The HCN: A Versatile Interconnection Network Based on Cubes

Kanad Ghose and Kiran Raghavendra Desai

Department of Computer Science
State University of New York
Binghamton, NY 13901

Abstract

This paper introduces a family of interconnection networks for loosely-coupled multiprocessors called Hierarchical Cubic Networks (HCNs). HCNs use the well-known hypercube network as their basic building block. Using a considerably lower number of links per node, HCNs realize lower network diameters than the hypercube. The performance of several well-known applications on a hypothetical system employing the HCN is identical to their performance on a hypercube. HCNs thus enjoy the same advantages as a hypercube, albeit with considerably simpler interconnections.

Keywords: Hypercube, Interconnection Networks, Loosely-Coupled Multiprocessor, Parallel Processing.

1. Introduction

In recent years, a modest number of multiprocessors using a Hypercube interconnection topology have been announced. These include the Intel iSPC, iSPC/2 [6], the NCUBE family, the Symult (Ametek) S series [1] and the Connection Machines, CM-1, CM-2 [12] and a variety of others. The hypercube interconnection has a number of features that make it attractive as an interconnection scheme in highly concurrent, loosely-coupled multiprocessors:

- The hypercube interconnection scheme subsumes a wide variety of interconnection patterns required by common parallel applications. More specifically, a hypercube of dimension n can emulate a wraparound mesh of dimension $p \times q$, where
$$\lceil \log (p) \rceil + \lceil \log (q) \rceil \leq n.$$
The Hypercube can thus emulate any interconnection topology that the mesh can emulate -- this includes trees, linear geometries, rings, X-trees and so on [6], [13]. In addition, it can emulate a variety of other interconnection patterns such as a logarithmic graph and the butterfly interconnection [13].
- The diameter (or maximum internode distance) is logarithmically related to the number of nodes.
- The number of links per node is also logarithmically related to the number of nodes.
- The algorithm for routing messages in a hypercube is fairly straightforward and simple.
- There are alternate paths between any pair of nodes, making possible fault-tolerance and alternate routings to avoid congestion.

A variety of interconnection networks based on the hypercube have also been proposed. These include the spanning bus hypercube [14], cube-connected-cycles [9], dual-bus hypercube [14], hypernets based on cubelets [5], generalized cubic networks [2], enhancements to the hypercube such as [4], and generalizations such as [8].

The major disadvantages of a hypercube network are:

- Incremental growth is not easy. If incremental growths are allowed, one needs complicated routing algorithms for the 'incomplete cube' [7].
- The layout of the hypercube is not planar, thus complicating its VLSI implementation [13].

In this paper, we will propose a new variety of interconnection network based on the hypercube that has the following advantages:

1. It would allow $2^{(m+n)}$ nodes to be connected, where $n \geq m$, such that the diameter is $\leq (m + n)$ (i.e., cubic), requiring only $(n + 1)$ links per node. In contrast, a hypercube network, connecting the same number of

nodes will require (m + n) links per node.

2. The network can be incrementally expanded, requiring no change in the interconnection strategy as long as a very simple but unconstraining rule is followed.

The proposed network will thus enjoy *all* of the advantages associated with the hypercube albeit with a relatively smaller degree of link connections per node and a lower number of links. Fallouts from the reduced number of links and links per node include a lower network cost, easier wireability, the possibility of integrating the communication controllers on the same chip as the processing node, relatively more planar layout and so on.

Section 2 of the paper describes the proposed network, called the Hierarchical Cubic Network (HCN). Section 3 of the paper describes routing algorithms for the proposed network and evaluates the diameter of the network under these algorithms. We also describe an algorithm that results in an optimal routing distance in Section 3. In Section 4, we describe the design of incomplete HCNs, relevant routing algorithms and the resulting diameters. Section 5 represents an assessment of the HCN, involving its comparison with similar networks on the basis of such things like physical characteristics, diameter and complexities of mapped algorithms. The last section represents the conclusions. An appendix to review the characteristics of a hypercube, especially its routing algorithm, is also provided.

2. The Hierarchical Cubic Network

We propose the design of a new cube based network, called the Hierarchical Cubic Network, HCN. The HCN is a hierarchical network that uses hypercube networks as its nodes. To prevent confusion, we will call each node of this hierarchical network a *cluster*. A cluster is a thus a hypercube network itself, of dimension *n*, say. The hierarchical network proposed has 2^m clusters, where $m \le n$. Such a network is characterized by the values of m and n -- we will thus use HCN (m, n) to describe the proposed network. When m = n, we have a *complete* HCN. When m < n, the HCN is *incomplete*.

Each node or processing element (PE) in a cluster has n *local links* (corresponding to the *n*-cube connection) and one additional link (called an *external link*) to a PE in some *other* cluster. In an incomplete HCN, for some PEs in a cluster, the external link may be unused. *For the time being, we will assume that m = n, i.e, the HCN is complete. We will incorporate appropriate changes that have to be made when the HCN is incomplete.*

We will number the PEs in the HCN as (i, j), where *i* is the binary address of the cluster in which this PE is located and *j* is the "local" address for the PE within the cluster. External links connecting PEs in two different clusters are set up as described in the following algorithm:

```
for i = 0 to (2^n - 1) do {
    for j = 0 to (2^n - 1) do {
        if ( i ≠ j ) then
            Connect((i, j), (j, i));
        else
            Connect((i, i), (ī, ī));
    }
}
```

Note:

1. A new link is not established between nodes (i, j), (k, l) if one already exists.

2. $\bar{i}$ is the complement (1s complement) of i; the external link between nodes (i, i) and $(\bar{i}, \bar{i})$, where $0 \le i \le (2^n - 1)$, is called a *diameter* link.

In what follows, we will use the term 'external link' to allude to external links that connect (i, j) to (j, i), such that $i \ne j$. We will also use the term 'external link' to describe external links that are diameter links. We now describe a set of routing algorithms which we will be using for routing between any two non-local nodes in a HCN(n, n) network.

Figures 1 and 2 depict examples of some complete HCNs. Table 1 depicts the characteristics of the HCN(n, n) network for some values of n. Note that the HCN exhibits diameters identical to hypercubes with the same number of nodes, but requires a considerably lower number of

n	Network	#nodes	#links	links/node	Dia.
1	HCN(1,1)	4	4	2	2
2	HCN(2,2)	16	24	3	4
3	HCN(3,3)	64	128	4	6
4	HCN(4,4)	256	640	5	8
5	HCN(5,5)	1024	3072	6	10
6	HCN(6,6)	4096	14336	7	12
7	HCN(7,7)	16K	64K	8	14

Table 1. Characteristics of some complete HCNs

interconnections per node than the hypercube. This difference becomes more dramatic as n increases, tending towards half the number of links per node compared to a hypercube. The number of links needed in a complete HCN is also considerably smaller than that needed by a hypercube (with the same number of nodes) and also approaches half the number of links in the hypercube, when n increases.

3. Routing Algorithms in a Complete HCN

In this section we will describe some possible routing algorithms for the complete HCN and estimate the diameter of the HCN under these routing algorithms.

3.1 Routing Algorithm A

Algorithm A routes a message from a node j in cluster i i.e., (i, j) to a node l in cluster k i.e. (k, l) as follows:

(i) Use the hypercube routing algorithm (Appendix) to route from node (i, j) to the node (i, k). (This amounts to *routing locally* within the hypercube i.) Node (i, k) is the node which has the external link to the cluster k.

(ii) Route from (i, k) to (k, i) to reach the cluster k. (This does an *inter-cluster* or *external routing*.)

(iii) Locally route from (k, i) to (k, l) in the hypercube k using the algorithm described in the Appendix.

Algorithm A thus routes from (i, j) to (k, l) using the external link that exists between cluster i and cluster k. This is the only external routing involved in Algorithm A -- all other routings are local.

3.2 Routing Algorithm B

Algorithm B routes a message from node (i, j) to node (k, l) as follows:

(i) Use local routing to route from (i, j) to (i, i) (Appendix). The node (i, j) has an external link to node ($\bar{i}$, $\bar{i}$) in cluster $\bar{i}$.

(ii) Use the diameter link from node (i, i) to go to node ($\bar{i}$, $\bar{i}$).

(iii) From node ($\bar{i}$, $\bar{i}$), use local routing to route to ($\bar{i}$, k). The node ($\bar{i}$, k) has the external link to cluster k.

(iv) From node ($\bar{i}$, k) take the external link to go to cluster k at node (k, $\bar{i}$).

(v) Now use local routing to go from node (k, $\bar{i}$) to node (k, l).

Algorithm B thus uses one diameter and one non-diameter link to do the routing between clusters. These are the only external routing involved in Algorithm B.

3.3 Routing Algorithm C

Algorithm C routes a message from node (i, j) to node (k, l) as follows:

(i) Use local routing to route from node (i, j) to node (i, l). Node (i, l) is the node with the external link to cluster l.

(ii) From node (i, l) take the external link to go to node (l, i).

(iii) From node (l, i) use local routing to route to node (l, k).

(iv) At node (l, k) take the external link to go to node (k, l), the destination.

Algorithm C thus uses two external links to do the routing between two clusters. In particular, it does not use any diameter links.

3.4 Routing Delays & Network Diameter

We now state some axioms and definitions which we will be using in the next few sections of our paper.

Axiom 1: The diameter of an n-cube is n. (This is clear from the discussions in the Appendix.)

Axiom 2: For every node i in an n-cube, there exists only one node $\bar{i}$ at a distance of n from it. All other nodes are at a distance less than n from i.

Definition: We define a term N(i, j) such that N(i, j) is the ordinality of the bit-string X obtained by the bit-wise Exclusive-OR of bit-strings i and j, i.e., N(i, j) is the number of 1's in the bit-string X.

Axiom 3: The best routing algorithm to route a message from node i to a node j in an n-cube takes N(i, j) steps. This is apparent from the routing algorithm given in the Appendix -- a message is routed from a source node to the destination via a number of nodes en route, and at each step in the routing, the ordinality of X is lowered by 1.

We now state the results pertaining to the routing delay and diameter as theorems.

Theorem 1: Algorithm A will route from (i, j) to (k, l) in less than or equal to (n + n) steps in an HCN(n, n) network iff $j \neq \bar{k}$ or $i \neq \bar{l}$.

Proof: Consider routing from node (i, j) to node (k, l) in an HCN(n, n). Algorithm A routes as follows:

$$(i, j) \xrightarrow{N(j,\,k)} (i, k) \xrightarrow{1} (k, i) \xrightarrow{N(i,\,l)} (k, l)$$

where the number on each arrow denotes the number of steps involved in the routing.

Hence the number of steps, R = N(j, k) + N(i, l) + 1. Now, for an n-cube, $N(j, k) \leq n$ and $N(i, l) \leq n$. Thus the number of routing steps $\leq (n + n + 1)$. In the worst case when $j = \bar{k}$ and $i = \bar{l}$, we have:

$$R = n + n + 1.$$

Using Axiom 2, we have:

$$R \leq n + n \qquad \text{when } j \neq \bar{k} \text{ OR } i \neq \bar{l}.$$

Hence the theorem. □

Theorem 2: Algorithm B will route from (i, j) to (k, l) in at most 2n steps in an HCN(n, n) network if

$$j = \bar{k} \text{ and } i = \bar{l} \text{ and } i \neq k$$

Proof: In general, algorithm B routes as follows:

$$(i, j) \xrightarrow{N(j,\,i)} (i, i) \xrightarrow{1} (\bar{i}, \bar{i}) \xrightarrow{N(\bar{i},\,k)} (\bar{i}, k)$$

$$\xrightarrow{1} (k, \bar{i}) \xrightarrow{N(\bar{i},\,l)} (k, l)$$

Thus the number of routing steps resulting from algorithm B is:

$$R = N(i, j) + 1 + N(\bar{i}, k) + 1 + N(\bar{i}, l)$$

Now consider routing from $(i, \bar{k})$ to $(k, \bar{l})$ in an HCN. For this particular source-destination node-pair, $j = \bar{k}$ and $i = \bar{l}$. Therefore, the number of routing steps following algorithm B is

$$2\,N(i, \bar{k}) + 2, \ \text{ since } N(i, \bar{k}) = N(k, \bar{i}).$$

Since $i \neq k$, from Axiom 2 we have

$$N(i, \bar{k}) \leq n - 1$$

Hence the number of steps is $\leq 2\,(n - 1) + 2$, i.e., the number of routing steps is $\leq 2n$. □

Theorem 3: The diameter of an HCN(n, n) network using an appropriate combination of Algorithm A and Algorithm B is 2n.

Proof: Follows directly from Theorems 1 and 2. □

3.5 An Optimal Routing Algorithm

We have discussed three routing algorithms among the many that are possible for the complete HCN. Any one or a combination of such algorithms can be chosen to satisfy one or more optimality criterion such as minimal intercluster traffic, minimum routing distance and so on. In this section, we will present a routing algorithm that provides the minimum routing distance between nodes (i, j) and (k, l).

For Algorithm A, we have:

$$R = N(j, k) + N(i, l) + 1$$

where R is the routing distance between nodes (i, j) and (k, l).

For Algorithm B,

$$R = N(i, j) + N(\bar{i}, k) + N(\bar{i}, l) + 2$$

For Algorithm C, the routing is done as follows

$$(i, j) \xrightarrow{N(j,\,l)} (i, l) \xrightarrow{1} (l, i) \xrightarrow{N(i,\,k)} (l, k) \xrightarrow{1} (k, l)$$

Thus, for Algorithm C, R = N(i, k) + N(j, l) + 2

Therefore the optimal-distance algorithm will choose one of these three algorithms that results in the minimum overall routing distance for a given set of values for i, j, k and l. The algorithm that results is as follows:

1. Compute:

   $R_A = \quad |j \oplus k| + |i \oplus l| + 1$

   $R_B = \quad |i \oplus j| + |\bar{i} \oplus k| + |\bar{i} \oplus l| + 1$

   $R_C = \quad |i \oplus k| + |j \oplus l| + 1$

   where $|b|$ is the ordinal operator that finds the number of 1s in the bit-string b.

2. Let R_Q = minimum (R_A, R_B, R_C), where Q is either A, B or C. Execute Algorithm Q.

Theorem 4: Under the optimal-distance algorithm, the diameter of the HCN(n, n) is bounded by 2n.

The proof is obvious. □

We have therefore managed to get cubic distances in the HCN with almost half the number of links per node and half the number of links as a hypercube. Further, the routing algorithm used to achieve this is almost as simple as the routing algorithm in a hypercube.

4. Incomplete HCNs

In this section, we will look at incomplete HCNs and derive their diameters under routing algorithms similar to the one we had described earlier.

4.1 Construction of Incomplete HCNs

An incomplete HCN can be described as HCN (m, n), where m < n, i.e., the number of nodes in a cluster exceed the number of clusters. Such a HCN can be designed as follows:

1. The 2^m n-cube clusters are now interconnected using inter-cluster links that are a subset of the inter-cluster connections in HCN(n, n), in the following manner: External connections are made from nodes in a cluster only if the *local* node number is less than 2^m, using the same interconnection algorithm we had used for an HCN(n, n) network. Consequently, some of the nodes do not have inter-cluster links. We will describe this network as $HCN_s(m, n)$, where the subscript indicates the subset-nature of the inter-cluster connections.
2. The HCN(m, n) network can also be built as multiple (precisely, $2^{(n-m)}$) HCN(m, m) networks which are interconnected with other similar (precisely, $2^{(n-m)} - 1$ others) using *local* links within the n-cube clusters. Now the connections will not be a subset of the HCN(n, n). We will allude to the resulting network as $HCN_m(m, n)$, where the subscript m designates 'multiple'. In a $HCN_m(m, n)$, every node has an external link, unlike the $HCN_s(m, n)$ -- this is a consequence of choosing 2^m nodes out of the 2^n nodes in each cluster to form a m-cube (which is used as a cluster for the HCN(m, m) networks). This also implies that the maximum distance between any two local nodes in the $HCN_m(m, n)$ that belong to two different HCN(m, m) s is at most (n - m).

4.2 Routing and Diameter Evaluations

In this section, due to the lack of space, we give an intuitive idea behind the routing algorithms and the assessment of the diameters for incomplete HCNs. A complete description of the routing algorithms and a formal approach to the derivation can be found in [3].

For a $HCN_s(m, n)$, the diameter links may be absent for some nodes. The routing algorithms take this fact into account. In the worst case, routing between a source-destination node-pair would involve a traversal of only local and non-diameter external links. Further, in this worst-case, in order to make use of these external links, 2(n - m) additional local links have to be traversed beyond the (2m + 1) links required by Algorithm A (which does not make use of diameter links). This implies that the resulting diameter cannot be worse than (2n + 1). With the addition of clusters, when the network grows towards HCN (n, n), the average inter-node distances would shrink and the diameter under the optimal-distance algorithm approaches 2n.

In the $HCN_m(m, n)$, all the constituent HCN(m, m) networks "meet" in every cluster. Routing a message from a source node to a destination node can be accomplished by routing locally to the HCN(m, m) that contains the destination node, and then routing within the HCN(m, m) so reached. Further, the maximum distance between the source node and the HCN(m, m) containing the destination node is (n - m) (local links). This is because nodes with the same local number in different clusters belong to a HCN(m, m), and the source and destinations can be in two different HCNs, spaced apart by at most (n - m) local links within any cluster. Thus, the diameter of the $HCN_m(m, n)$ is (n - m + 2m) = (n + m), under the optimal-distance routing algorithm.

5. An Assessment of the HCN

In this section, we evaluate some characteristics of the proposed network against similar networks, especially the hypercube network and its derivatives. We close this section by indicating some potentials of the HCN and describing further work in progress.

Network	#Nodes	#Links	#Links/Node	Diameter
HCN(m,n)	2^{m+n}	$2^{m+n-1}.(n+1)$	$n+1$	$m+n^{\dagger}$
(m + n)-Cube	2^{m+n}	$2^{m+n-1}.(m+n)$	$m+n$	$m+n$
2^m-CCC	2^{m+n}	3.2^{m+n-1}	3	5.2^{m-1}
Torus††	$(2^m)^{(1+n/m)}$	$2.(1+n/m).2^{m+n-1}$	$2.(1+n/m)$	$(1+n/m).2^{m-1}$
((m + n + 1)/2, 2)-Hypernet	2^{m+n}	$(m+n+3).2^{m+n-1}$	$(m+n+3)/2$	$(m+n+2)$

Table 2. A comparison of the topological characteristics of the HCN and some similar networks. Assumes the total number of connected nodes to be the same for all the networks.

Notes:

† Evaluated from the result cited in Theorem 3. The actual diameter under the optimal-distance routing algorithm is smaller than this value (Table 3).

†† One of the two possible tori network for connecting 2^{m+n} nodes.

5. 1 The HCN & Other Cube-Like Networks

The results of comparing the HCN against other cube-based and cube-like interconnection networks on the basis of typical attributes for a static interconnection network (number of links, number of links/processing element, diameter etc.) are depicted in Table 2. To allow a fair basis for comparison, we have assumed that *all* of these networks connect the same number of nodes. The HCN has characteristics similar to the hypercube -- however, it realizes all physical characteristics of the hypercube with only (n + 1) links/PE (as opposed to (m + n) links/PE for the hypercube). We have also reason to believe that the diameter of the HCN(m, n) is actually *less* than (m + n) under the optimal distance routing algorithm of Section 3.5. For a HCN(n, n), the network diameter under the optimal routing algorithm can be proved to be $(n + \lfloor n/2 \rfloor + 1)$ [3]. Although we do not present the formal proof in this paper for the sake of brevity, we report the results of simulation that evaluated the diameter of a HCN(m, n) following the optimal-distance routing algorithm in Table 3. Table 4 depicts the complexities of some well-known parallel algorithms on the HCN, the hypercube, the hypernet [5] and a few other static interconnection networks. All of the data for the networks other than the HCN was gleaned from [5].

Network	Diameter (Theorem 3)	Diameter (Simulated)
HCN(1, 1)	2	2
HCN(2, 2)	4	4
HCN(3, 3)	6	5
HCN(4, 4)	8	7
HCN(5, 5)	10	8

Table 3. The diameter of a HCN(n, n) computed using theorem 3 versus the diameter actually obtained from the simulation of the optimal-distance algorithm

Network	Vector Sum	Matrix Mult.	Convolution: NxN space	Bitonic Sort
HCN	log N	log N	M^2 + log N	$\log^2 N$
Hypercube	log N	log N	M^2 + log N	$\log^2 N$
Hypernet	log N	log N	M^2 + log N	$\log^2 N$
2-D Mesh	$N^{0.5}$	N	M^2	$N^{0.5}$
Bin. Tree	log N	N^2	N^2	N

Table 4. Algorithmic complexities using the HCN and other interconnections

5. 2 Potentials of the HCN

The network proposed in this paper seems very attractive from a number of standpoints:

- It is a hypercube-like network, in particular it has a lower diameter than a comparable hypercube, but requires almost half the number of interconnections per PE than the hypercube (for the same number of connected PEs and if the number of connected PEs is large). Also, the average inter-PE distance in a HCN is bounded by (m + n + 2)/2 -- this is one more than the corresponding number for a hypercube (viz., (m + n)/2). For large m and n, this difference becomes negligible.
- The HCN is expandable, if we realize it as a HCN_s(m, n), when m < n (Section 4.1). The diameter of the network is still 2n + 1. When m approaches n, the diameter drops to 2n. If lower diameters are needed for an incomplete HCN, it can be designed as a HCN_m(m, n), which gives a diameter of (m + n).
- The HCN is ideally suited for applications that require the use of groups of several identical processes that require infrequent communication among groups. Processes within a group can be allocated to nodes within a cluster, exploiting hypercube localities, while processes across groups can communicate efficiently using the external links. (The concept of a *task force*, as used in the Cm^* is identical to the process structure we just described.)
- The concept of the HCN can be extended to more levels of hierarchy. It is possible, for instance, to design a HCN that will use another HCN as its cluster and so on, leading to a family of hierarchical networks.

We are currently investigating the idea of building more hierarchy into the network proposed, the design of fail-safe routing algorithms, the mapping of algorithms to the HCN and some related issues. We have also formulated the design of a HCN that uses an additional connection between two clusters, resulting in a diameter of (m + 2) for a HCN connecting 2^{2m-1} nodes with (m + 1) links per node. This will be covered in a forthcoming paper.

6. Conclusions

We have introduced a new static interconnection network, the Hierarchical Cubic Network (HCN) and its routing algorithms. The HCN provides the advantages and the same diameter as a the well-known Hypercube interconnection, but requires almost half the number of connections per node, compared to a Hypercube that connects the same number of nodes. Further, the number of links needed by the HCN is also much smaller compared to a similar Hypercube interconnection. This implies, among other things, a lower cost for the network, reduced wiring efforts, increased integration possibilities and so on for the same performance as a comparable hypercube.

The HCN is incrementally expandable and algorithmic complexities on the HCN appear to be identical to that for a Hypercube. Simulations indicate that the HCN actually provides a *lower* diameter than the Hypercube under the optimal-distance routing algorithm presented in this paper. A comparison of the physical characteristics of the HCN with cube-based and similar networks reveals it to be extremely attractive as an interconnection network for large-scale (loosely-coupled) parallel computers. The HCN is ideally suited for the implementation of task forces of processes. The HCN also has the potential of accommodating multiple levels of hierarchy.

References

[1] Ametek Corp., Ametek Series 2010 Brochures.

[2] Bhuyan, L. N. and Agrawal, D. P. "Generalized Hypercube and Hyperbus Structures for a Computer Network," IEEE Trans. on Computers, Vol. C-33, No. 1, April 1984, pp. 323-333.

[3] Ghose, K. and Desai, K. R. "Hierarchical Cubic Networks", In Preparation.

[4] Hsu, W. T., Yew, P. and Zhu, C. "An Enhancement Scheme for Hypercube Interconnection Networks," Proc. Int'l. Conf. Parallel Processing, August 1987, pp. 820-823.

[5] Hwang, K. and Ghosh, J. "Hypernet: A Communication Efficient Architecture for Constructing Massively Parallel Computers," IEEE Trans. on Computers, Vol. C-36, No. 12, Dec 1987, pp 1450-1466.

[6] Intel Scientific Computers, iSPC System Overview Manual.

[7] Katseff, H. P. "Incomplete Hypercubes," IEEE Trans. on Computers, Vol. C-37, No. 5, May 1988, pp. 604-608.

[8] Lakshmivarahan, S. and Dhall, S. K. "A New Hierarchy of Hypercube Interconnection Schemes for Parallel Computers," Journal of Supercomputing, 2, (1988), pp. 81-108.

[9] Preparata, F.P. and Vuillemin, J. "The Cube-Connected Cycles: A Versatile Network for Parallel Computation," Comm. ACM, Vol. 24, No. 5, May 1981, pp. 300-309.

[10] Reed D. A. and Grunwald, D. C. "The Performance of Multicomputer Interconnection Networks," IEEE Computer, June 1987.

[11] Sullivan, H. and Bashkow, T. R. "A Large Scale Homogeneous, Fully Distributed Parallel Machine, I," Proc. Fourth Symp. Comp. Architecture, March 1977, pp. 105-117.

[12] Tucker, L. W. and Robertson, G. G. "Architecture and Applications of the Connection Machine," IEEE Computer, Vol. 21, No. 8, August 1988, pp. 26-38.

[13] Ullman, J. D. "Computational Aspects of VLSI," Ed. Computer Science Press.

[14] Wittie, L. D. "Communication Structures for Large Networks of Microcomputers," IEEE Trans on Computers, Vol. C-30, No. 4, April 81, pp 264-273.

APPENDIX. The Hypercube Interconnection: a Review

A hypercube network of dimension n, often called a n-cube, has 2^n nodes, each with a n-bit address ranging in value from 0 to $(2^n - 1)$. There is a link between nodes i and j iff their n-bit addresses differ in *exactly* one bit position. The total number of links in a n-cube is $n.2^{(n-1)}$.

The following routing algorithm describes how a message is routed from node i to node j in a hypercube where i and j are the values of the n-bit addresses. Each node i in the hypercube executes this algorithm when it has to send or relay a message to node j.

```
Compute rel = i ⊕ j; /* ⊕ = bit-wise ex-or */
if (rel == 0) then
    done /* message has reached its destination */
else
    begin
        Let K be the position of the most-significant 1
        in the n-bit entity, rel;
        /* bits are given positions (n - 1) through 0,
        msb onwards */
        Route message out through link # K
        /* Link number K connects two neighbouring
        nodes in the hypercube that differ in their
        addresses in the K-th bit position */
    end
```

This routing algorithm guarantees that the diameter of the hypercube (under this routing algorithm) is equal to the dimension of the hypercube, viz., n. In this paper, we used the term n-cube for a cube of dimension n.

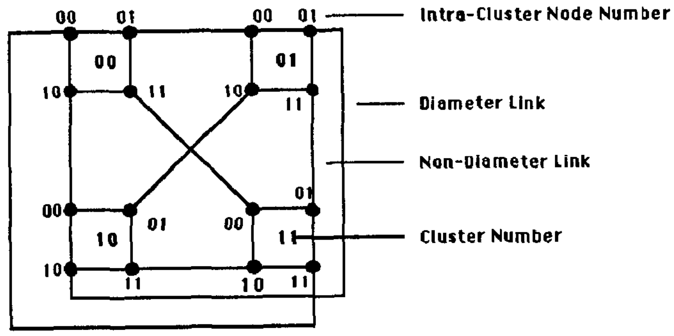

Figure 1. An HCN(2, 2) network.

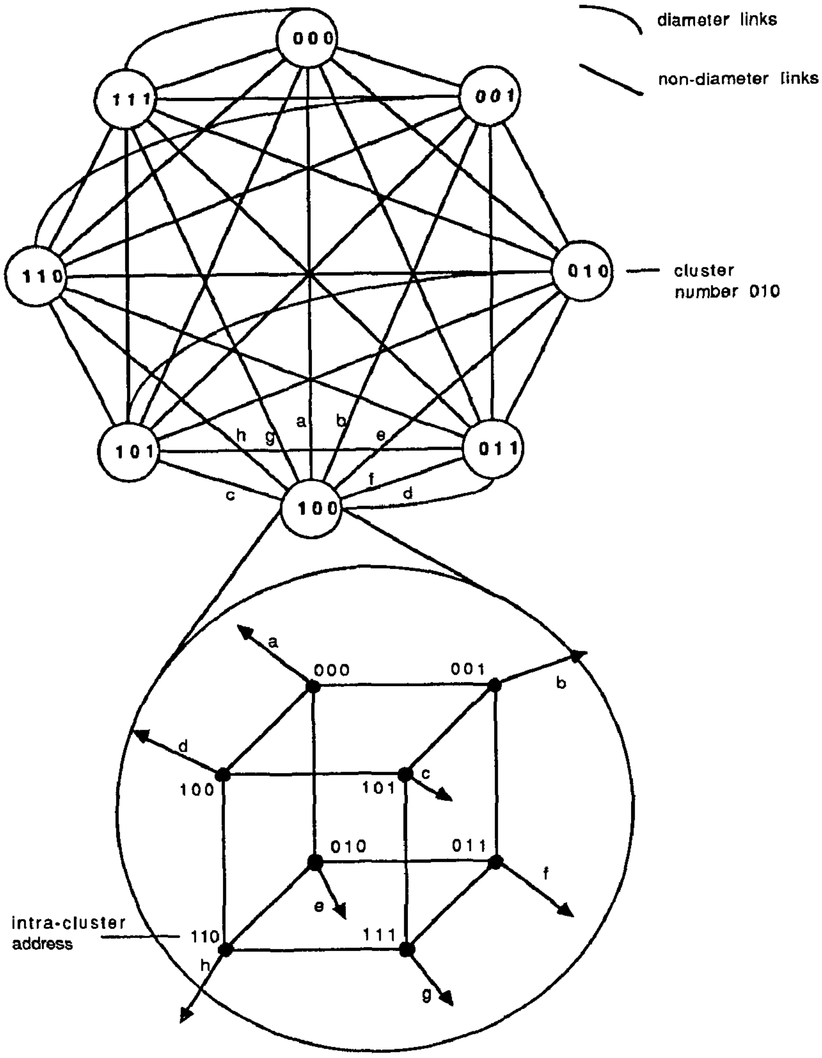

Figure 2. An HCN(3, 3) network.

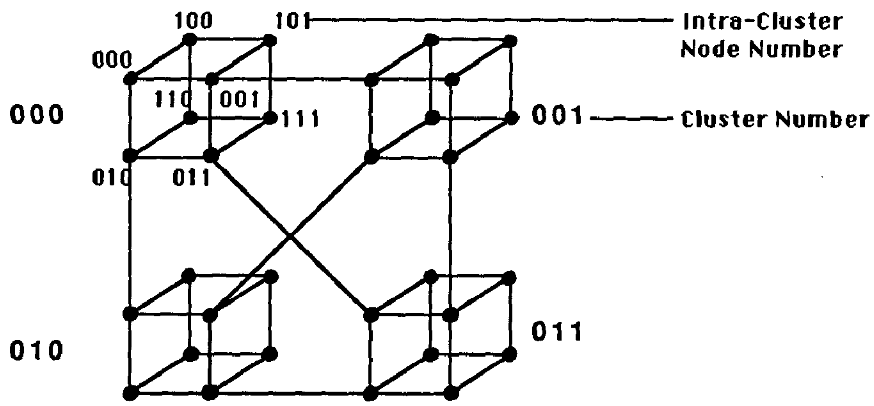

Figure 3.1. An HCNs(2, 3) network. (Diameter = 2n + 1 = 7)

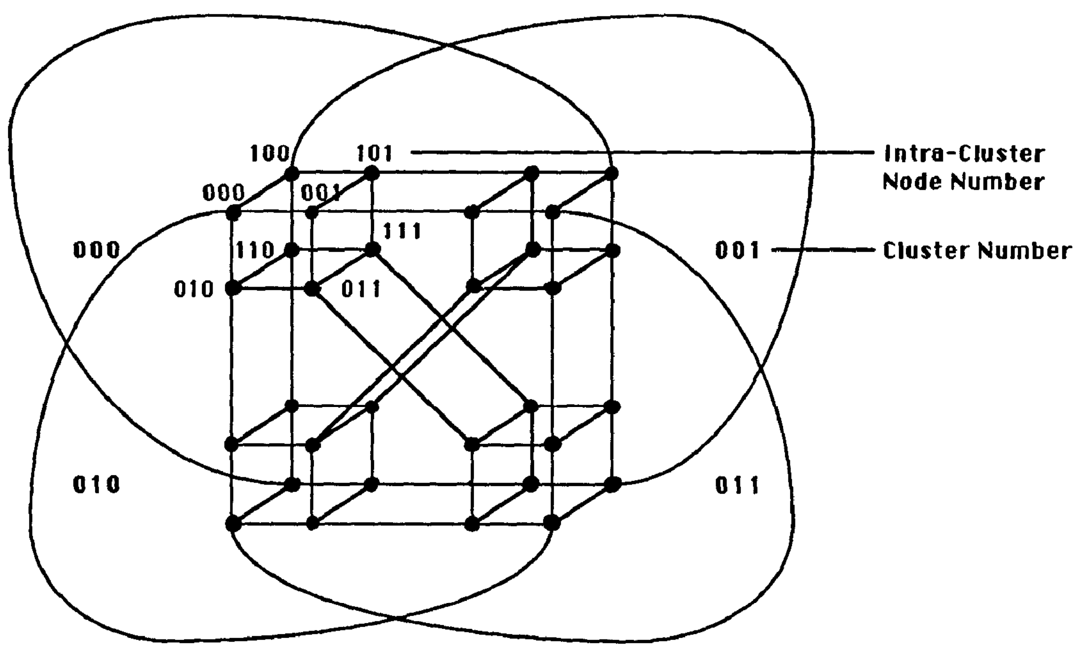

Figure 3.2. An HCNm(2, 3) network. (Diameter = m + n = 5)