Three Approaches to 3D-Orthogonal Box-Drawings*
(Extended Abstract)

Therese C. Biedl

McGill University, 3480 University St. #318, Montréal, Qc. H3A 2A7, Canada,
therese@cs.mcgill.ca.

Abstract. In this paper, we study orthogonal graph drawings in three dimensions with nodes drawn as boxes. The algorithms that we present can be differentiated as resulting from three different approaches to creating 3D-drawings; we call these approaches edge-lifting, half-edge-lifting, and three-phase-method.
Let G be a graph with n vertices, m edges, and maximum degree Δ. We obtain a drawing of G in an $n \times n \times \Delta$-grid where the surface area of the box of a node v is $\mathcal{O}(deg(v))$; this improves significantly on previous results. We also consider drawings with at most one node per grid-plane, and exhibit constructions in an $n \times n \times m$-grid and a lower bound of $\Omega(m^2)$; hence upper and lower bounds match for graphs with $\theta(n^2)$ edges.

1 Introduction

In this paper, we study *orthogonal drawings*, i.e., embeddings in the rectangular grid. We will mainly be concerned with drawings in dimension $k = 3$, but need the terminology for dimension $k = 2$ as well. A *grid-point* is a point in R^k whose coordinates are all integer. A *grid-box* is the set of all grid-points $(x_1, \ldots, x_k) \in R^k$ satisfying $x_i^l \leq x_i \leq x_i^u$ for some integers x_i^l, x_i^u, $i = 1, \ldots, k$. A *port* of a grid-box is a point of the box that is extremal in at least one direction.

Throughout this paper, a *kD orthogonal grid drawing* of a graph G is a drawing that satisfies the following. Distinct nodes are represented by disjoint k-dimensional grid-boxes. An edge $e = (v_1, v_2)$ is drawn as a simple path that follows grid-lines, possibly bending at grid-points; the endpoints of the path for e are ports on the boxes representing v_1 and v_2. The intermediate points along the path for an edge do not belong to any node box. For $k = 3$, the intermediate points along the path for an edge also do not belong to any other edge path; for $k = 2$, edge-paths may cross, but not touch or overlap.

The volume of a 3D-drawing is the volume of the smallest grid-box containing the drawing. Often we refer to this bounding box as an $X \times Y \times Z$-grid. In what

* These results were part of a PhD thesis at Rutgers University under the supervision of Prof. Endre Boros, and done, in part, while the author was working at Tom Sawyer Software and funded by the NIST under grant number 70NANB5H1162. Patent on these and related results is pending.

S.H. Whitesides (Ed.): GD'98, LNCS 1547, pp. 30–43, 1998.

follows, graph theoretic terms such as *node* are typically used to refer both to the graph theoretic object and to its representation in a drawing.

The orthogonal drawing style has received much attention in the graph drawing community, for example 11 out of 43 papers at the last graph drawing conference [7] pertained to them. The orthogonal drawing algorithms split into two major classes. If the maximum degree of the input graph is bounded by twice the dimension, then every node can be drawn as a point (we speak of a *point-drawings*). If the maximum degree exceeds this bounds, then one assigns a box to every node (we speak of a *box-drawing*). For aesthetical reasons, this box should be small relative to the degree of the node.

1.1 Existing Results

The problem of 2D-orthogonal drawings has been studied extensively, both for point-drawings and for box-drawings. See [1, 3, 6, 9, 10, 11, 12, 15] and the references therein.

3D-orthogonal drawings have been studied almost exclusively for point-drawings, so for graphs with maximum degree 6 (*6-graphs*). Rosenberg gave an algorithm to embed any 6-graph in a grid of volume $\mathcal{O}(n^{3/2})$, and showed that this is asymptotically optimal [14]. No bounds on the number of bends are given. Eades, Symvonis and Whitesides proposed drawings in an $\mathcal{O}(\sqrt{n}) \times \mathcal{O}(\sqrt{n}) \times \mathcal{O}(\sqrt{n})$-grid with at most 7 bends per edge [8]. They also gave a construction in a $3n \times 3n \times 3n$-grid with 3 bends per edge. This was later improved by Papakostas and Tollis to a grid of volume at most $4.66n^3$ with 3 bends per edge [13].

For 3D-orthogonal box-drawings, only two results are known to the author. Papakostas and Tollis presented an algorithm to embed any graph in a grid of volume $\mathcal{O}(m^3)$ with at most 2 bends per edge [13]. The author, together with Shermer, Whitesides and Wismath, studied how to embed the complete graph K_n, and established lower bounds [4].

1.2 Our Results

In this paper, we review old and present new algorithms for 3D-orthogonal box-drawings. These algorithms fit into three very different approaches to creating an orthogonal drawing. Two of these approaches have been used [4, 8] without being defined abstractly.

The first approach, which we call *edge-lifting*, yields drawings with excellent volume-bounds, but nodes may be disproportionally large. The second approach, which we call *half-edge-lifting*, yields drawings in which nodes are proportional to the degree of the node, i.e., they are in what we call the *degree-restricted model.* This approach makes it possible to convert many two-dimensional orthogonal graph drawings into three-dimensional ones. The third approach is called *three-phase method*, because it mirrors the three-phase method introduced for 2D-orthogonal drawings in [3].

The second and third approach result in new drawing algorithms with improved volume bounds. In particular, we improve the volume bound of $\mathcal{O}(m^3)$

[13] to $\mathcal{O}(n^3)$, while maintaining the property that the surface area of each node is proportional to the degree of the node. We also construct drawings in which the nodes are represented by cubes, at a slight increase of the volume to $\mathcal{O}(n^2 m)$. To our knowledge, this is the first algorithm that draws nodes as cubes.

The drawings created with the first two approaches can be considered two-dimensional in spirit, because they are created by starting with a 2D-orthogonal drawing and lifting it into 3D. The third approach works differently, by placing nodes directly in three-dimensional space. This enables us to study drawings with at most one node per grid-plane. We exhibit constructions that achieve a volume of $\mathcal{O}(n^2 m)$, and a volume of $\mathcal{O}(\sqrt{nm}^3)$, respectively; the latter construction again represents nodes as cubes. We also study lower bounds, and prove that no such drawing could have less than $\mathcal{O}(\max\{n^3, m^2\})$ volume, thus the smaller of our constructions is optimal for graphs with $\theta(n^2)$ edges.

2 Preliminaries

In the following, we clarify some terminology used for 3D-drawings. Recall that a *grid-box* in dimension k is the set of all grid-points $(x_1, \ldots, x_k) \in R^k$ satisfying $x_i^l \leq x_i \leq x_i^u$ for some integers x_i^l, x_i^u, $i = 1, \ldots, k$.[1] A grid-box is said to have *width* $w = x_1^u - x_1^l + 1$ and *height* $h = x_2^u - x_2^l + 1$; thus we measure the number of grid-points across, not the distance between the first and the last grid-point. In 3D, we also use the terms *depth* $d = x_3^u - x_3^l + 1$, *size* $w \times d \times h$, and *volume* whd. The volume is thus the number of grid-points contained in the grid-box, and can never be 0.

A *port* of a grid-box is a grid-point in the box that is extremal in at least one direction. Ports are classified by their direction of extremity as $+x$-ports, $-x$-ports, $+y$-ports, etc. An *x-port* is a port that is either a $+x$-port or a $-x$-port; y-ports and z-ports are defined similarly. Points that are extremal in more than one direction are counted repeatedly as ports. The total number of ports of a box is called the *surface area* of the box; for a box with width w, height h and depth d the surface area is thus $2(wh + hd + wd)$.

A *z-line* is a grid-line that is parallel to the z-axis. A *z-plane* is a grid-plane that is orthogonal to the z-axis. The *coordinate* of a z-plane is the fixed z-coordinate. A *z-segment* is a segment along a z-line. The terms are defined similarly for x and y. A *grid-plane* is an x-plane, y-plane or z-plane with integer fixed coordinate.

Let $G = (V, E)$ be a graph, $|V| = n$, $|E| = m$. Denote by Δ the maximum degree of G. Throughout this paper, we assume that G is connected (any two nodes are connected by a path) and simple (no loops and multiple edges), and has no nodes of degree 1; we call such a graph *normalized*.

[1] This paper allows degenerate boxes, i.e., boxes that have dimension 1 with respect to one or more coordinate directions. Such degenerateness can be removed by adding additional grid-lines, which increases the volume of the drawing by a multiplicative constant.

2.1 Models for Three-Dimensional Drawings

In the definition of an orthogonal drawing there are no restrictions on the size of node-boxes. However, typically one wants nodes resembling points, therefore their boxes should be approximately squares respectively cubes.

To achieve such drawings, a number of models have been introduced for 2D box-drawings. The *Unlimited-growth model* imposes no restrictions on the dimensions of the nodes. The *Proportional-growth model* demands that the dimensions of a node may be only as big as needed for the number of incident edges. The *Kandinsky-model* imposes special conditions on nodes that are placed in the same horizontal or vertical range. See [3, 9] for details.

The Unlimited-growth model transfers directly to 3D. The Kandinsky-model can likewise be transferred, but this will not be explored in this paper. There is no direct equivalent of the Proportional-growth-model in 3D. We could demand that the surface area on each side of the node is only as large as needed for the number of incident edges, but this may lead to problems because the number of incident edges on the sides may not admit a suitable factorization. Instead, we use the *Degree-restricted model* or *dr-model*: Informally, a drawing is in the degree-restricted model if the surface area of node v is $\mathcal{O}(deg(v))$ for all nodes v. More precisely, a drawing is *in the* (c_1, c_2)*-dr model* if for the surface area of a node v is at most $c_1 deg(v) + c_2$. A drawing is said to be in the *degree-restricted model* if it is in the (c_1, c_2)-dr model for some constants c_1, c_2 that are independent of the input graph.

The previous heuristics for 3D-orthogonal drawings worked in the degree-restricted model: The algorithm by Papakostas and Tollis yields drawings in the (6,0)-dr model [13], and the drawing of K_n in an $\frac{n}{2} \times \frac{n}{2} \times \frac{n}{2}$-grid in [4] is in the (2,4)-dr model. Also, every 2D-drawing in the proportional-growth model is in the $(2, 2)$-dr model.

Drawings in the degree-restricted model may still be unsatisfactory, for example, a drawing where node v has a $1 \times 1 \times deg(v)$-box is in the degree-restricted model, but the elongated boxes would be considered unpleasant by many users. Therefore, we propose another, more stringent model, which we call *cube-model*: The box of v must be a cube whose surface area is proportional to the degree of v. Similarly one could define the *square-model* for 2D-orthogonal drawings, but to our knowledge this has not been used. For an algorithm where the aspect ratio of node boxes is at most 2, see [2].

3 Three Approaches to 3D-Orthogonal Box-Drawings

3.1 Approach I: Lifting Edges

In this section, we present the first approach to 3D-orthogonal drawings, *lifting edges*, which has been used in [4].

The idea is to start with a 2D-orthogonal drawing Γ which is *semi-valid*; by that we mean that no nodes overlap and no edge crosses a node, but edges may overlap each other. Nodes may be boxes or points.

Split this drawing Γ into a number of drawings $\Gamma_1, \ldots, \Gamma_\theta$, such that each drawing Γ_i is a valid orthogonal drawing. Here, *splitting a drawing* means to find a partition $E_1 \cup \ldots \cup E_\theta$ of the edges; Γ_i is then the restriction of Γ to all nodes and the edges in E_i. See also Figure 1, where we split a semi-valid drawing of K_8 into four valid drawings. This is similar to the construction used in [4].

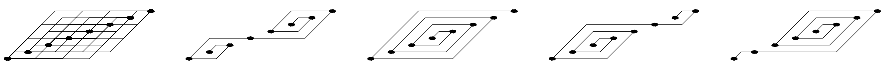

Fig. 1. We split a semi-valid drawing of K_8 into four crossing-free drawings.

Now place these crossing-free drawings $\Gamma_1, \ldots, \Gamma_\theta$ into θ z-planes $P_1, \ldots, P_\theta$. Extend every node to intersect all z-planes $P_1, \ldots, P_\theta$ to obtain an orthogonal box-drawing Γ', which has the same height and width as Γ, and depth θ.

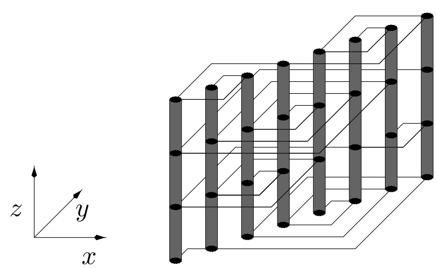

Fig. 2. Finishing the drawing started in Figure 1.

Using this method, drawings with excellent volume-bounds can be obtained. The first theorem results from the construction explained in Figure 1 and 2. For improvements, and the proof of the second theorem, we refer to [4].

Theorem 1. *[4] Every simple graph with n nodes, n even, has a box-drawing in an $n \times n \times \frac{n}{2}$-grid with 1 bend per edge.*

Theorem 2. *[4] Every simple graph with n nodes, $n = N^2$ a perfect square, has a box-drawing in a $2N \times 2N \times \frac{4}{3}N^3$-grid with 3 bends per edge.*

The main disadvantage of this approach is the size of the nodes. In the first construction, node v has a $1 \times 1 \times \frac{n}{2}$-box, which is in the degree-restricted model for the complete graph, but not for an arbitrary graph. In the second construction, node v has a $1 \times 1 \times \frac{4}{3}n^{1.5}$-box, which is not in the degree-restricted model.

Thus, even though this approach yields very small volume (in fact, the second construction matches asymptotically the lower bound [4]), its use is of rather theoretical nature to explore smallest-possible upper bounds.

3.2 Approach II: Lifting Half-Edges

In this section, we introduce a second approach to 3D-orthogonal box-drawings which we call *lifting half-edges*. Similar as in the previous approach, it starts with a 2D-orthogonal drawing and converts it into a 3D-orthogonal drawing. As opposed to the previous approach, it does so in a way that ensures that the drawings are in the degree-restricted model.

Again assume that we are given a semi-valid 2D-orthogonal drawing Γ. Split Γ into two drawings Γ_h and Γ_v, where Γ_h contains all horizontal edge-segments, while Γ_v contains all vertical edge-segments. Notice that Γ_h and Γ_v may have overlap, but they have no crossing. See Figure 3.

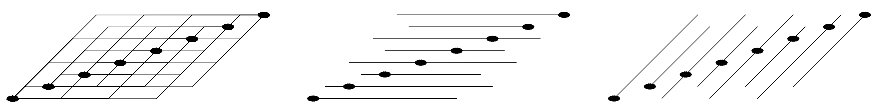

Fig. 3. We split a semi-valid drawing Γ of K_8 into Γ_h and Γ_v.

Then split Γ_v into θ_v drawings that have no overlap, and split Γ_h into θ_h drawings that have no overlap. As in the first approach, place each obtained drawing in z-plane of its own, and extend the nodes through all z-planes that contain an incident edge. At every bend of an edge in Γ, add a z-segment that connects the two endpoints of the horizontal and the vertical segment incident to this bend. Thus, if an edge had k bends in Γ, then it has $2k$ bends in the resulting three-dimensional drawing Γ'.

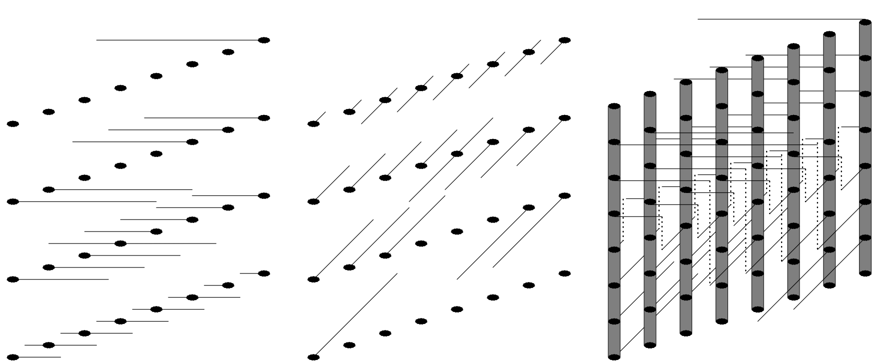

Fig. 4. Split Γ_h and Γ_v into four drawings each, extend nodes, and add z-segments. We show only a selected subset of the added z-segments.

The advantage of this approach lies in the fact that Γ_h and Γ_v have no crossings, and hence conflicts can be resolved much easier. However, care has to be taken which edge segment is assigned to which drawing, because the added z-segments must not cross. This is possible for a large class of input-drawings Γ.

We say that a semi-valid 2D-orthogonal drawing is *in general position* if no grid-line intersects more than one node. In particular, in a 2D-orthogonal drawing in general position, every edge has at least one bend, and no more than one bend per edge is needed [3]. For any such drawing, we can apply the half-edge lifting technique. Details are omitted.

Theorem 3. *Let Γ be a semi-valid 2D-orthogonal drawing in general position with exactly one bend per edge. Assume that Γ uses a $w \times h$-grid and at most θ edges overlap in any given place. Then there exists a 3D-orthogonal drawing in a $w \times h \times d$-grid, where $d \leq 2\theta$. The drawing is in the degree-restricted model.*

This theorem has far-reaching implications. For example, we can use it to generalize the interactive drawing results in [3], the results on orthogonal drawings with small area [2], and the results on incremental drawings with small area [2], because these drawings are in general position with one bend per edge.

In particular, using the semi-valid drawing used to create the 2D-drawing in an $\frac{m+n}{2} \times \frac{m+n}{2}$-grid in [2], we obtain the following results.

Corollary 1. *Every normalized graph has a 3D-orthogonal box-drawing in an $n \times n \times \Delta$-grid with 2 bends per edge. Node v is contained in a $1 \times 1 \times (\deg(v)/2+1)$-grid-box, so the drawing is in the $(2,6)$-dr model.*

This drawing improves on the result in [13] with respect to the volume ($\mathcal{O}(n^3)$ vs. $\mathcal{O}(m^3)$), and with respect to the surface area of nodes ((2,6)-dr model vs. (6,0)-dr model). (The drawings in [13] forbid degenerate node-boxes, but our results improve the volume even if we double all grid-dimensions to achieve non-degenerate node-boxes.)

One might criticize in the above construction that the nodes are highly degenerate in that they extend only in one direction. With a slightly different construction, the nodes become cubes, at the cost of an increase in volume.

Corollary 2. *Every normalized graph has a 3D-orthogonal box-drawing in a $(n+2\sqrt{nm}) \times (n+2\sqrt{nm}) \times \lceil\sqrt{\Delta}\rceil$-grid with 2 bends per edge. Node v is contained in a cube of side-length $2\lceil\sqrt{\deg(v)}/2\rceil$, so the drawing is in the cube-model.*

The main criticism of drawings created with the approach of lifting half-edges is that they are essentially two-dimensional. When looking at the drawing from the top, we see the input-drawing Γ. Moreover, a z-plane that is between the z-plane with the largest coordinate used for Γ_h and the z-plane with the smallest coordinate used for Γ_v intersects all edges and all nodes that have incident horizontal and vertical segments. Borrowing a term from computational geometry, one could call these drawings $2\frac{1}{2}$-dimensional. Whether such a drawing is advantageous or disadvantageous is debatable, but it is a puzzling question whether smaller drawings could be achieved by truly making use of the third dimension.

3.3 The Three-Phase Method

In this section, we explain a third approach to 3D-orthogonal drawings, which imitates the *three-phase method* for 2D [3].

In the first phase, *node placement*, nodes are drawn as points, not as boxes. In the second phase, *edge routing*, we assign bends to every edge. We continue to draw nodes as points, hence edges may overlap. In the third phase, *port assignment*, we replace each grid-plane by many grid-planes, and re-assign edges to ports of node-boxes such that all overlaps and crossings are removed.

The crux of the three-phase method is to find node placement and edge routing schemes that permit port assignment. For the three-phase method in 2D, a number of sufficient conditions have been found [3]. Unfortunately, they do not transfer easily to 3D, because we have to ensure additionally in 3D that there are no crossings between edges. We have found a set of sufficient conditions that ensure a crossing-free drawing. We state these conditions here, and explain the terms, and how to find the port assignment in the next few sections.

Condition 1 *We are given a node placement in line-free position and an edge routing using cube-routes. No two nodes coincide. No edge overlaps a node. Two edges may cross only it the two crossing segments attach to endpoints of the edges. Two edges may overlap only if the overlapping segments attach to a common endpoint of the two edges.*

Node Placement A node placement in a 3D-grid is said to be in *line-free position* if every grid-line intersects at most one node; it is said to be in *xy-general position* if every x-plane and every y-plane intersects at most one node, and to be in *general position* if every grid-plane intersects at most one node. The terms are defined similarly for an orthogonal drawing.

Any node placement in general position is also in xy-general position, and any node placement in xy-general position is also in line-free position. We can show the feasibility of port assignment (given that the other conditions are satisfied) for any node placement in line-free position. However, we managed to find an edge routing such that the other conditions are satisfied only in the case of a node placement in xy-general position or in general position.

Edge Routing In the three-phase method in 2D, each edge is routed with at most one bend; therefore there are at most two possible edge routings. In 3D, we allow two bends per edge. It is possible to draw each edge with one bend in the unlimited growth model [4], but we know of no such results in the degree-restricted model.

Allowing two bends per edge implies six possibilities of routing an edge, using the edges of the cube spanned by the two endpoints; we call these routes *cube-routes.* The three segments of an edge are called *x-segment*, *y-segment* and *z-segment.* If the two endpoints of an edge have one coordinate in common, then the cube degenerates to a rectangle. In this case, we place two bends at the same place (these will be expanded during port assignment). Note that the two endpoints cannot have two coordinates in common if the nodes are in line-free position.

We use three subclasses of these cube routes to ensure Condition 1:

- *directed z-routes*: Directed z-routes are those cube-routes for which the middle segment is the z-segment. We associate the two z-routes of an edge $e = (v, w)$ with a direction of e; thus $e = v \to w$ corresponds to the route that uses the x-line of v and the y-line of w.
- *color-routes*: Eades, Symvonis and Whitesides [8] gave an algorithm for 3D point-drawings which uses cube-routes. They restricted their attention to three out of the six routes, and associated them with colors. We call these routes the color-routes.
- *shortest-middle route*: The length of the x-segment, y-segment and z-segment is determined by the position of the endpoints. An edge is said to be routed using the shortest-middle route if the middle segment is the shortest segment, breaking ties arbitrarily.

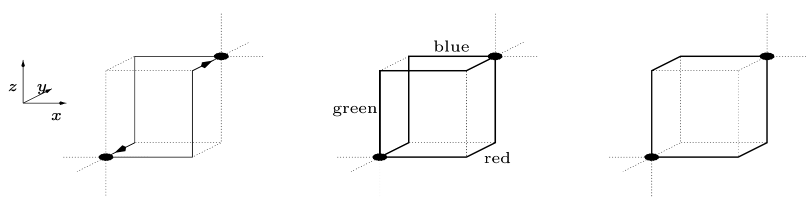

Fig. 5. The directed z-routes, the color-routes, and the shortest-middle routes.

The following lemmas are proved by straightforward case-analysis considering the grid-plane that contains an overlap or a crossing.

Lemma 1. *Let the node placement be in xy-general position, and let the edge routing be done with directed z-routes. Then Condition 1 is satisfied.*

Lemma 2. *Let the node placement be in general position, and let the edge routing be done with color-routes. Then Condition 1 is satisfied.*

Lemma 3. *Let the node placement be in general position, and let the edge routing be done with shortest-middle-routes. Then Condition 1 is satisfied.*

To analyze the edge routing, we introduce the following notation: For node v, define $A_x(v) = \max\{\#$ edges attaching on the $-x$-side of v, $\#$ edges attaching on the $+x$-side $v\}$; thus $A_x(v)$ is the number of x-ports needed at v. Similarly we define $A_y(v)$ and $A_z(v)$.

One can show that for any node placement there exists an edge routing using z-routes such that $A_x(v), A_y(v), A_z(v) \leq \lceil deg(v)/2 \rceil$ (see Lemma 3 in [3]). We leave as an open problem to find an edge routing that satisfies Condition 1 such that the bounds are roughly $A_x(v), A_y(v), A_z(v) \leq \lceil deg(v)/3 \rceil$.

Port Assignment In this section, we prove that port assignment is feasible if Condition 1 is satisfied.

Lemma 4. *Assume that we are given a node placement and edge routing that satisfies Condition 1. Then we can assign ports such that there is neither a crossing nor overlap.*

Proof. For every node v, choose arbitrary numbers $x_y(v), x_z(v), y_x(v), y_z(v), z_x(v), z_y(v)$ such that $x_z(v)y_z(v) \geq A_z(v)$, $x_y(v)z_x(v) \geq A_y(v)$ and $y_x(v)z_y(v) \geq A_x(v)$; and furthermore $x_y(v)+x_z(v) \geq 1$, $y_x(v)+y_z(v) \geq 1$ and $z_x(v)+z_y(v) \geq 1$. Good choices will be discussed later.

For every x-plane P_x after node placement, we add $\max_{v \in P_x}\{x_y(v)\}$ x-planes after P_x and $\max_{v \in P_x}\{x_z(v) - 1\}$ x-planes before P_x. For every y-plane P_y after node placement, we add $\max_{v \in P_y}\{y_x(v)\}$ y-planes after P_y and $\max_{v \in P_y}$ $\{y_z(v) - 1\}$ y-planes before P_y. For every z-plane P_z after node placement, we add $\max_{v \in P_z}\{z_x(v)\}$ z-planes after P_z and $\max_{v \in P_z}\{z_y(v) - 1\}$ z-planes before P_z. Here, "after" and "before" are taken with respect to the coordinate system.

Assume that node v was placed in the grid-planes P_x, P_y, P_z. Then we assign to v the grid-box that is the intersection of $x_z(v) + x_y(v)$ x-planes, $y_z(v) + y_x(v)$ y-planes, and $z_y(v) + z_x(v)$ z-planes; using the grid-planes before P_x, P_y, P_z, then P_x, P_y, P_z, and then the grid-planes after P_x, P_y, P_z. If two nodes were placed at different grid-points during node-placement, then their boxes are disjoint, because we added sufficiently many grid-planes around the point of each node.

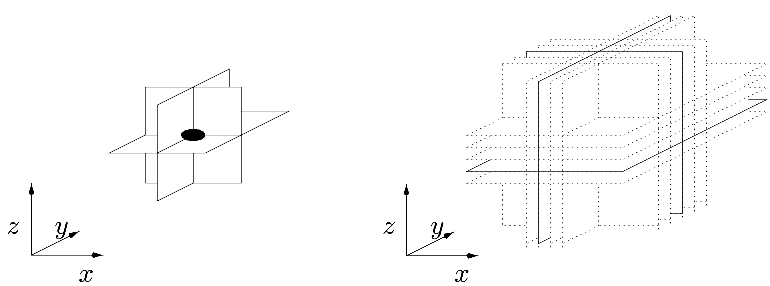

Fig. 6. We replace the point of v by the box that spans $x_y + x_z$ x-planes, $y_x + y_z$ y-planes and $z_x + z_y$ z-planes. In this example, $x_y = 2$, $x_z = 2$, $y_x = 2$, $y_z = 3$, $z_x = 3$ and $z_y = 2$; we show the added planes dashed.

Of the $z_x(v) + z_y(v)$ z-planes for node v, the $z_x(v)$ z-planes above P_z will be used to route edges that attach to v with an x-segment, while P_z and the $z_y(v) - 1$ z-planes below it will be used to route edges that attach to v with a y-segment. With this technique and by Condition 1, all crossings that occurred after edge routing are removed. See Figure 7 for an illustration of which ports are actually used at a node. By $x_z(v)y_z(v) \geq A_z(v)$, $x_y(v)z_x(v) \geq A_y(v)$ and $y_x(v)z_y(v) \geq A_x(v)$ there are sufficiently many z-ports, y-ports and x-ports, respectively.

Now we explain how to assign ports to edges attaching to the $-y$-side of v such that all overlaps are removed and no new crossings are introduced; port assignment for the other sides of v is done similarly. We group the edges attaching at the $-y$-side of v into four groups, depending on their direction of continuation. We assign sufficiently many $-y$-ports to each group such that no edges of two different groups could possibly cross. Then we sort the edges in each group by decreasing y-coordinate of the next bend, and assign them to a port of their group such that no two edges within one group cross. Details are omitted, see Figure 7 for some explanatory illustrations.

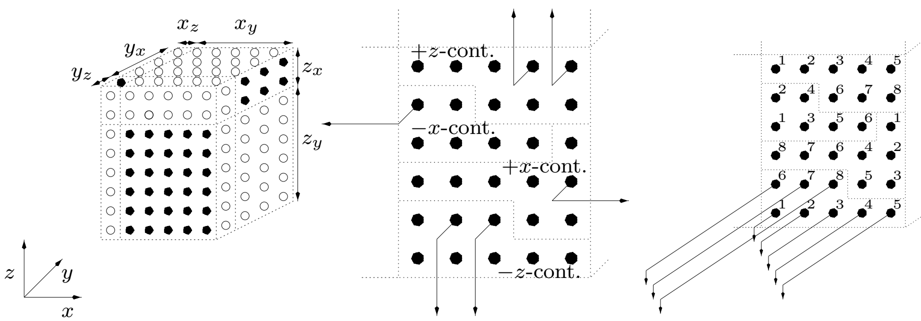

Fig. 7. On the left, we show the overall appearance of the node v, and indicate the used ports in black. On the middle and the right, we sketch the port assignment procedure: split the ports into four groups, and assign edges to ports in an appropriate order.

Lemma 5. *With a suitable choice of* $x_y(v), x_z(v), y_z(v), y_z(v), z_x(v)$*, and* $z_y(v)$*, the surface area of node* v *is at most* $12\deg(v) + 24$.[2]

Proof. Fix a node v and set $a_x = \max\{A_x(v), 1\}$, $a_y = \max\{A_y(v), 1\}$, $a_z = \max\{A_z(v), 1\}$. After possible renaming of coordinates, we may assume $a_x \leq a_y \leq a_z$. We will find integers x, y, z with $xy \geq a_z$, $xz \geq a_y$, and $zy \geq a_x$, and set $x_y(v) = x_z(v) = x \geq 1$, $y_z(v) = y_z(v) = y \geq 1$, and $z_x(v) = z_y(v) = z \geq 1$.

If $a_z \leq a_x a_y$ choose $x = \lceil \sqrt{a_y a_z / a_x} \rceil, y = \lceil \sqrt{a_x a_z / a_y} \rceil, z = \lceil \sqrt{a_x a_y / a_z} \rceil$. One can show that then the surface area of v is $4(xy+yz+zx) \leq 12(a_z+a_y+a_z)$.

If on the other hand $a_z > a_x a_y$, define $z = 1, y = a_x$ and $x = \lceil a_z/a_x \rceil > a_y$. By $xy + yz + zx < a_x + (a_z/a_x + 1) + a_x(a_z/a_x + 1) \leq 2(a_x + a_z) + 1$, the surface area of v is at most $8(a_x + a_z)$.

By $deg(v) \geq 2$, we have $a_x + a_y + a_z \leq A_x(v) + A_y(v) + A_z(v) + 2 \leq deg(v) + 2$, which finishes the proof.

[2] Our focus was on *that* the drawing is in the degree-restricted model, not on giving the smallest constants possible; therefore we chose the parameters for convenience of giving a simpler proof.

Results The main criticism of Approach II was that the resulting drawings are $2\frac{1}{2}$-dimensional, in that there exists a z-plane crossing all edges and nodes. Creating truly 3D-drawings is straightforward using the three-phase method. Place the nodes arbitrarily in general position, route the edges such that Condition 1 is satisfied (for example using color-routes), and then apply port assignment.

Using the node placement of [2], but assigning each node arbitrarily to a different z-plane, and routing the edges using directed z-routes, one obtains the following theorems.

Theorem 4. *Every normalized graph has a 3D-orthogonal box-drawing in general position in an $n \times n \times m$-grid with 2 bends per edge. Node v is contained in a $1 \times 1 \times (\deg(v)/2 + 1)$-grid-box, so the drawing is in the $(2, 6)$-dr model.*

Theorem 5. *Every normalized graph has a 3D-orthogonal box-drawing in general position in a grid of side-length $n + 2\sqrt{nm}$ with 2 bends per edge. Node v is contained in a cube of side-length $2\lceil\sqrt{\deg(v)}/2\rceil$, so the drawing is in the cube-model.*

However, these theorems are somewhat unsatisfactory, in that even for the smaller drawing we have a significant increase in the volume, from $\mathcal{O}(n^3)$ of Section 3.1 to $\mathcal{O}(n^2 m)$. So the question arises whether there exists a truly 3D-drawing with volume $\mathcal{O}(n^3)$. Note that the term *truly 3D-drawing* is not precisely defined. One could define it as a drawing with $o(n)$ nodes per grid-plane, or as a drawing with $\mathcal{O}(f(n))$ nodes per grid-plane, where one could argue the case for any of the functions $f(n) = \sqrt{n}$, $f(n) = n^{1/3}$, and $f(n) = 1$.

Our results are in this last and most stringent definition of a truly-3D drawing, which means drawings in general position. We can show that under these conditions, no drawing of volume better than $\mathcal{O}(m^2)$ is possible, thus we are optimal for graphs with $m = \theta(n^2)$.

Theorem 6. *Any 3D-drawing in general position has volume $\Omega(\max\{n^3, m^2\})$.*

Proof. Assume that we have a drawing with at most one node per grid-plane. Number the nodes as $1, \ldots, n$, and denote by x_i, y_i, z_i and d_i the dimensions and the degree of node i. Because no grid-plane intersects two nodes, the drawing must have width $\geq \sum x_i$, depth $\geq \sum y_i$ and height $\geq \sum z_i$ (all summations are from 1 to n). Furthermore, we have $2(x_i y_i + y_i z_i + z_i x_i) \geq d_i$ and $\sum d_i = 2m$.

By $x_i, y_i, z_i \geq 1$, the lower bound of $\Omega(n^3)$ follows immediately. For the lower bound of $\Omega(m^2)$, let $d_i^x = \min\{d_i/6, y_i z_i\}$, $d_i^y = \min\{d_i/6, x_i z_i\}$, and $d_i^z = \min\{d_i/6, x_i y_i\}$. In at least one of the three directions, the minimum must be attained at $d_i/6$, therefore $\sum(d_i^x + d_i^y + d_i^z) \geq \sum(d_i/6) = \frac{1}{3}m$ and (after possible renaming of the coordinate directions) $\sum d_i^z \geq \frac{1}{9}m$.

By $d_i^z \leq d_i \leq n$ we have $d_i^z/\sqrt{n} \leq \sqrt{d_i^z}$. Therefore $\sum \sqrt{d_i^z} \geq \frac{1}{9}m/\sqrt{n}$, and

$$(\sum x_i)(\sum y_i)(\sum z_i) \geq (\sum \sqrt{x_i}\sqrt{y_i})^2 n \geq (\sum \sqrt{d_i^z})^2 n \geq (\frac{1}{9}m/\sqrt{n})^2 n$$

by the Cauchy-Schwartz inequality, which yields the result.

We conjecture that this lower bound can be improved to $\Omega(n^2 m)$.

4 Conclusion and Open Problems

In this paper, we studied three-dimensional orthogonal drawings of graphs of arbitrarily high degrees. We presented three approaches, and obtained, among others, the following results:

Every normalized graph has a drawing in an $n \times n \times \Delta$-grid with 2 bends per edge. Thus, the resulting volume is $\mathcal{O}(n^3)$, which for the complete graph is a square-root factor better than the result of $\mathcal{O}(m^3)$ of Papakostas and Tollis [13]. Also, this result matches the $\mathcal{O}(n^3)$ construction for the complete graph [4], but is in the degree-restricted model for all graphs, not only graphs with degrees in $\theta(n)$.

This result falls short of the lower bound of $\Omega(n^{2.5})$ and the construction with volume $\mathcal{O}(n^{2.5})$ presented in [4]. However, the latter construction is not in the degree-restricted model, not even for the complete graph.

> Open Problem: In the degree-restricted model, is there a drawing of volume $\mathcal{O}(n^{2,5})$, or can the lower bound be raised to $\mathcal{O}(n^3)$?

Every normalized graph has an $(n+2\sqrt{nm}) \times (n+2\sqrt{nm}) \times \lceil\sqrt{\Delta}\rceil$-drawing with two bends per edge in the cube-model. This is, to our knowledge, the first result in the cube model that also maintains a small surface area of the cube, and a small overall volume. However, this result comes at an increase in the volume to $\mathcal{O}(n^2 m)$.

> Open Problem: Is there a drawing in the cube-model with volume $\mathcal{O}(n^3)$?

Every normalized graph G has an $n \times n \times m$-drawing in general position with two bends per edge. This result does not match the lower bound of $\Omega(\max\{n^3, m^2\})$, but we conjecture that no drawing can achieve an asymptotically smaller volume.

> Open Problem: For drawings in general position, is there a drawing of volume $\mathcal{O}(\max\{n^3, m^2\})$, or can the lower bound be raised to $\mathcal{O}(n^2 m)$?

Also, while having all nodes in one place is too two-dimensional, having each node in a separate grid-plane seems a waste. What is good middle ground? Do there exist small drawings with $\theta(n^{1/3})$ nodes in every grid-plane? What volume-bounds can be achieved?

References

[1] T. Biedl and G. Kant. A better heuristic for orthogonal graph drawings. *Computational Geometry: Theory and Applications*, 9:159–180, 1998.

[2] T. Biedl and M. Kaufmann. Area-efficient static and incremental graph drawings. In *5th European Symposium on Algorithms*, volume 1284 of *Lecture Notes in Computer Science*, pages 37–52. Springer-Verlag, 1997.

[3] T. Biedl, B. Madden, and I. Tollis. The three-phase method: A unified approach to orthogonal graph drawing. In DiBattista [7], pages 391–402.

[4] T. Biedl, T. Shermer, S. Whitesides, and S. Wismath. Orthogonal 3-D graph drawing. In DiBattista [7], pages 76–86.

[5] G. Di Battista, P. Eades, R. Tamassia, and I. Tollis. Algorithms for drawing graphs: an annotated bibliography. *Comp. Geometry: Theory and Applications*, 4(5):235–282, 1994.

[6] G. Di Battista, A. Garg, G. Liotta, R. Tamassia, E. Tassinari, and F. Vargiu. An experimental comparison of four graph drawing algorithms. *Computational Geometry: Theory and Applications*, 7(5-6), 1997.

[7] G. DiBattista, editor. *Symposium on Graph Drawing 97*, volume 1353 of *Lecture Notes in Computer Science*. Springer-Verlag, 1998.

[8] P. Eades, A. Symvonis, and S. Whitesides. Two algorithms for three dimensional orthogonal graph drawing. In S. North, editor. *Symposium on Graph Drawing 96*, volume 1190 of *Lecture Notes in Computer Science*. Springer-Verlag, 1997, pp. 139-154.

[9] U. Fößmeier and M. Kaufmann. Drawing high degree graphs with low bend numbers. In F. Brandenburg, editor. *Symposium on Graph Drawing 95*, volume 1027 of *Lecture Notes in Computer Science*. Springer-Verlag, 1996, pages 254–266.

[10] U. Fößmeier and M. Kaufmann. Algorithms and area bounds for nonplanar orthogonal drawings. In DiBattista [7], pages 134–145.

[11] A. Papakostas and I. Tollis. High-degree orthogonal drawings with small grid-size and few bends. In *5th Workshop on Algorithms and Data Structures*, volume 1272 of *Lecture Notes in Computer Science*, pages 354–367. Springer-Verlag, 1997.

[12] A. Papakostas and I. Tollis. Algorithms for area-efficient orthogonal drawings. *Computational Geometry: Theory and Applications*, 9:83–110, 1998.

[13] A. Papakostas and I. Tollis. Incremental orthogonal graph drawing in three dimensions. In DiBattista [7], pages 52–63.

[14] A. Rosenberg. Three-dimensional VLSI: A case study. *Journal of the Association of Computing Machinery*, 30(3):397–416, 1983.

[15] R. Tamassia. On embedding a graph in the grid with the minimum number of bends. *SIAM J. Computing*, 16(3):421–444, 1987.