Journal of Graph Algorithms and Applications
http://jgaa.info/ vol. 25, no. 1, pp. 59–96 (2021)
DOI: 10.7155/jgaa.00550

Minimal Embedding Dimensions of Rectangle k-Visibility Graphs

Espen Slettnes

MIT PRIMES-USA

Submitted: August 2019	Reviewed: May 2020	Revised: September 2020
Reviewed: December 2020	Revised: December 2020	Accepted: January 2021
Final: January 2021	Published: January 2021	
Article type: Regular paper	Communicated by: H. Meijer	

Abstract. Bar visibility graphs were adopted in the 1980s as a model to represent traces, e.g., on circuit boards and in VLSI chip designs. Two generalizations of bar visibility graphs, rectangle visibility graphs and bar k-visibility graphs, were subsequently introduced.

Here, we combine bar k- and rectangle visibility graphs to form rectangle k-visibility graphs (RkVGs), and further generalize these to higher dimensions. A graph is a d-dimensional RkVG if and only if it can be represented with vertices as disjoint axis-aligned hyperrectangles in d-space, such that there is an axis-parallel line of sight between two hyperrectangles that intersects at most k other hyperrectangles if and only if there is an edge between the two corresponding vertices.

For any graph G and a fixed k, we prove that given enough spatial dimensions, G has a rectangle k-visibility representation, and thus we define the minimal embedding dimension (MED) with k-visibility of G to be the smallest d such that G is a d-dimensional RkVG. We study the properties of MEDs and find upper bounds on the MEDs of various types of graphs. In particular, we find that the k-visibility MED of the complete graph on m vertices K_m is at most $\lceil m/(2(k+1)) \rceil$, of complete r-partite graphs is at most $r+1$, and of the m^{th} hypercube graph Q_m is at most $\lceil 2m/3 \rceil$ in general, and at most $\lfloor \sqrt{m} \rceil$ for $k=0,\ m \neq 2$.

E-mail address: espen@slett.net (Espen Slettnes)

1 Introduction

Bar visibility graphs were introduced in the 1980s as a way to model circuit traces in VLSI chip designs by Lodi and Pagli [16]. A graph G is a bar visibility graph if there is a one-to-one correspondence between its vertices and 2D horizontal bars, such that there is an unobstructed vertical line of sight between two bars (i.e., a vertical line segment between the two bars not intersecting other bars) if and only if there is an edge between the corresponding vertices in G. Note that the bars and visibility lines form a planar graph drawing of G.

In their 1997 paper *On Rectangle Visibility Graphs* [2], Bose et al. introduced rectangle visibility graphs as "a graph in the plane so that the vertices of the graph are rectangles that are aligned with the axes, and the edges of the graph are horizontal or vertical lines-of-sight". Previously, though using different terminology, Stephen Wismath established in his 1989 thesis [18] that all planar graphs are rectangle visibility graphs (i.e., have rectangle visibility representations).

Dean et al. introduced bar k-visibility graphs in 2007 [6] as a generalization in which visibility lines between the bars are relaxed from being unobstructed to being obstructed by at most k other bars. Hartke et al. published *Further Results on Bar k-Visibility Graphs* [13], and in combination these two papers established that the maximum number of edges in a bar k-visibility graph on n vertices is $(k+1)(3n-4k-6)$. Dean et al. also proved that the thickness* of every bar 1-visibility graph is at most 4, and Chang et al. [4] proved that the thickness of a bar k-visibility graph is at most $3k+3$.

Others, such as Babbitt et al. [1], have studied k-visibility on other types of visibility representations. Here we define a *rectangle k-visibility graph* to be a graph that can be represented with vertices as disjoint axis-parallel rectangles, where there is an edge between two vertices if and only if there is an axis-parallel line of sight, obstructed by at most k other rectangles, between the corresponding rectangles. By the above, as edges corresponding to horizontal as well as vertical visibility lines form bar k-visibility graphs, the number of edges and the thickness in a rectangle k-visibility graph are at most $2(k+1)(3n-4k-6)$ and $6k+6$, respectively. In particular, the respective thicknesses of rectangle 0- and 1-visibility graphs are at most 2 and 8.

Prior research has further generalized rectangle visibility graphs into 3 dimensions, where they are referred to as *box* visibility graphs [9]. Similar in spirit to the Euclidean dimension of a graph, which is the minimum number of dimensions for which it is a (strict) unit distance graph, here we consider a generalization of rectangle k-visibility graphs into higher dimensions and study the minimum dimension needed to represent various graphs with k-visibility for a fixed k. For example, as discussed above, the minimal embedding dimension (MED) of a planar graph given $k=0$ is at most 2.

We study such MEDs on general graphs in Section 3. Among other things, we show that the MED of a nonempty graph G on n vertices is at most $\lceil \frac{n}{2} \rceil$, that the MED of a disconnected graph G is the maximum of 2 and the MEDs of its connected components, and that MEDs are subadditive under the Cartesian product.

*The thickness of a graph is the minimum number of planar graphs into which its edges can be partitioned.

We then move on to specific graphs. We cover complete graphs in Section 4, where we establish that MEDs can be arbitrarily large and that the MED of the complete graph on m vertices K_m is at most $\max\left(3, \left\lceil \frac{m-22(\lfloor k/2 \rfloor+1)}{2(k+1)} \right\rceil + 1\right)$. In Section 5 on multipartite graphs we find that the MED of a complete r-partite graph is at most $r+1$. Section 6 is devoted to hypercubes; we show that the MED of Q_m is at most $\lceil \frac{2}{3}m \rceil$, and at most $\lfloor \sqrt{m} \rceil$ for $k = 0, m \neq 2$. Finally, in Appendix A we compare selected results from this paper to results about other graph dimensions.

2 Terminology

We define a d-dimensional *rectangle visibility graph* (RVG^s) to be a graph where vertices can be represented as disjoint (closed hyper-)rectangles in d dimensions, and edges as *all* axis-parallel lines of sight between (i.e., unobstructed line segments connecting) these (hyper-)rectangles. We also define these variants:

- An $\epsilon-$*visibility graph* (RVG$^\epsilon$) imposes a positive thickness to the line of sight between rectangles, such that the rectangles must overlap by a positive amount in all $(d-1)$ dimensions orthogonal to the line of sight. In contrast, a *strong visibility graph* (RVG^s) allows visibility lines with zero thickness, i.e. zero overlap along orthogonal directions.
- A *rectangle k-visibility graph* (RkVG^s, RkVG$^\epsilon$) allows the line of sight to be obstructed by up to k other rectangles.
- A *unit rectangle visibility graph* (URVG^s, URVG$^\epsilon$, URkVG^s, URkVG$^\epsilon$) imposes the restriction that all (hyper-)rectangles are unit hypercubes.

Unless explicitly stated, we use the term *rectangle* to mean *d-dimensional hyper-rectangle*. As a special case, a *box* is a 3-dimensional rectangle.

The *minimal embedding dimension* (MED) of a graph G is the smallest number of spatial dimensions d for the graph to be a specific one of the above. We denote by $M^s(G), \mu^s(G), M^s_k(G), \mu^s_k(G), M^\epsilon(G), \mu^\epsilon(G), M^\epsilon_k(G)$, and $\mu^\epsilon_k(G)$ the MEDs of G as a RVG^s, URVG^s, RkVG^s, URkVG^s, RVG$^\epsilon$, URVG$^\epsilon$, RkVG$^\epsilon$, and URkVG$^\epsilon$, respectively.

Example 1 *$\mu^s_1(C_4) = 2$ is the smallest number of dimensions in which we can represent C_4 as a unit rectangle 1-visibility graph with strong visibility.*

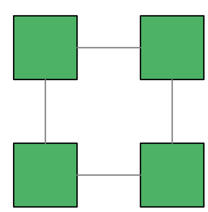

Additionally, we use the following conventions:

- G will be a simple graph. (We do not consider the null graph on zero vertices.)
- $n := |V(G)| \geq 1$ is the number of vertices (i.e., size) of G.
- A graph is *empty* if it has no edges.
- The $^\epsilon$ or ^{s} superscript may be omitted, in which case the strong and ϵ-visibility models can each be applied consistently.

Example 2 *"G is an $M(G)$-dimensional RVG" is always true because G is an $M^s(G)$-dimensional RVG^s and G is an $M^\epsilon(G)$-dimensional RVG$^\epsilon$.*

- All occurrences of $\left[\begin{smallmatrix}\mu\\M\end{smallmatrix}\right]$ can be consistently replaced by either μ or M.

  Example 3 *"G is a $\left[\begin{smallmatrix}\mu\\M\end{smallmatrix}\right](G)$-dimensional (U)RVG" is always true because G is a $\mu(G)$-dimensional URVG and G is an $M(G)$-dimensional RVG.*

3 General Graphs

3.1 Existence of the MEDs

Here we will prove that the minimal embedding dimension is well-defined, i.e. that every graph has a minimal embedding dimension. To that end, we first show how to think of a representation of a d-dimensional (U)RkVG (for large k) in terms of its projections to the axes.

Definition 4 *A graph G is an interval graph if there is a one-to-one correspondence between its vertices and a set of (closed) intervals, such that two intervals overlap if and only if there is an edge between the corresponding vertices in G.*

A unit interval graph, more commonly known as an indifference graph, is an interval graph that can be represented with unit intervals.

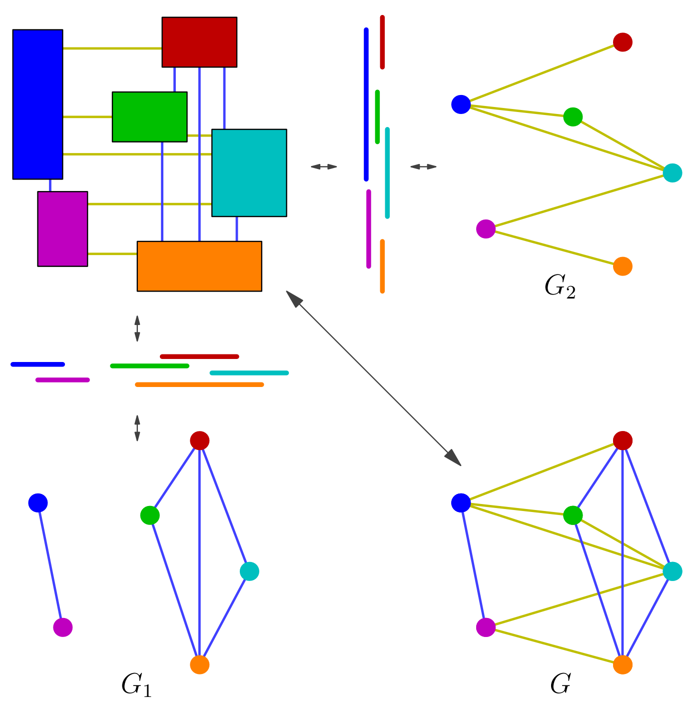

Figure 1: A graph G represented as a 2-dimensional RkVG, with projected intervals in each dimension corresponding to vertices in interval graphs G_1 and G_2

Lemma 1 *A graph G with n vertices is a d-dimensional (U)RkVG^s, where $k \geq n-2$, if and only if there exist d (unit) interval graphs $G_1, \dots, G_d$, each on the same vertex set as G, such that no edge is contained in all of $G_1, \dots, G_d$ and two vertices $u, v \in G$ are adjacent if and only if they are adjacent in all but exactly one of $G_1, \dots, G_d$.*

Proof: We note that $k \geq n-2$ is the same as infinite visibility, as at most $n-2$ rectangles can obstruct a visibility line between any two rectangles.

First we go from a d-dimensional (U)RkVG^s G to corresponding (unit) interval graphs $G_1, \dots, G_d$.

Consider the projections of all rectangles onto each of the axes of $\mathbb{R}^d$. Let G_i be the (unit) interval graph formed by the projection onto the i^{th} axis. Two rectangles cannot overlap in all of these projections, lest they would themselves overlap. In other words, no edge can be in all of $G_1, \dots, G_d$.

Two rectangles can see each other via a visibility line in the direction of the i^{th} axis $(1 \leq i \leq d)$ if and only if their respective projections do not overlap on the i^{th} axis, but overlap on all other axes $j \neq i$ for $(1 \leq j \leq d)$. In other words, two vertices G are adjacent if and only if they are adjacent in all but exactly one of $G_1, \dots, G_d$.

Then to construct a (U)RkVG^s representation if we have a set of (unit) interval graphs $G_1, \dots, G_d$, we can simply take the arrangement of rectangles for which the (unit) rectangle projections onto the axes correspond to the (unit) interval representations of G_1 through G_d. □

With this in mind, we now construct a representation of any graph G as a (U)RkVG by specifying its projections.

Theorem 1 *Every graph has a minimal embedding dimension as a (U)RkVG. Specifically, for a graph G on n vertices, $\left[{}^{\mu}_{M}\right]_k(G) \leq n$.*

Proof: Let G's vertex set be $[n] = \{1, 2, \dots, n\}$, and let

$$S_i(u) = \begin{cases} [0, 1] & \text{if } u = i \\ \left[\frac{2}{3}, \frac{5}{3}\right] & \text{if } i \sim_G u \text{ and } u > i \\ \left[\frac{4}{3}, \frac{7}{3}\right] & \text{if } i \not\sim_G u \text{ or } u < i \end{cases}$$

for $i, u \in [n]$. ("$\sim_G$" denotes the adjacency relation in G.)

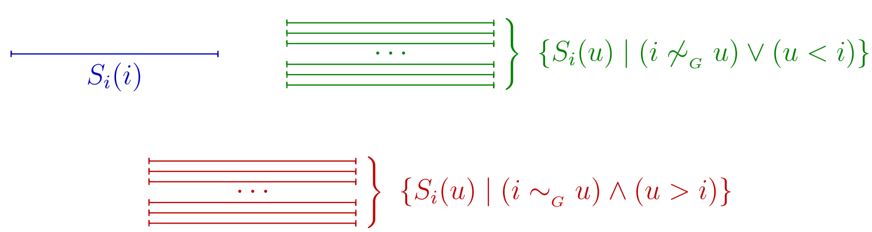

Figure 2: The n unit intervals $\{S_i(v) \mid v \in [n]\}$
(with artificial elevations added for illustration)

Let G_i be the (unit) interval graph formed by S_i. Note that

(a) In S_i's range, there is no interval strictly between two other intervals,
(b) any two intervals $S_i(u), S_i(v)$ where $u, v \neq i$ overlap,
(c) if $u \not\sim_G v$, intervals $S_i(u)$ and $S_i(v)$ do not overlap for $i \in \{u, v\}$,
(d) if $u \sim_G v$ and $u < v$, intervals $S_i(u)$ and $S_i(v)$ overlap for $i = u$ but not for $i = v$ and
(e) all overlaps are positive.

By (c) and (d), no edge is in all of the G_i's. By (b) and (d), if $u \sim_G v$, they are adjacent in all but one G_i representation. Finally, by (c), if $u \not\sim_G v$, they are not adjacent in two G_i's. Thus, by Lemma 1, G is a d-dimensional (U)R$(n-2)$VG^s.

By (a), no rectangle can block a visibility line between two others, and by (e), strong vs. ϵ-visibility doesn't matter, so G is also a d-dimensional URkVG. □

3.2 Basic Properties

We now make the following observations about minimal embedding dimensions:

Lemma 2 *Given a graph G on n vertices, $M_k(G) \leq \mu_k(G)$.*

Proof: Any representation of G as a URkVG in $\mu_k(G)$ dimensions is also a valid representation of G as RkVG, thus $M_k(G) \leq \mu_k(G)$. □

Lemma 3 *Given a graph G on n vertices, $\left[{\mu \atop M}\right]_k^\epsilon(G) \leq \left[{\mu \atop M}\right]_k^s(G)$.*

Proof: Given a representation of G as a (U)RkVG^s, let δ_i be the smallest nonzero difference between the i^{th} coordinates of any two of its hyperrectangles. In each dimension i, expand the rectangles by a margin of $\frac{\delta_i}{3}$.

No new strong-visibility lines have been created or destroyed, as the pairs of rectangles which overlapped have not changed in any dimension. Moreover, any two rectangles that previously had any overlap now have a positive overlap. Scaling the representation along the i^{th} axis by a factor of $1/\left(1 + 2\frac{\delta_i}{3}\right)$ so that the rectangles return to their original size, we now have a representation of G as a (U)RkVG$^\epsilon$, so $\left[{\mu \atop M}\right]_k^\epsilon(G) \leq \left[{\mu \atop M}\right]_k^s(G)$, as desired. □

Lemma 4 *A graph G is a d-dimensional (U)Rk VG if and only if $d \geq \left[{\mu \atop M}\right]_k(G)$.*

Proof: The former implies the latter by definition.

The latter implies the former because we can take a representation of G in $\left[{\mu \atop M}\right]_k(G)$ dimensions, place it in d-dimensional space, and thicken it by 1 unit in the remaining $\left(d - \left[{\mu \atop M}\right]_k(G)\right)$ dimensions. □

Lemma 5 *For any nonempty graph G, if $a, b > \omega(G)$,*

$$\begin{bmatrix}\mu \\ M\end{bmatrix}_{\omega(G)-2}(G) \leq \begin{bmatrix}\mu \\ M\end{bmatrix}_{a-2}(G) = \begin{bmatrix}\mu \\ M\end{bmatrix}_{b-2}(G),$$

where $\omega(G)$ denotes the size of G's maximum clique.

Proof: No representation of G can have a visibility line between two vertices that passes through $> \omega(G) - 2$ others, as that would form an $> \omega(G)$-clique. Thus, all visibility lines in representations of G are $(\omega(G) - 2)$-visibility lines, so any $(a - 2)$-visibility representation is a $(b - 2)$-visibility representation and vice versa, and additionally, any $(a - 2)$-visibility representation is an $(\omega(G) - 2)$-visibility representation, as desired. □

Because the chromatic number of G, $\chi(G)$, is at least $\omega(G)$, we get the following corollary:

Corollary 5 *Lemma 5 holds for $\chi(G)$ in place of $\omega(G)$.*

3.3 MEDs as RkVGs

Theorem 2 *Let G be a nonempty graph on n vertices. Then, $M_k(G) \leq \left\lceil \frac{n}{2} \right\rceil$.*

Proof: Let $S = \{v_1, \ldots, v_n\}$ be the vertices of G, where, WLOG, v_n shares an edge with v_{n-1} if $n \geq 2$. We divide S into subsets of at most 4 vertices, such that $S_m = \left\{v_{4(m-1)+1}, \ldots, v_{\min(n,4m)}\right\}$ for $m \in \left[1, \lceil \frac{n}{4} \rceil\right]$. Let G_m be the induced subgraph formed by vertices in S_m. Note that if $\left|S_{\lceil \frac{n}{4} \rceil}\right| \geq 2$, $G_{\lceil \frac{n}{4} \rceil}$ has at least one edge.

Let $\mathcal{S}_1, \ldots, \mathcal{S}_{\lceil \frac{n-2}{4} \rceil}$ be orthogonal 2-dimensional spaces, and if $n \equiv 1 \pmod 4$ or $n \equiv 2 \pmod 4$, let $\mathcal{S}_{\lceil \frac{n}{4} \rceil}$ be an additional orthogonal 1-dimensional space. We will construct a rectangle visibility representation of G by constructing its projections onto these spaces.

The projection of S_m onto $\mathcal{S}_m$ will be one of the arrangements in Figure 3, such that the visibility graph formed between the green rectangles is G_m (all possible values of G_m are covered in Figure 3).

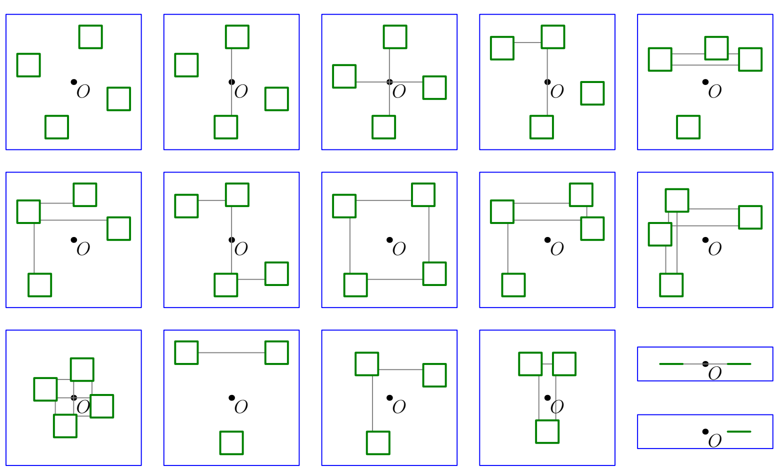

Figure 3: All possible projections of vertices $v_{4m-3}, \ldots, v_{\min(4m,n)}$ into $\mathcal{S}_m$

Let T_m be the complimentary set $S \setminus S_m$. We project every other vertex $v_i \in T_m$ onto the same 2-dimensional subspace $\mathcal{S}_m$ in such a way that each projection covers the central point O, and either overlaps or is adjacent to each vertex in $v_j \in S_m$. Some possible projections are illustrated as orange rectangles in Figure 4.

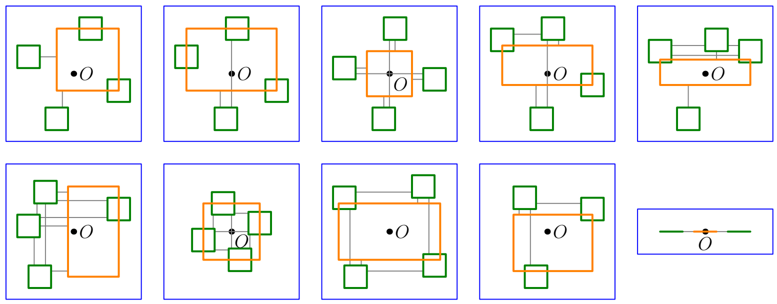

Figure 4: Sample projections of an additional vertex v_i, overlapping a central point O and either overlapping or adjacent to each of $v_{4m-3}, \ldots, v_{\min(4m,n)}$

We use the following rules:

- If $i < j$, the projections of v_i and v_j in $\mathcal{S}_m$ will *not* overlap; this counts as being disjoint in one dimension.
- If $i > j$, the projections of v_i and v_j in $\mathcal{S}_m$ will overlap if and only if $v_i \sim v_j$. If not, this counts as being disjoint in a second dimension, thus precluding any axis-parallel visibility line between the corresponding rectangles.

We note that every vertex $v_i \in T_m$ overlaps with point O in $\mathcal{S}_m$; thus there are no more disjoint projections than those described here.

By construction, we now have a representation of G where all pairs of vertices (v_i, v_j) are disjoint in one dimension if they are adjacent, and in two dimensions if they are non-adjacent. Moreover, there does not exist any third vertex v_k that blocks visibility between v_i and v_j (in particular in $\mathcal{S}_{\lceil \frac{k}{4} \rceil}$ at O for $k \notin S_i, S_j$), so no rectangle can block a visibility line. Thus by Lemma 1, this construction is a valid representation of G in $\lceil \frac{n}{2} \rceil$ dimensions. □

3.4 Graph Composition

We now look at relationships between the MEDs and various graph compositions.

3.4.1 Disjoint Union

We find the minimal embedding dimensions of the disjoint union of two graphs:

Lemma 6 *Let G_1, G_2 be graphs with disjoint vertex sets, and $D = \max\left(\left[{}^{\mu}_{M}\right]_k(G_1), \left[{}^{\mu}_{M}\right]_k(G_2)\right)$.*

If $D \geq 2$, the minimal embedding dimension of their disjoint union is $\left[\begin{smallmatrix}\mu\\M\end{smallmatrix}\right]_k(G_1 \sqcup G_2) = D$.

Proof: We will separately prove that $\left[\begin{smallmatrix}\mu\\M\end{smallmatrix}\right]_k(G_1 \sqcup G_2) \leq D$ and that $\left[\begin{smallmatrix}\mu\\M\end{smallmatrix}\right]_k(G_1 \sqcup G_2) \geq D$.

$\left[\begin{smallmatrix}\mu\\M\end{smallmatrix}\right]_k(G_1 \sqcup G_2) \leq D$

By Lemma 4, representations exist for each of G_1 and G_2 in D dimensions. By placing both of these representations in the same D-space in such a way that that they are non-overlapping in at least 2 dimensions, i.e., diagonally, we ensure that there exists no visibility lines between any vertex in G_1 and any vertex in G_2. Thus, this is a valid representation of $G_1 \sqcup G_2$ in D-space, as desired.

$\left[\begin{smallmatrix}\mu\\M\end{smallmatrix}\right]_k(G_1 \sqcup G_2) \geq D$

It suffices to show that $\left[\begin{smallmatrix}\mu\\M\end{smallmatrix}\right]_k(G_1) \leq \left[\begin{smallmatrix}\mu\\M\end{smallmatrix}\right]_k(G_1 \sqcup G_2)$, as this would by symmetry imply that $\left[\begin{smallmatrix}\mu\\M\end{smallmatrix}\right]_k(G_2) \leq \left[\begin{smallmatrix}\mu\\M\end{smallmatrix}\right]_k(G_1 \sqcup G_2)$, and these give $\left[\begin{smallmatrix}\mu\\M\end{smallmatrix}\right]_k(G_1 \sqcup G_2) \geq D$.

Take a representation of $G_1 \sqcup G_2$ in $\left[\begin{smallmatrix}\mu\\M\end{smallmatrix}\right]_k(G_1 \sqcup G_2)$ dimensions. By removing all vertices of G_2, we are not creating any new edges (unobstructing potential visibility lines) in G_1 as by definition no visibility line exists between two rectangles representing vertices in G_1 and G_2, respectively. This means that $\left[\begin{smallmatrix}\mu\\M\end{smallmatrix}\right]_k(G_1) \leq \left[\begin{smallmatrix}\mu\\M\end{smallmatrix}\right]_k(G_1 \sqcup G_2)$.

Thus, $\left[\begin{smallmatrix}\mu\\M\end{smallmatrix}\right]_k(G_1 \sqcup G_2) = D$, as desired. □

Because graphs with MED ≤ 1 must be connected, it follows that

Corollary 6 *Given graphs G_1 and G_2, the minimal embedding dimension of their disjoint union is*

$$\begin{bmatrix}\mu\\M\end{bmatrix}_k(G_1 \sqcup G_2) = \max\left(2, \begin{bmatrix}\mu\\M\end{bmatrix}_k(G_1), \begin{bmatrix}\mu\\M\end{bmatrix}_k(G_2)\right).$$

By repeatedly applying Corollary 6, we obtain

Corollary 7 *Given two or more graphs $G_1, \ldots, G_m$,*

$$\begin{bmatrix}\mu\\M\end{bmatrix}_k(G_1 \sqcup \ldots \sqcup G_m) = \max\left(2, \begin{bmatrix}\mu\\M\end{bmatrix}_k(G_1), \ldots, \begin{bmatrix}\mu\\M\end{bmatrix}_k(G_m)\right).$$

From Theorem 1 and Corollary 7, we obtain:

Corollary 8 *Let $m \leq n$ be the size of the largest connected component of a graph G on n vertices. Then,*

$$\mu_k(G) \leq \max(2, m).$$

From Theorem 2 and Corollary 7, we obtain:

Corollary 9 *For a graph G on n vertices, where the largest connected component has $m \leq n$ vertices,*

$$M_k(G) \leq \max\left(2, \left\lceil \frac{m}{2} \right\rceil\right).$$

3.4.2 Cartesian Product

We now show that MEDs are subadditive under the Cartesian product of graphs.

Theorem 3 *The minimal embedding dimension of the Cartesian product of two graphs G_1 and G_2 as (U)Rk VGs is bounded by*

$$\begin{bmatrix}\mu\\M\end{bmatrix}_k(G_1\Box G_2) \le \begin{bmatrix}\mu\\M\end{bmatrix}_k(G_1) + \begin{bmatrix}\mu\\M\end{bmatrix}_k(G_2).$$

Proof: Let $\mathcal{S}_1$ and $\mathcal{S}_2$ be orthogonal $\left[\begin{smallmatrix}\mu\\M\end{smallmatrix}\right]_k(G_1)$ and $\left[\begin{smallmatrix}\mu\\M\end{smallmatrix}\right]_k(G_2)$ dimensional spaces in $\left[\begin{smallmatrix}\mu\\M\end{smallmatrix}\right]_k(G_1) + \left[\begin{smallmatrix}\mu\\M\end{smallmatrix}\right]_k(G_2)$ dimensions. Take representations of G_1 and G_2 in $\mathcal{S}_1$ and $\mathcal{S}_2$, respectively. For any two rectangles r_1 and r_2 in these respective representations, let R_{r_1,r_2} be the rectangle in $\left[\begin{smallmatrix}\mu\\M\end{smallmatrix}\right]_k(G_1) + \left[\begin{smallmatrix}\mu\\M\end{smallmatrix}\right]_k(G_2)$ dimensions of which r_1 and r_2 are projections. Note that there is an immediate bijection between $\{R_{r_1,r_2} \mid r_1 \in \mathcal{S}_1, r_2 \in \mathcal{S}_2\}$ and vertices in $G_1\Box G_2$, namely, for any R_{r_1,r_2}, take the vertex in $G_1\Box G_2$ formed by the vertices corresponding to r_1 and r_2, respectively.

If R_{s_1,s_2} and R_{t_1,t_2} overlap then s_1 and t_1 overlap and s_2 and t_2 overlap, which is not possible unless $(s_1, s_2) = (t_1, t_2)$.

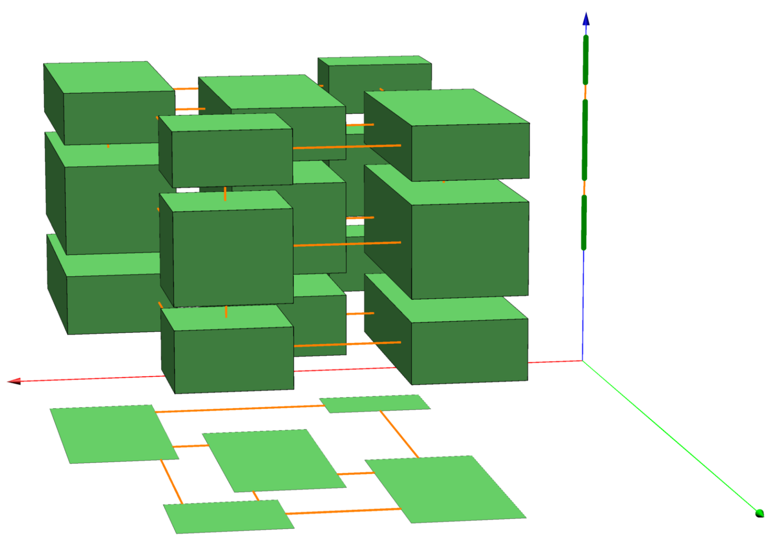

Figure 5: The Cartesian product of two graphs,
represented as the Cartesian product of their representations

Take the (U)RkVG of these rectangles. Given two adjacent rectangles, assume WLOG that the visibility line between these two rectangles is parallel to S_2. Then, in S_1 the projection of these two rectangles as well as any of the $\le k$ rectangles that obstruct the visibility line overlap, and thus are the same projected rectangle. The projection of these two rectangles onto S_2 are adjacent, obstructed by the projections of the same $\le k$ other rectangles. Conversely, if two rectangles in the projection onto S_1 are the same and in S_2 are adjacent or vice versa, the rectangles are adjacent. Therefore, this is a valid representation of $G_1 \square G_2$ in $\left[\begin{smallmatrix}\mu\\M\end{smallmatrix}\right]_k(G_1) + \left[\begin{smallmatrix}\mu\\M\end{smallmatrix}\right]_k(G_2)$ dimensions, as desired. □

By repeatedly applying Theorem 3 on multiple graphs, we obtain:

Corollary 10 *The minimal embedding dimension of the Cartesian product of multiple graphs $G_1, \ldots, G_m$ as a (U)RkVG is bounded by*

$$\begin{bmatrix}\mu\\M\end{bmatrix}_k (G_1 \square \cdots \square G_m) \leq \begin{bmatrix}\mu\\M\end{bmatrix}_k (G_1) + \cdots + \begin{bmatrix}\mu\\M\end{bmatrix}_k (G_m).$$

3.4.3 Rooted Product

We turn our attention to the rooted product, introduced in [11] by Godsil and McKay.

Definition 11 *Let G be a graph on n vertices, and let $\mathcal{H}$ be a sequence of n rooted graphs $H_1 \ldots H_n$. The rooted product of G by $\mathcal{H}$, denoted $G(\mathcal{H})$, is the (unrooted) graph obtained by identifying the root of H_i with the i^{th} vertex of G for all $i \in [n]$.*

Definition 12 *Given a representation of a graph as an RkVG, and an open half-space $\mathcal{S}$ with an axis-parallel $(d-1)$-dimensional hyperplane boundary, the expansion of the representation by a distance L is formed by moving all the hyperrectangles' corners in $\mathcal{S}$ by a distance L orthogonally away from the hyperplane.*

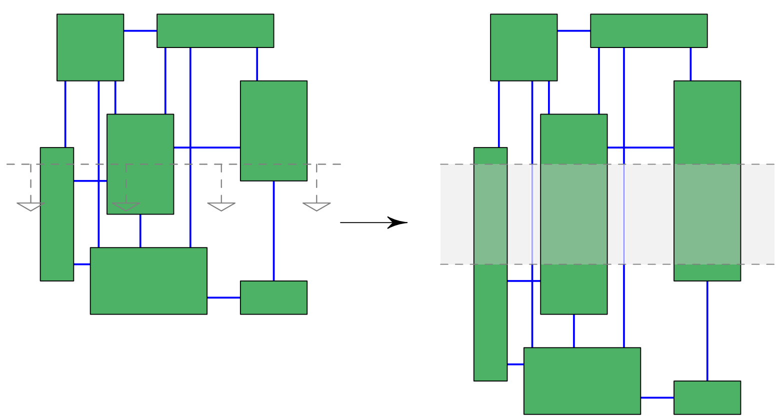

Figure 6: An expansion of an RVG representation

The expansion of a representation of a graph is another representation of the same graph, as all relationships are preserved.

Definition 13 *Given a representation of a graph G as an RkVG and a vertex $v \in G$ with rectangle R, the inflation of the representation at v by distance L is formed by expanding it on each half-space not containing R with boundary containing a face of R.*

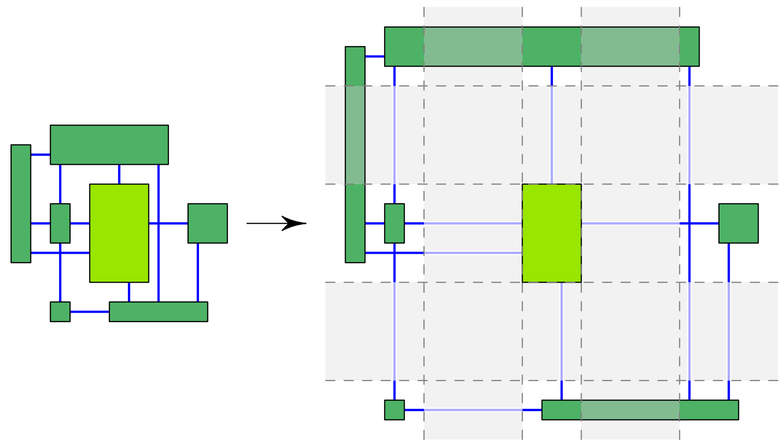

Figure 7: An inflation of a RVG representation

Theorem 4 *The minimal embedding dimension of the rooted product as a RkVG is bounded by*

$$\max\left(M_k(G), \max_{H\in\mathcal{H}}(M_k(H))\right) \leq M_k(G(\mathcal{H})) \leq M_k(G) + \max_{H\in\mathcal{H}}(M_k(H)).$$

Proof: For the lower bound, to establish that $M_k(G) \leq M_k(G(\mathcal{H}))$, we take any representation of $G(\mathcal{H})$ in $M_k(G(\mathcal{H}))$ dimensions. By definition, there's a naturally induced copy of G in $G(\mathcal{H})$.

Assume for the sake of contradiction that removing all rectangles representing vertices not in the induced G from the representation of $G(\mathcal{H})$ adds a visibility line segment between rectangles representing non-adjacent vertices in G. Let $v_1, \ldots, v_m$ be the vertices corresponding to the rectangles on this line segment, where v_1 and v_m are in the induced copy of G.

There is an i such that v_i is in the induced copy of G but v_{i+1} is not, as otherwise the path would not leave the induced G. Let $H \in \mathcal{H}$ be the rooted graph corresponding to v_i. v_{i+1} is in the induced copy of H, so removing v_i from $G(\mathcal{H})$ disconnects v_{i+1} from the induced G and in particular from v_m. However, $v_{i+1}, \ldots, v_m$ is a path connecting v_{i+1} to v_m that does not pass through v_i, a contradiction.

Then, to establish that $\forall H_i \in \mathcal{H},\ M_k(H_i) \leq M_k(G(\mathcal{H}))$, note that the natural copy of H_i in $G(\mathcal{H})$ is only connected to the rest of $G(\mathcal{H})$ at one vertex, so $G(\mathcal{H})$ can be expressed as $H_i(\mathcal{G}_i)$ for some sequence of rooted graphs $\mathcal{G}_i$. Thus by the above, $M_k(H_i) \leq M_k(H_i(\mathcal{G}_i)) = M_k(G(\mathcal{H}))$, as desired.

For the upper bound, by Lemma 4, we can take representations of $H_1, \ldots, H_n$ in $d = \max_{H\in\mathcal{H}}(M_k(H))$ dimensions. Rescale and translate all representations such that the rectangles corresponding to the roots are all unit size and centered at the origin. Let L be such that all representations fit in a $\underbrace{(2L+1)\times\cdots\times(2L+1)}_{d}$ bounding rectangle centered at the origin, i.e., such that all rectangle faces are within L of a parallel face of the central root rectangle.

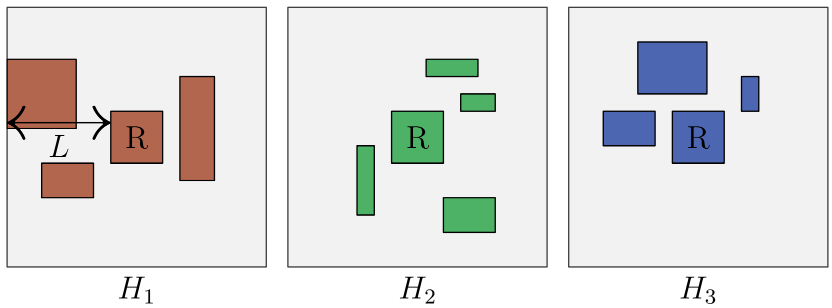

Figure 8: Representations of three rooted graphs H_1, H_2, H_3, with L depicted, and with roots indicated by "R"s

As described in Definition 13 and illustrated in Figure 9, now inflate the representation of H_i around the root vertex by $(i-1)\times L$ for all $i \in [n]$ so that no rectangles besides the root vertex overlap between the representations.

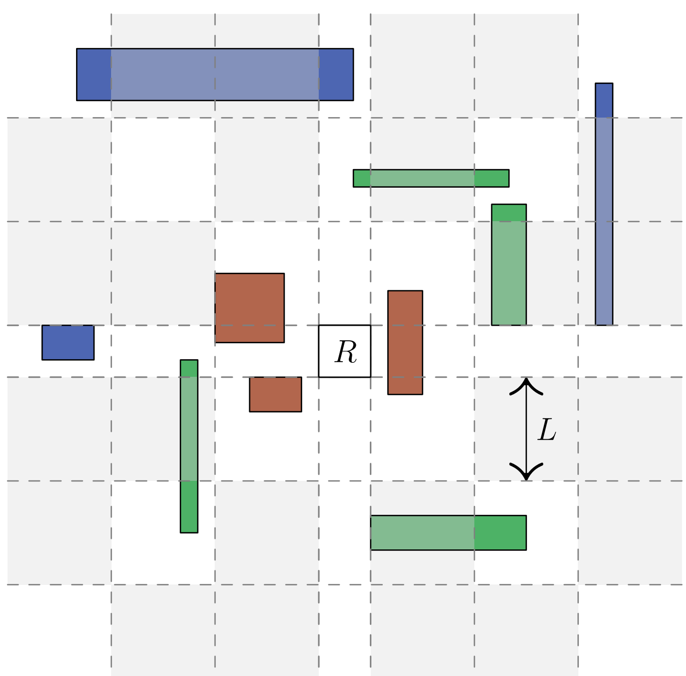

Figure 9: Respective inflations of H_1, H_2, H_3 at their roots by $0, L$, and $2L$, superimposed at their roots.

Finally, take a representation of G in $M_k(G)$ dimensions, as shown in Figure 10.

Figure 10: A representation of a graph G with vertices corresponding to $\mathcal{H}$

For all $i \in [n]$ and for $v \in H_i$, take the rectangle in $M_k(G) + \max_{H \in \mathcal{H}}(M_k(H))$ dimensions whose projection in the first $M_k(G)$ dimensions is the representation of the i^{th} vertex of G, and whose projection in the last $\max_{H \in \mathcal{H}}(M_k(H))$ dimensions is the representation of $v \in H_i$.

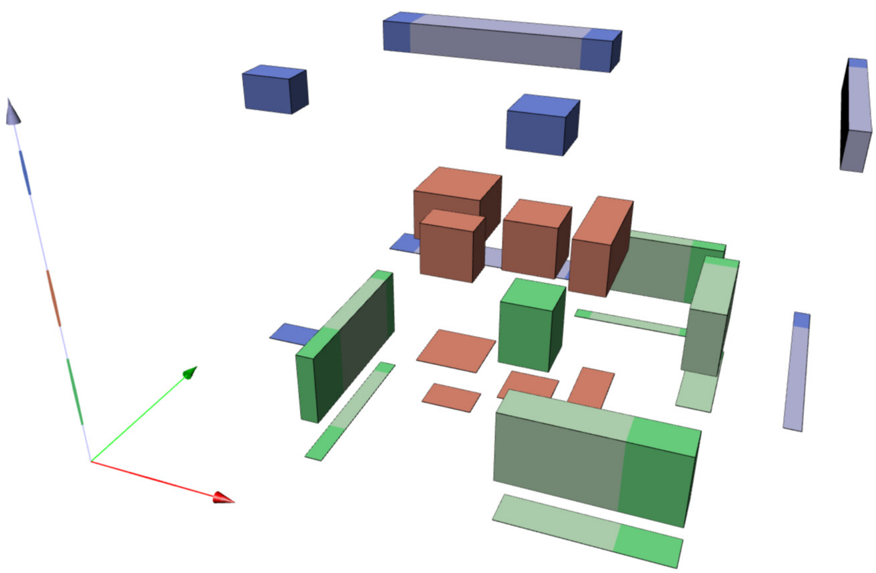

Figure 11: A representation of the rooted product $G(\mathcal{H})$

We claim that these rectangles form a representation of $G(\mathcal{H})$. Since all the roots of $\mathcal{H}$ have the same projection in the last $\max_{H \in \mathcal{H}}(M_k(H))$ dimensions, their visibilities are those of their projections in the first $M_k(G)$ dimensions; namely, the edges of G. Any rectangle that does not correspond to a root does not overlap with rectangles in the last $\max_{H \in \mathcal{H}}(M_k(H))$ by construction, and thus only sees those rectangles with which it overlaps in the first $M_k(G)$ dimensions and sees in the last $\max_{H \in \mathcal{H}}(M_k(H))$, as desired. $\square$

3.4.4 Corona Product

We now look at the corona product, introduced by Frucht and Harary [10].

Definition 14 *The corona product of two graphs G and H, denoted $G \odot H$, is obtained by taking* *one copy of G and $n = |V(G)|$ copies of H, and by connecting the i^{th} vertex of G to each vertex of the i^{th} copy of H for all $i \in [n]$.*

Remark 15 *For $\mathcal{H} = (H')_{i\in[n]}$ (i.e., H' repeated n times), where H' is H with an added universal root vertex (i.e., a root vertex connected to every other vertex of H), $G \odot H = G(\mathcal{H})$.*

Theorem 5 *The minimal embedding dimension of the corona product of two graphs G and H as a RkVG is bounded by*

$$M_k(G) \leq M_k(G \odot H) \leq \max(M_k(G), M_k(H)) + 1.$$

Proof:

By Remark 15 and Theorem 4, we have

$$M_k(G) \leq M_k(G(\mathcal{H})) = M_k(G \odot H),$$

where $\mathcal{H}$ is as in Remark 15.

We now show $M_k(G \odot H) \leq \max(M_k(G), M_k(H)) + 1$ by finding a $\max(M_k(G), M_k(H)) + 1$-dimensional representation of $G \odot H$.

By Lemma 4, we can take representations of G and H in $\max(M_k(G), M_k(H))$ dimensions.

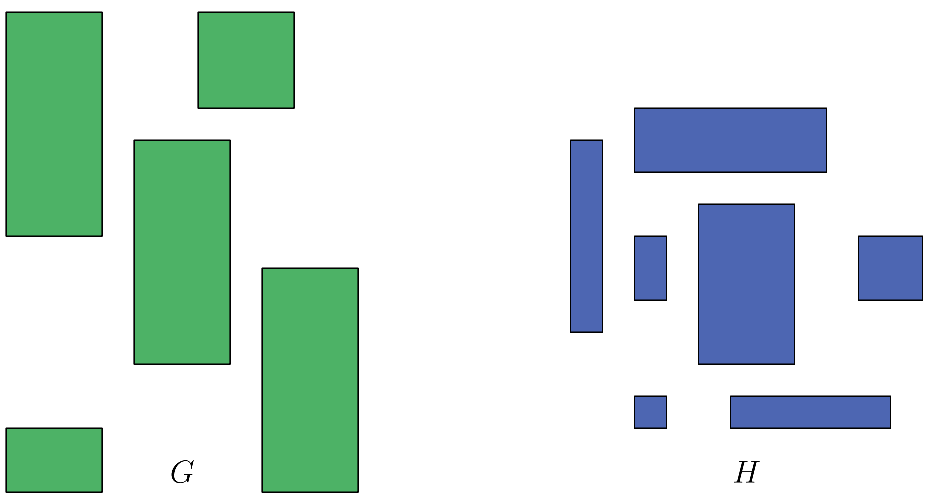

Figure 12: Representations of two graphs G and H

Shrink the representation of H until it is smaller than all of the rectangles in the representation of G, and thicken both representations orthogonally by one unit into the d^{th} dimension, where $d = \max(M_k(G), M_k(H)) + 1$. Take n copies of H's representation, corresponding to the n rectangles in the representation of G, and place them at different heights above the latter in the d^{th} dimension, such that each copy is exactly above its corresponding rectangle and no copies can see each other.

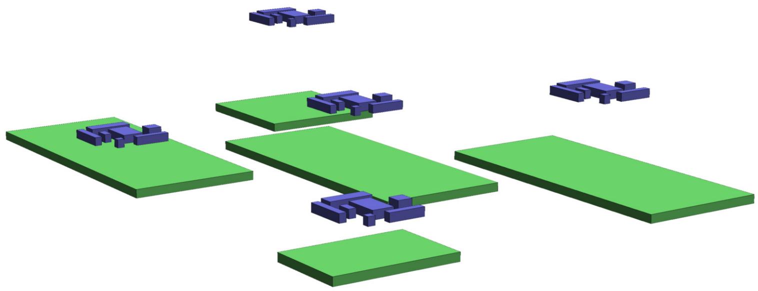

Figure 13: A representation of the corona product $G \odot H$

As desired, any rectangle in G's representation now has a visibility line to every rectangle in exactly one copy of H's representation, with no visibility lines to or between other copies of H; moreover, visibilities are maintained within each of the original representations. $\square$

4 Complete Graphs

We now construct arrangements of rectangles where every rectangle can see every other rectangle, thus giving the complete graph.

4.1 MEDs as URkVGs

Theorem 6 *For all $k \geq 0$, the minimal embedding dimension of the complete graph on m vertices, K_m, as a (U)RkVG is bounded by*

$$\begin{bmatrix}\mu\\M\end{bmatrix}_k (K_m) \leq \left\lceil \frac{3}{5}m \right\rceil.$$

Proof: Because we can remove a rectangle from any representation of K_{m+1} to get one of K_m, we need only prove that for $d \in \mathbb{N}$,

$$\begin{bmatrix}\mu\\M\end{bmatrix}_k (K_{5d}) \leq 3d,$$

$$\begin{bmatrix}\mu\\M\end{bmatrix}_k (K_{5d+1}) \leq 3d+1$$

and

$$\begin{bmatrix}\mu\\M\end{bmatrix}_k (K_{5d+3}) \leq 3d+2.$$

We imitate the proof of Theorem 2. For the sake of avoiding repetition, we simply show how the green and orange (this time, unit) rectangles are arranged. In the rest of the proof, the differences are:

- replace every occurrence of 4 with 5,
- let $\mathcal{S}_1, \ldots, \mathcal{S}_{\lceil \frac{n-3}{5} \rceil}$ be orthogonal three-dimensional spaces,
- for $n \equiv 1, 2, 3 \pmod 5$, let $\mathcal{S}_{\lceil \frac{n}{5} \rceil}$ be an additional $\left(\lceil \frac{3}{5} n \rceil - 3\lceil \frac{n-3}{5} \rceil\right)$-dimensional space.

Figures 14 through 19 depict all the relevant configurations. Note that since we are constructing K_m, we only need projections where all green rectangles are visible to each other and orange rectangles either intersect either all or none of the green rectangles.

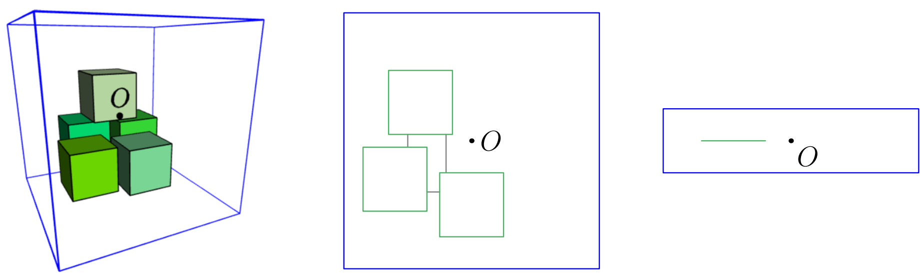

Figure 14: All possible projections of S_m into $\mathcal{S}_m$

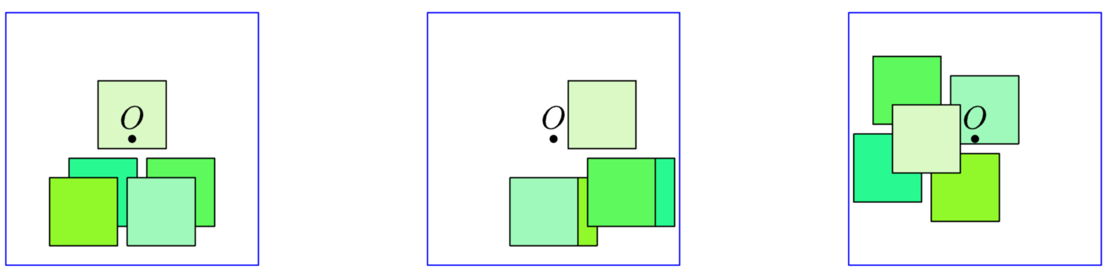

Figure 15: Side views of the 5 rectangles depicted on the left of Figure 14

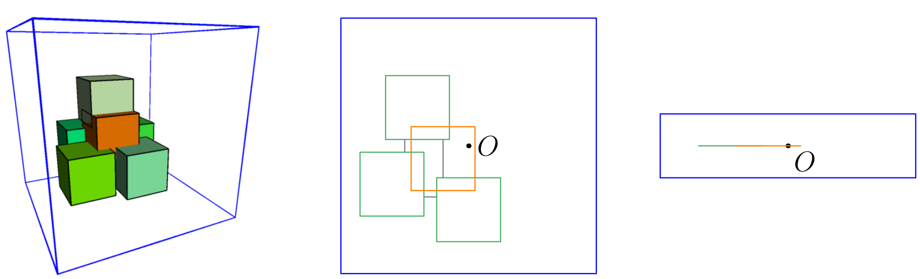

Figure 16: All possible projections of an additional vertex v_i,
overlapping a central point O and all rectangles in S_m

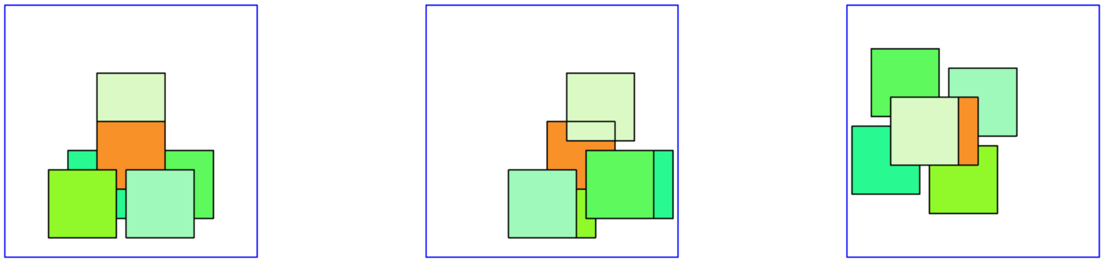

Figure 17: Side views of the 6 rectangles depicted on the left of Figure 16

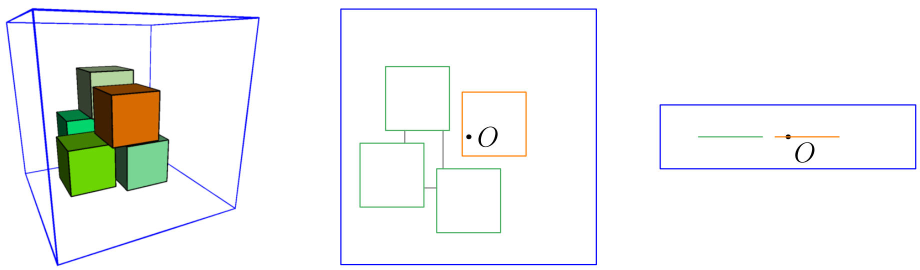

Figure 18: All possible projections of an additional vertex v_i, overlapping a central point O and adjacent to all rectangles in S_m

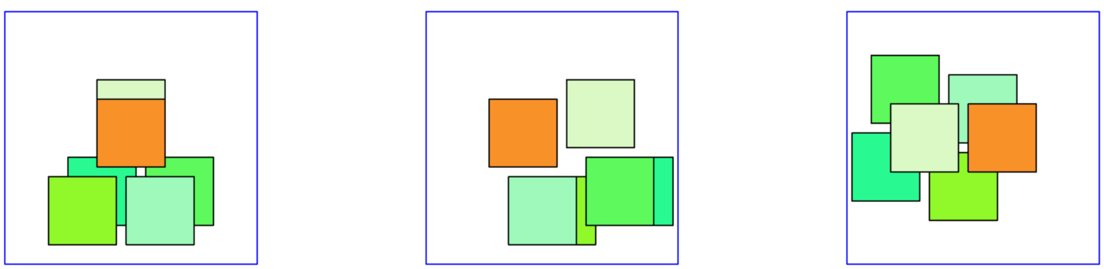

Figure 19: Side views of the 6 rectangles depicted on the left of Figure 18

□

4.2 MEDs as RkVGs

Theorem 7 *The complete graph on* $2(d-1)(k+1)+22(\lfloor k/2\rfloor+1)$ *vertices,* $K_{2(d-1)(k+1)+22(\lfloor k/2\rfloor+1)}$*, is a d-dimensional RkVG for* $d \geq 3$.

Proof: Figure 20, adapted from Figure 3 of [9], shows 22 rectangle projections.

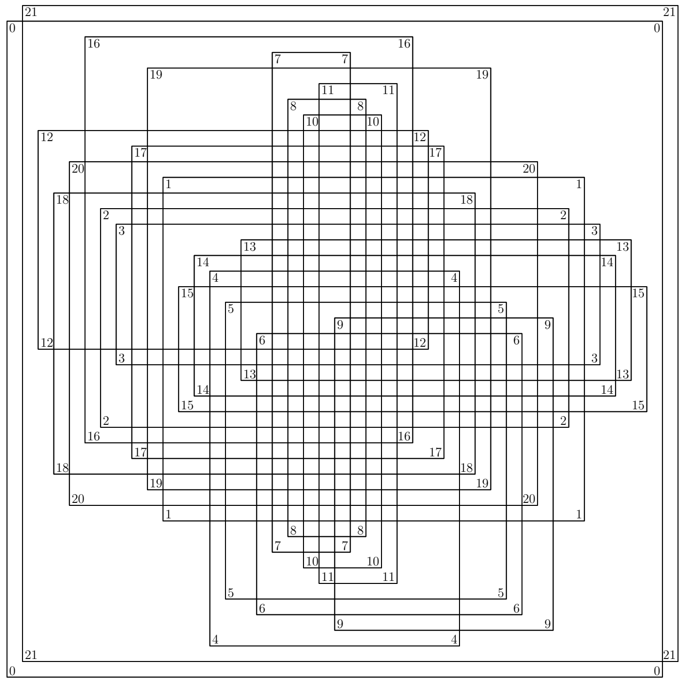

Figure 20: Projections of 22 rectangles, adapted from Figure 3 of [9]

If for i from 0 to 21 we place a corresponding rectangle with thickness $\delta \in (0, 1)$ in 3-dimensional space at height $z = i$ above this plane, such that its projection to the plane is the rectangle labeled i, we obtain a 0-visibility representation of K_{22}, as in Figure 21.

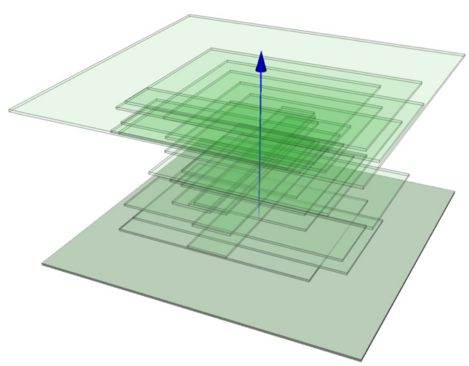

Figure 21: The 22 rectangles stacked on top of each other seen from above the top right relative to Figure 20

Replace the i^{th} rectangle with $\left\lfloor \frac{k}{2} \right\rfloor + 1$ duplicates of it with thickness (height) $\frac{\delta}{\left\lfloor \frac{k}{2} \right\rfloor + 1}$ at heights $i + \frac{j}{\left\lfloor \frac{k}{2} \right\rfloor + 1}$ for $j \in \left\{0, \dots, , \left\lfloor \frac{k}{2} \right\rfloor + 1\right\}$. We now have a k-visibility representation of $K_{22(\lfloor k/2 \rfloor + 1)}$, where all rectangles are visible from any of the four sides, as seen in Figure 22. (A visibility line between two rectangles with different projections passes through at most $\lfloor k/2 \rfloor$ other rectangles with the same projection as each of the former and the latter rectangle).

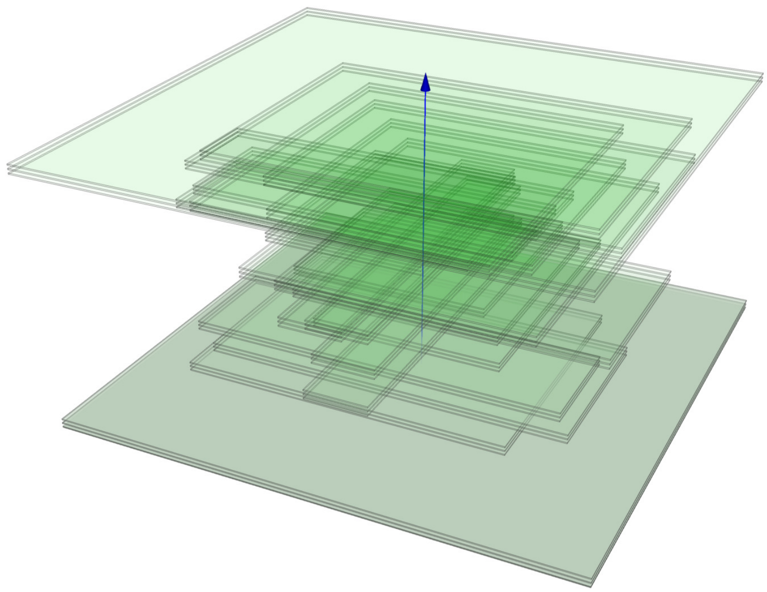

Figure 22: $22(\lfloor k/2 \rfloor + 1)$ rectangles stacked, representing $K_{22(\lfloor k/2 \rfloor+1)}$ with $k = 4$

We then thicken this representation by one unit into each of the remaining $(d-3)$ dimensions.

Finally, in each dimension except the 3^{rd} (along whose axis we stack our $22(\lfloor k/2 \rfloor + 1)$ rectangles), we add $k+1$ hyperrectangles in both directions from the center, at increasing distances and with increasingly large hyperfaces facing the center, such that each hyperrectangle has k-visibility to every other rectangle; i.e., such that the added rectangles in each dimension surround the entire representation up to that point. (See Figure 23 for an example).

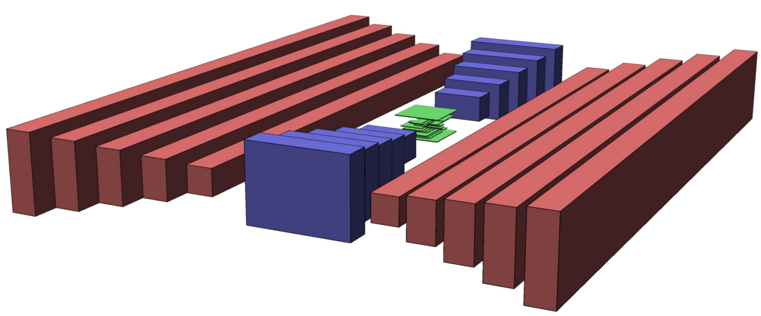

Figure 23: Representation of $K_{2(d-1)(k+1)+22(\lfloor k/2 \rfloor+1)}$ with $k = 4, \ d = 3$

Along the i^{th} axis, there are $2(k+1)$ rectangles surrounding the center and the rectangles corresponding to prior axes, for a total of $2(d-1)(k+1)$ rectangles surrounding the center. As all rectangles, big and small, are k-visible to each other, we have a representation of $K_{2(d-1)(k+1)+22(\lfloor k/2 \rfloor+1)}$. □

Because we can remove a rectangle from any representation of K_{m+1} to get one of K_m, $M_k(K_m)$ is non-decreasing, so this gives a bound for the minimal embedding dimension of the complete graph:

Corollary 16 *The minimal embedding dimension of the complete graph on m vertices K_m as a RkVG is bounded by*

$$M_k(K_m) \leq \max\left(3, \left\lceil \frac{m - 22(\lfloor k/2 \rfloor + 1)}{2(k+1)} \right\rceil + 1\right).$$

4.3 Growth of the MEDs

Lemma 7 *For some fixed k, let*

$$c_2 = 4k + 5$$
$$c_i = \binom{c_{i-1}}{2} + 1 \quad | \quad i \geq 3.$$

Then, $K_{c_{2d-2}}$ cannot be represented in d dimensions with all visibility lines parallel.

To prove this lemma, we apply a technique used by Fekete et al. in Theorem 4 of [8].

Proof: We use induction on d.

Base case: $d = 2$

Assume for the sake of contradiction that such a representation exists, and assume WLOG that all visibility lines are vertical. Flatten all rectangles so that they are horizontal line segments. We now have a bar k-visibility representation, as defined in the introduction, of $K_{c_2} = K_{4k+5}$. Hartke et al. have shown, however, that this is impossible [13].

Inductive step: $d - 1 \Rightarrow d$

We assume for the sake of contradiction that such a representation exists, and assume WLOG that all visibility lines are parallel to the first axis. As all rectangles overlap in every other dimension, there is then a line ℓ parallel to the d^{th} axis that passes through all c_{2d-2} rectangles.

Translate the coordinate system such that the origin lies on ℓ. Each rectangle has two faces orthogonal to the d^{th} axis, one on each side of ℓ. Let F_n and F'_n be the coordinates along the d^{th} axis of the corresponding faces for the n^{th} rectangle, where F_n is negative and F'_n is positive.

Chung showed that every sequence of $\binom{a}{2}+1$ numbers has a subsequence of length a with one local maximum [5]. Thus there exists a subsequence $\left(r_1,\ldots,r_{c_{2d-3}}\right)$ among our c_{2d-2} rectangles such that the sequence $\left(-F_{r_1},\ldots,-F_{r_{c_{2d-3}}}\right)$, has one local maximum.

Likewise, among these c_{2d-3} rectangles there is a sub-subsequence $\left(s_1,\ldots,s_{c_{2d-4}}\right)$ such that the distance from ℓ to the second face of each orthogonal rectangle, $\left(F'_{s_1},\ldots,F'_{s_{c_{2d-4}}}\right)$, form another unimaximal progression.

Note that in the d^{th} dimension, if rectangles s_i and s_k overlap for $i<j<k$, rectangle s_j contains their overlap. Thus, the visibility lines between these rectangles are those of their projections into the first $(d-1)$ dimensions. By the inductive hypothesis, these $c_{2d-4}=c_{2(d-1)-2}$ rectangles cannot form a complete graph, as desired.

□

Theorem 8 *The range of $\left[{\mu \atop M}\right]_k(K_m)$ over m for fixed k is the set of nonnegative integers, $\mathbb{Z}_{\geq 0}$.*

Proof: Let $r=R\left(\underbrace{c_{2d-2},c_{2d-2},\ldots,c_{2d-2}}_{d}\right)$ (adopting the notation from Lemma 7), where R denotes the multicolor Ramsey number function. Assume for the sake of contradiction that K_r is representable in d dimensions. Color each edge of K_r by the axis parallel to its visibility line. As this is a coloring with d colors of the edges of K_r, there is a monochromatic $K_{c_{2d-2}}$, contradicting Lemma 7. Thus, K_r is not representable in d dimensions.

Thus, no finite number of dimensions can represent K_m for all $m\in\mathbb{N}$, so $\left[{\mu \atop M}\right]_k(K_m)$ takes on arbitrarily large values. Since $\left[{\mu \atop M}\right]_k(K_1)=0$, it then suffices to show that $\left[{\mu \atop M}\right]_k(K_{m+1})\leq\left[{\mu \atop M}\right]_k(K_m)+1$.

Assume that we have a representation of K_m in $\left[{\mu \atop M}\right]_k(K_m)$ dimensions. Add an extra dimension, thicken all the rectangles by 1 unit in this dimension, and replace one rectangle with two copies shifted by $-\frac{2}{3}$ and $\frac{2}{3}$ into the new dimension, respectively. Then, as all visibilities are maintained and the two copies can see each other, we have a representation of K_{m+1} in $\left[{\mu \atop M}\right]_k(K_m)+1$ dimensions, as desired. □

5 Complete Multipartite Graphs

To construct complete multipartite graphs, we arrange the rectangles in a crosshatch, so to speak.

Theorem 9 *For all $k\geq 0$, the complete $(d-1)$-partite graph (which is the empty graph for $d=2$ and is $K_{m_1,\ldots,m_{d-1}}$ for $d>2$) is a d-dimensional RkVG.*

Proof: Take an $m_1\times m_2\times\cdots\times m_{d-1}$ lattice in $(d-1)$-space. For each of the $(d-1)$ axes, take all orthogonal $(d-2)$-spaces that pass through lattice points. Add a small thickness to each of these spaces in their respective orthogonal dimensions. Cut all these spaces off to get axis-orthogonal hyperrectangles surrounding the lattice points.

For example, given $d = 3, m_1 = 6, m_2 = 8$ we get the left hand side of Figure 24, and given $d = 4, m_1 = 6, m_2 = 8, m_3 = 5$ we get the configuration in Figure 25.

Note that any pair of rectangles corresponding to spaces orthogonal to the same axis do not intersect, but rectangles corresponding to different axes do.

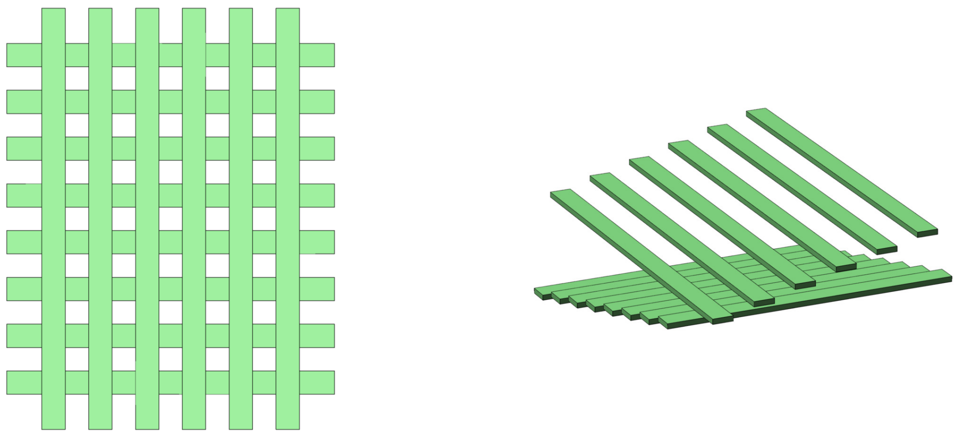

Figure 24: A representation of the 3-dimensional RVG $K_{6,8}$

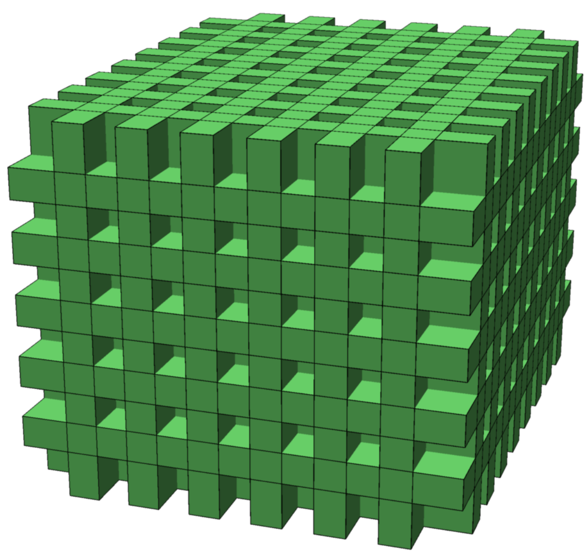

Figure 25: An overhead orthographic projection of a representation of the 4-dimensional RVG $K_{6,8,5}$

Now we extend the figure into the d^{th} dimension by adding a small thickness, and finally add a distinct height to each of them. (See the right hand side of Figure 24.)

As any two rectangles corresponding to the same axis are not k-visible to each other, but any other two rectangles are, we have a representation of $K_{m_1,\ldots,m_{d-1}}$. □

This gives a bound for the minimal embedding dimension of the complete multipartite graph:

Corollary 17 *The minimal embedding dimension of the complete r-partite graph as a RkVG is*

$$M_k(K_{m_1,\ldots,m_r}) \leq r+1$$

for $r > 1$.

6 Hypercubes

6.1 k-Visibility

Hypercubes are bipartite graphs, so by the proof of Corollary 5, 1-visibility lines need to be avoided.

Theorem 10 *For all $k \geq 0$, the minimal embedding dimension of the hypercube graph on 2^m vertices, Q_m, as a (U)RkVG is bounded by*

$$\begin{bmatrix}\mu\\ M\end{bmatrix}_k (Q_m) \leq \left\lceil \frac{2}{3}m \right\rceil.$$

Proof: Figure 26 shows a representation of Q_3 in 2 dimensions, so $M_k(Q_3) = \mu_k(Q_3) = 2$.

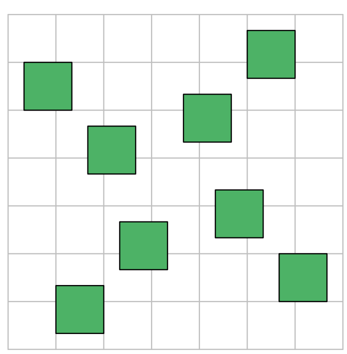

Figure 26: A (U)RkVG representation of the hypercube graph Q_3 in 2 dimensions

Since $\begin{bmatrix}\mu\\M\end{bmatrix}_k(Q_1) = 1$, by Corollary 10 we get

$$\begin{aligned}
\begin{bmatrix}\mu\\M\end{bmatrix}_k(Q_m) &= \mu_k\left(\underbrace{Q_3\square\cdots\square Q_3}_{\lfloor\frac{m}{3}\rfloor}\ \square\ \underbrace{Q_1\square\cdots\square Q_1}_{m-3\lfloor\frac{m}{3}\rfloor}\right)\\
&\le \underbrace{\mu_k(Q_3)+\cdots\mu_k(Q_3)}_{\lfloor\frac{m}{3}\rfloor} + \underbrace{\mu_k(Q_1)+\cdots+\mu_k(Q_1)}_{m-3\lfloor\frac{m}{3}\rfloor}\\
&= 2\left\lfloor\frac{m}{3}\right\rfloor + \left(m-3\left\lfloor\frac{m}{3}\right\rfloor\right)\\
&= \left\lceil\frac{2}{3}m\right\rceil
\end{aligned}$$

□

Remark 18 *In a 2-dimensional representation of a bipartite (U)Rk VG G with n vertices for $k > 0$, by Corollary 5, we can treat k as infinite, so G is the union of the interval graphs G_1 and G_2 (see Definition 4) formed by the horizontal and vertical projections, respectively. Order the vertices of G_1 by the starting points of their intervals and map each edge to its larger vertex in the ordering. If two edges were mapped to the same vertex v, the starting point of v would be contained in all intervals corresponding to vertices in the two edges, forming a triangle, a contradiction. Thus, this map is injective, so as no edge maps to the smallest vertex, there are $\le n-1$ edges in G_1, similarly $\le n-1$ in G_2, and in total, $\le 2(n-1)$ in G.*

As Q_4 is bipartite and has 16 vertices and $\frac{2^4\cdot 4}{2} > 2(16-1)$ edges, it cannot be represented in $d = 2$ dimensions. Thus, Theorem 10 is tight for $m \le 4, k > 0$.

6.2 0-Visibility

We now move on to 0-visibility, where as opposed to our previous construction, we do not have to worry about collinear rectangles.

Our 0-visibility representations of hypercubes will be arranged in grids, so to speak. For example, in the representation of Q_6 shown in Figure 27, the rectangles are organized in a $2^3 \times 2^3$ grid.

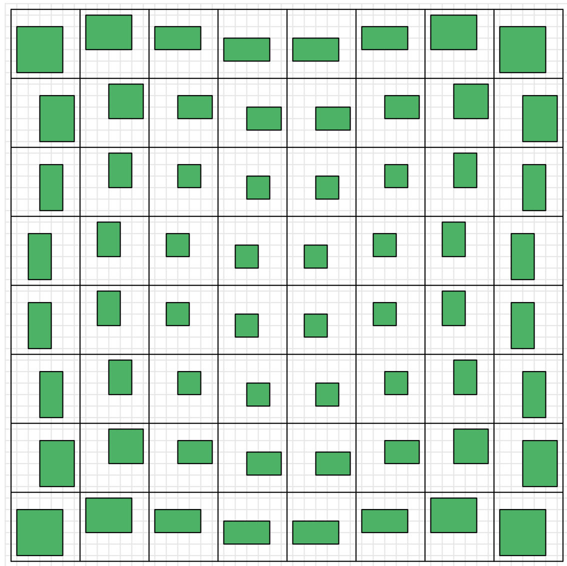

Figure 27: A 2-dimensional RVG representation of Q_6 with a grid-like structure

In order to construct representations of hypercube graphs, we will first show how to construct the columns, then show how to combine them into the full grid.

6.2.1 Gray Code

Before we proceed, we need to introduce the reflective binary Gray code, which we will simply refer to as Gray code.

Definition 19 *The Gray code is a reordering of the binary numeral system such that two successive values differ in only one bit (binary digit) [12].*

Like standard numbering systems (e.g., binary), Gray code representations of a number are implicitly padded with an infinite number of 0*'s on the left, and any number* i *is represented with a finite number of* 1*'s. The number zero is represented with only* 0*'s.*

Given the Gray code representation of a non-negative integer $i-1$*, the representation of* i *is formed by flipping the* j^{th} *digit from the right, where the rightmost digit is the* 0^{th} *digit and* 2^j *is the largest power of* 2 *that divides* i*.*

In the following discussion, we will denote by $G_{i,j}$ *digit #*j *of the Gray code representation of* i*, counting from the right such that* $G_{i,0}$ *is the least significant digit.*

Example 20 *The (four digit) Gray code representation of numbers* 0 *through* 15 *are shown in Figure 28.*

0	0000	1000	15
1	0001	1001	14
2	0011	1011	13
3	0010	1010	12
4	0110	1110	11
5	0111	1111	10
6	0101	1101	9
7	0100	1100	8

Figure 28: The Gray code representations of numbers 0 through 15

We will make use of the following properties of Gray code:

- It is the reflective binary code, where the representation of numbers $0, \ldots, (2^k - 1)$ are repeated in reverse order for numbers $2^k, \ldots, (2^{k+1} - 1)$, except that the k^{th} digit is 1 instead of 0 (with digit #0 being the rightmost). In other words, for all $i < 2^n$ and $j < n$,

$$G_{i,j} = G_{2^{n+1}-1-i,j}.$$

- The map from nonnegative integers to their Gray code representations is a bijection.

- The parity of a number is the parity of the number of 1's in its Gray code representation. As a consequence, any two numbers whose Gray code representations differ in exactly one bit have different parities.

- $G_{i,0} = 0$ iff $i \equiv \{0, 3\} \pmod 4$.

- For all non-negative integers i, j, $G_{2i,j+1} = G_{2i+1,j+1} = G_{i,j}$.

6.2.2 MEDs as URVGs

First we construct the unit rectangle columns:

Lemma 8 *The d-dimensional RVG formed by cubes of side length 2 centered at points of the form*

$$((d+2)i,\ G_{i,0},\ G_{i,1},\ G_{i,2},\ \ldots,\ G_{i,d-2})$$

for $0 \le i < 2^d$ is Q_d.

This construction for $d = 3$ is shown in Figure 29.

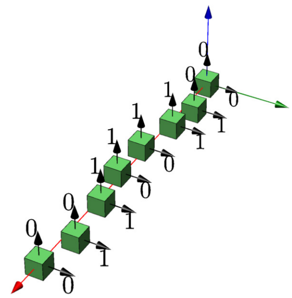

Figure 29: A 3-dimensional URVG representation of Q_3
(the x, y, and z axes are colored red, green, and blue, respectively,
and the offset of the center of each cube from the x axis is indicated)

Note that the constant $(d+2)$ does not effect the validity of this lemma as anything sufficiently large for the cubes to be disjoint (more than 2) would work. Choosing $(d+2)$ becomes useful later on.

Proof: We use induction on d.

Base case: $d = 1$

This case trivially holds, as a two segments form a valid representation of Q_1 in 1-dimensional space.

Inductive step: $d \Rightarrow d+1$

For the remainder of this proof, we refer to the j^{th} coordinate of the center of a rectangle as its j^{th} coordinate.

Q_{d+1} is formed by two induced copies of Q_d with an edge between every pair of corresponding vertices. In light of this, we split the representation into two sets: $i < 2^d$ and $i \geq 2^d$, and biject each of these sets to the representation of Q_d.

For $0 \leq i < 2^d$, map the i^{th} rectangle of the domain to the i^{th} rectangle of the range, and for $2^d \leq i < 2^{d+1}$, map the i^{th} rectangle of the domain to the $(2^{d+1}-1-i)^{\text{th}}$ rectangle of the range. We first show that these maps are graph isomorphisms. All visibility lines are parallel to the 1^{st} axis and the order of the first coordinates is preserved (or reversed). In addition, since $G_{i,j} = G_{2^{d+1}-1-i,j}$ for $j < d$, this map preserves the 2^{nd} to $(d-1)^{\text{th}}$ coordinates.

For $i < 2^d$, $G_{i,d-1}$ is 0 iff $i < 2^{d-1}$, and for $i \geq 2^d$, $G_{i,d-1}$ is 0 iff $i \geq 2^{d+1} - 2^{d-1}$. Thus after introducing the d^{th} coordinate, within each group, all rectangles between two others with d^{th} coordinate 0 have d^{th} coordinate 0, and the same is true of d^{th} coordinate 1. Thus, any obstructing rectangles in the image are obstructing rectangles in the domain, as desired.

It remains to show that the correct visibility lines are drawn between the two groups; more specifically, that two rectangles in different groups share an edge if and only if they are mapped to the same rectangle. Since the two rectangles mapping to rectangle i of the image are the only two rectangles intersecting the (parameterized) line

$$(t,\ 2G_{i,0}-1/2,\ 2G_{i,1}-1/2,\ 2G_{i,2}-1/2,\ \ldots,\ 2G_{i,d-2}-1/2),$$

it indeed holds that there is an edge between them. In addition, for any two rectangles in different groups whose images are not the same, say rectangles $i_0 < 2^d \leq i_1$ with $i_0+i_1 \neq 2^{d+1}-1$, if $i_0+i_1 < 2^{d+1}-1$ we would have $i_0 < 2^{d+1}-1-i_1 < 2^d \leq i_1$, so rectangle $N = (2^{d+1}-1-i_1)$ is between the two rectangles, whereas if $i_0+i_1 > 2^{d+1}-1$ we would have $i_1 > 2^{d+1}-1-i_0 \geq 2^d > i_0$, so rectangle $N = (2^{d+1}-1-i_0)$ is between the two rectangles. In the former and latter cases rectangle N has the same coordinates (besides the first) as rectangles i_1 and i_0, respectively, and thus blocks any possible visibility line between i_0 and i_1.

Combining the above, we find that this is indeed a valid representation.

□

Now we arrange such columns into a grid:

Theorem 11 *The hypercube graph on 2^{d^2} vertices, Q_{d^2}, is a d-dimensional URVG.*

Proof: For any tuple $(i_1, i_2, \ldots, i_d)$ with $0 \leq i_j \leq 2^d-1$ for all $j \in \{1,\ldots,d\}$, take cubes of side length 2, where each cube is centered at the sum of the vectors

$$\begin{gathered}\big((d+2)i_1,\ G_{i_1,0},\ G_{i_1,1},\ \ldots,\ G_{i_1,d-3},\ G_{i_1,d-2}\big),\\ \big(G_{i_2,d-2},\ (d+2)i_2,\ G_{i_2,0},\ \ldots,\ G_{i_2,d-4},\ G_{i_2,d-3}\big),\\ \big(G_{i_3,d-3},\ G_{i_3,d-2},\ (d+2)i_3,\ \ldots,\ G_{i_3,d-5},\ G_{i_3,d-4}\big),\\ \vdots\\ \big(G_{i_{d-1},1},\ G_{i_{d-1},2},\ G_{i_{d-1},3},\ \ldots,\ (d+2)i_{d-1},\ G_{i_{d-1},0}\big),\\ \big(G_{i_d,0},\ G_{i_d,1},\ G_{i_d,2},\ \ldots,\ G_{i_d,d-2},\ (d+2)i_d\big).\end{gathered}$$

There are 2^{d^2} such tuples.

If we fix all but one of $i_1, i_2, \ldots, i_d$, the corresponding rectangles form a Q_d by Lemma 8.

In addition, if WLOG $i_1 > i'_1$ the smaller 1^{st} coordinate of the rectangle corresponding to $(i_1, i_2, \ldots, i_d)$ is more than the larger 1^{st} coordinate of the rectangle corresponding to $(i'_1, i'_2, \ldots, i'_d)$, because

$$\begin{aligned}\left((d+2)i_1 + \sum_{j=2}^{d} G_{i_j,d-j}\right) - 1 &\geq \left(((d+2)i'_1 + (d+2)(i_1-i'_1)) + \sum_{j=2}^{d} 0\right) - 1\\ &\geq ((d+2)i'_1 + (d+2)\cdot 1 + 0) - 1\end{aligned}$$

$$> \left((d+2)i'_1 + \sum_{j=2}^{d} 1 \right) + 1$$
$$\geq \left((d+2)i'_1 + \sum_{j=2}^{d} G_{i'_j, d-j} \right) + 1.$$

Thus, these are the only visibility lines, and no rectangles intersect.

If we shrink this construction by a factor of 2, all the cubes become unit rectangles. We see this construction applied to $d = 0$ through $d = 3$ in figures 30, 31, 32, and 33, respectively. □

Figure 30: A 0-dimensional URVG representation of Q_0

Figure 31: A 1-dimensional URVG representation of Q_1

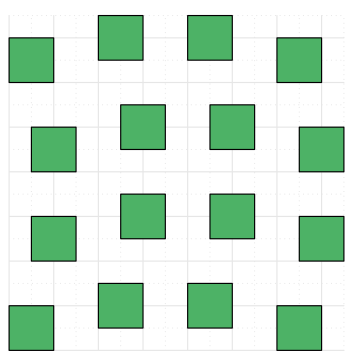

Figure 32: A 2-dimensional URVG representation of Q_4

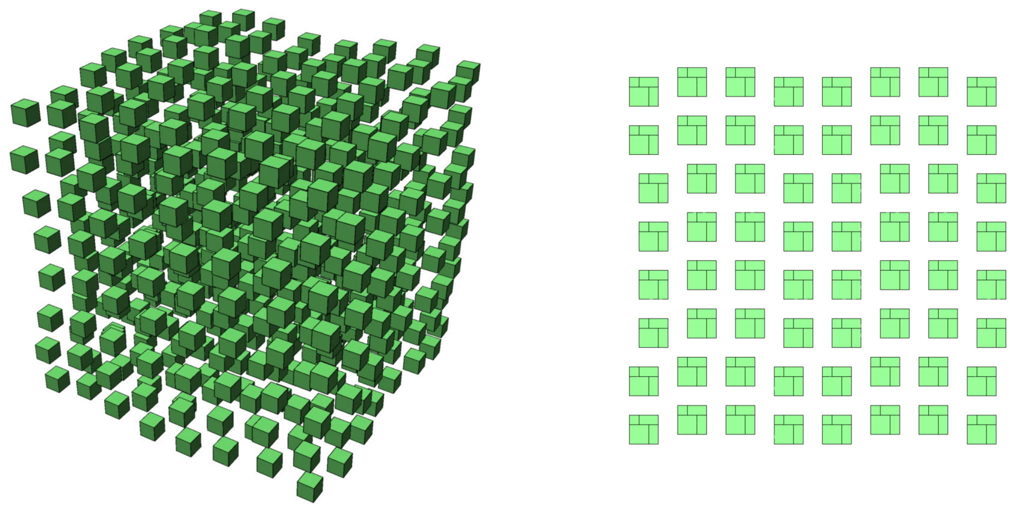

Figure 33: A 3-dimensional URVG representation of Q_9

By cutting the group of rectangles in half repeatedly, we then obtain the following corollary:

Corollary 21 *The minimal embedding dimension of the m-dimensional hypercube graph as a URVG is bounded by*

$$\mu(Q_m) \leq \lceil \sqrt{m} \rceil.$$

6.2.3 MEDs as RVGs

We proceed similarly for normal rectangles, again by first constructing rectangle columns:

Lemma 9 *The d-dimensional RVG formed by rectangles with opposite vertices a_i and b_i, where $d \geq 2$, $0 \leq i < 2^{d+1}$, and*

$$a_i = \begin{bmatrix} (d+4)i \\ 2G_{i,d-2} + G_{i,d-1} - G_{i,d-2}G_{i,d-1} \\ G_{i,0} \\ G_{i,1} \\ \vdots \\ G_{i,d-3} \end{bmatrix}, \ b_i = \begin{bmatrix} (d+4)i+4 \\ G_{i,d-2} - G_{i,d-1} + 4 \\ G_{i,0}+4 \\ G_{i,1}+4 \\ \vdots \\ G_{i,d-3}+4 \end{bmatrix},$$

is Q_{d+1}.

Again, the constant $d+4$ does not effect the validity.

Proof: We use induction on d.

<u>Base case: $d = 2$</u>

The resulting representation is shown in Figure 34.

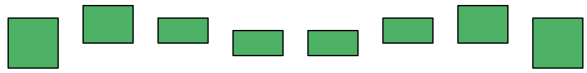

Figure 34: A 2-dimensional RVG representation of Q_3

Inductive step: $d \Rightarrow d+1$

Again, we show the validity of our construction for Q_{d+2} by splitting it into two groups and bijecting each group to the representation of Q_{d+1}. This time, our two groups will be the rectangles corresponding to 0 or 3 $\pmod 4$, as well as those corresponding to 1 or 2 $\pmod 4$. The maps will both send rectangle i of the domain to rectangle $\left\lfloor \frac{i}{2} \right\rfloor$ of the range.

All potential visibility lines are parallel to the 1st axis, and are thus contained in some plane of the form $x_3 = c$, where c is a constant and x_3 is the coordinate along the 3rd axis. Since the rectangles' projections along the 3rd axis are all of the form $[0,4]$ or $[1,5]$, the three distinct cross-sections where a visibility line could be are $x_3 \in (0,1)$, $x_3 \in (1,4)$, or $x_3 \in (4,5)$.

For $x_3 \in (0,1)$, the rectangles in the cross-section are those with $G_{i,0} = 0$, i.e., the first group. In addition, $G_{i,j+1} = G_{\left\lfloor \frac{i}{2} \right\rfloor, j}$, so as the order of the first coordinates is preserved, the map on the first group sends two rectangles in the cross-section to two adjacent rectangles iff they are adjacent.

Similarly, for $x_3 \in (4,5)$, the rectangles in the cross section are those with $G_{i,0} = 1$, i.e., the second group, so similarly the visibility lines here are exactly the desired ones.

All rectangles appear in the cross-section $x \in (1,4)$, so there are no additional visibility lines between two rectangles in the same group. Taking a line that passes through all rectangles, we see that any rectangles $i_0 \neq j_0$ with $\left\lfloor \frac{i_0}{2} \right\rfloor = \left\lfloor \frac{j_0}{2} \right\rfloor$ are adjacent. It thus suffices to show that the converse holds, i.e. that if rectangles i_0 in the first group and j_0 in the second group are adjacent, $\left\lfloor \frac{i_0}{2} \right\rfloor = \left\lfloor \frac{j_0}{2} \right\rfloor$.

Assume for the sake of contradiction that $\left\lfloor \frac{i_0}{2} \right\rfloor \neq \left\lfloor \frac{j_0}{2} \right\rfloor$. Then, the rectangles in the image corresponding to $\left\lfloor \frac{i_0}{2} \right\rfloor$ and $\left\lfloor \frac{j_0}{2} \right\rfloor$ must be adjacent, as otherwise, there would be some rectangle blocking i_0 and j_0 corresponding to the rectangle blocking $\left\lfloor \frac{i_0}{2} \right\rfloor$ and $\left\lfloor \frac{j_0}{2} \right\rfloor$. Thus, by construction, the Gray code representations of $\left\lfloor \frac{i_0}{2} \right\rfloor$ and $\left\lfloor \frac{j_0}{2} \right\rfloor$ differ in exactly one bit, and thus have different parity.

If $\left\lfloor \frac{i_0}{2} \right\rfloor$ is even and $\left\lfloor \frac{i_0}{2} \right\rfloor$ odd, $i_0 \equiv 0 \pmod 4$ and $i_1 \equiv 2 \pmod 4$, so rectangles $i_0 + 1$ and $i_1 + 1$ have the same cross-section as respective rectangles i_0 and i_1 except along the first axis, and thus depending on whether $i_0 < i_1$ or $i_0 > i_1$, one of these blocks any possible visibility line between i_0 and i_1, a contradiction.

On the other hand, if $\left\lfloor \frac{i_0}{2} \right\rfloor$ is odd and $\left\lfloor \frac{i_0}{2} \right\rfloor$ even, $i_0 \equiv 3 \pmod 4$ and $i_1 \equiv 1 \pmod 4$, so one of $i_0 - 1$ and $i_1 - 1$ blocks any possible visibility line between i_0 and i_1, a contradiction.

□

Again we arrange such columns into a grid:

Theorem 12 *The hypercube graph Q_{d^2+d} is a d-dimensional RVG for $d \geq 2$.*

Proof: For any tuple $(i_1, i_2, \ldots, i_d)$ with $0 \leq i_j \leq 2^{d+1} - 1$ for all $j \in \{1, 2, \ldots, d\}$, take a rectangle with opposite corners at

$$\sum_{j=1}^{d} \begin{bmatrix} 0 & & & & & 1 \\ 1 & 0 & & & \mathbf{0} & \\ & 1 & \ddots & & & \\ & & \ddots & \ddots & & \\ & \mathbf{0} & & 1 & 0 & \\ & & & & 1 & 0 \end{bmatrix}^{j-1} \begin{bmatrix} (d+4)i_j \\ 2G_{i_j,d-2} + G_{i_j,d-1} - G_{i_j,d-2}G_{i_j,d-1} \\ G_{i_j,0} \\ G_{i_j,1} \\ \vdots \\ G_{i_j,d-3} \end{bmatrix}$$

and

$$\sum_{j=1}^{d} \begin{bmatrix} 0 & & & & & 1 \\ 1 & 0 & & & \mathbf{0} & \\ & 1 & \ddots & & & \\ & & \ddots & \ddots & & \\ & \mathbf{0} & & 1 & 0 & \\ & & & & 1 & 0 \end{bmatrix}^{j-1} \begin{bmatrix} (d+4)i_j + 4 \\ G_{i_j,d-2} - G_{i_j,d-1} \\ G_{i_j,0} \\ G_{i_j,1} \\ \vdots \\ G_{i_j,d-3} \end{bmatrix}.$$

There are 2^{d^2+d} such tuples. Note that subtracting the opposite coordinates of each rectangle gives that the side lengths are from 2 to 4.

If we fix all but one of $i_1, i_2, \ldots, i_d$, the corresponding rectangles form a Q_{d+1} by Lemma 8. In addition, if WLOG $i_1 > i'_1$ the smaller 1st coordinate of the rectangle corresponding to $(i_1, i_2, \ldots, i_d)$ is more than the larger 1st coordinate of the rectangle corresponding to $(i'_1, i'_2, \ldots, i'_d)$, because

$$\begin{aligned}
&(d+4)i_1 + (2G_{i_d,d-2} + G_{i_d,d-1} - G_{i_d,d-2}G_{i_d,d-1}) + \sum_{j=2}^{d-1} G_{i_j,d-1-j} \\
&= (d+4)i_1 + (2 - (2 - G_{i_d,d-1})(1 - G_{i_d,d-2})) + \sum_{j=2}^{d-1} G_{i_j,d-1-j} \\
&\geq ((d+4)i'_1 + (d+4)(i_1 - i'_1)) + (2 - (2-0)(1-0)) + \sum_{j=2}^{d-1} 0 \\
&\geq (d+4)i'_1 + (d+4) \cdot 1 + 0 + 0 \\
&> (d+4)i'_1 + 4 + 1 + (d-2)
\end{aligned}$$

$$= ((d+4)i'_1 + 4) + (1-0) + \sum_{j=2}^{d-1} 1$$

$$\geq ((d+4)i'_1 + 4) + \left(G_{i'_d, d-2} - G_{i'_d, d-1}\right) + \sum_{j=2}^{d-1} G_{i'_j, d-1-j}.$$

Thus, these are the only visibility lines, and no rectangles intersect.

This construction is applied to $d = 2$ and $d = 3$ in figures 35 and 36, respectively. □

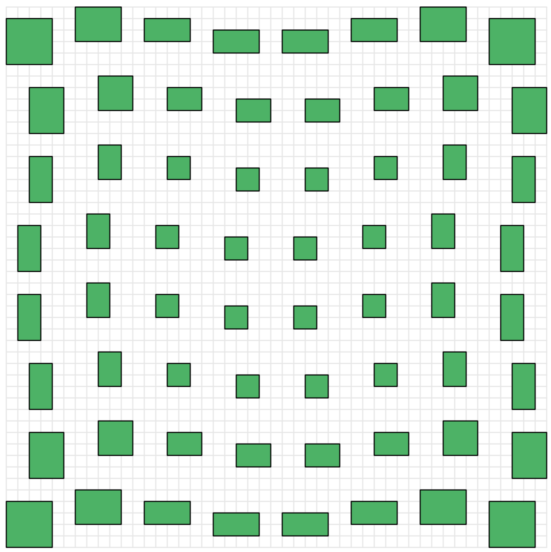

Figure 35: A 2-dimensional RVG representation of Q_6

Because we can cut the group of rectangles in half repeatedly, and because $d^2 + d < \left(d + \frac{1}{2}\right)^2 < \left(d^2 + d\right) + 1$, this gives

Corollary 22 *For all $m \neq 2$, the minimal embedding dimension of the m-dimensional hypercube graph as a RVG is bounded by*

$$M(Q_m) \leq \lfloor \sqrt{m} \rceil$$

(where $\lfloor x \rceil$ denotes x rounded to the nearest integer).

Remark 23 *The minimal embedding dimension d of hypercube graphs Q_m as RVGs include the following:*

- $M(Q_0) = 0$. *A representation of Q_0 in 0 dimensions is shown in Figure 30.*
- $M(Q_1) = 1$. *A representation of Q_1 in 1 dimension is shown in Figure 31. Q_1 cannot be represented in 0 dimensions because visibility lines are 1-dimensional.*

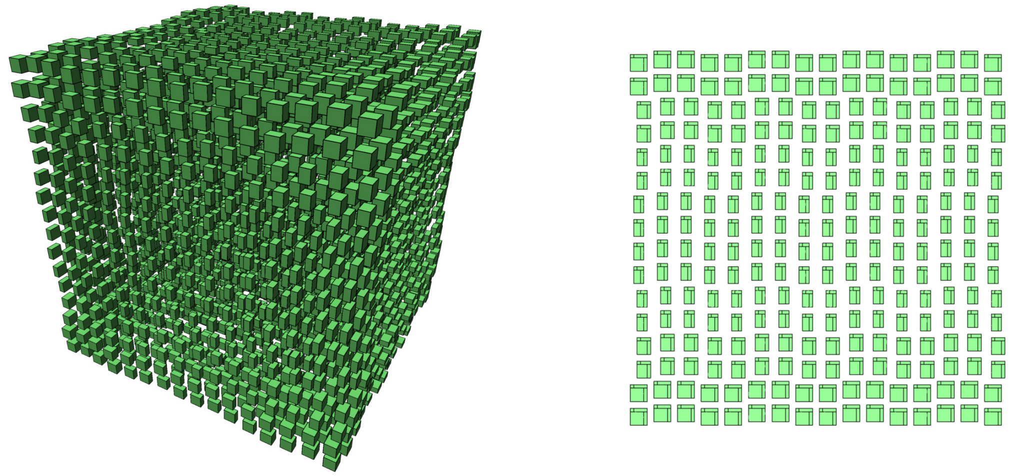

Figure 36: A 3-dimensional RVG representation of Q_{12}

- *$M(Q_i) = 2$ for all $i \in \{2, \ldots, 6\}$. A Q_6 in 2-space is shown in Figure 35. Q_5, Q_4, Q_3 and Q_2 can be obtained by repeatedly removing the right or top half of the rectangles in these configurations, thus ending up with $2^5, 2^4, 2^3$, and 2^2 rectangles in each respective representation.*

  None of these graphs can be represented in 1-space, where the only graphs that can be represented are paths. Thus, the minimal embedding dimensions of Q_2 through Q_6 are $d = 2$.

- *$M(Q_i) = 3$ for all $i \in \{8, \ldots, 12\}$. A representation of Q_{12} in 3-space is shown in Figure 36. There are $16^3 = 2^{4 \cdot 3} = 2^{12}$ boxes in this representation, which correspond to the $n = 2^{12}$ vertices in Q_{12}.*

  Q_{11}, Q_{10}, Q_9 and Q_8 can be obtained by repeatedly removing the top half of the boxes in these configurations, thus ending up with $n = 2^{11}, 2^{10}, 2^9$, and 2^8 boxes in each respective representation.

  Dean and Hutchinson found that a bipartite 2-dimensional RVG on $n \geq 4$ vertices has at most $4n - 12$ edges [7]. Given $m \geq 8$, the number of edges in Q_m is $\frac{m}{2} \cdot 2^m \geq 4 \cdot 2^m = 4n$. Since Q_m is bipartite, and $4n > 4n - 12$, it follows that $Q_{m \geq 8}$ is not representable in 2 dimensions. Thus, the minimal embedding dimensions of Q_8 through Q_{12} are all $d = 3$.

Open Questions

Some questions not answered in this paper but possibly worth exploring in future works include:

1. For fixed n, d, k, what is the maximum number of edges in a d-dimensional (U)RkVG on n vertices?
2. For fixed n, k, what is the maximum MED as a (U)RkVG over all graphs on n vertices?

3. Which of $M_k(K_m), \mu_k(K_m), M_k(Q_m), \mu_k(Q_m)$ are sublinear for fixed k? For those $f(m)$ not sublinear, what is $\lim_{m\to\infty} \frac{f(m)}{m}$ (if it is defined)?
4. For fixed r, k and sufficiently large m, what is $M_k\left(K_{\underbrace{m,\ldots,m}_{r}}\right)$ (presuming it is eventually constant in the first place)?
5. For fixed k, is the MED of Q_m as a (U)RkVG unbounded?
6. What is $M(Q_7)$?
7. For fixed k, can the MED as a (U)RkVG of the composition (under e.g., the tensor product or the strong product) of two graphs with bounded MEDs be unbounded?

Acknowledgments

I am grateful for my wonderful mentor, Dr. Jesse Geneson of San José State University, who inspired me to delve into the topic of visibility graphs in the first place and has guided me throughout this project. I would also like to thank Dr. Pavel Etingof, Dr. Slava Gerovitch, and Dr. Tanya Khovanova of the MIT PRIMES-USA program, and the Mathematics Department of the Massachusetts Institute of Technology for the opportunity to work on this research project. Finally, I thank the anonymous JGAA referees for their thorough and useful comments.

References

[1] M. Babbitt, J. Geneson, and T. Khovanova. On k-visibility graphs. *Journal of Graph Algorithms and Applications*, 19(1):345–360, 2015. doi:10.7155/jgaa.00362.

[2] P. Bose, A. Dean, J. Hutchinson, and T. Shermer. On rectangle visibility graphs. In S. North, editor, *Graph Drawing*, pages 25–44, Berlin, Heidelberg, 1997. Springer Berlin Heidelberg. doi:10.1007/3-540-62495-3_35.

[3] J. Cáceres, C. Hernando, M. Mora, I. M. Pelayo, M. L. Puertas, C. Seara, and D. R. Wood. On the metric dimension of cartesian products of graphs. *SIAM Journal on Discrete Mathematics*, 21(2):423–441, 1 2007. URL: http://dx.doi.org/10.1137/050641867, doi:10.1137/050641867.

[4] H.-C. Chang, H.-I. Lu, and Y. Sung. Asymptotically optimal thickness bounds of generalized 1 bar visibility graphs. 2016. URL: https://users.cs.duke.edu/~hc252/papers/pdf/bark.pdf.

[5] F. Chung. On unimodal subsequences. *Journal of Combinatorial Theory, Series A*, 29(3):267 – 279, 1980. URL: http://www.sciencedirect.com/science/article/pii/0097316580900217, doi:10.1016/0097-3165(80)90021-7.

[6] A. M. Dean, W. Evans, E. Gethner, J. D. Laison, M. Safari, and W. T. Trotter. Bar k-visibility graphs. *Journal of Graph Algorithms and Applications*, 11(1):45–59, 2007. doi:10.7155/jgaa.00136.

[7] A. M. Dean and J. P. Hutchinson. Rectangle-visibility representations of bipartite graphs. *Discrete Applied Mathematics*, 75(1):9 – 25, 1997. URL: http://www.sciencedirect.com/science/article/pii/S0166218X96000297, doi:10.1016/S0166-218X(96)00029-7.

[8] S. P. Fekete, M. E. Houle, and S. Whitesides. New results on a visibility representation of graphs in 3d. In F. J. Brandenburg, editor, *Graph Drawing*, pages 234–241, Berlin, Heidelberg, 1996. Springer Berlin Heidelberg. doi:10.1007/BFb0021807.

[9] S. R. Fekete and H. Meijer. Rectangle and box visibility graphs in 3d. 09, Feb. 1999. URL: https://www.ibr.cs.tu-bs.de/users/fekete/hp/publications/PDF/1999-Rectangle_and_Box_Visibility_Graphs_in_3D.pdf, doi:10.1142/S0218195999000029.

[10] R. Frucht and F. Harary. On the corona of two graphs. *Aequationes Mathematicae*, 4, 10 1970. URL: http://deepblue.lib.umich.edu/bitstream/2027.42/44326/1/10_2005_Article_BF01844162.pdf, doi:10.1007/BF01817769.

[11] C. D. Godsil and B. D. McKay. A new graph product and its spectrum. 18(1):21–28. doi:10.1017/S0004972700007760.

[12] F. Gray. Pulse code communication, Nov. 1947.

[13] S. G. Hartke, J. Vandenbussche, and P. Wenger. Further results on bar k-visibility graphs. 21:523–531, June 2007. doi:10.1137/050644240.

[14] B. Horvat and T. Pisanski. Products of unit distance graphs. *Discrete Mathematics*, 310(12):1783 – 1792, 2010. Algebraic and Topological Graph Theory. URL: http://www.sciencedirect.com/science/article/pii/S0012365X09005949, doi:https://doi.org/10.1016/j.disc.2009.11.035.

[15] B. Lindström. On a combinatorial problem in number theory. *Canadian Mathematical Bulletin*, 8(4):477–490, 1965. doi:10.4153/CMB-1965-034-2.

[16] E. Lodi and L. Pagli. A vlsi solution to the vertical segment visibility problem. *IEEE Transactions on Computers*, 35:923–928, 10 1986. doi:10.1109/TC.1986.1676685.

[17] H. Maehara. On the euclidean dimension of a complete multipartite graph. *Discrete Mathematics*, 72(1):285 – 289, 1988. URL: http://www.sciencedirect.com/science/article/pii/0012365X88902178, doi:https://doi.org/10.1016/0012-365X(88)90217-8.

[18] S. K. Wismath. Bar-representable visibility graphs and a related network flow problem, 08 1989. URL: https://open.library.ubc.ca/media/download/pdf/831/1.0051983/1.

A MEDs vs. Other Dimensions

Here, we compare MEDs of graphs as (U)RVGs with other dimensions of graphs, in particular:

Definition 24 *The Euclidean dimension* $\operatorname{Edim}(G)$ *of a graph* G *is the minimum number of dimensions required to represent its vertices as points such that two points have distance* 1 *if and only if they share an edge.*

Definition 25 *The metric dimension* $\beta(G)$ *of a connected graph* G *is the minimum number of vertices that need to be selected in* G *for every vertex to be uniquely identified by its distance to each selected vertex.*

In the table shown in Figure 38, G_1, G_2, and G are nonempty graphs.

	Edim	β	μ	M
G	$\leq n-1$	$\leq n-1^*$	$\leq n$	$\leq \lceil \frac{n}{2} \rceil$
$G_1 \sqcup G_2$	$\max\begin{pmatrix}\operatorname{Edim}(G_1),\\ \operatorname{Edim}(G_2)\end{pmatrix}$	N/A	$\max\begin{pmatrix}2,\\ \mu(G_1),\\ \mu(G_2)\end{pmatrix}$	$\max\begin{pmatrix}2,\\ M(G_1),\\ M(G_2)\end{pmatrix}$
$G_1 \square G_2$	$\max\begin{pmatrix}2,\\ \operatorname{Edim}(G_1),\\ \operatorname{Edim}(G_2)\end{pmatrix}$ [14]	Unbounded [3]	$\leq \mu(G_1)+\mu(G_2)$	$\leq M(G_1)+M(G_2)$
K_m	$m-1$	$m-1$	$\leq \lceil \frac{3}{5}m \rceil$	$\leq \max\begin{pmatrix}3,\\ \lceil \frac{m-20}{2} \rceil\end{pmatrix}$
$K_{m_1,\ldots,m_r}$	$\leq 2r$ [17]	$n-r^\dagger$	$\leq n$	$\leq r+1$
Q_m	≤ 2	$\sim \frac{m \log 4}{\log m}$ [15]	$\leq \lceil \sqrt{m} \rceil$	$\leq \lfloor \sqrt{m} \rceil$
$\dim(G)=0$	K_1	K_1	K_1	K_1
$\dim(G)\leq 1$	$\bigsqcup_{j=1}^{r} P_{m_j}$	P_m	P_m	P_m

Figure 37: Comparison of minimal embedding dimensions to other dimensions

*Defined only if G is connected.

†n is the number of vertices, $\sum_{j=1}^{r} m_j$. Holds if and only if at most one m_j is under 2.