

Nearly Optimal Three Dimensional Layout of Hypercube Networks ^{*}

Tiziana Calamoneri Annalisa Massini

Department of Computer Science,
University of Rome “La Sapienza”
via Salaria 113, 00198 Roma, Italy.
e-mail:{calamo, massini}@di.uniroma1.it

Abstract. In this paper we consider the three-dimensional layout of hypercube networks. Namely, we study the problem of laying hypercube networks out on the three-dimensional grid with the properties that all nodes are represented as rectangular slices and lie on two opposite sides of the bounding box of the layout volume. We present both a lower bound and a layout method providing an upper bound on the layout volume and the maximum wire-length of the hypercube network.

Keywords: Hypercube Network, Three Dimensional Layout, VLSI layout volume.

1 Introduction

The layout of interconnection networks has important cost and performance implications for multiprocessors and parallel and distributed systems based on components. Thus, there is currently renewed interest in finding efficient VLSI layouts for various interconnection networks (see for instance [3, 5, 11, 16–19]).

Recent hardware advances have allowed three-dimensional circuits to have a cost low enough to make them commonly available. For this reason three-dimensional layouts of interconnection topologies on rectilinear grids are becoming of wide interest. The importance of efficiently representing interconnection networks in three dimensions has already been stated in the 80’s by Rosenberg

^{*} A preliminary version of this paper appeared on the Proc. of the 11th International Symposium on Graph Drawing (GD 2003).

[14], and the most relevant aims are to shorten wires and to save in material (i.e. to minimize volume).

The three-dimensional layout problem is related both to the study of the VLSI layout problem for integrated circuits and to the study of algorithms for drawing graphs. Indeed, the tie between VLSI layout studies and theoretical graph drawing is very strong since laying out a network on a grid is equivalent to orthogonally drawing the underlying graph, except for the fact that wire crossings are allowed. Nevertheless, almost all the results known in the literature about three-dimensional grid drawing of graphs are valid for very general graphs (see, the proceedings of the past editions of the *Graph Drawing* Symposia and, for instance, [1, 6, 15]) and therefore they do not work efficiently for structured and regular graphs such as the most commonly used interconnection topologies. In particular, either the produced volume is too large (in order of magnitude or in the multiplicative constant), or the length of each wire is not kept of the same order of magnitude and sufficiently small, or the shape of nodes is not feasible to represent a processor, a switch, etc. Hence, usually, lower and upper bounds on the layout volume of specific interconnection topologies must be established and the layout technique is the more preferred the more these bounds are closer; of course, the layout method should try to highlight all the interesting properties of the network and to make as fast as possible the flow through it by shortening the maximum wire length.

The hypercube network has been widely considered as a network for parallel computing because of its desirable parameters, e.g. regularity, logarithmic diameter, fault tolerance. A further parameter distinguishing ‘good’ networks is whether they can be laid out compactly in the orthogonal grid. This issue has widely been studied in two dimensions [2, 7, 8, 12].

In this paper we study the three-dimensional layout of hypercube networks with the property that nodes lie on two planes, corresponding to opposite sides of the bounding box, partitioned into two equally sized sets, and they are represented as rectangular slices. The model of keeping all nodes on the surface of the bounding box is motivated by two main reasons: (i) it decreases the problems due to power consumption and resulting heat generation and (ii) it makes easier the connection of the topologies to other part of the architecture, as the wires are necessarily connected to nodes. We present upper and lower bounds for the layout volume of hypercube networks that are very close, although not coinciding (in fact they differ by a factor of $\log^{1/2} N$). The maximum wire-length of the presented layout is $O(\log N \sqrt{N})$.

The rest of this paper is organized as follows. In Section 2 we give some preliminary definitions and recall some useful previous results. Section 3 is the core of the work and we present both a lower bound and a layout method providing an upper bound on the layout volume of the hypercube network. Finally, in Section 4 we address some conclusions and list some interesting open problems.

2 Preliminaries

In this section we recall some useful definitions and state some preliminary results.

A *hypercube* Q_n of dimension n is an interconnection network with $N = 2^n$ nodes, each labeled with an n -bit binary string. Two nodes are connected by an edge if their labels differ in exactly one bit.

It is easy to see that a Q_n can be constructed by considering two copies of Q_{n-1} , adding an edge for each pair of equally labeled nodes and completing the labels of nodes by adding one bit to the left of the string (0 in the first copy and 1 in the second copy of Q_{n-1}). In the following we will call *homonym nodes*

in Q_n two nodes connected by this inductive step, i.e. two nodes whose labels differ only in the leftmost bit.

By iterating this recursive construction, we deduce that Q_n can be constructed starting from 2^{n-4} copies of Q_4 , adding some edges and completing the labels. In the following we will extensively use this observation and we will clarify the reason why Q_4 is chosen.

Another interesting property of Q_n we will exploit in the following is its bipartiteness; a feasible partition of the nodes can be obtained by dividing nodes having an even number of 0s in their label from nodes having an odd number of 0s in their label. The bipartiteness property is particularly evident when Q_n is visualized as a lattice (see Fig. 2).

Finally, observe that the hypercube Q_n is a regular graph of degree n , and when we want to lay the hypercube out on a three-dimensional grid we can not represent the nodes with dots if the degree n is greater than 6, so we must use an alternative representation.

Given a certain *shape* (e.g. box, prism, cube, rectangle, segment, dot, etc.), a *three-dimensional grid shape-layout* of an interconnection network \mathcal{G} is a mapping of \mathcal{G} to the three-dimensional grid such that nodes are mapped to portions of the grid of the given shape and edges are mapped to independent grid-paths satisfying the following conditions:

- distinct grid-paths are edge-disjoint (then at most three paths can cross at a grid-node);
- grid-paths that share an intermediate grid-node must cross at that node (that is ‘knock-knee’ paths [10] are not allowed);
- a grid-path may touch no mapped node, except at its end-points.

If the layout can be enclosed in a $H \times W \times L$ three-dimensional box, its *volume* is the product $H \times W \times L$. Observe that not all layout models avoid knock-knees

(e.g. they are allowed in [16].) We decide to forbid them in order to obtain a more realistic physical model.

In this paper we will consider the rectangular shape. In particular, we represent nodes as rectangular slices with area proportional to their degree (dimension $1 \times \text{degree}$). Furthermore we impose that all nodes lie on two planes, corresponding to opposite sides of the bounding box, partitioned into two equally sized sets

A k -channel routing is introduced in [13] to describe the way of connecting two sets of nodes of a graph:

Definition 1. A k -channel routing involves a bidimensional grid and two sets S and S' each consisting of k nodes to be connected by a 1-1 function. S and S' are arranged onto opposite sides of the grid.

The grid involved in any k -channel routing is not greater than $(k + 1) \times (\frac{3}{2}k + 2)$ and S and S' lie on the shortest sides [13]. See Fig. 1 for an example.




Fig. 1. A channel routing with $k = 6$ requiring a grid of maximum area, i.e. $(k + 1) \times (3/2k + 2)$.

We define the $a \times b$ -3D double channel routing as a slight modification of the notion of k -3D channel routing introduced in [4], that is the three-dimensional version of the k -channel routing.




Fig. 2. Q_5 represented as a lattice.

Definition 2. An $a \times b$ -3D double channel routing involves a three-dimensional grid (the channel) and two sets S and S' , each containing $a \cdot b$ nodes, arranged on the grid nodes of opposite faces of dimension $a \times b$, to be connected by a 1-1 function f . Function f associates to a node (x, y) of S a node (x', y') of S' such that $x' = g(x)$ and $y' = h(y)$, where functions g and h are two-dimensional a - and b -channel routings, respectively.

In the following, for the sake of brevity, when no confusion arises, we refer to the $a \times b$ -3D double channel routing as *3D channel routing*.

The following theorem states the necessary volume for any 3D channel routing.

Theorem 1. A three-dimensional grid of size $(a + 1) \times (b + 1) \times \max\{\frac{3}{2}a + 2, \frac{3}{2}b + 2\}$ is enough to realize a 3D channel routing.

Proof. This proof is very similar to the proof of Theorem 14 of [4], but we detail it here for the sake of completeness.

Project the three-dimensional grid of the 3D channel routing on plane xz (see Fig. 3). It is easy to see that function g maps rows of S in rows of S' and it can

be considered as a two-dimensional a -channel routing on plane xz . Therefore, a $(a+1) \times (\frac{3}{2}a+2)$ two-dimensional grid is enough to realize such a channel routing. Lines drawn on plane xz to represent the channel routing relative to function g become (bent) planes in three dimensions (see Fig. 3). Each of such planes has on opposite horizontal sides a length b row x of S and its corresponding row $g(x)$ of S' . On each such plane we have to realize a two-dimensional b -channel routing given by function h , and this requires a length of the plane of at least $\frac{3}{2}b + 2$. It follows that the size of a three-dimensional grid containing the 3D channel routing has dimensions at most $(a+1) \times (b+1) \times \max\{\frac{3}{2}a+2, \frac{3}{2}b+2\}$.




Fig. 3. A 3D channel routing.

3 Three Dimensional Layout of Hypercube Network

In this section we discuss the problem of laying the hypercube network out; in particular, we determine lower and upper bounds for the volume.

3.1 Lower Bound

In order to compute a lower bound on the layout volume of the hypercube network, we can use the general formula provided in [4]:

lower bound on the layout volume of $\mathcal{H} \geq (\sqrt{BW(\mathcal{H})} - 1)^3$

where \mathcal{H} is a general network and $BW(\mathcal{H})$ is its bisection width (i.e. the smallest number of edges whose removal partitions \mathcal{H} into two disjoint subgraphs, each containing half of the nodes of \mathcal{H}).

It is well known that the bisection width of Q_n is $\frac{N}{2}$ [9], so a lower bound for the layout volume for Q_n is $\Omega(N^{\frac{3}{2}})$. This lower bound does not take into account the fact that, when $n > 6$, nodes cannot be represented as dots (i.e. with null area), but it is necessary to use a space (either area or volume) proportional to the degree for each of them. Furthermore, the previous formula assumes that the nodes can be distributed anywhere inside the volume of the layout. On the contrary, our aim is to find a three-dimensional layout of Q_n with some special properties; under these assumptions, we can better specify the value of the lower bound on the volume by exploiting a proof technique presented in [14] for one-active-level layouts: we obtain the value $\Omega(N^{3/2} \log^{1/2} N)$ as follows.

Theorem 2. *A lower bound for a 3D grid layout of Q_n is $\Omega(N^{3/2} \log^{1/2} N)$, if the layout has the following properties:*

- *all nodes are represented as rectangular slices with area proportional to their degree,*
- *all nodes lie on two planes, corresponding to opposite sides of the bounding box,*
- *nodes are partitioned into two equally sized sets.*

Proof. Let W, L and H be the dimensions of the bounding box of any three-dimensional layout of Q_n with nodes partitioned into two equally sized sets and laid out on two opposite sides. Without loss of generality, assume the nodes lie on the sides of dimension $W \times L$ of the bounding box and that $W \geq L$. Since

each side contains $\frac{N}{2}$ nodes, each of degree $\log N$, it is straightforward to see that $W \times L = \Omega(\frac{N}{2} \log N)$. Two cases can occur, according to whether $H \geq W$ or not.

If $H \geq W$, then $H \geq (\frac{N}{2} \log N)^{1/2}$. It follows that the volume must be $\Omega((\frac{N}{2})^{3/2} \log^{3/2} N)$, which is even greater than we claim.

If H is not the maximum dimension of the bounding box, we can consider a plane π orthogonal to the sides of the bounding box containing nodes and parallel to the pair of parallel sides of the bounding box of dimension $L \times H$. The intersection of the bounding box with plane π has area $L \times H$. We want to place π in such a way that π cuts the layout into two portions, and partitions the set of nodes into two equally sized sets leaving $\frac{N}{2}$ nodes to its left and $\frac{N}{2}$ nodes to its right. Since the nodes are not represented as dots but are rectangular slices, π could intersect some nodes. To avoid this, we can stretch the layout, lengthen dimension W to $W' \geq W$ and include the intersected nodes in one of the two sides. In this way π partitions the set of nodes into two equally sized sets (see Fig. 4). Observe that, since nodes are represented by slices having area proportional to their degree, each dimension of each slice is $O(\log N)$. It follows that $W' = W + O(\log N)$. In view of this equality, the hypotheses that $W' \geq W \geq L$ and that plane π cuts a number of edges at least equal to the bisection width of Q_n , i.e. $\frac{N}{2}$, it follows that the area of the portion of π inside the bounding box is at least $\frac{N}{2}$, i.e. $H \times L \geq \frac{N}{2}$. Since $L \geq N^{1/2} \log^{1/2} N$ we have that the volume is $\Omega(N^{3/2} \log^{1/2} N)$.

3.2 Layout of Q_n

In this subsection we provide a method to lay the hypercube network out for each $n \geq 4$ (indeed, for $n = 3$ the layout is trivial and the volume is equal to 1).

In our layout we choose to represent nodes as rectangular slices of area $1 \times \lceil \frac{n}{2} - 1 \rceil$ (see Fig. 5). Each dot at integer coordinates on the slice is a pin where




Fig. 4. A 3D layout where dimension W is stretched in order to guarantee that plane π cuts the set of nodes into equally sized sets.

one wire is connected. Observe that when n is odd, the last row of each slice has only one active pin, whereas when n is even, all pins of the slice are active (see Fig. 5). In the following we will call *columns* a and b the two vertical lines of pins on each slice.




Fig. 5. Slices representing nodes of even and odd degree n ; the white dot is not active.

As we highlighted in the previous section, the basic building block of our result is the layout of Q_4 , so we show it, before generalizing the construction to Q_n .

Since Q_4 has 2^4 nodes and its degree is 4, we represent each node as a unit area slice and we arrange these slices in two rectangles of dimension 3×7 , as

shown in Fig. 6. Our layout puts nodes onto opposite sides of the bounding box, and all edges are routed inside. Hence, we exploit the property that Q_4 is bipartite and put on one side (let it be the *upper side*) all nodes having an even number of 0s in their label, and on the opposite side (let it be the *lower side*) the remaining nodes (see Fig. 6).




Fig. 6. Placement of Q_4 's nodes on a three-dimensional grid.




Fig. 7. Q_4 represented as a butterfly-like network. Grey nodes are the same as those in the leftmost node column.

Furthermore, we want to use a 3D channel routing to arrange all edges on the grid. To this aim, we need to assign to each node a position in such a way that the rules of the 3D channel routing (cf. Def. 2) are kept; namely, we have to

map columns of the upper side to columns of the lower side. In other words, all nodes of each column i on the upper side must be mapped to nodes of a column j on the lower side. The way to arrange nodes on the rectangle derives from the representation of Q_4 as a Butterfly-like network, as shown in Fig. 7. More precisely, we can map the leftmost node column of Fig. 7 to the leftmost slices of the upper side (see Fig. 6), the second node column in Fig. 7 in the leftmost slices of the lower side, and so on. This arrangement allows us to use pins in column a (respectively, b) to connect edges in the odd (respectively even) numbered edge-stages of the representation in Fig. 7. Such an assignment respects the rules of a 3D channel routing, and a possible final three-dimensional layout is sketched in Fig. 8, where the 4-channel routings inside the bent planes are not drawn for the sake of clarity. The volume of the built layout for Q_4 is $8 \times 4 \times 4$.




Fig. 8. 3D layout of Q_4 , where the 4-channel routings inside the bent planes are not drawn.

In order to generalize our layout to each $n \geq 4$, we exploit the layout of Q_4 and the recursive definition of Q_n . Namely, given a hypercube Q_n , we will arrange it on the three-dimensional grid as 2^{n-4} copies of hypercubes of dimension 4 plus a set of edges.

Claim. The set of edges connecting the 2^{n-4} copies of hypercubes of dimension 4 constituting a Q_n is a 3D channel routing.

Proof. We give the construction and the proof of our claim by induction on n . Our layout of Q_n is built as the arrangement of the two copies of the layout of Q_{n-1} one beside the other, according to some rules specified in the following. This implies that all nodes are arranged in rows both on the upper and the lower side (see Fig. 9).

Base. The claim is trivially true when $n = 4$.




Fig. 9. Arrangement in rows of nodes, where the lowest pins of each slice are highlighted.

Inductive Hypothesis. Q_r , $4 \leq r < n$, can be arranged on the grid as 2^{r-4} copies of Q_4 s plus a set of edges constituting a 3D channel routing.

Inductive Step. Q_n can be obtained by considering two copies of Q_{n-1} and connecting with an edge each pair of homonym nodes. In view of the inductive hypothesis, the claim is true for Q_{n-1} . So, we can represent Q_n on the three-dimensional grid by putting two copies of 3D layouts of Q_{n-1} s one beside the other. Since the copies of Q_{n-1} s consist of 2^{n-5} copies of Q_4 s each, by the inductive hypothesis, Q_n will consist of 2^{n-4} copies of Q_4 s. It remains to prove that the edges connecting homonym pairs plus the existing 3D channel routings inside the two copies of Q_{n-1} constitute a new 3D channel routing.

Let us distinguish two different cases according to the parity of n :




Fig. 10. Construction of Q_n starting from two copies of Q_{n-1} s; a. n odd; b. n even. Double circled nodes represent homonym nodes in the two copies of Q_{n-1} s. Simply circled nodes are homonym nodes of those in the upper part of the layout at the previous step, i.e. in the two copies of Q_{n-2} s inside the same copy of Q_{n-1} .

– n **odd**

Arrange the 3D layouts of Q_{n-1} s one beside the other but vertically flipping one of them, in such a way that homonym nodes lie on opposite sides of the bounding box (see Fig. 10a). Observe that the flipping operation can be performed by simply swapping upper and lower sides, in view of the butterfly-like structure of the hypercube.

The set of edges connecting the copies of Q_4 s is a 3D channel routing. Indeed, it is constituted by a set of edges inside the two copies of Q_{n-1} s and by a set

of edges connecting homonym nodes of the two copies of Q_{n-1} s. The first set forms two disjoint 3D channel routings by the inductive hypothesis.

The second set connects homonym nodes located on the upper side of the left (right) Q_{n-1} with nodes located on the lower side of the right (left) Q_{n-1} (for an example, see nodes highlighted in Fig. 10a). Each node has one more active pin on column a and it corresponds to the edge that must be still laid out. As all nodes are arranged in rows by construction, the second set is completely disjoint from the first one and it has even stronger properties than a 3D channel routing, namely it ordinally connects nodes of the same pin row with nodes on the same pin row and nodes on the same pin column with nodes on the same pin column.

From the previous observations it follows that these two 3D channel routings together form a unique 3D channel routing.

– **n even**

In this case each node has both columns constituted by the same number of pins and the lower pin in column b corresponds to the edges to be still laid out (see Fig. 5). The lower pin in column a corresponds to edges laid out in the previous step (when connecting two copies of Q_{n-2} s in a Q_{n-1}). Hence, when we route the edges connecting homonym nodes using the lower pins of column b , we must take into account the existing 3D channel routing used by the lower pins of column a . Again, arrange the 3D layouts of Q_{n-1} s one beside the other, flipping both vertically and horizontally one of them, as shown in Fig. 10b.

This arrangement guarantees that the rules of the 3D channel routing are respected by the set of all edges connecting the copies of Q_4 s. Actually, this set is constituted by a set of edges inside the two copies of Q_{n-1} s and a set of edges connecting homonym nodes of the two copies of Q_{n-1} s.

The first set forms two disjoint 3D channel routings by the inductive hypothesis; inside this set, we want to highlight the subset of edges outgoing from the lower pin of column a of each node. In view of the 3D channel routing of the previous step $n - 1$ (odd), they connect each node in general position (i, j) of the upper face, left (right) side of each copy of Q_{n-1} with its homonym in position (i, j) of the lower face, right (left) side of the same copy (for an example, see nodes indicated by the arrows in Fig. 10b).

Now consider the second set: it connects each node in general position (i, j) of the upper face, left (right) side of the first copy, with its homonym in position (i, j) of the lower face, right (left) side of the second copy by means of the lower pin of column b (for an example, see double circled nodes in Fig. 10b). So all connections coming out from the lower pins of the slices go to the same pin column and we have a 3D channel routing.

3.3 Upper Bound

Starting from the construction of the layout of Q_n presented in the previous subsection, we deduce an upper bound on the volume of the layout of Q_n .

Theorem 3. *There exists a three-dimensional layout of Q_n having volume $\Theta(N^{3/2} \log N)$.*

Proof. Let W_n, L_n and H_n be the three dimensions of the bounding box of the layout of Q_n presented in the previous subsection. We define some recurrence equations in order to compute these values for each $n \geq 4$.

If $n = 4$ we have already obtained $W_4 = 8, L_4 = 4$ and $H_4 = 4$ (see Fig. 8), and we use these values as base of the recurrence.

As our layout method is conceived, the equations are different according to whether n is odd or even.

If n is **odd**, we have

W_n : The width of the layout of Q_n is the same as the width of Q_{n-1} (see Fig. 10a), hence $W_n = W_{n-1}$.

L_n : $L_n = 2L_{n-1} + \sqrt{N/2}$ as we use two copies of Q_{n-1} s and we need to lengthen the slice columns by one. The contribution of $\sqrt{N/2}$ follows from the way of positioning the slices.

H_n : The height is increased by the 3D channel routing introduced to connect the two copies. As this 3D channel routing has $a = b = \sqrt{N/2}$, its contribution to the height is $\frac{3}{2}\sqrt{N/2} + 2$, hence $H_n = H_{n-1} + \frac{3}{2}\sqrt{N/2} + 2$.

If n is **even**, we have

W_n : $W_n = 2W_{n-1}$ as we use two copies of Q_{n-1} s.

L_n : The length of the layout of Q_n is the same as the length of Q_{n-1} (see Fig. 10b), hence $L_n = L_{n-1}$. Indeed, the newly inserted edges do not introduce any new contribution to this dimension.

H_n : Also in this case the height is increased by the 3D channel routing introduced to connect the two copies and $a = \sqrt{N}$ and $b = \sqrt{N}/2$. Nevertheless, the contribution to the height is due only to the connections outgoing from the lower pins of columns b in the slices. Indeed, the contribution of the other connections has been already computed in H_{n-1} . Hence, we have to add only the difference between the contributions of the two 3D channel routings, i.e. $H_n = H_{n-1} + \frac{3}{2}\sqrt{N/2}(\sqrt{2} - 1)$.

By solving the previous recurrence equations, we get

$$W_n = \Theta(\sqrt{N})$$

$$L_n = \Theta(\sqrt{N} \log N)$$

$$H_n = \Theta(\sqrt{N})$$

both when n is even and when n is odd.

It follows that the volume is $\Theta(N^{\frac{3}{2}} \log N)$.

Observe that this upper bound differs from the lower bound only by a $\log^{1/2} N$ factor.

Lemma 1. *The 3D layout of Q_n presented in the previous section has maximum wire length $O(\sqrt{N} \log N)$.*

Proof. In order to compute the maximum wire length, let us consider the worst case in the 3D channel routing inside the layout.

As highlighted in Fig. 11a, each plane constituting the 3D channel routing has length at most $H_n + 2L_n - 1$. Inside this plane (see Fig. 11b), the longest wire-length is $H_n + 2L_n - 1 + 2W_n - 1$.




Fig. 11. a. The longest plane of the 3D channel routing inside the layout of Q_n and b. the longest wire inside the same plane.

From this formula and from the values of H_n , W_n and L_n found in the previous proof, the claim follows.

Observe that in the two-dimensional layout of Q_n , the maximum wire length is $\Theta(N)$ [8], so this 3D layout saves both in material and in wire-length.

We conclude this subsection by summarizing all the results in the following theorem:

Theorem 4. *The method described in the previous subsection provides a 3D grid layout of Q_n in $\Omega(N^{3/2} \log N)$ volume, with $O(\sqrt{N} \log N)$ maximum wire-length, if the layout has the following properties:*

- *all nodes are represented as rectangular slices with area proportional to their degree,*
- *all nodes lie on two planes, corresponding to opposite sides of the bounding box,*
- *nodes are partitioned into two equally sized sets.*

4 Conclusions and Open Problems

In this paper we have considered the three-dimensional layout of hypercube networks. Namely, we have studied the problem of laying out hypercube networks on the three-dimensional grid with the properties that all nodes are represented as rectangular slices and lie on two opposite sides of the bounding box of the layout volume. We have presented upper and lower bounds on the volume that are very close, although not coinciding (in fact they differ by a factor of $\log^{1/2} N$). The maximum wire-length of the presented layout is $O(\log N \sqrt{N})$.

Two open problems arise from this work. The first problem derives from the differing upper and lower bounds, and consists of understanding which bound is tight. Secondly, it would be interesting to study the more general 3D layout problem, i.e. the problem where nodes are not constrained to be on two opposite sides of the bounding box. In this case, it is possible to consider also other shapes for nodes (i.e. boxes). Concerning this problem, note that it is quite easy to find an $O(N^{3/2})$ volume 3D layout of the CCC_n network (i.e. the network obtained from Q_n by substituting each Q_n 's node with a cycle of length n) from the technique used in [4]. We have wondered if it is possible to deduce a 3D layout of Q_n from the layout of CCC_n with the same volume exploiting the definition

of CCC_n . This does not seem possible, because the nodes constituting the cycles are not clustered in the layout.

References

- [1] T. Biedl, T. Thiele, and D.R. Wood, Three-dimensional orthogonal graph drawing with optimal volume, Proc. Graph Drawing (GD'00), Lecture Notes in Computer Science 1984, 2001, pp. 284–295.
- [2] G. Brebner, Relating routing graphs and two dimensional grids, Proc. VLSI: Algorithms and Architectures, 1985, pp. 221–231.
- [3] T. Calamoneri and A. Massini, An optimal layout of multigrid networks, Information Processing Letters 72 (1999), 137–141.
- [4] T. Calamoneri and A. Massini, Optimal three-dimensional layout of interconnection networks, Theoretical Computer Science 255 (2001), 263–279.
- [5] A. DeHon, Compact, multilayer layout for butterfly fat-tree, ACM Symp. on Parallel Algorithms and Architectures (SPPA), 2000, pp. 206–215.
- [6] P. Eades, A. Symvonis, and S. Whitesides, Two algorithms for three dimensional orthogonal graph drawing, Proc. Graph Drawing '96 (GD '96), Lectures Notes in Computer Science 1190, 1996, pp. 139–154.
- [7] S. Even and R. Kuperstock, Layout area of the hypercube, Journal of Interconnection Networks 4 (2003), 395–417.
- [8] R.I. Greenberg and L. Guan, On the area of hypercube layouts, Information Processing Letters 84 (2002), 41–46.
- [9] F.T. Leighton, Introduction to parallel algorithms and architectures: arrays, trees and hypercubes, Morgan Kaufmann Publ., San Mateo, CA, 1992.
- [10] K. Melhorn, F.P. Preparata, and M. Sarrafzadeh, Channel routing in knock-knee mode: simplified algorithms and proofs, Algorithmica 1 (1996), 213–221.
- [11] S. Muthukrishnan, M. Paterson, S.C. Sahinalp, and T. Suel, Compact grid layouts of multi-level networks, Proc. 31st ACM Symp. on Theory of Computing, 1999, pp. 455–463
- [12] A. Patel, A. Kusalik, and C. McCroskey, Area-efficient layouts for binary hypercubes, IEEE Trans. on Computers 49 (2000), 160–169.
- [13] R.Y. Pinter, On routing two-point nets across a channel, Proc. 19th ACM-IEEE Design Automation Conf., 1982, pp. 894–902.
- [14] A.L. Rosenberg, Three-dimensional VLSI: a case study, Journal of the ACM 30(3) (1983), 397–416.
- [15] D.R. Wood, A new algorithm and open problems in three-dimensional orthogonal graph drawing, Proc. 10th Australasian Workshop Combinatorial Algorithms, (AWOCA), 1999, pp. 157–167.
- [16] T. Yamada and S. Ueno, On three-dimensional layout of De Bruijn networks, Proc. 2002 IEEE International Symposium on Circuits and Systems (ISCAS '02), Vol. III, 2002, pp. 779–782.
- [17] T. Yamada and S. Ueno, On three-dimensional layout of pyramid networks, Proc. 2002 IEEE Asia Pacific Conference on Circuits and Systems (APCCAS '02), Vol. 1, 2002, pp. 159–164.
- [18] C.-H. Yeh, B. Parhami, E.A. Varvarigos, and H. Lee, VLSI layout and packaging of butterfly networks, Proc. ACM Symp. on Parallel Algorithms and Architectures (SPPA), 2000, pp. 196–205.

- [19] C.-H. Yeh, E.A. Varvarigos, and B. Parhami, Multilayer VLSI layout for interconnection networks, Proc. 2000 International Conference on Parallel Processing (ICPP), 2000, pp. 33-40.