1990

A Lower Bound on Embedding Large Hypercubes into Small Hypercubes

Ajay K. Gupta

Alfred J. Boals

Naveed A. Sherwani

Susanne E. Hambrusch
Purdue University, seh@cs.purdue.edu

Report Number:

90-984

A LOWER BOUND ON EMBEDDING LARGE
HYPERCUBES INTO SMALL HYPERCUBES

Ajay K. Gupta
Alfred J. Boals
Naveed A. Sherwani
Susanne E. Hambrusch

CSD-TR-984
May 1990

A Lower Bound on Embedding Large Hypercubes into Small Hypercubes

Ajay K. Gupta[1]
Alfred J. Boals
Naveed A. Sherwani[2]
Department of Computer Science
Western Michigan University
Kalamazoo, Michigan.

Susanne E. Hambrusch[3]
Department of Computer Sciences
Purdue University
West Lafayette, Indiana

[1]Research supported in part by a Fellowship from the Faculty Research and Creative Activities Support Fund, WMU-FRCASF 90-15 and by the National Science Foundation under grant MIP-87-15652.

[2]Research supported in part by a Fellowship from the Faculty Research and Creative Activities Support Fund, WMU-FRCASF 89-225274.

[3]Research supported by the Office of Naval Research under Contracts N00014-84-K-0502 and N00014-86-K-0689, and the National Science Foundation under Grant MIP-87-15652.

Abstract

The problem of embedding an N-processor architecture G into an M-processor architecture H for $N > M$ arises when algorithms designed for architectures of an ideal size are simulated on existing architectures which are of a fixed size. In this paper we present solutions to this embedding problem for the case when both architectures are hypercubes and the embeddings are to achieve a balanced load. An embedding achieves a balanced load if every processor of H simulates at most $\lceil \frac{N}{M} \rceil$ processors of G. We show that in this case hypercube G can be embedded into hypercube H with a dilation of 1 and an optimal congestion of $\frac{N}{M}$. The main contribution of the paper is the lower bound on the congestion.

1 Introduction

Almost all parallel algorithms developed are architecture specific, that is, the architecture of the target parallel machine plays an important role in algorithm design and communication techniques used. Furthermore, it is often assumed that parallel machines are of ideal size; that is, $\mathcal{F}(N)$ processors are available for a problem of size N. In practice, however, an algorithm may be required to run on various architectures having different interconnection network and/or number of processors. Thus, understanding the relationships between various architectures is critical for development of portable parallel algorithms.

In this paper we consider the problem of mapping an algorithm designed for a hypercube G of N processors to that of a hypercube H of M processors with $N > M$. We phrase the problem as a graph embedding problem. The concept of graph embeddings has proven to be a successful one in understanding relationships between different architectures [1, 3, 6, 8, 11, 10, 12]. When $N > M$, one processor of H simulates a number of processors of G and the load of the processors of H becomes a crucial quantity. We concentrate on embeddings that achieve a balanced load; i.e., every processor of H simulates at most $\lceil \frac{N}{M} \rceil$ processors of G. Embeddings achieving a balanced load are of practical importance, since they make every processor of H share an equal load. Before describing our results in detail, we give the necessary definitions.

When the architectures of both machines are viewed as graphs, an embedding $< f, g >$ of G into H is defined by a surjective mapping f from the processors of G to the processors of H together with a mapping g that maps every edge $e = (v, w)$ of G onto a path $g(e)$ connecting $f(v)$ and $f(w)$. We refer to f as the assignment. Since architectures G and H are viewed as graphs, henceforth, we will refer to the processors of G and H as nodes. Two commonly and extensively studied cost measures of an embedding are the dilation and the congestion [1, 3, 6, 7, 8, 11, 12]. The *dilation* δ is defined as the maximum distance in H between two adjacent nodes in G, and the *congestion* λ is defined as the maximum number of paths over an edge in H, where every path represents an edge in G. The *load* μ is defined as the maximum number of nodes of G assigned to any node of

H. We say that an embedding achieves a *balanced load* when $\mu = \lceil \frac{N}{M} \rceil$.

Embeddings of a guest network G into a host network H of the same topology, but smaller size have previously been studied in [2, 5, 6, 7]. In [5], Fishburn and Finkel consider various architectures for specific values of N and M. Berman and Snyder [2] present embeddings by performing contractions which guarantee a dilation of 1, but do not achieve a balanced load. In [7], Gupta and Hambrusch present efficient balanced load embeddings of complete binary trees for all the values of N and M. More recently, Sang and Sudborough [13] have investigated the problem of achieving a balanced load for meshes.

As stated earlier, in this paper we investigate the problem of embedding a large hypercube into a small hypercube. We first give the definition of a hypercube. An n-dimensional hypercube Q_n has $N = 2^n$ nodes and every node in Q_n is labeled as $b_{n-1}b_{n-2}\dots b_0$, where $b_s \in \{0,1\}$ for $0 \leq s \leq n-1$. A node with label $b_{n-1}\dots b_0$ is connected to n nodes having labels $b_{n-1}\dots b_{s+1}\bar{b}_s b_{s-1}\dots b_0$, for $0 \leq s \leq n-1$. When viewing the hypercube Q_n as a graph, we let V_n be its node set and E_n be its edge set.

In section 2, we give a simple embedding of an n-dimensional hypercube Q_n into an m-dimensional hypercube Q_m, for $n > m$. This embedding achieve an optimal dilation of 1, a congestion of 2^{n-m}, and a balanced load of 2^{n-m}. In Section 3 we show that the congestion achieved is optimal. We show that any embedding of Q_n into Q_m which achieves a balanced load must have a congestion of at least 2^{n-m}, that is, $\lambda = \Omega(2^{n-m})$. In order to prove our lower bound we define a hypercube-like architecture called a *compact hypercube*. Compact hypercubes contrast with hypercubes in the sense that while hypercubes can only be defined for powers of 2, compact hypercubes can be defined for any integer $N > 0$. Intuitively, a compact hypercube is a union of disjoint hypercubes together with the additional edges such that it forms an N-node induced graph in a hypercube. Compact hypercubes have been shown to retain many hypercube properties [4].

From gupta@sol.cs.wmich.edu Thu May 24 14:53:15 1990

2 Hypercube Embedding

In this section we present an efficient embedding of an n-dimensional guest hypercube Q_n into an m-dimensional host hypercube Q_m, for $n > m$. The embedding achieves an optimal dilation of 1, a congestion of 2^{n-m}, and a balanced load of 2^{n-m}. The embedding strategy is rather straight forward. The main idea is to contract 2^{n-m} nodes of Q_n forming a subhypercube of dimension $n-m$ and assign them to a node of Q_m.

Let $g_{n-1}, g_{n-2}, \ldots, g_0$ be the labels of nodes in Q_n for $g_i = 0, 1$ and let $h_{m-1}, h_{m-2}, \ldots, h_0$ be the labels of nodes in Q_m for $h_i = 0, 1$. We assign nodes[1] $\overbrace{\star \ldots \star}^{n-m} g_{m-1}, g_{m-2}, \ldots, g_0$ of Q_n to node $h_{m-1}, h_{m-2}, \ldots, h_0$ of Q_m where $h_i = g_i$ for $0 \leq i \leq m-1$. Obviously adjacent nodes in Q_n are either assigned to a single node or to adjacent nodes in Q_m. Thus, the dilation achieved by the embedding is 1 which is optimal. Furthermore, 2^{n-m} nodes of Q_n are assigned to every node of Q_m. Given two adjacent nodes v and v' in Q_m, each node of Q_n which is assigned to v is adjacent to exactly one node of Q_n which is assigned to v'. Since there are 2^{n-m} nodes of Q_n assigned to every node of Q_m, the congestion achieved is 2^{n-m}. We prove in Section 3 that this congestion is indeed optimal.

We conclude this section by pointing out that even though we described only one embedding of Q_n into Q_m, there are $\frac{n!}{(n-m)!\, m!}$ such embeddings. We can easily obtain different embeddings by choosing different combination of m bits in the labels of nodes in Q_n.

3 Lower Bound on the Congestion

We now show that any balanced load embedding of an n-dimensional hypercube Q_n into an m-dimensional hypercube Q_m must have a congestion of at least 2^{n-m}. The main idea of the proof is as follows. In any embedding the edges of Q_n that do not contribute to the congestion are the edges having both end points (i.e., nodes of Q_n) assigned to the same node in Q_m. Thus our goal is to obtain an upper bound on the number of such

[1] $\star$ in the label indicates a wild card character that could be either 0 or 1

edges. Let this upper bound be U. Since, Q_n contains $n2^{n-1}$ edges, $n2^{n-1} - U$ edges of Q_n must contribute to the congestion. Hence, by using U and the number of edges in Q_m, we can determine a lower bound for the congestion λ. The main thrust of the lower bound proof is the computation of U.

In a graph $G = (V, E)$, we refer to $|V|$, the number of nodes in G, as the *order* of G and to $|E|$, the number of edges in G, as the *size* of G. An induced subgraph on a set S of nodes in graph G is the graph whose node set is S and whose edge set consists of those edges in G having both ends in S.

In order to compute U, we determine an upper bound for the size, say $\mathcal{E}$, of an induced subgraph of order 2^{n-m} in Q_n. The size $\mathcal{E}$ multiplied by 2^m gives us U, since in a balanced load embedding exactly 2^{n-m} nodes of Q_n are assigned to every node of Q_m. The most important property that allows us to compute the value $\mathcal{E}$ is the following property of the hypercube: Any k-dimensional hypercube contains two disjoint $k-1$ dimensional subhypercubes such that every node in one subhypercube is adjacent to exactly one node in the other subhypercube. We first give the recursive definitions of a compact set and a compact hypercube which are crucial in determining the value of $\mathcal{E}$.

Definition 1 *A set of nodes S in an n-dimensional hypercube Q_n is compact if*

i. $|S| \leq 1$, *or*

ii. *For a positive integer p, $2^{p-1} \leq |S| < 2^p$, there exists a p-dimensional hypercube Q_p with S as a subset of its nodes. In addition, Q_p contains two disjoint $(p-1)$-dimensional hypercubes Q^0_{p-1} and Q^1_{p-1} such that the node set V^0_{p-1} of Q^0_{p-1} is a subset of S and $S - V^0_{p-1}$ is a compact set in Q^1_{p-1}.*

Given a compact set S, we define a *compact hypercube* to be the subgraph of a hypercube which is induced by the set S. Figure 1 shows an example of a 13 node compact hypercube. It is important to note that compact hypercubes are a generalization of hypercubes in the sense that they are defined for any arbitrary number of nodes. Furthermore, compact hypercubes share many properties with hypercubes [4, 9, 14]. We denote the number of edges in a compact hypercube of r nodes by $\mathcal{E}(r)$. Lemma 1 gives a reduction formula for $\mathcal{E}(r)$ and Lemma 2 gives the precise value of $\mathcal{E}(r)$.

Lemma 1 *If $G = (V, E)$ is a compact hypercube of hypercube Q_d and p is a positive integer such that $2^{p-1} \leq |V| < 2^p$ then*

$$\mathcal{E}(|V|) = \mathcal{E}(|V| - 2^{p-1}) + \mathcal{E}(2^{p-1}) + |V| - 2^{p-1}, \text{ with } \mathcal{E}(0) = \mathcal{E}(1) = 0.$$

Proof: Follows immediately from the definition of a compact set. ∎

Lemma 2 *If $G = (V, E)$ is a compact hypercube of hypercube Q_d and the binary representation of $|V|$ is $b_{p-1}b_{p-2}\cdots b_0$, then*

$$\mathcal{E}(|V|) = |E| = \sum_{i=0}^{p-1} ib_i 2^{i-1} + \sum_{i=0}^{p-2}(b_{i+1}\sum_{j=0}^{i} b_j 2^j).$$

Proof: Intuitively, the first term in the formula for $|E|$ counts the number of edges in the disjoint hypercubes Q_i's that are contained in G. The second term counts number of edges that connect a hypercube Q_i with a hypercube Q_j, for all possible pairs i and j. We prove the lemma using induction. It is easy to see that the formula is correct for $|V| \leq 2$. Assume that the formula is correct for $|V| < l$. Suppose now that $|V| = l$, where the binary representation of l is $b_{q-1}b_{q-2}\ldots b_0$ with $b_{q-1} \neq 0$ and $q \leq p$, i.e., $b_q = b_{q+1} = \cdots = b_{p-1} = 0$. Thus, we have

$$|E| = \sum_{i=0}^{p-1} ib_i 2^{i-1} + \sum_{i=0}^{p-2}(b_{i+1}\sum_{j=0}^{i} b_j 2^j) = \sum_{i=0}^{q-1} ib_i 2^{i-1} + \sum_{i=0}^{q-2}(b_{i+1}\sum_{j=0}^{i} b_j 2^j).$$

By using Lemma 1

$$\begin{aligned}
\mathcal{E}(|V|) &= \mathcal{E}(|V| - 2^{q-1}) + \mathcal{E}(2^{q-1}) + |V| - 2^{q-1} \\
&= \sum_{i=0}^{q-2} ib_i 2^{i-1} + \sum_{i=0}^{q-3}(b_{i+1}\sum_{j=0}^{i} b_j 2^j) + \mathcal{E}(2^{q-1}) + |V| - 2^{q-1} \qquad \text{(by induction)} \\
&= \sum_{i=0}^{q-2} ib_i 2^{i-1} + \sum_{i=0}^{q-3}(b_{i+1}\sum_{j=0}^{i} b_j 2^j) + (q-1)2^{q-2} + \sum_{i=0}^{q-1} b_i 2^i - 2^{q-1} \\
&\qquad \text{(by induction and the value of } |V|\text{)} \\
&= \sum_{i=0}^{q-1} ib_i 2^{i-1} + \sum_{i=0}^{q-2}(b_{i+1}\sum_{j=0}^{i} b_j 2^j).
\end{aligned}$$

This completes the proof. ∎

We now prove a lemma and a theorem which show that any induced graph of order r in a hypercube has its size bounded by $\mathcal{E}(r)$. The result of the theorem provides the desired upper bound for the maximum size of the induced graph of order 2^{n-m} in a hypercube.

Lemma 3 *Suppose Q^0_{k-1} and Q^1_{k-1} are disjoint $(k-1)$-dimensional subhypercubes in a k-dimensional hypercube Q_k. If $G_i = (V_i, E_i)$ is a compact hypercube of Q^i_{k-1} for $i = 0, 1$ then*

$$\mathcal{E}(|V_0| + |V_1|) \geq \mathcal{E}(|V_0|) + \mathcal{E}(|V_1|) + \min\{|V_0|, |V_1|\}.$$

Proof: We proceed by induction on $|V_0| + |V_1|$. Note that if $|V_0| + |V_1| \leq 2$, then the theorem holds. Assume now that the theorem holds for $|V_0| + |V_1| < l$. Suppose $|V_0| + |V_1| = l$ and the binary representation of $|V_0|$ is $b_{p-1}b_{p-2}\ldots b_0$ and that of $|V_1|$ is $c_{q-1}c_{q-2}\ldots c_0$, where $b_{p-1} = c_{q-1} = 1$. Without loss of generality we assume $|V_0| \geq |V_1| > 1$, since the theorem holds for $\min\{|V_0|, |V_1|\} \leq 1$. We consider three cases depending on the relative values of p and q.

Case 1: $q < p$ and $|V_0| + |V_1| \leq 2^p$.

By definition, G_0 contains a subgraph Q_{p-1} which is a $(p-1)$-dimensional hypercube such that $G_0 - Q_{p-1}$ is a compact hypercube of Q^0_{k-1}. Thus by using induction

$$\mathcal{E}(|V_0| + |V_1| - 2^{p-1}) \geq \mathcal{E}(|V_0| - 2^{p-1}) + \mathcal{E}(|V_1|) + \min\{|V_0| - 2^{p-1}, |V_1|\}. \quad (1)$$

Using Lemma 1, we have

$$\mathcal{E}(|V_0| + |V_1|) = \mathcal{E}(|V_0| + |V_1| - 2^{p-1}) + \mathcal{E}(2^{p-1}) + |V_0| + |V_1| - 2^{p-1}, \text{and} \quad (2)$$

$$\mathcal{E}(|V_0|) = \mathcal{E}(|V_0| - 2^{p-1}) + \mathcal{E}(2^{p-1}) + |V_0| - 2^{p-1} \quad (3)$$

Substituting $\mathcal{E}(|V_0| + |V_1| - 2^{p-1})$ from equation 1 into equation 2, we get

$$\begin{aligned}\mathcal{E}(|V_0| + |V_1|) \geq\ & \mathcal{E}(|V_0| - 2^{p-1}) + \mathcal{E}(|V_1|) + \min\{|V_0| - 2^{p-1}, |V_1|\} \\ & + \mathcal{E}(2^{p-1}) + |V_0| + |V_1| - 2^{p-1}\end{aligned}$$

$$
\begin{aligned}
&\geq \mathcal{E}(|V_0|) + \mathcal{E}(|V_1|) + \min\{|V_0| - 2^{p-1}, |V_1|\} + |V_1| \\
&\qquad\qquad \text{(by substituting for } \mathcal{E}(|V_0| - 2^{p-1}) \text{ from equation 3)} \\
&\geq \mathcal{E}(|V_0|) + \mathcal{E}(|V_1|) + |V_1|.
\end{aligned}
$$

Case 2: $q < p$ and $|V_0| + |V_1| > 2^p$.

As in Case 1, G_0 contains a subgraph Q_{p-1} which is a $(p-1)$-dimensional hypercube such that $G_0 - Q_{p-1}$ is a compact hypercube. Now using Lemma 1

$$
\mathcal{E}(|V_0| + |V_1|) = \mathcal{E}(|V_0| + |V_1| - 2^p) + \mathcal{E}(2^p) + |V_0| + |V_1| - 2^p, \tag{4}
$$

$$
\mathcal{E}(|V_0| + |V_1| - 2^{p-1}) = \mathcal{E}(|V_0| + |V_1| - 2^{p-1} - 2^{p-1}) + \mathcal{E}(2^{p-1}) + |V_0| + |V_1| - 2^p, \tag{5}
$$

$$
\mathcal{E}(2^p) = \mathcal{E}(2^p - 2^{p-1}) + \mathcal{E}(2^{p-1}) + 2^{p-1}, \text{ and} \tag{6}
$$

$$
\mathcal{E}(|V_0|) = \mathcal{E}(|V_0| - 2^{p-1}) + \mathcal{E}(2^{p-1}) + |V_0| - 2^{p-1}. \tag{7}
$$

Combining equations 4, 5, and 6, we have

$$
\mathcal{E}(|V_0| + |V_1|) = \mathcal{E}(|V_0| + |V_1| - 2^{p-1}) + \mathcal{E}(2^{p-1}) + 2^{p-1} \tag{8}
$$

By the induction hypothesis

$$
\mathcal{E}((|V_0| - 2^{p-1}) + |V_1|) \geq \mathcal{E}(|V_0| - 2^{p-1}) + \mathcal{E}(|V_1|) + \min\{|V_0| - 2^{p-1}, |V_1|\} \tag{9}
$$

Now, using equations 8 and 9, we have

$$
\begin{aligned}
\mathcal{E}(|V_0| + |V_1|) &\geq \mathcal{E}(|V_0| - 2^{p-1}) + \mathcal{E}(|V_1|) + \min\{|V_0| - 2^{p-1}, |V_1|\} + \mathcal{E}(2^{p-1}) + 2^{p-1} \\
&\geq \mathcal{E}(|V_0|) + \mathcal{E}(|V_1|) + \min\{|V_0| - 2^{p-1}, |V_1|\} + 2^p - |V_0| \\
&\qquad\qquad \text{(by using equation 7)} \\
&\geq \mathcal{E}(|V_0|) + \mathcal{E}(|V_1|) + |V_1|.
\end{aligned}
$$

Case 3: $q = p$.

In this case both V_0 and V_1 contain $(p-1)$-dimensional hypercubes Q_{p-1} and Q'_{p-1}, respectively. The following three formulae result from Lemma 1.

$$\begin{aligned}
\mathcal{E}(|V_0|+|V_1|) &= \mathcal{E}(|V_0|+|V_1|-2^p)+\mathcal{E}(2^p)+|V_0|+|V_1|-2^p && (10)\\
\mathcal{E}(|V_i|-2^{p-1}) &= \mathcal{E}(|V_i|)-\mathcal{E}(2^{p-1})-|V_i|+2^{p-1} \quad \text{for } i=0,1. && (11)\\
\mathcal{E}(2^p) &= \mathcal{E}(2^p-2^{p-1})+\mathcal{E}(2^{p-1})+2^{p-1} && (12)
\end{aligned}$$

By the induction hypothesis, we have

$$\mathcal{E}(|V_0|-2^{p-1}+|V_1|-2^{p-1}) \geq \mathcal{E}(|V_0|-2^{p-1})+\mathcal{E}(|V_1|-2^{p-1})+|V_1|-2^{p-1}$$

and by using equations 10, 11, and 12, we obtain the desired result of the lemma; i.e., $\mathcal{E}(|V_0|+|V_1|) \geq \mathcal{E}(|V_0|)+\mathcal{E}(|V_1|)+|V_1|$. This proves the lemma. ∎

Theorem 4 *If $G=(V,E)$ is an induced subgraph of a d-dimensional hypercube Q_d then $|E| \leq \mathcal{E}(|V|)$; i.e., the size of G is bounded above by the size of a compact hypercube of the same order.*

Proof: Let k be the dimension of a smallest hypercube that contains G; i.e., $k=\min\{j \mid$ there exists a j-dimensional subhypercube of Q_d containing $G\}$. Let Q_k be such a subhypercube of Q_d and let Q^0_{k-1} and Q^1_{k-1} be two disjoint $(k-1)$-dimensional hypercubes contained in Q_k. Let $Q^i_{k-1}=(V^i_{k-1},E^i_{k-1})$. We define $G_i=(V_i,E_i)=(V\cap V^i_{k-1},E\cap E^i_{k-1})$ for $i=0$ and 1. Since every node in V^0_{k-1} is adjacent to exactly one node in V^1_{k-1}, the number of edges between V_0 and V_1 is at most $\min\{|V_0|,|V_1|\}$. It now follows that

$$|E| \leq |E_0|+|E_1|+\min\{|V_0|,|V_1|\}.$$

Because the size of G_i is at least 1, by induction we have $|E_i| \leq \mathcal{E}(|V_i|)$ for $i=0,1$. Now, using these formulae:

$$|E| \leq \mathcal{E}(|V_0|)+\mathcal{E}(|V_1|)+\min\{|V_0|,|V_1|\}.$$

We now obtain the result of the theorem by applying Lemma 3 and thus

$$|E| \leq \mathcal{E}(|V_0|+|V_1|)=\mathcal{E}(|V|).$$

This completes the proof of Theorem 4. ∎

The above theorem gives us an upper bound on the size of an induced subgraph as a function of the order of the induced subgraph in a hypercube. We are now ready to complete our lower bound proof for the congestion. We actually need the result of Theorem 4 only for the case when the order of the induced subgraph is 2^{n-m} (since exactly 2^{n-m} nodes of Q_n are assigned to every node of Q_m in a balanced embedding). Interestingly, Theorem 4 holds for any arbitrary order.

Theorem 5 *Any embedding of an n-dimensional hypercube Q_n into an m-dimensional hypercube Q_m with a balanced load must achieve a congestion of at least 2^{n-m}; i.e.,* $\lambda = \Omega(2^{n-m})$.

Proof: Let $\phi = < f, g >$ be an embedding of Q_n into Q_m that achieves a balanced load. Every node of Q_m is assigned 2^{n-m} nodes of Q_n by ϕ. An edge e fails to contribute to the congestion if it is an edge whose end points belong to $f^{-1}(v)$, for some node v in Q_m; i.e., e is an edge in the induced graph on $f^{-1}(v)$. We know from Theorem 4 that the size of the induced graph on $f^{-1}(v)$ is bounded above by $\mathcal{E}(|f^{-1}(v)|) = (n-m)2^{n-m-1}$, since $|f^{-1}(v)| = 2^{n-m}$. Thus, the maximum number of edges of Q_n that fail to contribute to the congestion is $(n-m)2^{n-m-1}2^m$ because there are 2^m nodes in Q_m. Remaining edges in Q_n must contribute to the congestion and in order to minimize the congestion over the edges of Q_m, these edges must be mapped evenly by g in ϕ. Hence

$$\begin{aligned} \lambda &\geq \frac{n2^{n-1} - (n-m)2^{n-m-1}2^m}{m2^{m-1}} \\ &\geq 2^{n-m} \end{aligned}$$

Note that if the dilation is greater than one, then the above inequality is strict. This completes the proof of Theorem 5. ∎

4 Conclusions

In this paper we presented a lower bound proof showing that any embedding of an $N = 2^n$-node hypercube Q_n into an $M = 2^m$-node hypercube Q_m which achieves a balanced load must have a congestion of at least $\frac{N}{M}$, i.e., $\lambda = \Omega(\frac{N}{M})$. We also presented a simple embedding of Q_n into Q_m that achieves this lower bound on the congestion together with an optimal dilation of 1 and a balanced load of $\frac{N}{M}$. The lower bound proof makes use of compact hypercubes which are a generalization of hypercubes. Compact hypercubes are defined for any arbitrary value of N as compared to a complete hypercube of N nodes where N is a power of two. We conjecture that compact hypercubes can be used to obtain lower bound proofs on the congestion of embeddings of k-ary hypercubes (the hypercubes used in this paper are binary hypercubes). Furthermore, since our lower bound proof of this paper is constructive, the lower bound proofs would result in efficient embeddings of large k-ary hypercubes into small k-ary hypercubes.

References

[1] R. Aleliunas and A. Rosenberg. On embedding rectangular grids into square grids. *IEEE Trans. on Computers*, V-31:907–913, 1982.

[2] F. Berman and L. Snyder. On mapping parallel algorithms into parallel architectures. *Journal of Parallel and Distributed Computing*, 4:439–458, 1987.

[3] S. Bhatt, F. Chung, F. T. Leighton, and A. Rosenberg. Optimal simulations of tree machines. In *the Proceedings of* 27^{th} *Symposium on Foundation of Computer Science*, pages 274–282, 1986.

[4] A. J. Boals, A. K. Gupta, and N. A. Sherwani. Properties of Compact Hypercubes. . Manuscript 1990.

[5] J. P. Fishburn and R. A. Finkel. Quotient networks. *IEEE Transactions on Computers*, C-31(4):288–295, April 1982.

[6] A. K. Gupta. *On the Relationship between Parallel Computation and Graph Embeddings.* PhD thesis, Purdue University, August 1989.

[7] A. K. Gupta and S. E. Hambrusch. Embedding Large Tree Machines into Small Ones. In *the Proceedings of the* 5^{th} *MIT Conference on Advanced Research in VLSI*, pages 179-199, March 1988.

[8] J. W. Hong, K. Mehlhorn, and A. Rosenberg. Cost trade-offs in graph embeddings, with applications. *JACM*, pages 709–728, 1983.

[9] H. P. Katseff. Incomplete Hypercubes. *IEEE Transactions on Computers,* vol. 37, no. 5, May 1988, pages 604-608.

[10] R. Koch, F. T. Leighton, B. Maggs, S. Rao, and A. Rosenberg. Work-preserving emulations of fixed-connection networks. In *the Proceedings of* 21^{th} *STOC*, pages 227–240, 1989.

[11] S. R. Kosaraju and M. Atallah. Optimal simulations between mesh connected array of processors. *Journal of the Association of Computing Machinery*, 35(3):635–650, July 1988.

[12] A. L. Rosenberg. Preserving proximity in arrays. *SIAM Journal on Computing*, pages 443–460, 1979.

[13] F. C. Sang and I. H. Sudborough. Embedding large meshes into small ones. To appear in the *Proceedings of the IEEE Symposium on Circuits and Systems*, 1990.

[14] P. K. Srimani and B. P. Sinha. Message Broadcasting in Point-to-Point Computer Networks. In *the Proceedings of the International Symposium of Circuits and Systems*, pages 189-192, 1988.

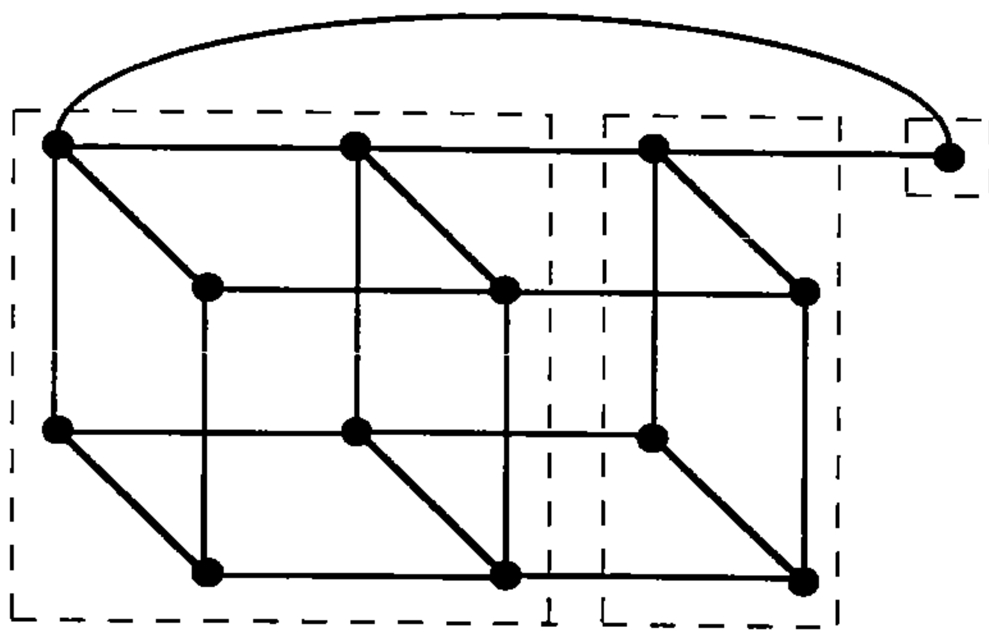

Figure 1: A compact hypercube of 13 nodes. (Complete hypercubes forming the compact hypercube are shown in dashed boxes.)