

Reliable Broadcasting in Product Networks

Feng Bao¹, Yoshihide Igarashi¹, Sabine R. Öhring²

鮑 豊

五十嵐 善英

ザビネ オーリング

Abstract

In this paper we study the reliable broadcasting in product networks. We suppose that the faulty nodes and faulty links may arbitrarily change the messages that pass through them, and may even fabricate messages. An n -channel network can tolerate $\lfloor (n-1)/2 \rfloor$ such arbitrary faults in broadcasting in the worst case. We prove that the product network of any n component networks is an n -channel network, and hence it can tolerate $\lfloor (n-1)/2 \rfloor$ faults in the worst case. If there are f faulty nodes randomly distributed in the n -product networks, the broadcasting succeeds with a probability higher than $1 - (4b^3nf/N)^{\lceil n/2 \rceil}$, where N is the node number of the n -product network and b is the upper bound of the node numbers of the n component networks. If only links may fail while all the nodes are healthy, then $\Theta(L)$ faulty links that are randomly distributed in the n -product network can be tolerated with high probability, where L is the link number of the network.

1 Introduction

1.1 Product Networks

Let $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ be two finite undirected graphs. The *cartesian product* of G_1 and G_2 is defined as $G = G_1 \times G_2$ with the node-set $V = V_1 \times V_2 = \{(x, y) \mid x \in V_1, y \in V_2\}$. There is an edge $\{(x, y), (u, v)\}$ in G iff either $x = u$ and $\{y, v\} \in E_2$, or $\{x, u\} \in E_1$ and $y = v$. The graphs G_1 and G_2 are called the *factors* or *component networks* of G . G consists of $|G_2|$ copies of G_1 , namely the subgraphs G_1x_2 with node-set $\{(x_1, x_2) \mid x_1 \in V_1\}$ and edge-set $\{(x, x_2), (y, x_2) \mid \{x, y\} \in E_1\}$. Analogously, G has $|G_1|$ copies x_1G_2 of G_2 induced by the node-set $\{(x_1, x_0) \mid x_2 \in V_2\}$.

Figure 1 shows an example of a product network.




Figure 1: Example of a product network.

This definition can be generalized to a product of n graphs $G = (V, E) = G_1 \times G_2 \times \dots \times G_n$ with $G_i = (V_i, E_i)$, $1 \leq i \leq n$. It holds $V = V_1 \times \dots \times V_n$ and $E = \{(x_1 \dots x_n, y_1 \dots y_n) \mid \exists i \in \{1, \dots, n\} \text{ with } \{x_i, y_i\} \in E_i \text{ and } x_j = y_j \text{ for } i \neq j\}$. An interconnection topology derived from several component networks by this product operation is called a *product network*.

Examples for product networks include the $(m_1 \times \dots \times m_n)$ -mesh (respectively torus) defined as $L_{m_1} \times \dots \times L_{m_n}$ (respectively $R_{m_1} \times \dots \times R_{m_n}$) (for a linear array L_j or a ring R_j of length j), the n -dimensional binary hyper-cube is $Q_n = Q_{n-1} \times K_2$, the generalized hyper-cube $GQ_b^n = GQ_b^{n-1} \times K_b$, where K_b is the complete graph of order b , $b \geq 2$, the hyper de Bruijn network $HD(m, n) = Q_m \times DG(n)$ (for the binary de Bruijn graph $DG(n)$ of order n) [5] and the *hyper Petersen* network $HP_n = Q_{n-3} \times P$ [4]. Here an n -dimensional *binary hypercube*, Q_n , has the node-set $V_n = Z_2^n = \{x_1 \dots x_n \mid x_i = 0 \text{ or } 1, 1 \leq i \leq n\}$, which is the set of binary strings of length n . There exists an edge between two nodes iff

¹Department of Computer Science, Gunma University, Kiryu, 376 Japan

²Department of Computer Science, University of North Texas, Denton, TX 76203-3886, USA

their binary labels differ in exactly one bit. A binary de Bruijn graph of order n , $DG(n)$, has the same node-set as Q_n and the edge-set $\{(x_1 \dots x_n, x_2 \dots x_n p) | p, x_i \in Z_2, 1 \leq i \leq n\}$.

Youssef has proven in [9] that for two graphs G_0 and G_1 , the product network $G = G_1 \times G_2$ has the diameter $d(G) = d(G_1) + d(G_2)$, the degree $\deg(G) = \deg(G_1) + \deg(G_2)$, the average distance $d_{avg}(G) = d_{avg}(G_1) + d_{avg}(G_2)$, and the node-connectivity $c(G) = c(G_1) + c(G_2)$.

1.2 Fault-Tolerant Broadcasting

Broadcasting is the process of information dissemination in a communication network by which a message originated at one node (source node) is transmitted to all other nodes in the network [7]. If there exist faulty links and faulty nodes in the network, the task of fault-tolerant broadcasting is to disseminate the information from the source node (source node is supposed to be always healthy) to all the healthy nodes in the network. We say that a broadcasting succeeds if after the broadcasting procedure all the healthy nodes in the network obtain correct message held by the source node. Recently a lot of attention has been devoted to fault-tolerant broadcasting [1],[2],[6],[8]. In this paper we study fault-tolerant broadcasting in product networks.

There are usually two assumptions of fault type. One is to assume that only fail-stop faults take place, i.e., a faulty node or link does not transmit any message. It just stops the message. The other one is to assume that a faulty node or link may behave in arbitrarily harmful performance, i.e., it may not only stop a message, but also arbitrarily change the message that pass through it, and even fabricate a message. The faults we consider in this paper are of such arbitrary type.

In the study of fault-tolerant broadcasting, two situations are usually considered. One is to consider the maximum number of faults which can be tolerated in the worst case. Apparently, in this situation the maximum number of faults cannot exceed the degree (respectively half degree) of any node in the presence of fail-stop faults (respectively arbitrary faults). The other situation is that faults are randomly distributed in the network, and the relationship between the number of faults and the probability of successful broadcasting is considered.

We call a graph G an n -channel graph at node u , if there are n spanning trees of G rooted at u , $T_1, T_2, \dots, T_n$ such that for any node v of G , paths from u to v in different T_i are node disjoint. If a graph G is n -channel graph at every node u , we call G an n -channel graph. We show that an n -channel network can tolerate $\lfloor (n-1)/2 \rfloor$ arbitrarily-faulty nodes/links in broadcasting in the worst case. In this paper we do following work:

(1) We prove that the product network of any n component networks $G_1 \times G_2 \times \dots \times G_n$ is an n -channel network. Hence, it can tolerate $\lfloor (n-1)/2 \rfloor$ arbitrarily-faulty nodes/links in broadcasting in the worst case.

(2) If every component network G_i has n_i nodes and $n_i \leq b$ for some constant b , the product network $G_1 \times G_2 \times \dots \times G_n$ can tolerate $\frac{N}{4b^3nk}$ faulty nodes of arbitrary type that are randomly distributed in the network with probability larger than $1 - k^{-\lceil n/2 \rceil}$. Here N is the node number of $G_1 \times G_2 \times \dots \times G_n$.

(3) We exploit the fact that there exist n disjoint paths of length ≤ 3 between any pair of adjacent nodes in $G_1 \times G_2 \times \dots \times G_n$ and construct a reliable broadcasting, which tolerates $\Theta(L)$ arbitrarily-faulty links that are randomly distributed in the network (L is the number of the links in $G_1 \times G_2 \times \dots \times G_n$).

2 Broadcasting in Product of n Networks

Let G be a graph and v be a node of G . We call G to be n -channel at node v if there exist n spanning trees of G rooted at v , denoted by $T_1, T_2, \dots, T_n$, which satisfy the following condition:

For any node u of G , the paths $p_1(v, u)$, $p_2(v, u)$, $\dots$, $p_n(v, u)$ are node-disjoint except for v and u where $p_i(v, u)$ denotes the path from v to u in T_i , $1 \leq i \leq n$. See Figure 2.

We call G an n -channel graph if G is n -channel at every node.

Theorem 1 *If G is an n -channel network, then G can tolerate $\lfloor (n-1)/2 \rfloor$ arbitrarily-faulty nodes/links in the worst case in broadcasting.*

Proof: We suppose that every node u of the network G is a processor and has the knowledge about the topology of the network.

Let node s be the source node and $T_1, T_2, \dots, T_n$ be the n spanning trees rooted at s . For any node u , the path $p_i(s, u)$ from s to u in T_i is node-disjoint from the path $p_j(s, u)$ in T_j if $1 \leq i \neq j \leq n$. Node s holds a message m which is needed to be disseminated to all the healthy nodes in the network.




Figure 2: G is n -channel at node v if $p_i(v, u)$'s are node-disjoint for any u .

At first, s transmits the message (i, m) to all its sons in T_i , for $1 \leq i \leq n$. Then every node u in the network works concurrently in the following way:

When receiving a message (i', m') from node v , u checks whether v is the father of u in $T_{i'}$. If yes, then u saves the message (i', m') and transmits it to all its sons in $T_{i'}$. Otherwise, u does nothing.

(Note: (1) If the message received by u is not in form of (i', m') , then u does nothing. (2) If u receives messages more than one times from a same adjacent node, it only accept the message in the first time. (3) Since there may exist faults, the message (i', m') received by u is not necessary to be (i, m) , the correct message. But u regards (i', m') as correct one if it comes from the father of u in $T_{i'}$.)

After the broadcasting is completed, each node u in G obtains at most n copies of the message, each from one of $T_1, T_2, \dots, T_n$. If there are no more than $\lfloor (n-1)/2 \rfloor$ faulty nodes/links, then at least $\lceil (n+1)/2 \rceil$ paths among $p_1(s, u), p_2(s, u), \dots, p_n(s, u)$ are fault free. Hence, more than half of the copies of message m obtained by u are correct. By *majority voting*, u can pick out the correct message m . $\square$

In the proof of above Theorem, there is an implicit assumption that every node u knows that the source node is s . Actually, this assumption can be removed by modifying the broadcasting in the following way:

At first, s transmits the message (s, i, m) to all its sons in $T_i(s)$, for $1 \leq i \leq n$. When receiving a message (s', i', m') from node v , u checks whether v is the father of u in $T_{i'}(s')$. If yes, then u saves the message (s', i', m') and transmits it to all its sons in $T_{i'}(s')$. Otherwise, u does nothing. Here $T_i(v)$'s denote the spanning trees rooted at node v , and for any node w , $p_i(v, w)$'s are node-disjoint.

Now we consider the product network of n component networks. Let $G_1, G_2, \dots, G_n$ be n basic networks, i.e., each G_i is a relatively simple graph such as an array, a ring or a small complete graph etc. In general, we let each G_i be a small graph. Denote the product of $G_1, G_2, \dots, G_n$ by $P(n, G_i) = G_1 \times G_2 \times \dots \times G_n$. Each node u of $P(n, G_i)$ can be written as $u = \langle u_1, u_2, \dots, u_n \rangle$ where u_i is a node of G_i , for $1 \leq i \leq n$. In the next Theorem, we prove that $P(n, G_i)$ is an n -channel network. We only need that each G_i is connected and a spanning tree of each G_i is used.

Theorem 2 $P(n, G_i)$ is an n -channel network for any n networks $G_1, G_2, \dots, G_n$.

Proof: To prove $P(n, G_i)$ being an n -channel network, we only need to prove that $P(n, G_i)$ is n -channel at every node. Let $\langle s_1, s_2, \dots, s_n \rangle$ be a node of $P(n, G_i)$ and let BT_i be a spanning tree of G_i rooted at s_i for $i = 1, 2, \dots, n$. Then from $BT_1, BT_2, \dots, BT_n$, we can construct spanning trees $T_1, T_2, \dots, T_n$ of $P(n, G_i)$ rooted at $\langle s_1, s_2, \dots, s_n \rangle$. For $1 \leq i \leq n$, we construct T_i as following:

Let $V_1 = \{ \langle s_1, s_2, \dots, s_{i-1}, x_i, s_{i+1}, \dots, s_n \rangle \mid x_i \in G_i \}$. For any two nodes of V_1 , $\langle s_1, s_2, \dots, s_{i-1}, y_i, s_{i+1}, \dots, s_n \rangle$ and $\langle s_1, s_2, \dots, s_{i-1}, y'_i, s_{i+1}, \dots, s_n \rangle$, add a link between $\langle s_1, s_2, \dots, s_{i-1}, y_i, s_{i+1}, \dots, s_n \rangle$ and $\langle s_1, s_2, \dots, s_{i-1}, y'_i, s_{i+1}, \dots, s_n \rangle$ if and only if there is a link between y_i and y'_i in BT_i .

Let $V_2 = \{ \langle s_1, s_2, \dots, s_{i-1}, x_i, x_{i+1}, s_{i+2}, \dots, s_n \rangle \mid x_i \in G_i - \{s_i\}, x_{i+1} \in G_{i+1} \}$. For the two nodes of V_2 , $\langle s_1, s_2, \dots, s_{i-1}, x_i, y_{i+1}, s_{i+2}, \dots, s_n \rangle$ and $\langle s_1, s_2, \dots, s_{i-1}, x_i, y'_{i+1}, s_{i+2}, \dots, s_n \rangle$, add a link between $\langle s_1, s_2, \dots, s_{i-1}, x_i, y_{i+1}, s_{i+2}, \dots, s_n \rangle$ and $\langle s_1, s_2, \dots, s_{i-1}, x_i, y'_{i+1}, s_{i+2}, \dots, s_n \rangle$ if and only if there is a link between y_{i+1} and y'_{i+1} in BT_{i+1} .

Let $V_3 = \{ \langle s_1, s_2, \dots, s_{i-1}, x_i, x_{i+1}, x_{i+2}, s_{i+3}, \dots, s_n \rangle \mid x_i \in G_i - \{s_i\}, x_{i+1} \in G_{i+1}, x_{i+2} \in G_{i+2} \}$. For the two nodes of V_3 , $\langle s_1, s_2, \dots, s_{i-1}, x_i, x_{i+1}, y_{i+2}, s_{i+3}, \dots, s_n \rangle$ and $\langle s_1, s_2, \dots, s_{i-1}, x_i, x_{i+1}, y'_{i+2}, s_{i+3}, \dots, s_n \rangle$, add a link between $\langle s_1, s_2, \dots, s_{i-1}, x_i, x_{i+1}, y_{i+2}, s_{i+3}, \dots, s_n \rangle$ and $\langle s_1, s_2, \dots, s_{i-1}, x_i, x_{i+1}, y'_{i+2}, s_{i+3}, \dots, s_n \rangle$ if and only if there is a link between y_{i+2} and y'_{i+2} in BT_{i+2} .

$\vdots$

Let $V_{n-1} = \{ \langle x_1, \dots, x_{i-2}, s_{i-1}, x_i, \dots, x_n \rangle \mid x_i \in G_i - \{s_i\}, x_j \in G_j \text{ for } j = i+1, i+2, \dots, n, 1, 2, \dots, i-2 \}$. For the two nodes of V_{n-1} , $\langle x_1, \dots, x_{i-3}, y_{i-2}, s_{i-1}, x_i, \dots, x_n \rangle$ and $\langle x_1, \dots, x_{i-3}, y'_{i-2}, s_{i-1}, x_i, \dots, x_n \rangle$, add a link between $\langle x_1, \dots, x_{i-3}, y_{i-2}, s_{i-1}, x_i, \dots, x_n \rangle$ and $\langle x_1, \dots, x_{i-3}, y'_{i-2}, s_{i-1}, x_i, \dots, x_n \rangle$ if and only if there is a link between y_{i-2} and y'_{i-2} in BT_{i-2} .

Let $V_n = \{ \langle x_1, \dots, x_{i-1}, x_i, \dots, x_n \rangle \mid x_i \in G_i - \{s_i\}, x_j \in G_j, \text{ for } j \neq i \}$. For the two nodes of V_n , $\langle x_1, \dots, x_{i-2}, y_{i-1}, x_i, \dots, x_n \rangle$ and $\langle x_1, \dots, x_{i-2}, y'_{i-1}, x_i, \dots, x_n \rangle$, add a link between $\langle x_1, \dots, x_{i-2}, y_{i-1}, x_i, \dots, x_n \rangle$ and $\langle x_1, \dots, x_{i-2}, y'_{i-1}, x_i, \dots, x_n \rangle$ if and only if there is a link between y_{i-1} and y'_{i-1} .

Finally, denote the leftmost son of s_i in BT_i by t_i . Let $V_{n+1} = \{ \langle x_1, x_2, \dots, x_{i-1}, s_i, x_{i+1}, \dots, x_n \rangle \mid x_j \in G_j, j \neq i \} - \{ \langle s_1, s_2, \dots, s_n \rangle \}$. For any node of V_{n+1} , $\langle x_1, \dots, x_{i-1}, s_i, x_{i+1}, \dots, x_n \rangle$, add a link between $\langle x_1, \dots, x_{i-1}, s_i, x_{i+1}, \dots, x_n \rangle$ and $\langle x_1, \dots, x_{i-1}, t_i, x_{i+1}, \dots, x_n \rangle$.

Apparently, $V_1 - \{s\} \subset V_2 \subset \dots \subset V_n$ and $V_n \cup V_{n+1} \cup \{s\} = V(P(i, G_i))$. It is not difficult to verify that above connection makes a spanning tree T_i of $P(n, G_i)$ rooted at $\langle s_1, s_2, \dots, s_n \rangle$.

By the method of above construction, we can construct n spanning trees of $P(n, G_i)$, $T_1, T_2, \dots, T_n$ rooted at $s = \langle s_1, s_2, \dots, s_n \rangle$. For any node $u = \langle u_1, u_2, \dots, u_n \rangle$ of $P(n, G_i)$, we denote the path from $s = \langle s_1, s_2, \dots, s_n \rangle$ to $u = \langle u_1, u_2, \dots, u_n \rangle$ in T_j by $p_j(s, u)$ for $j = 1, 2, \dots, n$. Next we prove that $p_1(s, u), p_2(s, u), \dots, p_n(s, u)$ are node-disjoint. Hereafter in this paper, when we talk about the nodes of $p_j(s, u)$, we do not include s and u .

Suppose that $u_j = s_j$ for $j = i_1, i_2, \dots, i_k$, $1 \leq i_1 < i_2 < \dots < i_k \leq n$, and $u_j \neq s_j$ for $1 \leq j \neq i_1, i_2, \dots, i_k \leq n$.

(1) If $j \in \{i_1, i_2, \dots, i_k\}$, then $p_j(s, u)$ is node-disjoint with all other $p_i(s, u)$ since the j -th component of all the nodes on $p_j(s, u)$ is t_j , the leftmost son of s_j while for $i \neq j$ the j -th component of all the nodes on $p_i(s, u)$ is s_j .

(2) Now we only need to prove that for $j, j' \neq i_1, i_2, \dots, i_k$ and $j \neq j'$, $p_j(s, u)$ and $p_{j'}(s, u)$ are node-disjoint. We denote the number of different components between two nodes $w = \langle w_1, \dots, w_n \rangle$ and $w' = \langle w'_1, \dots, w'_n \rangle$ by $ND(w, w')$, i.e., $ND(w, w') =$ the number of w_i 's such that $w_i \neq w'_i$. Let v be a node on $p_j(s, u)$, v' be a node on $p_{j'}(s, u)$ and $j \neq j'$. If $ND(v, s) \neq ND(v', s)$, of course $v \neq v'$. If $ND(v, s) = ND(v', s)$, it is not difficult to see from the construction of $T_1, T_2, \dots, T_n$ that $v \neq v'$. $\square$

We suppose that each transmission of message (i, m) via a link takes a unit time, or we say, takes one step. The time needed by a broadcasting is measured as the number of concurrent steps in the broadcasting. The quantity of the broadcasting is measured as the total number of transmissions. We can not give the quantity of the broadcasting if there exist faults in the network since faults may fabricate message. But if no faults exist, the quantity of the broadcasting should be $n(N-1)$. It is the sum of the transmission numbers of n spanning trees.

If at each step, each node can transmit a message to all its adjacent nodes, the broadcasting is called all-port broadcasting. If at each step, each node can transmit a message to only one of its adjacent nodes, the broadcasting is called one-port broadcasting. Let $s = \langle s_1, s_2, \dots, s_n \rangle$ be the source node of $P(n, G_i)$ and hold a message to be disseminated to all healthy nodes in $P(n, G_i)$. Let BT_i be a spanning tree of G_i rooted at s_i for $1 \leq i \leq n$. Let T_i be the n spanning trees of $P(n, G_i)$ rooted at $s = \langle s_1, s_2, \dots, s_n \rangle$ as described in Theorem 2.

Theorem 3 gives the time needed by all-port broadcasting and the time needed by one-port broadcasting. Here we suppose that in one time unit (or one step), a message (i, m) is allowed to be transmitted via a link forth and back once, or we say the two adjacent nodes communicate once.

Theorem 3 (1) The all-port broadcasting which tolerates $\lfloor (n-1)/2 \rfloor$ faulty nodes/links needs concurrent steps not exceeding $1 + \sum_{i=1}^n a_i$, where a_i is the number of concurrent steps needed by the all-port broadcasting from s_i in G_i via BT_i .

(2) The one-port broadcasting which tolerates $\lfloor (n-1)/2 \rfloor$ faulty nodes/links needs concurrent steps not exceeding $2 \sum_{i=1}^n o_i$, where o_i is the number of concurrent steps needed by the one-port broadcasting from s_i in G_i via BT_i .

Proof: (1) The fault-tolerant broadcasting described in Theorem 1 is actually the broadcasting which consists of n concurrent broadcastings via $T_1, T_2, \dots, T_n$ respectively. Here each T_i is a spanning tree rooted at the source node s , and for any node u of $P(n, G_j)$ paths from s to u in different T_i 's are node-disjoint. If a transmission from v to w via the link (v, w) appears in T_i , then there is no transmission from v to w in T_j for $j \neq i$ (but it is possible that a transmission from w to v appears in T_j). Hence, the number of concurrent steps needed by the fault-tolerant broadcasting is equal to the maximum height of $T_1, T_2, \dots, T_n$.

The height of BT_i is a_i , for $1 \leq i \leq n$. From the construction in Theorem 2, the height of T_i is equal to $\sum_{1 \leq j \neq i \leq n} a_j + \max(a_i, 2)$. Since $a_i \geq 1$ for $1 \leq i \leq n$, the height of $T_i \leq 1 + \sum_{i=1}^n a_i$.

(2) The fault-tolerant one-port broadcasting works in the following way:

There are $2n$ rounds in the broadcasting. In round 1, o_1 steps are needed. In round 2, o_2 steps are needed. ... In round n , o_n steps are needed. In round $n+1$, o_1 steps are needed. ... In round $2n$, o_n steps are needed.

Let $1 \leq r \leq n$, $1 \leq j \leq o_r$ and $u = \langle u_1, \dots, u_r, \dots, u_n \rangle$ be a node of $P(n, G_i)$. Let u transmit a message to $u' = \langle u_1, \dots, u_{r-1}, u'_r, u_{r+1}, \dots, u_n \rangle$ at the j -th step in round r or round $n+r$ if and only if the following two conditions are satisfied.

(i) In the one-port broadcasting of G_r from s_r via BT_r , the j -th step is the transmission from u_r to u'_r .

(ii) The transmission from u to u' is legal in the broadcasting specified in Theorem 1, i.e., the message (k, m) that u want to transmit to u' is from the father of u in T_k and u' is one of the sons of u in T_k .

There can be at most one legal transmission from u to u' at the j -th step in round r and round $n+r$.

From round 1 to round $n+1$, the sub-broadcasting via T_1 is completed. From round 2 to round $n+2$, the sub-broadcasting via T_2 is completed. ...

From round n to round $2n$, the sub-broadcasting via T_n is completed. The whole broadcasting needs $2n$ rounds, totally $2 \sum_{i=1}^n o_i$ steps. (Note: the sub-broadcastings via T_i 's do not conflict over any link since the paths from any node u in different T_i 's are node-disjoint. $\square$)

3 Broadcasting with Random Faults

In the above Section, we proved that the product network $P(n, G_i)$ can tolerate $\lfloor (n-1)/2 \rfloor$ faults in the worst case. However, in the reality the worst case appear with very small probability. The more reasonable assumption is that faults are randomly distributed in the network. In this case the broadcasting succeeds with high probability even much more than $\lfloor (n-1)/2 \rfloor$ faults take place.

In this Section, we suppose that faulty nodes are randomly distributed in the network. We study the relation between the number of faulty nodes and the probability with which the broadcasting succeeds.

Suppose that there are f faulty nodes randomly distributed in the network, i.e., we suppose that each configuration of the network with f faulty nodes is equally probable. We denote a configuration of network G with f faulty nodes by C_G^f , the set of all the configurations with f faulty node by $\mathbf{C}_G^f$. For any $C_G^f \in \mathbf{C}_G^f$, if the broadcasting succeeds in presence of C_G^f , then we say C_G^f is a successful configuration. Otherwise C_G^f is called a failed configuration. The probability of successful broadcasting in presence of f random faulty nodes is measured as the ratio of the number of successful C_G^f 's to $|\mathbf{C}_G^f|$. We denote the set of all the successful configurations by $\mathbf{SC}_G^f$ and the set of all the failed configurations by $\mathbf{FC}_G^f$. Hence, the probability of successful broadcasting = $|\mathbf{SC}_G^f|/|\mathbf{C}_G^f| = 1 - |\mathbf{FC}_G^f|/|\mathbf{C}_G^f|$.

Now we consider the product network $P(n, G_i)$. We require each G_i be a small graph. Suppose that there is a bound to the node numbers of all the G_i 's, i.e., $|G_i| \leq b$ for some constant b , $1 \leq i \leq n$.

Let u be a node of $P(n, G_i)$. There are n node-disjoint paths $p_1(s, u)$, $p_2(s, u)$, ..., $p_n(s, u)$ from s to u . The message is disseminated from s to u through these n paths in the broadcasting. It is easy to see from the construction of Theorem 2 that there are less than nb nodes on each $p_i(s, u)$ for any $1 \leq i \leq n$. If in a configuration $C_{P(n, G_i)}^f$ more than $\lfloor (n-1)/2 \rfloor$ paths among $p_1(s, u)$, $p_2(s, u)$, ..., $p_n(s, u)$ have faulty nodes, then $C_{P(n, G_i)}^f$ is said to be a failed configuration on u . Denote the set of all the failed configuration on u by $\mathbf{FC}_{P(n, G_i)}^f(u)$. We have

$$|\mathbf{FC}_G^f| < \sum_{v \in P(n, G_i)} |\mathbf{FC}_{P(n, G_i)}^f(v)|.$$

Theorem 4 If there are f faulty nodes randomly distributed in $P(n, G_i)$, the broadcasting succeeds with a probability higher than $1 - \left(\frac{4b^3nf}{N}\right)^{\lfloor n/2 \rfloor}$, where N is the node number of $P(n, G_i)$.

Proof: The probability of successful broadcasting in $P(n, G_i)$ is

$$\frac{|\mathbf{SC}_{P(n, G_i)}^f|}{|\mathbf{C}_{P(n, G_i)}^f|} = 1 - \frac{|\mathbf{FC}_{P(n, G_i)}^f|}{|\mathbf{C}_{P(n, G_i)}^f|} > 1 - \frac{\sum_{v \in P(n, G_i)} |\mathbf{FC}_{P(n, G_i)}^f(v)|}{|\mathbf{C}_{P(n, G_i)}^f|}$$

For any node v of $P(n, G_i)$, there are n node-disjoint paths $p_1(s, v), p_2(s, v), \dots, p_n(s, v)$ of length $< bn$. By repeatedly counting, we have

$$|\mathbf{FC}_{P(n, G_i)}^f(v)| < (bn)^{\lceil n/2 \rceil} \binom{n}{\lceil n/2 \rceil} \binom{N - \lceil n/2 \rceil}{f - \lceil n/2 \rceil}$$

Hence

$$\begin{aligned} \frac{|\mathbf{FC}_{P(n, G_i)}^f|}{|\mathbf{C}_{P(n, G_i)}^f|} &< \frac{N(bn)^{\lceil n/2 \rceil} \binom{n}{\lceil n/2 \rceil} \binom{N - \lceil n/2 \rceil}{f - \lceil n/2 \rceil}}{\binom{N}{f}} \\ &= N(bn)^{\lceil n/2 \rceil} \binom{n}{\lceil n/2 \rceil} \frac{f(f-1) \cdots (f+1 - \lceil n/2 \rceil)}{N(N-1) \cdots (N+1 - \lceil n/2 \rceil)} \\ &< b^n (bn)^{\lceil n/2 \rceil} 2^n \left(\frac{f}{N}\right)^{\lceil n/2 \rceil} \leq \left(\frac{4b^3 nf}{N}\right)^{\lceil n/2 \rceil} \end{aligned}$$

Hence, we obtain

$$\frac{|\mathbf{SC}_{P(n, G_i)}^f|}{|\mathbf{C}_{P(n, G_i)}^f|} > 1 - \left(\frac{4b^3 nf}{N}\right)^{\lceil n/2 \rceil}$$

□

Corollary 1 For any $k > 1$, if there are $\frac{N}{4b^3 nk}$ faulty nodes randomly distributed in $P(n, G_i)$, the broadcasting succeeds with a probability higher than $1 - k^{-\lceil n/2 \rceil}$.

In reference [1], a broadcasting in face of randomly distributed faults is said to be ϵ -safe if the probability with which the broadcasting succeeds is higher than $1 - N^{-\epsilon}$. From Corollary 1, our broadcasting is ϵ -safe if there are $\frac{N}{4b^3 nk}$ faulty nodes randomly distributed in $P(n, G_i)$ for any $k > 1$, where ϵ is dependent on both k and b .

Similarly, we can also suppose that there are f faulty links randomly distributed in the network $P(n, G_i)$. By the same method, we have the following Theorem.

Theorem 5 If there are f faulty links randomly distributed in $P(n, G_i)$, the broadcasting succeeds with a probability higher than $1 - \left(\frac{4b^3 nf}{L}\right)^{\lceil n/2 \rceil}$, where L is the link number of $P(n, G_i)$.

4 Tolerate $\Theta(L)$ Faulty Links

In Theorem 5, we consider the situation where f faulty links are randomly distributed in the network $P(n, G_i)$ while all the nodes are healthy. By Theorem 5, the broadcasting succeeds with a probability higher than $1 - k^{-\lceil n/2 \rceil}$ if $f = \Theta(L/kn)$. Here the coefficient of Θ is a constant dependent on b . Next we show that in this situation we can modify the broadcasting such that the broadcasting succeeds with a probability higher than $1 - k^{-\lceil n/2 \rceil}$ if $f = \Theta(L/k)$.

The network $P(n, G_j)$ has a spanning tree T . For example we let $T = T_1$. For any pair of adjacent nodes of $P(n, G_j)$ there exist at least n paths of length ≤ 3 between them. Let $u = \langle u_1, u_2, \dots, u_n \rangle$ and $u' = \langle u_1, \dots, u_{i-1}, u'_i, u_{i+1}, \dots, u_n \rangle$ where u_i is adjacent with u'_i in BT_i . The n disjoint paths between u and u' are described as follows:

For $k = 1, 2, \dots, n$, we use $p_k(u, u')$ to denote the k -th path from u to u' . (But this time the meaning of $p_k(u, u')$ is different from the above.)

If $k = i$, $p_k(u, u')$ is the link between u and u' .

If $k \neq i$, $p_k(u, u')$ is the path

$$\begin{aligned} u &\longleftrightarrow \\ &< u_1, \dots, u_{k-1}, t_k, u_{k+1}, \dots, u_{i-1}, u_i, u_{i+1}, \dots, u_n \rangle \longleftrightarrow \\ &< u_1, \dots, u_{k-1}, t_k, u_{k+1}, \dots, u_{i-1}, u'_i, u_{i+1}, \dots, u_n \rangle \longleftrightarrow \\ &< u_1, \dots, u_{k-1}, u_k, u_{k+1}, \dots, u_{i-1}, u'_i, u_{i+1}, \dots, u_n \rangle = u' \end{aligned}$$

Here t_k is any neighbor of u_k in G_k . Actually, It is not difficult to see that there exist $1 + \sum_{1 \leq j \neq i \leq n} \text{degree}(u_j)$ node-disjoint paths of length ≤ 3 between $u = \langle u_1, u_2, \dots, u_n \rangle$ and $u' = \langle$

$u_1, \dots, u_{i-1}, u'_i, u_{i+1}, \dots, u_n >$. In this section, we only consider the results obtained by exploiting n node-disjoint paths of length ≤ 3 between any pair of adjacent nodes. It is easy to generalize our results to the situation where $1 + \sum_{1 \leq j \neq i \leq \text{degree}(u_j)}$ node-disjoint paths of length ≤ 3 are exploited.

The modified broadcasting is simple: Consider the broadcasting on $P(n, G_j)$ via the spanning tree T . We replace each transmission in the broadcasting via T by n transmissions via the n paths of length ≤ 3 . We describe it more formally as in follows.




Figure 3: Each link of the spanning tree is replaced by n disjoint paths of length ≤ 3 .

Let u and u' be adjacent in T . In the broadcasting via T , the message is transmitted from u to u' . In the modified broadcasting, the message is transmitted from u to u' through $p_1(u, u')$, $p_2(u, u')$, ..., $p_n(u, u')$.

The modified broadcasting needs about $3n$ times of transmissions as the broadcasting via T needs. Hence, The modified broadcasting needs about 3 times of transmissions as needed by the broadcasting described in Section 2. The time (concurrent steps) needed by the modified broadcasting does not exceed $3n$ times as needed by the the broadcasting in Section 2.

The modified broadcasting behaves much better in tolerating randomly distributed faulty links although it also tolerates at most $\lfloor (n-1)/2 \rfloor$ faulty links in the worst case.

Theorem 6 *If there are f faulty links randomly distributed in $P(n, G_i)$ and $|G_i| \leq b$ for any $1 \leq i \leq n$, the modified broadcasting succeeds with a probability higher than $1 - (12b^2 f/L)^{\lceil n/2 \rceil}$, where L denotes the link number of $P(n, G_i)$.*

Proof: We denoted a configuration of f faulty links in $P(n, G_i)$ by $C_{P(n, G_i)}^f$ and the set of all the $C_{P(n, G_i)}^f$ by $\mathbf{C}_{P(n, G_i)}^f$. Similarly as before, we denote the set of all the $C_{P(n, G_i)}^f$ which makes the modified broadcasting succeed by $\mathbf{SC}_{P(n, G_i)}^f$, and the set of all the $C_{P(n, G_i)}^f$ which makes the modified broadcasting fail by $\mathbf{FC}_{P(n, G_i)}^f$. Let (u, v) be a link of T . Denote the set of all the $C_{P(n, G_i)}^f$ in which at least $\lceil n/2 \rceil$ paths among $p_1(u, v)$, $p_2(u, v)$, ..., $p_n(u, v)$ are faulty by $\mathbf{FC}_{P(n, G_i)}^f(u, v)$. Apparently,

$$|\mathbf{FC}_{P(n, G_i)}^f| < \sum_{(u, v) \in T} |\mathbf{FC}_{P(n, G_i)}^f(u, v)|$$

Since there are at most 3 links along each $p_i(u, v)$, still by repeatedly counting, we have

$$|\mathbf{FC}_{P(n, G_i)}^f(u, v)| < \binom{n}{\lceil n/2 \rceil} 3^{\lceil n/2 \rceil} \binom{L - \lceil n/2 \rceil}{f - \lceil n/2 \rceil}$$

Hence

$$\begin{aligned} \frac{|\mathbf{FC}_{P(n, G_i)}^f|}{|\mathbf{C}_{P(n, G_i)}^f|} &< N \binom{n}{\lceil n/2 \rceil} 3^{\lceil n/2 \rceil} \frac{\binom{L - \lceil n/2 \rceil}{f - \lceil n/2 \rceil}}{\binom{L}{f}} \\ &= N \binom{n}{\lceil n/2 \rceil} 3^{\lceil n/2 \rceil} \frac{f(f-1) \cdots (f+1 - \lceil n/2 \rceil)}{L(L-1) \cdots (L+1 - \lceil n/2 \rceil)} \\ &< b^n 2^n 3^{\lceil n/2 \rceil} \left(\frac{f}{L} \right)^{\lceil n/2 \rceil} \leq \left(\frac{12b^2 f}{L} \right)^{\lceil n/2 \rceil} \end{aligned}$$

Hence, we have

$$\frac{|\mathbf{SC}_{P(n,G_i)}^f|}{|\mathbf{C}_{P(n,G_i)}^f|} = 1 - \frac{|\mathbf{FC}_{P(n,G_i)}^f|}{|\mathbf{C}_{P(n,G_i)}^f|} > 1 - \left(\frac{12b^2f}{L}\right)^{\lceil n/2 \rceil}$$

□

Corollary 2 *The modified broadcasting succeeds with a probability higher than $1 - k^{-\lceil n/2 \rceil}$ with $\Theta(L)$ faulty links randomly distributed in $P(n, G_i)$, where the coefficient of Θ is dependent on k and b .*

Corollary 2 shows that the modified broadcasting is ϵ -safe even if there are cL faulty links randomly distributed in the network, where c is a constant dependent on both ϵ and k .

5 Conclusions

In this paper we study the reliable broadcasting in product networks. We prove that an n -product network is an n -channel graph by constructing n spanning trees $T_1, T_2, \dots, T_n$. The reliable broadcasting is naturally based on these n spanning trees. It can tolerate $\lfloor (n-1)/2 \rfloor$ faulty nodes/links of arbitrary type in the worst case. The relation between the number of randomly distributed faults and the probability with which the broadcasting succeeds is analyzed. Actually, what we give in this paper is a lower bound of the successful probability. The results are obtained under the assumption $|G_i| \leq b$. However, if the concrete number of nodes for each G_i is given, our method can be applied to deriving a tighter bound. For the situation where only faulty links exist, we give another broadcasting which can tolerate $\Theta(L)$ randomly distributed faulty links with high probability.

References

- [1] D. Bienstock, Broadcasting with random faults, Discrete Applied Mathematics, Vol. 20, pp.1-7, 1988.
- [2] D.B. Blough, A. Pelc, Optimal communication in networks with randomly distributed faults, Network, Vol. 23, pp.691-701, 1993.
- [3] G. Chartrand and L. Lesniak, Graphs and Digraphs (second Edition), Wadsworth & Brooks, 1986.
- [4] S.K. Das, S. Ohring, and A.K. Banerjee, Embeddings into hyper Petersen networks: Yet another hypercube-like interconnection topology, Journal of VLSI Design, Special Issue on Interconnection networks, 1994.
- [5] E. Ganesan and D.K. Pradhan, The hyper-debrujin networks: Scalable versatile architecture, IEEE Trans. on Parallel and Distributed Systems, 4(9), pp. 962-978, 1993.
- [6] L. Gargano, A. Rescigno, U. Vaccaro, Fault-tolerant hypercube broadcasting via information dispersal, Networks, Vol. 23, pp.271-282, 1993.
- [7] S.M. Hedetniemi, S.T. Hedetniemi, A.L. Liestman, A survey of gossiping and broadcasting in communication networks, Network Vol.18, pp.1240-1268, 1988.
- [8] P. Ramanathan, K.G. Shin, Reliable broadcasting in hypercube multiprocessors, IEEE Trans. on Computers, Vol. 37, pp.1654-1657, 1988.
- [9] A. Youssef, Cartesian product networks, In Proceedings of the 1991 International Conference on Parallel Processing, Vol. I, pp. 684-685, 1991.