Proceedings of International Conference
on Algorithms

Efficient Broadcasting Algorithms in Faulty Hypercubes*

Qian-Ping Gu and Shietung Peng

Department of Computer Software
The University of Aizu
Aizu-Wakamatsu, Fukushima, 965-80 Japan
E-mail:{qian/s-peng}@u-aizu.ac.jp

Abstract

In this paper, we consider fault-tolerant broadcasting in MIMD hypercubes in which each node can send message to all its neighbors simultaneously, however, each node can receive message from only one of its neighbors at each time step. Let H_n denote the n-dimensional hypercube. H_n is called d-safe if each node of H_n has at least d fault-free neighbors. We give efficient broadcasting algorithms for d-safe ($d \geq 1$) H_n with up to $2^d(n-d)-1$ faulty nodes. For the case of $d=1$, our algorithm completes the broadcasting in optimal time steps. For $d>1$, we show an algorithm which completes the broadcasting in $n+O(d^2)$ time steps. All the algorithms have optimal traffic steps.

1 Introduction

Broadcasting is the process of transmitting messages from one processor, called source, to all other processors, and is one of the key communication shemes in interconnected multi-processor systems. Broadcasting provides basic functions to implement distributed agreement, clock synchronization, and broadcast-and-aggregate type of algorithms. Time steps and traffic steps are the main criteria used to measure the performance of broadcasting at the system level. The maximum number of links the message traverses to reach one of the destinations is defined as time steps, and the total number of distinct links the message traverses to reach all destinations is measured in traffic steps. In this paper, we consider fault-tolerant broadcasting in hypercube-connected multicomputers. There are several communication models proposed for broadcasting in multicomputers. A restricted one is single-port SIMD model in which only the communication along the same dimension is allowed at each time step. While in multi-port MIMD model, a node can send message to every neighbor in one time step. In this paper, we adopt the multi-port MIMD model in which every node can receive message from only one of the neighbor nodes at each time step. Fault-tolerant broadcasting schemes can be classified by the fault information used at each node. In local-information based schemes, each node knows only the status of its adjacent links and nodes. On the other hand, global-information based schemes assume that each node knows the distribution of faults in the network. In this paper, we develop broadcasting schemes based on global-information.

Hypercubes are interesting interconnection networks and have been adopted in many commercial machines. Let H_n denote the n-dimensional hypercube. H_n is n-connected and can tolerate as many as $n-1$ arbitrary faulty nodes for broadcasting. Many fault-tolerant broadcasting schemes for H_n have been developed [6, 9, 7, 1, 8, 11, 9, 10]. An $n+12$ time steps local-information, source independent broadcasting algorithm for SIMD H_n with at most $n-1$ faults was given in [7]. A source dependent broadcasting algorithm for SIMD H_n with better time steps was given in [1]. Global information of faulty nodes could improve the time for broadcasting. An optimal $n+1$ time steps broadcasting scheme for an MIMD H_n with at most $n-1$ faults was given in [8]. If the number of faulty nodes is beyond $n-1$, H_n could be disconnected. However, the connectivity is a worst-case measure in the sense that n node failures can disconnect H_n only if they are all neighbors of a particular node. This is unlikely to happen in practice. It is known that H_n is still connected if H_n has at most $2^d(n-d)-1$ faulty nodes and each node of H_n has at least d fault-free neighbors [3, 5]. We call H_n *d-safe* if each node of H_n has at least d fault-free neighbors.

In this paper, we give efficient broadcasting algorithms for d-safe ($d \geq 1$) MIMD H_n with at most $2^d(n-d)-1$ faulty nodes. Let F be the set of faulty nodes in a d-safe H_n. We first construct a spanning tree of $H_n - F$ with the source as the root. The broadcasting is then done on the spanning tree that implies the traffic steps are optimal. For $d=1$, a spanning tree of $H_n - F$ with height at most $n+2$ (optimal) can be found in $O(n^2)$ time. For $d>1$, a spanning tree of $H_n - F$ with height $n+O(d^2)$ can be found in $O(n^2+|F|)$ time. The broadcasting on the spanning tree takes h time steps where h is the height of the spanning tree. Thus, for $d=1$, our broadcasting algorithm is time steps optimal as well. Our algorithms are source independent.

The rest of this paper is organized as follows: Section 2 gives the preliminaries of the paper. The broadcasting algorithms for $d=1$ and $d>1$ are given in

*This work was supported by the Founding of Group Research Projects at The University of Aizu

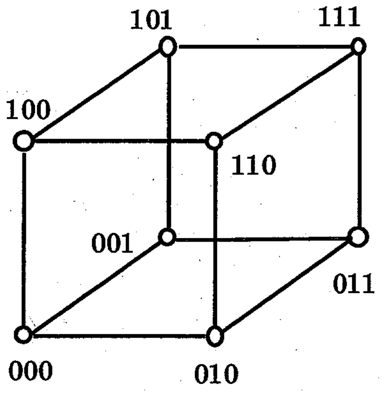

Figure 1: A 3-dimensional hypercube.

Sections 3 and 4, respectively. Section 5 concludes the paper.

2 Preliminaries

A path in a graph is a sequence of edges of the form $(s_1, s_2)(s_2, s_3)\ldots(s_{k-1}, s_k)$, $s_i \in V$, $1 \le i \le k$, and $s_i \ne s_j$, $i \ne j$. The length of a path is the number of edges in the path. We sometimes denote the path from s_1 to s_k by $s_1 \to s_k$. A path P is called *fault-free* if all nodes in P are non-faulty. For any two nodes $s, t \in G$, $d_G(s,t)$ denotes the distance between s and t in G, i.e., the length of the shortest path connecting s and t. The diameter of G is defined as $d(G) = \max\{d_G(s,t) | s, t \in G\}$.

An n-dimensional hypercube H_n is an undirected graph on the node set $H_n = \{0,1\}^n$ such that there is an edge between $u \in H_n$ and $v \in H_n$ if and only if u and v differ exactly in one bit position. Figure 1 gives an H_3. H_n is n-connected, has 2^n nodes, and has diameter $d(H_n) = n$. For $n \ge 1$, the 0-subcube of H_n on dimension i, denoted as H^0_{n-1}, is defined to be the subgraph of H_n induced by the set of nodes whose ith bit are 0. Define similarly the 1-subcube H^1_{n-1}. H^0_{n-1} and H^1_{n-1} are both isomorphic to H_{n-1} and are connected to each other by edges in dimension i of H_n. For a node $s = a_1a_2\ldots a_n \in H_n$, $s^{(i)}$, $1 \le i \le n$, denotes the node $a_1\ldots a_{i-1}\bar{a_i}a_{i+1}\ldots a_n$, where $\bar{a_i}$ is the logical negation of a_i. Similarly, $s^{(i_1,i_2,\ldots,i_k)}$ denotes the node $b_1\ldots b_n$, where $b_{i_j} = \bar{a_{i_j}}$, $1 \le j \le k$, and $b_l = a_l$ for $l \in <n> -\{i_1,\ldots,i_k\}$, where $<n> = \{1,2,...,n\}$ and $<n> -\{i_1,...,i_k\} = \{j | j \in <n>, j \notin \{i_1,...,i_k\}\}$.

Binomial spanning trees are powerful tools for broadcasting in hypercubes [4, 2]. A *binomial spanning tree* T of H_n with root s is recursively defined as follows: For $n = 0$, $T = \emptyset$. For $n > 0$, let H^0_{n-1} and H^1_{n-1} be the subcubes partitioned on some dimension i. Assume $s \in H^0_{n-1}$ and $s' \in H^1_{n-1}$ is the neighbor of s. Then $T = \{(s,s')\} \cup T_0 \cup T_1$, where T_0 is a binomial spanning tree of H^0_{n-1} with root s and T_1 is a binomial spanning tree of H^1_{n-1} with root s'. The above binomial spanning tree of H_n has height n. Figure 2 gives a binomial spanning tree of the H_3 with root $s = 000$. The procedure in Figure 3 generates a binomial spanning tree of H_n. Notice that there are n choices to partition H_n in procedure BST for each n. Based on this fact, when $|F| \le n-1$, a spanning tree of $H_n - F$ with optimal height $n+1$ can be generated as shown in Figure 4 [8]. It is easy to see that the procedure ST0 constructs a spanning tree of $H_n - F$ with root s and height at most $n+1$ in $O(n^2)$ time.

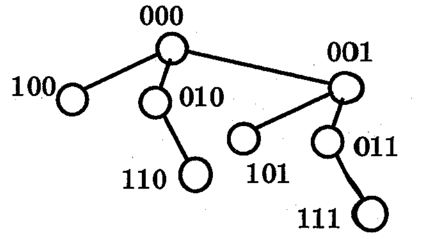

Figure 2: A binomial spanning tree of H_3.

Procedure BST(s, H_n)
Input: A source node s in H_n.
Output: A binomial spanning tree of H_n with root s.
begin
  if $(n = 0)$ **then** return;
  Partition H_n into H^0_{n-1} and H^1_{n-1};
/*Assume $s \in H^0_{n-1}$ and $s' \in H^1_{n-1}$ is the neighbor of s.*/
  keep the edge (s, s') at s;
  BST(s, H^0_{n-1}); BST(s', H^1_{n-1});
end

Figure 3: A binomial spanning tree generator.

Lemma 1 *For $|F| \le n-1$, let T be the spanning tree of $H_n - F$ with root s and height $n+1$ generated by ST0. Then there is exactly one node t in T with $d_T(s,t) = n+1$.*

Proof: Assume T is generated in Case 1. Since the height of T is $n+1$, T is generated in the case that $|F_1| = 0$. In this case, a binomial spanning tree T' of H^1_{n-1} with root s' and height $n-1$ is generated, and each fault-free node of H^0_{n-1} is connected to its neighbor in H^1_{n-1}. Obviously, there is exactly one node t' in the binomial spanning tree with $d_{T'}(s',t') = n-1$. Let $t \in H^0_{n-1}$ be the neighbor of t'. Then t is the only node with $d_T(s,t) = d_{T'}(s',t') + 2 = n+1$.

If T is generated in Case 2, then by recursion we can get a spanning tree T_0 of $H^0_{n-1} - F$ with root s and height n and a spanning tree T_1 of $H^1_{n-1} - F$ with root s' and height n. There is exactly one node t' in T_1 with $d_{T_1}(s',t') = n$. Thus, this t' is the only node that $d_T(s,t') = n+1$. □

Procedure ST0(s, F, H_n)
Input: A source node s and a set F ($|F| \le n-1$) of fault nodes in H_n.
Output: A spanning tree of $H_n - F$ with root s and height at most $n+1$.
begin
  if $(n = 0)$ **then** return;
  Find a fault-free neighbor s' of s and partition H_n into H^0_{n-1} and H^1_{n-1} with $s \in H^0_{n-1}$ and $s' \in H^1_{n-1}$;
  Keep the edge (s, s') at s;
  $F_0 := F \cap H^0_{n-1}$; $F_1 := F \cap H^1_{n-1}$;
  Case 1: $|F_0| = 0$ or $|F_1| = 0$.
    if $(|F_0| = 0)$ **then** {BST(s, H^0_{n-1});
    $\forall(t' \in H^1_{n-1} - (F \cup \{s'\}))$, keep edge (t, t') at t;}
    else {BST(s', H^1_{n-1});
    $\forall(t \in H^0_{n-1} - (F \cup \{s\}))$, keep edge (t', t) at t';}
  Case 2: $|F_0| > 1$ and $|F_1| > 1$.
    ST0(s, F_0, H^0_{n-1}); ST0(s', F_1, H^1_{n-1});
end

Figure 4: An algorithm which produces a spanning tree of $H_n - F$ for $|F| \le n-1$.

We will use procedures BST and ST0 as subroutines in our algorithms.

3 Broadcasting in 1-safe hypercubes

In this section, we show an algorithm which constructs a spanning tree of $H_n - F$, where H_n is 1-safe and $|F| \le 2n-3$, with the optimal height. The idea of the algorithm is as follows: First, we find a fault-free neighbor s' of the source s and partition H_n into H^0_{n-1} and H^1_{n-1} such that s and s' are separated into different subcubes. Assume $s \in H^0_{n-1}$ and $s' \in H^1_{n-1}$. Since $|F| \le 2n-3$, at least one subcube has at most $n-2$ faults. Assume H^1_{n-1} does (the case of $|F_0| \le n-2$ is done similarly). Then it is known that a spanning tree T_1 of $H^1_{n-1} - F$ with root s' can be found. We recursively find a spanning tree of $H^0_{n-1} - F$ with root s if H^0_{n-1} is 1-safe. If H^0_{n-1} is not 1-safe, then there is a node $u \in H^0_{n-1}$ such that all the neighbors of u in H^0_{n-1} are faulty. If $u = s$, we connect each node $t \in (H^0_{n-1} - (F \cup \{s\}))$ to T_1 by a fault-free path. If $u \ne s$, we connect u to T_1 by a fault-free path and find a spanning tree of $H^0_{n-1} - (F \cup \{u\})$ with root s.

To show the algorithm, we first give two procedures which handle some special cases of the problem. The first procedure ST1 (see Figure 5), given F with $|F| \le n-2$ and s in $H_n - F$, finds a spanning tree of $H_n - F$ with root s and height n. Node that procedure ST1 has some independent interests as well.

Lemma 2 *Given F with $|F| \le n-2$ and $s \in H_n - F$, procedure ST1 constructs a spanning tree T of $H_n - F$ with root s and height n.*

Procedure ST1(s, F, H_n)
Input: A set F of faulty nodes in H_n with $|F| \le n-2$, and node $s \in H_n - F$.
Output: A spanning tree of $H_n - F$ with root s and height at most n.
begin
  Find a fault-free neighbor s' of s and partition H_n into H^0_{n-1} and H^1_{n-1} with $s \in H^0_{n-1}$ and $s' \in H^1_{n-1}$;
  $F_0 := F \cap H^0_{n-1}$; $F_1 \cap H^1_{n-1}$;
  Case 1: $|F_0| = 0$ or $|F_1| = 0$.
    if $(|F_0| = 0)$ **then** {BST(s, H^0_{n-1});
    $\forall t' \in (H^1_{n-1} - (F \cup \{s'\}))$, keep edge (t, t') at t;}
    else {ST0(s, F, H^0_{n-1}); BST(s', H^1_{n-1});}
  Case 2: $|F_0| \ge 1$ and $|F_1| \ge 1$.
    ST1(s, F_0, H^0_{n-1}); ST1(s', F_1, H^1_{n-1});
end

Figure 5: An algorithm which produces a spanning tree of $H_n - F$ for $|F| \le n-2$.

Proof: Assume T is constructed in Case 1. If $|F_0| = 0$ then a binomial spanning tree of H^0_{n-1} with root s and height $n-1$ is constructed first and then the nodes of $H^1_{n-1} - F$ are connected to the binomial tree by paths of length 1. That is, the height of T is n. If $|F_1| = 0$, then a binomial spanning tree of H^1_{n-1} with root s' and height $n-1$ and a spanning tree of $H^0_{n-1} - F$ with root s and height n are generated. Obviously the height of T in this case is n as well.

If T is constructed in Case 2, each $(n-1)$-dimensional subcube has at most $n-3$ faults, and by the recursion, a spanning tree of height $n-1$ can be obtained. This implies the height of T is n. □

Let F be the set of faulty nodes in H_n. We say a node u of H_n is disconnected from H_n by F, if u is not faulty and all neighbors of u are in F. The second procedure ST2 (see Figure 6), given F with $|F| \le 2n-3$ in H_n such that F disconnects u from H_n and a node $s \in (H^n - (F \cup \{u\}))$, finds a spanning tree of $H^n - (F \cup \{u\})$ with root s and height at most $n+1$.

Lemma 3 *Given a set F of faulty nodes in H^n such that $|F| \le 2n-3$ and F disconnects u from H_n and a source node $s \in (H^n - (F \cup \{u\}))$, procedure ST2 finds a spanning tree T of $H^n - (F \cup \{u\})$ with root s and height at most $n+1$.*

Proof: For any node $v \in H_n$, define $N(v)$ the set of neighbor nodes of v. Let u be the node disconneted by F. Then $N(u) \subseteq F$ and $|F - N(u)| \le n-3$. Obviously, $|N(s) \cap N(u)| \le 2$. Therefore, s has at least one fault-free neighbor s'. Partition H_n into H^0_{n-1} and H^1_{n-1} with $s \in H^0_{n-1}$ and $s' \in H^1_{n-1}$. Assume $u \in H^0_{n-1}$. The case that $u \in H^1_{n-1}$ can be done similarly. Since $|F| \le 2n-3$, $|N(u) \cap H^0_{n-1}| = n-1$, and $|N(u) \cap H^1_{n-1}| = 1$, we have $|F_0| \ge n-1$ and $1 \le |F_1| \le n-2$. If $|F_1| = 1$, then the only

Procedure ST2(s, F, H_n)
Input: A set F of faulty nodes in H_n with
$|F| \le 2n-3$, F disconnects u from H_n,
and $s \in H^n - (F \cup \{u\})$.
Output: A spanning tree of $H_n - (F \cup \{u\})$ with
root s and height at most $n+1$.
begin
Find a fault-free neighbor s' of s and partition H_n
into H^0_{n-1} and H^1_{n-1} with $s \in H^0_{n-1}$ and $s' \in H^1_{n-1}$;
Connect s' to s; $F_0 := F \cap H^0_{n-1}$; $F_1 \cap H^1_{n-1}$;
/*Assume $u \in H^0_{n-1}$. $|F_0| \ge n-1$
and $1 \le |F_1| \le n-2$.*/
if $(|F_1| = 1)$ **then**
{ST1(s', F_1, H^1_{n-1});
/*Finds a spanning tree T_1 of $H^1_{n-1} - F$*/
$\forall t \in H^0_{n-1} - (F \cup \{s\})$, connect t to T_1
by a path of length 1;}
else {ST0(s', F_1, H^1_{n-1}); ST2(s, F_0, H^0_{n-1});}
end

Figure 6: An algorithm which produces a spanning tree of $H_n - F$ where F disconnects u.

faulty node of H^1_{n-1} is the neighbor of u in H^1_{n-1}. By procedure ST1, we can find a spanning tree T_1 of $H^1_{n-1} - F$ with root s' and height $n-1$. Obviously, each node $t \in (H^0_{n-1} - (F \cup \{u\}))$ has a fault-free neighbor in H^1_{n-1}. The height of T found in this case is $(n-1)+1+1 = n+1$.

Assume $|F_1| \ge 2$. Then $|F_0| \le 2n-5$ and we can recursively find a spanning tree T_0 of $H^0_{n-1} - (F \cup \{u\})$ with root s and height at most $(n-1)+1$. Since $|F_1| \le n-2$, we can find a T_1 of $H^1_{n-1} - F$ with root s' and height at most $(n-1)+1$. The height of T found in this case is at most $n+1$ as well. □

Now, we show our algorithm for the case that $|F| \le 2n-3$ in Figure 7.

Theorem 4 *Given F with $|F| \le 2n-3$ and $s \notin F$ in 1-safe H_n, algorithm STree1 finds a spanning tree T of $H_n - F$ with root s and height at most $n+2$ in $O(n^2)$ time.*

Proof: Assume T is found in Case 1. From Lemma 2, we can find a spanning tree T_1 of $H^1_{n-1} - F$ with root s' and height $n-1$. Since H_n is 1-safe, each node of $H_n - F$ has at least 1 fault-free neighbor. For each $t \in (H^0_{n-1} - (F \cup \{s\}))$, if t has a fault-free neighbor in H^1_{n-1}, then t can be connected to T_1 by a fault-free path of length 1. So we assume the neighbor of t in H^1_{n-1} is faulty. Let u be the fault-free neighbor of t. Since H^1_{n-1} has at most one faulty node, the neighbor of u in H^1_{n-1} must be fault-free. Therefore, t can be connected to T_1 by a fault-free path of length 2. Thus, the height of T is at most $(n-1)+1+2 = n+2$.

Assume T is found in Case 2. By procedure ST0, we can find a spanning tree T_1 of $H^1_{n-1} - F$ with root

Algorithm STree1(s, F, H_n)
Input: A set F of faulty nodes in 1-safe H_n
with $|F| \le 2n-3$, and node $s \in H_n - F$.
Output: A spanning tree of $H_n - F$ with height
at most $n+2$.
begin
Find a fault-free neighbor s' of s;
Partition H_n s.t. $s \in H^0_{n-1}$ and $s' \in H^1_{n-1}$;
$F_0 := F \cap H^0_{n-1}$; $F_1 := F \cap H^1_{n-1}$;
/*Assume $|F_1| \le n-2$.*/
Case 1: $|F_1| \le 1$
ST1(s', F_1, H^1_{n-1});
/*Finds a spanning tree T_1 of $H^1_{n-1} - F$.*/
$\forall t \in H^0_{n-1} - (F \cup \{s\})$, connect t to T_1
by a fault-free path of length at most 2;
Case 2: $|F_1| \ge 2$.
ST0(s', F_1, H^1_{n-1});
/*Finds a spanning tree T_1 of $H^1_{n-1} - F$.*/
if (H^0_{n-1} is 1-safe) **then** STree1(s, F_0, H^0_{n-1})
else
/*$\exists$ node $u \in H^0_{n-1}$ which is disconnected by F_0.*/
if $(u \ne s)$ **then** $\{ST2(s, F_0, H^0_{n-1})$;
connect u to T_1 by the path of length 1;}
else $\forall t \in H^0_{n-1} - (F \cup \{s\})$, connect t to T_1
by a fault-free path of length at most 3;
end.

Figure 7: An algorithm which produces a spanning tree of $H_n - F$ for $|F| \le 2n-3$.

s' and height at most $(n-1)+1 = n$. If we can find a spanning tree T_0 of $H^0_{n-1} - F$ with root s recursively, the height of T_0 is at most $(n-1)+2 = n+1$. And the height of T is at most $n+1$.

So, we assume that H^0_{n-1} is not 1-safe, i.e., F_0 disconnects a node u from H^0_{n-1}. If $u \ne s$, then from the fact that H_n is 1-safe and all the $n-1$ neighbors of u in H^0_{n-1} are faulty, u has a fault-free neighbor in H^1_{n-1}. Therefore, u can be connected to T_1 by the path by length 1. By procedure ST2, we can find a spanning tree of $H^0_{n-1} - (F_0 \cup \{u\})$ with root s and height n.

Assume $u = s$. In this case, we connect each $t \in (H^0_{n-1} - (F \cup \{s\}))$ to T_1 by a fault-free path of length at most 3 as follows. First, we observe that for partitioning H_n on dimension i and $t \in H^0_{n-1}$, the path $P_i : (t, t^{(i)})$ (of length 1) and paths $P_j : (t, t^{(j)})(t^{(j)}, t^{(j,i)})$ (of length 2), $1 \le j \le n$ and $j \ne i$, are the n paths from t to n distinct nodes of H^1_{n-1}. For each $t \in (H^0_{n-1} - (F \cup \{s\}))$ with $d_{H_n}(s,t) \ge 3$, the $n-1$ faulty nodes which are also the neighbors of s can not block any of the n paths for t. Since $|F| \le 2n-3$, there are at least two fault-free paths of length at most 2 which connect t to T_1. Note that T_1 may have height n. However, from Lemma 1 if T_1 has height n, there is only one node w with $d_{T_1}(s', w) = n$. Therefore, we can always connect a t with $d_{H_n}(s,t) \geq 3$ to a node $v \in T_1$ with $v \neq w$.

For each $(t \in H^0_{n-1} - (F \cup \{s\}))$ with $d_{H_n}(s,t) = 2$, the neighbors of s blocks two of the n paths for t, and thus, all the n paths may be blocked by the faulty nodes. If this happens, all the $n-2$ faulty nodes which are not the neighbors of s appear in the $n-2$ of the n paths for t. From this, we can find a fault-free path $t \rightarrow v \in H^0_{n-1}$ of length 2 such that $d_{H_n}(s,v) = 2$ and the neighbor of v in H^1_{n-1} is fault-free. Therefore, t can be connected to a node $t' \in T_1$ with $d_{T_1}(s',t') = 2$ by a fault-free path of length 3 (which implies $d_T(s,t) = 6$).

Summarizing the above, we can find a spanning tree T for $H_n - F$ with root s and height at most $n+2$ for $n \geq 4$. Note that it is trivial to prove the lemma for the case of $n \leq 3$.

Since each recursive step takes $O(n)$ time, it is easy to see that constructing the spanning tree of $H_n - F$ takes $O(n^2)$ time. This completes the proof. □

Since the broadcasting in $H_n - F$ is done on the spanning tree constructed in algorithm STree1, we have the following result.

Theorem 5 *In a 1-safe H_n containing up to $2n-3$ faulty nodes, broadcasting from the source node to all non-faulty nodes can be done in $n+2$ time steps and $2^n - |F|$ traffic steps.*

The traffic steps are optimal because the broadcasting is done on the spanning tree. We now show that the diameter of 1-safe $H_n - F$ with $|F| = 2n-3$ is at least $n+2$ that implies the time steps of our broadcasting algorithm are optimal as well. Let $s = 110\ldots0$, $s_1 = 100\ldots0$, $s_2 = 000\ldots0$, and $t = 11\ldots1$. Let $N(u) = \{v | d_{H_n}(u,v) = 1\}$. Choose $F = N(s) \cup N(s_1) - \{s, s_1, s_2\}$. Then we have $d_{H_n - F}(s,t) = n+2$ which implies $d(H_n - F) \geq n+2$.

4 Broadcasting in d-safe H_n with $d > 1$

To show the broadcasting algorithm for d-safe H_n with $d > 1$, we first give some important properties of H_n.

Lemma 6 *For a node $s \in H^0_{n-1}$, let G be a connected subgraph of H^0_{n-1} such that $s \in G$ and the degree of each node u of G with $d_{H_n}(s,u) \leq d$ is at least d. Then we can find a $G' \subseteq G$ with $|G'| = 2^d$ and $s \in G'$ such that for any $t \in G'$, $d_{H_n}(s,t) \leq d$, and there are at least $2^d(n-d)$ node-disjoint paths (except the common end nodes in G') of length at most 2 from the nodes of G' to $2^d(n-d)$ distinct nodes of H^1_{n-1}.*

Proof: We prove the lemma by induction on d. For $d = 0$, let $G' = \{s\}$. Assume H_n is partitioned into H^0_{n-1} and H^1_{n-1} on dimension i. Then the path $P_i : (s, s^{(i)})$ (of length 1) and paths $P_j : (s, s^{(j)})(s^{(j)}, s^{(j,i)})$ (of length 2), $1 \leq j \leq n$ and $j \neq i$, are the n paths from s to n distinct nodes of H^1_{n-1}.

Assume the lemma is true for $d = k$ and we prove it for $d = k+1$. Pick up any edge of G and assume

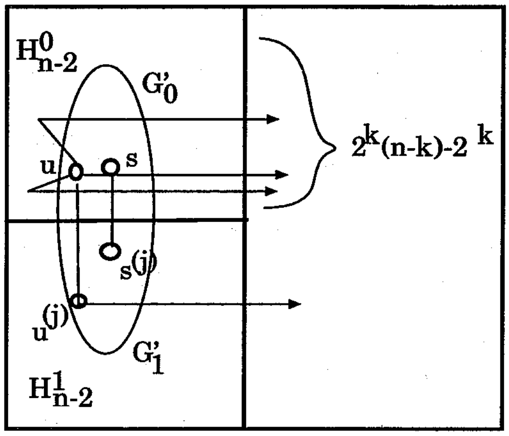

Figure 8: The paths from G' to H^1_{n-1}.

the edge is in dimension j. We further partition H^0_{n-1} into H^0_{n-2} and H^1_{n-2} on dimension j (see Figure 8). Let $G_0 = G \cap H^0_{n-2}$ and $G_1 = G \cap H^1_{n-2}$. Assume $s \in G_0$. Then $s^{(j)} \in G_1$. Then, each node u of G_0 with $d(s,u) \leq k$ has degree at least k. By the induction hypothesis, we can find a $G'_0 \subseteq G_0$ with $|G'_0| = 2^k$ and $s \in G'_0$ such that for any $t \in G'_0$, $d_{H_n}(s,t) \leq k$, and there are at least $2^k(n-k)$ disjoint paths from G'_0 to $2^k(n-k)$ distinct nodes of H^1_{n-1}. We can also find a $G'_1 \subseteq G_1$ with the similar properties. Let $G' = G'_0 \cup G'_1$. $|G'| = 2^{k+1}$ and $s \in G'$. Obviously, for any $t \in G'$, $d_{H_n}(s,t) \leq k+1$. For each node u of G'_0 (G'_1), there is exactly one path $u \rightarrow u^{(j)} \rightarrow u^{(j,i)} \in H^1_{n-1}$ which passes through H^1_{n-2} (H^0_{n-2}). Deleting the paths which pass through H^1_{n-2} (H^0_{n-2}), there are at least $2^k(n-k) - 2^k$ disjoint paths from G'_0 (G'_1) to H^1_{n-1} (see Figure 8). Therefore, the total number of disjoint paths of length at most 2 from G' to H^1_{n-1} is at least $2 \times (2^k(n-k) - 2^k) = 2^{k+1}(n-(k+1))$. Thus, the lemma holds. □

Lemma 7 *For $0 \leq d \leq n-1$, give a set F of faulty nodes in H_n such that $|F| \leq 2^d(n-d) - 1$ and each node of $H_n - F$ has degree at least d. Then for any $s \in H^0_{n-1} - F$, there is a fault-free path $s \rightarrow t \in H^1_{n-1}$ of length at most $d+2$.*

Proof: For $s \in H^0_{n-1} - F$, we check the nodes $u \in H^0_{n-1} - F$ such that $d_{H_n}(s,u) \leq d$ and u is connected to s. Let $u' \in H^1_{n-1}$ be the neighbor of u. If there is a fault-free u' then the lemma is proven. So we assume all such u' are faulty. From this and the fact that each of those nodes u has degree at least d in $H_n - F$, u has degree at least d in $H^0_{n-1} - F$. Therefore, there is

Algorithm STree2(s, d, F, H_n)
Input: A source node s and F with $|F| \leq 2^d(n-d)-1$ in d-safe H_n.
Output: A spanning tree of $H_n - F$ with root s and height $n + O(d^2)$.
begin
  if $(d = 1)$ **then** STree1(s, F, H_n);
  Find a fault-free neighbor s' of s and partition H_n into H^0_{n-1} and H^1_{n-1} with $s \in H^0_{n-1}$ and $s' \in H^1_{n-1}$;
  Keep edge (s, s') at s;
  $F_0 := F \cap H^0_{n-1}$; $F_1 := F \cap H^1_{n-1}$;
/*Assume $|F_1| \leq |F|/2$ ($|F|/2 \leq 2^{(d-1)}(n-d+1)-1$). The other case can be done similarly.*/
  STree2$(s', d-1, F_1, H^1_{n-1})$;
/*Finds a spanning tree T_1 of $H^1_{n-1} - F$.*/
  $\forall u \in (H^0_{n-1} - (F_0 \cup \{s\}))$, find a fault-free path $u \rightarrow t \in T_1$ of length at most $d+2$;
end

Figure 9: An algorithm which produces a spanning tree of $H_n - F$ for $|F| \leq 2^d(n-d)-1$.

a subgraph G of H^0_{n-1} such that $s \in G$ and for $u \in G$ with $d_{H_n}(s,u) \leq d$, u has degree at least d. From this and Lemma 6, we can find at least $2^d(n-d)$ paths of length at most $d+2$ from s to $2^d(n-d)$ distinct nodes of H^1_{n-1}. In addition, one faulty node can block at most one of the $2^d(n-d)$ paths. From $|F| \leq 2^d(n-d)-1$, the lemma holds. □

Theorem 8 *If H_n is d-safe and $|F| \leq 2^d(n-d)-1$, where F is the set of faulty nodes, then we can find a spanning tree of $H_n - F$ of height $n + O(d^2)$ in $O(n^2 + |F|)$ time.*

Proof: We show first that the algorithm STree2 generates a spanning tree T_n of $H_n - F$ of height $n + O(d^2)$. From the algorithm, it is easy to see that $h(T_n) \leq h(T_{n-1}) + d + 2$, where $h(T)$ is the height of the tree T. This implies $h(T_n) \leq h(T_{n-d}) + \sum_{i=1}^{d}(i+2) = n - d + 1 + \sum_{i=1}^{d}(i+2) = n + O(d^2)$. Since each iteration takes $O(|F|/2^i + n)$ time, the time complexity of the algorithm is then $O(n^2 + \sum_{i=1}^{d}(|F|/2^i + n)) = O(n^2 + |F|)$. □

Theorem 9 *In a d-safe H_n containing up to $2^d(n-d)-1$ faulty nodes, broadcasting from the source node to all non-faulty nodes can be done in $n + O(d^2)$ time steps and $2^n - |F|$ traffic steps.*

5 Conclusional Remarks

In this paper, we have shown an optimal broadcasting algorithm for 1-safe faulty hypercubes with up to $2n-3$ faulty nodes. However, the broadcasting algorithm for d-safe faulty hypercubes with $d > 1$ is not optimal. It can be shown that the lower bound of height of the spanning tree for $H_n - F$ is $n + d + 1$. Whether or not we can find a broadcasting algorithm which matches this lower bound remains open. To design efficient broadcasting algorithms in other interconnection networks under the d-safe fault tolerance model is also worth further research.

References

[1] Y. Chang. Fault tolerant broadcasting in simd hypercubes. In *Proc. IEEE Symposium on Parallel and Distributed Processing (SPDP'93)*, pages 348–351, 1993.

[2] K. Day and A. Tripathi. A comparative study of topological properties of hypercubes and star graphs. *IEEE Trans. on Parallel and Distributed Systems*, 5(1):31–38, 1994.

[3] A. H. Esfahanian. Generalized measures of fault tolerance with application to n-cube networks. *IEEE Trans. on Computers*, 38(11):1586–1591, 1989.

[4] S.L. Johnson and C.T. Ho. Optimum broadcasting and personalized communication in hypercubes. *IEEE Trans. on Computers*, 38(9):1249–1268, 1989.

[5] S. Latifi, M. Hedge, and M. Naraghi-Pour. Conditional connectivity measures for large multipprocessor systems. *IEEE Trans. on Computers*, 43(2):218–222, 1994.

[6] T.C. Lee and J.P. Hayes. A fault-tolerant communication scheme for hypercube computers. *IEEE Trans. on Computers*, pages 1242–1256, 1992.

[7] C.S. Raghavendra and M.A. Sridhar. Broadcasting in faulty simd hypercubes. In *Proc. IEEE Symposium on Parallel and Distributed Processing (SPDP'92)*, pages 4–11, 1992.

[8] C.S. Raghavendra, P.J. Yang, and S.B. Tien. Free dimensions—an efficient approach to archieving fault tolerance in hypercube. *IEEE Trans. on Computers*, Vol. 44:1152–1157, 1995.

[9] P. Ramanathan and K.G. Shin. Reliable broadcast in hypercube multiprocessors. *IEEE Trans. on Computers*, Vol. 37:1654–1657, 1988.

[10] P.W. Sui, S.D. Wang, and Y.Y. Lee. Broadcasting on faulty hypercubes. In *Proc. of the 1994 Inthernational Conference on Parallel and Distributed Systems (ICAPDS'94)*, pages 128–133, 1994.

[11] J. Wu. Safty levels-an efficient mechanism for achieving reliable broadcasting in hypercubes. *IEEE Trans. on Computers*, 44:702–706, 1995.