Optimal Simulation of Full Binary Trees on Faulty Hypercubes

Bethany M. Y. Chan, Francis Y. L. Chin, *Senior Member, IEEE,* and Chung-Keung Poon

Abstract—We study the problem of running full binary tree based algorithms on a hypercube with faulty nodes. The key to this problem is to devise a method for embedding a full binary tree into the faulty hypercube. Based on a novel embedding strategy, we present two results for embedding an $(n - 1)$-tree (a full binary tree with $2^{n-1} - 1$ nodes) into an n-cube (a hypercube with 2^n nodes) with unit dilation and load. For the problem where the root of the tree must be mapped to a specified hypercube node (*specified root embedding problem*), we show that up to $n - 2$ (node or edge) faults can be tolerated. This result is optimal in the following sense: 1) it is time-optimal, 2) $(n - 1)$-tree is the largest full binary tree that can be embedded in an n-cube, and 3) $n - 2$ faults is the maximum number of worst-case faults that can be tolerated in the specified root problem. Furthermore, we also show that any algorithm for this problem cannot be totally recursive in nature. For the problem where the root can be mapped to any ·nonfaulty hypercube node (*variable root embedding problem*), we show that up to $2n - 3 - \lceil \log n \rceil$ faults can be tolerated. Thus we have improved upon the previous result of $n - 1 - \lceil \log n \rceil$. In addition, we show that the algorithm for the variable root embedding problem is optimal within a class of algorithms called *recursive embedding algorithms* as far as the number of tolerable faults is concerned. Finally, we show that when an $O(1/\sqrt{n})$ fraction of nodes in the hypercube are faulty, it is not always possible to have an $O(1)$-load variable root embedding no matter how large the dilation is.

Index Terms—Embedding, hypercubes, full binary trees, dilation, simulation, faulty architecture.

I. INTRODUCTION

THE HYPERCUBE parallel multiprocessor architecture has been the topic of much recent research. It has been shown to be a very versatile architecture [15] capable of efficiently simulating networks such as rings [10], grids [6], [7], [8], [11] and trees [3], [4], [9], [14], [17], [18]–[21]. At the same time, the hypercube has been shown to be very robust [1], [2], [5], [12], [13], [16].

This paper attempts to further demonstrate the versatility and robustness of the hypercube by showing how a full binary tree can be embedded into a faulty hypercube. We consider a strong fault model in which a faulty node can neither compute nor communicate with its neighbors. Thus the embedding should avoid all the faulty nodes, faulty links as well as links which are nonfaulty but adjacent to a faulty node. There are two different versions of this embedding problem. The

specified root embedding problem is to embed an $(n - 1)$-tree into an n-cube with unit dilation and load. Moreover, the root of the tree must be mapped to a *specified* hypercube node. Note that even when there are no faults, an $(n - 1)$-tree is the largest full binary tree that can be embedded into an n-cube [4], [18], [19], [20]. The *variable root embedding problem* is similar but the root can be mapped to *any* nonfaulty node. Thus there are fewer restrictions on the embedding and hence more faults can be tolerated. It was shown in [9] that up to $n - 1 - \lceil \log n \rceil$ faults can be tolerated in the variable root problem. Recently, [21] showed the existence of a variable root embedding which can tolerate up to $\Omega(n^2 / \log n)$ faults. However, their method is non-recursive and non-constructive. For the specified root embedding, [21] also derived a method to tolerate $\lceil n/4 \rceil$ faults.

This paper presents three new results in regards to both the specified and variable root embedding problems:

i) for the specified root embedding problem, a novel algorithm which is optimal in the following sense:

- time-*optimal*—$O(|T|)$ where $|T|$ = size of the full binary tree,
- *largest* full binary tree that can be embedded—an $(n - 1)$-tree into an n-cube,
- *maximum* number of worst-case faults that can be tolerated—$(n - 2)$ faults,
- totally recursive in nature for all but one situation. It can be shown that that situation is impossible to be solved by any total recursive embedding algorithm.

ii) for the variable root embedding problem, another novel method which is constructive and is based on combinatorial counting and recursive embedding. This method can tolerate up to $2n - 3 - \lceil \log n \rceil$ faults, which is asymptotically the optimal (maximum) number of faults that can be tolerated by any recursive embedding algorithm.

iii) for the variable root embedding problem, no more than $O(1/\sqrt{n}) \times 2^n$ faults can be tolerated in the worst case. Furthermore, the result holds even if we allow unbounded dilation and $O(1)$ load.

II. DEFINITIONS AND NOTATIONS

A *hypercube* of n dimensions (called an n-*cube*) is an undirected graph of 2^n nodes each having a unique n-bit label. Two nodes are connected by a link if and only if their labels differ in exactly one bit position. We shall refer to nodes by

their labels. Furthermore, a node having a 0 (or 1) as the kth bit of its label is said to have a 0 (or 1) in dimension k, with the first bit or dimension taken to mean the leftmost bit. Two nodes differing in the ith dimension only are said to be neighbors of each other on dimension i.

To specify subcubes of the n-cube, we use strings of length n consisting of 0's, 1's and *'s only. A string of length n with exactly m *'s describes an m-cube within an n-cube. For example, *01** denotes the 3-cube comprised of the eight nodes 00100, 00101, 00110, 00111, 10100, 10101, 10110 and 10111 of a 5-cube.

A *full binary tree* with n levels of nodes (or $2^n - 1$ nodes) is called an *n-tree*. The levels are numbered from 0 to $n-1$, with the root being at level 0. Nodes at level i are denoted by strings of length i consisting of L's and R's only. In particular, the empty string ϵ specifies the root and if T is a string denoting a tree node, LT and RT specify respectively the left and right son of that node.

An *embedding* of a binary tree into a hypercube is a mapping from nodes of the tree to nodes of the hypercube. It is said to have *unit dilation and load* if and only if at most one tree node is mapped to each hypercube node and adjacent nodes in the tree are mapped to adjacent nodes in the hypercube. From now on, *embedding* refers to one with unit dilation and load unless otherwise stated.

To specify the (unit dilation and load) embedding, we define $H[T]$ as the hypercube node to which tree node T is mapped. As our embedding methods are recursive, it will be convenient to have some notations for specifying the subproblems. Thus we define $C[T]$ as the subcube in which we want to embed the subtree rooted at T. Furthermore, T-*embedding* refers to the embedding of the subtree rooted at T within $C[T]$, with T mapped to $H[T]$.

III. SPECIFIED ROOT TREE EMBEDDING

In this section, we consider the problem of specified root tree embedding. Our main result is the following.

Theorem 3.1: For all $n \geq 2$ and $0 \leq f \leq n-2$, there exists an embedding of an $(n-1)$-tree into an n-cube containing f faulty nodes/links with the root of the tree mapped to a specified nonfaulty node S in the hypercube.

Note that $n-2$ faults are the maximum that can be tolerated when the root is specified. Thus, $n-2$ is the *fault-tolerance capacity* of an n-cube. For example, if there were $n-1$ faulty nodes neighboring the specified node S, embedding would be impossible. On the other hand, Theorem 3.2 shows that faults can be ignored if they are far away from the root.

Theorem 3.2: The embedding of any m-tree with root mapped to hypercube node S can never be affected by faulty node F if the Hamming distance between S and F is $\geq m$.

Proof: A level i tree node can only be mapped to a hypercube node at Hamming distance $\leq i - 1$ from S. Thus, the whole m-tree can only be mapped to hypercube nodes at Hamming distance $\leq m - 1$ from S. Hence, it is impossible to use F in any embedding. □

To simplify the proof for Theorem 3.1, we assume, without loss of generality, that i) exactly $n - 2$ faults with Hamming
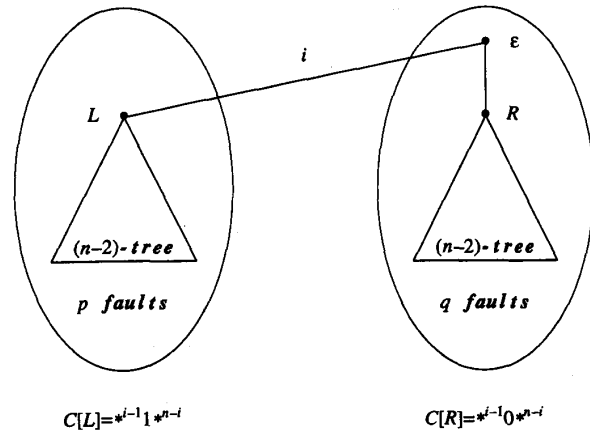

$C[L]=*^{i-1}1*^{n-i}$ $C[R]=*^{i-1}0*^{n-i}$

Fig. 1. Splitting $C[\epsilon] = *^n$ on dimension i.

distance $\leq n - 2$ from S are in the n-cube, ii) all of the faults are node faults, as link faults can be handled by treating either of the two nodes connected by the link as faulty, and iii) $S = 0^n$. We prove the theorem by actually constructing an embedding, i.e., performing the ϵ-embedding with $H[\epsilon] = 0^n$ and $C[\epsilon] = *^n$. The construction is based on two techniques: *cube splitting* and *node borrowing*.

A. Cube Splitting and Node Borrowing

The idea of cube splitting is to map L and R to two nonfaulty neighbors of $H[\epsilon]$ and then embed the subtrees rooted at L and R into two disjoint subcubes. With $H[\epsilon] = 0^n$ and $C[\epsilon] = *^n$, we say that $C[\epsilon]$ is *split* on dimension i if we set $C[L] = *^{i-1}1*^{n-i}$, $C[R] = *^{i-1}0*^{n-i}$ and $H[L] = 0^{i-1}10^{n-i}$. (In other words, $H[L]$ is the neighbor of $H[\epsilon]$ on dimension i, $C[L]$ is the subcube not containing $H[\epsilon]$, and $C[R]$ is the other subcube. See Fig. 1.)

Thus by splitting $C[\epsilon]$ on dimension i and mapping R to a nonfaulty neighbor of $H[\epsilon]$ in $C[R]$, we reduce the original problem to two subproblems, i.e., finding the L- and the R-embeddings. (Note that $H[L]$ lies in $C[L]$ while $H[R] = 0^{j-1}10^{n-j}$ for some $j \neq i$ must be in $C[R]$.) By recursively splitting the subcubes on suitable dimensions, we can derive an embedding in most circumstances. This type of embedding, in which the left and right subtrees of every internal node of the $(n-1)$-tree are mapped to disjoint subcubes, is called a *recursive embedding* in [21].

[21] also showed that there are certain fault patterns in which recursive embeddings are impossible. Fortunately, we find that only some leaves of the subtree rooted at R cannot be mapped within $C[R]$. We overcome this problem by *borrowing* some nodes from $C[L]$, i.e., mapping these leaves to nonfaulty nodes in $C[L]$. While performing the L-embedding, the borrowed nodes are treated as faults. This avoids mapping the subtree of L to these nodes. As only a few nodes in $C[L]$ are borrowed (will be shown later), the L-embedding is not much affected and hence the whole embedding can be done.

B. Constructing the Embedding

We construct the ϵ-embedding for a n-cube with $n-2$ faults recursively. The base cases, where $n \leq 5$, are considered in Section III-E.

For $n \geq 6$, the first step is to find a suitable dimension to split $C[\epsilon]$. Suppose we split on some dimension i such that $C[L]$ contains p faults and $C[R]$ contains q faults (see Fig. 1). Then the subtree in $C[L]$ has to avoid p faults while that in $C[R]$ has to avoid q faults and the assigned node $H[\epsilon]$. Ideally, we want to split $C[\epsilon]$ so that $p \leq n-3$ and $q \leq n-4$. Then we can perform the L- and R-embeddings recursively. However, this is not always achievable. Thus our strategy is to ensure $p \leq n-3$ first. More precisely, we split $C[\epsilon]$ on dimension i provided

A1) $H[L]$ is nonfaulty,

A2) $C[L]$ has at most $n-3$ faults and

A3) $C[L]$ has as many faults as possible while satisfying conditions A1) and A2).

Theorem 3.3 below shows that a dimension i which satisfies conditions A1) and A2) can always be found. Among all the different possibilities for i, we choose the one which *maximizes* p. Hence, the splitting can always be done.

Theorem 3.3: There exists a dimension i such that $*^{i-1}1*^{n-i}$ has at most $n-3$ faults (i.e., $*^{i-1}0*^{n-i}$ has at least one fault) and $0^{i-1}10^{n-i}$ is a nonfaulty hypercube node.

Proof: The faults must differ in some dimension j, i.e., both $*^{j-1}0*^{n-j}$ and $*^{j-1}1*^{n-j}$ have at least one fault. If $0^{j-1}10^{n-j}$ is nonfaulty, let i be j. Otherwise, let i be a dimension other than j where $0^{i-1}10^{n-i}$ is nonfaulty. Note that this dimension always exists as faulty node $0^{j-1}10^{n-j}$ belongs to $*^{i-1}0*^{n-i}$ and there are at most $n-2$ faults. □

To illustrate the idea of conditions A1)–A3) about the splitting dimension, let us consider the following examples.

Example 3.1: $n = 6$ and there are 4 faults

$$\beta_1 = 100000$$
$$\beta_2 = 010010$$
$$\beta_3 = 111000$$
$$\beta_4 = 101100.$$

By condition A1), we cannot split on dimension 1 because of β_1. Also, all the remaining dimensions satisfy condition A2). Among them, we can choose dimensions 2 or 3 because both $*1****$ and $**1***$ contains 2 faults. Hence we have $p = 2$ in this case. □

Without loss of generality, let the cube the Example 3.1 be split on dimension 2, we have $H[\epsilon] = 000000$, $H[L] = 010000$, $H[R] = 000001$, $C[L] = *1****$ and $C[R] = *0****$. Since $C[L]$ and $C[R]$ are 5-cubes, each containing 2 faults with $H[\epsilon]$ in $C[R]$, the embedding of the 4-trees in $C[L]$ and $C[R]$ are possible by recursion by treating $H[E]$ as an extra fault in $C[R]$.

Example 3.2: $n = 6$ and the 4 faults are

$$\beta_1 = 100011$$
$$\beta_2 = 010011$$
$$\beta_3 = 001011$$
$$\beta_4 = 000011.$$

This time, no neighbor of $H[\epsilon]$ is faulty, thus condition A1) is always satisified. However, we cannot split on dimensions 5 and 6 for violation of condition A2). Among the remaining dimensions, we can choose dimensions 1, 2 or 3 because each of the subcubes $1****, *1***$ and $**1***$ contains 1 fault. Hence for this example, we have $p = 1$. □

Example 3.3: $n = 6$ and the 4 faults are

$$\beta_1 = 010000$$
$$\beta_2 = 001000$$
$$\beta_3 = 000100$$
$$\beta_4 = 011000.$$

In this case, we cannot split on dimensions 2, 3 and 4 because of condition A1). We can, however, split on dimensions 1, 5 or 6. Thus $p = 0$ in this case. □

Based on p, we have two cases in general: (1) $0 \leq p \leq 1$ and (2) $2 \leq p \leq n-3$.

For case (2) where $2 \leq p \leq n-3$, we have $q = n-2-p \leq n-4$. Hence both the L- and R-embeddings can be done recursively as also illustrated in Example 3.1.

For case (1) where $0 \leq p \leq 1$, the L-embedding is still easy but the R-embedding is more complicated. There are effectively $q + 1 = n - 1 - p > n - 3$ faults in $C[R]$ (the extra fault is from $H[\epsilon]$) and we cannot perform the R-embedding as simply as in case (2). The details for this case are described in next few sections. Section III-C will be totally devoted for a particular case when $p = 0$ for easy understanding of the general case to be described in Section III-D when $0 \leq p \leq 1$.

C. Case (1) where $n \geq 6, p = 0$

1) Patterns of the Faults: Before describing the embedding strategy in detail, we want to show that when $p = 0$, the faults are having certain properties and can always be arranged so that the faults are of a standard pattern for easy description of the algorithm. When $p = 0$, $*^{k-1}1*^{n-k}$ must contain no faults unless $0^{k-1}10^{n-k}$ is faulty.

Referring to Example 3.3, the number of faults in $*^{k-1}1*^{n-k}$ is 0 when $k = 1$, 5 or 6.

Lemma 3.1: Given that $0^{k-1}10^{n-k}$ is nonfaulty and $p = 0$, the number of faults in $*^{k-1}1*^{n-k}$ must be 0.

Proof: Note that $*^{k-1}1*^{n-k}$ cannot contain more than 0 and less than $n-2$ faults. Otherwise we could split on dimension k and have $p > 0$. Assume that $*^{k-1}1*^{n-k}$ contains $n-2$ faults (i.e., all faults have a 1 in dimension k). Then, for any dimension $j \neq k, 0^{j-1}10^{n-j}$ must not be faulty. Hence $*^{j-1}1^{n-j}$ contains 0 or $n-2$ faults (i.e., all faults agree in dimension j) because $p = 0$. This implies that all faults would be identical which is a contradiction. Thus the number of faults in $*^{k-1}1*^{n-k}$ must be 0. □

$$x \begin{cases} 0100 \ \dots \ 000 \ \dots \ 0 \\ 0010 \ \dots \ 000 \ \dots \ 0 \\ \quad \vdots \\ 0000 \ \dots \ 010 \ \dots \ 0 \\ 0zzz \ \dots \ zz0 \ \dots \ 0 \\ \quad \vdots \\ 0zzz \ \dots \ zz0 \ \dots \ 0 \end{cases}$$
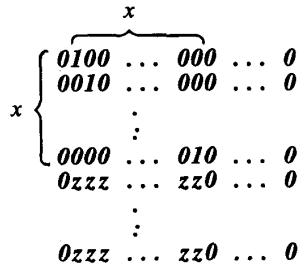
Fig. 2. The pattern of faults when $p = 0$ where z can be 0 or 1.

Suppose x is the number of faulty neighbors of $H[\epsilon] = 0^n$. For example $x = 3$ in Example 3.3. Then the faults can be arranged to the pattern shown in Fig. 2 where z can be 0 or 1.

The following lemma shows some bounds on the value of x.

Lemma 3.2:

1) $x \le n - 2$ and
2) when $p = 0, x \ge \lceil \log(n - 1) \rceil$.

Proof: (a) is obvious as the number of faults is no more than $n - 2$.

(b) For all dimensions k where $0^{k-1}10^{n-k}$, neighbor of $H[\epsilon]$, is nonfaulty, the number of faults in $*^{k-1}1*^{n-k}$, must be 0 (Lemma 3.1). Therefore all faults have 0 in dimension k. Moreover, $H[\epsilon] = 0^n$ also has 0 in this dimension. Thus the number of such dimensions, k, i.e. $n - x$, cannot be too large in order for the $n - 2$ faults and $H[\epsilon]$ to be distinct. In fact, it can be shown easily that $x \ge \lceil \log(n - 1) \rceil$. □

Note that $n - 2 \ge x \ge \lceil \log(n - 1) \rceil$ and it is not possible that all the faults have a 1 in a particular dimension, i.e., column of 1's.

2) An Example of the Embedding: First at all, we have to choose the dimension on which $C[\epsilon]$ is split with care. As will be shown later, that dimension has to satisfy more contraints besides conditions A1)–A3). After that, $C[R]$ contains all the $n - 2$ faults and the assigned node $H[\epsilon]$, and we have to split $C[R]$ again so as to dichotomize the $n - 2 + 1 > n - 3$ faults/assigned nodes. In particular, we ensure that $C[LR]$ contains as many faults as possible while not exceeding its *fault-tolerance capacity* (dimensionality of subcube minus 2). Then the LR-embedding can be done recursively. For the RR-embedding, we split $C[RR]$ again if it still contains too many faults. In general, we keep on splitting $C[R^i]$ until the R^i-embedding can be done recursively. Note that in addition to faults, $C[R^i]$ contains a number of assigned nodes, $H[\epsilon], H[R], \cdots, H[R^{i-1}]$. While after each splitting, the size of the cube and the number of faults are reduced, but the number of assigned nodes will be increased, and thus, the embedding might not be more easily done than before splitting.

Consider the following example in which $n = 7$ and the 5 faults are as below

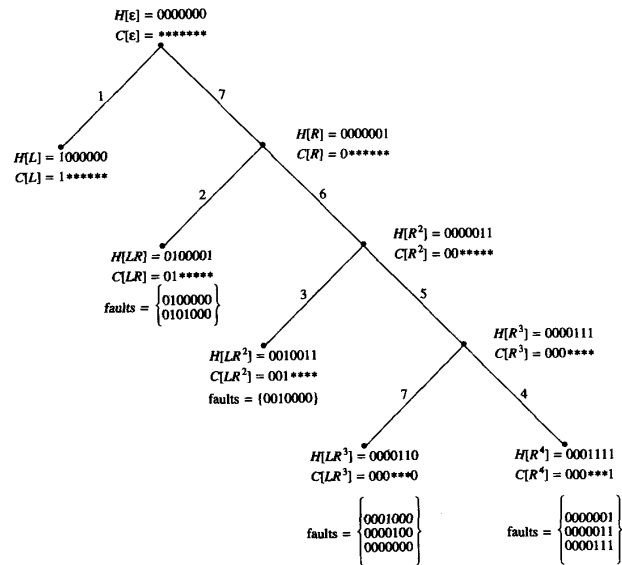0100000
0010000
0001000 .
0000100
0101000

Fig. 3. An example for a 7-cube.

In this example, $p = 0$ (as the cube cannot be split on dimensions 2–5 because of condition [A1)] and $x = 4$. The main ideas of our embedding strategy will be illustrated in this example.

The embedding can be defined as follows (Fig. 3)

$$\begin{aligned} H[\epsilon] &= 0000000 & & \\ H[L] &= 1000000 & H[R] &= 0000001 \\ C[L] &= 1****** & C[R] &= 0****** \\ H[LR] &= 0100001 & H[RR] &= 0000011 \\ C[LR] &= 01***** & C[RR] &= 00***** \\ H[LR^2] &= 0010011 & H[RR^2] &= 0000111 \\ C[LR^2] &= 001**** & C[RR^2] &= 000****. \end{aligned}$$

Without loss of generality, the cube is split on dimension 1. Thus we have $C[L] = 1******$ containing no faults while $C[R] = 0******$ containing all the faults.

The L-embedding of a 5-tree in the subcube $C[L]$ can be done without problem. On the other hand, the R-embedding of a 5-tree in the subcube $C[R]$ rooted at $H[R]$ has to avoid six nodes, namely, all the five faults and the node $H[\epsilon]$. In fact there are several choices for $H[R]$. Besides being a nonfaulty neighbor of $H[\epsilon], H[R]$ is chosen to be further away from the faults. Since $C[R]$ has to avoid more nodes than its fault tolerant capacity, subcube $C[R]$ has to be split again on a dimension such that $C[LR]$ contains as many faults as possible. Thus dimension 2 is chosen such that $C[LR]$ contains two faults whereas $C[R^2]$ contains the remaining three faults and two assigned nodes, $H[\epsilon]$ and $H[R]$. As the number of faults in $C[LR]$ is within the fault-tolerant capacity, the LR-embedding is possible by recursion. On the other hand, $C[R^2]$ has again to avoid more nodes than its fault-tolerant capacity. $C[R^2]$ has to be split accordingly into $C[LR^2]$ and $C[R^3]$ in the same way as $C[R]$.

$C[LR^2] = 001****$ containing only one fault allows an LR^2-embedding of a 3-tree in $C[LR^2]$ rooted at $H[LR^2] = 0010011$.

On the other hand, the R^3-embedding of a 3-tree rooted at $H[R^3] = 0000111$ into the $C[R^3] = 000****$ has to avoid the faults 0001000, 0000100, and the assigned nodes 0000000, 0000001 and 0000011. Among these five nodes, only 0001000 and 0000000 are at Hamming distance ≥ 3 from 0000111 and can be ignored by the embedding (Theorem 3.2). Thus the R^3-embedding cannot be performed by straightforward recursion. Fortunately, the faults and the assigned hypercube nodes can be somewhat separated by splitting $C[R^3]$ on dimension n (i.e., dimension 7). That is, let $H[LR^3] = 0000110, C[LR^3] = 000***0$ and $H[RR^3] = 0001111, C[RR^3] = 000***1$.

$C[LR^3]$ will contain faults 0001000, 0000100 and the designated root 0000000. Since 0000000 and 0001000 are at Hamming distance ≥ 2 from 0000110, they can be ignored by the embedding according to Theorem 3.2. Thus, the LR^3-embedding can be performed successfully.

On the other hand, $C[R^4]$ contains the assigned nodes 0000001, 0000011 and 0000111. As 0000001 and 0000011 are of Hamming distance ≥ 2 from 00001111 and can be ignored, the R^4-embedding can also be done recursively by considering 0000111 as the only fault in the $C[R^4]$. Thus, the embedding can be done.

Note that the splittings of the n-cube are done on certain dimensions with specific properties. In this example, the cube is initially split on dimension 1 and then 2, 3 etc, in a sequential order. This phenomenon is not by coincidence but by an arrangement of dimensions into a standard pattern.

Also note that except $H[L]$ and $H[LR^3]$, all the assigned nodes are in $******1$. Since all faults are in $******0$, these assigned nodes cannot be faulty. In general, we can also show that $H[L]$ and $H[LR^3]$ are nonfaulty by arranging the dimensions properly.

Finally, one should also observe that $C[LR^3]$ and $C[R^4]$ somewhat split the nodes to be avoided into two groups, one contains faults and the other assigned nodes. But unfortunately $C[LR^3]$ and $C[R^4]$ may still have to avoid more nodes that their fault tolerant capacity. However, the embedding is performed such that $H[LR^3]$ and $H[R^4]$ are assigned to hypercube nodes which are sufficiently far away from the faults and the assigned nodes. As a result, the numbers of effective nodes to be avoided are less than their fault-tolerant capacity.

3) Description of the Embedding for $p = 0$: As discussed before, we try to split subcube $C[\epsilon]$ on a dimension satisfying conditions A1–A3. Without loss of generality, assume the dimensions are arranged into the standard pattern (as shown in Fig. 2) so that we first split $C[R^{i-1}]$ on dimension i for $i = 1, \cdots, \lfloor n/2 \rfloor$. The splitting of $C[R^{i-1}]$ for some $i, 1 \leq i \leq \lfloor n/2 \rfloor$, might not be necessary if $C[R^{i-1}]$ contains no more than $n - 2i - 2$ faults as the total number of faults and assigned nodes to be avoided in the R^{i-1}-embedding will be less than $C[R^{i-1}]$'s fault-tolerant capacity. However, for easy description of the algorithm, it is assumed that $C[R^{i-1}]$ is split on all dimensions $i = 1, \cdots, \lfloor n/2 \rfloor$ in sequence.
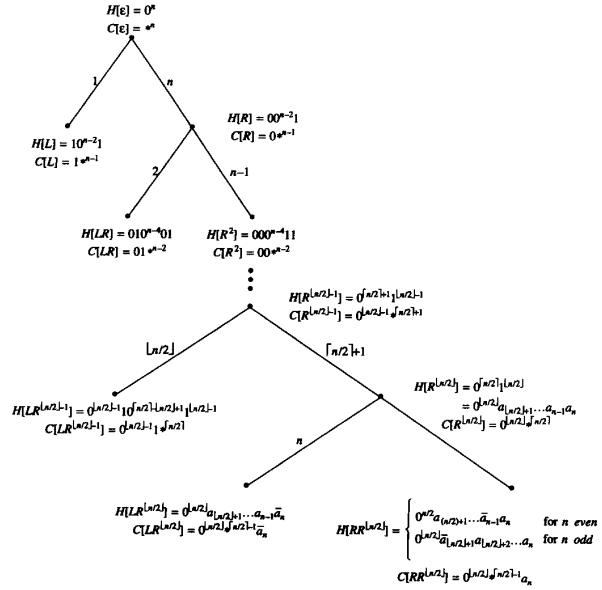

Fig. 4. Description of the embedding for $p = 0$.

Thus $C[LR^{i-1}]$ and $C[RR^{i-1}]$ are $(n-i)$-cubes, $0^{i-1}1*^{n-i}$ and $0^{i-1}0*^{n-i}$ respectively for any particular $i, 1 \leq i \leq \lfloor n/2 \rfloor$. That means the subtree rooted at LR^{i-1} is to be embedded within $0^{i-1}1*^{n-i}$. When splitting $C[R^{i-1}]$, some of the faults within the subcube fall into $C[LR^{i-1}]$. We shall show that $C[LR^{i-1}]$ gets as many faults as possible but no more than $n - i - 2$ (the fault-tolerance capacity of $C[LR^{i-1}]$). Moreover, it does not contain any assigned nodes except $H[LR^{i-1}]$ itself. Hence, the LR^i-embedding can be done recursively.

On the other hand, $C[RR^{i-1}] = 0^{i-1}0*^{n-i}$ will get the remaining faults together with the assigned nodes: $H[\epsilon], H[R], \cdots, H[R^{i-1}]$. Note that these are the only assigned nodes we need to take care of while considering the RR^{i-1}-mapping. It is because the subtrees rooted at $L, LR, \cdots, LR^{i-1}$ will not be embedded within $C[RR^{i-1}]$ as $C[L], C[LR], \cdots, C[LR^{i-1}]$ and $C[RR^{i-1}]$ are all disjoint subcubes. Consider dimensions $i + 1$ to n. We must map RR^{i-1} to a neighbor of $H[R^{i-1}]$ on one of these dimensions. Among them, we choose a dimension (i.e., dimension $n+1-i$) so that the Hamming distances between $H[RR^{i-1}]$ and most of the faults are greater than those between $H[R^{i-1}]$ and the corresponding faults. This way, the roots of lower subtrees will be further and further away from most of the faults.

By the time $i = \lfloor n/2 \rfloor$, there will be approximately $n/2$ assigned nodes and at most $(n/2) - 2$ faults within $C[R^{\lfloor n/2 \rfloor}]$. By splitting on a proper dimension, (i.e., dimension n) and choosing a suitable $H[RR^{\lfloor n/2 \rfloor}]$ (depending on whether n is even or odd), it can be shown that $C[LR^{\lfloor n/2 \rfloor}]$ contains all the remaining faults while $C[RR^{\lfloor n/2 \rfloor}]$ contains all the assigned nodes. It can be shown that the $LR^{\lfloor n/2 \rfloor}$-, and $RR^{\lfloor n/2 \rfloor}$-embeddings can be done recursively. The whole ϵ-embedding is then completed.

Refer to Fig. 4 for a pictorial view of the embedding for $p = 0$ after the dimension arrangement.

/* defining the tree as outlined in Fig. 4 */

1) For $i = 1, \cdots, \lfloor n/2 \rfloor$, let

$$H[LR^{i-1}] = 0^{i-1}10^{n-2i}01^{i-1} \quad C[LR^{i-1}] = 0^{i-1}1 *^{n-i}$$
$$H[RR^{i-1}] = 0^{i-1}00^{n-2i}11^{i-1} \quad C[RR^{i-1}] = 0^{i-1}0 *^{n-i}.$$

2) For $i = 1, \cdots, \lfloor n/2 \rfloor$, recursively perform the LR^{i-1}-embedding, i.e., embed the $(n-1-i)$-tree rooted at LR^{i-1} into the $(n-i)$-cube $C[LR^{i-1}]$.

3) Suppose

$$H[R^{\lfloor n/2 \rfloor}] = 0^{\lceil n/2 \rceil}1^{\lfloor n/2 \rfloor} = 0^{\lfloor n/2 \rfloor}a_{\lfloor n/2 \rfloor+1} \cdots a_n.$$

Then, we let

$$H[LR^{\lfloor n/2 \rfloor}] = 0^{\lfloor n/2 \rfloor}a_{\lfloor n/2 \rfloor+1} \cdots a_{n-1}\bar{a}_n,$$
$$C[LR^{\lfloor n/2 \rfloor}] = 0^{\lfloor n/2 \rfloor} *^{\lceil n/2 \rceil-1} \bar{a}_n,$$

$$H[RR^{\lfloor n/2 \rfloor}] = \begin{cases} 0^{n/2}a_{n/2+1} \cdots a_{n-2}\bar{a}_{n-1}a_n \\ \quad \text{for } n \text{ even} \\ 0^{\lfloor n/2 \rfloor}\bar{a}_{\lfloor n/2 \rfloor+1}a_{\lfloor n/2 \rfloor+2} \cdots a_n \\ \quad \text{for } n \text{ odd} \end{cases}$$

and

$$C[RR^{\lfloor n/2 \rfloor}] = 0^{\lfloor n/2 \rfloor} *^{\lceil n/2 \rceil-1} a_n.$$

4) Recursively perform the $RR^{\lfloor n/2 \rfloor}$-embedding, i.e., embed the $(\lceil n/2 \rceil - 2)$-tree rooted at $RR^{\lfloor n/2 \rfloor}$ in the $(\lceil n/2 \rceil - 1)$-cube, $C[RR^{\lfloor n/2 \rfloor}]$.

5) Recursively perform the $LR^{\lfloor n/2 \rfloor}$-embedding, i.e., embed the $(\lceil n/2 \rceil - 2)$-tree rooted at $LR^{\lfloor n/2 \rfloor}$ in the $(\lceil n/2 \rceil - 1)$-cube, $C[LR^{\lfloor n/2 \rfloor}]$.

D. Case (1) when $n \geq 6, 0 \leq p \leq 1$

After describing the embedding algorithm for $p = 0$, we shall generalize the method to the case when $0 \leq p \leq 1$.

1) Patterns of the Faults: This section can be considered as a part of the correctness proof. Basically, we want to show that when $p \leq 1$, the faults are having certain properties and can always be arranged (Section III-D2) so that the faults are of a standard pattern for easy description of the algorithm.

Note that the faults must exhibit one of the following three patterns. One of the patterns is shown in Fig. 2 when $p = 0$.

When $p = 1, *^{k-1}1*^{n-k}$ can contain either 0, 1 or $n - 2$ faults if $0^{k-1}10^{n-k}$ is nonfaulty. Otherwise, if $*^{k-1}1*^{n-k}$ contains more than 1 but less than $n - 2$ faults, we would have split on dimension k according to condition A3) and have $p > 1$. The two other possible patterns after dimensions arrangement are shown in Fig. 5.

Thus, for all $1 \leq i \leq n, *^{i-1}1*^{n-i}$ cannot contain more than p and less than $n - 2$ faults unless $0^{i-1}10^{n-i}$ is faulty (as dimension 1 in Example 3.1.). As in Example 3.2., $p = 1, 0^{k-1}10^{n-k}$ is nonfaulty for every k. Hence the number of faults in $*^{k-1}1*^{n-k}$ is 0 (for $k = 4$), 1 (for $k = 1, 2,$ or 3) or $n - 2$ (for $k = 5$ or 6).
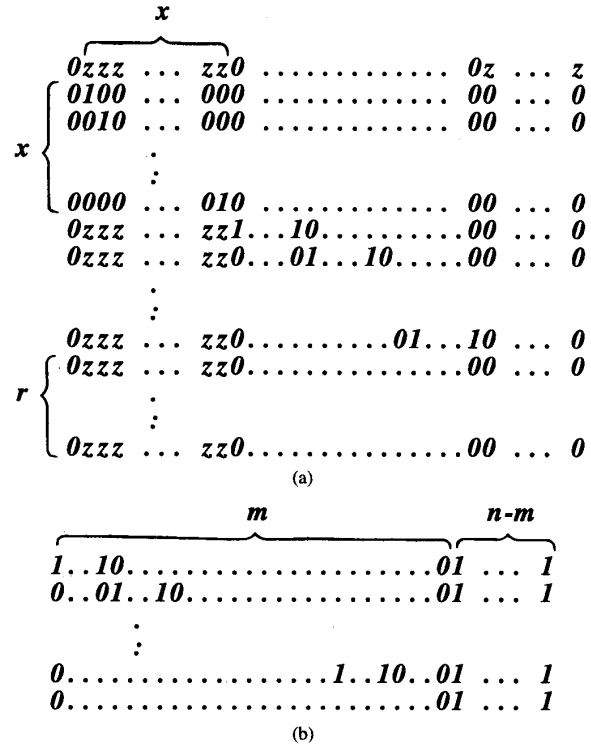

(a)

(b)

Fig. 5. The 2 possible patterns when $p = 1$, where z can be 0 or 1 and r can be 0.

2) Arranging Dimensions: We choose dimension 1 so that $1*^{n-1}$ contains p faults and 10^{n-1} is nonfaulty. This allows us to split $C[\epsilon]$ on dimension 1 to arrive at 2 subcubes $C[L]$ and $C[R]$, where $C[L] = 1*^{n-1}$ and $C[R] = 0*^{n-1}$.

$C[L]$ contains p faults and $C[R] = 0*^{n-1}$ contains $n-2-p$ faults. For the benefit of easy isolation of nonfaulty nodes in $C[R]$ during embedding, we would like to identify certain dimensions such that all these $n - 2 - p$ faults have the same bit pattern. Consider a dimension k such that $2 \leq k \leq n$ and $0*^{k-2}1*^{n-k}$ contains all the $n-2-p$ faults in $0*^{n-1}$, i.e., all the faults have a 1 in their k-bit position. Let there be $n - m$ such dimensions. If $n - m > 0$, we let them be the rightmost $n - m$ dimensions (Example 3.2 and Fig. 5(a) and 6(b)) so that we can easily choose nonfaulty nodes from $0*^{m-1}0^{n-m}$ during the embedding. In the case where $n - m = 0$, we would instead find a dimension k such that $0*^{k-2}0*^{n-k}$ contains all $n - 2 - p$ faults and swap it to the rightmost dimension (see Example 3.3. and Figs. 2, 5(a) and 6(b)). Again, this allows us to find nonfaulty nodes from $0*^{n-2}1$ for embedding. $\square$

Note that this dimension k may not exist if we have not chosen dimension 1 carefully. Let us consider the following example.

Example 3.4: $n = 6$ and the 4 faults are

$$\beta_1 = 111000$$
$$\beta_2 = 010000$$
$$\beta_3 = 001000$$
$$\beta_4 = 010111. \qquad \square$$

dimension

$n > m$

(a)

dimension

$n = m$

(b)

Fig. 6. The bit pattern of faults after arranging dimension 1 and the last $\max(n - m, 1)$ dimensions. The faults are listed row by row. z is 0 or 1 depending on p is 0 or 1. The value of x and each of the y's can be 0 or 1 subject to other conditions.

However, it can be proved (in Appendix A) that such a dimension k can always be found by properly choosing dimension 1. In Example 3.4, we can swap dimensions 1 and 4 (and also exchange row 1 with row 4) so that the faults have the following pattern (like the pattern in Fig. 5(a) and 6(b)). Then we can choose nonfaulty nodes from $0*^{n-2}1$.

$$\beta_4 = 110011$$

$$\beta_2 = 010000$$

$$\beta_3 = 001000$$

$$\beta_1 = 011100.$$

The next step is to arrange dimensions 2 through m such that for all $i, 2 \leq i \leq m, 0^{i-1}1*^{n-i}$ contains as many faults as that in $0^{i-1}*^{k-i}1*^{n-k}$ for all k where $i \leq k \leq m$. The arrangement allocates more faults to larger subcubes and guarantees that $0^{i-1}1*^{n-i}$ contains no more than $n - i - 2$ faults (Appendix B). Hence we can recursively embed an $(n-i-1)$-tree within this $(n - i)$-cube.

More precisely, the dimensions are arranged such that

B1) subcube $1*^{n-1}$ contains p faults and 10^{n-1} is non-faulty;

B2) subcube

$$\begin{cases} 0*^{m-1}1^{n-m} & \text{if } n > m \\ 0*^{m-2}0 & \text{if } n = m \end{cases}$$

contains $n - 2 - p$ faults;

B3) for all $i, 2 \leq i \leq m$, subcube $0^{i-1}1*^{n-i}$ contains at least as many faults as in $0^{i-1}*^{k-i}1*^{n-k}$ for all k such that $i \leq k \leq m$.

Lemma A.1 (in Appendix A) shows that conditions (B1) and (B2) are achievable by carefully choosing dimension 1. Lemma A.2 (in Appendix A) shows that dimensions 2 through m can be arranged according to the number of faults in the subcubes $0^{i-1}1*^{n-i}$ so as to satisfy condition (B3) without violating conditions (B1) and (B2).

3) Description of the Embedding for $0 \leq p \leq 1$ (Case 1): As discussed in Section III-C, we first try to split $C[\epsilon]$ on a dimension satisfying conditions A1-A3. If $p > 1$, it is done and so we shall concentrate on the case where $0 \leq p \leq 1$ (Case 1). The embedding for this case is almost the same as the embedding for the case when $p = 0$ (Section III-C). The only difference is that the embedding in performed as if $C[\epsilon]$ is an m-cube instead of an n-cube (thus, the embedding would be exactly the same if $n = m$).

Assume the dimensions of the n-cube $C[\epsilon]$ have been arranged such that the conditions B1, B2 and B3 are satisfied. $C[R^{i-1}]$ is split on dimension i for $i = 1, \cdots, \lfloor m/2 \rfloor$ (instead of $\lfloor n/2 \rfloor$ for the case when $p = 0$) into $C[LR^{i-1}] = 0^{i-1}1*^{n-i}$ and $C[RR^{i-1}] = 0^{i-1}0*^{n-i}$. $C[LR^{i-1}]$ will get as many faults as possible but no more than $n - i - 2$ (the fault-tolerance capacity of $C[LR^{i-1}]$). Moreover, it does not contain any assigned nodes except $H[LR^{i-1}]$ itself. Hence, the LR^i-embedding can be done recursively. On the other hand, $C[RR^{i-1}] = 0^{i-1}0*^{n-i}$ will get the remaining faults together with the assigned nodes: $H[\epsilon], H[R], \cdots, H[R^{i-1}]$.

By the time $i = \lfloor m/2 \rfloor$, there will be approximately $m/2$ assigned nodes and at most $n - 2 - m/2$ faults within $C[R^{\lfloor m/2 \rfloor}]$. By splitting on dimension n and choosing a suitable $H[RR^{\lfloor m/2 \rfloor}]$ (depending on whether m is even or odd), it can be shown that $C[LR^{\lfloor m/2 \rfloor}]$ contains all the remaining faults while $C[RR^{\lfloor m/2 \rfloor}]$ contains all the assigned nodes. However, many of the faults will be very far from $H[LR^{\lfloor m/2 \rfloor}]$. In fact, when m is even, most of the faults will be too far to affect the embedding. When m is odd, there are also several faults which can be ignored. Anyway, the number of *effective* faults/assigned nodes within each subcube does not exceed their fault-tolerance capacities. Hence the $LR^{\lfloor m/2 \rfloor}$-, and $RR^{\lfloor m/2 \rfloor}$-embeddings can also be done recursively. The whole ϵ-embedding is then completed.

Refer to Fig. 7 for a pictorial view of the general embedding for $0 \leq p \leq 1$ after the dimension arrangement. Note that the following algorithm and the tree outlined in Fig. 7 are the generalization of the algorithm and the corresponding tree for $p = 0$ given in Section III-C3. The proof of correctness is given in Section III-D4.

$H[\varepsilon] = 0^n$
$C[\varepsilon] = *^n$

$H[R] = 00^{m-2}10^{n-m}$
$C[R] = 0*^{n-1}$

$H[L] = 10^{m-2}10^{n-m}$
$C[L] = 1*^{n-1}$

$H[LR] = 010^{m-4}010^{n-m}$
$C[LR] = 01*^{n-2}$

$H[R^2] = 000^{m-4}110^{n-m}$
$C[R^2] = 00*^{n-2}$

$H[R^{\lfloor m/2 \rfloor - 1}] = 0^{\lfloor m/2 \rfloor}1^{\lfloor m/2 \rfloor -1}0^{n-m}$
$C[R^{\lfloor m/2 \rfloor -1}] = 0^{\lfloor m/2 \rfloor -1}*^{n-\lfloor m/2 \rfloor +1}$

$H[LR^{\lfloor m/2 \rfloor -1}] = 0^{\lfloor m/2 \rfloor -1}10^{m-2\lfloor m/2 \rfloor +1}1^{\lfloor m/2 \rfloor -1}0^{n-m}$
$C[LR^{\lfloor m/2 \rfloor -1}] = 0^{\lfloor m/2 \rfloor -1}1*^{n-\lfloor m/2 \rfloor}$

$H[R^{\lfloor m/2 \rfloor}] = 0^{\lfloor m/2 \rfloor}1^{\lfloor m/2 \rfloor}0^{n-m}$
$= 0^{\lfloor m/2 \rfloor}a_{\lfloor m/2 \rfloor +1}\cdots a_{n-1}a_n$
$C[R^{\lfloor m/2 \rfloor}] = 0^{\lfloor m/2 \rfloor}*^{n-\lfloor m/2 \rfloor}$

$H[LR^{\lfloor m/2 \rfloor}] = 0^{m/2}a_{\lfloor m/2 \rfloor +1}\cdots a_{n-1}\bar{a}_n$
$C[LR^{\lfloor m/2 \rfloor}] = 0^{m/2}*^{n-\lfloor m/2 \rfloor -1}\bar{a}_n$

$H[RR^{\lfloor m/2 \rfloor}] = \begin{cases} 0^{m/2}a_{\lfloor m/2 \rfloor +1}\cdots \bar{a}_{n-1}a_n & \text{for } m \text{ even} \\ 0^{m/2}\bar{a}_{\lfloor m/2 \rfloor +1}a_{\lfloor m/2 \rfloor +2}\cdots a_n & \text{for } m \text{ odd} \end{cases}$

$C[RR^{\lfloor m/2 \rfloor}] = 0^{m/2}*^{n-\lfloor m/2 \rfloor -1}a_n$

Fig. 7. Description of the embedding for $0 \le p \le 1$.

/* defining the tree as outlined in Fig. 7. */

1) For $i = 1, \cdots, \lfloor m/2 \rfloor$, let

$$H[LR^{i-1}] = 0^{i-1}10^{m-2i}01^{i-1}0^{n-m}$$
$$C[LR^{i-1}] = 0^{i-1}1*^{n-i}$$
$$H[RR^{i-1}] = 0^{i-1}00^{m-2i}11^{i-1}0^{n-m}$$
$$C[RR^{i-1}] = 0^{i-1}0*^{n-i}.$$

2) For $i = 1, \cdots, \lfloor m/2 \rfloor$, recursively perform the LR^{i-1}-embedding, i.e., embed the $(n-1-i)$-tree rooted at LR^{i-1} into the $(n-i)$-cube $C[LR^{i-1}]$.

3) Suppose $H[R^{\lfloor m/2 \rfloor}] = 0^{\lceil m/2 \rceil}1^{\lfloor m/2 \rfloor}0^{n-m} = 0^{\lfloor m/2 \rfloor}a_{\lfloor m/2 \rfloor +1} \cdots a_n$.
   Then, we let

$$H[LR^{\lfloor m/2 \rfloor}] = 0^{\lfloor m/2 \rfloor}a_{\lfloor m/2 \rfloor +1} \cdots a_{n-1}\bar{a}_n,$$
$$C[LR^{\lfloor m/2 \rfloor}] = 0^{\lfloor m/2 \rfloor}*^{n-\lfloor m/2 \rfloor -1}\bar{a}_n,$$
$$H[RR^{\lfloor m/2 \rfloor}] = \begin{cases} 0^{m/2}a_{m/2+1} \cdots a_{n-2}\bar{a}_{n-1}a_n \\ \quad \text{for } m \text{ even} \\ 0^{\lfloor m/2 \rfloor}\bar{a}_{\lfloor m/2 \rfloor +1}a_{\lfloor m/2 \rfloor +2} \cdots a_n \\ \quad \text{for } m \text{ odd} \end{cases}$$

and

$$C[RR^{\lfloor m/2 \rfloor}] = 0^{\lfloor m/2 \rfloor}*^{n-\lfloor m/2 \rfloor -1}a_n.$$

4) Recursively perform the $RR^{\lfloor m/2 \rfloor}$-embedding, i.e., embed the $(n - \lfloor m/2 \rfloor - 2)$-tree rooted at $RR^{\lfloor m/2 \rfloor}$ in the $(n - \lfloor m/2 \rfloor - 1)$-cube, $C[RR^{\lfloor m/2 \rfloor}]$.

5) Recursively perform the $LR^{\lfloor m/2 \rfloor}$-embedding, i.e., embed the $(n - \lfloor m/2 \rfloor - 2)$-tree rooted at $LR^{\lfloor m/2 \rfloor}$ in the $(n - \lfloor m/2 \rfloor - 1)$-cube, $C[LR^{\lfloor m/2 \rfloor}]$.

4) Correctness: There are a number of observations we can make to argue the validity of the above embedding approach. However, this section contains a lot of details and can be skipped in the initial reading of the paper.

1) $m \ge 3$ (Lemma 3.3)
2) For $i = 1, \cdots, \lfloor m/2 \rfloor +1$, $H[LR^{i-1}]$ and $H[RR^{i-1}]$ are adjacent to $H[R^{i-1}]$ (Lemmas 3.4(a), 3.5(a));
3) For $i = 1, \cdots, \lfloor m/2 \rfloor +1$, $H[LR^{i-1}]$ and $H[RR^{i-1}]$ are nonfaulty (Lemmas 3.4(c), 3.5(c));
4) For $i = 1, \cdots, \lfloor m/2 \rfloor +1$, $H[LR^{i-1}]$ and $H[RR^{i-1}]$ are distinct (Lemmas 3.4(d), 3.5(d));
5) It is always possible to perform the $RR^{\lfloor m/2 \rfloor}$-embedding (Lemma 3.6);
6) It is always possible to perform the LR^{i-1}-embedding for $i = 1, \cdots, \lfloor m/2 \rfloor$ (Lemma B.3 in Appendix B);
7) It is always possible to perform the $LR^{\lfloor m/2 \rfloor}$-embedding with node borrowing (Lemma B.4 in Appendix B, Lemmas 3.7 and 3.8).

The above observations allow us to conclude that the embedding of an $(n-1)$-tree into an n-cube containing $n-2$ faults, with the root of the tree mapped to a specified nonfaulty hypercube node S, can be done when $n \ge 6, p \le 1$.

Lemma 3.3: For $n \ge 6, n - m \le 3$ and $m \ge 3$.

Proof: Assume to the contrary that $n - m \ge 4$. Note that all faults in 0^{n-1} lie in $0*^{m-1}1^{n-m}$ (condition (B2)). In other words, each fault is either in $0*^{k-2}1*^{m-k}1^{n-m}$ for some $2 \le k \le m$ or $0^m 1^{n-m}$. We can show that $0*^{k-2}1*^{m-k}1^{n-m}$ contains at most 1 fault (Lemma B.2(a) in Appendix B). Therefore the maximum possible number of faults in $0*^{n-1}$ is $\le m - 1 + 1 = m \le n - 4 < n - 2 - p$ which is impossible. Therefore $n - m \le 3$ and hence $m \ge 3$ (as $n \ge 6$). $\square$

Lemma 3.3 ensures that $\lfloor m/2 \rfloor \ge 1$ and hence the above embedding strategy is well-defined.

Lemma 3.4: For $i = 1, \cdots, \lfloor m/2 \rfloor$,

1) $H[LR^{i-1}]$ and $H[RR^{i-1}]$ are neighbors of $H[R^{i-1}]$ on dimensions i and $m + 1 - i$ respectively,
2) $C[LR^{i-1}]$ and $C[RR^{i-1}]$ are disjoint subcubes obtained by splitting $C[R^{i-1}]$ on dimension i,
3) $H[LR^{i-1}]$ and $H[RR^{i-1}]$ are nonfaulty,
4) $H[LR^{i-1}]$ and $H[RR^{i-1}]$ avoid all the assigned nodes above level i (i.e., those $H[T]$ where T is a tree node at level $j < i$) that lie in $C[R^{i-1}]$.

Proof: (a) and (b) are obvious by noting that

$$H[R^{i-1}] = 0^{i-1}00^{m-2i}01^{i-1}0^{n-m}, H[LR^{i-1}]$$
$$= 0^{i-1}10^{m-2i}01^{i-1}0^{n-m}, H[RR^{i-1}]$$
$$= 0^{i-1}00^{m-2i}11^{i-1}0^{n-m}$$

and that

$$C[R^{i-1}] = 0^{i-1}*^{n-i+1},$$
$$C[LR^{i-1}] = 0^{i-1}1*^{n-i},$$
$$C[RR^{i-1}] = 0^{i-1}0*^{n-i}.$$

(c) Since $H[LR^{i-1}]$ and $H[RR^{i-1}]$ (except $H[L]$) lie in

$$\begin{cases} 0*^{m-1}0^{n-m} & n > m \\ 0*^{m-1}1 & n = m \end{cases}$$

while all faults in $0*^{n-1}$ are in

$$\begin{cases} 0*^{m-1}1^{n-m} & n > m \\ 0*^{m-1}0 & n = m \end{cases}$$

by condition (B2), these nodes are nonfaulty. Also, $H[L]$ is chosen as nonfaulty by condition (B1).

(d) Since $C[L], C[LR], \cdots, C[LR^{i-2}]$ and $C[RR^{i-2}] = C[R^{i-1}]$ are disjoint subcubes (a corollary of Lemma 3.4(b)), the subtrees rooted at $L, LR, \cdots, LR^{i-2}$ will not be mapped to nodes in $C[R^{i-1}]$. Hence, $H[R^j] = 0^j 0^{m-2j} 1^j 0^{n-m}$ for $j = 0, \cdots, i-1$ are the only assigned nodes above level i that lie in $C[R^{i-1}]$. After splitting $C[R^{i-1}]$ on dimension i, they will fall into $C[RR^{i-1}]$ (hence distinct from $H[LR^{i-1}]$) and are at Hamming distance $i - j > 0$ from $H[RR^{i-1}]$. □

Lemma 3.5:

1) $H[LR^{\lfloor m/2 \rfloor}]$ and $H[RR^{\lfloor m/2 \rfloor}]$ are neighbors of $H[R^{\lfloor m/2 \rfloor}]$ on dimension n and dimension

$$\begin{cases} n - 1 & m \text{ even} \\ \lfloor m/2 \rfloor + 1 & m \text{ odd} \end{cases}$$

respectively;
2) $C[LR^{\lfloor m/2 \rfloor}]$ and $C[RR^{\lfloor m/2 \rfloor}]$ are disjoint subcubes obtained by splitting $C[R^{\lfloor m/2 \rfloor}]$ on dimension n;
3) $H[LR^{\lfloor m/2 \rfloor}]$ and $H[RR^{\lfloor m/2 \rfloor}]$ are nonfaulty;
4) $H[LR^{\lfloor m/2 \rfloor}]$ and $H[RR^{\lfloor m/2 \rfloor}]$ avoid all the assigned nodes above level $\lfloor m/2 \rfloor + 1$ that lie in $C[R^{\lfloor m/2 \rfloor}]$.

Proof: (a) and (b) are obvious from Step (3) of our embedding algorithm.

(c) $H[LR^{\lfloor m/2 \rfloor}]$ is, in fact, the hypercube node

$$\begin{cases} 0^{\lceil m/2 \rceil} 1^{\lfloor m/2 \rfloor} 0^{n-m-1} 1 & \text{if } n > m \\ 0^{\lceil m/2 \rceil} 1^{\lfloor m/2 \rfloor - 1} 0 & \text{if } n = m \end{cases}$$

and can be made nonfaulty (Lemma A.3 in Appendix A). $H[RR^{\lfloor m/2 \rfloor}]$ is nonfaulty using the same argument as in Lemma 3.4(c).

(d) The only assigned nodes above level $\lfloor m/2 \rfloor + 1$ that lie in $C[R^{\lfloor m/2 \rfloor}]$ are $H[R^j]$ for $j = 1, \cdots, \lfloor m/2 \rfloor$ and if $n > m, H[\epsilon]$ also. By splitting $C[R^{\lfloor m/2 \rfloor}]$ on dimension n, they will be in $C[RR^{\lfloor m/2 \rfloor}]$ (hence distinct from $H[LR^{\lfloor m/2 \rfloor}]$) and at Hamming distance

$$\begin{cases} (m/2) - 2 & j = 1 \text{ and } n = m \text{ and } m \text{ even} \\ \lfloor m/2 \rfloor - j + 1 & \text{otherwise} \end{cases}$$

from $H[RR^{\lfloor m/2 \rfloor}]$. □

Lemma 3.6 proves that the $RR^{\lfloor m/2 \rfloor}$-embedding can be done recursively by showing that the number of faults/assigned nodes in the subcube $C[RR^{\lfloor m/2 \rfloor}]$ is no more than its fault-tolerance capacity, $n - \lfloor m/2 \rfloor - 3$.

Lemma 3.6: It is always possible to perform the $RR^{\lfloor m/2 \rfloor}$-embedding.

Proof: Note that $C[RR^{\lfloor m/2 \rfloor}] = 0^{\lfloor m/2 \rfloor} *^{n-\lfloor m/2 \rfloor - 1} a_n$ where

$$a_n = \begin{cases} 0 & \text{if } n > m \\ 1 & \text{if } n = m \end{cases}$$

and hence does not contain any fault (condition (B2)). However, it contains the assigned nodes $H[R^j]$ for $j = 1, \cdots, \lfloor m/2 \rfloor$ and if $n > m, H[\epsilon]$ also.

Moreover,

$$H[RR^{\lfloor m/2 \rfloor}] = \begin{cases} 0^{(m/2)} 1^{(m/2)} 0^{n-m-2} 10 \\ \quad \text{if } m \text{ even and } n - m \geq 2 \\ 0^{(m/2)} 1^{(m/2)-1} 00 \\ \quad \text{if } m \text{ even and } n - m = 1 \\ 0^{(m/2)} 1^{(m/2)-2} 01 \\ \quad \text{if } m \text{ even and } n - m = 0 \\ 0^{\lfloor m/2 \rfloor} 1^{\lceil m/2 \rceil} 0^{n-m} \\ \quad \text{if } m \text{ odd}. \end{cases}$$

By rewriting $H[R^j] = 0^j 0^{m-2j} 1^j 0^{n-m}$ as $0^{\lfloor m/2 \rfloor} 0^{\lceil m/2 \rceil - j} 1^j 0^{n-m}$, it is easy to see that $H[R^j]$ (except $H[\epsilon]$ and $H[R]$ when $n = m$ and m is even) is at Hamming distance $\lfloor m/2 \rfloor - j + 1$ from $H[RR^{\lfloor m/2 \rfloor}]$. When $n = m$ and m is even, $H[R]$ is at Hamming distance $\lfloor m/2 \rfloor - 2$ from $H[RR^{\lfloor m/2 \rfloor}]$ and $H[\epsilon]$ is outside $C[RR^{\lfloor m/2 \rfloor}]$. Thus they will not affect the $RR^{\lfloor m/2 \rfloor}$-embedding. Since all the other assigned nodes have distinct Hamming distances from $H[RR^{\lfloor m/2 \rfloor}]$, there are $\leq n - \lfloor m/2 \rfloor - 3$ assigned nodes at Hamming distance $\leq n - \lfloor m/2 \rfloor - 3$ from $H[RR^{\lfloor m/2 \rfloor}]$. Hence the $RR^{\lfloor m/2 \rfloor}$-embedding can be done recursively. □

In a similar fashion, Lemma B.3 (in Appendix B) shows that the LR^{i-1}-embedding for $i = 1, \cdots, \lfloor m/2 \rfloor$ can be done recursively and Lemma B.4 (in Appendix B) shows that the $LR^{\lfloor m/2 \rfloor}$-embedding is also possible when n is even or $x < n - 2$ (where x is the number of faulty neighbors of $H[\epsilon]$). Basically, the idea is to argue that the number of faults/assigned nodes in each of the subcubes is no more than its fault-tolerance capacity.

When n is odd and $x = n - 2$, for the $LR^{\lfloor m/2 \rfloor}$-embedding, we have to use node borrowing. Before we describe the embedding strategy for this case, let us consider the embedding of an $(n - 1)$-tree into an n-cube with $n - 1$ faults having a specific pattern and the tree root is mapped to a specific node (Lemma 3.7). The bit patterns of the specified root and faults can be described as follows.

$$\begin{array}{ll} \text{root} & 0111\cdots 11 \\ \text{faults} & 0100\cdots 00 \\ & 0010\cdots 00 \\ & \quad \vdots \\ & 0000\cdots 10 \\ & 0000\cdots 01. \end{array}$$

Lemma 3.7: For all $n \geq 3$, it is always possible to embed an $(n-1)$-tree (except the leaf L^{n-2}) into an n-cube with $n-1$ faults, $0^i 10^{n-1-i}$ for $i = 1, \cdots, n-1$ such that $H[\epsilon] = 01^{n-1}$ and $H[L^{n-3}] = 0^{n-2} 11$.
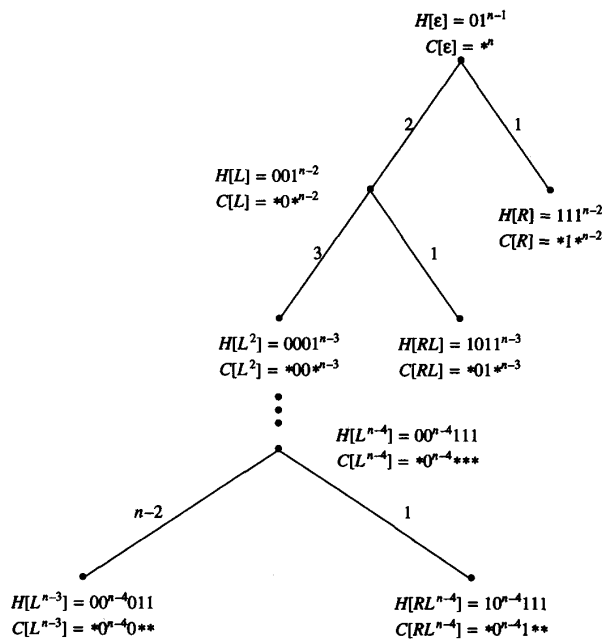
Fig. 8. Description of the embedding.

Proof: The embedding can be done by having, for $i = 1, \cdots, n - 3$,

$$H[LL^{i-1}] = 00^{i-1}01^{n-1-i}$$
$$C[LL^{i-1}] = *0^{i-1}0*^{n-1-i}$$
$$H[RL^{i-1}] = 10^{i-1}11^{n-1-i}$$
$$C[RL^{i-1}] = *0^{i-1}1*^{n-1-i}.$$

Refer to Fig. 8 for a pictoral representation of the embedding. It can be shown that for $i = 1, \cdots, n - 3, H[LL^{i-1}]$ and $H[RL^{i-1}]$ are neighbors of $H[L^{i-1}]$ on dimension $i + 1$ and 1 respectively. In addition, since these nodes have at least two 1's in their labels, they are nonfaulty.

Since for each $i = 1, \cdots, n - 3$, the subcube $C[RL^{i-1}] = *0^{i-1}1*^{n-i-1}$ contains only one fault 0^i10^{n-i-1} (which can be ignored by Theorem 3.2) and an assigned node $H[L^{i-1}] = 0^i1^{n-i}$, the RL^{i-1}-embedding, for $i = 1, \cdots, n-3$, is possible by recursion. For the L^{n-3}-embedding, $H[L^{n-3}] = 00^{n-3}11$ and $C[L^{n-3}] = *0^{n-3}*^2$ contains two faults $00^{n-3}10$ and $00^{n-3}01$. Hence, we can let $H[RL^{n-3}] = 10^{n-3}11$. ☐

Note that in the above embedding, the tree node L^{n-2} cannot be embedded within $C[L^{n-3}]$.

Lemma 3.8: It is always possible to perform the $LR^{\lfloor m/2 \rfloor}$-embedding with node borrowing when n is odd and $x = n-2$. (Note that $n = m$ in this case.)

Proof: Note that $H[LR^{\lfloor m/2 \rfloor}] = 0^{\lfloor n/2 \rfloor}01^{\lfloor n/2 \rfloor -1}0$ and $C[LR^{\lfloor m/2 \rfloor}] = 0^{\lfloor n/2 \rfloor} *^{\lceil n/2 \rceil -1} 0$. Also the $n - \lfloor m/2 \rfloor - 2 = \lceil n/2 \rceil - 2$ faults in $C[LR^{\lfloor m/2 \rfloor}]$ are $0^{\lfloor n/2 \rfloor}00^{i-1}10^{\lceil n/2 \rceil -2-i}0$ for $i = 1, \cdots, \lceil n/2 \rceil - 2$. Hence by Lemma 3.7, we can embed the subtree rooted at $LR^{\lfloor m/2 \rfloor}$ except the leaf L^{n-2}. As L^{n-3} is mapped to $H[L^{n-3}] = 0^{\lfloor n/2 \rfloor}0^{\lfloor n/2 \rfloor -2}110$, we can borrow the node $10^{\lfloor n/2 \rfloor}-10^{\lfloor n/2 \rfloor -2}110$ from $C[L]$, i.e., we can map L^{n-2} to that node.

Note that in this case where $x = n - 2, p$ must be 0. Therefore $C[L]$ contains $p = 0$ faults at Hamming distance $\leq n - 3$ from $H[L]$. Hence the node $10^{\lfloor n/2 \rfloor -1}0^{\lfloor n/2 \rfloor -2}110$, which is at Hamming distance 2 from $H[L] = 10^{n-1}$, must be nonfaulty. Finally, the L-embedding has to avoid the node $10^{\lfloor n/2 \rfloor -1}0^{\lfloor n/2 \rfloor -2}110$. As $n \geq 6$, this embedding is still possible by recursion. ☐

E. Base Cases

When $n = 2$, there are $n - 2 = 0$ faults. When $n = 3$, we can assume without loss of generality that the only fault is 001. Thus we can set $H[L] = 100$ and $H[R] = 010$.

When $n = 4$, we can enumerate all the sets of faults with Hamming distance at most 2 from $H[\epsilon] = 0000$, i.e., faults with one or two 1's. It is easy to see that any set of faults is equivalent to one of the following distinct sets by renaming of dimensions: $\{0100, 0010\}, \{0100, 1100\}, \{0100, 1010\}, \{0110, 1100\}$ and $\{0110, 1001\}$. In all cases, we can embed the 3-tree by setting $H[L] = 1000, C[L] = 1***, H[R] = 0001$ and $C[R] = 0***$. Note that $C[L]$ has at most 1 fault and $C[R]$ has 1 assigned node but no fault with Hamming distance ≤ 1 from $H[R]$. Hence both the L- and R-embeddings can be embedded easily.

When $n = 5$, we try to split the 5-cube according to conditions A1)-A3).

- If $p = 2, C[L]$ contains 2 faults and $C[R]$ contains 1 fault and $H[\epsilon]$. Hence the L- and R-embeddings can be done recursively.
- If $p = 1$, we can still apply the method in Section III-D to find the embedding.
- If $p = 0$, there are only 2 possible sets of faults: $\{01000, 00100, 00010\}$ and $\{01000, 00100, 01100\}$. Other sets of faults can be shown to be equivalent to one of them by arranging the dimensions according to Section III-D2. We can still apply the algorithm in Section III-D3. Although when the set of faults is $\{01000, 00100, 01100\}$, the 3-cube $01 ***$ contains 2 faults, 01000 and 01100, the second one can be ignored (by Theorem 3.2). Hence the embedding can be done.

IV. VARIABLE ROOT TREE EMBEDDING

In this section, we consider the problem of variable root tree embedding. Because we can choose any nonfaulty node as root, more faults can be tolerated. For example, consider the case of n faults $0^{i-1}10^{n-i}$, for $i = 1, 2, \cdots, n$, in the n-cube. While it is impossible to embed an $(n - 1)$-tree rooted at 0^n, it is easy to embed the tree rooted at 1^n. In fact, we have the following result:

Theorem 4.1: For all $n \geq 5$ and $0 \leq f \leq \min(2n-7, 2n - 3 - \lceil \log n \rceil)$, there exists an embedding of an $(n - 1)$-tree into an n-cube having f faulty nodes/links. (In this case, the root can be mapped to any nonfaulty hypercube node.)

We still use cube splitting as our major technique for this problem. Suppose, for some dimension $i, *^{i-1}1*^{n-i}$ and $*^{i-1}0*^{n-i}$ contain p and q faults respectively. Without loss of generality, we assume $p \leq q$. Then we set $C[R] = *^{i-1}1*^{n-i}$
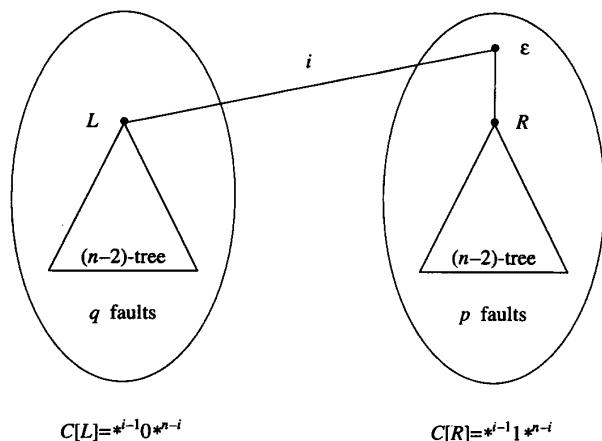
Fig. 9. Splitting $C[\epsilon] = *^n$ on dimension i.

and $C[L] = *^{i-1}0 *^{n-i}$. Moreover, we shall map the root, ϵ, into $C[R]$. (See Fig. 9 and note the difference between Figs. 1 and 9.)

Thus the R-embedding has to avoid $p+1$ faults effectively while the L-embedding has to avoid q faults. As $p \leq q$ and $f \leq 2n-7, p \leq \lfloor f/2 \rfloor \leq n-4$. Therefore we can map R to any nonfaulty node in $C[R]$ as we wish and then perform the R-embedding using the specified root method. For the L-embedding, there may be more than $n-3$ faults in the $(n-1)$-cube $C[L]$. Hence we recursively apply the variable root method to perform the L-embedding. If we can show that the number of possible L-embeddings, each with a different $H[L]$, is greater than p, then there will always exist a choice for $H[L]$ in $C[L]$ whose neighbor in $C[R]$ is nonfaulty. Hence we can map ϵ to this neighbor and map R to a nonfaulty neighbor of $H[\epsilon]$ in $C[R]$. (Remember that the R-embedding is always possible no matter where we place $H[R]$.)

Let $v(n, f)$ be the number of nodes, S, in an n-cube with f faulty nodes so that an $(n-1)$-tree can be embedded in the hypercube with the root mapped to S.

Lemma 4.1:

a) $v(n, f) = 2^n - f$ for $f = 0, 1, \cdots, n-2$ and $n \geq 2$,

b) $v(n, f) \geq 2^{2n-2-f} - f$ for $f = n-2, \cdots, 2n-7$ and $n \geq 5$.

Proof: (a) By Section III, the embedding of an $(n-1)$-tree into an n-cube with at most $n-2$ faults is always possible as long as $H[\epsilon]$ is nonfaulty. (b) To be given in the rest of this section. $\square$

A corollary of Lemma 4.1(b) is that $v(n, f) \geq 1$ as long as $f \leq 2n-3 - \lceil \log n \rceil$. Therefore, for the variable root embedding, we can tolerate up to $\min(2n-7, 2n-3-\lceil \log n \rceil)$ faults, and hence Theorem 4.1 follows. We shall prove Lemma 4.1(b) by induction on n with base case $n = 5$. When $n = 5, n-2 \leq f \leq 2n-7$ implies that $f = 3$. By Lemma 4.1(a), $v(n, f) = 2^5 - 3 = 2^{2n-2-f} - f$. Hence the base case is true. For $n \geq 6$, if $f = n-2$, then by Lemma 4.1(a) again, $v(n, f) = 2^n - f = 2^{2n-2-f} - f$. Hence for the induction step, we just need to consider $n-1 \leq f \leq 2n-7$ only.

For $n \geq 6$ and $f = n-1, \cdots, 2n-7$, we split $C[\epsilon]$ on a dimension such that p is as large as possible. There are two cases:

1) $p \geq 2$ (Section IV-A)
2) $p = 1$ (Section IV-B)

A. Case (1) where $p \geq 2$

Note that $q = f - p \leq f - 2 \leq 2(n-1) - 7$.

- If $(n-1) - 2 = n-3 \leq q$, the L-embedding has $v(n-1, q) \geq 2^{2(n-1)-2-q} - q$ possible nodes for $H[L]$ by the induction hypothesis. Of these nodes, at most p have faulty neighbors in $C[R]$. For the other nodes, the root ϵ can be mapped to their respective nonfaulty neighbors in $C[R]$ and the R-embedding can be done easily since $p \leq n-4$. Hence $v(n, f) \geq v(n-1, q) - p \geq (2^{2n-2-(q+2)} - q) - p \geq 2^{2n-2-f} - f$ as $p \geq 2$.

- If $q \leq (n-1) - 3 = n-4$, then by Lemma 4.1(a), the L-embedding has $v(n-1, q) = 2^{n-1} - q$ possible nodes for $H[L]$. Hence $v(n, f) \geq v(n-1, q) - p = (2^{n-1} - q) - p \geq 2^{2n-2-f} - f$ as $f \geq n-1$.

B. Case (2) where $p = 1$

In this case, we shall prove that $v(n, f) \geq 2^{n-1} - f \geq 2^{2n-2-f} - f$ for $f \geq n-1$. We first prove that if $f \geq n+2$, we can always find a dimension to split $C[\epsilon]$ such that $q \geq p \geq 2$. Hence $p = 1$ implies $n-1 \leq f \leq n+1$.

Lemma 4.2: If $n \geq 4$ and $f \geq n+2$, the $C[\epsilon]$ can be split into two $(n-1)$-cubes, $C[L]$ and $C[R]$, such that $q \geq p \geq 2$, where q and p are the number of faults in $C[L]$ and $C[R]$, respectively.

Proof: Assume to the contrary that $p \leq 1$ for any dimension on which $C[\epsilon]$ is split, in particular $C[L] = *^{i-1}0*^{n-1}$ and $C[R] = *^{i-1}1*^{n-1}$. (Note that if $q < p$, the ith bit of all faults can be complemented in order to have $q \geq p$.) In other words, at most one fault has a 1 in its ith dimension. Thus, the total number of 1's in the labels of all the faults would be $\leq n$. On the other hand, each fault (except 0^n) has at least one 1 in its label. If there are $\geq n + 2$ faults, this would imply that there are $\geq (n+1)$ 1's in the labels of all the faults. Thus, we can conclude that $p \geq 2$. $\square$

Since no less than $f - 1$ faults (except 0^n) have at least one 1 in their labels, it is easy to show that the set of faults possesses the following properties:

P1) If $f = n + 1$, the faults are 0^n, and $0^{i-1}10^{n-i}$ for $1 \leq i \leq n$.

P2) If $f = n$, there is at most one fault with two 1's. Should there be a fault with two 1's, there must exist fault 0^n with the remaining $n - 2$ faults having a single 1 each.

P3) If $f = n - 1$, there are at most two faults with two 1's or (exclusive) one fault with three 1's.

P4) There are at least $n - 4$ faults with a single 1.

Let S be an arbitrary nonfaulty hypercube node with 0's in m dimensions. It is sufficient to show that the embedding is possible with $H[\epsilon] = S$ provided m is even and $\leq n - 1$.

If $m = 0$, i.e., $S = 1^n$, at least $f - 2$ faults are of Hamming distance $\geq n - 1$ from S and thus can be ignored (Theorem 3.2). As a consequence, the embedding of an $(n - 1)$-tree into this n-cube is possible by the method in Section III.

For $2 \leq m \leq n - 1$ where m is even and $n \geq 6$, we shall arrange dimensions so that

C1) S is of the form $0^m 1^{n-m}$ with $n-m \geq 1$, and subcube $0^{m/2} *^{n-(m/2)-1} 1$ does not contain a fault with more than two 1's;

C2) 10^{n-1} and $0^{m-1} 10^{n-m}$ are faulty; and

C3) subcube $0^{i-1} 1 *^{n-i}, 1 \leq i \leq m - 2$, contains at least as many faults as subcube $0^i 1 *^{n-i-1}$ does. More precisely, subcube $0^{i-1} 1 *^{n-i}$ contains exactly one fault for all i such that $1 \leq i \leq j$ for some $j \leq m$.

Note that conditions (C1)–(C3) can always be satisfied through arrangement of dimensions. Consider condition (C1). There is at most one fault with more than two 1's (by property (P3)). Moreover it should be different from S in at least one dimension because S is nonfaulty. Suppose $S = 0^m 1^{n-m}$. Then this fault must either have a 1 in a dimension from 1 to m or a 0 in dimension from $m + 1$ to n. By swapping this dimension to a dimension from 1 to $m/2$ or to dimension n, this fault will not lie in $0^{m/2} *^{n-(m/2)-1} 1$. Hence condition (C1) is satisfied. Since there are at least two faults with a single 1 (from property (P4) and $n \geq 6$), condition (C2) can always be satisfied with arrangement of dimensions 1 and m. Furthermore, it is straightforward to arrange dimensions to satisfy condition (C3) without affecting dimension 1 and dimensions m to n.

Our next step is to apply the embedding strategy for even m and $2 \leq m \leq n - 1$ as described in Section III-D.

Step (1) $H[\epsilon] = 0^m 1^{n-m}$

Step (2) For $i = 1, \cdots, m/2$,

$$H[LR^{i-1}] = 0^{i-1} 10^{m-2i} 01^{i-1+n-m}$$
$$C[LR^{i-1}] = 0^{i-1} 1 *^{n-i}$$
$$H[RR^{i-1}] = 0^{i-1} 00^{m-2i} 11^{i-1+n-m}$$
$$C[RR^{i-1}] = 0^{i-1} 0 *^{n-i}.$$

Step (3)

$$H[LR^{m/2}] = 0^{m/2} 1^{n-(m/2)-1} 0$$
$$C[LR^{m/2}] = 0^{m/2} *^{n-(m/2)-1} 0$$
$$H[RR^{m/2}] = 0^{m/2} 1^{n-(m/2)-2} 01$$
$$C[RR^{m/2}] = 0^{m/2} *^{n-(m/2)-1} 1.$$

By condition (C2) and $p = 1, 10^{n-1}$ and $0^{m-1} 10^{n-m}$ are the only faults in $1 *^{n-1}$ and $*^{m-1} 1 *^{n-m}$ respectively. Thus $H[L], H[R], H[LR]$ and $H[RR]$ cannot be faulty. $H[LR^{i-1}]$ and $H[RR^{i-1}]$ would not be faulty for $i = 3, \cdots, (m/2) + 1$ as these nodes have more than three 1's. As each $C[LR^{i-1}] = 0^{i-1} 1 *^{n-i}$ has exactly 1 fault for $i = 1, \cdots, m/2$, these LR^{i-1}-embeddings can be done with recursion.

By step (3), $C[RR^{m/2}] = 0^{m/2} *^{n-(m/2)-1} 1$ contains all the assigned hypercube nodes $H[R^j], j = 0, \cdots, m/2$, and possibly at most another faulty node, whereas $C[LR^{m/2}]$ contains only faulty nodes (at least $(m/2) - 1$ of them).

The following lemmas (similar to Lemmas 3.6, 3.8 and Lemma B.4 in Appendix B) will show that the $RR^{m/2}$- and $LR^{m/2}$-embeddings are always possible as long as conditions (C1)–(C3) are satisfied.

Lemma 4.3: It is always possible to perform the $RR^{m/2}$-embedding.

Proof: If $C[RR^{m/2}]$ contains a fault (i.e., that fault with a 1 in dimension n), because of condition (C1), that fault must have at most one 1 in dimensions $(m/2) + 1$ to $n - 1$ and have Hamming distance $\geq n - (m/2) - 2$ away from $H[RR^{m/2}]$. Thus it can be ignored in the embedding (Theorem 3.2).

Rewriting the assigned node $H[R^j]$ as $0^{m/2} 0^{(m/2)-j} 1^j 1^{n-m}$, the Hamming distance between $H[R^j]$ and $H[RR^{m/2}] = 0^{m/2} 1^{m/2} 1^{n-m-2} 01$ is $(m/2) - j + 1$ except when $n - m = 1$ and $j = 0$. Note that when $n - m = 1$, the Hamming distance between $H[\epsilon]$ and $H[RR^{m/2}]$ is $n - (m/2) - 2$, thus $H[\epsilon]$ is too far to affect the embedding. Since the other nodes have unique Hamming distances from $H[RR^{m/2}]$, there are $n - (m/2) - 3$ of them with Hamming distance $\leq n - (m/2) - 3$ away from $H[RR^{m/2}]$. Thus $RR^{m/2}$-embedding within $C[RR^{m/2}]$ is always possible by recursion. □

Lemma 4.4: It is always possible to perform the $LR^{m/2}$-embedding.

Proof: Note that all faults in $C[LR^{m/2}]$ having exactly one 1 in their labels are of Hamming distance $n - (m/2) - 2$ from the $H[LR^{m/2}] = 0^{m/2} 1^{n-(m/2)-1} 0$. Since at most two faults cannot be ignored (by property (P3) and Theorem 3.2), the $LR^{m/2}$-embedding is possible. □

From the above lemmas, it follows that the embedding is possible with $H[\epsilon] = S$ for all nonfaulty hypercube node S having m 0's where m is even and $\leq n - 1$. Hence

$$v(n, f) \geq \binom{n}{0} + \binom{n}{2} + \cdots + \binom{n}{2\lfloor (n-1)/2 \rfloor} - f$$

where the $\binom{n}{i}$'s are the binomial coefficients. When n is odd,

$$v(n, f) \geq \binom{n-1}{0} + \binom{n-1}{1} + \cdots + \binom{n-1}{n-1} - f$$
$$= 2^{n-1} - f.$$

When n is even, by property (P4), there are at least $n - 4$ faults having an odd number of 0's. Hence

$$v(n, f) \geq \binom{n-1}{0} + \binom{n-1}{1} + \cdots + \binom{n-1}{n-2}$$
$$- (f - n + 4)$$
$$\geq 2^{n-1} - f$$

for $n \geq 5$. Thus no matter n is even or odd, $v(n, f) \geq 2^{n-1} - f \geq 2^{2n-2-f} - f$ (as $f \geq n - 1$).

Note that embedding is also possible with the root mapped to nonfaulty S such that $m \leq n - 1$ and m is odd. However, the embedding is quite tedious and handling these cases does not improve the lower bound on $v(n, f)$. Hence the embedding is not shown here.

C. A Bound on the Number of Faults

In this section, we establish an upper bound on the least number of faulty nodes present in an n-cube which can make the embedding of an $(n-1)$-tree impossible. We allow the embedding to have unbounded dilation, $O(1)$ load and the freedom to choose any nonfaulty node as the root. However, the embedding should never use any faulty hypercube node or any of its n edges. We find that there exists a set of $O(1/\sqrt{n}) \times 2^n$ faulty nodes which can make such an embedding impossible.

We first consider the case where load = 1. For convenience, we consider the n-cube as a graph of $n+1$ layers where layer i includes all nodes with exactly i 1's in their labels. A node is *above* (or *below*) layer i if its label has less (or more) than i 1's.

Assume n is even and consider the case that layer $n/2$ is faulty. Then there are $\binom{n}{n/2}$ faulty nodes. By Stirling's formula, this is $\leq (\sqrt{2/\pi n}) 2^n e^{1/12n}$ which is an $O(1/\sqrt{n})$ fraction of 2^n. Moreover, it is impossible to embed an $(n-1)$-tree into such an n-cube with unit load. It is because the nodes above layer $n/2$ are disconnected from those below and both groups have less than $2^{n-1} - 1$ nodes when $n \geq 4$.

The situation is similar when n is odd. We can make layers $\lceil n/2 \rceil$ and $\lfloor n/2 \rfloor$ faulty. Then there are

$$\binom{n}{\lceil n/2 \rceil} + \binom{n}{\lfloor n/2 \rfloor} = \binom{n+1}{(n+1)/2}$$
$$\leq \sqrt{\frac{2}{\pi(n+1)}} 2^{n+1} e^{1/12(n+1)}$$

faults which is again an $O(1/\sqrt{n})$ fraction of 2^n. Also, the number of nodes above layer $\lfloor n/2 \rfloor$ is less than $2^{n-1} - 1$ for $n \geq 3$ (and so is the number of nodes below layer $\lceil n/2 \rceil$). Hence the embedding of an $(n-1)$-tree is impossible.

Thus, for $n \geq 3$, it is not always possible to embed with unit load, an $(n-1)$-tree into an n-cube having an $O(1/\sqrt{n})$ fraction of faults even though we are free to choose any nonfaulty node as the root of the tree.

If the load $= c$, we can make layers $l_1, \cdots, l_{2c}$ faulty where $0 < l_1 < l_2 \cdots < l_{2c} \leq n$. Moreover, for $i = 1, \cdots, 2c$, the number of nodes between layers $l_{i-1} + 1$ and l_i (inclusively) is $\geq \lfloor (2^{n-1} - 1)/c \rfloor$ and that between layers $l_{i-1} + 1$ and $l_i - 1$ (inclusively) is $< \lfloor (2^{n-1} - 1)/c \rfloor$. (For convenience, we let $l_0 = -1$.) Note that nodes between layers $l_{i-1} + 1$ and $l_i - 1$ are nonfaulty. Thus the nonfaulty nodes of the n-cube are separated into $2c$ or $2c+1$ groups (depending on $l_{2c} = n$ or not), each having $< \lfloor (2^{n-1} - 1)/c \rfloor$ nodes. Note that the group of nodes below layer l_{2c} has $\leq 2^n - (2c \times \lfloor (2^{n-1} - 1)/c \rfloor) \leq 2c < \lfloor (2^{n-1} - 1)/c \rfloor$ nodes for large enough n. Hence an $(n-1)$-tree cannot be embedded (wholly or partly) within any group. Finally, the total number of nodes in all the faulty layers is at most $2c \times O(1/\sqrt{n}) \times 2^n$ or $O(1/\sqrt{n}) \times 2^n$. Hence our claim follows.

V. RECURSIVE EMBEDDINGS

In this section, we study an interesting relationship between our embedding strategies and the *recursive embedding* mentioned in [21].

Definition 5.1: A *recursive embedding* is one that maps the left and right subtrees of every internal node of the binary tree into disjoint subcubes.

It was shown in [21] that any recursive (variable root) embedding of an $(n-1)$-tree into an n-cube with unit dilation and load can tolerate no more than $2n - 3$ faults in the worst case. Before we prove this result, let us have the following definitions.

Definition 5.2: Let $a_1 \cdots a_n$ be the label of a hypercube node A. Then parity of A is defined as $parity(A) = (\Sigma_{i=1}^n a_i)$ mod 2.

In other words, $parity(A) = 0$ if A has even number of 1's in its label and $parity(A) = 1$ otherwise.

Lemma 5.1: For any tree node T, $parity(H[T]) = parity(H[\epsilon])$ if and only if T is at an even level.

Proof: It is easily proved by induction on n, the level number of T. Suppose the lemma is true for $n - 1$. Then $parity(H[S]) = parity(H[\epsilon])$ if and only if S is at an even level, where S is the parent of T. Since the number of 1's in $H[T]$ must differ from that in $H[S]$ by one (dilation one embedding), $parity(H[T]) \neq parity(H[S])$. Moreover, T is at even level if and only if S is not. Hence $parity(H[T]) = parity(H[\epsilon])$ if and only if T is at an even level. □

Lemma 5.2: For any leaf nodes S and T, $parity(H[S]) = parity(H[T])$.

Proof: It follows directly from Lemma 5.1 as S and T are on the same level. □

The following lemma from [21] gives an important property about recursive embedding.

Lemma 5.3: [21] For all $n \geq 2$ and for every hypercube node A in an n-cube, there exists a leaf, T_A, of the $(n-1)$-tree, such that the Hamming distance between A and $H[T_A]$ is at most 2 where H is a recursive embedding.

Proof: By definition, the recursive embedding cuts the n-cube into 2^{n-2} disjoint 2-cubes, each containing one leaf. Hence every hypercube node, lying in some 2-cube, is at Hamming distance at most two from the leaf in that 2-cube. □

The following theorem from [21] gives an upper bound on the least number of faults such that no recursive embedding is possible.

Theorem 5.1: [21] For $n \geq 3$, there exists a set of $2n - 2$ faults such that no recursive embedding can avoid all the faults.

Proof: Consider the set of $2n - 2$ faults: $\{0^{i-1}10^{n-i-1}0, 0^{i-1}10^{n-i-1}1$ for $i = 1, \cdots, n - 1\}$. They isolate the subcube $0^{n-1}*$ from the rest of nonfaulty nodes. Hence $0^{n-1}*$ cannot contain part of the $(n-1)$-tree. As $n \geq 3$, it cannot contain the whole $(n-1)$-tree too. By Lemma 5.3, there must exist two leaves, U and V, such that $H[U]$ is at Hamming distance exactly 2 from $0^{n-1}0$ and so is $H[V]$ from $0^{n-1}1$. Hence $parity(H[U]) = parity(0^{n-1}0) = 0$ and $parity(H[V]) = parity(0^{n-1}1) = 1 \neq parity(H[V])$, thus contradicting Lemma 5.2. □

It is easy to see that our specified root embedding is basically a recursive embedding. The only violation is the occasional use of node borrowing. We shall show in Theorem 5.2 that node borrowing is necessary for the specified root embedding to tolerate $n - 2$ faults. Moreover, there is only one case in which the technique is required. On the other

hand, our variable root embedding does not require node borrowing at all (to be proved in Theorem 5.3) and therefore is a recursive embedding. Thus we have achieved their $2n - 3$ bound asymptotically.

Theorem 5.2: For the specified root embedding problem, recursive embedding fails if and only if $n \geq 5$, n is odd and $x = n - 2$ where x is the number of faulty neighbors of $H[\epsilon]$.

Proof: *(if)* The parity of $H[\epsilon]$ is 0. For odd n, the parity of the leaves of an $(n - 1)$-tree is 1 (by Lemma 5.1). By Lemma 5.3, $H[\epsilon]$ must be at Hamming distance 1 from a leaf. However, among the n neighbors of $H[\epsilon]$, $n - 2$ of them are faulty and the other 2 are $H[L]$ and $H[R]$ which are internal nodes for $n \geq 4$. Thus, recursive embedding is not possible.

(Only if) Referring to our embedding algorithm, excluding the recursive calls, node borrowing is only used when n is odd and $x = n - 2$. Thus, it is sufficient to prove that node borrowing is also not needed in any of the recursive calls. We shall prove this by induction on n. For the base cases where $n \leq 4$, Section III-E shows that node borrowing is not needed. For the induction step, there are 4 cases to be considered.

A. Case (1) $0 \leq p \leq 1$

We need to argue that the LR^{i-1}-embedding for $i = 1, \cdots, \lfloor m/2 \rfloor + 1$ and the $RR^{\lfloor m/2 \rfloor}$-embedding will not require node borrowing. Basically, we try to prove that at each recursive step, not many faults can be neighbors of the specified root.

First, $C[L]$ contains only $p \leq 1$ faults. Second, at most 1 fault in $C[LR^{i-1}]$ for $i = 2, \cdots, \lfloor m/2 \rfloor$ is adjacent to $H[LR^{i-1}]$ since the $H[LR^{i-1}]$'s lie in

$$\begin{cases} 0 *^{m-1} 0^{n-m} & \text{if } n > m \\ 0 *^{n-2} 1 & \text{if } n = m \end{cases}$$

while the faults in $0*^{n-1}$ lie in

$$\begin{cases} 0 *^{m-1} 1^{n-m} & \text{if } n > m \\ 0 *^{n-2} 0 & \text{if } n = m \end{cases}.$$

For the $LR^{\lfloor m/2 \rfloor}$-embedding, we shall show that there are less than $n - \lfloor m/2 \rfloor - 3$ faults adjacent to

$$H[LR^{\lfloor m/2 \rfloor}] = \begin{cases} 0^{\lceil m/2 \rceil} 1^{\lfloor m/2 \rfloor} 0^{n-m-1} 1 & \text{if } n > m \\ 0^{\lceil m/2 \rceil} 1^{\lfloor m/2 \rfloor - 1} 0 & \text{if } n = m \end{cases}$$

provided $n - \lfloor m/2 \rfloor - 1 \geq 5$. Consider a fault A, which is adjacent to $H[LR^{\lfloor m/2 \rfloor}]$. There are 2 cases.

- If $n - m \geq 2$, then A must disagree with $H[LR^{\lfloor m/2 \rfloor}]$ in at least $n - m - 1 \geq 1$ dimensions. (Note that A lies in $0^{\lfloor m/2 \rfloor} *^{\lceil m/2 \rceil} 1^{n-m}$ by condition (B2)). Hence A must agree with $H[LR^{\lfloor m/2 \rfloor}]$ in any dimension from $\lceil m/2 \rceil + 1$ to m. In other words, A has 1's in all these dimensions. Then, since at most one fault in

$$C[LR^{\lfloor m/2 \rfloor}] = \begin{cases} 0^{\lfloor m/2 \rfloor} *^{n-\lfloor m/2 \rfloor - 1} 1 & \text{if } n > m \\ 0^{\lfloor m/2 \rfloor} *^{n-\lfloor m/2 \rfloor - 1} 0 & \text{if } n = m \end{cases}$$

can have a 1 in any of these dimensions (refer to the proof in Lemma B.4), no other faults can be adjacent to $H[LR^{\lfloor m/2 \rfloor}]$.

- If $n - m \leq 1$, A can disagree with $H[LR^{\lfloor m/2 \rfloor}]$ in at most one dimension from $\lceil m/2 \rceil + 1$ to m. By similar argument as above, no other faults can disagree with $H[LR^{\lfloor m/2 \rfloor}]$ in at most one of these dimensions unless $\lfloor m/2 \rfloor - 1 \leq 2$ but this implies $n - \lfloor m/2 \rfloor - 1 \leq 4$.

Consider the $RR^{\lfloor m/2 \rfloor}$-embedding. Referring to Lemma 3.6, $C[RR^{\lfloor m/2 \rfloor}]$ contains the assigned nodes $H[R^j]$ for $j = 1, \cdots, \lfloor m/2 \rfloor$ and if $n > m$, $H[\epsilon]$ also. However, it does not contain any fault by condition (B2). Since the Hamming distance between $H[R^j]$ (except $H[\epsilon]$ and $H[R]$ when $n = m$ and m is even) and $H[RR^{\lfloor m/2 \rfloor}]$ is $\lfloor m/2 \rfloor - j + 1$, only $H[R^{\lfloor m/2 \rfloor}]$ is adjacent to $H[RR^{\lfloor m/2 \rfloor}]$. Furthermore, $H[\epsilon]$ and $H[R]$, when $n = m$ and m is even, are either outside $C[RR^{\lfloor m/2 \rfloor}]$ or not adjacent to $H[RR^{\lfloor m/2 \rfloor}]$ (refer to Lemma 3.6).

B. Case (2) $p = 2$

$C[L]$ and $C[R]$ contains 2 and $n - 3$ faults/assigned nodes respectively. Consider the $n - 4$ faults (excluding the assigned node) in $C[R]$. We shall show that at most 2 of them can be adjacent to $H[R]$ if we have chosen $H[R]$ carefully.

Suppose A is a fault adjacent to $H[R] (= 0^{j-1} 1 0^{n-j}$, say). Then it must have 1's in exactly two dimensions and one of them must be dimension j. However, if we have chosen j (i.e., $H[R]$) carefully such that $0^{j-1} 1 0^{n-j}$ is nonfaulty and $*^{j-1} 1 *^{n-j}$ contains $< n - 2$ faults, then by conditions A1)–A3), $*^{j-1} 1 *^{n-j}$ contains at most 2 faults. Hence there are at most 2 faults have a 1 in dimension j. Consequently, at most 2 faults are adjacent to $H[R]$. Note that such a dimension j can always be found because there are at most $n - 3$ dimensions k such that $0^{k-1} 1 0^{n-k}$ is faulty or $*^{k-1} 1 *^{n-k}$ contains all the faults.

C. Case (3) $3 \leq p \leq n - 4$

In this case, both $C[L]$ and $C[R]$ do not have $n - 3$ faults/assigned nodes.

D. Case (4) $p = n - 3$

In this case, $C[L]$ and $C[R]$ contains $n - 3$ and 2 faults/assigned nodes respectively. When the $n - 3$ faults in $C[L]$ are all adjacent to $H[L] (= 10^{n-1}$, say), the $n - 3$ faults in $C[L]$ must be $10^{i-2} 10^{n-i-2} 00$ for $i = 2, \cdots, n - 2$ and the only fault in $C[R]$ must be $0a_2 \cdots a_n$ where the a's can be 0 or 1. (See Fig. 10)

$$1100 \cdots 0000$$
$$1010 \cdots 0000$$
$$\vdots$$
$$1000 \cdots 0100$$
$$0aaa \cdots aaaa.$$

This time we need to split $C[\epsilon]$ on another dimension. Consider the new value of p. Obviously, it cannot be in cases (3) or (4) by referring to Fig. 10. If it is in case (2), then node borrowing is not required as shown above. If it is in case (1), we proceed to find the embedding using the method in Section III–C. In particular, if a_{n-1} or a_n or both are 0, we

1100...0000
1010...0000

.
.
.

1000...0100
0aaa...aaaa

Fig. 10. The $n - 2$ faults. The value of the a's can be 0 or 1, but not all 0.

have the case $n = m$ by splitting on any dimension from 2 to $n - 2$. Otherwise (i.e. both a_{n-1} and a_n are 1), we have the case $n - m = 1$ by splitting on dimensions $n - 1$ or n. □

Theorem 5.3: Our variable root embedding strategy is a recursive embedding.

Proof: Note that for $p = 1$, no node borrowing is needed. For $2 \leq p \leq n - 5$, $C[R]$ contains $p + 1 \leq n - 4$ faults effectively and hence, the specified root embedding in $C[R]$ does not require node borrowing. The variable root embedding in $C[L]$ also does not require node borrowing (by induction hypothesis).

For $p = n - 4$, $C[R]$ contains $n - 3$ faults. However, we can choose R carefully so that $H[R]$ does not have $n - 3$ faults/assigned nodes as neighbors in $C[R]$. Then the specified root embedding in $C[R]$ does not require node borrowing also.

VI. CONCLUSIONS

In this paper, we present three new results on finding a unit dilation and load embedding of an $(n - 1)$-tree into an n-cube that contain some faulty processors and/or links.

For the specified root embedding problem, we show that up to $n - 2$ faults can be tolerated. This is optimal in the sense that $n - 2$ faults are the maximum number of faults that can be tolerated when the root is specified. Moreover, the $(n - 1)$-tree is the largest full binary tree that can be embedded into an n-cube even when there are no faults.

For the variable root embedding problem, we show that up to $2n - 3 - \lceil \log n \rceil$ faults can be tolerated. It is not surprising to see that more faults can be tolerated in this problem because there are fewer restrictions on the embedding. Also, our method is classified as a recursive embedding in [21]. It was shown in [21] that no recursive embedding can tolerate more that $2n - 3$ faults in the worst case. Hence our result achieves their bound asymptotically.

Finally, we prove that when an $O(1/\sqrt{n})$ fraction of nodes in the n-cube is faulty, it is not always possible to embed an $(n - 1)$-tree into an n-cube with unbounded dilation and $O(1)$ load even if we are free to choose any nonfaulty node as the root.

As the variable root embedding has not achieved optimality in the broad sense, more work can be done to tighten the upper and lower bounds on the number of faults that can be tolerated. In particular, non-recursive embedding strategies should be studied in order to tolerate more than $2n - 3$ faults. Recently, [21] has derived an embedding method that can tolerate $\Omega(n^2/\log n)$ faults. Their method is, of course, non-recursive. Furthermore, embeddings with greater dilation and/or load are interesting and useful variations for investigation. Special

patterns of faults that better model the actual fault scenario are also worth studying.

APPENDIX A

This section contains the technical details for the arrangement of dimensions described in Section III–D2.

Lemma A.1: There exists two dimensions i and j such that $0^{i-1}10^{n-i}$ is nonfaulty, $*^{i-1}1*^{n-i}$ contains exactly p faults and either $*^{j-1}0*^{n-j}$ or $*^{j-1}1*^{n-j}$ contains all the remaining $n - 2 - p$ faults, i.e., all the faults in $*^{i-1}0*^{n-i}$.

Proof: When $p = 0$, there exists two dimensions i and j such that $0^{i-1}10^{n-i}$ and $0^{j-1}10^{n-j}$ are nonfaulty (Lemma 3.2(a)). Then by Lemma 3.1, both $*^{i-1}1*^{n-i}$ and $*^{j-1}1*^{n-j}$ contain 0 faults.

When $p = 1$, by conditions A1)–A3), there must exist a dimension i' such that $0^{i'-1}10^{n-i'}$ is nonfaulty and $*^{i'-1}1*^{n-i'}$ contains $p = 1$ faults. We have the following cases:

- If there also exists a dimension $j' \neq i'$ such that $*^{j'-1}1*^{n-j'}$ contains zero or all the $n - 2 - p$ faults in $*^{i'-1}0*^{n-i'}$, then the lemma is proved.

- If such a dimension j' does not exist, we will try to prove that there exists a fault with at least two 1's, say in dimensions i'' and j'' where $0^{i''-1}10^{n-i''}$ and $0^{j''-1}10^{n-j''}$ are nonfaulty. Then since $p = 1$, $*^{i''-1}1*^{n-i''}$ and $*^{j''-1}1*^{n-j''}$ will contain $p = 1$ fault, namely this fault, and $*^{j''-1}0*^{n-j''}$ will contain all the remaining faults.

To find such a fault and dimensions i'' and j'', consider those $k \neq i'$ where $0^{k-1}10^{n-k}$ is nonfaulty. With the assumption that j' does not exist, the subcube $*^{k-1}1*^{n-k}$ must contain exactly $p = 1$ fault in $*^{i'-1}0*^{n-i'}$. There are $n - 1 - x$ such subcubes as $*^{k-1}1*^{n-k}$. (Remember that x is the number of faulty neighbors of $H[\epsilon]$.) Since there are only $n - 2 - p - x$ faults which can belong to these subcubes, there must be at least two different (out of these $n - 1 - x$) subcubes, $*^{i''-1}1*^{n-i''}$ and $*^{j''-1}1*^{n-j''}$, containing the same single fault. Hence $*^{i''-1}1*^{n-i''}$ and $*^{j''-1}1*^{n-j''}$ contain $p = 1$ fault and contain none of the rest. Then we can let i be i'' and j be j'' (or vice-versa). □

It follows from Lemma A.1 that conditions (B1) and (B2) can be satisfied simultaneously as we can swap dimension i to dimension 1 and dimension j to dimension n respectively.

Lemma A.2: Condition (B3) can be achieved by arranging dimensions 2 through m while preserving conditions (B1) and (B2).

Proof: Note that all dimensions except dimensions 2 to m are unaffected. When $m = n$, $*^{n-1}1$ contains 0 faults in $0*^{n-1}$ and hence dimension n will not be affected too. Therefore conditions (B1) and (B2) are preserved.

We exchange the dimensions in the following manner. Initially, $0*^{n-1}$ contains $n - 2 - p$ faults. Let $j, 2 \leq j \leq m$, be the dimension such that the subcube $0*^{j-2}1*^{n-j}$ contains the largest number of faults. Then, dimensions j and 2 are exchanged. Thus $01*^{n-2}$ has the largest number of faults among all subcubes $0*^{k-2}1*^{n-k}$ where $2 \leq k \leq m$. The same procedure is then repeated on the subcube 0^2*^{n-2}, i.e., find dimension $j, 3 \leq j \leq m$, such that the subcube $0^2*^{j-3}1*^{n-j}$

has the largest number of faults and then exchange this dimension j with dimension 3 so that $0^2 1*^{n-3}$ has the largest number of faults among all subcubes $0^2 *^{k-3} 1*^{n-k}$ where $3 \leq k \leq m$. Note that dimension 2 will not be affected. This procedure is repeated in a similar fashion on $0^3 *^{n-3}, 0^4 *^{n-4}$, and so on to $0^{m-1} * *^{n-m}$. With similar arguments, the dimensions are exchanged to yield the property stated above.   □

Lemma A.3: Hypercube node

$$\begin{cases} 0^{\lceil m/2 \rceil} 1^{\lfloor m/2 \rfloor} 0^{n-m-1} 1 & \text{if } n > m \\ 0^{\lceil m/2 \rceil} 1^{\lfloor m/2 \rfloor -1} 0 & \text{if } n = m \end{cases}$$

can be made nonfaulty by arranging dimensions $\lceil m/2 \rceil$ and $\lceil m/2 \rceil + 1$ while preserving conditions (B1), (B2) and (B3).

Proof: For $n - m \geq 2$, the above mentioned node becomes $0^{\lceil m/2 \rceil} 1^{\lfloor m/2 \rfloor} 0^{n-m-1} 1$ which is never faulty by condition (B2). So nothing has to be done.

For $n - m \leq 1$, if the above mentioned node is faulty, we show that dimensions $\lceil m/2 \rceil$ and $\lceil m/2 \rceil + 1$ can be swapped so that the node becomes nonfaulty and at the same time, conditions (B1)–(B3) are preserved.

We claim that both $0^{\lceil m/2 \rceil -1} 1*^{\lfloor m/2 \rfloor +n-m}$ and $0^{\lceil m/2 \rceil} 1*^{\lfloor m/2 \rfloor +n-m-1}$ contain exactly 1 fault. Since the above mentioned node is faulty, $0^{\lceil m/2 \rceil} 1*^{\lfloor m/2 \rfloor +n-m-1}$ contains $r \geq 1$ faults. On the other hand, $0^{\lceil m/2 \rceil -1} 1*^{\lfloor m/2 \rfloor +n-m}$ contains ≤ 1 faults. Otherwise, the total number of faults in $C[\epsilon]$ is at least $2(\lceil m/2 \rceil - 1) + p + r$ (by condition (B3)) and hence at least $n - 3 + p + r$ (as $n - m \leq 1$). If $p = 1, p + r \geq 2$. If $p = 0$, then by Lemma 3.1, $0^{\lceil m/2 \rceil} 1*^{\lfloor m/2 \rfloor +n-m-1}$ must also contain the fault $0^{\lceil m/2 \rceil} 1 0^{\lfloor m/2 \rfloor +n-m-1}$. Hence $r \geq 2$. Consequently, no matter whether p is 0 or 1, the total number of faults in $C[\epsilon]$ is greater than $n - 2$ which is impossible. Thus by condition (B3), we can conclude that both $0^{\lceil m/2 \rceil -1} 1*^{\lfloor m/2 \rfloor +n-m}$ and $0^{\lceil m/2 \rceil} 1*^{\lfloor m/2 \rfloor +n-m-1}$ contain exactly 1 fault.

It follows from the claim that conditions (B1)–(B3) are preserved even if we swap dimensions $\lceil m/2 \rceil$ and $\lceil m/2 \rceil + 1$. Moreover, the fault in the former subcube, $0^{\lceil m/2 \rceil -1} 1*^{\lfloor m/2 \rfloor +n-m}$, must be

$$\begin{cases} 0^{\lceil m/2 \rceil -1} 1 0^{\lfloor m/2 \rfloor +n-m-1} 1 & \text{if } n - m = 1 \\ 0^{\lceil m/2 \rceil -1} 1 0^{\lfloor m/2 \rfloor +n-m-1} 0 & \text{if } n - m = 0 \end{cases}$$

before the dimensions are swapped because the node

$$\begin{cases} 0^{\lceil m/2 \rceil} 1^{\lfloor m/2 \rfloor} 0^{n-m-1} 1 & \text{if } n > m \\ 0^{\lceil m/2 \rceil} 1^{\lfloor m/2 \rfloor -1} 0 & \text{if } n = m \end{cases}$$

is faulty and $p \leq 1$. By exchanging dimensions $\lceil m/2 \rceil$ and $\lceil m/2 \rceil + 1$, it is changed to

$$\begin{cases} 0^{\lceil m/2 \rceil} 1 0^{\lfloor m/2 \rfloor +n-m-2} 1 & \text{if } n - m = 1 \\ 0^{\lceil m/2 \rceil} 1 0^{\lfloor m/2 \rfloor +n-m-2} 0 & \text{if } n - m = 0. \end{cases}$$

Also, the fault in latter subcube, $0^{\lceil m/2 \rceil} 1*^{\lfloor m/2 \rfloor +n-m-1}$, is changed to

$$\begin{cases} 0^{\lceil m/2 \rceil -1} 1 0 1 0^{\lfloor m/2 \rfloor +n-m-2} 1 & \text{if } n - m = 1 \\ 0^{\lceil m/2 \rceil -1} 1 0 1 0^{\lfloor m/2 \rfloor +n-m-2} 0 & \text{if } n - m = 0. \end{cases}$$

□

APPENDIX B

This section contains the lemmas which show the feasibility of the LR^{i-1}-embeddings for $i = 1, \cdots, \lfloor m/2 \rfloor + 1$ and the $RR^{\lfloor m/2 \rfloor}$-embedding.

Before we study these embeddings, let us consider some properties of the faults in $0*^{n-1}$. We shall show that there exists a subset E of the faults in $0*^{n-1}$ such that for all k where $2 \leq k \leq m$, $0*^{k-2} 1*^{m-k} 1^{n-m}$ contains at most one fault in E, i.e., at most one fault in E has a 1 in dimension k. Hence each $C[LR^i] = 0^{i-1} 1*^{n-i}$ contains at most one fault in E, i.e., at least $e - 1$ faults are not in $C[LR^i]$. Lemma B.1 below shows that e, the size of E, is not too small. This ensures that the remaining faults in $0*^{n-1}$ can be separated (by splitting the $C[R^i]$'s on suitable dimensions) so that the $C[LR^i]$'s will not have too many faults. In the case where $n > m$, Lemma B.2 shows that E can even include all the faults in $0*^{n-1}$. For instance, in Example 3.2, we can have $E = \{\beta_1, \beta_2, \beta_3, \beta_4\}$ while in Example 3.3, we can have $E = \{\beta_1, \beta_2, \beta_3\}$ or $\{\beta_3, \beta_4\}$.

Lemma B.1:

1) when $p = 0$, there exists a set of faults, E, with cardinality $e \geq 3$;
2) when $p = 1$, there exists a set of faults, E, with cardinality $e \geq 2$.

Proof: (a) When $p = 0$, the number of faulty neighbors of $H[\epsilon], x$, is $\geq \lceil \log(n-1) \rceil \geq 3$ (by Lemma 3.2(b) and $n \geq 6$). Also, these faults satisfy the conditions for E. Therefore, $e \geq x \geq 3$.

(b) When $p = 1$,

- if there is no faulty neighbor of $H[\epsilon]$ on dimensions 2 to m:

   Then $0*^{k-2} 1*^{m-k} 1^{n-m}$ must contain at most 1 fault in $0*^{n-1}$ for all $2 \leq k \leq m$. Therefore, the set of all faults in $0*^{n-1}$ satisfies the conditions of E. Hence $e = n - 2 - p \geq 3$, as $n \geq 6$.

- if there is at least one neighboring fault, $0^{k-1} 1 0^{n-k}$ where $2 \leq k \leq m$:

   Then the required set can be formed by having the fault $00^{k-2} 1 0^{n-k}$ together with one more fault which does not belong to the subcube $0*^{k-2} 1*^{n-k}$. Then $e \geq 2$. Note that we can always find such an additional fault; otherwise, the subcube $0*^{k-2} 1*^{n-k}$ would contain all the faults in $0*^{n-1}$ and dimension k would have been moved to one of the rightmost $n - m$ dimensions to satisfy condition (B2).   □

Lemma B.2: If $n - m > 0$, then

1) the set of all faults in $0*^{n-1}$ satisfies the condition of E and
2) $p = 1$.

Proof:

1) When $n - m > 0$, there does not exist a k such that $2 \leq k \leq m$ and $0^{k-1} 1 0^{n-k}$ is faulty (by condition (B2)). Similar to the proof for Lemma B.1(b) where there is no faulty neighbors of $H[\epsilon]$, the set of all faults in $0*^{n-1}$ satisfies the condition of E.

2) If all subcubes $0*^{k-2}1*^{n-k}$ contains 0 faults, all the faults would be identical to 0^m1^{n-m}. Hence there must be some of them containing 1 fault. Then $p = 1$. □

Lemma B.3: It is always possible to perform the LR^{i-1}-embedding for $i = 1, \cdots, \lfloor m/2 \rfloor$.

Proof: First, the L-embedding can always be performed since $C[L]$ contains only $p \leq 1$ faults and zero assigned nodes above level 1.

Now, consider the LR^{i-1}-embedding for $i = 2, \cdots, \lfloor m/2 \rfloor$. The subcubes $0^{j-1}1*^{n-j}$ where $2 \leq j \leq i$ can, in total, contain at most $i - 1$ faults of the set E and all faults in $0*^{n-1}$ not in E. In other words, they contain at most $n-2-(p+e)+(i-1)$ faults where e is the size of E as defined in Lemma B.1. From condition (B3), $C[LR^{j-1}] = 0^{j-1}1*^{n-j}$ contains at least as many faults as in $C[LR^{i-1}] = 0^{i-1}1*^{n-i}$ for $2 \leq j \leq i$. Therefore, the number of faults in $C[LR^{i-1}]$ is at most $\lfloor n-2-(p+e)+(i-1)/i-1 \rfloor$ which is $\leq n - i - 2$ provided $n-2-(p+e)+(i-1)-(n-i-2)(i-1) \leq i-2$ or $(i-2)(n-3-i)+(e+p)-3 \geq 0$. This condition is always satisfied as $p + e \geq 3$ (Lemma B.1) and $(i-2)(n-3-i) \geq 0$ ($n \geq 6$ and $2 \leq i \leq \lfloor m/2 \rfloor$). Moreover, $C[LR^{i-1}]$ does not contain any assigned nodes above level i. Hence the LR^{i-1}-embedding can always be done. □

Lemma B.4: It is always possible to perform the $LR^{\lfloor m/2 \rfloor}$-embedding unless n is odd and $x = n - 2$.

Proof: Consider $C[LR^{\lfloor m/2 \rfloor}] = 0^{\lfloor m/2 \rfloor} *^{n-\lfloor m/2 \rfloor -1} \bar{a}_n$ where

$$a_n = \begin{cases} 0 & n > m \\ 1 & n = m \end{cases}.$$

It may contain some faults but no assigned nodes above level $\lfloor m/2 \rfloor + 1$ except $H[\epsilon]$ when $n = m$. Now, we claim that at most one fault in $C[LR^{\lfloor m/2 \rfloor}]$ has a 1 in dimensions $\lfloor m/2 \rfloor + 1$ to $\min(m, n - 1)$, i.e., for $i = 1, \cdots, \min(\lceil m/2 \rceil, n - \lfloor m/2 \rfloor - 1)$, the subcube $0^{\lfloor m/2 \rfloor} *^{i-1} 1 *^{n-\lfloor m/2 \rfloor -1-i} \bar{a}_n$ contains at most 1 fault. This is obvious when $n > m$ (Lemma B.2a). Consider $n = m$. If for some i, the subcube $0^{\lfloor m/2 \rfloor} *^{i-1} 1 *^{n-\lfloor m/2 \rfloor -1-i} \bar{a}_n$ contains ≥ 2 faults, then by condition (B3), the total number of faults in the n-cube is $\geq 2\lfloor m/2 \rfloor = 2\lfloor n/2 \rfloor \geq n - 1$ which is impossible. Hence our claim follows.

Next, we consider the number of faults with Hamming distance $\leq n - \lfloor m/2 \rfloor - 3$ from

$$H[LR^{\lfloor m/2 \rfloor}] = \begin{cases} 0^{\lceil m/2 \rceil} 1^{\lfloor m/2 \rfloor} 0^{n-m-1} 1 & \text{if } n > m \\ 0^{\lceil m/2 \rceil} 1^{\lfloor m/2 \rfloor -1} 0 & \text{if } n = m \end{cases}.$$

A. Case (i) even m

A fault in $C[LR^{\lfloor m/2 \rfloor}]$ is at Hamming distance $\leq n - \lfloor m/2 \rfloor - 3$ from $H[LR^{\lfloor m/2 \rfloor}]$ if and only if it agrees with $H[LR^{\lfloor m/2 \rfloor}]$ in at least 2 dimensions from dimensions $\lfloor m/2 \rfloor + 1$ to $\min(m, n - 1)$. (They must disagree in dimensions $\min(m, n - 1) + 1$ to $n - 1$ by condition (B2).) In other words, it must have 1 in at least 2 of these dimensions. Hence there are $\leq \lfloor \min(m, n-1) - (m/2)/2 \rfloor = \lfloor \min((m/2), n - (m/2) - 1)/2 \rfloor$ such faults. Note that this is $\leq n - \lfloor m/2 \rfloor - 3$ provided $n + (n - m) \geq 8$ or $n + 3(n - m) \geq 10$.

- When $n - m \geq 2$, the first inequality holds (as $n \geq 6$).
- When $n - m = 1, n \geq 7$ (so that m is even). Hence the second inequality is satisfied.
- When $n - m = 0$, the first inequality is true if $n \geq 8$. It is also impossible to have $n = 7$ (for m is even). If $n = 6$, $\lfloor \min((m/2), n - (m/2) - 1)/2 \rfloor = 1$. The only fault here is $000110 = 0^{\lceil m/2 \rceil} 1^{\lfloor m/2 \rfloor -1} 0$ which is impossible by condition (B4). Hence the $LR^{\lfloor m/2 \rfloor}$-embedding can be done recursively.

B. Case (ii) odd m

If $C[LR^{\lfloor m/2 \rfloor}] = 0^{\lfloor m/2 \rfloor} *^{n-\lfloor m/2 \rfloor -1} \bar{a}_n$ contains at least 1 fault, then so does $0^{\lfloor m/2 \rfloor} 1 *^{n-\lfloor m/2 \rfloor -2} \bar{a}_n$ (condition (B3)). Consider a fault A in this subcube (there can be only one fault in this subcube). There are two cases to be considered.

1) If fault A is at Hamming distance $\leq n - \lfloor m/2 \rfloor - 3$ from $H[LR^{\lfloor m/2 \rfloor}]$, it has a 1 in at least 2 dimensions from $\lfloor m/2 \rfloor + 2$ to $\min(m, n - 1)$. Hence the total number of faults in $C[LR^{\lfloor m/2 \rfloor}]$ that can affect the embedding is $\leq 1 + \min(m, n - 1) - \lfloor m/2 \rfloor - 3$ or $\leq \min(\lceil m/2 \rceil - 2, n - \lfloor m/2 \rfloor - 3)$ or $\leq n - \lfloor m/2 \rfloor - 3$.

2) If fault A does not affect the embedding, then it has at most one 1 in dimensions from $\lfloor m/2 \rfloor + 2$ to $\min(m, n - 1)$.

   a) When $n - m \geq 1$, there should be at most $n - \lfloor m/2 \rfloor - 3$ faults that can affect the embedding. Otherwise the total number of faults in $C[\epsilon]$ (including the fault A which cannot affect the embedding) would be $\geq (n - \lfloor m/2 \rfloor - 2) + (p + \lfloor m/2 \rfloor) \geq n - 1$ (as $p = 1$ by Lemma B.2(b)).

   b) When $n - m = 0$, it is impossible to have $n - \lfloor m/2 \rfloor - 2$ faults. Otherwise, the $n - 2$ faults must be $00^{i-2}10^{n-i}$ for $i = 2, \cdots, n-1$. In other words, they are all neighbors of $H[\epsilon] = 0^n$, i.e., $x = n - 2$. Since $n = m$ and m is odd, n is odd. The case where n is odd and $x = n - 2$ is excluded from consideration here. □

REFERENCES

[1] F. Annexstein, "Fault tolerance of hypercube derivative networks," in *Proc. 1st Annual ACM Symp. Parallel Algorithms and Architectures*, 1986, pp. 179–188.

[2] B. Becker and H. U. Simon, "How robust is the n-cube?" *Information and Computation 77*, no. 2, pp 162–178, 1988.

[3] S. Bhatt, F. Chung, T. Leighton, and A. Rosenberg, "Optimal simulations of tree machines," in *Proc. 27th Annual Symp. Foundations Comput. Sci.*, 1986, pp. 274–282.

[4] S. N. Bhatt and I. C. F. Ipsen, "How to embed trees in hypercubes," Yale University, Res. Rep.YALEU/DCS/RR-443, Dec. 1985.

[5] J. Bruck, R. Cypher, and D. Soroker, "Running algorithms efficiently on faulty hypercubes," in *Proc. 2nd Annual ACM Symp. on Parallel Algorithms and Architectures*, 1990, pp. 37–44.

[6] M. Y. Chan, "Embedding of grids into optimal hypercubes," *SIAM J. Comput., SICOMP 20-5* pp. 834–864, Oct. 1991.

[7] M. Y. Chan and F. Chin, "On embedding rectangular grids in hypercubes," *IEEE Trans. Comput. 37*, pp. 1285–1288, 1988.

[8] ———, "A parallel algorithm for an efficient mapping of grids into hypercubes," *IEEE Trans. Parallel Distrib. Syst.*, vol. 4, pp. 933–946, Aug. 1993.

[9] M. Y. Chan and S. J. Lee, "Fault-tolerant embeddings of complete binary trees in hypercubes," *IEEE Trans. Parallel Distrib. Syst.*, vol. 4, pp. 277–288, Mar. 1993.
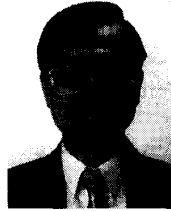
[10] ____, "On the existence of Hamiltonian circuits in faulty hypercubes," *SIAM J. Discrete Mathematics*, vol. 4, no. 4, pp. 511–527, Nov. 1991.

[11] C. T. Ho and S. L. Johnsson, "Embedding meshes in Boolean cubes by graph decomposition," *J. Parallel Distrib. Comput.*, pp. 325–339, Apr. 1990.

[12] J. Hastad, T. Leighton, and M. Newman, "Reconfiguring a hypercube in the presence of faults," in *Proc. 19th Annual ACM Symp. Theory Comput.*, 1987, pp. 274–284.

[13] J. Hastad, T. Leighton, and M. Newman, "Fast computation using faulty hypercubes," *Proc. 21st Annual ACM Symp. Theory Comput.*, 1989, pp. 251–263.

[14] T. Leighton, M. Newman, A. G. Ranade, and E. Schwabe, "Dynamic tree embeddings in butterflies and hypercubes," *Proc. 2nd Annual ACM Symp. Parallel Algorithms and Architectures*, 1989, pp. 224–234.

[15] M. Livingston and Q. Stout, "Embeddings in hypercubes," *Mathematical Computer Modelling*, vol. 11, pp. 222–227.

[16] M. Livingston, Q. Stout, N. Graham, and F. Harary, "Subcube fault-tolerance in hypercubes," Univ. of Michigan Computing Research Laboratory, Tech. Rep. CRL-TR-12-87, Sept. 1987.

[17] B. Monien and I. H. Sudborough, "Simulating binary trees on hypercubes," *Proc. 3rd Aegean Workshop on Comput.*, 1988, pp. 170–180.

[18] A. Wagner, "Embedding arbitrary binary trees in a hypercube," *J. Parallel Distrib. Comput.* 7, no. 3, pp. 503–520, 1989.

[19] A. Wagner and D. G. Corneil, "Embedding trees in a hypercube is NP-complete," *SIAM J. Comput. 19*, no. 3, pp. 570–590, 1990.

[20] A. Y. Wu, "Embedding of tree networks into hypercubes," *J. Parallel Distrib. Comput. 2*, pp. 238–249, 1985.

[21] A. Wang, R. Cypher, and E. Mayr, "Embedding complete binary trees in faulty hypercubes," Tech. Rep. RJ 7821 (72203), Nov. 1990.

Francis Y. L. Chin (S'71–M'76–SM'85) received the B.Sc. degree in engineering science from the University of Toronto, Toronto, Canada, in 1972, and the M.S., M.A., and Ph.D. degrees on electrical engineering and computer science from Princeton University, NJ, in 1974, 1975, and 1976, respectively.

Since 1975, he has taught at the University of Maryland, Baltimore County, University of California, San Diego, University of Alberta, and Chinese University of Hong Kong. He is currently Head of the Department of Computer Science, the University of Hong Kong. His current research interests include algorithm design and analysis, and parallel and distributed computing.

Dr. Chin has served on the program committees of numerous international workshops and conferences including the International Workshop on Statistical and Scientific Database Management (1981, 1983, 1986, and 1988), the Annual International Symposium on Algorithms (1991 and 1992), the International Conference on Computer Processing of Chinese and Oriental Languages (1988, 1989, and 1991), the International Computer Symposium (Taiwan, 1992) and the Fourth IEEE Symposium on Parallel and Distribution Processing (Dallas, TX, 1992). He has also served as Conference Chairman of the International Symposium on Algorithms and Computation (1993).

Bethany M. Y. Chan received the the B.A. and M.S. degrees in computer science from the University of California, San Diego, in 1981 and 1980, respectively, and the Ph.D. degree in computer science from the University of Hong Kong in 1988.

In 1982, she was an Assistant Lecturer at the Chinese University of Hong Kong. From 1982–1987, she was an Assistant Lecturer and Lecturer at the University of Hong Kong. From 1987–1990, she was an Assistant Professor at the University of Texas at Dallas. She is currently an Honorary Lecturer at the University of Hong Kong and an assistant director of investment research at Worldsec International Ltd.

Chung-Keung Poon received the B.Sc.(Elec.) and M.Phil. degrees from the Department of Electrical Engineering and the Department of Computer Science, the University of Hong Kong, in 1989 and 1991, respectively. He is working toward the Ph.D. degree at the University of Toronto, Toronto, Ontario, Canada. His research interests include fault tolerant computing, parallel processing, and computation complexity.