

Algorithms for Drawing Media

David Eppstein*

Computer Science Department,
School of Information & Computer Science,
University of California, Irvine
eppstein@uci.edu

Abstract. We describe algorithms for drawing media, systems of states, tokens and actions that have state transition graphs in the form of partial cubes. Our algorithms are based on two principles: embedding the state transition graph in a low-dimensional integer lattice and projecting the lattice onto the plane, or drawing the medium as a planar graph with centrally symmetric faces.

1 Introduction

Media [7, 8] are systems of states, tokens, and actions of tokens on states that arise in political choice theory and that can also be used to represent many familiar geometric and combinatorial systems such as hyperplane arrangements, permutations, partial orders, and phylogenetic trees. In view of their importance in modeling social and combinatorial systems, we would like to have efficient algorithms for drawing media as state-transition graphs in a way that makes the action of each token apparent. In this paper we describe several such algorithms.

Formally, a *medium* consists of a finite set of *states* transformed by the actions of a set of *tokens*. A string of tokens is called a *message*; we use upper case letters to denote states, and lower case letters to denote tokens and messages, so Sw denotes the state formed by applying the tokens in message w to state S . Token t is *effective* for S if $St \neq S$, and message w is *stepwise effective* for S if each successive token in the sequence of transformations of S by w is effective. A message is *consistent* if it does not contain the reverse of any of its tokens. A set of states and tokens forms a medium if it satisfies the following axioms:

1. Each token t has a unique *reverse* \tilde{t} such that, for any states $S \neq Q$, $St = Q$ iff $Q\tilde{t} = S$.
2. For any states $S \neq Q$, there exists a consistent message w with $Sw = Q$.
3. If message w is stepwise effective for S , then $Sw = S$ if and only if the number of copies of t in w equals the number of copies of \tilde{t} for each token t .
4. If $Sw = Qz$, w is stepwise effective for S , z is stepwise effective for Q , and both w and z are consistent, then wz is consistent.

The states and state transitions of a medium can also be viewed as a graph, and it can be shown that these graphs are *partial cubes* [12]: that is, their vertices can be mapped to a hypercube $\{0, 1\}^d$ in such a way that graph distance equals L_1 distance in

* Supported in part by NSF grant CCR-9912338.




Fig. 1. 11 of the 12 pentominoes represent isometric lattice embeddings of media. The twelfth, the U pentomino, does not, because a pair of vertices that are three edges apart in the graph have placements that are only one unit apart.

the hypercube. For media, we can find such a mapping by choosing arbitrarily state S , and assigning any state S' a coordinate per token t that is 1 when a consistent path from S to S' contains t and 0 otherwise. Conversely, any d -dimensional partial cube gives rise to a medium with its vertices as states and with $2d$ tokens; the action of any token is to change one of the partial cube coordinates to a zero or to a one, if it does not already have that value and if such a change would produce another vertex of the partial cube.

We assume throughout, as in [7], that we are given as input an explicit description of the states, tokens, and actions of a medium. However, our algorithms are equally applicable to any partial cube or family of partial cube graphs such as the median graphs. If a partial cube representation is not given, it can be found (and the corresponding medium constructed) in time $O(mn)$ via known algorithms [1, 11, 12, 15].

2 Lattice Dimension

As we have seen, media can be embedded *isometrically* (that is, in a distance-preserving way) into hypercubes $\{0, 1\}^d$ (with L_1 distance), and hypercubes can be embedded isometrically into integer lattices \mathbb{Z}^d , so by transitivity media can be embedded isometrically onto integer lattices. Conversely any finite isometric subset of an integer lattice forms a partial cube and corresponds as described above to a medium.

If the dimension of the lattice in which a medium is embedded is low, we may be able to use the embedding as part of an effective drawing algorithm. For instance, if a medium M can be embedded isometrically onto the planar integer lattice \mathbb{Z}^2 , then we can use the lattice positions as vertex coordinates of a drawing in which each edge is a vertical or horizontal unit segment (Figure 1). If M can be embedded isometrically onto the cubic lattice \mathbb{Z}^3 , in such a way that the projection onto a plane perpendicular to the vector $(1, 1, 1)$ projects different vertices to distinct positions in the plane, then this projection produces a planar graph drawing in which the edges are unit vectors at 60° and 120° angles (Figure 10, center).

Recently, we showed that the *lattice dimension* of a medium or partial cube, that is, the minimum dimension of a lattice \mathbb{Z}^d into which it may be isometrically embedded, may be determined in polynomial time [6]. We now briefly our algorithm for finding low-dimensional lattice embeddings.

Suppose we are given an undirected graph G and an isometry $\mu : G \mapsto \{0, 1\}^\tau$ from G to the hypercube $\{0, 1\}^\tau$ of dimension τ . Let $\mu_i : G \mapsto \{0, 1\}$ map each vertex v of




Fig. 2. A medium, left, and its semicube graph, right. From [6].




Fig. 3. A matching in the semicube graph (left, solid edges) completed to a set of paths by adding edges from each semicube to its complement (left, dashed edges), and the corresponding lattice embedding of the original medium (right). From [6].

G to the i th coordinate of $\mu(v)$, and assume that each coordinate μ_i takes on both value 0 and 1 for at least one point. From G and μ we can define 2τ distinct *semicubes* $S_{i,\chi} = \{v \in V(G) \mid \mu_i(v) = \chi\}$, for any pair i, χ with $0 \leq i < \tau$ and $\chi \in \{0, 1\}$. We now construct a new graph $Sc(G)$, which we call the *semicube graph* of G . We include in $Sc(G)$ a set of 2τ vertices $u_{i,\chi}$, $0 \leq i < \tau$ and $\chi \in \{0, 1\}$. We include an edge in $Sc(G)$ between $u_{a,b}$ and $u_{c,d}$ whenever $S_{a,b} \cup S_{c,d} = V(G)$ and $S_{a,b} \cap S_{c,d} \neq \emptyset$; that is, whenever the corresponding two semicubes cover all the vertices of G non-disjointly. Although defined from some particular isometry μ , the semicube graph turns out to be independent of the choice of μ . An example of a partial cube G and its semicube graph $Sc(G)$ is shown in Figure 2. The main result of [6] is that the lattice dimension of G can be determined from the cardinality of a maximum matching in $Sc(G)$:

Theorem 1 (Eppstein [6]). *If G is a partial cube with isometric dimension τ , then the lattice dimension of G is $d = \tau - |M|$ where M is any maximum matching in $Sc(G)$.*

More specifically, we can extend a matching in $Sc(G)$ to a collection of d paths by adding to the matching an edge from each semicube to its complement. The d th coordinate of a vertex in the lattice embedding equals the number of semicubes that contain the vertex in even positions along the d th path.

We can use this result as part of a graph drawing system, by embedding our input medium in the lattice of the lowest possible dimension and then projecting that lattice onto the plane. For two-dimensional lattices, no projection is needed, and we have already discussed projection of certain three-dimensional integer lattices onto two-dimensional triangular lattices. We discuss more general techniques for lattice projec-

tion in the next section. It is essential for this result that we require the embedding to be isometric. Even for trees it is NP-complete to find an embedding into \mathbb{Z}^2 with unit length edges that is not required to be distance-preserving [2]. However a tree embeds isometrically in \mathbb{Z}^2 if and only if it has at most four leaves [14].

3 Drawing High-Dimensional Lattice Graphs

Two-dimensional lattice embeddings of media, and some three-dimensional embeddings, lead to planar graph drawings with all edges short and well separated by angles. However, we are also interested in drawing media without low dimensional embeddings. We describe here a method for finding drawings with the following properties:

1. All vertices are assigned distinct integer coordinates in \mathbb{Z}^2 .
2. All edges are drawn as straight line segments.
3. No edge passes closer than unit distance to a vertex that is not one of its endpoints.
4. The line segments representing two edges of the drawing are translates of each other if and only if the two edges are parallel in the lattice embedding.
5. The medium corresponding to a Cartesian product of intervals $[a_0, b_0] \times [a_1, b_1] \times \dots \times [a_{d-1}, b_{d-1}]$ is drawn in area $O(n^2)$, where n is the number of its states.

Because of property 4, the lattice embedding and hence the medium structure of the state transition graph can be read from the drawing. To achieve these properties, we map \mathbb{Z}^d to \mathbb{Z}^2 linearly, by choosing two vectors X and $Y \in \mathbb{Z}^d$, and mapping any point $p \in \mathbb{Z}^d$ to the point $(X \cdot p, Y \cdot p) \in \mathbb{Z}^2$. We now describe how these vectors X and Y are chosen. If $L \subset \mathbb{Z}^d$ is the set of vertex placements in the lattice embedding of our input medium, define a *slice* $L_{i,j} = \{p \in L \mid p_i = j\}$ to be the subset of vertices having i th coordinate equal to j . We choose the coordinates X_i sequentially, from smaller i to larger, so that all slices $L_{i,j}$ are separated from each other in the range of x -coordinates they are placed in. Specifically, set $X_0 = 0$. Then, for $i > 0$, define

$$X_i = \max_j \left(\min_{p \in L_{i,j}} \sum_{k=0}^{i-1} X_k p_k - \max_{q \in L_{i,j-1}} \sum_{k=0}^{i-1} X_k q_k \right),$$

where the outer maximization is over all j such that $L_{i,j}$ and $L_{i,j-1}$ are both nonempty. We define Y similarly, but we choose its coordinates in the opposite order, from larger i to smaller: $Y_{d-1} = 0$, and

$$Y_i = \max_j \left(\min_{p \in L_{i,j}} \sum_{k=i+1}^{d-1} X_k p_k - \max_{q \in L_{i,j-1}} \sum_{k=i+1}^{d-1} X_k q_k \right).$$

Theorem 2. *The projection method described above satisfies the properties 1–5 listed above. The method’s running time on a medium with n states and τ tokens is $O(n\tau^2)$.*

Proof. Property 2 and property 4 follow immediately from the fact that we our drawing is formed by projecting \mathbb{Z}^d linearly onto \mathbb{Z}^2 , and from the fact that the formulas used to calculate X and Y assign different values to different coordinates of these vectors.

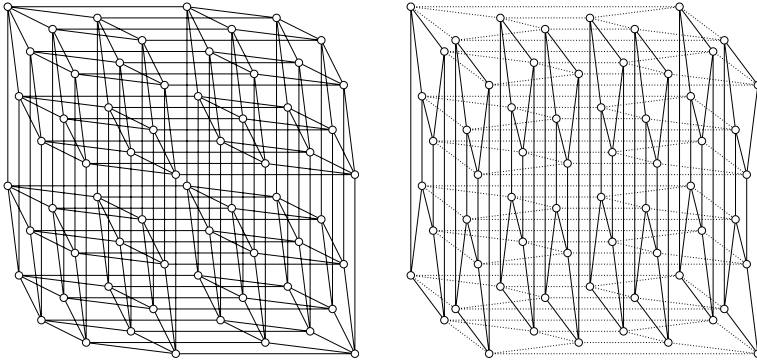


Fig. 4. Left: lattice drawing of six-dimensional hypercube; right: a drawing with geometric thickness two is possible, but the vertex placement is less regular and edges formed by actions of the same token are not all drawn parallel.

All vertices are assigned distinct coordinates (property 1): for, if vertices p and q differ in the i th coordinates of their lattice embeddings, they belong to different slices $L_{i,j}$ and $L_{i,j'}$ and are assigned X coordinates that differ by at least X_i (unless $i = X_i = 0$ in which case their Y coordinates differ by at least Y_i).

The separation between vertices and edges (property 3) is almost equally easy to verify: consider the case of three vertices p , q , and r , with an edge pq to be separated from r . Since p and q are connected by an edge, their lattice embeddings must differ in only a single coordinate i . If r differs from p and q only in the same coordinate, it is separated from edge pq by a multiple of (X_i, Y_i) . Otherwise, there is some coordinate $i' \neq i$ in which r differs from both p and q . If $i' > i$, the construction ensures that the slice $L_{i',j}$ containing pq is well separated in the x -coordinate from the slice $L_{i',j'}$ containing r , and if $i' < i$ these slices are well separated in the y coordinate.

Finally, we consider property 5. For Cartesian products of intervals, in the formula for X_i , the value for the subexpression $\min_{p \in L_{i,j}} \sum_{k=0}^{i-1} X_k p_k$ is the same for all j considered in the outer maximization, and the value for the subexpression $\max_{q \in L_{i,j-1}} \sum_{k=0}^{i-1} X_k q_k$ is also the same for all j considered in the outer maximization, because the slices are all just translates of each other. Therefore, there is no gap in x -coordinates between vertex placements of each successive slice of the medium. Since our drawings of these media have vertices occupying contiguous integer x coordinates and (by a symmetric argument) y coordinates, the total area is at most n^2 .

The time for implementing this method is dominated by that for finding a minimum-dimension lattice embedding of the input graph, which can be done in the stated time bound [6]. \square

When applied to a hypercube, the coordinates X_i become powers of two, and this vertex placement algorithm produces a uniform placement of vertices (Figure 4, left) closely related to the Hammersley point set commonly used in numerical computation and computer graphics for its low discrepancy properties [16]. Other examples of drawings produced by this method can be seen in Figures 6, 9, and 10(left).




Fig. 5. Left: a graph with a face-symmetric planar drawing; center: connecting opposite pairs of edge midpoints produces a weak pseudoline arrangement; right: the arrangement.

4 Face-Symmetric Planar Drawings

Our two-dimensional and projected three-dimensional lattice drawings are planar (no two edges cross) and each internal face is symmetric (squares for two-dimensional lattices, 60° - 120° rhombi and regular hexagons for projected three-dimensional lattices). We now describe a different type of drawing of the state-transition graphs of media as planar graphs, generalizing this symmetry property. Specifically, we seek straight-line planar drawings in which each internal face is strictly convex and centrally symmetric; we call such a drawing a *face-symmetric planar drawing*.

A *weak arrangement of pseudolines* [9] is a collection of curves in the plane, each homeomorphic to a line, such that any pair of curves in the collection has at most one point of intersection, and such that if any two curves intersect then they cross properly at their intersection point. Weak arrangements of pseudolines generalize pseudoline arrangements [10] and hyperbolic line arrangements, and are a special case of the *extendible pseudosegment arrangements* defined by Chan [3]. Any weak pseudoline arrangement with n pseudolines partitions the plane into at least $n + 1$ and at most $n(n + 1)/2 + 1$ cells, connected components of the set of points that do not belong to any pseudoline. Any pseudoline in the arrangement can be partitioned into *nodes* (crossing points) and *arcs* (connected components of the complement of the crossing points); we use this terminology to avoid confusion with the vertices and edges of the medium state-transition graphs we hope to draw. Each arc is adjacent to two cells and two nodes. We define the *dual* of a weak pseudoline arrangement to be the graph having a vertex for each cell of the arrangement and an edge connecting the vertices dual to any two cells that share a common arc; this duality places the graph's vertices in one-to-one correspondence with the arrangement's cells, and the graph's edges in one-to-one correspondence with the arrangement's arcs.

Lemma 1. *If G has a face-symmetric planar drawing, then G is the dual of a weak pseudoline arrangement.*

Lemma 2. *If G is the dual of a weak pseudoline arrangement, then G is the state transition graph of a medium.*




Fig. 6. Media with planar state-transition graphs but with no face-symmetric planar drawing.




Fig. 7. Converting a weak pseudoline arrangement into a face-symmetric planar drawing. Left: arrangement drawn inside a circle O such that crossings with O are equally spaced around the circle. Right: edges dual to arcs of ℓ_i are drawn as unit length and perpendicular to the chord through the points where ℓ_i crosses O .

By these lemmas (the proofs of which we omit due to lack of space), every face-symmetric planar drawing represents the state transition graph of a medium. However, not every medium, and not even every medium with a planar state transition graph, has such a drawing; see for instance Figure 6, the medium in Figure 9(right), and the permutahedron in Figure 10(left) for media that have planar state transition graphs but no face-symmetric planar drawing.

Lemma 3. *If G is the dual of a weak pseudoline arrangement, then G has a face-symmetric planar drawing.*

Proof. Let G be dual to a weak pseudoline arrangement \mathcal{A} ; the duality fixes a choice of planar embedding of G as well as determining which faces of that embedding are internal and external. Denote by $|\mathcal{A}|$ the number of pseudolines in \mathcal{A} . Let O be a circle (the size and placement of the circle within the plane being irrelevant to our construction), and deform \mathcal{A} as necessary so that each pseudoline crosses O , with all nodes interior to O , and so that the $2|\mathcal{A}|$ points where pseudolines cross O are spaced at equal distances around the perimeter of O (Figure 7, left). Then, for each pseudoline ℓ_i of \mathcal{A} , let c_i be the chord of O connecting the two points where ℓ_i crosses O . We will draw G in such a way that the edges of G that are dual to arcs of ℓ_i are drawn as unit length segments perpendicular to c_i (Figure 7, right). To do so, choose an arbitrary starting vertex v_0 of G , and place it arbitrarily within the plane. Then, the placement of any other vertex v_i of G can be found by following a path from v_0 to v_i in G , and for each edge of the path

moving unit distance (starting from the location of v_0) in the direction determined for that edge as described above, placing v_i at the point reached by this motion when the end of the path is reached. It is straightforward to show from Lemma 2 and the axioms defining a medium that this vertex placement does not depend on the choice of the path from v_0 to v_i , and that if all vertices are placed in this way then all edges of G will be unit length and perpendicular to their corresponding chords c_i . Thus, we have a drawing of G , in which we can identify sets of edges corresponding to the faces of G . We omit the proof that this drawing is face-symmetric planar due to lack of space. \square

Lemma 4. *If G is biconnected, at most one planar embedding of G is dual to a weak pseudoline arrangement. This embedding (if it exists) can be found in time $O(n)$.*

Proof. We use a standard technique in graph drawing and planar embedding problems, the SPQR tree [4, 13]. Each node v in the SPQR tree of G has associated with it a multigraph G_v consisting of some subset of vertices of G , edges of G , and *virtual edges* representing contracted parts of the remaining graph that can be separated from the edges of G_v by a *split pair* of vertices (the endpoints of the virtual edge). The non-virtual edges of G are partitioned among the nodes of the SPQR tree. If two nodes are connected by an edge in the SPQR tree, each has a virtual edge connecting two vertices shared by both nodes. We root the SPQR tree arbitrarily; let (s_v, t_v) denote the split pair connecting a non-root node v to its parent, and let H_v denote the graph represented by the SPQR subtree rooted at v . We work bottom up in the rooted tree, showing by induction on tree size that the following properties hold for each node of the tree:

1. Each graph H_v has at most one planar embedding that can be part of an embedding of G dual to a weak pseudoline arrangement.
2. If v is a non-root node, and G is dual to a weak pseudoline arrangement, then edge $s_v t_v$ belongs to the outer face of the embedding of H_v .
3. If v is a non-root node, form the path p_v by removing virtual edge $s_v t_v$ from the outer face of H_v . Then p_v must lie along the outer face of any embedding of G dual to a weak pseudoline arrangement.

SPQR trees are divided into four different cases (represented by the initials S, P, Q, and R) and our proof follows the same case analysis, in each case showing that the properties at each node follow from the same properties at the descendant nodes. We omit the details of each case due to lack of space. \square

Theorem 3. *Given an input graph G , we can determine whether G is the dual of a weak pseudoline arrangement, and if so construct a face-symmetric planar drawing of G , in linear time.*

Proof. If G is biconnected, we choose a planar embedding of G by Lemma 4. Otherwise, each articulation point of G must be on the outer face of any embedding. We find biconnected components of G , embed each component by Lemma 4, and verify that these embeddings place the articulation points on the outer faces of each component. We then connect the embeddings together into a single embedding having as its outer face the edges that are outer in each biconnected component; the choice of this embedding may not be unique but does not affect the correctness of our algorithm.




Fig. 8. Face-symmetric planar drawings of three irregular media.

Once we have an embedding of G , we must verify that we have the dual of a weak pseudoline arrangement and construct a face-symmetric planar drawing. We first make sure all faces of G are even, and apply Lemma 1 to construct an arrangement of curves \mathcal{A} dual to G . We test that \mathcal{A} has no closed curves, then apply the construction of Lemma 3 to produce vertex placements for a drawing of G , test for each edge of G that the endpoints of that edge are placed at unit distance apart with the expected slope, and test that each internal face of G is drawn as a correctly oriented strictly convex polygon. If our input passes all these tests we have determined that it is the dual of a weak pseudoline arrangement and found a face-symmetric planar drawing. \square

Our actual implementation is based on a simpler but less efficient algorithm that uses the known medium structure of the input to construct the dual weak pseudoline arrangement one curve at a time, before applying the construction of Lemma 3 to produce a face-symmetric planar drawing from the weak pseudoline arrangement. Examples of drawings produced by our face-symmetric planar drawing code are shown in Figure 8.

5 Implementation and Examples

We implemented our algorithms in Python, with drawings output in SVG format. Our code allows various standard combinatorial media (such as the collection of permutations on n items) to be specified on the command line; irregular media may be loaded from a file containing hypercube or lattice coordinates of each state. We have seen already examples of our implementation's output in Figures 4, 6, 8, and 9. Figure 10 provides additional examples. All figures identified as output of our code have been left unretouched, with the exception that we have decolorized them for better printing.

6 Conclusions and Open Problems

We have shown several methods for drawing the state transition graphs of media. There are several interesting directions future research in this area could take.




Fig. 9. Lattice drawings of four irregular media with three-dimensional lattice embeddings, from [6]. The bottom left drawing is of a medium isomorphic to the weak ordering medium shown in Figure 10(right).




Fig. 10. Media defined by orderings of n -item sets. Left: Lattice drawing of total orderings (permutations) on four items. Center: Projected three-dimensional lattice drawing of partial orderings on three items. Right: Face-symmetric planar drawing of weak orderings on three items.

- If a three-dimensional lattice embedding can be projected perpendicularly to the vector $(1, 1, 1)$ (or more generally $(\pm 1, \pm 1, \pm 1)$) without placing two vertices in the same point, the projection produces a planar drawing with all edges having equal lengths and angles that are multiples of 60° (e.g., Figure 10, center). Our lattice dimension algorithm can find a three-dimensional embedding, if one exists, and it is trivial to test the projection property. However, a medium may have more than one three-dimensional embedding, some of which have the projection property and some of which don't. For instance, the medium in the lower right of Figure 9 is the same weak ordering medium as the one in Figure 10(right), however the former drawing is from a lattice embedding without the projection property. Is it possible to efficiently find a projectable three-dimensional lattice embedding, when one exists? More generally, given an arbitrary dimension lattice embedding of a medium, can we find a planar projection when one exists?
- Hypercubes have projected lattice drawings with $O(n^2)$ area and unit separation between vertices and nonadjacent edges. Can similar area and separation bounds be achieved for projected lattice drawings of more general classes of media?
- Our lattice and face-symmetric planar drawings have several desirable qualities; for instance, all edges corresponding to a single token are drawn as line segments with the same slope and length, and our lattice drawings have good vertex-vertex

and vertex-edge separation. However, we have not seriously examined the extent to which other important graph drawing properties may be achieved. For instance, d -dimensional hypercubes (and therefore also media with up to $2d$ tokens) may be drawn with geometric thickness [5] at most $\lceil d/3 \rceil$ (Figure 4, right) however our lattice projection methods achieve geometric thickness only $\lceil d/2 \rceil$ while the only way we know how to achieve the better $\lceil d/3 \rceil$ bound is to use a more irregular drawing in which edges coming from the same token are no longer parallel.

References

1. F. Aurenhammer and J. Hagauer. Recognizing binary Hamming graphs in $O(n^2 \log n)$ time. *Mathematical Systems Theory* 28:387–395, 1995.
2. S. Bhatt and S. Cosmodakis. The complexity of minimizing wire lengths in VLSI layouts. *Inform. Proc. Lett.* 25:263–267, 1987.
3. T. M. Chan. On levels in arrangements of curves. *Discrete & Comput. Geom.* 29(3):375–393, April 2003.
4. G. Di Battista and R. Tamassia. Incremental planarity testing. *Proc. 30th IEEE Symp. Foundations of Computer Science (FOCS 1989)*, pp. 436–441, 1989.
5. D. Eppstein. Layered graph drawing. <http://www.ics.uci.edu/~eppstein/junkyard/thickness/>.
6. D. Eppstein. The lattice dimension of a graph. To appear in *Eur. J. Combinatorics*, arXiv:cs.DS/0402028.
7. D. Eppstein and J.-C. Falmagne. Algorithms for media. ACM Computing Research Repository, June 2002, arXiv:cs.DS/0206033.
8. J.-C. Falmagne and S. Ovchinnikov. Media theory. *Discrete Applied Mathematics* 121(1–3): 103–118, September 2002.
9. H. de Fraysseix and P. Ossona de Mendez. Stretching of Jordan arc contact systems. *Proc. 11th Int. Symp. Graph Drawing (GD 2003)*, pp. 71–85. Springer-Verlag, Lecture Notes in Computer Science 2912, 2003.
10. J. E. Goodman and R. Pollack. Allowable sequences and order types in discrete and computational geometry. *New Trends in Discrete and Computational Geometry*, chapter V, pp. 103–134. Springer-Verlag, Algorithms and Combinatorics 10, 1993.
11. W. Imrich and S. Klavžar. On the complexity of recognizing Hamming graphs and related classes of graphs. *Eur. J. Combinatorics* 17:209–221, 1996.
12. W. Imrich and S. Klavžar. *Product Graphs*. John Wiley & Sons, 2000.
13. P. Mutzel. The SPQR-tree data structure in graph drawing. *Proc. 30th Int. Coll. Automata, Languages and Computation (ICALP 2003)*, pp. 34–46. Springer-Verlag, Lecture Notes in Computer Science 2719, June 2003.
14. S. Ovchinnikov. The lattice dimension of a tree. arXiv.org, February 2004, arXiv:math.CO/0402246.
15. P. Winkler. Isometric embeddings in products of complete graphs. *Discrete Applied Mathematics* 7:221–225, 1984.
16. T.-T. Wong, W.-S. Luk, and P.-A. Heng. Sampling with Hammersley and Halton points. *J. Graphics Tools* 2(2):9–24, 1997.