the underlying undirected graph and, for each directed edge from s to t, the rectangle corresponding to s is below the rectangle corresponding to t.

It is not hard to see that any $K_{m,n}$ admits a VR-representation (see Figure 11).

Theorem 6.1 *$K_{m,n}$ admits a VR-representation for any m, n.*

We now turn our attention to complete bipartite graphs minus a matching.

Theorem 6.2 *The directed version of $K_{4,4} - M$ is VR-representable.*

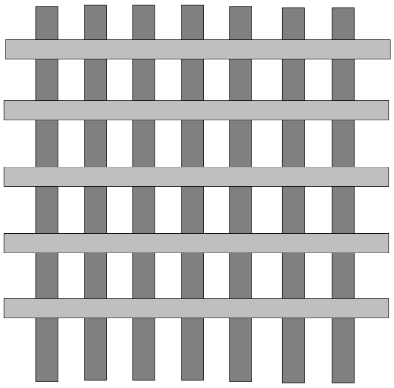

Figure 11: Any $K_{n,m}$ has a VR-representation.

Proof: The graph is represented as follows. Let one partition consist of the rectangles 1–4, and let the other consist of 5–8. (Recall that the label on a rectangle is its z–coordinate). For $1 \leq i \leq 4$, the rectangle i will be visible to all of the rectangles 5–8 except $i+4$. Figure 12 shows the construction, showing two projections of rectangles on the same area of the x–y plane. By inspection, this is a representation of $K_{4,4} - M$. As all of the rectangles in one partition (5–8) are above the rectangles in the other (1–4), it is also a representation of the directed $K_{4,4} - M$. ■

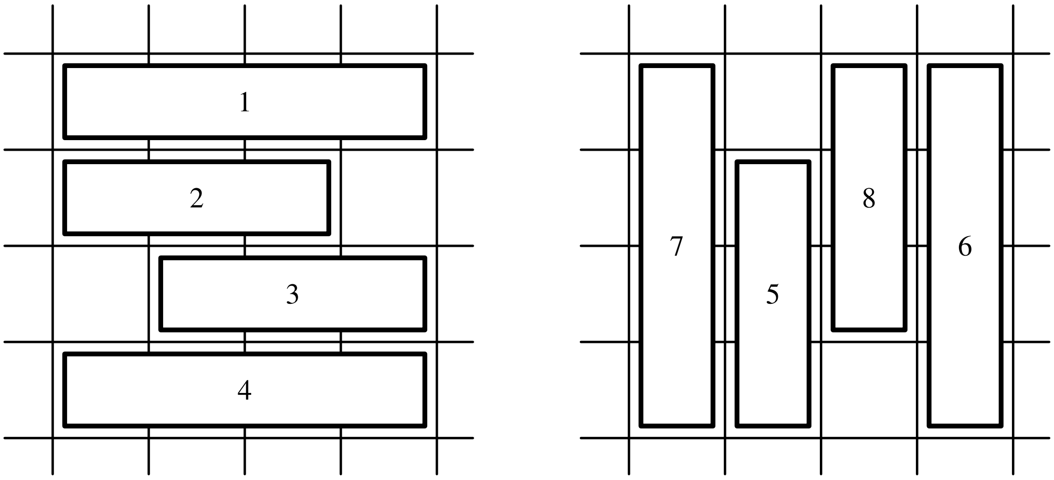

Figure 12: Representing the directed $K_{4,4} - M$

Theorem 6.3 *$K_{5,5} - M$ is VR-representable.*

Proof: We first modify the construction of Theorem 6.2 by expanding the rectangles of each set until their projections onto $z = 0$ touch. We then add a rectangle 0 to the partition 5–8, and a rectangle 9 to the partition 1–4. The rectangle 0 will not see the rectangle 9, but will see the other rectangles (1–4) in the partition containing 9. Similarly, the rectangle 9 will not see the rectangle 0, but will see the other rectangles (5–8) in the partition containing 0. The construction is shown in Figure 13, which shows four projections of rectangles onto the same area of the x–y plane. ■

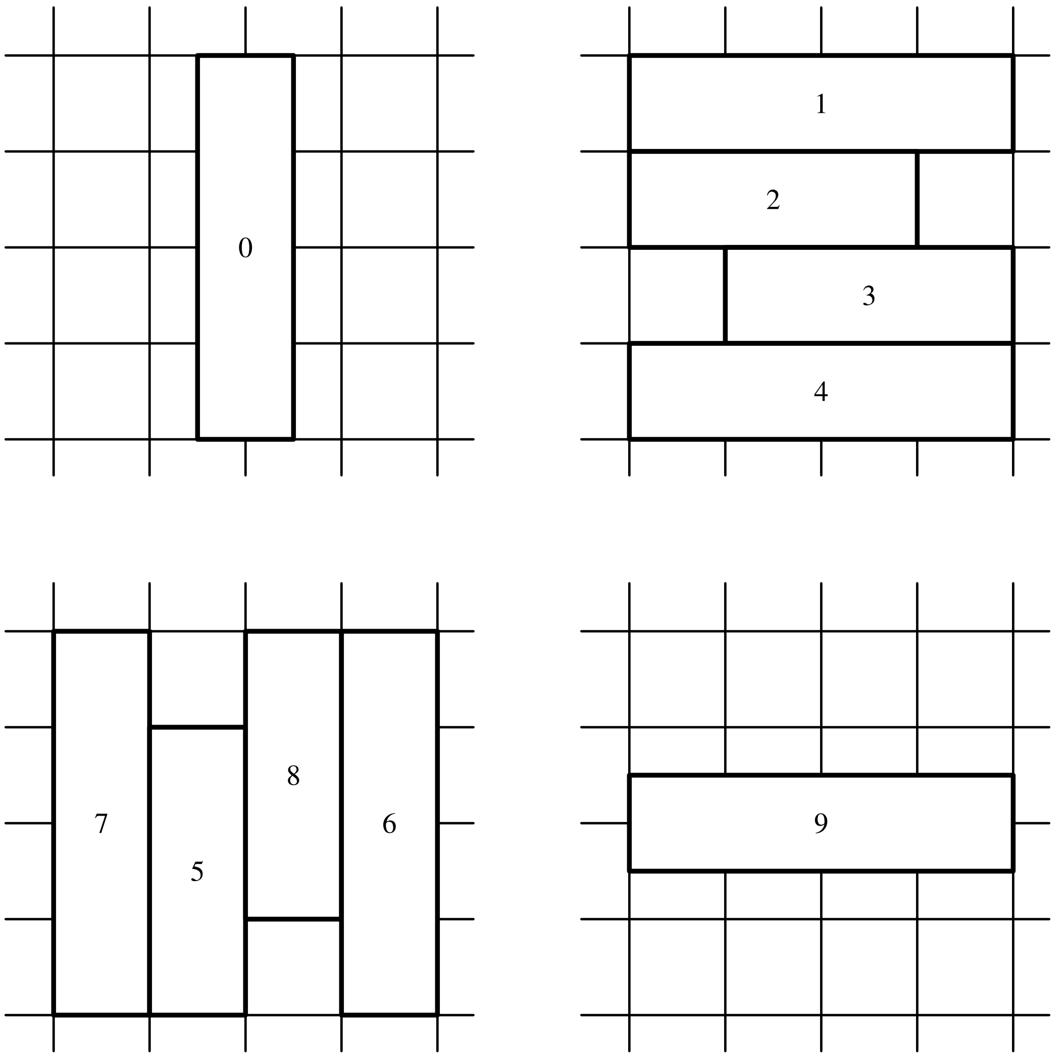

Figure 13: Representing $K_{5,5} - M$ using ϵ–visbility

Theorem 6.4 *The directed version of $K_{5,5} - M$ is not VR-representable.*

Proof: Consider the partition that is the source of the edges. We call the corresponding rectangles α, β, γ, δ, ϵ. Let the other class consist of the rectangles A, B, C, D, and E, where α does not see A, β does not see B, etc.

Suppose that there were a directed VR-representation. As there are no edges directed from $A \dots E$ to $\alpha \dots \epsilon$, no rectangle of $A \dots E$ lies down from any rectangle of $\alpha \dots \epsilon$ and we can thus pull the rectangles $\alpha \dots \epsilon$ down until all of $\alpha \dots \epsilon$ are down from all of $A \dots E$.

The rectangles α, β, γ, and δ are all intersected by E. The rectangle ϵ is not; we may assume without loss of generality that ϵ is entirely above E. We now shrink E as much as possible by moving its bottom edge abovewards, while retaining the property that the rectangles represent $K_{5,5} - M$.

Suppose that E shrank until it became a line segment. The four rectangles that intersect E would then occur without loss of generality in left-to-right order α, β, γ, and δ. We may also (by symmetry) assume that the top edge of β is above the top edge of γ (see Figure 14.) The rectangle B goes as far left as the right edge of α, as far right as the left edge of δ, as far down as the top edge of γ, and as far up as the bottom edge of ϵ. Such a B must intersect (and thus be seen by) the rectangle β, and thus E did not shrink to a segment.

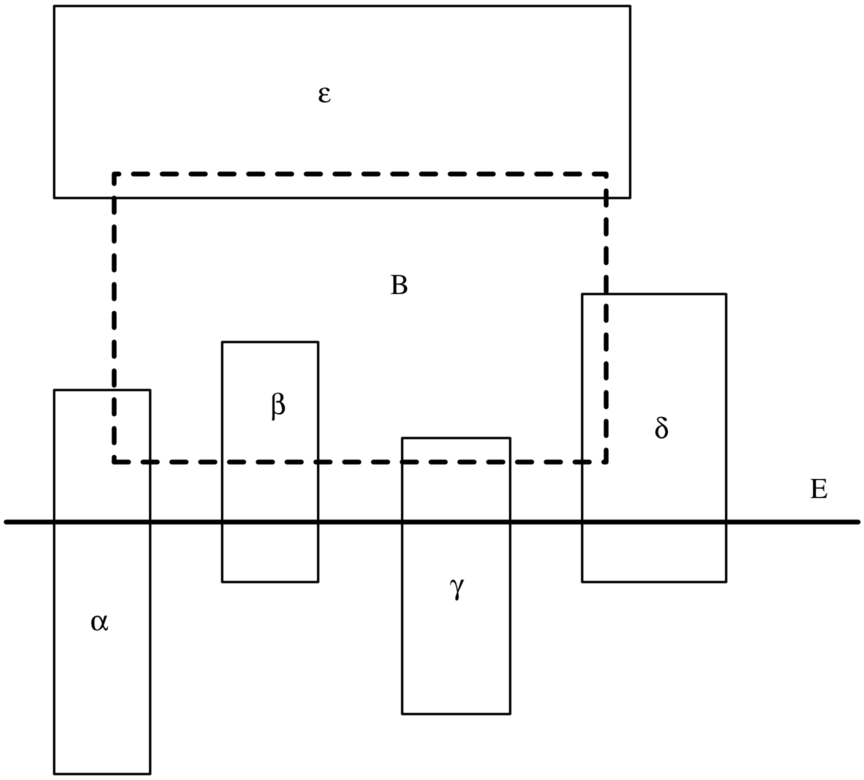

Figure 14: The rectangle E shrank to a segment

Thus E must have stopped shrinking because its bottom edge became collinear with the top edge of one of α, β, γ, δ. Without loss of generality, we may assume that it was α (see Figure 15). The rectangles B, C, and D all intersect α and ϵ, but not E. Each of these rectangles therefore must lie either left of E or right of E. Either two of them are to the left of E or at least two are to the right. Without loss of generality, we assume that the two are B and C and they are to the right of E, and that C is to the right of B (see Figure 16). However, in this situation, any rectangle that intersects both E and C must also intersect B. The rectangle β, in particular, must therefore intersect B, which is a contradiction. Thus, as each case has led to a contradiction, $K_{5,5} - M$ cannot be represented in the directed fashion. ∎

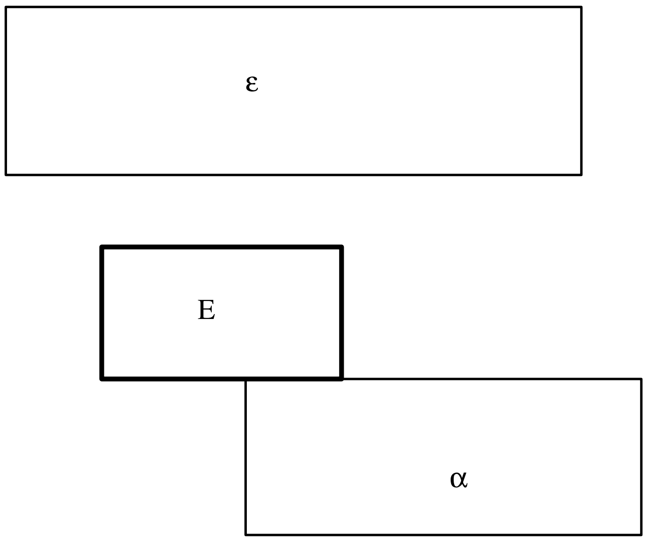

Figure 15: The rectangle E did not shrink to a segment

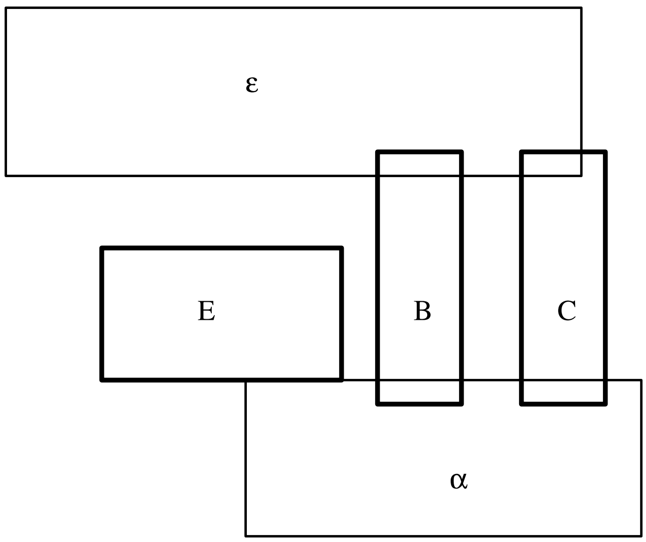

Figure 16: The rectangles E and C are separated by B

Conjecture 6.5 *$K_{6,6} - M$ is not VR-representable.*

7 VR-representations of tripartite graphs

In this section we examine VR-representations of tripartite graphs. Recall that a *tripartite graph* is a graph whose vertex set can be partitioned into three subsets such that each edge is between vertices in different subsets. A *complete tripartite graph* is a tripartite graph with tripartition (L, M, N) such that each vertex of L is joined to each vertex of M and N and also each vertex of M is joined to each vertex of N. If $|L| = l, |M| = m, |N| = n$, then the graph is denoted $K_{l,m,n}$.

In Figure 17, we have outlined a construction giving a VR-representation of $K_{m,n,2}$ for all m, n. In the figure, adjacent rectangles in M and N are not disjoint, but this can be changed by simply raising every other rectangle in M and N slightly in the z direction.

Theorem 7.1 *$K_{m,n,2}$ has a VR-representation for all m, n.*

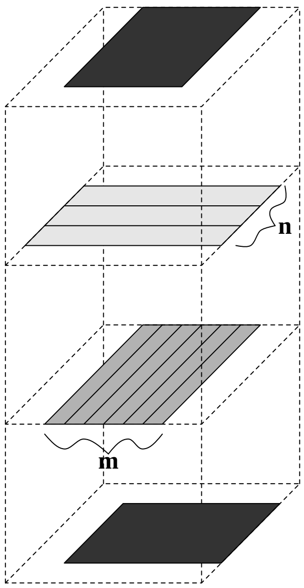

Figure 17: The complete tripartite graph $K_{m,n,2}$

However, for $m, n \geq 3$ this construction cannot be used to represent $K_{m,n,3}$. If an additional dark rectangle is placed in the construction, then it must be placed in between the set M of m rectangles and the set N of n rectangles; otherwise either it will not be visible to one of these rectangles or it will be visible to one of the other two dark rectangles. However, since this rectangle must span both the M rectangles and the N rectangles, it will block the visibility between the middle rectangles in M and N.

Using a different construction it is possible to represent $K_{m,4,3}$, as Figure 18 illustrates.

In this figure the z axis is horizontal. We believe, however, that this construction cannot be extended to values larger than 4 and 3.

Theorem 7.2 *$K_{m,4,3}$ has a VR-representation for all m.*

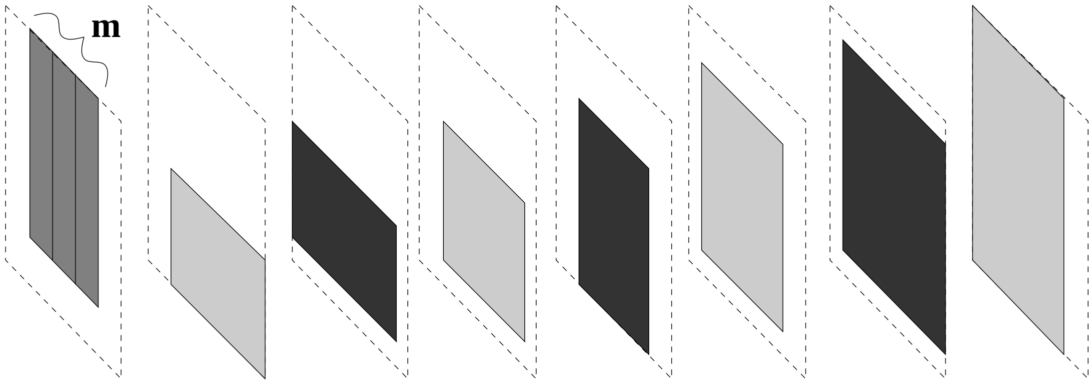

Figure 18: The complete tripartite graph $K_{m,4,3}$

The previous examples suggest the following open problem: What is the largest $K_{m,n,3}$ graph that is VR-representable?

8 Nonclosure under graph minors

In this section, we show that the family of VR-representable graphs is not closed under graph minors. To prove that the class of VR-representable graphs is not closed under graph minors, we consider the complete bipartite graph $K_{103,103}$, which is VR-representable as illustrated in Figure 11.

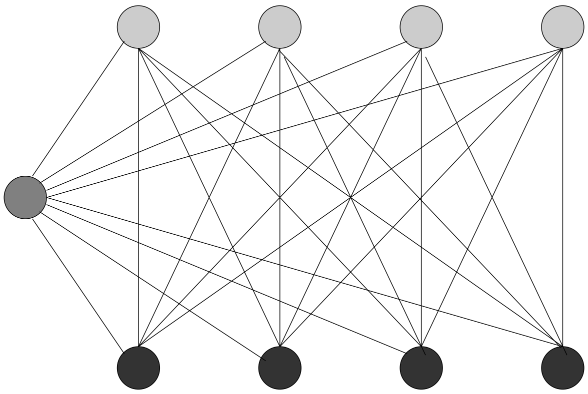

Figure 19: Contracting an edge of $K_{5,5}$

When an edge of a complete bipartite graph is contracted, it yields a vertex that is adjacent to all the vertices of the graph (see Figure 19 for an example of contracting an edge of $K_{5,5}$). If we contract the 103 edges of a perfect matching in $K_{103,103}$, then we obtain the complete graph K_{103}, which is not VR-representable by Theorem 5.4. This gives us the following result.

Theorem 8.1 *The class of VR-representable graphs is not closed under graph minors.*

9 Variations

In this section, we study some variations of VR-representations where rectangles are replaced by squares, discs, unit squares and unit discs. We study the effects these replacements have on the representation of complete graphs and complete bipartite graphs.

9.1 VR-representations by squares and unit squares

In this subsection, we deal with VR-representations using squares and unit squares instead of rectangles. While the upper and lower bounds for the largest possible n allowing a VR-representation for K_n remain the same even if we consider squares instead of rectangles, it is no longer clear that any $K_{n,m}$ can be represented using squares.

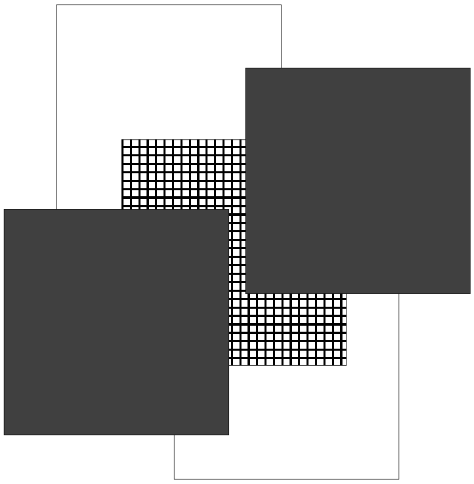

Figure 20: Representing $K_{3,3}$ by unit squares.

Theorem 9.1 *$K_{3,3}$ has a VR-representation using unit squares.*

Proof: Consider Figure 20. The two classes of squares are shown dark and light. The plaid square in the middle corresponds to a dark square that is exactly covered by an equally sized light square. ∎

Conjecture 9.2 *$K_{4,3}$ does not have a VR-representation by squares.*

The above conjecture and theorem suggest that the representation of $K_{n,m}$, $n, m \geq 3$, by squares seems to be the same whether or not we add the requirement that all squares be of the same size, i.e. unit squares. On the other hand, the K_{20} construction no longer holds for unit squares. The largest complete graph that we can represent by unit squares is of considerably smaller size. This should not come as a surprise as unit squares are much more restrictive than abitrary squares.

Theorem 9.3 *K_7 has a VR-representation using unit squares.*

Proof: See Figure 21. The six phases indicate how the seven squares are placed on top of each other. As each square is added, it is visible from all the squares below it. ∎

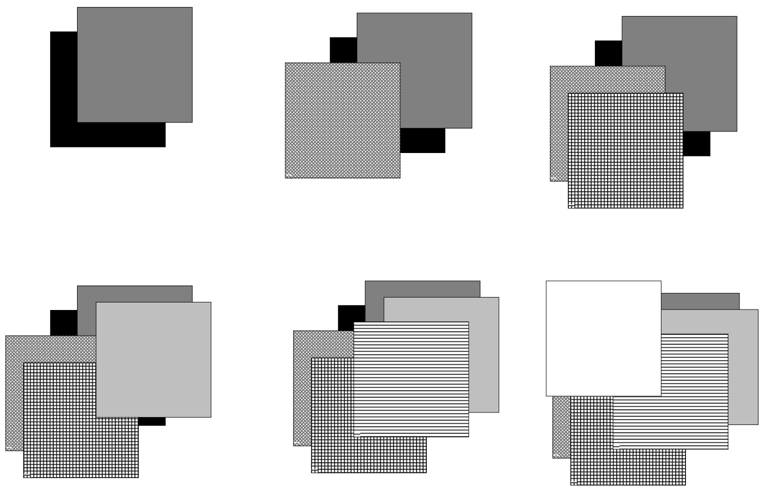

Figure 21: Representing K_7 by unit squares.

We now prove an upper bound on the size of the largest clique admitting a VR-representation using unit squares. Consider an arrangement of unit squares S that has a clique for its visibility graph.

Lemma 9.4 *Let S_i, S_j, S_k, S_l be four unit squares of S with $i < j < k < l$. It cannot be the case that $S_i <_l S_j <_l S_k <_l S_l$.*

Proof: Note that with unit squares, if $S_i <_l S_j <_l S_k <_l S_l$, then also $S_i <_r S_j <_r S_k <_r S_l$. In order for S_i to see S_k, S_i and S_k must both extend below (or, symmetrically, above) S_j (see Figure 22). But then S_l cannot see both S_i and S_j because S_k blocks it. Since this leads to a contradiction, the lemma is proved. ∎

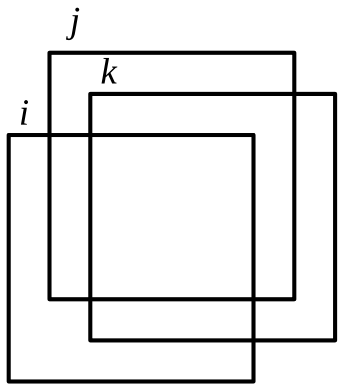

Figure 22: Trying to lay out $S_i <_l S_j <_l S_k <_l S_l$

Theorem 9.5 *K_{10} does not have a VR-representation using unit squares.*

Proof: Suppose there were some collection S of 10 unit squares that represented K_{10}. Sort the squares of S in increasing order according to $<_l$ and list their z-coordinates. By Theorem 5.1 (with $j = k = 3$), there must be an ascending or descending subsequence of length 4. By symmetry of z-coordinates, we may assume that this sequence is ascending. This sequence corresponds to a collection of 4 squares that have the same order in both z-coordinate and $<_l$, but this situation is prohibited by Lemma 9.4. Therefore, our assumption that K_{10} has a VR-representation is false. ∎

There is a gap between our upper and lower bounds; we conjecture that the correct bound for unit squares is the following.

Conjecture 9.6 *K_8 does not have a VR-representation of unit squares.*

9.2 VR-representations with discs and unit discs

We now consider the situation for discs. We show that any K_n can be represented when rectangles are replaced by discs, while there seem to be strict bounds on possible $K_{n,m}$ that can be represented by discs. This is just the opposite situation as for rectangles.

Theorem 9.7 *Any K_n has a VR-representation using unit discs.*

Proof: See Figure 23. Let $\frac{1}{2} < r < 1$ and consider the n points $m_1 = (x_1, y_1, 0), \ldots, m_n = (x_n, y_n, 0)$, equally spaced on a semicircle of radius r around the point $M_0 = (0, 0, 0)$. Let c_i be the unit disc orthogonal to the z-axis around the point m_i and let $C_i = c_i + (0, 0, i)$ as well as $M_i = m_i + (0, 0, i)$. We claim that the C_i form a VR-representation using discs for K_n.

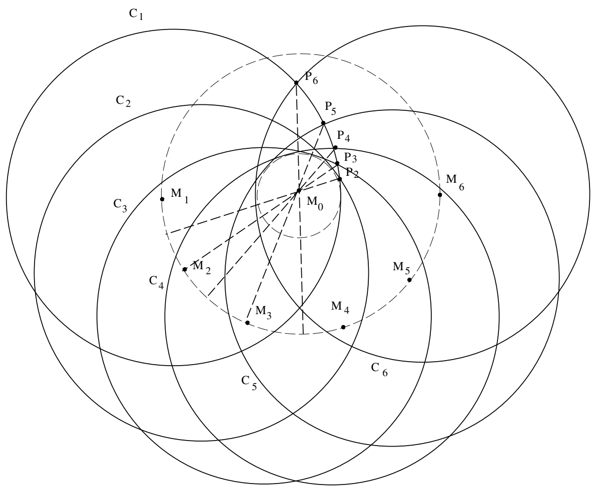

Figure 23: Representing K_n by n unit discs.

By symmetry, it suffices to show that any $C_i, i \geq 2$ can see C_1. For this purpose, consider the points p_i that we get by selecting one of the two intersection points of the boundary of c_i with c_1, as indicated in Figure 23. The m_i form a counterclockwise sequence of points around M_0. Since they lie on the orthogonal bisector of the line $\overline{m_1, m_i}$ that runs through M_0, the p_i also form a counterclockwise sequence, implying that none of the points $P_i = p_i + (0, 0, i)$ on the disc C_i can be obstructed by a disc C_j, $1 < j < i$ from seeing C_1. ■

With respect to bipartite graphs, we present the following constructions for discs and unit discs. First we show how to represent $K_{3,3}$ by unit discs and then $K_{n,3}$ by discs of arbitrary size.

Theorem 9.8 *$K_{3,3}$ has a VR-representation using unit discs.*

Proof: See Figure 24. The two sets of discs are drawn "light" and "dark". As in Figure 20, the plaid disc corresponds to a degenerate pair of discs, one dark, one light that are placed such that the light one exactly covers the dark one, when seen from above. ■

Theorem 9.9 *$K_{n,3}$ has a VR-representation using discs.*

Proof: See Figure 25. (For the sake of clarity, we have only shown parts of two larger discs.) As before, the plaid disc corresponds to a degenerate pair of discs, one dark, one light that are placed such that the dark one exactly covers the light one, when seen from above. In addition, we line up $n-1$ smaller dark discs below, such that the interior of all n dark discs can be intersected by one y-orthogonal hyperplane. Then we can place two

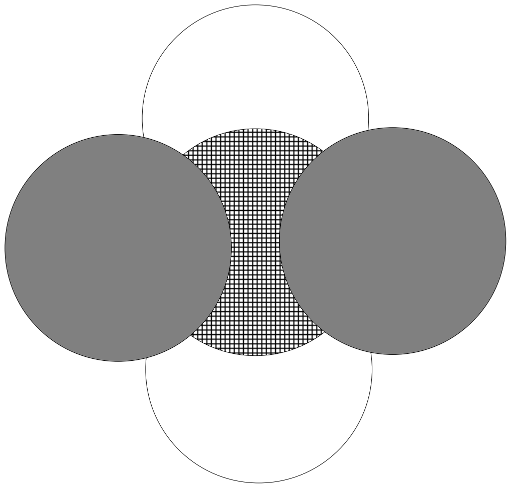

Figure 24: Representing $K_{3,3}$ by unit discs.

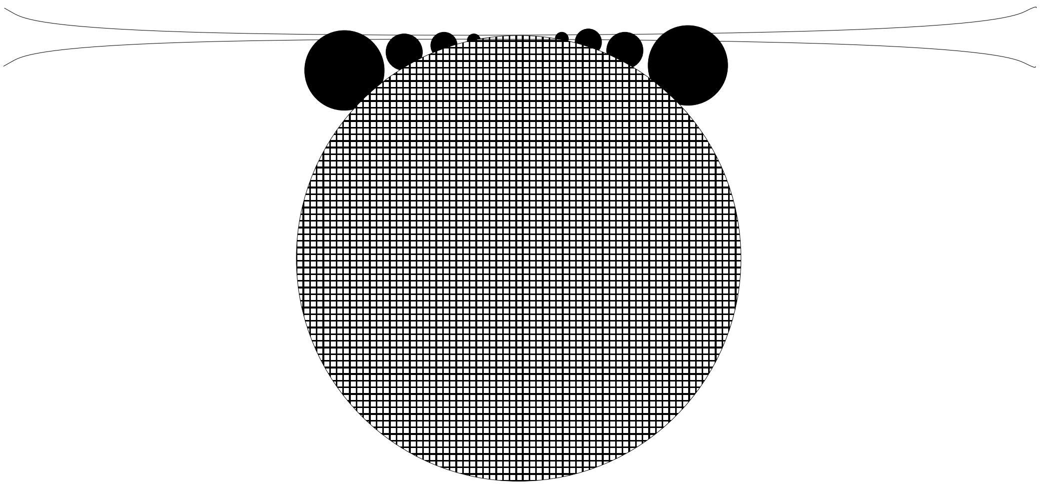

Figure 25: Representing $K_{n,3}$ by discs.

sufficiently large light discs above the arrangement, such that they both can see all the dark discs, while not being visible to each other. ∎

We end this section with the following two conjectures.

Conjecture 9.10 *$K_{4,4}$ does not have a VR-representation of discs.*

Conjecture 9.11 *$K_{4,3}$ does not have a VR-representation of unit discs.*

10 Conclusions

We have studied the visibility representations of graphs where each vertex of the graph maps to a closed, disjoint rectangle in 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 directions, and edges are expressed by ϵ-visibility restricted to the z-direction.

We showed that all planar graphs are VR-representable, K_n is VR-representable for values of $n \leq 20$, and K_n is not VR-representable for $n \geq 103$. Concerning bipartite graphs, we showed that $K_{m,n}$ is VR-representable for all m and n, and provided bounds on the representability of complete bipartite graphs minus a perfect matching. We also provided bounds on the representability of tripartite graphs. Furthermore, we showed that the family of VR-representable graphs is not closed under graph minors. Finally, we considered variant VR-representations, where the rectangles are replaced by squares and discs and provide bounds on the complete graphs and complete bipartite graphs that admit such representations.

Many directions for further research have been generated from this initial investigation as evidenced by the open problems and conjectures mentioned throughout the paper.

11 Acknowledgments

Our study of VR-representations began at Bellairs Research Institute of McGill University during the 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, Steve Skiena, Roberto Tamassia, Yanni Tollis, and Godfried Toussaint.

References

[1] Dean, A.M. and Hutchinson, J.P., "Rectangular visibility representations of bipartite graphs", Submitted for publication, 1994.

[2] Di Battista, G., Eades, P., Tamassia, R., and Tollis, I.G., Algorithms for Automatic Graph Drawing: An Annotated Bibliography, Technical Report, Dept. of Computer Science, Brown University, 1993.

[3] Erdös, P. and Szekeres, A., "A combinatorial problem in geometry, " *Compositio Mathematica* v. 2 (1935), 463-470.

[4] Kant, G., Liotta, G., Tamassia, R., and Tollis, I.G., "Area Requirements of Visibility Representations of Trees", *Proceedings of the 5th Canadian Conf. on Comp. Geom.*, Waterloo, Ontario, (1993) 192-197.

[5] Tamassia, R. and Tollis, I. G., "A unified approach to visibility representations of planar graphs," *Discrete Comput. Geom.* v. 1 (1986), 321-341.

[6] Wismath, Stephen Kenneth. Bar-Representable Visibility Graphs and a Related Network Flow Problem. University of British Columbia, Dept. of Computer Science Technical Report 89-24, August 1989.