

Communication in Random Geometric Radio Networks with Positively Correlated Random Faults

Evangelos Kranakis¹, Michel Paquette¹, and Andrzej Pelc²

¹ School of Computer Science, Carleton University, Ottawa, Ontario, K1S 5B6, Canada. kranakis@scs.carleton.ca, michel.paquette@polymtl.ca

² Département d'informatique et d'ingénierie, Université du Québec en Outaouais, Gatineau, Québec, J8X 3X7, Canada. pelc@uqo.ca

Abstract. We study the feasibility and time of communication in random geometric radio networks, where nodes fail randomly with positive correlation. We consider a set of radio stations with the same communication range, distributed in a random uniform way on a unit square region. In order to capture fault dependencies, we introduce the *ranged spot* model in which damaging events, called *spots*, occur randomly and *independently* on the region, causing faults in all nodes located within distance s from them. Node faults within distance $2s$ become dependent in this model and are positively correlated. We investigate the impact of the spot arrival rate on the feasibility and the time of communication in the fault-free part of the network. We provide an algorithm which broadcasts correctly with probability $1 - \epsilon$ in faulty random geometric radio networks of diameter D in time $O(D + \log 1/\epsilon)$.

Keywords: Fault-tolerance, dependent faults, broadcast, crash faults, random, geometric radio network.

1 Introduction

Wireless networks have received much attention in recent years because of applications where wired networks are impractical or impossible to deploy. These networks are now so common that the idea of large scale wireless networks has become natural. However, as they grow in size, complexity, and area, wireless networks become increasingly vulnerable to component failures and damaging environmental phenomena. Nodes of a network may fail and the communication medium may become too noisy to support correct message transmissions. These failures often result in delaying, blocking, or even distorting transmitted messages. Hence, it becomes important that the desired tasks may be accomplished efficiently in spite of these faults, usually without knowing their location ahead of time. Networks with this property are called fault-tolerant.

An important type of wireless network is obtained from a set of stations in the plane where each station u has communication range r_u . The resulting network is modeled as a directed graph in which stations are nodes and a directed edge

from u to v exists if v is at distance at most r_u from u . Such networks are called geometric radio networks (GRN).

One of the most important communication tasks is broadcasting. In this process, a source node attempts to transmit a message to all other nodes of the network. This process is successful if, upon termination, all functional nodes, connected to the source by a fault-free path, have received the source message. Although the question of fault-tolerant broadcasting has been widely studied for faulty point-to-point networks, few results are known about this process in geometric radio networks. To the best of our knowledge, all existing analytic results examine the general problem of broadcasting in networks where the number of faults is bounded above (cf., e.g., [8]), or faults are distributed randomly and independently (cf., e.g., [9, 14]). Hence, the present paper is the first to address the problem of broadcasting in GRNs in the presence of positively correlated faults.

1.1 Model and Problem Definitions

We seek to model a network composed of mobile stations moving under the *Random Waypoint* mobility model [7] inside an open region, e.g., a train station or a plaza. Under this mobility model, mobile stations alternately move and pause for random amounts of time, choosing a direction, distance and speed at random at every movement phase. Here, we assume that the mobile stations move at low, pedestrian-like speeds, making the network appear static for the short duration of communication processes; the distance and directions are chosen in some uniform way. We further assume that neither the boundaries of the open region nor the other mobile stations have any effect on the mobile station movements. Hence, any snapshot of the graph is a set of stations distributed on a plane by a Poisson process. Due to the short duration of the communication processes, we consider that the faults are permanent. The proposed static model also applies to networks of sensors spread randomly in hostile environments where individual placement and replacement of units is not possible.

We focus attention on a unit square region of the plane. Node locations occur with Poisson arrival rate n . We fix a parameter r , called the communication range. Any two nodes at Euclidean distance at most r from one another can communicate directly. We now define the *ranged spot* fault model. Damaging phenomena, called *spots*, occur on the plane with Poisson arrival rate λ . Some examples of damaging phenomena are lightning strikes (and other electrostatic discharges), electromagnetic pulses and explosions. We fix a parameter $s > 0$, called the *spot range*. Each spot causes permanent crash faults in all nodes within distance s of it, i.e., inside the disk of radius s centered at it, which we call the *spot disk*. For a fixed spot i , we denote the corresponding spot disk by D_i . Faulty nodes can neither send nor receive messages for the entire communication process. More formally, consider the undirected fault-free graph $G(V, E)$, where V is the random set of nodes occurring on the unit square with Poisson arrival rate n , and E is the set of all node pairs $\{u, v\}$ for which the Euclidean distance is at most r . Let S be the set of spots which occur on the unit square with

Poisson arrival rate λ . Let F be the set of faulty nodes, i.e., all the nodes in V whose location is within distance s from at least one spot in S . We consider the graph $G[V']$ induced on G by the set $V' = V \setminus F$ of all functional nodes. To remind the reader how it is built, throughout this paper, we will denote the graph $G[V']$ by $U(n, r, \lambda, s)$.

As usual in wireless network algorithms, communication in $U(n, r, \lambda, s)$ is assumed synchronous; nodes have synchronized clocks and the communication process is executed in fixed time steps, called *rounds*. All communication is done using the same base frequency, modulation and encoding, hence using a single channel. In each round, each node either sends a message or listens to the channel. In the first case, we say that the node is a sender, otherwise, it is a receiver. In a fixed round t , a node v receives a message if and only if it is a receiver and precisely one of its neighbors is a sender. If no neighbor of v is a sender, then there is no message on the channel which v can receive. If more than one neighbor of v sends a message, we say that a *collision* occurs at v and v can only perceive noise on the channel. Nodes do not have collision detection abilities, i.e., they cannot distinguish collision noise from background noise (which is apparent when no messages are heard).

We say that an event occurs in the graph with high probability (w.h.p.) if its probability converges to 1 as the node arrival rate n grows to infinity. We say that an event occurs on the graph with constant (positive) probability if its probability p is bounded away from 0 and from 1 as n grows to infinity, i.e., if there exist positive constants ϵ_1, ϵ_2 such that $0 < \epsilon_1 < p < \epsilon_2 < 1$ for all n . Specifically, we say that a graph is connected w.h.p. when the event that it is connected occurs w.h.p. On the other hand, we say that a graph is not connected w.h.p. when the event that it is disconnected occurs at least with constant probability.

In this paper, we study the question of feasibility and efficiency of communication in the fault-free graph $U(n, r, \lambda, s)$.

1.2 Our Results

We first give answers to the question for which parameters $s = s(n)$ and $\lambda = \lambda(n, s)$ there exist any fault-free nodes in the unit square, i.e., when the fault-free graph $U(n, r, \lambda, s)$ is non-empty, w.h.p. For $s \in o(1)$, we find a threshold function $l(n, s)$ and constants L_1, L_2 such that, for $\lambda \geq L_1 \cdot l(n, s)$ fault-free nodes do not exist, w.h.p., while for $\lambda \leq L_2 \cdot l(n, s)$ they do exist w.h.p. For $s \in \Omega(1)$, we show that, for $\lambda \in \omega(1/s^2)$ fault-free nodes do not exist, w.h.p., and for $\lambda \in o(1/s^2)$ they do exist, w.h.p.

We then give answers to the question for which parameters $s = s(n)$, $r = r(n)$ and $\lambda = \lambda(n, s, r)$, the fault-free graph $U(n, r, \lambda, s)$ is connected, w.h.p. Connectivity is equivalent to feasibility of communication in our setting. We restrict attention to the case of small spot range, more precisely, we work under the assumption $s \in o(r)$. In the case $r \in o(1)$, we find a threshold function $c(n, s, r)$ and constants C_1, C_2 such that, for $\lambda \geq C_1 \cdot c(n, s, r)$ the graph $U(n, r, \lambda, s)$ is not connected w.h.p., and for $\lambda \leq C_2 \cdot c(n, s, r)$ it is connected, w.h.p. Then, in the

case $r \in \Omega(1)$, and for $\lambda \in o(1/s^2)$, we show that for the values of λ for which the graph $U(n, r, \lambda, s)$ contains at least one node w.h.p., it is also connected w.h.p.

Finally, under the additional restriction on spot range, when $s \in o(1/\sqrt{n})$, we show an algorithm which accomplishes broadcast with probability at least $1 - \epsilon$ in time $O(D + \log 1/\epsilon)$ in the graph $U(n, r, \lambda, s)$ of diameter D .

Due to lack of space, the proofs of several lemmas and theorems are omitted.

1.3 Related Work

The fundamental questions of network reliability have received much attention in the context of point-to-point networks, under the assumption that components fail randomly and independently (cf., e.g. [1–3, 11] and the survey [12]). On the other hand, empirical work has shown that positive correlation of faults is a more reasonable assumption for networks [6, 15, 16]. In particular, in [16], the authors provide empirical evidence that data packets losses are spatially correlated in networks. Moreover, in [6], the authors simulate failures in a sensor network using a model similar to that of the present paper; according to these authors, the environment provides many spatially correlated phenomena resulting in such fault patterns. More recently, in [10], a gap was demonstrated between the fault-tolerance of networks when faults occur independently as opposed to when they occur with positive correlation. To the best of our knowledge, this was the first paper to provide analytic results concerning network fault-tolerant communication in the presence of positively correlated faults.

In contrast, few results are known about fault-tolerant communication in geometric radio networks. To the best of our knowledge, all existing analytic results examine the problem of broadcasting in networks where, either the number of faults is bounded above (cf., e.g., [8]), or faults occur randomly and independently (cf., e.g., [9, 14]). In particular, in [14], the authors consider the problem of connectivity of a square grid of n sensors with communication range r on a unit square when faults occur at the nodes randomly and independently with probability $1 - p$. They show that if $pr^2 \approx \frac{\log n}{n}$, then the functional nodes are all part of a connected component w.h.p. In [8], the authors consider the problem of broadcasting in a fault-free connected component of a radio network whose nodes are located at grid points of square grids and can communicate within a square of size r . For an upper bound t on the number of faulty nodes, in worst-case location, the authors propose a $\Theta(D+t)$ -time oblivious broadcast algorithm and a $\Theta(D + \log(\min(r, t)))$ -time adaptive broadcast algorithm, both operating on a connected fault-free component of diameter D .

The question of communication in networks of unknown topology has been widely studied in recent years. In fact, in [4], the authors state that broadcasting algorithms which function in unknown GRNs also function in the resulting fault-free connected components of faulty GRNs. A basic performance evaluation criterion of broadcasting algorithms is the time necessary for the algorithm to terminate; in synchronous networks, this time is measured as the number of communication rounds. For networks whose fault-free part has a diameter D , $\Omega(D)$ is a trivial lower bound on broadcast time, but optimal running time

is a function of the information available to the algorithms (cf., e.g., [5]). For instance, in [5], an algorithm was obtained which accomplishes broadcast in arbitrary GRNs in time $O(D)$ under the assumption that nodes have a large amount of knowledge about the network, i.e. given that all nodes have a *knowledge radius* larger than R , the largest communication radius. The authors also show that algorithms broadcasting in time $O(D + \log n)$ are asymptotically optimal, for unknown GRNs when nodes can communicate spontaneously (before receiving the source message) and either can detect collisions or have knowledge of node locations at some positive distance δ , arbitrarily small. In the present paper, we assume that nodes communicate spontaneously, but know nothing of the network, other than their own location, and cannot detect collisions. Under these assumptions, we show an $O(D + \log 1/\epsilon)$ -time algorithm which correctly broadcasts in the random graph $U(n, r, \lambda, s)$ with probability at least $1 - \epsilon$.

2 Liveness of the Graph

In this section, we show bounds on the spot arrival rate λ for which functional nodes exist in the unit square, i.e., the graph $U(n, r, \lambda, s)$ contains at least one node, w.h.p. We say that the graph $U(n, r, \lambda, s)$ is *alive* if it contains at least one node; otherwise, we say that it is *dead*.

Theorem 1. *For $s \in o(1)$, there exist two positive constants, L_1 and L_2 , such that if the spot arrival rate $\lambda \geq \frac{L_1 \ln(\min\{n, 1/s^2\})}{s^2}$, then the graph $U(n, r, \lambda, s)$ is dead, w.h.p., and if $\lambda \leq \frac{L_2 \ln(\min\{n, 1/s^2\})}{s^2}$, then $U(n, r, \lambda, s)$ is alive, w.h.p.*

Theorem 2. *For $s \in \Omega(1)$, the graph $U(n, r, \lambda, s)$ is dead w.h.p. if $\lambda \in \omega(1/s^2)$ and alive w.h.p. if $\lambda \in o(1/s^2)$.*

Remark 1. For $s \in o(1/\sqrt{n})$ and $\lambda = \frac{\ln(cn)}{\pi s^2}$, where c is a positive constant, the graph $U(n, r, \lambda, s)$ is dead with constant probability.

3 Connectivity of $U(n, r, \lambda, s)$

In the preceding section, we gave a threshold for the spot arrival rate for which the graph $U(n, r, \lambda, s)$ is non-empty w.h.p. We now answer the next natural question: for which spot arrival rate is the graph $U(n, r, \lambda, s)$ connected w.h.p.?

It has been shown, in [13], that for any real number c , if $r \geq \sqrt{\frac{\ln n + c}{\pi n}}$, then the probability that the graph $U(n, r, \lambda, s)$, with $\lambda = 0$, is connected is at least $e^{-e^{-c}}$, as $n \rightarrow \infty$. If we substitute $e^{-c} = f(n)$, assume that $f(n) \in o(1)$ and recall that $e^{-f(n)} = 1 - f(n) + f(n)^2/2 - \dots \geq 1 - f(n)$, then we see that if

$$r \geq \sqrt{\frac{\ln n + \ln 1/f(n)}{\pi n}}$$

then

$$\Pr[U(n, r, 0, s) \text{ is connected}] \geq 1 - f(n).$$

Hence, it is natural to investigate the connectivity of the graph $U(n, r, \lambda, s)$ under the assumption $r^2 \geq \frac{\ln n + \ln 1/f(n)}{\pi n}$, for some $f(n) \in o(1)$, when we know that connectivity is guaranteed w.h.p. without faults. In what follows we make this assumption.

The main results of this section are Theorems 3 and 4. In Theorem 3, for spot range s of lower order of magnitude than the communication range r and for $r \in o(1)$, we show a threshold for the spot arrival rate λ below which the graph $U(n, r, \lambda, s)$ is connected w.h.p. and above which it is not. For the case $r \in \Omega(1)$, the separation is different: in Theorem 4, we show thresholds for the spot arrival rate λ below which the graph $U(n, r, \lambda, s)$ is connected w.h.p. and above which it is *dead* w.h.p.

Theorem 3. *For $s \in o(r)$ and $r \in o(1)$, there exist two positive constants, C_1 and C_2 , such that if spots appear with arrival rate $\lambda \geq C_1 \ln \left(\frac{r^2 \min\{n, 1/s^2\}}{\ln(1/r^2)} \right) / s^2$, then the graph $U(n, r, \lambda, s)$ is not connected w.h.p., and if the spot arrival rate $\lambda \leq C_2 \ln \left(\frac{r^2 \min\{n, 1/s^2\}}{\ln(1/r^2)} \right) / s^2$, then the graph $U(n, r, \lambda, s)$ is connected, w.h.p.*

Theorem 3 will follow from Lemmas 2, 3, 4, and 5.

Theorem 4. *For $s \in o(r)$ and $r \in \Omega(1)$,*

1. *if $s \in o(1)$, then there exist two positive constants, C_3 and C_4 , such that*
 - (a) *for $\lambda \leq C_3 \ln(\min\{n, 1/s^2\})/s^2$, $U(n, r, \lambda, s)$ is connected, w.h.p.,*
 - (b) *for $\lambda \geq C_4 \ln(\min\{n, 1/s^2\})/s^2$, the graph $U(n, r, \lambda, s)$ is dead w.h.p.,*
2. *if $s \in \Omega(1)$, then*
 - (a) *for $\lambda \in o(1/s^2)$, the graph $U(n, r, \lambda, s)$ is connected, w.h.p.*
 - (b) *for $\lambda \in \omega(1/s^2)$, the graph $U(n, r, \lambda, s)$ is dead w.h.p.,*

Theorem 4 will follow from Theorems 1 and 2 and from Lemmas 6 and 7.

3.1 Non-Connectivity Results

In this section, we show conditions on spot arrival rate implying, w.h.p., non-connectivity of the graph $U(n, r, \lambda, s)$ by the existence of two functional nodes which cannot communicate with one another in the unit square.

Denote by \mathcal{P}_{left} and \mathcal{P}_{right} the two rectangular halves of the unit square. Partition \mathcal{P}_{left} and \mathcal{P}_{right} respectively into meshes of $r \times r$ squares. Group these squares in matrices of 5×5 squares, called blocks; let \mathcal{B}_{left} and \mathcal{B}_{right} be the sets of these blocks. For each block b , denote by c_b the central square in this block and by p_b the union of 8 squares adjacent to c_b . Let *alive* _{b} be the event that c_b contains at least one functional node. Let *surround* _{b} be the event that p_b contains no functional node. Let *isolation* _{b} be the intersection of events *surround* _{b} and *alive* _{b} . If *isolation* _{b} occurs, and there is at least one functional node outside b , then nodes in c_b have no functional path to this external functional node,

and then, the graph $U(n, r, \lambda, s)$ is disconnected. In particular, we show non-connectivity w.h.p. by proving that events $isolation_{b_1}$ and $isolation_{b_2}$, $b_1 \in \mathcal{B}_{left}$ and $b_2 \in \mathcal{B}_{right}$, occur w.h.p. Note that, for distinct blocks b_1 and b_2 , events $surround_{b_1}$ and $surround_{b_2}$ are independent.

We first examine non-connectivity in the case when $r \in o(1)$ and $s \in o(1/\sqrt{n})$, in Lemma 2. Non-connectivity for $r \in o(1)$ and for larger values of $s \in o(r)$ will be addressed in Lemma 3. The case $s \in \Omega(1)$ is treated in the next section. We show that for these values of r , the graph is connected w.h.p. for those spot arrival rates for which it is alive w.h.p.

Let F_v be the event that a fixed node v is faulty, i.e., that there exists at least one spot within distance s of it. Then, for spot arrival rate λ we have

$$\Pr[F_v] = 1 - e^{-\lambda\pi s^2}.$$

While distant faults are independent, the presence of a faulty node within distance $2s$ from a fixed node v implies that there is a spot which might be close enough to v to make it faulty, i.e., the occurrence of a fault at a node can never decrease the probability of a fault on another node. This is why faults are positively correlated. Hence, the following fact applies to the events F_v .

Fact 1 *For any set Z of nodes,*

$$\Pr\left[\bigcap_{v \in Z} F_v\right] \geq \prod_{v \in Z} \Pr[F_v].$$

A set S of nodes whose elements have a distance greater than $2s$ from one another is called *sparse*. Such a set has the property that the events F_v , for $v \in S$, are independent. The following lemma states that there exist large sparse sets, w.h.p.

Lemma 1. *A square A with area $|A|$ contains a sparse set S of size at least $k|A| \min\{n, 1/s^2\}$, for some positive constant k , w.h.p., if $|A| \min\{n, 1/s^2\} \rightarrow \infty$ as $n \rightarrow \infty$.*

Lemma 2. *Fix any constants $\alpha > 8$ and $\beta > 1$. For $s \in o(1/\sqrt{n})$ and $r \in o(1)$, the graph $U(n, r, \lambda, s)$ is not connected w.h.p. when $\lambda = \beta \ln\left(\frac{\alpha nr^2}{\ln(1/r^2)}\right) / \pi s^2$.*

Proof. Consider $f(n) \in \omega(1)$ and the set Λ of spot arrival rates of the form $\lambda = \ln\left(\frac{\alpha r^2 n}{\ln(1/(r^2 f(n)))}\right) / (\pi s^2)$. Consider two subsets of Λ : Λ_1 consisting of these λ of the form $\lambda = \ln(g(n)r^2 n) / (\pi s^2)$, with $g(n) \in O(1)$ and Λ_2 consisting of these λ of the same form with $g(n) \in \Omega(1)$. In each case, we show that there exists at least one occurrence of the event $isolation_b$ in each set \mathcal{B}_{left} and \mathcal{B}_{right} and thus, that the graph $U(n, r, \lambda, s)$ is disconnected.

Case 1: $\lambda = \ln(g(n)r^2 n) / (\pi s^2)$, with $g(n) \in O(1)$

Fix a block b and consider the event $alive_b$. Consider the subsquare $c'_b \subset c_b$ whose points are at distance greater than $2s$ from p_b , i.e. for which the contained nodes become faulty independently from nodes in p_b . For $s/r \rightarrow 0$ as $n \rightarrow \infty$,

$|c'_b| > 0.9r^2$, for large n . From Lemma 1, since $0.9nr^2 > 0.9 \log n \in \omega(1)$, it follows that, w.h.p., there is a sparse set of nodes S_b , in c'_b , of size at least knr^2 , for some positive constant k . Events F_v , $v \in S_b$, occur independently. Let A be the event that the above lower bound on the size of the sparse set S_b holds. Assume A . Then,

$$\begin{aligned} \Pr[alive_b] &= 1 - \Pr[\forall v \in c_b F_v] \geq 1 - \Pr[\forall v \in c'_b F_v] \\ &\geq 1 - \Pr[\forall v \in S_b F_v] \geq 1 - (\Pr[F_v])^{knr^2} = 1 - (1 - e^{-\lambda\pi s^2})^{knr^2} \\ &= 1 - (1 - e^{-\ln(g(n)r^2n)})^{knr^2} \\ &= 1 - (1 - 1/g(n)r^2n)^{knr^2} \geq c' \in \Theta(1) \end{aligned}$$

since $g(n) \in O(1)$. Since $\Pr[A] \rightarrow 1$ for large n , this implies that the probability of event $alive_b$ is at least a positive constant c . Let \mathcal{A}_{left} be the set of all blocks b in \mathcal{B}_{left} for which the event $alive_b$ occurs. Since the probability of the event $alive_b$ is a constant, the expected size of the set \mathcal{A}_{left} is a constant fraction of $|\mathcal{B}_{left}|$. The number of blocks in \mathcal{B}_{left} is $|\mathcal{B}_{left}| = 1/50r^2$. For $r \in o(1)$, $|\mathcal{B}_{left}|$ grows to infinity as $n \rightarrow \infty$ and thus, under the preceding assumptions, we use Chernoff bounds to show that $|\mathcal{A}_{left}| \geq (0.9)c/(50r^2)$ w.h.p. Assume this bound on $|\mathcal{A}_{left}|$ and let $k/r^2 = (0.9)c/(50r^2)$ for the remainder of the proof.

Fix a block b and consider the event $surround_b$. Using Chernoff Bounds adapted to Poisson distributions, we can show that, w.h.p., at most αnr^2 nodes are in p_b ; let E be the event that this bound holds. Assume E . Then, we have, by Fact 1,

$$\Pr[surround_b] = \Pr\left[\bigcap_{v \in p_b} F_v\right] \geq \prod_{v \in p_b} \Pr[F_v] \geq (1 - e^{-\lambda\pi s^2})^{\alpha nr^2}$$

and since $\Pr[E] \rightarrow 1$ for large n , we have $\Pr[surround_b] \geq (0.9)(1 - e^{-\lambda\pi s^2})^{\alpha nr^2}$, for large n . Then, the probability that there exists a block $b \in \mathcal{B}_{left}$ for which event $isolation_b$ occurs is

$$\begin{aligned} \Pr[\exists b \in \mathcal{B}_{left} isolation_b] &= \Pr[\exists b \in \mathcal{A}_{left} surround_b] \\ &= 1 - \Pr[\forall b \in \mathcal{A}_{left} \neg surround_b] \\ &= 1 - (\Pr[\neg surround_b])^{|\mathcal{A}_{left}|} \\ &\geq 1 - \left(1 - (0.9) \left(1 - \frac{\ln(1/(r^2 f(n)))}{\alpha r^2 n}\right)^{\alpha r^2 n}\right)^{k/r^2} \\ &= 1 - (1 - (0.9)r^2 f(n))^{k/r^2} \rightarrow 1 \text{ as } n \rightarrow \infty. \end{aligned}$$

The same calculations apply to the second half of the unit square, thus showing the occurrence of at least 2 events $isolation_b$ w.h.p. This concludes the argument in the first case.

Case 2: $\lambda = \ln(g(n)r^2n)/(\pi s^2)$, with $g(n) \in \Omega(1)$

Consider again the event $surround_b$. For $\lambda = \ln(g(n)r^2n)/(\pi s^2)$, with $g(n) \in$

$\Omega(1)$, the same argument as in case 1 implies

$$\begin{aligned}\Pr[\text{surround}_b] &\geq (0.9)(1 - e^{-\lambda\pi s^2})^{\alpha nr^2} = (0.9)(1 - 1/(g(n)r^2n))^{\alpha nr^2} \\ &\geq c' \in \Theta(1).\end{aligned}$$

Let \mathcal{S}_{left} be the set of all blocks in \mathcal{B}_{left} for which the event surround_b occurs. Since the probability of surround_b is constant, the expected size of the set \mathcal{S}_{left} is a constant fraction of $|\mathcal{B}_{left}|$. The number of blocks in \mathcal{B}_{left} is $|\mathcal{B}_{left}| = 1/50r^2$. For $r \in o(1)$, $|\mathcal{B}_{left}|$ grows to infinity as $n \rightarrow \infty$ and thus, under the preceding assumptions, we use Chernoff bounds to show that $|\mathcal{S}_{left}| \geq (0.9)c'/(50r^2)$ w.h.p. Assume this bound on $|\mathcal{S}_{left}|$ and let $k/r^2 = (0.9)c'/(50r^2)$ for the remainder of the proof.

From Remark 1, if the spot arrival rate is $\lambda = \ln(nh(n))/(\pi s^2)$, $h(n) \in \Omega(1)$, we find a positive constant probability that the graph $U(n, r, \lambda, s)$ is dead. Hence, consider the subset of spot arrival rates of the form $\lambda = \ln(nh(n))/(\pi s^2)$, $h(n) \in o(1)$. Then, for these values of λ , the probability that there exists a block $b \in \mathcal{B}_{left}$ for which event isolation_b occurs is

$$\begin{aligned}\Pr[\exists b \in \mathcal{B}_{left} \text{ isolation}_b] &= \Pr[\exists b \in \mathcal{S}_{left} \text{ alive}_b] = 1 - \Pr[\forall b \in \mathcal{S}_{left} \neg \text{alive}_b] \\ &= 1 - (\Pr[\neg \text{alive}_b])^{|\mathcal{S}_{left}|} \\ &\geq 1 - (1 - (1 - (1 - e^{-\lambda\pi s^2})^{k'nr^2}))^{k/r^2} \\ &= 1 - ((1 - e^{-\ln(nh(n))})^{k'nr^2})^{k/r^2} \\ &= 1 - (1 - 1/(nh(n)))^{k'kn} \rightarrow 1 \text{ as } n \rightarrow \infty.\end{aligned}$$

The same calculations apply to the second half of the unit square, thus showing the occurrence of at least 2 events isolation_b w.h.p. This concludes the argument in the second case.

To conclude the proof, fix the function $f(n) = 1/r$. Since $r \in o(1)$, we have $f(n) \in \omega(1)$. Hence the corresponding $\tilde{\lambda} = \ln\left(\frac{\alpha r^2 n}{\ln(1/(r^2 f(n)))}\right)/(\pi s^2) = \ln\left(\frac{\alpha r^2 n}{\ln(1/r)}\right)/(\pi s^2)$ is in Λ . We show that $\tilde{\lambda} < \beta \ln\left(\frac{\alpha r^2 n}{\ln(1/r^2)}\right)/(\pi s^2)$, for any constant $\beta > 1$. Indeed,

$$\begin{aligned}\tilde{\lambda} &= \ln\left(\frac{\alpha r^2 n}{\ln(1/r)}\right)/(\pi s^2) = \ln\left(\frac{\alpha r^2 n}{0.5 \ln(1/r^2)}\right)/(\pi s^2) \\ &= \left(\ln\left(\frac{\alpha r^2 n}{\ln(1/r^2)}\right) + \ln 2\right)/(\pi s^2) \leq \beta \ln\left(\frac{\alpha r^2 n}{\ln(1/r^2)}\right)/(\pi s^2).\end{aligned}$$

It follows that all $\lambda = \beta \ln\left(\frac{\alpha r^2 n}{\ln(1/r^2)}\right)/(\pi s^2)$, for any constant $\beta > 1$, are also in Λ which proves the lemma. Note that, under the assumption $s \in o(1/\sqrt{n})$, we have $\min\{n, 1/s^2\} = n$. \blacksquare

The proof of Lemma 3 differs from the proof of Lemma 2 only in the use of a partition to obtain the probability of the event surround_b .

Lemma 3. Fix any constant $\beta > 1$. For $s \in o(r)$ and $r \in o(1)$, the graph $U(n, r, \lambda, s)$ is disconnected w.h.p. when $\lambda = 4\beta \ln \left(\frac{8r^2/s^2}{\ln(1/r^2)} \right) / \pi s^2$.

The preceding lemmas concern only the case when $r \in o(1)$. As stated in Theorem 4, for $r \in \Omega(1)$, thresholds on spot arrival rate separate the case of connected $U(n, r, \lambda, s)$ from the case when it is dead. Hence, we do not provide any non-connectivity result for $r \in \Omega(1)$ and defer this case to the next section.

3.2 Connectivity Results

In this section, we show conditions on spot arrival rate guaranteeing connectivity of the graph $U(n, r, \lambda, s)$ w.h.p. We show connectivity of $U(n, r, \lambda, s)$ by proving the existence of a fault-free node in each square of a sufficiently fine partition of the unit square w.h.p. This implies the existence of a fault-free path between any pair of nodes of the graph $U(n, r, \lambda, s)$ and hence this graph is connected.

Partition the unit square into a mesh of $r/\sqrt{5} \times r/\sqrt{5}$ squares, called blocks. Let \mathcal{B} be the set of all blocks. The distance between any two points in blocks which are adjacent by an edge (edge-adjacent) is at most r . Hence, functional nodes in adjacent blocks can communicate with each other.

Partition each block $b \in \mathcal{B}$ into a mesh of $3s \times 3s$ squares called tiles. Let T_b be the set of all these $r^2/(45s^2)$ tiles for the block b . For a fixed tile $t \in T_b$, let $free_t$ be the event that it contains no spot. Under the event $free_t$, the central $s \times s$ square $c_t \subset t$ is at distance greater than s from all spots. Let a_t be the event that c_t contains at least one node. Since node arrivals and spot arrivals are independent, the events $free_t$ and a_t are independent. Moreover, for all $t \neq s \in T_b$, the events a_t, a_s ($free_t, free_s$) are independent since they are respectively the result of arrivals inside non-overlapping tiles t and s .

Consider the event $alive_b$ that a fixed block b contains at least one functional node. The event $alive_b$ is implied by the existence of a tile $t \in T_b$ where both the events $free_t$ and a_t occur. Let $alive'_b = \{\exists t \in T_b \text{ s.t. } free_t \cap a_t\}$ be this sub-event of $alive_b$. Hence, the probability of event $alive_b$ that a fixed block b contains at least one functional node is

$$\begin{aligned} \Pr[alive_b] &\geq \Pr[alive'_b] = \Pr[\exists t \in T_b \text{ } free_t \cap a_t] = 1 - \Pr[\forall t \in T_b \neg free_t \cup \neg a_t] \\ &= 1 - (\Pr[\neg free_t \cup \neg a_t])^{|T_b|} = 1 - (1 - \Pr[free_t \cap a_t])^{|T_b|} \\ &= 1 - (1 - \Pr[free_t] \Pr[a_t])^{|T_b|} = 1 - (1 - e^{-\lambda 9s^2} (1 - e^{-ns^2}))^{r^2/(45s^2)}. \end{aligned}$$

Let $connect$ be the event that each block b in \mathcal{B} contains at least one functional node. We have $\Pr[connect] \geq \Pr[\forall b \in \mathcal{B} \text{ } alive'_b]$. The next two lemmas are easily derived from the above estimates.

Lemma 4. For $s \in o(1/\sqrt{n})$ and any constant $\alpha < 1$, the graph $U(n, r, \lambda, s)$ is connected, w.h.p., when the spot arrival rate is $\lambda = \alpha \ln \left(\frac{nr^2}{45 \ln(1/r^2)} \right) / 9s^2$.

Lemma 5. For $s \in \Omega(1/\sqrt{n}) \cap o(r)$ and any constant $\alpha < 1$, $U(n, r, \lambda, s)$ is connected, w.h.p., when the spot arrival rate is $\lambda = \alpha \ln \left(\frac{r^2/s^2}{45 \ln(1/r^2)} \right) / 9s^2$.

For large values of r , we show connectivity for the same range of λ for which we have shown the graph $U(n, r, \lambda, s)$ to be alive w.h.p.

Lemma 6. *For $r \in \Omega(1)$ and $s \in o(1)$, the graph $U(n, r, \lambda, s)$ is connected, w.h.p., when the spot arrival rate is $\lambda = \alpha \frac{\ln(\min\{n, 1/s^2\})}{\pi s^2}$, for any constant $\alpha < 1$.*

For $r \in \Omega(1)$ and $s \in \Omega(1) \cap o(r)$, we observe that if $r \in \Theta(1)$, then the condition $s \in o(r)$ is impossible. Hence, necessarily, $r \in \omega(1)$. Since the unit square has a diameter of $\sqrt{2}$, if it is alive, then it is also connected for $r \in \omega(1)$ and sufficiently large n . Hence Lemma 7 follows from Theorem 2.

Lemma 7. *For $r \in \omega(1)$ and $s \in \Omega(1) \cap o(r)$, the graph $U(n, r, \lambda, s)$ is connected, w.h.p., when the spot arrival rate is $\lambda \in o(1/s^2)$.*

4 Broadcasting Algorithm

We propose a deterministic algorithm which completes broadcast with probability $1 - \epsilon$ in time $O(D + \log 1/\epsilon)$, in the fault-free graph $U(n, r, \lambda, s)$ for $s \in o(1/\sqrt{n})$. The algorithm consists of two parts: a preprocessing part called *spokesman election*, and a message transmission part. In the spokesman election part a unique node, called the *spokesman*, is selected in each square of a partition defined below. Only the spokesman of a square relays messages in the following part.

Partition the unit square into a mesh of $r/\sqrt{5} \times r/\sqrt{5}$ squares called boxes and let S be the set of these boxes. Group the boxes in 5×5 matrices, called blocks and let B be the set of all these blocks. For all blocks, label its boxes 1 through 25, row by row. Further partition each box into a mesh of $1/\sqrt{n} \times 1/\sqrt{n}$ squares, called tiles. Let T_i be the set of all tiles in a box i . For all boxes, label the tiles 1 through $t = r^2 n/5$, row by row.

Algorithm \mathcal{A}^* ***Spokesman Election Part***

Nodes know their location and hence, they can compute the labels i, j of their box, and tile, respectively. Nodes label themselves (i, j) accordingly.

In parallel for all blocks, the algorithm executes rounds $i = 1, 2, \dots, 25$. In a round i , the algorithm sequentially goes through steps $j = 1, 2, \dots, t$. In a round i , at step j , all nodes with label (i, j) (in box i and tile j) transmit their label and the list of labels heard from adjacent boxes. At any given step j' , when only one node transmits its label (i, j') , the message is heard by all other nodes in the box i and all edge-adjacent boxes; The first node whose message is heard is chosen as the spokesman for box i by all other nodes in the box i (the node itself does not know it yet) and in edge-adjacent boxes. In subsequent steps in round i , nodes in the box i containing this node (i, j') are silent. The node (i, j') will learn that it is the spokesman for the box i when, in an edge adjacent box, a unique node transmits its own label and the list of labels heard from adjacent boxes. Since all boxes, except box 25, are edge-adjacent to a box with a larger label, by the end of round 25, if a spokesman is chosen for each box, then all

spokesmen, with the exception of the spokesman in box 25 are confirmed, i.e., they know that they are spokesmen. Hence, after round 25, a single transmission from the spokesman in box 24 is sufficient to confirm the spokesman of box 25. This transmission is done in parallel by all spokesmen in boxes labeled 24, right after the end of round 25.

Hence, the spokesman election part chooses and confirms one spokesman in every box if there is, in every box, a tile which contains exactly one functional node.

Message Transmission Part

In the first step of this part, the source transmits its message. Then, in parallel for all blocks, the algorithm is executed in identical phases $\rho = 1, 2, \dots$. In phase ρ , steps $j = 1, 2, \dots, 25$ are executed sequentially. In a step j , a spokesman of box j which has received the source message but has not relayed it yet, transmits the message. This completes the description of algorithm \mathcal{A}^* . See Figure 1.

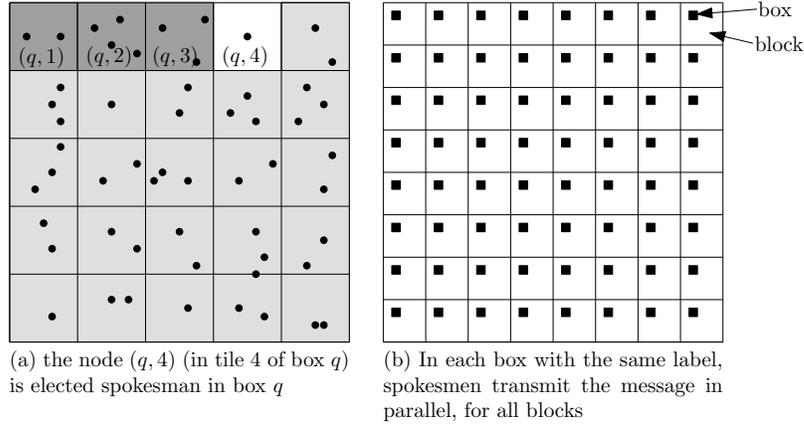


Fig. 1. Algorithm \mathcal{A}^* : (a) Spokesman Election part. (b) Message Transmission part.

Let ϵ be the tolerated error probability for the algorithm, i.e., we wish to broadcast with probability at least $1 - \epsilon$. Let \mathcal{A} be the algorithm \mathcal{A}^* modified so that the spokesman election part uses only the first $\frac{\ln(2D^2/\epsilon)}{\ln(1/(1-(0.9)e^{-(c+1)}))}$ tiles of each box.

Theorem 5. *Let c be a positive constant and $d = \ln(1/(1 - (0.9)e^{-(c+1)}))$. For $s \in o(1/\sqrt{n})$, $r^2 \geq \frac{5 \ln(5D^2/\epsilon)}{dn}$, and $\lambda \leq c/(\pi s^2)$, the algorithm \mathcal{A} broadcasts a message in time $O(D + \log 1/\epsilon)$, with probability at least $1 - \epsilon$.*

Proof. Consider a tile t . There exists a subsquare a of t of area $(1/\sqrt{n} - s)^2 = 1/n - 2s/\sqrt{n} + s^2$ whose nodes are not affected by spots in other tiles; the remaining subset a' of the tile has area $2s/\sqrt{n} - s^2$. Let $good_t$ be the event that there exists exactly one node in a , no node in a' , and that the node in a is not

within distance s of a spot. We have

$$\begin{aligned}\Pr[\text{good}_t] &= e^{-(1/n - 2s/\sqrt{n} + s^2)n} \cdot e^{-(2s/\sqrt{n} - s^2)n} \cdot e^{-\lambda\pi s^2} \\ &\geq (1 - 2s\sqrt{n} + s^2n)e^{-1 + (2s\sqrt{n} - s^2n) - (2s\sqrt{n} - s^2n) - \frac{c}{\pi s^2}\pi s^2} \geq 0.9e^{-(c+1)}\end{aligned}$$

for large n . Let spokesman_q be the event that the spokesman election part is successful in a fixed box q . Since $r^2 \geq \frac{5\ln(5D^2/\epsilon)}{dn}$, there are at least $nr^2/5 = \frac{\ln(5D^2/\epsilon)}{d}$ tiles in each box. Hence, the algorithm \mathcal{A} can execute its spokesman election part. Then, we have

$$\begin{aligned}\Pr[\text{spokesman}_q] &= 1 - (1 - \Pr[\text{good}_t])^{\ln(5D^2/\epsilon)/d} \\ &\geq 1 - (1 - 0.9e^{-(c+1)})^{\ln(5D^2/\epsilon)/d} \\ &= 1 - \left(\frac{5D^2}{\epsilon}\right)^{\ln(1 - 0.9e^{-(c+1)})/d} \\ &= 1 - \left(\frac{5D^2}{\epsilon}\right)^{-\frac{\ln(1/(1 - 0.9e^{-(c+1)}))}{\ln(1/(1 - 0.9e^{-(c+1)}))}} = 1 - \frac{\epsilon}{5D^2}.\end{aligned}$$

There are at most $5D^2$ boxes. Hence, the event spokesmen that each box contains one spokesman occurs with probability

$$\Pr[\text{spokesmen}] \geq \left(1 - \frac{\epsilon}{5D^2}\right)^{5D^2} \geq 1 - \epsilon.$$

We now show that, assuming the event spokesmen , all functional nodes are informed and we estimate the total time of the algorithm. Consider the Message Transmission part. For each phase, 25 time steps are elapsed. We say that a box with label j is active if the algorithm step is j , i.e., when its spokesman may transmit. All boxes with the same label are located at distance at least $4r/\sqrt{5}$ from each other. Only spokesmen in active boxes (with the same label j , at a step j) transmit. Hence all nodes in boxes adjacent to active boxes will receive the message correctly at every time step when a spokesman transmits in this active box (due to large distances between boxes with the same label, there are no collisions in adjacent boxes). It follows that if a message is received by any box in a block i at time t , then there exists a positive constant δ such that at time $t + \delta$ all nodes in the block i will know the message. Moreover, at time $t + \delta$, the nodes in boxes outside the block i , but adjacent to the boxes in block i also have received the message. Consider two nodes in different blocks i and j such that there is a sequence of edge-adjacent blocks of length $k - 1$ between them. If all nodes in block i have received the message by time t , it follows from the above that, at the time $t + k\delta$, the message will also be received by all nodes in block j . Since the unit square is partitioned in rows and columns of $\sqrt{5}/(5r)$ blocks, there is a sequence of, at most, $2\sqrt{5}/(5r)$ blocks between any two blocks i and j , so that consecutive blocks are edge-adjacent. Hence, the total broadcast time is at most $2\delta\sqrt{5}/(5r)$. Since the diameter of the graph is at least $1/r$, the

message transmission part is completed in time $O(D)$. The spokesman election part of the algorithm terminates in $O(\log(5D^2/\epsilon)) = O(1 + \log D + \log 1/\epsilon)$ time steps. Hence, the total execution time of the algorithm is $O(D + \log 1/\epsilon)$, and the algorithm is correct with probability at least $1 - \epsilon$. ■

Acknowledgments. Evangelos Kranakis and Michel Paquette were supported by MITACS and NSERC. Andrzej Pelc was supported by the Research Chair in Distributed Computing of the Université du Québec en Outaouais and NSERC.

References

1. D. Bienstock. Broadcasting with random faults. *Discr. Appl. Math*, 20:1–7, 1988.
2. B. S. Chlebus, K. Diks, and A. Pelc. Sparse networks supporting efficient reliable broadcasting. *Nordic Journal of Computing*, 1:332–345, 1994.
3. B. S. Chlebus, K. Diks, and A. Pelc. Reliable broadcasting in hypercubes with random link and node failures. *Comb., Prob. and Computing*, 5:337–350, 1996.
4. A. E.F. Clementi, A. Monti, and R. Silvestri. Round robin is optimal for fault-tolerant broadcasting on wireless networks. *J. Par. Distrib. Comp.*, 64:89–96, 2004.
5. A. Dessmark and A. Pelc. Broadcasting in geometric radio networks. *Journal of Discrete Algorithms*, 5:187–201, 2007.
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. D. B. Johnson and D. A. Maltz. *Dynamic Source Routing in Ad Hoc Wireless Networks*, chapter 5, pages 153–181. Kluwer Academic Publishers, 1996.
8. E. Kranakis, D. Krizanc, and A. Pelc. Fault-tolerant broadcasting in radio networks. *Journal of Algorithms*, 39:47–67, 2001.
9. E. Kranakis, D. Krizanc, and J. Urrutia. Coverage and connectivity in networks with directional sensors. In *Proc. Euro-Par 2004*, pages 917–924, 2004.
10. E. Kranakis, M. Paquette, and A. Pelc. Communication in networks with random dependent faults. In *Proc. 32nd Int. Symp. MFCS*, pages 418–429, 2007.
11. M. Paquette and A. Pelc. Fast broadcasting with byzantine faults. *International Journal of Foundations of Computer Science*, 17(6):1423–1439, 2006.
12. A. Pelc. Fault-tolerant broadcasting and gossiping in communication networks. *Networks*, 28(6):143–156, 1996.
13. M. D. Penrose. On k -connectivity for a geometric random graph. *Random Struct. Alg.*, 15:145–164, 1999.
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.
15. M. Thottan and C. Ji. Using network fault predictions to enable IP traffic management. *J. Network Syst. Manage.*, 9(3):327–346, 2001.
16. M. Yajnik, J. Kurose, and D. Towsley. Packet loss correlation in the MBone multicast network. In *Proceedings of IEEE Global Internet*, pages 94–99, 1996.