ELSEVIER

Discrete Applied Mathematics 83 (1998) 3–20

DISCRETE APPLIED MATHEMATICS

Reliable broadcasting in product networks

Feng Bao [a], Yoshihide Igarashi [a,*], Sabine R. Öhring [b]

[a] *Department of Computer Science, Gunma University, Kiryu, 376 Japan*

[b] *Department of Computer Science, University of North Texas, Denton, TX 76203-3886, USA*

Received 26 May 1995; received in revised form 15 April 1996; accepted 24 June 1996

Abstract

We investigate the reliability of broadcasting in product networks containing faulty nodes and/or links. Faults considered in this paper are mainly of the Byzantine type, i.e., a faulty node or a faulty link may not only stop sending a message but also arbitrarily change a message passing through the faulty place or even fabricate a false message. We assume that no nodes have a priori information about faults in a network. Hence, the key problem of reliable broadcasting in our model is how to control the message transmission so that any corrupted message cannot affect the result of the broadcasting too much. We propose the concept of an n-channel graph which has n-independent spanning trees rooted at each node. The fault tolerance can be achieved by sending n copies of the message along the n-independent spanning trees rooted at the source node. In this paper we show how to construct n-independent spanning trees of a product network from spanning trees of n-component graphs. Furthermore, we can design an efficient and reliable broadcasting scheme based on independent spanning trees for a product network from simple broadcasting schemes for component networks. The degrees of fault tolerance against crash faults and Byzantine faults of nodes and/or links are, respectively, $n-1$ and $\lfloor (n-1)/2 \rfloor$ in the worst case. We can successfully broadcast with a probability higher than $1-k^{-\lceil n/2 \rceil}$ in any product network of order N consisting of n-component graphs of order b or less, if at most $N/4b^3nk$ faulty nodes are randomly distributed in the network. We can also successfully broadcast with a probability higher than $1-k^{-\lceil n/2 \rceil}$ in any product network of size L, of n component graphs of size b or less, if at most $L/12b^2k$ faulty links are randomly distributed in the network.

Keywords: Broadcast; Byzantine faults; Channel graphs; Fault tolerance; Independent spanning trees; Product networks

1. Introduction

A processor network is expressed as a graph, where a node is a processor and an edge is a communication link. Broadcasting is the process of sending a message from the source node to all other nodes in a network. It can be accomplished by message dissemination in such a way that each node repeatedly receives and forwards messages.

* Corresponding author. Fax: 81 277 30 1801; e-mail: igarashi@comp.cs.gunma-u.ac.jp.

PII S0166-218X(97)00100-5

Some of the nodes and/or links may be faulty. However, multiple copies of messages can be disseminated through disjoint paths. We say that the broadcasting succeeds if all the healthy nodes in the network finally obtain the correct message from the source node within a certain limit of time. Recently, a lot of attention has been devoted to fault-tolerant broadcasting in networks [1–4, 6–9, 11–17]. The reader may find more related references on this topic in good survey papers by Hedetniemi et al. [9] and by Pelc [14]. A paper by Hadzilacos and Toueg [8] is also a good survey on fault-tolerant broadcasts, but it emphasizes various types of failures including timing failures and the consensus problem.

Two types of faults are usually supposed in the study of broadcasting in networks. One is the crash type, and the other is the Byzantine type. For the former type, a faulty node or link does not transmit any message (i.e., it stops sending any message). For the latter type, a faulty node or link may not only stop sending messages but also arbitrarily change messages passing through the faulty place, or even fabricate malicious messages. Fault tolerance of broadcasting against faults of the crash type has been much studied. However, fault-tolerant broadcasting against Byzantine faults has not been much studied. So far we know the following work related to this subject. In [1], Bagchi and Hakimi studied reliable broadcasting in a network with faulty nodes of the Byzantine type where each healthy node knows local information about faults, i.e., it knows which neighboring nodes are faulty. In [13], Pelc studied all-port broadcasting in a network with faulty links of the Byzantine type. He proposed a broadcasting scheme such that each healthy node sends a message after confirming it by majority voting. In [4], Blough and Pelc discussed broadcasting in complete graphs, core graphs and fat-ring graphs with faults of the Byzantine type, where no nodes know any information about faults. Their broadcasting scheme is based on the principle of *vote-and-send* in networks with the property that each node has a large number of neighboring nodes. In [2], Bao et al. discussed the reliability and efficiency of one-port broadcasting schemes in hypercubes with faulty nodes and links of the Byzantine type under the assumption that no nodes have a priori information about faults.

In this paper we study fault-tolerant broadcasting against faulty nodes and/or links of the Byzantine type in product networks. We assume that the source node is always healthy and that no nodes know a priori information about faults. Hence, no information about the locations of faults can be used in the broadcasting. We also assume that each node in a network has a priori information about the network topology. For communication efficiency we should avoid unnecessary message transmissions, although some multiple copies of the message are necessary to tolerate faulty nodes and links. For this purpose we need a good strategy to control message transmissions so that every transmission contributes to the majority voting to obtain the correct message. Our fault-tolerant broadcasting scheme is non-adaptive in the sense that for each destination node, fixed internally node-disjoint paths from the source node to the destination node are used as its transmission channels no matter where faults exist. It does not require fault detection. The idea of such a kind of broadcasting scheme may be originally by Ramanathan and Shin [16]. Cristian et al. [6] discussed authentication-detectable

Byzantine faults that can be detected by error-correcting techniques or cryptographic protocols such as digital signatures. However, Byzantine faults considered in this paper belong to non-authentication-detectable Byzantine faults.

The paper consists of five more sections. In Section 2, we introduce the concept of an n-channel graph, i.e., a graph that has n independent spanning trees rooted at each node. For an n-channel graph, we propose a general broadcasting scheme that sends messages along n independent spanning trees rooted at the source node. Fault tolerance can be achieved since we send n copies of the message along different channels (i.e., independent spanning trees), so that a message corrupted at a place can spoil only one channel for each destination node. The degree of fault tolerance means the number of faults that can be tolerated. For the broadcasting scheme, the degrees of tolerance to crash and Byzantine node (and/or link) faults are, respectively, $n-1$ and $\lfloor (n-1)/2 \rfloor$. It was proved in [10] by Itai and Rodeh that a graph is 2-channel if and only if it is 2-connected. It was also proved that any 3-connected graph is a 3-channel graph by Cheriyan and Maheshari [5], and by Zehavi and Itai [19]. However, it remains open whether any k-connected graph is always a k-channel graph for any $k>3$.

In Section 3, we show that a product graph of n-component graphs is an n-channel graph if each component graph is connected. We give a construction of n independent spanning trees of such a product network. Moreover, we derive a one-port broadcasting scheme that uses n independent spanning trees, and analyze its efficiency. The efficiency here stands for both the running time and the number of transmissions used by the broadcasting. The broadcasting discussed in this paper is optimal in the number of transmissions while its running time depends on the structure of the n independent spanning trees. In Section 4, we discuss the probability of successful broadcasting under the assumption that faults are randomly distributed in a network. We can successfully broadcast with a probability higher than $1-k^{-\lceil n/2 \rceil}$ in any product network of order N consisting of n-component graphs of order b or less, if at most $N/(4b^3nk)$ faulty nodes are randomly distributed in the network. Here the order (size) of a graph means the number of its nodes (links).

In Section 5, we propose a broadcasting scheme for a product network with only faulty links of the Byzantine type, where we use the same principle *vote-and-send* as in [13]. The broadcasting in any product network of size L, consisting of n component graphs of order b or less succeeds with a probability higher than $1-k^{-\lceil n/2 \rceil}$, if at most $L/12b^2k$ faulty links are randomly distributed in the network. This means that broadcasting in any product network of size L with $O(L)$ faulty links randomly distributed, consisting of n-component graphs, succeeds with a high probability, if b is a constant independent of L and n. Conclusions are given in Section 6.

2. Broadcasting in n-channel graphs

We first introduce the definitions of independent spanning trees and n-channel graphs.

Definition 1. Two spanning trees of a graph $G=(V,E)$ are called independent if they are rooted at the same node r, and for each node $v \in V$, the two paths from r to v, one path in each spanning tree, are internally node-disjoint. The n spanning trees of G are said to be independent if they are pairwise independent.

Definition 2. A graph G is called an n-channel graph at node r if there are n independent spanning trees rooted at r of G. If G is an n-channel graph at each and every node, G is said to be an n-channel graph.

Itai and Rodeh [10] used independent spanning trees as reliable communication channels in networks, and they showed a linear-time algorithm for finding two independent spanning trees in a biconnected graph. However, it is open whether for any $k>3$, any k-connected graph is a k-channel graph [5, 19].

Theorem 1. *For broadcasting in any n-channel graph, the degrees of tolerance to crash and Byzantine node (and/or link) faults are $n-1$ and $\lfloor (n-1)/2 \rfloor$, respectively.*

Proof. For the case of faults of the crash type the assertion of the theorem is immediate.

We consider faults of the Byzantine type and a fault-tolerant broadcasting scheme sending messages through n independent spanning trees of an n-channel graph G. For simplicity, we first assume that each node in a network knows the source node. However, this restriction can be removed as stated later in this proof. Let s be the source node of G, and let $T_1(s), T_2(s), \ldots, T_n(s)$ be n independent spanning trees rooted at s of G. We assume that every node of G knows its parent and sons in $T_i(s)$ for each i $(1 \leqslant i \leqslant n)$. Let m be the original message in the source node s. For each i $(1 \leqslant i \leqslant n)$, s sends message (i,m) to the sons of s in $T_i(s)$. Then each node u of G works concurrently in the following fashion:

> When u receives a message $(\bar{i},\bar{m})$ from an adjacent node v, u checks whether v is the father of u in $T_{\bar{i}}(s)$. If yes, then u saves $(\bar{i},\bar{m})$ and transmits $(\bar{i},\bar{m})$ to all its sons in $T_{\bar{i}}(s)$; otherwise, u does nothing.

Note that $\bar{m}$ may not be m since a message may be altered at a faulty place. If the message received by u is not in the form of $(\bar{i},\bar{m})$, then u ignores it. If u receives messages more than once from the same node, it accepts only the message that arrived first. The situation of this process is depicted in Fig. 1.

When the broadcasting is completed, each node u of G obtains at most n copies of the message m, each of which comes through one of the n-independent spanning trees $T_1(s), T_2(s), \ldots, T_n(s)$. Although any faulty node can affect all the n-independent spanning trees, for each node u the faulty node can affect just one of the n paths from s to u, each path in its corresponding independent spanning tree. Hence, if at most

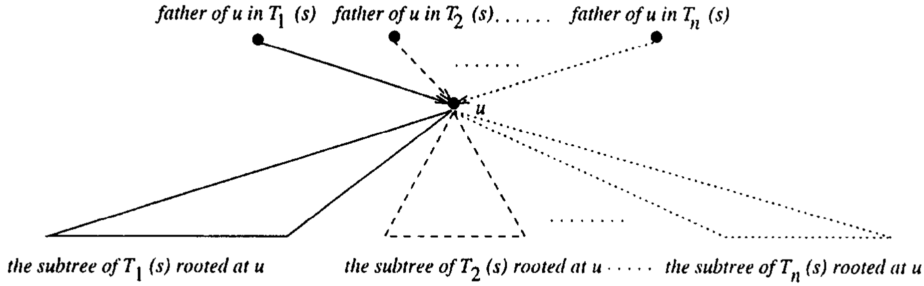

Fig. 1. Message transmissions through u.

$\lfloor (n-1)/2 \rfloor$ faulty nodes and/or links exist, then every node u of G can obtain the correct message m by majority voting.

In the discussion above, we assumed that every node u of G knows s to be the source node. However, this assumption is too restrictive, and it can be removed by slightly modifying the broadcasting procedure as follows. Suppose that node s wants to broadcast a message m in the network where other nodes do not know that s is the source node. Let G be an n-channel graph, and let $T_1(r), T_2(r), \ldots, T_n(r)$ be n independent spanning trees rooted at node r. Every node u of G knows $T_1(r), T_2(r), \ldots, T_n(r)$ for each node r since we assume that every node has a priori information about the network topology. In the modified version of the broadcasting, for each $i (1 \leqslant i \leqslant n)$, s sends message (s, i, m) to its sons in $T_i(s)$. Then every node u of G works concurrently in the following fashion:

> When u receives a message $(\bar{s}, \bar{i}, \bar{m})$ from an adjacent node v, u checks whether v is the father of u in $T_{\bar{i}}(\bar{s})$. If yes, and u did not receive any message $(\bar{s}, \bar{i}, \bar{m})$ from v before, then u saves $(\bar{s}, \bar{i}, \bar{m})$ and transmits $(\bar{s}, \bar{i}, \bar{m})$ to all the sons of u in $T_{\bar{i}}(\bar{s})$. Otherwise, u does nothing.

When the broadcasting is completed, for some node x of G, each node u of G may obtain at most n messages which are claimed to be originally from x, namely, $(x, 1, m_1), (x, 2, m_2), \ldots, (x, n, m_n)$. If the message through the jth independent spanning tree did not reach u, or it was rejected by u because of its illegal format (e.g., (x, j, m_j) is rejected by u if it reached u through an edge not in $T_j(x)$), we interpret m_j as an empty value. Then u chooses the most common non-empty value of $\{m_1, m_2, \ldots, m_n\}$. If at least $\lceil (n+1)/2 \rceil$ non-empty common values exist in the set, u decides the most common value to be the message from the source node x. If at most $\lfloor (n-1)/2 \rfloor$ faults exist in the network and only s broadcasts message m, then every node u knows correctly that the message disseminated from s is m by majority voting. Note that if a malicious node x pretends to be the source node and sends a message through n-independent spanning trees rooted at x, in general we cannot remove such a message from the malicious node x. However, even in such a case every healthy node can obtain the correct message from the correct source node if the number of faults is at most $\lfloor (n-1)/2 \rfloor$. □

In order to tolerate $\lfloor (n-1)/2 \rfloor$ faults of the Byzantine type in G under the assumption that no nodes have a priori information about faults, each node should receive at least n copies of the message originally from the source node. Hence, for this purpose at least $(|V(G)| - 1)n$ transmissions are required, where $|V(G)|$ denotes the order of G (i.e., the number of nodes of G). As described in the proof of Theorem 1, the broadcasting based on n-independent spanning trees uses exactly $(|V(G)| - 1)n$ transmissions if no faults exist. We therefore can say that the broadcasting through the n-independent spanning trees in this fashion is optimal for communication complexity. It is hard to estimate the number of transmissions when faults of the Byzantine type exist. The running time of the broadcasting depends on various factors of the network. These factors are the structures of the n-independent spanning trees, communication mode (one-port or all-port), timing of transmissions, and so on.

3. Broadcasting in product networks

For a pair of graphs $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$, the product of G_1 and G_2, denoted by $G = G_1 \times G_2$, is a graph with the node set $V_1 \times V_2 = \{(x, y) \mid x \in V_1, y \in V_2\}$ and the edge set such that two nodes (u_1, u_2) and (v_1, v_2) are adjacent in $G_1 \times G_2$ if and only if either $u_1 = v_1$ and $(u_2, v_2) \in E_2$, or $u_2 = v_2$ and $(u_1, v_1) \in E_1$. The definition of the product of two graphs can be generalized to the product of n graphs in the natural way. That is, $G_1 \times G_2 \times \cdots \times G_n$ means $(G_1 \times \cdots \times G_k) \times (G_{k+1} \times \cdots \times G_n)$ for some k $(1 \leqslant k \leqslant n-1)$. Each G_i $(1 \leqslant i \leqslant n)$ is called a component graph of $G_1 \times G_2 \times \cdots \times G_n$. The product of graphs is also called a product network. To avoid trivial cases, throughout this paper we assume that each component graph is connected and has at least two nodes. An edge in a network is often called a link in this paper.

Some popular interconnection networks are product networks. For example, the n-dimensional hypercube (n-cube for short) Q_n is $Q_{n-1} \times K_2 = Q_{n-2} \times K_2 \times K_2 = \cdots = K_2 \times K_2 \times \cdots \times K_2$, and an n-dimensional generalized hypercube Q_n^t is $Q_{n-1}^t \times K_t = Q_{n-2}^t \times K_t \times K_t = K_t \times K_t \times \cdots \times K_t$, where K_t is the complete graph of order t. The $(m_1 \times \cdots \times m_n)$-mesh is $L_{m_1} \times \cdots \times L_{m_n}$, and the $(m_1 \times \cdots \times m_n)$-torus is $R_{m_1} \times \cdots \times R_{m_n}$, where L_{m_i} and R_{m_i} are a linearly linked graph of order m_i and a ring of order m_i, respectively.

Let $d(G)$, $deg(G)$, $d_{\mathrm{avg}}(G)$ and $c(G)$ denote the diameter, the degree, the average distance and the node connectivity of G, respectively. Youssef [18] proved that for a pair of graphs G_1 and G_2, $d(G_1 \times G_2) = d(G_1) + d(G_2)$, $deg(G_1 \times G_2) = deg(G_1) + deg(g_2)$, $d_{\mathrm{avg}}(G_1 \times G_2) = d_{\mathrm{avg}}(G_1) + d_{\mathrm{avg}}(G_2)$, and $c(G_1 \times G_2) = c(G_1) + c(G_2)$.

In this section, we construct n-independent spanning trees in a product network of n-component graphs, $G = G_1 \times G_2 \times \cdots \times G_n$. We use just one spanning tree of each G_i, i.e., we only require that each G_i is connected. The construction of independent spanning trees given in this section is suitable for a product network such that each component graph cannot be expressed as a product of graphs and has a small connectivity. For example, we construct n-independent spanning trees of the n-cube Q_n

from nK_2's (i.e., we use the structure $Q_n = K_2 \times K_2 \cdots \times K_2$ instead of the structure $Q_n = Q_i \times Q_{n-i}$). Note that from the structure $Q_n = Q_i \times Q_{n-i}$ we can obtain only two independent spanning trees of Q_n by our construction. The construction given in this section is relatively simple, and the proof of its correctness is clear. The basic idea of this construction can be generalized to a construction using more spanning trees of each component graph. However, that construction is complicated and will not be discussed in this paper.

Theorem 2. *For any node $s = (s_1, s_2, \ldots, s_n)$ of product $G = G_1 \times G_2 \times \cdots \times G_n$, there exist n-independent spanning trees of G rooted at s, where G_i is connected and s_i is a node of G_i for each i $(1 \leqslant i \leqslant n)$.*

Proof. Let T_j^s be a spanning tree of G_j rooted at s_j for each j $(1 \leqslant j \leqslant n)$. From these spanning trees, T_j^s $(1 \leqslant j \leqslant n)$, we construct n independent spanning trees $T_1, \ldots, T_n$ rooted at $(s_1, \ldots, s_n)$ of G.

For each i $(1 \leqslant i \leqslant n)$ we construct T_i as follows: We first develop the ith component along T_i^s, then develop the $(i+1)$th component along T_{i+1}^s, the $(i+2)$th component along $T_{i+2}^s, \ldots,$ the nth component along T_n^s, the first component along $T_1^s, \ldots,$ and finally the $(i-1)$th component along T_{i-1}^s (if $i = 1$, the final development is along T_n^s). Each node in $\{(x_1, \ldots, x_{i-1}, s_i, x_{i+1}, \ldots, x_n) \mid x_j \in G_j \text{ for } j = 1, 2, \ldots, i-1, i+1, \ldots, n\}$ is connected to the node obtained by replacing the ith component s_i by one of its sons in T_i^s, say t_i. Each of these edges is used as a link to reach a node with s_i in the ith component. The construction of T_i is more formally described as follows:

(1) Let $V_1 = \{(s_1, \ldots, s_{i-1}, x_i, s_{i+1}, \ldots, s_n) \mid x_i \in V(G_i)\}$. Each pair of nodes in V_1, $(s_1, \ldots, s_{i-1}, y_i, s_{i+1}, \ldots, s_n)$ and $(s_1, \ldots, s_{i-1}, y_i', s_{i+1}, \ldots, s_n)$ are connected by an edge if and only if y_i and y_i' are adjacent in T_i^s.

(2) For each k $(1 \leqslant k \leqslant n-i)$, let $V_{1+k} = \{(s_1, \ldots, s_{i-1}, x_i, \ldots, x_{i+k}, s_{i+k+1}, \ldots, s_n) \mid x_i \in V(G_i) - \{s_i\}$, and $x_j \in V(G_j)$ for $j = i+1, i+2, \ldots, i+k\}$. Each pair of nodes in V_{1+k}, $(s_1, \ldots, s_{i-1}, x_i, \ldots, x_{i+k-1}, y_{i+k}, s_{i+k+1}, \ldots, s_n)$ and $(s_1, \ldots, s_{i-1}, x_i, \ldots, x_{i+k-1}, y_{i+k}', s_{i+k+1}, \ldots, s_n)$ are connected by an edge if and only if y_{i+k} and y_{i+k}' are adjacent in T_{i+k}^s.

(3) For each k $(n - i + 1 \leqslant k \leqslant n-1)$, let $V_{1+k} = \{(x_1, \ldots, x_{k-n+i}, s_{k-n+i+1}, \ldots, s_{i-1}, x_i, \ldots, x_n) \mid x_i \in V(G_i) - \{s_i\}$, and $x_j \in (G_j)$ for $j = i+1, \ldots, n, 1, 2, \ldots, k-n+i\}$. Each pair of nodes in V_{1+k}, $(x_1, \ldots, x_{k-n+i-1}, y_{k-n+i}, s_{k-n+i+1}, \ldots, s_{i-1}, x_i, \ldots, x_n)$ and $(x_1, \ldots, x_{k-n+i-1}, y_{k-n+i}', s_{k-n+i+1}, \ldots, s_{i-1}, x_i, \ldots, x_n)$ are connected by an edge if and only if y_{k-n+i} and y_{k-n+i}' are adjacent in T_{k-n+i}^s.

(4) Let $V_{n+1} = \{(x_1, \ldots, x_{i-1}, s_i, x_{i+1}, \ldots, x_n) \mid x_j \in V(G_j)$ for $j \neq i\}$. For each node $(x_1, \ldots, x_{i-1}, s_i, x_{i+1}, \ldots, x_n)$ in V_{n+1}, it is connected to $(x_1, \ldots, x_{i-1}, t_i, x_{i+1}, \ldots, x_n)$ by an edge, where t_i is the leftmost son of s_i in T_i^s.

Apparently, $V_1 - \{s\} \subset V_2 \subset \cdots \subset V_n$, and $V_n \cup V_{n+1}$ is equal to $V(G)$. Let T_i be the graph with the node set $V(G)$ and the edge set specified in (1)–(4) above. From the construction of T_i, for each node u in T_i there is a path from s to u. Hence, T_i is connected. Let b_i denote $|V(G_i)|$ for each i $(1 \leqslant i \leqslant n)$. The numbers of edges constructed in $V_1, V_2, V_3, \ldots, V_n, V_{n+1}$ are $b_i - 1, (b_i - 1)(b_{i+1} - 1), (b_i - 1)b_{i+1}(b_{i+2} - 1), \ldots, (b_i - 1)b_{i+1} \cdots b_n b_1 \cdots b_{i-2}(b_{i-1} - 1), b_1 \cdots b_{i-1} b_{i+1} \cdots b_n$, respectively. The edge between $(s_1, s_2, \ldots, s_n)$ and $(s_1, \ldots, s_{i-1}, t_i, s_{i+1}, \ldots, s_n)$ appears twice, both in V_1 and in V_{n+1}. All other edges appear exactly once. Hence, the number of edges of T_i is $(b_i-1)+(b_i-1)(b_{i+1}-1)+(b_i-1)b_{i+1}(b_{i+2}-1)+\cdots+(b_i-1)b_{i+1}\cdots b_n b_1 \cdots b_{i-2}(b_{i-1} - 1) + b_1 \cdots b_{i-1} b_{i+1} \cdots b_n - 1 = b_1 b_2 \cdots b_n - 1 = |V(G)| - 1$. Hence, T_i is a spanning tree rooted at $s = (s_1, \ldots, s_n)$ of G.

For each node $u = (u_1, \ldots, u_n)$ of G and each i $(1 \leqslant i \leqslant n)$, a path from s to u in T_i excluding s and u is denoted by $p_i(s,u)$. We show that $p_1(s,u), \ldots, p_n(s,u)$ are node-disjoint. Suppose that $s_j = u_j$ if $j \in \{i_1, \ldots, i_k\}$, and otherwise $s_j \neq u_j$, where $1 \leqslant i_1 < \cdots < i_k \leqslant n$.

We first consider the case where $j \in \{i_1, \ldots, i_k\}$. Then the jth component of every node on $p_j(s,u)$ is t_j, the leftmost son of s_j in T_j^s, whereas for any $i \neq j$, the jth component of every node on $p_i(s,u)$ is s_j. Hence, for any $i \neq j$, $p_i(s,u)$ and $p_j(s,u)$ are node-disjoint.

We next assume that $j \notin \{i_1, \ldots, i_k\}$. Then we need only to prove that for $i \notin \{i_1, \ldots, i_k\}$ and $i \neq j$, $p_i(s,u)$ and $p_j(s,u)$ are node-disjoint. Suppose that $i \notin \{i_1, \ldots, i_k\}$. Without loss of generality, we may assume that $i < j$. Let r be the largest integer such that $r < j$ and $u_r \neq s_r$ (i.e., $r \notin \{i_1, \ldots, i_k\}$). Since $u_i \neq s_i$, such r exists. Then the jth component of any node on $p_i(s,u)$ is s_j unless its rth component is u_r. On the other hand, the rth component of any node on $p_j(s,u)$ is not u_r, and the jth component of any node on $p_j(s,u)$ is not s_j. Hence, $p_i(s,u)$ and $p_j(s,u)$ are node-disjoint.

Thus, the n spanning trees $T_1, T_2, \ldots, T_n$ we constructed above are independent spanning trees rooted at s of $G = G_1 \times G_2 \times \cdots \times G_n$. □

Note: "Disjoint spanning trees" of a product network constructed in [12] are just edge disjoint. Those spanning trees are much weaker than independent spanning trees considered in this paper, i.e., the fault tolerance of broadcasting along the spanning trees in [12] is much weaker than ours in this paper against faulty nodes.

The next theorem is immediately from Theorem 1 and the discussion in Section 2.

Theorem 3. *For broadcasting in the product network $G_1 \times \cdots \times G_n$, the degrees of tolerance to crash and Byzantine node (and/or link) faults are $n-1$ and $\lfloor (n-1)/2 \rfloor$, respectively.*

We next consider the efficiency of the broadcasting scheme based on n-independent spanning trees. The time interval for a message transmission through a link is assumed to be a unit time. This time unit is also called a step. Broadcasting time is measured as the number of concurrent steps from the start to the end of the broadcasting. The number of messages transmitted during the broadcasting is measured as the total number of link transmissions in the network. In general, we cannot evaluate the number of link transmissions in a network with faults of the Byzantine type, since faulty nodes and/or links may arbitrarily fabricate messages and send them. However, if there are no faults, the number of link transmissions during the broadcasting through the n-independent spanning trees of a product network is $n(N-1)$, where N is the order of the product network and n is the number of its component graphs.

We can consider two types of broadcasting with respect to message transmission methods. One is the all-port broadcasting, and the other is the one-port broadcasting. For the former type each node can transmit messages to all adjacent nodes in a step, whereas for the latter type each node can transmit a message to only one of the adjacent nodes or receive a message from one of the adjacent nodes in a step.

Theorem 4. *Let $s=(s_1,\ldots,s_n)$ be the source node in the product network $G=G_1\times\cdots\times G_n$. Suppose that α_i is the minimum time of all-port broadcasting from s_i throughout G_i and that β_i is the minimum time of one-port broadcasting from s_i throughout G_i, for $i=1,2,\ldots,n$. Then the following two assertions hold true*:

(1) *There exists a fault-tolerant all-port broadcasting from s throughout G such that its running time is $1+\sum_{i=1}^{n}\alpha_i$, and its degree of tolerance to Byzantine faults is $\lfloor (n-1)/2\rfloor$.*

(2) *There exists a fault-tolerant one-port broadcasting from s throughout G such that its running time is $2\sum_{i=1}^{n-1}\beta_i+\beta_n+1$, and its degree of tolerance to Byzantine faults is $\lfloor (n-1)/2\rfloor$.*

Proof. From Theorem 3 both types of broadcasting can tolerate up to $\lfloor (n-1)/2\rfloor$ faults of the Byzantine type.

(1) Since α_i is the minimum time of all-port broadcasting from s_i throughout G_i, there exists a spanning tree T_i^s rooted at s_i of G such that its height is α_i for each i $(1\leqslant i\leqslant n)$. Let $T_1,T_2,\ldots,T_n$ be n-independent spanning trees rooted at $s=(s_1,\ldots,s_n)$ of G constructed from $T_1^s,T_2^s,\ldots,T_n^s$. From the construction described in the proof of Theorem 1, the height of T_i is $\sum_{j\in\{1,\ldots,i-1,i+1,\ldots,n\}}\alpha_j+\max\{\alpha_i,2\}$. Hence, the all-port broadcasting through the n-independent spanning trees $T_1,T_2,\ldots,T_n$ finishes in $1+\sum_{i=1}^{n}\alpha_i$ steps.

(2) Again, we suppose that for each i $(1\leqslant i\leqslant n)$, T_i^s is the spanning tree rooted at s_i of G_i and that the one-port broadcasting through T_i^s needs β_i steps. Let $T_1,T_2,\ldots,T_n$ be the n-independent spanning trees rooted at $s=(s_1,\ldots,s_n)$ of G constructed from $T_1^s,T_2^s,\ldots,T_n^s$.

We consider a particular implementation of the broadcasting to estimate its running time. We divide the broadcasting process into $2n$ rounds, say $round_1, round_2,\ldots, round_{2n}$. For each i $(1\leqslant i\leqslant n-1)$, each of $round_i$ and $round_{n+i}$ consists of β_i steps, $round_n$ consists of β_n steps, and $round_{2n}$ consists of just one step. Consider broadcasting through T_i $(1\leqslant i\leqslant n)$. From the construction of the n-independent spanning trees the following implementation is possible: the development of the jth component can be done along T_j^s during $round_j$ if $j>i$, and can be done during $round_{j+n}$ if $j<i$. The development of the ith component can be done along T_i^s during $round_i$ except for one-step moves from $(x_1,\ldots,x_{i-1},t_i,x_{i+1},\ldots,x_n)$ to $(x_1,\ldots,x_{i-1},s_i,x_{i+1},\ldots,x_n)$, where for each k, x_k is any node in the kth component graph. These one-step moves can be done

in a step of $round_{n+i}$. We can see that this implementation is a one-port broadcasting by considering the following restriction on transmissions: Let $u=(u_1,\dots,u_{i-1},u_i,u_{i+1},\dots,u_n)$ and $u'=(u_1,\dots,u_{i-1},u_i',u_{i+1},\dots,u_n)$ be a pair of adjacent nodes in $G_1\times\cdots\times G_n$. We allow node u to send message to u' only at the jth step of $round_i$ or $round_{i+n}$ if in the one-port broadcasting in G_i through the spanning tree T_i^s, the transmission from u_i to u_i' is at the jth step. Since $T_1,\dots,T_n$ are independent, u has at most one message to send at each step. Hence, the one-port broadcasting finishes in $2\sum_{i=1}^{n-1}\beta_i+\beta_n+1$ steps. □

4. Randomly distributed faults

For large networks the worst case rarely appears. That is, broadcasting succeeds with a high probability even if many more than $\lfloor (n-1)/2\rfloor$ faults exist in a product network of n component graphs.

We assume that f faulty nodes of the Byzantine type are randomly distributed in a product network and that all the links are healthy. We consider broadcasting from the source node through n-independent spanning trees of a product network. The purpose of this section is to show the relation between the number of faulty nodes and the probability of successful broadcasting, i.e., we show that the broadcasting in a product network G can tolerate a constant fraction of $N/(b^3n)$ Byzantine faults randomly distributed in G with a high probability, if n is sufficiently large, where N is the order of G, n is the number of component graphs of G, and b is the maximum among the orders of the component graphs.

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, and denote the set of all configurations of G with f faulty nodes by $\mathbf{C}_G^f$. If broadcasting in c_G^f succeeds, then c_G^f is called a successful configuration, otherwise c_G^f is called a failed configuration (see Fig. 2). We denote the set of all successful configurations of G with f faulty nodes and the set of all failed configurations of G with f faulty nodes by

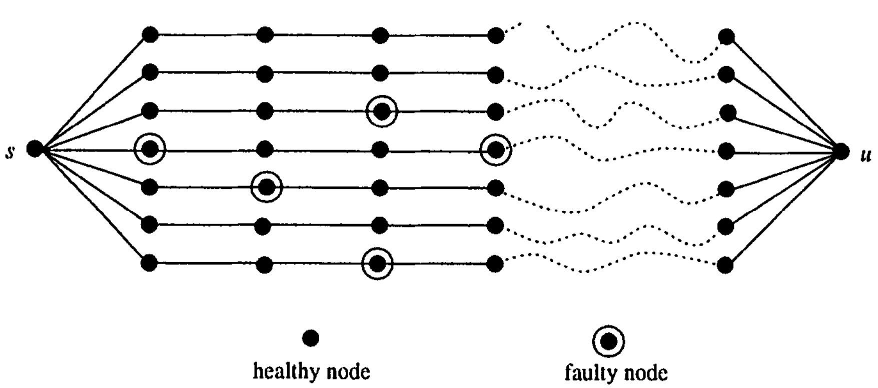

Fig. 2. A failed configuration.

$\mathbf{SC}_G^f$ and $\mathbf{FC}_G^f$, respectively. Then the probability of successful broadcasting in G with f faulty nodes is equal to $|\mathbf{SC}_G^f|/|\mathbf{C}_G^f| = 1 - |\mathbf{FC}_G^f|/|\mathbf{C}_G^f|$.

Let us consider a product network $G = G_1 \times \cdots \times G_n$, where the order of each G_i $(1 \leqslant i \leqslant n)$ is not greater than b. Let s be the source node of G, and let $q_1(s,u), \dots, q_n(s,u)$ be n paths from s to u in the n-independent spanning trees $T_1, T_2, \dots, T_n$, respectively. As shown in the previous section, these n paths are internally node-disjoint. By the broadcasting scheme the message from the source node s is disseminated to each node u of G through the n-independent spanning trees. From the construction of the n-independent spanning trees, for each i $(1 \leqslant i \leqslant n)$ there are at most nb nodes on $q_i(s,u)$. If in a configuration c_G^f, more than $\lfloor (n-1)/2 \rfloor$ paths among $q_1(s,u)$, $q_2(s,u), \dots$, $q_n(s,u)$ contain faulty nodes, then c_G^f is called a failed configuration for u. We denote the set of all failed configurations for u by $\mathbf{FC}_G^f(u)$. Then we have

$$|\mathbf{FC}_G^f| < \sum_{u \in V(G)} |\mathbf{FC}_G^f(u)|.$$

Theorem 5. *If f faulty nodes of the Byzantine type are randomly distributed in a product network $G = G_1 \times \cdots \times G_n$ with no faulty links, then the broadcasting through the n-independent spanning trees in G is successful with a probability higher than $1 - (4b^3nf/N)^{\lceil n/2 \rceil}$, where N is the order of G and b is the maximum order of the component graphs.*

Proof. The probability of successful broadcasting through the n-independent spanning trees of G is

$$\frac{|\mathbf{SC}_G^f|}{|\mathbf{C}_G^f|} = 1 - \frac{|\mathbf{FC}_G^f|}{|\mathbf{C}_G^f|} > 1 - \frac{\sum_{u \in V(G)} |\mathbf{FC}_G^f(u)|}{|\mathbf{C}_G^f|}.$$

For any node u of G, there are n internally node-disjoint paths $q_1(s,u), \dots, q_n(s,u)$ of length at most bn. Let us consider the configurations where exactly $\lceil n/2 \rceil$ faulty nodes exist and they are on the different paths among the n node-disjoint paths from s to u. The number of such configurations is apparently no more than $(bn)^{\lceil n/2 \rceil} \binom{n}{\lceil n/2 \rceil}$. By repeatedly counting, we have

$$|\mathbf{FC}_G^f(u)| < (bn)^{\lceil n/2 \rceil} \binom{n}{\lceil n/2 \rceil} \binom{N - \lceil n/2 \rceil}{f - \lceil n/2 \rceil}.$$

Hence,

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

Then we have

$$\frac{|\mathbf{SC}_G^f|}{|\mathbf{C}_G^f|} > 1 - \left(\frac{4b^3nf}{N}\right)^{\lceil n/2\rceil}. \qquad \square$$

The probabilistic assumption in Theorem 5 is that there are exacly f faulty nodes randomly distributed in the network. We can also more realistically consider all configurations with up to f faults to estimate the probability of successful broadcasting. In this case, we assume that the probabilities of configurations with the same number of faults are equal. Then we can derive a similar result for such a distributed model of faults as described in the next theorem.

Theorem 6. *If up to f faulty nodes of the Byzantine type are randomly distributed in a product network $G = G_1 \times \cdots \times G_n$ with no faulty links, then the the broadcasting through the n-independent spanning trees in G is successful with a probability higher than $1 - (4b^3nf/N)^{\lceil n/2\rceil}$, where N is the order of G and b is the maximum order of the component graphs.*

Proof. We assume that for any i $(i \leqslant f)$, every configuration with i faulty nodes is equally probable. Let the probability that exactly i faulty nodes exist be p_i. Since we assume that there are up to f faulty nodes exist in the network, $\sum_{i=0}^{f} p_i = 1$. From Theorem 5 we have

$$\frac{|\mathbf{SC}_G^i|}{|\mathbf{C}_G^i|} > 1 - \left(\frac{4b^3ni}{N}\right)^{\lceil n/2\rceil} \geqslant 1 - \left(\frac{4b^3nf}{N}\right)^{\lceil n/2\rceil}.$$

Hence, when there are up to f faulty nodes randomly distributed in the network, the probability of successful broadcasting is

$$\sum_{i=0}^{f} \frac{|\mathbf{SC}_G^i|}{|\mathbf{C}_G^i|} p_i > 1 - \left(\frac{4b^3nf}{N}\right)^{\lceil n/2\rceil}. \qquad \square$$

The next corollary is immediately from Theorem 6.

Corollary 1. *For any $k > 1$, if at most $N/4b^3nk$ faulty nodes of the Byzantine type are randomly distributed in a product network $G = G_1 \times \cdots \times G_n$, then the broadcasting through the n-independent spanning trees in G is successful with a probability higher than $1 - k^{-\lceil n/2\rceil}$, where N is the order of G and b is the maximum order of the component graphs.*

From Corollary 1, we can say that if b is a constant independent of N and n, and if n is sufficiently large, then the product network can tolerate a fraction of N/n faulty nodes of the Byzantine type randomly distributed in the product network with a high probability.

We can similarly prove the next corollary.

Corollary 2. *For any $k>1$, if at most N/b^2nk faulty nodes of the crash type are randomly distributed in a product network $G=G_1\times\cdots\times G_n$, then the broadcasting through the n-independent spanning trees in G is successful with a probability higher than $1-k^{-n}$, where N is the order of G and b is the maximum order of the component graphs.*

In the case where faulty links exist in a product network with no faulty nodes, we can derive a similar result to Theorem 6.

Theorem 7. *If up to f faulty links of the Byzantine type are randomly distributed in a product network $G=G_1\times\cdots\times G_n$ with no faulty nodes, then the broadcasting through the n-independent spanning trees in G is successful with a probability higher than $1-(4b^3nf/L)^{\lceil n/2\rceil}$, where L is the size of G and b is the maximum order of the component graphs.*

The following two corollaries can be also similarly derived.

Corollary 3. *For any $k>1$, if at most $L/4b^3nk$ faulty links of the Byzantine type are randomly distributed in a product network $G=G_1\times\cdots\times G_n$, then the broadcasting through the n-independent spanning trees in G is successful with a probability higher than $1-k^{-\lceil n/2\rceil}$, where L is the size of G and b is the maximum order of the component graphs.*

Corollary 4. *For any $k>1$, if at most L/b^2nk faulty links of the crash type are randomly distributed in a product network $G=G_1\times\cdots\times G_n$, then the broadcasting through the n-independent spanning trees in G is successful with a probability higher than $1-k^{-n}$, where L is the size of G and b is the maximum order of the component graphs.*

5. Tolerate O(L) faulty links

From Theorem 7 or Corollary 3 in the previous section, if the number of faulty links is at most $L/4b^3nk$, then broadcasting through the n-independent spanning trees succeeds with a probability higher than $1-k^{-\lceil n/2\rceil}$. This means that the broadcasting can tolerate, with a high probability, $O(L/n)$ faulty links of the Byzantine type randomly distributed in a product network of n-component graphs. In this section, we give another broadcasting scheme that can tolerate, with a high probability, $O(L)$ faulty links of the Byzantine type randomly distributed in a product network of n-component graphs. This result supports the intuition that the malicious results caused by faulty links of the Byzantine type are fewer than those caused by faulty nodes of the Byzantine type.

Let us consider broadcasting from the source node $s=(s_1,\dots,s_n)$ in a product network $G=G_1\times\cdots\times G_n$. Let T be one of the n-independent spanning trees rooted at s of G, say $T=T_1$, where $T_1,\dots,T_n$ are specified in the proof of Theorem 2. Consider a pair of nodes of G, $u=(u_1,\dots,u_{i-1},u_i,u_{i+1},\dots,u_n)$ and $u'=(u_1,\dots,u_{i-1},u_i',u_{i+1},\dots,u_n)$, where u_i and u_i' are adjacent in G_i. Then there exist the following n internally node-disjoint paths from u to u', $q_1(u,u'),\dots,q_n(u,u')$: Let $q_i(u,u')$ be the link between u and u'. Let for each j $(1\leqslant j\neq i\leqslant n)$, $q_j(u,u')$ be the path,

$$
\begin{aligned}
u &\leftrightarrow (u_1,\dots,u_{j-1},t_j,u_{j+1},\dots,u_{i-1},u_i,u_{i+1},\dots,u_n)\\
&\leftrightarrow (u_1,\dots,u_{j-1},t_j,u_{j+1},\dots,u_{i-1},u_i',u_{i+1},\dots,u_n)\\
&\leftrightarrow (u_1,\dots,u_{j-1},u_j,u_{j+1},\dots,u_{i-1},u_i',u_{i+1},\dots,u_n)=u',
\end{aligned}
$$

where t_j is an adjacent node of u_j in G_j.

For each j $(1\leqslant j\neq i\leqslant n)$, there are $deg(u_j)$ choices of $q_j(u,u')$, or say t_j, where $deg(u_j)$ is the degree of u_j. These $deg(u_j)$ paths, $q_j(u,u')$'s are internally node-disjoint, and each of $q_j(u,u')$ and any $q_k(u,u')$ $(k\neq j)$ are also internally node-disjoint. Hence, there exist at least $1+\sum_{1\leqslant j\neq i\leqslant n} deg(u_j)$ node-disjoint paths of length less than or equal to 3 between $u=(u_1,u_2,\dots,u_n)$ and $u'=(u_1,\dots,u_{i-1},u_i',u_{i+1},\dots,u_n)$. In this section, for each pair of adjacent nodes u and u' we only consider the n paths $q_1(u,u'),\dots,q_n(u,u')$ instead of all the $1+\sum_{1\leqslant j\neq i\leqslant n} deg(u_j)$ paths. However, it is not difficult to modify the broadcasting scheme so that we use all the $1+\sum_{1\leqslant j\neq i\leqslant n} deg(u_j)$ paths for connecting u and u'. The analysis of its fault tolerance will be similar.

We modify the broadcasting scheme as follows: Consider the broadcasting from s in G through the spanning tree T rooted at s. We replace the transmission between each pair of adjacent nodes in T, say, from u to u' by n transmissions through the n paths, $q_1(u,u'),\dots,q_n(u,u')$ of length less than or equal to 3. The majority voting is taken by u' on the n messages transmitted through $q_1(u,u'),\dots,q_n(u,u')$, respectively.

The difference between the principles of the modified broadcasting and the original one proposed in Section 2 is as follows: in the original broadcasting, the majority voting is taken at every node after all the transmissions have finished while in the modified broadcasting, each node takes the majority voting on the n messages received through the n paths of length at most 3 as described above. Then it transmits the message chosen from the majority ones to other nodes along the modified graph from T. The principle of the modified broadcasting is shown in Fig.3. The transmission graph obtained by the modification of the spanning tree in this way is called the transmission pseudo-tree.

The number of transmissions used by the modified one is not larger than three times of the number of transmissions by the original one. For the case of all-port broadcasting, the running time of the modified one is about three times as much as the running time of the original one. For the case of one-port broadcasting, it is difficult to estimate the running time of the modified one.

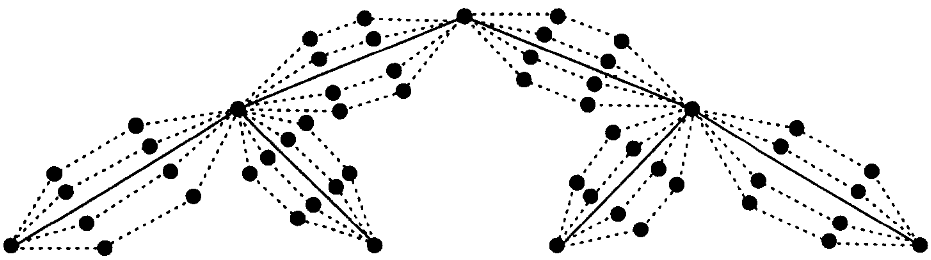

Fig. 3. Each link of T is replaced with n paths of length equal or less than 3.

Theorem 8. *If up to f faulty links of the Byzantine type are randomly distributed in a product network $G = G_1 \times \cdots \times G_n$ with no faulty nodes, then the broadcasting through the transmission pseudo-tree in G is successful with a probability higher than $1 - (12b^2 f/L)^{\lceil n/2 \rceil}$, where L is the size of G and b is the maximum order of the component graphs.*

Proof. We assume that for any i ($i \leqslant f$), every configuration with i faulty links is equally probable. We first consider the case where there exist exactly f faulty links of the Byzantine type randomly distributed in the network. We denote a configuration of G with f faulty links by $\bar{c}_G^f$, and the set of all configurations of G with f faulty links by $\bar{\mathbf{C}}_G^f$. For the broadcasting through the transmission pseudo-tree in G, we denote the set of successful configurations of G with f faulty links by $\mathbf{S}\bar{\mathbf{C}}_G^f$, and the set of failed configurations of G with f faulty links by $\mathbf{F}\bar{\mathbf{C}}_G^f$.

Let T be T_1 (i.e., the first spanning tree specified in the proof of Theorem 2). Let for each link (u,v) of T, $\mathbf{F}\bar{\mathbf{C}}_G^f(u,v)$ be the set of configurations of G such that at least $\lceil n/2 \rceil$ paths among $q_1(u,v), \ldots, q_n(u,v)$ contain faulty links. Then the following inequality holds:

$$|\mathbf{F}\bar{\mathbf{C}}_G^f| < \sum_{(u,v) \in E(T)} |\mathbf{F}\bar{\mathbf{C}}_G^f(u,v)|.$$

Since the length of each $q_i(u,v)$ ($1 \leqslant i \leqslant n$) is at most 3, we have

$$|\mathbf{F}\bar{\mathbf{C}}_G^f(u,v)| < \binom{n}{\lceil n/2 \rceil} 3^{\lceil n/2 \rceil} \binom{L - \lceil n/2 \rceil}{f - \lceil n/2 \rceil}.$$

Hence,

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

where N is the order of G. Thus, we have

$$\frac{|\mathbf{S}\bar{\mathbf{C}}_G^f|}{|\bar{\mathbf{C}}_G^f|} = 1 - \frac{|\mathbf{F}\bar{\mathbf{C}}_G^f|}{|\bar{\mathbf{C}}_G^f|} > 1 - \left(\frac{12b^2 f}{L}\right)^{\lceil n/2 \rceil}$$

The rest of the proof is similar to the proof of Theorem 6, i.e., using the inequality shown above and the same argument as the proof of Theorem 6, we can show that the probability of the successful broadcasting is higher than $1 - (12b^2 f/L)^{\lceil n/2 \rceil}$ in the case where there exist up to f faulty links of the Byzantine type randomly distributed in the network. □

The next corollary is immediate from Theorem 8.

Corollary 5. *For any $k > 1$, if at most $L/12b^2k$ faulty links of the Byzantine type are randomly distributed in a product network $G = G_1 \times \cdots \times G_n$ with no faulty nodes, then the broadcasting through the transmission pseudo-tree in G is successful with a probability higher than $1 - k^{-\lceil n/2 \rceil}$, where L is the size of G and b is the maximum order of the component graphs.*

We can similarly prove the next corollary.

Corollary 6. *For any $k > 1$, if at most $L/3bk$ faulty links of the crash type are randomly distributed in a product network $G = G_1 \times \cdots \times G_n$ with no faulty nodes, then the broadcasting through the transmission pseudo-tree in G is successful with a probability higher than $1 - k^{-n}$, where L is the size of G and b is the maximum order of the component graphs.*

6. Conclusions

We studied the reliable broadcasting in product networks. We showed that any product network of n-component graphs is an n-channel graph. The reliability of our broadcasting scheme is based on message transmissions through the n-independent spanning trees rooted at the source node. For the case where all the nodes are healthy but some links may be faulty, we proposed a modified broadcasting scheme which is more reliable against faulty links but less efficient than the original one. Since we assume that no nodes in a network have a priori information about faults, the method using independent spanning trees is the most efficient and the most reliable against Byzantine faults. However, the following two problems arise from our approach if we consider a general network:

(1) How can we construct independent spanning trees for an arbitrary graph? This is a very hard problem. In fact, it is open whether every n-connected graph has n-independent spanning trees with the same root. The problem has been solved only for k-connected graphs, $k \leqslant 3$ in [5, 19]. Furthermore, even if we know a method of

constructing independent spanning trees for some graphs, the independent spanning trees obtained by the method may not have good properties. It is important to construct independent spanning trees with good properties, e.g. with low heights and regular structures.

(2) How can we devise efficient broadcasting schemes, in particular, for one-port broadcasting, based on message transmissions through independent spanning trees if they are available? Since such broadcasting consists of sub-broadcastings, each through one of the independent spanning trees, there are few hints about how each node uses an efficient strategy at each step. A good scheme has a short broadcasting time, but a poor scheme may need a long broadcasting time.

These problems would be worthy of further investigation.

Acknowledgements

The authors would like to thank the anonymous referees who suggested a number of ways of improving the presentation of this paper.

References

[1] A. Bagchi, L. Hakimi, Information dissemination in distributed systems with faulty units, IEEE Trans. Comput. 43 (1994) 698–710.
[2] F. Bao, Y. Igarashi, K. Katano, Broadcasting in hypercubes with randomly distributed Byzantine Faults, Proc. WDAG'95, vol. 972, 1995, pp. 215–229.
[3] D. Bienstock, Broadcasting with random faults, Discrete Appl. Math. 20 (1988) 1–7.
[4] D.B. Blough, A. Pelc, Optimal communication in networks with randomly distributed Byzantine faults, Networks 23 (1993) 691–701.
[5] J. Cheriyan, S.N. Maheshwari, Finding nonseparating induced cycles and independent spanning trees in 3-connected graphs, J. Algorithms 9 (1988) 507–537.
[6] F. Cristian, H. Aghili, R. Strong, D. Dolev, Atomic broadcasting: from simple message diffusion to Byzantine agreement, Inform. Comput. 118 (1995) 158–179.
[7] L. Gargano, A. Rescigno, U. Vaccaro, Fault-tolerant hypercube broadcasting via information dispersal, Networks 23 (1993) 271–282.
[8] V. Hadzilacos, S. Toueg, A modular approach to fault-tolerant broadcasts and related problems, Technical Report, Department of Computer Science, Cornell University, 1994.
[9] S.M. Hedetniemi, S.T. Hedetniemi, A.L. Liestman, A survey of gossiping and broadcasting in communication networks, Networks 18 (1988) 1240–1268.
[10] A. Itai, M. Rodeh, The multi-tree approach to reliability in distributed networks, Inform. Comput. 79 (1988) 43–59.
[11] P. Jalote, Fault Tolerance in Distributed Systems, Prentice-Hall, Englewood Cliffs, NJ, 1994.
[12] S.R. Öhring, D.H. Hohndel, Optimal fault-tolerant communication algorithms on product networks using spanning trees, 6th IEEE Symposium on Parallel and Distributed Processing, 1994.
[13] A. Pelc, Reliable communication in networks with Byzantine link failures, Networks 22 (1992) 441–459.
[14] A. Pelc, Fault-tolerant broadcasting and gossiping in communication networks, Technical Report, Department d'Informatique, Universite du Quebec, 1995.
[15] P. Ramanathan, D.D. Kandlur, K.G. Shin, Hardware-assisted software clock synchronization for homogeneous distributed systems, IEEE Trans. Comput. 39 (1990) 514–524.
[16] P. Ramanathan, K.G. Shin, Reliable broadcasting in hypercube multiprocessors, IEEE Trans. Comput. 37 (1988) 1654–1657.

[17] P. Ramanathan, K.G. Shin, R.W. Butler, Fault-tolerant clock synchronization in distributed systems, Computer 23 (1990) 33–42.
[18] A. Youssef, Cartesian product networks, Proceedings of the 1991 Internat. Conference on Parallel Processing, vol. I, 1991, pp. 684–685.
[19] A. Zehavi, A. Itai, Three tree-paths, J. Graph Theory 13 (1989) 175–188.