

Communication in Networks with Random Dependent Faults

Evangelos Kranakis^{*†§} Michel Paquette^{*§¶} Andrzej Pelc^{†‡||}

Abstract

The aim of this paper is to study communication in networks where nodes fail in a random dependent way. In order to capture fault dependencies, we introduce the *neighborhood fault* model, where damaging events, called *spots*, occur randomly and *independently* with probability p at nodes of a network, and cause faults in the given node and all of its neighbors. Faults at distance at most 2 become dependent in this model and are positively correlated. We investigate the impact of spot probability on feasibility and time of communication in the fault-free part of the network. We show a network which supports fast communication with high probability, if $p \leq 1/c \log n$. We also show that communication is not feasible with high probability in most classes of networks, for constant spot probabilities. For smaller spot probabilities, high probability communication is supported even by bounded degree networks. It is shown that the torus supports communication with high probability when p decreases faster than $1/n^{1/2}$, and does not when $p \in 1/O(n^{1/2})$. Furthermore, a network built of tori is designed, with the same fault-tolerance properties and additionally supporting fast communication. We show, however, that networks of degree bounded by a constant d do not support communication with high probability, if $p \in 1/O(n^{1/d})$. While communication in networks with independent faults was widely studied, this is the first analytic paper which investigates network communication for random dependent faults.

Keywords: Fault-tolerance, dependent faults, communication, crash faults, network connectivity.

^{*}School of Computer Science, Carleton University, Ottawa, Ontario, K1S 5B6, Canada. kranakis@scs.carleton.ca, michel.paquette@polymtl.ca

[†]Département d'informatique et d'ingénierie, Université du Québec en Outaouais. Gatineau, Québec, J8X 3X7, Canada. pelc@uqo.ca

[‡]Research supported by Natural Sciences and Engineering Research Council of Canada (NSERC)

[§]Research supported by Mathematics of Information Technology and Complex Systems (MITACS).

[¶]Research supported by NSERC Postgraduate Scholarship.

^{||}Research supported by the Research Chair in Distributed Computing of the Université du Québec en Outaouais.

1 Introduction

As interconnection networks grow in size and complexity, they become increasingly vulnerable to component failures. Links and nodes of a network may fail, and these failures often result in delaying, blocking, or even distorting transmitted messages. It becomes important to design networks in such a way that the desired communication task be accomplished efficiently in spite of these faults, usually without knowing their location ahead of time. Such networks are called fault-tolerant.

The fundamental questions of network reliability have received much attention in past research under the assumption that components fail randomly and independently (cf., e.g. [1, 3, 4, 12, 14] and the survey [13]). On the other hand, empirical work has shown that positive correlation of faults is a more reasonable assumption for networks [6, 16, 18]. In [18], the authors provide empirical evidence that data packets losses are spatially correlated in networks, and in [16], the authors use the assumption of failure spatial correlation to enhance network traffic management. Furthermore, in [6], the authors simulate failures in a sensor network using a model much like that of the present paper; according to these authors, the environment provides many spatially correlated phenomena resulting in such fault patterns. Physical and logical phenomena generally affect physical components, causing failures in a positively correlated way. E.g., on August 14, 2003, faults cascaded on the power distribution network and deprived part of North America of electricity. Logical phenomena, like computer viruses and worms, also cause dependent faults. Lightning strikes hitting one node of an electric network cause power outages in entire city blocks.

As our society is increasingly dependent on information networks, it becomes essential to study questions relating to tolerance of dependent positively correlated faults. However, no analytic work has been done for communication networks under this assumption about faults.

In this paper, we consider the problem of feasibility and time of communication in networks with dependent positively correlated faults. To the best of our knowledge, this is the first analytic paper which provides this type of results for network communication.

1.1 Model and Problem Definition

A communication network is modeled as an undirected graph $G = (V, E)$ with a set of nodes V connected by a set of undirected links E . We say that two nodes are *adjacent* (or *neighbors*) if they share a link. The distance between nodes $u, v \in V$ is the minimum number of links which must be traversed from u to reach v ; it is denoted by $dist(u, v)$. In a network, $\Gamma(u)$ is the set of nodes adjacent to u ; $\Gamma_i(u)$ is the set of nodes $v \in V$ whose distance from u is i ; we also denote by $\Gamma_{\leq i}(u)$ the set of nodes $v \in V$ whose distance from u is at most i . A node is said to be *functional*, or *fault-free*, when it executes only its predefined algorithm without any deviation, and doing so, transmits all messages correctly, in a timely manner and without any loss; a node which is not functional is said to be *faulty*. Faults can be of different types: at opposite ends of the spectrum are *crash* and *Byzantine* faults. Faults of the crash type cause faulty components to stop all communication; these components can neither send, receive nor relay any message. Faulty nodes of the Byzantine type may behave arbitrarily (even maliciously) as transmitters. We say that faults are permanent when they affect the nodes for the entire duration of a communication process; otherwise, the faults are said to be transient. In this paper, we assume that faults are permanent and of crash type. Throughout the paper, \log means logarithm with base 2 and \ln means the natural logarithm.

We consider communication in the fault-free part of the network, where all nodes exchange messages with each other. Communication among functional nodes is feasible if the fault-free part of the network is connected and contains at least two nodes. We measure communication time under the all-port message passing model, where nodes can communicate with all their neighbors in one round, and under the 1-port model, in which every node can send a message to at most one neighbor in one round. Under the all-port model, communication can be completed in time D if the fault-free part of the network has diameter D . Hence, we study the connectivity and diameter of the fault-free part of the network. Moreover, we seek networks of low maximum degree Δ . Since in the 1-port model communication can be completed in time $D\Delta$, networks of low maximum degree and low diameter of the fault-free part support fast communication also in the 1-port model.

In order to capture fault dependencies, we introduce the *neighborhood fault* model, where damaging events, called *spots*, occur randomly and *independently* at nodes of a network, with probability p , and cause permanent crash faults in the given node and all of its neighbors. Faults at distance at most 2 become dependent in this model and are positively correlated. We investigate the impact of spot probability on feasibility and time of communication in the fault-free part of the network.

We design general networks and bounded degree networks which support fast and highly reliable communication despite relatively high spot probabilities. We also prove bounds on spot probability such that highly reliable communication is not supported.

We focus attention on the problem of feasibility and time of communication, guaranteed with *high probability*, i.e., with probability converging to 1 as the size of the network grows. Under the all-port model, in which time of communication is proportional to the diameter, this problem reduces to the question for what spot probability the fault-free part of the network is connected and when it has diameter at most D , with high probability. Under the 1-port model, the same reduction is valid for networks of a given degree.

1.2 Related Work

Dependent fault models were introduced in the study of integrated circuit manufacturing yields. This research models defects as the result of impurities, positioned randomly and independently, affecting nearby circuit components in a dependent way. Results were proposed mainly according to the *quadrat-based* and *center-satellite* approaches. In [15], the author proposed a coarse approach to analyzing production yields based on the assumption that faults occurred in clusters inside a defined grid pattern on Very Large Scale Integration (VLSI) wafers; this quadrat-based model offered provably good results and ease of use required by the industry. Then, in [10], the authors introduced a detailed model of manufacturing defects in VLSI wafers based on the *center-satellite concept* for ecological sampling [17]. Later on, in [2], the authors proposed a simplified center-satellite model of manufacturing defects on VLSI wafers for the study of the memory array reconfiguration problem. In fact, both the center-satellite and quadrat-based approaches are still in use for System on Chip (SoC) (cf., e.g., [8, 9]) and VLSI (cf., e.g., [5, 19]) applications. Throughout this field of literature, the consensus is that results originating from the center-satellite approach, as opposed to quadrat-based approaches, are more difficult to apply but provide better prediction quality.

The above approach should be contrasted with the literature on fault-tolerant communication in networks. Many results concerned random link and/or node failures (cf., e.g. [1, 3, 4, 12, 14] and the survey [13]) but, to the best of our knowledge, in all cases faults

were assumed to be independent. In [1], the author shows the existence of networks in which $O(\log n)$ -time broadcast can be done, under the 1-port model, with high probability, despite links which fail randomly and independently with positive constant probability. In [3], the authors design a network of logarithmic degree which can support high probability communication in time $O(\log n)$ when faults occur randomly and independently on links and nodes with any constant probabilities smaller than 1. In [12], the authors design a similar network which can support communication with high probability in time $O(\log^2 n)$ with Byzantine faults.

Our present research focuses on communication network failures which occur in a dependent way. We consider networks modeled by arbitrary graphs, hence the geometry-dependent, quadrat-based approach to fault dependencies is not appropriate. Our neighborhood fault model, more appropriate for general graphs, is a simplified version of the center-satellite approach.

1.3 Our Results

All our results address the general problem for which spot probabilities p there exist networks supporting communication with high probability, and if so, if this communication is fast in the all-port and 1-port models. Hence we ask for which spot probabilities the fault-free part of the network is connected of size larger than 1, and if so, does it have a small diameter. Moreover, in our positive results we seek networks of low maximum degree. The results are summarized in Table 1. In what follows, we discuss them in detail.

In Section 2, we address the questions regarding general networks. We first show that there exists a constant c , such that for spot probability $p \leq 1/c \log n$, there exists an n -node graph whose fault-free part has logarithmic diameter and logarithmic degree, with high probability. Hence it supports high probability communication in time $O(\log n)$ in the all-port model and in time $O(\log^2 n)$ in the 1-port model. On the negative side, we show that for constant spot probability p , there exist constants c_1 and c_2 such that: if all degrees in a graph are at most $c_1 \log n$ then the graph is disconnected with high probability; if all degrees in a graph are at least $c_2 \log n$ then the graph has all nodes faulty with high probability. In either case, highly reliable communication is not possible. This leaves some very particular networks undecided. For example, this negative result does not cover the important case of the hypercube, for some constant spot probabilities. Therefore, we study the hypercube separately and prove that, for any constant spot probability $0 < p \leq 1$, this network does not support high probability communication. The above should be contrasted with the results from [3, 4] showing that, for *independent* faults, fast highly reliable communication is possible for arbitrary constant fault probabilities in some graphs and for small constant fault probability, even in the hypercube.

In Section 3, we investigate communication in bounded degree networks. We show that the torus supports communication with high probability when $p \in 1/\omega(n^{1/2})$. (As usual, $\omega(f)$ denotes the set of functions g such that $g/f \rightarrow \infty$.) However, the diameter of an n -node torus is at least $\Theta(\sqrt{n})$ and the fault-free part has the same large diameter. Hence we seek networks with the same fault-tolerance properties, but with small diameter. We construct a bounded degree network built of tori, whose fault-free part has diameter $O(\log n)$ whenever $p \in 1/\omega(n^{1/2})$. Hence this network supports high probability communication in logarithmic time, both in the all-port and in the 1-port models. On the negative side, we show that neither the torus nor the above network can support highly reliable communication when $p \in 1/O(n^{1/2})$. Finally, we prove that networks of degree bounded by a constant d cannot

support communication with high probability when $p \in 1/O(n^{1/d})$.

Spot probability	Network	Degree	Diameter	Theorem
$p \leq \frac{1}{c \log n}$	<i>RBT</i>	$\Theta(\log n)$	$O(\log n)$	Theorem 1
$p \in \Theta(1)$	Any	$\leq c_1 \log n$ $\geq c_2 \log n$	not feasible	Theorem 2
$p \in \Theta(1)$	Hypercube	$\log n$	not feasible	Theorem 3
$p \in 1/\omega(n^{1/2})$	Torus	4	$\Omega(\sqrt{n})$	Theorem 4
$p \in 1/O(n^{1/2})$	Torus	4	not feasible	Theorem 6
$p \in 1/\omega(n^{1/2})$	Toroidal Tree	8	$O(\log n)$	Theorem 5
$p \in 1/O(n^{1/2})$	Toroidal Tree	8	not feasible	Theorem 7
$p \in 1/O(n^{1/d})$	Any	$d \in \Theta(1)$	not feasible	Theorem 8

Table 1: Summary of results. “Diameter” means the diameter of the fault-free part of the network. “Not feasible” in the Diameter column means that communication with high probability is not feasible.

2 General Networks

In this section, we focus on general networks. We first design a network which supports communication with high probability when spot probability $p \leq 1/c \log n$, for some positive constant c . We then establish two bounds on node degrees showing that a large class of networks cannot support communication with high probability when spot probability is a positive constant.

2.1 Upper Bounds

This section is dedicated to proving the following result.

Theorem 1 *There exists a n -node graph whose fault-free part has diameter $O(\log n)$ and logarithmic degree, with high probability, for spot probability $p \leq 1/c \log n$, where c is some positive constant.*

The network construction is based on a binary tree structure where each node of the tree represents a group of nodes and each link of the tree represents a random set of links between nodes in adjacent groups. To be more precise, for a fixed m , we define a random n -node graph $G(n, m)$. Let $x = \lceil n/m \rceil$. Partition the set of all nodes into subsets S_1, \dots, S_x , of size m , (S_x of size at most m) called *supernodes*. Let $\mathcal{S} = \{S_1, \dots, S_x\}$ be the set of all supernodes.

Let $L = \lfloor \log x \rfloor$. Arrange all supernodes into a binary tree T with $L + 1$ levels $0, 1, \dots, L$, placing each supernode S_i on level $\lfloor \log i \rfloor$. Level 0 contains the root and levels $L - 1$ and L contain leaves of T . The supernode S_1 , is the root of T . For every $1 \leq i \leq \lfloor x/2 \rfloor$, S_{2i} is the left child of S_i and S_{2i+1} is the right child of S_i in T (S_{2i+1} exists if $x \geq 2i + 1$). For every $1 < i \leq x$, supernode $S_{\lfloor i/2 \rfloor}$ is the parent of S_i . If a supernode is a parent or a child of another supernode, we say that these supernodes are adjacent in T .

The set of edges of $G(n, m)$ is defined as follows. If supernodes S_i and S_j are adjacent in T , then there is an edge in $G(n, m)$ between any node in S_i and any node in S_j with probability p_l . Moreover, supernodes have no interior links. The graph $G(n, m)$ is called a Random Binary Thick Tree (*RBTT*). See Figure 1.




Figure 1: Random Binary Thick Tree

In the remainder of this section, we analyze *RBTT* and show that, if $p \leq 1/c \log n$, for some constant $c > 0$ to be defined below, then it supports communication with high probability in time $O(\log n)$. Before we are ready to prove Theorem 1, we need to establish an upper bound on the number of spots in each supernode, and a lower bound on the number of functional nodes in each supernode. These two results will be used in the proof of two connectivity lemmas needed in the proof of the theorem. We consider the n -node *RBTT* with link probability $p_l = 1/18 \ln n$ and $m = \lceil 1152 \ln^2 n \rceil$ nodes per supernode, when spot probability is $p \leq 1/(768 \ln n)$. Hence, we take $c = 768/\ln 2$.

Let C_1 be the event that all supernodes in *RBTT* contain less than $6 \ln n + 1$ spots.

Lemma 1 *The event C_1 occurs with probability at least $1 - 1/n$.*

Proof. Let $p = k/768 \ln n$, with $0 < k \leq 1$. Let N_i be the random variable which counts the number of spots in S_i . Since $1152 \ln^2 n \leq m < 1152 \ln^2 n + 1$, we have that $(3/2)k \ln n \leq E[N_i] = mp < k((3/2) \ln n + 1/4)$. Using Chernoff bounds (see [7, 11]) with parameter $(1 + \epsilon) = 4/k$, we can show that, for a fixed S_i , we have $N_i < 6 \ln n + 1$ with high probability. More precisely,

$$\begin{aligned}
 \Pr[C_1] &= \Pr[(\forall S_i \in \mathcal{S}) N_i < 6 \ln n + 1] \\
 &\geq 1 - \sum_{S_i \in \mathcal{S}} \Pr[N_i \geq (4/k)k((3/2) \ln n + 1/4)] \\
 &\geq 1 - n \left(\frac{e^{(4/k-1)}}{(4/k)^{(4/k)}} \right)^{k((3/2) \ln n + 1/4)} \\
 &\geq 1 - n \left(\frac{ke}{4} \right)^{6 \ln n} \\
 &\geq 1 - nn^{-6 \ln(4/e)} \\
 &\geq 1 - 1/n,
 \end{aligned}$$

for $0 < k \leq 1$. ■

Let C_2 be the event that all supernodes in $RBTT$ have more than $288(1-\epsilon)\ln^2 n$ functional nodes, for a given constant $0 < \epsilon \leq 1$.

Lemma 2 *The event C_2 occurs with probability at least $1 - 1/n^{d \log n}$, for some positive constant d .*

Proof. Fix the set of spots in the graph. Pick a spot-free node u . The event FF_u that node u is fault-free occurs if none of its links has a spot as an endpoint. Since u is in a supernode adjacent to at most 3 other supernodes, it may have a link to at most $18 \ln n + 3$ spots, if event C_1 holds. It follows, for large enough n , that

$$\begin{aligned} \Pr[FF_u] &\geq \Pr[FF_u \wedge C_1] \\ &= \Pr[FF_u \mid C_1] \Pr[C_1] \\ &\geq \left(1 - \frac{1}{18 \ln n}\right)^{18 \ln n + 3} \Pr[C_1] \\ &\geq \frac{1}{4} \left(1 - \frac{1}{n}\right). \end{aligned}$$

Let N_i be the random variable which counts the number of non-faulty nodes in S_i . Since, under event C_1 , at most $6 \ln n + 1$ spots occur in any supernode, at least $m - 6 \ln n - 1$ nodes are fault-free, if C_1 occurs. Thus,

$$\begin{aligned} E[N_i] &\geq \Pr[C_1] \cdot (m - 6 \ln n - 1) \cdot \Pr[FF_u] \\ &\geq \left(1 - \frac{1}{n}\right) \cdot (1152 \ln^2 n - 6 \ln n - 1) \cdot \frac{1}{4} \left(1 - \frac{1}{n}\right) \\ &\geq 288 \left(1 - \frac{6}{1152 \ln n} - \frac{1}{1152 \ln^2 n}\right) \left(1 - \frac{1}{n}\right)^2 \ln^2 n \end{aligned}$$

and therefore $E[N_i] \geq 288(1 - \epsilon') \ln^2 n$, for any constant $0 < \epsilon' \leq 1$, if n is sufficiently large. We use a Chernoff bound to show that, with high probability, $N_i \geq 288(1 - \epsilon')(1 - \epsilon'') \ln^2 n$ for any constant $0 < \epsilon'' \leq 1$. It follows that the probability of this event is

$$\begin{aligned} \Pr[C_2] &= \Pr[(\forall S_i \in \mathcal{S}) N_i > 288(1 - \epsilon')(1 - \epsilon'') \ln^2 n] \\ &\geq 1 - \sum_{S_i \in \mathcal{S}} \Pr[N_i \leq 288(1 - \epsilon')(1 - \epsilon'') \ln^2 n] \\ &\geq 1 - \sum_{S_i \in \mathcal{S}} e^{-(\epsilon'')^2 288(1 - \epsilon') \ln^2 n / 2} \\ &\geq 1 - nn^{-144(1 - \epsilon') \epsilon''^2 \ln n} \\ &= 1 - n^{1 - 144(1 - \epsilon') \epsilon''^2 \ln n}. \end{aligned}$$

Using the previous results, we now present two connectivity lemmas in preparation for the proof of the main theorem of this section. ■

Lemma 3 *All functional nodes are connected to at least one functional node in each supernode adjacent to their own, with probability exceeding $1 - 1/n^{13}$.*

Proof. Fix a node u . Let $N(u)$ denote the set of supernodes adjacent to the supernode containing u . Consider the event γ_{u,S_k} that u has a link to at least one functional node in a given supernode $S_k \in N(u)$. The event γ_{u,S_k} occurs unless all links from u to functional nodes in S_k do not exist. From Lemma 2, we get for any constants $0 < \epsilon', \epsilon'' \leq 1$

$$\begin{aligned}
\Pr[\gamma_{u,S_k}] &\geq \Pr[\gamma_{u,S_k} \wedge C_2] \\
&= \Pr[C_2] \Pr[\gamma_{u,S_k} \mid C_2] \\
&\geq \Pr[C_2] \left(1 - (1 - 1/18 \ln n)^{288(1-\epsilon')(1-\epsilon'') \ln^2 n}\right) \\
&\geq \Pr[C_2] \left(1 - e^{-\frac{288(1-\epsilon')(1-\epsilon'') \ln^2 n}{18 \ln n}}\right) \\
&= \Pr[C_2] \left(1 - n^{-16(1-\epsilon')(1-\epsilon'')}\right) \\
&\geq \left(1 - n^{-d \ln n}\right) \left(1 - n^{-16(1-\epsilon')(1-\epsilon'')}\right),
\end{aligned}$$

and hence $\Pr[\gamma_{u,S_k}] \geq 1 - n^{-15}$. Furthermore, since the graph contains at most n functional nodes, which should be connected to at least one functional node in at most 3 supernodes, the estimated probability is at least

$$\begin{aligned}
\Pr[(\forall u \in V (\forall S_k \in N(u)) \gamma_{u,S_k})] &\geq 1 - \sum_{u \in V} \sum_{S_k \in N(u)} \Pr[\neg \gamma_{u,S_k}] \\
&\geq 1 - 3n \Pr[\neg \gamma_{u,S_k}] \\
&\geq 1 - 3nn^{-15} \\
&> 1 - n^{-13}.
\end{aligned}$$

■

Lemma 4 *All functional node pairs in supernodes at distance 3 are connected by a fault-free path with probability at least $1 - 1/n^{1.9}$.*

Proof. This lemma is proven in steps, defining connection probabilities and lower bounds on the number of connected nodes at distances 1, 2, and 3.

Fix 4 supernodes, S_u, S_i, S_j, S_k , which form a simple path in RBTT. I.e., S_u is adjacent to S_i , which is adjacent to S_j , which is adjacent to S_k .

Fix a node u in S_u . Let X_i be the random variable which counts the number of functional nodes $i \in \Gamma(u)$ located in S_i . From Lemma 2, each supernode contains more than $288(1 - \epsilon) \ln^2 n$ fault-free nodes, for any $0 < \epsilon \leq 1$ with probability $1 - 1/n^{d \log n}$, for some positive constant d . Since the link probability is $p_l = 1/18 \ln n$,

$$\begin{aligned}
E[X_i] &\geq \Pr[C_2] 288(1 - \epsilon) \ln^2 n / 18 \ln n \\
&= (1 - 1/n^{d \log n}) 16(1 - \epsilon) \ln n \\
&\geq 16(1 - \epsilon') \ln n,
\end{aligned}$$

with some $1 > \epsilon' > \epsilon$. The probability that a functional node has at most $16(1 - \sqrt{3/8(1 - \epsilon')})(1 - \epsilon') \ln n$ such neighbors is

$$\begin{aligned} \Pr \left[X_i \leq 16 \left(1 - \sqrt{3/8(1 - \epsilon')} \right) (1 - \epsilon') \ln n \right] &\leq e^{-\left(\sqrt{3/8(1 - \epsilon')} \right)^2 16(1 - \epsilon') \ln n / 2} \\ &= n^{-3}. \end{aligned}$$

Let A be the event that node u has at least $16 \left(1 - \sqrt{3/8(1 - \epsilon')} \right) (1 - \epsilon') \ln n$ functional neighbors in S_i .

Assume event A occurs. Now, fix a node x in S_j . Fix a subset $S \subseteq \Gamma(u) \cap S_i$ of functional nodes, with size $16 \left(1 - \sqrt{3/8(1 - \epsilon')} \right) (1 - \epsilon') \ln n$. Denote by P_{Sx} the event that there exists a link between the node x and any node from S . This event occurs unless x has no link to some node in S . Hence,

$$\begin{aligned} \Pr[P_{Sx}|A] &= 1 - (1 - 1/18 \ln n)^{16 \left(1 - \sqrt{3/8(1 - \epsilon')} \right) (1 - \epsilon') \ln n} \\ &\geq 1 - e^{-16 \left(1 - \sqrt{3/8(1 - \epsilon')} \right) (1 - \epsilon') \ln n / (18 \ln n)} \\ &\geq \left(1 - e^{-8 \left(1 - \sqrt{3/8(1 - \epsilon')} \right) (1 - \epsilon') / 9} \right) \\ &\geq 1/4 \end{aligned}$$

for some small ϵ' .

Let X_j be the random variable which counts the number of functional nodes $j \in S_j$ which are adjacent to some node in S . We have that $E[X_j] \geq (1/4) \cdot 288(1 - \epsilon') \ln^2 n$, assuming that A holds. Let B be the event that $X_j \geq 72(1 - \epsilon'') \ln^2 n$, for some small $\epsilon'' > \epsilon'$. Since all events P_{Sx} , for fixed S and varying x , are independent, we use a Chernoff bound to show that, if event A occurs, event B occurs with probability $1 - 1/n^{k' \log n}$, for some positive constant k' .

Assume event $A \cap B$. Fix a functional node k in S_k . Fix a subset $S' \subseteq S_j$ of functional nodes, each of which is a neighbor of some element of S , with size $72(1 - \epsilon'') \ln^2 n$. Denote by $P_{S'k}$ the event that there exists a link between node k and some node in S' . This event occurs unless k has no link to any node in S' . Hence,

$$\begin{aligned} \Pr[P_{S'k}|B \cap A] &= \left(1 - (1 - 1/18 \ln n)^{72(1 - \epsilon'') \ln^2 n} \right) \\ &\geq 1 - e^{-72(1 - \epsilon'') \ln^2 n / (18 \ln n)} \\ &\geq 1 - n^{-4(1 - \epsilon'')}. \end{aligned}$$

Consider the event P_{uijk} that there exists a fault-free path of the form $uijk$ from a fixed node u to a fixed node k . Clearly, P_{uijk} is a subset of the event detailed in the above argument. Hence,

$$\begin{aligned} \Pr[P_{uijk}] &\geq \Pr[P_{S'k} \cap B \cap A] \\ &= \Pr[A] \Pr[B|A] \Pr[P_{S'k}|B \cap A] \\ &\geq (1 - n^{-3}) \left(1 - 1/n^{k' \log n} \right) \left(1 - n^{-4(1 - \epsilon'')} \right) \\ &\geq \left(1 - n^{-3 + \epsilon'''} \right), \end{aligned}$$

for some $0 < \epsilon''' < 0.1$.

There are at most n functional nodes in $RBTT$, each with at most $O(\log^2 n)$ functional nodes in supernodes at distance 3. Hence, there are at most $O(n \log^2 n)$ functional node pairs in supernodes at distance 3. It follows that all node pairs in supernodes at distance 3 are connected with probability at least $1 - n^{-1.9}$. ■

Combining the previous lemmas, we are now ready to prove Theorem 1.

Proof. of Theorem 1 The $RBTT$ contains $O(n/\log^2 n)$ supernodes connected in a binary-tree structure of diameter $D \in O(\log n)$. It follows from the construction that the maximum degree of the $RBTT$ is $O(\log n)$, with high probability. By Lemma 4, all functional node pairs in supernodes at distance 3 are connected by at least one fault-free path of length 3 with probability greater than $1 - 1/n^{1.9}$. Therefore, all functional nodes in the subgraph $RBTT'$ composed of the root supernode S_1 and of all supernodes at distances multiple of 3 from S_1 are connected with this probability. Clearly, functional nodes not in $RBTT'$ are in supernodes adjacent to supernodes in $RBTT'$. Thus, by Lemma 3, all these functional nodes are also connected to at least one functional node in $RBTT'$ with probability exceeding $1 - 1/n^{1.3}$. Hence, with probability exceeding $1 - 1/n^{1.8}$, the fault-free part of $RBTT$ is connected.

We now investigate the diameter of the fault-free part of $RBTT$. From the above argument, we observe that 1) nodes in supernodes at distances multiple of 3 are connected with high probability by a path of length equal to the distance of the supernodes; 2) functional nodes in all other supernodes are connected with high probability by a path of length at most 2 longer than the distance of the supernodes. This leads to the conclusion that the diameter of the fault-free part of $RBTT$ is also in $O(\log n)$, with high probability. ■

2.2 Lower Bounds

We have shown that it is possible to build a logarithmic-degree graph which supports communication with high probability in spite of spot probabilities $p \leq 1/c \log n$, for some positive constant c . The natural question then is whether it is possible to build arbitrarily large networks which can support communication with high probability despite larger spot probabilities. In what follows, we show that for constant spot probabilities, most networks do not have this property. More formally, we demonstrate Theorem 2.

Theorem 2 *For any constant spot probability $p > 0$, there exist constants c_1 and c_2 such that: if all degrees in a graph are at most $c_1 \log n$ then the fault-free part of the graph is disconnected with high probability; if all degrees in a graph are at least $c_2 \log n$ then the graph has all nodes faulty with high probability. In either case, highly reliable communication is not possible.*

In order to prove Theorem 2, we need several lemmas. First, we describe an event which implies that a small fault-free part of the graph is surrounded by faulty nodes, and bound its probability. We then define the size of a set of nodes for which these events are independent in the graph. With these lemmas, we bound the degree of graphs for which the fault-free part is not connected, with high probability. We then bound the degree of graphs such that all nodes are faulty, with high probability.

A node u is said to be *insular* if it is not faulty and all its neighbors who have neighbors at distance 2 from u are faulty. Denote by I_u the event that u is insular.

Lemma 5 *In a graph of maximum degree Δ , $\Pr[I_u] \geq (1 - p)^{\Delta+1} p^\Delta$, for spot probability p .*

Proof. An equivalent definition of event I_u is as follows: No node in $\Gamma_{\leq 1}(u)$ is a spot and, for each node $i \in \Gamma(u)$, either there exists a node $g \in \Gamma(i) \cap \Gamma_2(u)$ which is a spot, or $\Gamma(i) \cap \Gamma_2(u) = \emptyset$. In order to obtain a lower bound on the probability of event I_u , we now describe the construction of an elementary event $\varphi_u \subseteq I_u$: Sequentially, for each node $i \in \Gamma(u)$, mark one unmarked neighbor node which is not in $\Gamma_{\leq 1}(u)$ and label it m_i . If no such node exists, then mark nothing. Denote by \mathcal{M} the set of nodes marked by this process. Clearly, $|\mathcal{M}| \leq |\Gamma(u)| \leq \Delta$. The event φ_u occurs if all nodes in \mathcal{M} are spots and all nodes in $\Gamma_{\leq 1}(u)$ are not.

By definition, if all nodes in \mathcal{M} are spots, then each node $i \in \Gamma(u)$ for which there exists a node $m_i \in \mathcal{M}$, is faulty. Still by definition, all nodes $i \in \Gamma(u)$ for which there is no $m_i \in \mathcal{M}$ are adjacent to a $m_u \in \mathcal{M}$, $i \neq u$ or have no neighbor at distance 2 from u . Therefore, if all marked nodes are spots and all nodes in $\Gamma_{\leq 1}(u)$ are not spots, the event I_u is implied. See Figure 2.




Figure 2: Building an elementary event φ_u : for node a , no node is marked because its neighbors are in $\Gamma_{\leq 1}(u)$; for node b , node e is marked and labeled m_b ; for node c no node is marked because its neighbor outside $\Gamma_{\leq 1}(u)$, e , is already marked and labeled m_b . If m_b is a spot and u, a, b, c are not spots, then u is insular.

It follows from the preceding discussion that

$$\begin{aligned} \Pr[I_u] \geq \Pr[\varphi_u] &= (1-p)^{|\Gamma_{\leq 1}(u)|} p^{|\mathcal{M}|} \\ &\geq (1-p)^{\Delta+1} p^\Delta. \end{aligned}$$

■

Lemma 6 For all n -node graphs with maximum degree bounded above by Δ , there exist more than n/Δ^{2h+1} nodes c_i with disjoint h -hop neighborhoods $\Gamma_{\leq h}(c_i)$.

Proof. Take any graph whose maximum degree is bounded above by some Δ . Algorithm 1 constructs a sequence $\{c_0, c_1, \dots\}$ of nodes with disjoint h -hop neighborhoods $\Gamma_{\leq h}(c_i)$.

We now prove by induction that the sequence $\{c_0, c_1, \dots\}$ satisfies the statement of the lemma. For the sequence $\{c_0\}$, this is trivial. By inductive hypothesis, assume $\{c_0, c_1, \dots, c_i\}$ has been constructed such that there is no pair $c_j \neq c_k$ with intersecting sets $\Gamma_{\leq h}(c_j)$ and $\Gamma_{\leq h}(c_k)$. Since, by construction, all nodes less than $2h+1$ hops away from any node c_j , $j \leq i$, are marked and c_{i+1} is chosen from the unmarked node set, it is at least $2h+1$ hops away from

Algorithm 1 Construction

- 1: $i \leftarrow 0$
 - 2: $c_0 \leftarrow$ pick a node $v \in V$
 - 3: mark nodes in $\Gamma_{\leq 2h}(c_0)$
 - 4: **repeat**
 - 5: $i \leftarrow i + 1$
 - 6: $c_i \leftarrow$ pick an unmarked node adjacent to a marked node
 - 7: mark nodes in $\Gamma_{\leq 2h}(c_i)$
 - 8: **until** all $v \in V$ are marked
-

any c_j , $j < i + 1$. Furthermore, for any two nodes u, v with distance at least $2h + 1$, $\Gamma_{\leq h}(u)$ and $\Gamma_{\leq h}(v)$ are disjoint. Therefore, the inductive hypothesis holds for $\{c_0, c_1, \dots, c_{i+1}\}$.

Finally, for a graph whose maximum degree is at most Δ , less than Δ^{2h+1} nodes can be marked at every step of the construction. Therefore, the above construction defines a sequence C of more than $\frac{n}{\Delta^{2h+1}}$ nodes. \blacksquare

We now prove a bound on the degree of networks whose fault-free part is disconnected with high probability; thus they cannot support communication with high probability.

Lemma 7 *For any constant $c_1 < \frac{1}{\log(1/(p(1-p)))}$, any graph with maximum degree lower than $c_1 \log n$ has the fault-free part disconnected, with high probability, if spot probability p is a positive constant.*

Proof. In Lemma 5, we bound the probability of event I_u , that a fixed node u is insular. With an upper bound $c_1 \log n$ on the degree of the nodes,

$$\begin{aligned} \Pr[I_u] &\geq (1-p)^{c_1 \log n + 1} p^{c_1 \log n} \\ &\geq (1-p)n^{c_1 \log(1-p)} n^{c_1 \log p} \\ &\geq \frac{1-p}{n^{c_1 \log(1/(p(1-p)))}}. \end{aligned}$$

Let \mathcal{S} be a set of nodes in the graph for which the events I_u are independent. Since event I_u occurs if spots are located at distance 2 from u , the occurrence of this event is independent for nodes whose neighborhoods of radius at most 2 are disjoint. Furthermore, since all nodes have degree at most $c_1 \log n$, by Lemma 6, we have $|\mathcal{S}| \geq n/(c_1 \log n)^5$. Thus,

$$\begin{aligned} \Pr[\exists u \in V I_u] &= 1 - \Pr[\forall u \in V \neg I_u] \\ &\geq 1 - \Pr[\neg I_u]^{|\mathcal{S}|} \\ &\geq 1 - \left(1 - \frac{1-p}{n^{c_1 \log(1/(p(1-p)))}}\right)^{n/(c_1 \log n)^5} \\ &\geq 1 - e^{-\left(\frac{(1-p)n^{1-c_1 \log(1/(p(1-p)))}}{(c_1 \log n)^5}\right)} \end{aligned}$$

and, for $c_1 < \frac{1}{\log(1/(p(1-p)))}$, this event occurs with high probability.

We now fix an insular node $u \in \mathcal{S}$ and show that, in the set \mathcal{S} , there is at least one other node v which is fault-free. Nodes u and v are clearly not connected by a fault-free path. Let

$I(u)$ be the set of nodes connected to node u by a fault-free path, including node u . Let FF_v be the event that a node v is fault-free.

$$\begin{aligned}
\Pr[(\exists v \in V \setminus I(u)) FF_v] &= 1 - \Pr[(\forall v \in V \setminus I(u)) \neg FF_v] \\
&\geq 1 - \Pr[\neg FF_v]^{|S|-1} \\
&\geq 1 - \left(1 - (1-p)^{c_1 \log n + 1}\right)^{(1-\epsilon)n/(c_1 \log n)^5} \\
&= 1 - \left(1 - \frac{(1-p)}{n^{c_1 \log 1/(1-p)}}\right)^{(1-\epsilon)n/(c_1 \log n)^5} \\
&\geq 1 - e^{-\left(\frac{(1-p)(1-\epsilon)n^{1-c_1 \log(1/(1-p))}}{(c_1 \log n)^5}\right)}
\end{aligned}$$

and, for $c_1 < \frac{1}{\log(1/(p(1-p)))}$, this event occurs with high probability. It follows that the fault-free part of the graph is disconnected with high probability. \blacksquare

Lemma 8 *For any constant $c_2 > \frac{1}{\log(1/(1-p))}$, any graph with minimum degree at least $c_2 \log n$ has all nodes faulty with high probability, if spot probability p is a positive constant.*

Proof. Consider a node u with at least $c_2 \log n$ neighbors. Let F_u denote the event that u is faulty. F_u does not occur if none of u and its neighbors is a spot. Let F_V denote the event that all nodes in the graph are faulty. It follows that we can bound the probability of this event as

$$\begin{aligned}
\Pr[F_V] &= \Pr[\forall u \in V F_u] \\
&= 1 - \Pr[\exists u \in V \neg F_u] \\
&\geq 1 - \sum_{\forall u \in V} \Pr[\neg F_u] \\
&= 1 - n \cdot (1-p)^{c_2 \log n + 1} \\
&= 1 - n \cdot (1-p) \cdot n^{c_2 \log(1-p)} \\
&= 1 - \frac{(1-p)}{n^{c_2 \log(1/(1-p)) - 1}}.
\end{aligned}$$

Take $c_2 > \frac{1}{\log(1/(1-p))}$, then we have $c_2 \log(1/(1-p)) - 1 > 0$ and $\Pr[F_V]$ converges to 1 as n grows. \blacksquare

Lemmas 7 and 8, imply Theorem 2.

The preceding theorem leads to the conclusion that high probability communication is not possible, for a large class of graphs, when spot probability is a positive constant. However, the bounds $c_1 \log n$ and $c_2 \log n$ do not coincide. Since $c_1 < \frac{1}{\log(1/(p(1-p)))}$ and $c_2 > \frac{1}{\log(1/(1-p))}$, we have $c_1 < c_2$ for all positive values of p . It remains open whether or not there exists an arbitrarily large graph which supports reliable communication despite constant spot probabilities.

We will now attempt to provide insight into the question of what happens when node degrees lie between these bounds. For example, when $p = 1/2$, we have $c_1 < 1/2$ and $c_2 > 1$. Thus, with degree $\log n$, the important case of the n -node hypercube is not covered by Theorem 2. We will investigate this case in the following section.

2.3 Communication in the Hypercube

The hypercube H_k of dimension k is a 2^k -node graph with the set of nodes with identifiers from $\{0, 1\}^k$ and the set of links between nodes whose identifiers have a Hamming distance of 1. Hence the n -node hypercube H_k has dimension $\log n$. This section is dedicated to proving Theorem 3.

Theorem 3 *The n -node hypercube H_k does not support high probability communication for any constant spot probability $0 < p \leq 1$.*

We first show that for constant $0 < p < 1/2$, the fault-free part of the graph is disconnected with high probability. We then show that for $1/2 < p \leq 1$, the graph has all nodes faulty with high probability, and that for $p = 1/2$, the graph has all nodes faulty with constant probability. This will prove Theorem 3.

In order to prove that the fault-free part of the hypercube is disconnected for constant $0 < p < 1/2$, we first show that there exists at least one pair of functional nodes with some minimum constant distance in the hypercube. We will then show that there exists no fault-free path of sufficient length to connect this pair of nodes.

Lemma 9 *Assume constant spot probability $0 < p < 1/2$. For any constant c the n -node hypercube H_k contains at least one pair of fault-free nodes at distance at least c , with high probability.*

Proof. Consider event FF_u that a node u is fault-free. FF_u occurs if and only if there is no spot on nodes in $\Gamma_{\leq 1}(u)$. Furthermore, if two nodes u and v are at distance at least 3, then FF_u and FF_v are independent. On the hypercube, a result by Hamming implies that there exists a set \mathcal{S}_3 of nodes at distance at least 3 from each other with size $|\mathcal{S}_3| \in \Omega(\frac{2^k}{k})$. (Hamming codes are subsets of $\{0, 1\}^L$ of size $\Omega(\frac{2^L}{L})$ such that any two words have Hamming distance at least 3.) It follows from the preceding discussion that

$$\begin{aligned} \Pr[(\exists u \in V) FF_u] &= 1 - \Pr[\forall u \in V \neg FF_u] \\ &\geq 1 - (1 - \Pr[FF_u])^{|\mathcal{S}_3|} \\ &\geq 1 - \left(1 - (1-p)^{k+1}\right)^{\frac{2^k}{k}} \\ &= 1 - \left(1 - (1-p)(1-p)^{\log n}\right)^{\frac{n}{\log n}} \\ &= 1 - \left(1 - \frac{(1-p)}{n^{\log 1/(1-p)}}\right)^{\frac{n}{\log n}}. \end{aligned}$$

For $p < 1/2$, $\log 1/(1-p) < 1$ and therefore, we have at least one functional node with high probability.

Pick one functional node u and, w.l.o.g., label the nodes of H_k such that u has the label 0^k . Pick a positive integer c and consider the subset of nodes with labels $1^c\{0, 1\}^{k-c}$. These nodes induce a sub-hypercube H'_{k-c} of H_k with nodes at distance at least c from 0^k . Hamming's result applied on H'_{k-c} shows a set of nodes \mathcal{S}'_3 at distance at least 3 from each other with size $|\mathcal{S}'_3| \in \Omega(\frac{2^{k-c}}{k-c}) \in \Omega(\frac{n}{\log n})$. The above probabilistic argument applied to H'_{k-c} proves the lemma. ■

Lemma 10 *The fault-free part of the n -node hypercube H_k is disconnected with high probability, for constant spot probability $0 < p < \frac{1}{2}$.*

Proof. Pick a constant $c > 1/\log(1/(1-p))$. Since $0 < p < 1/2$, by Lemma 9, there is a pair (u, v) of functional nodes at distance $l \geq \lceil c \rceil$, with high probability. For nodes u and v to be connected, there must exist a fault-free path of length at least l between them.

Consider any path $x = x_1x_2 \dots x_l$ of length l in H_k . For the path x to be fault-free, all nodes x_1, x_2, \dots, x_l and all their neighbors must be spot-free. In any graph, only nodes at distance 2 or less can have common neighbors. In H_k , nodes at distance 2 share 2 neighbors; nodes at distance 1 are mutual neighbors and do not share other neighbors. Thus, for a path of length l to be fault-free, there must be more than $l(k-3)$ nodes which are spot-free. Furthermore, each node in H_k is an endpoint to, at most, $\log^l n$ paths of length l ; there are at most $n \log^l n$ paths of length l in H_k . Let \mathcal{P}_l be set of all paths of length l . Let P_x be the event that a path x is fault-free. It follows from the preceding discussion that

$$\begin{aligned} \Pr[(\exists x \in \mathcal{P}_l) P_x] &\leq \sum_{x \in \mathcal{P}_l} \Pr[P_x] \\ &\leq n \log^l n \cdot (1-p)^{l(k-3)} \\ &= \frac{\log^l n}{(1-p)^{3l}} \cdot n(1-p)^{l \log n} \\ &= \frac{\log^l n}{(1-p)^{3l}} \cdot n^{1-l \log(1/(1-p))} \\ &\rightarrow 0 \end{aligned}$$

since $l \geq \lceil c \rceil > \frac{1}{\log(1/(1-p))}$.

Since all paths of length greater than l include path segments of length l , these also contain faulty nodes with high probability. Thus, with high probability, no fault-free path can connect fault-free nodes u and v and the lemma follows. \blacksquare

It remains to see what occurs in the hypercube when spot probability is $p \geq 1/2$.

Lemma 11 *The n -node hypercube H_k has all nodes faulty with high probability, for spot probability $p > 1/2$, and all nodes faulty with constant probability, for spot probability $p = 1/2$.*

Proof. Consider event FF_u that a node u is fault-free. This event is equivalent to u and $\Gamma(u)$ being spot-free. For $p > 1/2$, the event that all nodes in H_k are faulty has probability

$$\begin{aligned} \Pr[(\forall u \in V) \neg FF_u] &\geq 1 - \sum_{u \in V} \Pr[FF_u] \\ &\geq 1 - n(1-p)^{k+1} \\ &\geq 1 - n(1-p)n^{\log(1-p)} \\ &\geq 1 - \frac{1-p}{n^{\log(1/(1-p))-1}}. \end{aligned}$$

Thus, for $p = 1/2$, $\log(1/(1-p)) = 1$ and all nodes are faulty with constant probability and, for $p > 1/2$, $\log(1/(1-p)) > 1$ and all nodes are faulty with high probability. \blacksquare

Lemmas 10 and 11, imply Theorem 3.

3 Bounded Degree Networks

The RBTT presented in Section 2 remained connected despite relatively high spot probabilities. However, its degree is unbounded. For certain applications, smaller-degree networks may be preferred as they are easier to implement and give shorter communication time in the 1-port model. Therefore, it is natural to ask if bounded-degree networks can also support high-probability communication with comparable spot probabilities.

In this section we construct bounded-degree networks which tolerate inverse polynomial spot probabilities and which support high-probability communication with optimal time complexity. Furthermore, we prove that bounded-degree networks can tolerate at most inverse polynomial spot probabilities.

3.1 Upper Bounds

We now study the properties of two networks: the torus and a torus-based tree-like network that we call the *toroidal tree*. We show that the torus supports high-probability communication for spot probability in $1/\omega(n^{1/2})$. However, the diameter of the torus is quite large, which prohibits fast communication. Thus we design a tree-like structure based on the torus which provides the same fault-tolerance properties and supports communication in time $O(\log n)$, even in the 1-port model.

3.1.1 The Torus

In this section, we show an upper bound on the spot probability such that the fault-free part of the torus remains connected. Denote by $\mathcal{T}_{m \times k}$ the $m \times k$ torus with $m, k \geq 4$. The torus has the set of nodes $\{u = (u_x, u_y) : u_x \in \{0, 1, \dots, m-1\}, u_y \in \{0, 1, \dots, k-1\}\}$ and the set of links $\{(u, v) : |u_x - v_x| \bmod m + |u_y - v_y| \bmod k = 1\}$.

Theorem 4 *The fault-free part of the n -node torus $\mathcal{T}_{m \times k}$ is connected with high probability for $p \in 1/\omega(n^{1/2})$.*

To prove Theorem 4, we show that, if spots are far enough from each other, fault-free node pairs are connected by paths which circumvent all faulty nodes.

Lemma 12 *In $\mathcal{T}_{m \times k}$, for each node u , there exists a cycle in the set $\Gamma_2(u) \cup \Gamma_3(u)$. This cycle includes all elements of $\Gamma_2(u)$.*

Proof. Consider the node with coordinates (x, y) . The following cycle around node (x, y) lies in the set $\Gamma_2(u) \cup \Gamma_3(u)$:

$$\begin{aligned} & \{(x, y + 2), (x + 1, y + 2), (x + 1, y + 1), (x + 2, y + 1), (x + 2, y), (x + 2, y - 1), \\ & (x + 1, y - 1), (x + 1, y - 2), (x, y - 2), (x - 1, y - 2), (x - 1, y - 1), (x - 2, y - 1), \\ & (x - 2, y), (x - 2, y + 1), (x - 1, y + 1), (x - 1, y + 2), (x, y + 2)\}. \end{aligned}$$

See Figure 3. Clearly, all nodes of $\Gamma_2((x, y))$ are included in this cycle. Since the torus is regular, the proof generalizes to all nodes of $\mathcal{T}_{m \times k}$. ■

Lemma 13 *The fault-free part of $\mathcal{T}_{m \times k}$ is connected if all spots are at distance at least 5 from each other.*




Figure 3: Cycle on $\Gamma_2((x, y)) \cup \Gamma_3((x, y))$

Proof. Pick any pair of spots $i \neq j$. Since $dist(i, j) \geq 5$, all nodes $u \in \Gamma_2(j)$ are at least at distance 3 from i and all nodes $u \in \Gamma_3(j)$ are at least at distance 2 from i . Hence, for all spots i , nodes $u \in \Gamma_2(i) \cup \Gamma_3(i)$ are functional. By Lemma 12, for each spot i , there exists a fault-free cycle \mathcal{C}_i included in $\Gamma_2(i) \cup \Gamma_3(i)$ which includes all nodes in $\Gamma_2(i)$.

Fix a spot i . Pick nodes $u, v \notin \Gamma_{\leq 1}(i), u \neq v$. We say that a path between functional nodes u and v is disconnected by spot i if the path contains at least one faulty node in $\Gamma_{\leq 1}(i)$. Let $\mathcal{P}_{(u,v),\Gamma_{\leq 1}(i)}$ be the set of paths disconnected by spot i between nodes u and v . We say that a path between functional nodes u and v circumvents spot i if it contains nodes in \mathcal{C}_i and excludes nodes in $\Gamma_{\leq 1}(i)$. Let $\mathcal{P}_{(u,v),\mathcal{C}_i}$ be the set of paths between nodes u and v which circumvent spot i . All paths in $\mathcal{P}_{(u,v),\Gamma_{\leq 1}(i)}$ have a prefix $u \dots a$ and a suffix $b \dots v$ of elements not in $\Gamma_{\leq 1}(i)$ such that the successor of a and the predecessor of b in the path are both in $\Gamma_{\leq 1}(i)$; elements a and b must be in $\Gamma_2(i)$. Similarly, all elements of $\mathcal{P}_{(u,v),\mathcal{C}_i}$ include a prefix $u \dots a$ and a suffix $b \dots v$ of elements such that a, b are respectively the first and last elements of $\mathcal{C}_i \cap \Gamma_2(i)$ in the path. Since \mathcal{C}_i includes all nodes in $\Gamma_2(i)$, for each path in $\mathcal{P}_{(u,v),\Gamma_{\leq 1}(i)}$, there is at least one path in $\mathcal{P}_{(u,v),\mathcal{C}_i}$ with matching prefix $u \dots a$ and a suffix $b \dots v$. Thus, there exists a function $f_i : \mathcal{P}_{(u,v),\Gamma_{\leq 1}(i)} \rightarrow \mathcal{P}_{(u,v),\mathcal{C}_i}$ which maps all paths disconnected by spot i to paths which circumvent spot i . See Figure 4.

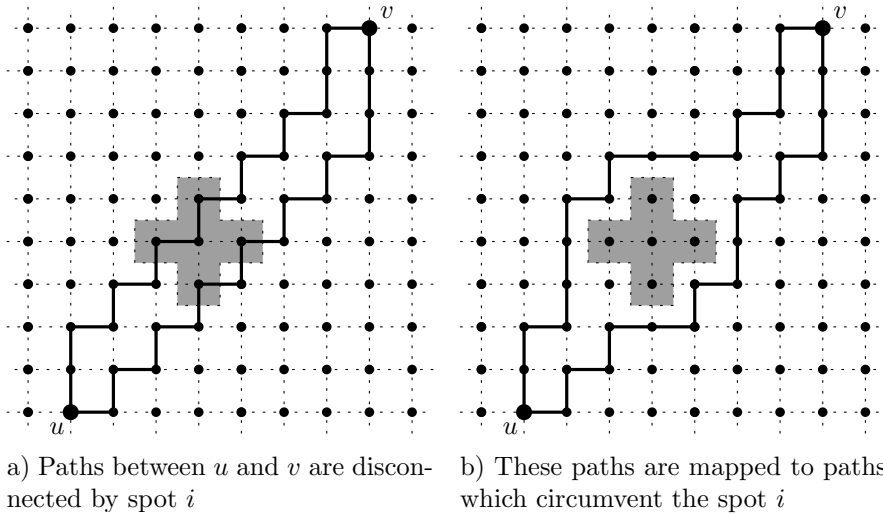


Figure 4: Path Mapping

Since we have shown that there is a fault-free cycle \mathcal{C}_i surrounding each spot i , the reasoning extends to any number of spots. In fact, since the sets $\Gamma_{\leq 1}(i)$ are disjoint for all spots i , the composition of all functions f_i maps all paths disconnected by spots to paths which circumvent all spots, for all functional pairs of nodes (u, v) , i.e. maps all paths which contain faulty nodes to fault-free paths, for all functional pairs of nodes (u, v) . See Figure 5. \blacksquare




Figure 5: Path Mapping

We are now ready to prove Theorem 4.

Proof. of Theorem 4 Let D_5 be the event that all spots are at distance at least 5 from each other. By Lemma 13, the event D_5 implies that the fault-free part of the torus is connected. Moreover, it has more than one element. Alternately, D_5 is the event that for all spots i , no set $\Gamma_{\leq 5}(i)$ contains a spot $j \neq i$. We note that $|\Gamma_{\leq 5}(i)| \leq 61$.

Let \mathcal{S} be the set of spots. Let $D_{5,i}$ be the event that for spot i , the set $\Gamma_{\leq 5}(i)$ contains no spot $j \neq i$. The event D_5 is the intersection of all $D_{5,i}$, $i \in \mathcal{S}$. For a fixed spot i , the probability of the event $D_{5,i}$ is $\Pr[D_{5,i}] = (1 - p)^{|\Gamma_{\leq 5}(i)|} > 1 - 61p$. Let the spot probability be $p \in 1/\omega(n^{1/2})$. Hence $p = 1/(n^{1/2}f(n))$, where $f(n) \rightarrow \infty$. Let A be the event that there are at most $2n^{1/2}$ spots. Using Chernoff bounds, with parameter $(1 + \epsilon) = 2f(n)$ we can show $\Pr[A] \geq 1 - ((e/(2f(n)))^{2n^{1/2}})$. It follows from the preceding discussion that

$$\begin{aligned}
\Pr[D_5] &\geq \Pr[A] \Pr[D_5 \mid A] \\
&= \Pr[A] \Pr \left[\bigcap_{i \in \mathcal{S}} D_{5,i} \mid A \right] \\
&= \Pr[A] \left(1 - \Pr \left[\bigcup_{i \in \mathcal{S}} \neg D_{5,i} \mid A \right] \right) \\
&\geq \Pr[A] (1 - 2n^{1/2} \Pr[\neg D_{5,i}]) \\
&> \Pr[A] \left(1 - 2n^{1/2} \left(1 - \left(1 - \frac{61}{n^{1/2}f(n)} \right) \right) \right)
\end{aligned}$$

$$\begin{aligned} &\geq \left(1 - \left(\frac{e}{2f(n)}\right)^{2n^{1/2}}\right) \left(1 - \frac{2 \cdot 61}{f(n)}\right) \\ &\rightarrow 1 \text{ as } n \rightarrow \infty. \end{aligned}$$

■

3.1.2 The Toroidal Tree

We now design a network which provides the same fault-tolerance as the torus, while also providing optimal-order communication time for bounded-degree graphs. Since the diameter of a bounded-degree graph is at least logarithmic, our aim is to construct a network whose fault-free part has logarithmic diameter. Such a network supports highly reliable communication in optimal time $O(\log n)$, even in the 1-port model. The network construction is based on two binary trees, T and T' , connected by a link between their root nodes. Each node of T, T' represents a group of nodes, and groups adjacent in T, T' have a subset of nodes in common. More precisely, for constant $k \geq 4$, we define a n -node graph $G(n, k)$. Assume that the set of nodes can be partitioned exactly as described below; this is easy to modify in the general case, by adding nodes to a leaf group.

Let the sets $\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_x$ and $\mathcal{T}'_1, \mathcal{T}'_2, \dots, \mathcal{T}'_{x'}$ be tori with $2k$ rows $\{0, 1, \dots, 2k - 1\}$ and k columns $\{0, 1, \dots, k - 1\}$; $|x - x'| \leq 1$. We describe the construction for the tree T ; the same construction is applied for the tree T' . Arrange all \mathcal{T}_i as the nodes of T , with $L + 1$ levels $0, 1, 2, \dots, L$, placing each \mathcal{T}_i on level $\lfloor \log i \rfloor$ of T . Level 0 contains the root of T and levels $L - 1$ and L contain the leaves. For every $1 \leq i \leq \lfloor x/2 \rfloor$, \mathcal{T}_{2i} is the left child of \mathcal{T}_i in T and \mathcal{T}_{2i+1} is the right child of \mathcal{T}_i in T (\mathcal{T}_{2i+1} exists if $x \geq 2i + 1$). For every $1 < i \leq x$, $\mathcal{T}_{\lfloor i/2 \rfloor}$ is the parent of \mathcal{T}_i . Use row 0 of each child torus to connect it to its parent in T . Use row k of each parent torus to connect it to both its children in T . Use row 0 of both roots in T, T' to connect them together. Connections between tori adjacent in T, T' are done by identifying the respective rows.

It follows from the above description that $x + x' = \lfloor (n - 2k^2) / ((2k - 1)k) \rfloor + 1$ tori are located on $L = \lfloor \log(x + x' + 1) \rfloor$ levels in $G(n, k)$. The graph $G(n, k)$ is called a Toroidal Tree. It has bounded maximal degree. See Figure 6.




Figure 6: The toroidal tree

Lemma 14 *The fault-free part of the Toroidal Tree is connected if all spots are at distance at least 5.*

Proof. Since the toroidal tree is build of $k \times 2k$ tori \mathcal{T}_i , $k \geq 4$, and since faults are at least at distance 5 from each other, by Lemma 13 the fault-free part of each \mathcal{T}_i is connected. Since $k \geq 4$, at least one common node is functional for each pair of adjacent tori. Therefore, the fault-free part of each adjacent pair of tori is connected and the lemma follows. ■

Theorem 5 For $p \in 1/\omega(n^{1/2})$, the n -node Toroidal Tree supports high probability communication in time $O(\log n)$.

Proof. To prove connectivity of the fault-free part of the graph, we repeat the proof of Theorem 4, substituting Lemma 13 with Lemma 14. Since in the toroidal tree $|\Gamma_{\leq 5}(i)| \in \Theta(1)$, connectivity follows.

Since every torus \mathcal{T}_i forming the Toroidal Tree has a constant number of nodes, the diameter of its fault-free part is bounded. Since these tori are arranged in a binary tree, it follows that the diameter of the fault-free part of the Toroidal Tree is $O(\log n)$. By construction, the Toroidal Tree has bounded maximum degree. ■

3.2 Lower Bounds

In this section, we show that bounded-degree graphs do not support high probability communication even for relatively small spot probabilities. We first show that the bounds on spot probability provided in Theorem 4 and Theorem 5 are tight for tori and toroidal trees. We then show that for general bounded-degree networks, if spot probability is the inverse of some polynomial, then high probability communication is not supported.

3.2.1 The Torus and Toroidal Tree

We first show a lower bound on spot probability such that the torus does not support high probability communication. Then, we extend the proof to show a similar lower bound for the toroidal tree. The following lower bounds match the upper bounds from Theorem 4 and Theorem 5, thus showing that the results are tight.

Theorem 6 For spot probability $p \in 1/O(n^{1/2})$, the n -node torus $\mathcal{T}_{m \times k}$ does not support high probability communication.

Proof. Let Com denote the event that the fault-free part of the n -node torus $\mathcal{T}_{m \times k}$ supports communication. In order for Com to occur, the fault-free part of the graph must be connected and have size at least 2. Let J_u be the event that a node u is functional and has 2 spots close to it, such that all neighbors of node u are faulty (see Figure 7). Clearly, $\Pr[J_u] \geq (1-p)p^2$. If an event J_u occurs then, either the fault-free part of the graph is disconnected, or the fault-free part of the graph has size less than 2; events J_u and Com are mutually exclusive. Furthermore, the events J_u are independent for nodes which are at distance 5 from each other. There exists a set \mathcal{S} of nodes at distance at least 5 from each other, such that $|\mathcal{S}| = kn$, for some constant k . Let $p \geq c/n^{1/2}$, with c a positive constant. It follows that

$$\begin{aligned} \Pr[Com] &\leq \Pr[\forall u \in V \neg J_u] \\ &\leq \Pr[\forall u \in \mathcal{S} \neg J_u] \\ &= (1 - \Pr[J_u])^{|\mathcal{S}|} \\ &\leq (1 - (1-p)p^2)^{kn} \end{aligned}$$




Figure 7: Two spots disconnect a torus

$$\begin{aligned}
 &\leq (1 - (1/2)c^2/n)^{kn} \\
 &\leq e^{-kc^2/2} \\
 &< 1.
 \end{aligned}$$

■

Theorem 7 *For spot probability $p \in 1/O(n^{1/2})$, the n -node Toroidal Tree does not support high probability communication.*

Proof. We use the argument from Theorem 6. Inside the set of constant size component tori forming the Toroidal Tree, there is a set \mathcal{S} of linear size with the same property as before. Thus the result follows. ■

3.2.2 General Bounded Degree Graphs

We showed in the preceding section that in the case of the torus and the Toroidal Tree, spot probabilities at most $1/\omega(n^{1/2})$ can be tolerated if these graphs support high-probability communication. In the following theorem, we show that a similar phenomenon occurs for all graphs whose degree is bounded by a constant.

Theorem 8 *For spot probability $p \in 1/O(n^{1/d})$, no n -node graph with degree bounded above by $d \in \Theta(1)$ supports high probability communication.*

Proof. Let $p \geq c/n^{1/d}$, with c a positive constant. Consider the event I_u that a node u is insular, as defined in Section 2. Clearly, a graph is disconnected if the event I_u occurs for some node u and that there is a functional node v which has no fault-free path which connects it to the node u . By Lemma 5, a fixed node u is insular with probability $\Pr[I_u] \geq (1-p)^{d+1}p^d$. From Lemma 6 we know that any bounded degree graph has at least n/d^{2h+1} nodes with disjoint h -radius neighborhoods. Since event I_u depends only on the status of nodes at distance 2 from u , the occurrence of these events is independent for nodes u which have radius 2 disjoint neighborhoods; i.e. for $h = 2$. It follows that

$$\begin{aligned}
 \Pr[\exists u \in V I_u] &= 1 - \Pr[\forall u \in V \neg I_u] \\
 &\geq 1 - \Pr[\forall u \in \mathcal{S} \neg I_u] \\
 &= 1 - (1 - \Pr[I_u])^{|\mathcal{S}|} \\
 &\geq 1 - \left(1 - (1-p)^{d+1}p^d\right)^{\frac{n}{d^5}}
 \end{aligned}$$

$$\begin{aligned}
&\geq 1 - \left(1 - \frac{1}{2} \left(cn^{-1/d}\right)^d\right)^{\frac{n}{d^5}} \\
&= 1 - \left(1 - \frac{1}{2} \cdot \frac{c^d}{n}\right)^{\frac{n}{d^5}} \\
&> 1 - e^{-\frac{c^d}{2d^5}}.
\end{aligned}$$

We now fix an insular node $u \in \mathcal{S}$ and show that, in the set \mathcal{S} , there is at least one other node v which is fault-free. Nodes u and v are clearly not connected by a fault-free path. Let $I(u)$ be the set of nodes connected to node u by a fault-free path, including node u . Let FF_v be the event that a node v is fault-free. It follows that

$$\begin{aligned}
\Pr[(\exists v \in V \setminus I(u)) FF_v] &= 1 - \Pr[(\forall v \in V \setminus I(u)) \neg FF_v] \\
&\geq 1 - \Pr[(\forall v \in \mathcal{S} \setminus u) \neg FF_v] \\
&= 1 - \Pr[\neg FF_v]^{|\mathcal{S}|-1} \\
&\geq 1 - \left(1 - (1-p)^{d+1}\right)^{(1-\epsilon)\frac{n}{d^5}} \\
&\geq 1 - \left(1 - \left(1 - cn^{-1/d}\right)^{d+1}\right)^{(1-\epsilon)\frac{n}{d^5}} \\
&\geq 1 - \left((d+1)cn^{-1/d}\right)^{(1-\epsilon)\frac{n}{d^5}} \\
&\rightarrow 1.
\end{aligned}$$

Hence the fault-free part of the graph is disconnected with positive constant probability and the lemma follows. ■

4 Conclusion

We provided what is, to the best of our knowledge, the first set of analytic results related to fault-tolerance of networks in the presence of dependent, positively correlated faults. To do so, we introduced the *neighborhood fault* model where damaging events, called *spots*, occur randomly and independently at nodes of a network with probability p , and cause faults in the affected node and its neighbors.

We addressed questions regarding the connectivity and diameter of the fault-free part of networks in this fault model, as these characteristics of the network are responsible for the feasibility of communication and for its time. Our results show clear differences between the assumption of independent faults and that of the neighborhood fault model. For example, while under independent faults with small constant fault probability $p > 0$ the fault-free part of the hypercube remains connected with high probability [4], this is not the case under the neighborhood fault model with any positive constant spot probability. Likewise, the fault-free part of the torus is connected with high probability for fault probability $p \in 1/\Omega(n^{1/4})$ when faults are independent, but this is not the case for such spot probabilities under the neighborhood fault model.

It remains open whether or not there exists a network, which, under the neighborhood fault model, has the fault-free part connected with high probability despite constant spot probabilities. Our results support the conjecture that this is not the case.

The neighborhood fault model is the first step in modeling dependent positively correlated faults in networks. It would be interesting to analyze more precise center-satellite based models in which independent spots yield faults in nodes with probability decreasing with the distance of the node from the spot.

References

- [1] D. Bienstock. Broadcasting with random faults. *Discrete Appl. Math.*, 20(1):1–7, 1988.
- [2] D. M. Blough and A. Pelc. A clustered failure model for the memory array reconfiguration problem. *IEEE Trans. Computers*, 42(5):518–528, 1993.
- [3] B. S. Chlebus, K. Diks, and A. Pelc. Sparse networks supporting efficient reliable broadcasting. *Nordic Journal of Computing*, 1:332–345, 1994.
- [4] B. S. Chlebus, K. Diks, and A. Pelc. Reliable broadcasting in hypercubes with random link and node failures. *Combinatorics, Probability and Computing*, 5:337–350, 1996.
- [5] A. Choi, N. Park, F. J. Meyer, F. Lombardi, and V. Piuri. Reliability measurement of fault-tolerant onboard memory system under fault clustering. In *Proceedings of 19th Instrumentation and Measurement Technology Conference, 2002. IMTC*, volume 2, pages 1161–1166. IEEE Computer Society, 2002.
- [6] D. Ganesan, R. Govindan, S. Shenker, and D. Estrin. Highly-resilient, energy-efficient multipath routing in wireless sensor networks. *ACM SIGMOBILE Mobile Computing and Communications Review*, 5(4):11–25, 2001.
- [7] T. Hagerup and C. Rüb. A guided tour of Chernoff bounds. *Inf. Proc. Letters*, 33:305–308, 1989/90.
- [8] B. Jang, A. Choi, N. Park, Y. B. Kim, V. Pliuri, and F. Lombardi. Spare line borrowing technique for distributed memory cores in soc. In *Proceedings of 22nd Instrumentation and Measurement Technology Conference, 2005. IMTC*. IEEE Computer Society, May, 17–19 2005.
- [9] F. J. Meyer and N. Park. Predicting defect-tolerant yield in the embedded core context. *IEEE Trans. Computers*, 52(11):1470–1479, 2003.
- [10] F. J. Meyer and D. K. Pradhan. Modeling defect spatial distribution. *IEEE Trans. Computers*, 38(4):538–546, 1989.
- [11] R. Motwani and P. Raghavan. *Randomized Algorithms*. Cambridge University Press, 1995.
- [12] M. Paquette and A. Pelc. Fast broadcasting with byzantine faults. *International Journal of Foundations of Computer Science*, 17(6):1423–1439, 2006.
- [13] A. Pelc. Fault-tolerant broadcasting and gossiping in communication networks. *Networks*, 28(6):143–156, 1996.

- [14] S. Shakkottai, R. Srikant, and N. Shroff. Unreliable sensor grids: Coverage, connectivity and diameter. In *INFOCOM 2003. Twenty-Second Annual Joint Conference of the IEEE Computer and Communications Societies*, volume 2, pages 1073–1083, March–April 2003.
- [15] C. H. Stapper. On yield, fault distributions and clustering of particles. *IBM Journal of Research and Development*, 30(3):326–338, 1986.
- [16] M. Thottan and C. Ji. Using network fault predictions to enable IP traffic management. *J. Network Syst. Manage*, 9(3):327–346, 2001.
- [17] W. Warren. The center-satellite concept as a basis for ecological sampling. *Stat. Ecol.*, 2:87–118, 1971.
- [18] M. Yajnik, J. Kurose, and D. Towsley. Packet loss correlation in the Mbone multicast network, May 27 1996.
- [19] F. Yu, C.-H. Tsai, Y.-W. Huang, D. T. Lee, H.-Y. Lin, and S.-Y. Kuo. Efficient exact spare allocation via boolean satisfiability. In *20th IEEE International Symposium on Defect and Fault Tolerance in VLSI Systems (DFT'05)*, pages 361–370. IEEE Computer Society, 2005.