Reliable Broadcast in Wireless Networks with Probabilistic Failures

Vartika Bhandari
Dept. of Computer Science, and
Coordinated Science Laboratory
University of Illinois at Urbana-Champaign
vbhandar@uiuc.edu

Nitin H. Vaidya
Dept. of Electrical and Computer Eng., and
Coordinated Science Laboratory
University of Illinois at Urbana-Champaign
nhv@uiuc.edu

***Abstract*— We consider the problem of reliable broadcast in a wireless network in which nodes are prone to failure. Each node can fail independently with probability p. Failures are permanent. The primary focus is on Byzantine failures, but we also handle crash-stop failures. We consider two network models: a regular grid, and a random network. Our necessary and sufficient conditions for the Byzantine failure model indicate that p should be less than $\frac{1}{2}$, and the *critical* node degree is $\Theta\left(d_{min} + \frac{\ln n}{\ln \frac{1}{2p} + \ln \frac{1}{2(1-p)}}\right)$ (where d_{min} is the minimum node degree associated with a non-empty neighborhood, and is a small constant). For a random network we prove that, for failure probability less than $\frac{1}{2}$, the *critical* average degree for reliable broadcast is $O(\frac{\ln n}{\frac{1}{2}-p+\frac{1}{2}\ln \frac{1}{2(1-p)}})$. We briefly discuss the issue of crash-stop failures for which we have results that improve upon previously existing results for this model, when p approaches 0. We also identify an interesting similarity in the structure of various known results in the literature pertaining to a set of related problems in the realm of connectivity and reliable broadcast.**

***Index Terms*— Probabilistic failure, Byzantine faults, crash-stop faults, broadcast, fault-tolerance, reliability**

I. Introduction

Reliable broadcast in the presence of Byzantine and crash-stop failures has been studied under different network and failure models. A reliable broadcast mechanism may be of significant utility in large-scale sensor network deployments. While the shared nature of the wireless medium is conducive to the broadcast operation, the unreliability of the wireless channel, and the possibility of collisions can make it a difficult problem to solve. As a first step towards addressing the issue, it is useful to focus on an idealized wireless channel. We consider the problem of reliable broadcast in a such an idealized wireless network. We primarily focus on Byzantine failures, but have also considered the case of crash-stop failures. The failures are permanent and are assumed to occur probabilistically, i.e., each node can fail independently with a certain probability p. However, once failure has happened, the faulty nodes can exhibit worst-case behavior.

For the Byzantine failure model, we show that reliable broadcast in a grid network of n nodes requires that $p < \frac{1}{2}$, and the *critical* node degree (defined in Section III) is $\Theta\left(d_{min} + \frac{\ln n}{\ln \frac{1}{2p} + \ln \frac{1}{2(1-p)}}\right)$. This may alternatively be stated as $\Theta\left(d_{min} + \frac{\ln n}{D(Q_{\frac{1}{2}}||P)}\right)$ where $Q_{\frac{1}{2}}$ denotes the $Bernoulli(\frac{1}{2})$ distribution, P denotes the $Bernoulli(p)$ distribution, and $D(Q||P)$ denotes the *relative entropy* (or Kullback-Leibler distance) between distributions Q and P. We also prove that in a randomly deployed network with Byzantine failures, the critical average node degree for reliable broadcast is $O\left(\frac{\ln n}{\frac{1}{2}-p+\frac{1}{2}\ln \frac{1}{2(1-p)}}\right)$ when $p < \frac{1}{2}$.

We briefly discuss the case of crash-stop failures in a grid network, in Section X. For crash-stop failures, the problem of reliable broadcast is equivalent to connectivity. For this case, we have results showing that the critical node degree is $\Theta\left(d_{min} + \frac{\ln n}{\ln \frac{1}{p}}\right)$ with $p < 1$, or alternatively stated, $\Theta\left(d_{min} + \frac{\ln n}{D(Q_1||P)}\right)$, where Q_1 is the $Bernoulli(1)$ distribution. Our results improve upon previous results proved in [1] when the failure probability p approaches 0.

We also identify an interesting but intuitive similarity in the structure of results (previously known results, as well as the results derived in this paper) for a set of mutually related problems pertaining to connectivity and reliable broadcast. This is discussed in Section XI.

II. Problem Model

We consider a two network models, viz., a regular grid, where nodes are located on a two-dimensional square grid (each grid unit is a 1×1 square), and a random network, where node locations are i.i.d. uniformly distributed over the deployment region. In both models, the network is assumed to be deployed over a $\sqrt{n}$ x $\sqrt{n}$ square region. The pre-failure topology (i.e., node locations) of the deployed network is assumed to be known to all nodes.

Formal Definition of Reliable Broadcast: Any node in the network can originate a broadcast message. The goal is to ensure that all nodes receive the valid broadcast value *with high probability* [1] . In the Byzantine failure model, this source

This research is supported in part by NSF grant CNS 05-19817, and a Vodafone Graduate Fellowship.

[1] We use the term *with high probability* (w.h.p.) to mean with probability 1 as $n \to \infty$.

node may be faulty. Thus the goal is to ensure that if the source is non-faulty, every non-faulty node in the network correctly receives and determines the broadcast value; if the source is faulty, all non-faulty nodes should agree on some common value. In the crash-stop failure model, a message can only be originated by a non-faulty node (as faulty nodes cease to function), and the goal is to ensure that all non-faulty nodes receive this value. If even one non-faulty node fails to make a valid value determination, the broadcast is deemed to have failed. Reliable broadcast is said to fail in a given fault configuration, if it fails for at least one possible broadcast origin/source.

For a given broadcast instance, once an origin/source is designated, it is identified as $(0,0)$. All nodes can then be uniquely identified by their coordinate location (x,y) w.r.t. this origin. In the grid network model, the node coordinates are always *integers*, while for random networks they are *real* numbers. All nodes have a common transmission radius $r(n,p)$. For grid networks, we assume that $r(n,p)$ is an integer, and for random networks it is allowed to be any real number. A message transmitted by a node (x,y) is heard by all nodes within distance $r(n,p)$ from it (where distance is defined in terms of the assumed metric). The set of these nodes is termed the neighborhood of (x,y).

In this paper, we consider two distance metrics: L_∞ and L_2. The L_∞ metric is the metric induced by the L_∞ norm [2], such that the distance between points (x_1,y_1) and (x_2,y_2) is given by $\max\{|x_1-x_2|,|y_1-y_2|\}$ in this metric. Thus the neighborhood of (a,b) comprises a square of side $2r$ with its centroid at (a,b), and the degree of a node is $4r^2+4r$. In this metric, the minimum node degree $d_{min}=8$ corresponding to $r=1$. The L_2 metric is induced by the L_2 norm [2], and is the Euclidean distance metric. The L_2 distance between points (x_1,y_1) and (x_2,y_2) is given by $\sqrt{(x_1-x_2)^2+(y_1-y_2)^2}$, and the neighborhood of (a,b) comprises nodes within a circle of radius r centered at (a,b). The L_∞ metric (which was also used in [3], [4], and [5]) enables more tractable analysis, from which necessary and sufficient conditions for the L_2 (Euclidean) metric proceed. In Section VIII, we further elaborate on this.

A random failure mode is assumed, wherein each node can fail with probability p independently of other nodes. Failures are permanent. We primarily focus on Byzantine failures. In the Byzantine failure mode, a faulty node can behave arbitrarily, in contrast to crash-stop failures, where a faulty node simply stops functioning. However, in our model, the Byzantine nodes cannot spoof addresses or cause collisions, i.e., the MAC layer is assumed fault-free, and the Byzantine faults reside only in higher layers of the protocol stack. We assume that the channel is perfectly reliable, and a local broadcast is correctly received by all neighbors. The same *reliable local broadcast* assumption underlies the results in [3] and [4] for a locally bounded adversarial fault model. While the *occurrence* of the permanent failures is probabilistic, the failed Byzantine nodes can thereafter choose to behave in a worst-case manner (i.e. modulate the messages they send to cause most confusion to non-faulty nodes). The non-faulty nodes do not know which nodes have failed.

III. Notation and Terminology

We briefly describe the notation and terminology used in this paper.

Nodes are identified by their coordinate location i.e. (x,y) denotes the node at (x,y). The neighborhood of (x,y) comprises all nodes within distance r of (x,y) and is denoted as $nbd(x,y)$. For succint description of grid network proofs, we define a term $pnbd(x,y)$ where $pnbd(x,y) = nbd(x-1,y) \cup nbd(x+1,y) \cup nbd(x,y-1) \cup nbd(x,y+1)$. Intuitively $pnbd(x,y)$ denotes the *perturbed neighborhood* of (x,y) obtained by perturbing the center of the neighborhood to one of the nodes immediately adjacent to (x,y) on the grid. Also $faults(\mathcal{S})$ denotes the number of faulty nodes in the set of nodes $\mathcal{S}$. The term *qnbd* is sometimes used as an abbreviation for quarter-neighborhood (defined later in the paper). Transmission range is referred to as $r(n,p)$ and sometimes as just r. The node degree is referred to as $d(n,p)$ or just d.

We use standard asymptotic notation [6]. Besides, we denote by $D(Q_{\frac{1}{2}}||P)$ the relative entropy between the $Bernoulli(\frac{1}{2})$ and $Bernoulli(p)$ distributions. Thus $D(Q_{\frac{1}{2}}||P) = \frac{1}{2}\ln\frac{1}{2p} + \frac{1}{2}\ln\frac{1}{2(1-p)}$.

By *critical* transmission range for reliable broadcast, we imply the minimum transmission range $r_{critical}$, required to guarantee that broadcast is achievable w.h.p.

In a grid network, with the considered L_∞ metric, the node degree is exactly determined by specifying the transmission range. Hence, we can define the notion of *critical* degree $d_{critical}$ correponding to the transmission range $r_{critical}$.

Thus:

$$d_{critical} = \Omega(g(n,p)) \implies \exists c_1 > 0, \text{ such that:}$$
$$d \le c_1 g(n,p) \implies \lim_{n\to\infty} Pr[\text{reliable broadcast achievable}] < 1$$

This yields a *necessary* condition. If $\lim_{n\to\infty} Pr[\text{reliable broadcast achievable}] = 0$, it is a *strong* necessary condition.

$$d_{critical} = O(f(n,p)) \implies \exists c_2 > 0, \text{ such that:}$$
$$d \ge c_2 f(n,p) \implies \lim_{n\to\infty} Pr[\text{reliable broadcast achievable}] = 1$$

This yields a *sufficient* condition.

Thus $d_{critical}$ is $\Theta(f(n,p))$ implies that $d_{critical}$ is $\Omega(f(n,p))$ and $O(f(n,p))$.

In a random network, the degrees of individual nodes can vary; however, it is possible to define a notion of *critical* average degree $d^{avg}_{critical}$, which is the average degree corresponding to the range $r_{critical}$. Then $d^{avg}_{critical}$ can be expressed in asymptotic notation, similar to $d_{critical}$ for a grid network.

IV. Related Work

Reliable broadcast in radio networks has been studied in [3], [7], [4] for a locally bounded adversarial model in which the adversary may choose fault locations so long as no neighborhood has more than t faulty nodes. The issue of

achieving broadcast when the (locally bounded) adversary can cause bounded a bounded number of collisions or address spoofing is handled in [5].

However, in many practical situations, nodes may fail randomly with a certain probability. It is therefore of interest to determine the conditions under which reliable broadcast is achievable under such a probabilistic fault model. In [8], reliable broadcast under probabilistic *transient* failures has been studied. Our results pertain to reliable broadcast in the presence of *permanent* random Byzantine failures, e.g., when a Byzantine adversary launching a remote attack has an independent probability p of compromising each node.

For crash-stop faults, the reliable broadcast problem reduces to the connectivity problem. Conditions for connectivity and coverage have been formulated in the context of different network models. A grid network model similar to ours was considered in [1] where nodes are located at grid locations on a square grid, but may fail independently. Nodes have a common transmission range r. The probability of not failing is q (where $q = 1-p$), and it is shown that a sufficient condition for connectivity and coverage is that transmission range r must be set to ensure that node degree is $c_1(\frac{\log n}{q})$ (for some constant c_1). It is also shown that a necessary condition for coverage (and hence for joint coverage and connectivity) is that node degree be at least $c_2(\frac{\log n}{q})$ (for another constant c_2). A fallacy in the above necessary condition was pointed out by [9], and a subsequent correction [10] by the authors of [1] presents examples illustrating that the necessary condition may fail to hold for certain subranges of q. We have also derived results for crash-stop failures/connectivity that yield a different expression than [1], and while our results are within a constant factor of their results for most values of p, our results are more accurate when $p \to 0$. We discuss this further in Section X.

Recently, necessary and sufficient conditions for asymptotic connectivity in a random network with low duty cycle sensors have been formulated in [11]. This is equivalent to the problem of crash-stop failures in a random network.

V. Some Useful Mathematical Results

We state some mathematical results that have been used in our proofs:

Fact 1: $\forall x \in [0,1] : \ln \frac{1}{1-x} \geq x$

Fact 2: If $|f(n)| \leq n^{\frac{1}{2}-\varepsilon} (0 < \varepsilon < \frac{1}{2})$:

$$\lim_{n\to\infty} \left(1 + \frac{f(n)}{n}\right)^n = e^{(\lim_{n\to\infty} f(n))}$$

Lemma 1: (Jogdeo & Samuels [12]) Given $X = Y_1 + Y_2 + ..., +Y_n$ where $\forall i, Y_i = Bernoulli(p_i)$, and $\sum p_i = np$, the median m of the distribution is either $\lfloor np \rfloor or \lceil np \rceil$, i.e., $Pr[X \leq m] \geq \frac{1}{2}$ and $Pr[X \geq m] \geq \frac{1}{2}$.

Lemma 2: (Chernoff Bound) If $X = \sum_{i=1}^{n} X_i$, where X_i's are i.i.d. $Bernoulli(p)$, then for $0 < \beta < 1$:

$$Pr[X \leq (1-\beta)E[X]] \leq \exp(-\frac{\beta^2}{2}E[X]) \quad (1)$$

Lemma 3: (Relative Entropy Form of Chernoff-Hoeffding Bound[13]) If $X = \sum_{i=1}^{n} X_i$, where X_i's are i.i.d. $Bernoulli(p)$, then for $p \leq \beta \leq 1$:

$$Pr[X \geq \beta n] \leq e^{-n(\beta \ln \frac{\beta}{p} + (1-\beta)\ln \frac{1-\beta}{1-p})} \quad (2)$$

Lemma 4: [14] If X_1, X_2,..., X_n are drawn i.i.d. from alphabet χ according to $Q(x)$, then probability of the observed sequence being $\mathbf{x}$ is given by:

$$Q^{(n)}(\mathbf{x}) = e^{-n(H(P_x)+D(P_\mathbf{x}||Q))} \quad (3)$$

where H and D denote the entropy and relative entropy functions (here considered w.r.t base e), and $P_\mathbf{x}$ is the empirical distribution of sequence $\mathbf{x}$.

Let $T(P)$ denote the type class corresponding to distribution P, i.e., the set of sequences $\mathbf{x}$ whose empirical probability distribution is P. Then, for any distribution P belonging to the set of possible types with denominator n, and any distribution Q, the size of type class $T(P)$ satisfies:

$$\frac{1}{(n+1)^{|\chi|}} e^{nH(P)} \leq |T(p)| \leq e^{nH(P)} \quad (4)$$

and, the probability of the type class $T(P)$ under $Q^{(n)}$ is governed by:

$$\frac{1}{(n+1)^{|\chi|}} e^{-n(D(P||Q))} \leq Q^{(n)}(T(p)) \leq e^{-n(D(P||Q))} \quad (5)$$

Lemma 5: (Vapnik-Chervonenkis Theorem) Let S be a set with finite VC dimension VCdim($\mathcal{S}$). Let $\{X_i\}$ be i.i.d. random variables with distribution P. Then for $\varepsilon, \delta > 0$:

$$Pr\left(\sup_{D\in\mathcal{S}} |\frac{1}{N}\sum_{i=1}^{N} I_{X_i \in D} - P(D)| \leq \varepsilon\right) > 1-\delta$$

$$\text{whenever } N > \max\left(\frac{8\text{VCdim}(\mathcal{S})}{\varepsilon}\log_2\frac{16e}{\varepsilon}, \frac{4}{\varepsilon}\log_2\frac{2}{\delta}\right)$$

Lemma 6: Suppose we are given a region of area n, with n nodes located uniformly at random. Consider all axis-parallel rectangles of area $a(n)$. If $a(n) \geq 100\alpha \ln n, 1 \leq \alpha \leq \frac{n}{100 \ln n}$, then each such rectangle has at least $100\alpha \ln n - 50 \ln n$ nodes, with probability at least $1 - \frac{50 \ln n}{n}$.

Proof: Please see [15]. ■

VI. Reliable Broadcast with Probabilistic Byzantine Failures

We present necessary and sufficient conditions for achievability of reliable broadcast in a grid network. Note that node degree $d(n,p) = 4(r^2(n,p) + r(n,p))$ for nodes not near the edges, and the minimum number of neighbors of any node (even one located in a corner) is at least $\frac{1}{4}d(n,p)$. In the following proofs, we shall assume a toroidal network for ease of explanation. However this assumption can be relaxed without affecting the results. This is discussed further in Section IX.

A. Sufficient Condition for Reliable Broadcast

We now present a sufficient condition for the asymptotic achievability of reliable broadcast.

Theorem 1: In the grid network model, when $p < \frac{1}{2}$, and $d(n,p) = 4r(n,p)(r(n,p)+1) \geq \max\{d_{min}, 16\frac{\ln n}{\ln\frac{1}{2p}+\ln\frac{1}{2(1-p)}}\} = \max\{d_{min}, 8\frac{\ln n}{D(Q_{\frac{1}{2}}||P)}\}$, reliable broadcast is asymptotically achievable with probability 1.

Note that trivially $r(n,p)$ must be at least 1, else nodes would have no neighbors. Also when $\ln\frac{1}{2p} + \ln\frac{1}{2(1-p)} \leq \frac{16\ln n}{n-1}$, all network nodes are neighbors of the source, and thus the sufficient condition degenerates to merely indicating that having everyone in direct range suffices for reliable broadcast (which is the trivial sufficient condition for the assumed network and fault model). Thus the sufficient condition is of interest only so long as $\ln\frac{1}{2p} + \ln\frac{1}{2(1-p)} > \frac{16\ln n}{n-1}$.

a) $p = o(\frac{1}{n})$*:* When $p = o(\frac{1}{n})$, i.e., $np \to 0$, the probability of even a single node failing approaches 0 asymptotically, and thus reliable broadcast is trivially ensured even with $r(n,p) = 1$, i.e., degree d_{min}. This may be seen thus:

$$Pr[\text{No failures}] = (1-p)^n \tag{6}$$

$$\lim_{n\to\infty} Pr[\text{No failures;trivial broadcast}] \geq \lim_{n\to\infty}(1-p)^n \tag{7}$$

$$= e^{-\lim(np)} = 1 \text{ from Fact 2} \tag{8}$$

b) $p = \Omega(\frac{1}{n})$*:* We define a term called quarter-neighborhood (*qnbd*) of a node (x,y), and denote it by $qnbd(x,y)$. We associate eight quarter-neighborhoods with each node: $qnbd_A$, $qnbd_B$, $qnbd_C$, $qnbd_D$, $qnbd_{A'}$, $qnbd_{B'}$, $qnbd_{C'}$, $qnbd_{D'}$. The quarter-neighborhoods for a node (a,b) are depicted in Fig. 1 and 2, and their spatial extents are tabulated in Table I. Observe that $qnbd_B(a,b) = qnbd'_A(a-r-1,b)$, $qnbd_C(a,b) = qnbd_A(a-r,b+r+1)$, and $qnbd_D(a,b) = qnbd'_A(a,b+r+1)$. Similarly, $qnbd_{B'}(a,b) = qnbd_A(a-r-1,b)$, $qnbd_{C'}(a,b) = qnbd_{A'}(a-r-1,b+r)$, and $qnbd_{D'}(a,b) = qnbd_A(a,b+r+1)$. Thus if we simply consider $qnbd_A(u)$ and $qnbd_{A'}(u)\forall$ nodes u, we will have considered all quarter-neighborhoods, i.e. the number of distinct (but *not disjoint*) quarter-neighborhoods is $2n$. Henceforth, we shall sometimes use $Q(x,y)$ to refer to $qnbd_A(x,y)$, and $Q'(x,y)$ to refer to $qnbd_{A'}(x,y)$. The population of any *qnbd* is $r(r+1)$, and since $d = 4r^2+4r = 4r(r+1)$, the *qnbd* population = $\frac{d}{4}$. We now state and prove the following result which is crucial to proving our sufficient condition for reliable broadcast:

Region	x-extent	y-extent
$qnbd_A(a,b)$	$a \leq x \leq (a+r)$	$(b-r) \leq y \leq (b-1)$
$qnbd_B(a,b)$	$(a-r) \leq x \leq (a-1)$	$(b-r) \leq y \leq b$
$qnbd_C(a,b)$	$(a-r) \leq x \leq a$	$(b+1) \leq y \leq (b+r)$
$qnbd_D(a,b)$	$(a+1) \leq x \leq (a+r)$	$b \leq y \leq (b+r)$
$qnbd_{A'}(a,b)$	$(a+1) \leq x \leq (a+r)$	$(b-r) \leq y \leq b$
$qnbd_{B'}(a,b)$	$(a-r) \leq x \leq a$	$(b-r) \leq y \leq (b-1)$
$qnbd_{C'}(a,b)$	$(a-r) \leq x \leq (a-1)$	$b \leq y \leq (b+r)$
$qnbd_{D'}(a,b)$	$a \leq x \leq (a+r)$	$(b+1) \leq y \leq (b+r)$

TABLE I
SPATIAL EXTENTS OF QUARTER NEIGHBORHOODS

Theorem 2: If $p < \frac{1}{2}$, and $d(n,p) = 4r(n,p)(r(n,p)+1) \geq \max\{d_{min}, 16\frac{\ln n}{\ln\frac{1}{2p}+\ln\frac{1}{2(1-p)}}\} = \max\{d_{min}, 8\frac{\ln n}{D(Q_{\frac{1}{2}}||P)}\}$, then

$$\lim_{n\to\infty} Pr[\ \forall(x,y)\, faults(Q(x,y)) < \frac{d}{8} \text{ and } faults(Q'(x,y)) < \frac{d}{8}] \to 1$$

Proof: As shown above, the population of any *qnbd* is $\frac{d}{4}$. Each node may fail independently with probability p. Let $Y_{(x,y)}$ be a random variable denoting the number of faulty nodes in $Q(x,y)$. Then $E[Y_{(x,y)}] = p\frac{d}{4}$. Using $\delta = \frac{1}{2p} - 1$, we may then apply the relative entropy form of the Chernoff bound (Lemma 3) to $Y_{(x,y)} = \sum_{j\in Q(x,y)} I_j$, where I_j is an indicator variable that takes value 1 if node j is faulty. Note that $d \geq \max\{d_{min}, 16\frac{\ln n}{\ln\frac{1}{2p}+\ln\frac{1}{2(1-p)}}\} \geq 16\frac{\ln n}{\ln\frac{1}{2p}+\ln\frac{1}{2(1-p)}}$.

Thus, we obtain:

$$Pr[Y_{(x,y)} \geq \frac{d}{8}] \leq e^{-\frac{d}{4}(\frac{1}{2}\ln\frac{1}{2p}+\frac{1}{2}\ln\frac{1}{2(1-p)})} \tag{9}$$

$$\leq e^{-(\frac{16\ln n}{4(\ln\frac{1}{2p}+\ln\frac{1}{2(1-p)})})(\frac{1}{2}\ln\frac{1}{2p}+\frac{1}{2}\ln\frac{1}{2(1-p)})} = e^{-2\ln n} = \frac{1}{n^2} \tag{10}$$

Similarly, setting $Y'_{(x,y)}$ be a random variable denoting the number of faulty nodes in $Q'(x,y)$, we obtain that:

$$Pr[Y'_{(x,y)} \geq \frac{d}{8}] \leq \frac{1}{n^2} \tag{11}$$

By application of union bound over all $2n$ distinct quarter-neighborhoods:

$$\therefore \lim_{n\to\infty} Pr[\forall(x,y), Y(x,y) < \frac{d}{8} \text{ and } Y'(x,y) < \frac{d}{8}] \tag{12}$$

$$\geq 1 - 2n\left(\frac{1}{n^2}\right) = 1 - \frac{2}{n} \to 1 \tag{13}$$

■

We now consider a simple broadcast protocol that is fairly similar to the protocol described in [3] for the locally-bounded adversarial model:

- Initially, the source s does a local broadcast of the message.
- Each neighbor i of the source immediately commits to the the *first* value v it heard from the source, and then locally broadcasts it once in a $COMMITTED(i,v)$ message.

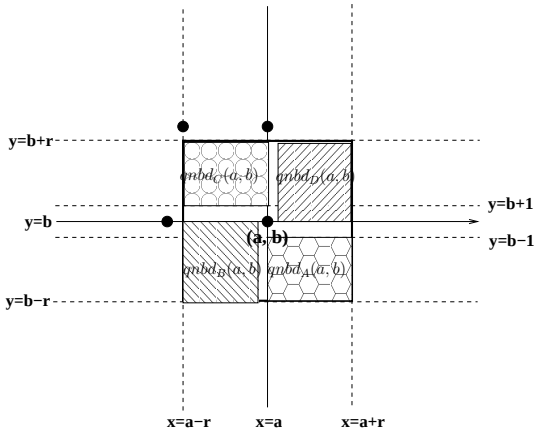

Fig. 1. Depiction of $qnbd_A$, $qnbd_B$, $qnbd_C$, $qnbd_D$

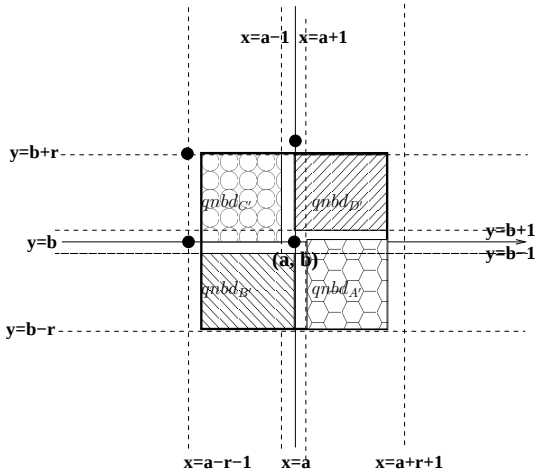

Fig. 2. Depiction of $qnbd_{A'}$, $qnbd_{B'}$, $qnbd_{C'}$, $qnbd_{D'}$

- Hereafter, the following protocol is followed by all nodes $j \notin nbd(s)$:
  If $\frac{1}{2}r(r+1)+1 = \frac{d}{8}+1$ $COMMITTED(i,v)$ messages are received for a certain value v, from neighbors i all lying within a single *qnbd*, commit to v, and locally broadcast a $COMMITTED(j,v)$ message.

Theorem 3: (Probabilistic Correctness) The probability that a node shall commit to a wrong value by following the above protocol diminishes to 0 asymptotically.

Proof: If all $Q(x,y)(Q'(x,y))$ have strictly less than $\frac{d}{8}$ faults, the correctness of the protocol proceeds as follows:

By the *reliable local broadcast* assumption, fault-free nodes in $nbd(s)$ are guaranteed to be able to commit to the correct value. The proof for the remaining nodes is by contradiction. Consider the first fault-free node, say j, that makes a wrong decision to commit to a value v. This implies that $\frac{d}{8}+1$ of its neighbors within some *qnbd* broadcast a $COMMITTED$ message for v (the $COMMITTED$ messages were directly heard, leaving no place for doubt). All of these nodes cannot be faulty, as less than $\frac{d}{8}$ nodes in any *qnbd* are faulty. Thus there was at least one fault-free node that committed to v. Since j is the first fault-free node to make a wrong decision, none of the fault-free nodes amongst the $\frac{d}{8}+1$ nodes could have made a wrong decision. Thus v must indeed be the correct value.

We know that all *qnbd* have less than $\frac{d}{8}$ faults with probability 1 asymptotically, and hence the protocol also functions correctly with probability 1 asymptotically. ■

Theorem 4: (Probabilistic Completeness) Each node is eventually able to commit to the (probabilistically) correct value.

Proof:

The proof proceeds by induction.

Base Case:

All non-faulty nodes in $nbd(0,0)$ are able to commit to the correct value. This follows trivially since they hear the origin directly, and we assume that address-spoofing is impossible.

Inductive Hypothesis:

If all non-faulty neighbors of a node located at (a,b) i.e. all non-faulty nodes in $nbd(a,b)$ are able to commit to the correct value, then all non-faulty nodes in $pnbd(a,b)$ are able to commit to the correct value.

Proof of Inductive Hypothesis:

We show that each node P in $pnbd(a,b) - nbd(a,b)$ has one of $qnbd_A(P)$, $qnbd_B(P)$, $qnbd_C(P)$, $qnbd_D(P)$, $qnbd_{A'}(P)$, $qnbd_{B'}(P)$, $qnbd_{C'}(P)$, $qnbd_{D'}(P)$ fully contained in $nbd(a,b)$. Since less than $\frac{d}{8}$ of the nodes in a *qnbd* are faulty with probability 1 (asymptotically), this guarantees that the node will become aware of $\frac{d}{8}+1$ nodes in $nbd(a,b)$ having committed to a (the correct) value, and will also commit to it. The situation is depicted in Fig. 3 for $P \in \{(a-r+l, b+r+1) | 1 \le l \le r\}$, for which $qnbd_A(P)$ lies in $nbd(a,b)$. For all other locations, a similar argument holds. ■

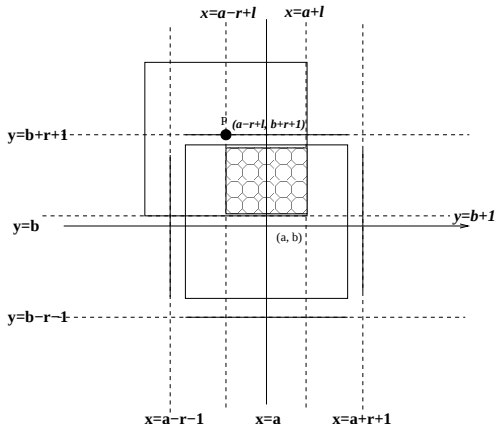

Fig. 3. Node at P has a *qnbd* in $nbd(a,b)$

B. Necessary Conditions for Reliable Broadcast

Theorem 5: If a node not in $nbd(s)$ has at least half faulty neighbors, it can be made to commit to an erroneous value

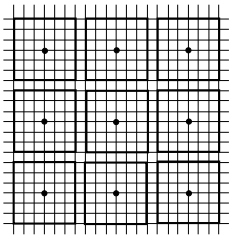

Fig. 4. Division of network into disjoint neighborhoods

with probability at least $\frac{1}{2}$.

Proof: Assume that the message is drawn from $\{0,1\}$. A node which is not an immediate neighbor of the source must rely on messages received from its neighbors.

First, consider any function that takes as argument messages received from all neighbors and outputs one of 0 or 1. Then corresponding to each fault configuration C_1 with $t \geq \frac{d}{2}$ in $nbd(u)$, there is another configuration C_2 with t faults, such that all non-faulty nodes in C_1 are faulty in C_2, while the non-faulty nodes in C_2 were all faulty in C_1. Then, the faulty nodes can modulate their message-sending behavior so that u is unable to distinguish between the case where the correct broadcast value was 0 and configuration was C_1 and the case when the correct value was 1 and the configuration was C_2 (recall that once failure has happened, the faulty nodes can exhibit worst-case behavior). Thus, there are two equally likely possibilities for a given set of received messages, and u cannot expect to choose the correct one with a probability greater than half. If the message can have more than two possible values, it cannot increase the probability of correct choice.

For a more detailed proof, please see [15]. ■

Theorem 6: When failure probability p satisfies: $\frac{1}{2} \leq p \leq 1-\varepsilon$, and $\lim_{n\to\infty} \frac{n}{d} \to \infty$ (this happens when $d = o(n)$), $\lim_{n\to\infty} Pr[$ reliable broadcast fails$] \to 1$.

Proof: Note that in this case, ε can be an arbitrarily small constant, but must be independent of n. Consider a particular node j in the network. Then, if j is non-faulty, but more than half of its neighbors are faulty, reliable broadcast fails with probability at least half. Given that there are d neighbors, and each may fail independently with probability p, let Y_j denote the number of failed neighbors of j. Then, Y_j takes values from $0,1,...,d$, and $E[Y_j] \geq \frac{d}{2}$. Thus $\lfloor E[Y] \rfloor \geq \lfloor \frac{d}{2} \rfloor = \frac{d}{2}$ (since $d = 4r^2+4r$ is always even). Thus, $Pr[Y \geq \frac{d}{2}] \geq Pr[Y \geq \lfloor E[Y] \rfloor] \geq \frac{1}{2}$ (from Lemma 1). Let us call this probability q. When $p \leq 1-\varepsilon$, we have $1-p \geq \varepsilon > 0$. Thus:

$$Pr[\ j \text{ alive; has half+ faulty neighbors}] \geq (1-p)q \geq \frac{\varepsilon}{2} > 0$$

Let us mark out a subset of nodes j such that the neighborhoods of these nodes are all disjoint, as in Fig. 4. Then the number of such nodes that we may obtain is at least $(\lfloor \frac{\lfloor\sqrt{n}\rfloor}{2r+1} \rfloor)^2$ (where $d = 4r^2+4r$), which for large n is at least $\frac{n}{2d}$. Let I_j be an indicator variable that takes value 1 if j is non-faulty but has at least half faulty neighbors, and commits to the wrong value. Then $Pr[I_j = 1] \geq \frac{1}{2}(\frac{\varepsilon}{2}) = \frac{\varepsilon}{4} > 0$, and all I_j's are independent. Note that $\frac{n}{d} \to \infty$. Let X be a random variable indicating the number of non-faulty nodes with at least half faulty neighbors that resultantly commit to the wrong value. Then $E[X] = \sum_j Pr[I_j = 1] \geq \frac{\varepsilon}{4}(\frac{n}{2d}) \to \infty$ (since $\varepsilon > 0$ is a constant independent of n). Thus from the Chernoff Bound in Lemma 2, for any $0 < \beta < 1$ (e.g., set $\beta = \frac{1}{2}$):

$$Pr[X \leq (1-\beta)E[X]] \leq e^{-\frac{\beta^2 E[X]}{2}} (0 < \beta < 1)$$

$$\lim_{n\to\infty} Pr[X > (1-\beta)E[X]] \geq \lim_{n\to\infty} 1 - e^{-\frac{\beta^2 E[X]}{2}} = 1 \quad (\because E[X] \to \infty)$$

Thus, as $n \to \infty$, the number of non-faulty nodes isolated by half or more faulty neighbors, and which commit to the wrong value, will also tend to infinity with probability 1.

■

Theorem 7: In a grid network, when $0 \leq p \leq \frac{1}{2} - \varepsilon$ ($\varepsilon > 0$), and node degree $d(n,p) = 4r(n,p)(r(n,p)+1) < \max\{d_{min}, c\frac{\ln n}{\ln\frac{1}{2p} + \ln\frac{1}{2(1-p)}}\}$ (for suitable constant $c < 1$), reliable broadcast asymptotically fails with probability 1. [2]

Proof: It is immediately obvious that $r(n,p)$ must be at least 1 (i.e. d must at least be d_{min}), else nodes will have no neighbors. We therefore focus on the case $d \geq d_{min} = 8$, and consider the $c\frac{\ln n}{\ln\frac{1}{2p} + \ln\frac{1}{2(1-p)}}$ term. Suppose failure probability $p \leq \frac{1}{2} - \varepsilon$, where ε is an arbitrarily small constant independent of n. Take $f(n) = (\ln n)^2$, and n to be large enough so that $\ln\frac{1}{2p} + \ln\frac{1}{2(1-p)} \geq \frac{1}{\ln n}$. Choose a suitable *constant* $0 < c < 1$ such that $\frac{c}{2}\ln n \leq \ln n - 3\ln\ln n - 2\ln f(n)$, i.e., $\frac{c}{2}\ln n \leq \ln n - 7\ln\ln n$, for sufficiently large n (e.g. $c = 0.9$ would work). Setting $d < c\frac{\ln n}{\ln\frac{1}{2p} + \ln\frac{1}{2(1-p)}}$, for this choice of c, and large enough n, we obtain $d < c(\ln n)^2 < (\ln n)^2$.

Consider a particular node j in the network. Then, if j is non-faulty, but more than half of its neighbors are faulty, reliable broadcast fails with probability at least half (as j commits to an erroneous value with probability at least $\frac{1}{2}$, from Theorem 5). Given that there are d neighbors, and each may fail independently with probability p, let $I_{jk} (1 \leq k \leq d)$ denote the indicator variable corresponding to neighbor k of j (enumerated in some order), such that $I_{jk} = 1$ if k is faulty, and 0 otherwise. Then $Y_j = \sum I_{jk}$ denotes the number of failed neighbors of j. Y takes values from $0,1,...,d$, and $E[Y] = pd$. $Pr[Y_j \geq \frac{d}{2}] = \sum_{i=\frac{d}{2}}^{d} \binom{d}{i} p^i (1-p)^{(d-i)}$. Let us simply consider the event $Y_j = \frac{d}{2}$. Then we can apply the lower bound from Lemma 4. The variables $I_{jk} (1 \leq k \leq d)$ are drawn from $\chi = \{0,1\}$ as per distribution $P = Bernoulli(p)$, and the distribution Q corresponding to $Y_j = \frac{d}{2}$ is $Bernoulli(\frac{1}{2})$ (we shall refer to this as $Q_{\frac{1}{2}}$).

[2] We have a new result extending this theorem to a larger range of p values. For specific details, please see [15].

$|\chi| = 2$, and $\frac{1}{(d+1)^{|\chi|}} = \frac{1}{(d+1)^2} > \frac{1}{\frac{3}{2}d^2} = \frac{2}{3}e^{-2\ln d}$ (since $d \geq 8$). Thus, we obtain:

$$\begin{aligned} q = Pr[Y_j \geq \frac{d}{2}] \geq Pr[Y_j = \frac{d}{2}] &\geq \frac{1}{(d+1)^{|\chi|}} e^{-d(D(Q_{\frac{1}{2}}||P))} \\ = \frac{1}{(d+1)^2} e^{-d(D(Q_{\frac{1}{2}}||P))} &= \frac{2}{3} e^{-d(D(Q_{\frac{1}{2}}||P))-2\ln d} \\ > \frac{2}{3} e^{-(c\frac{\ln n}{\ln\frac{1}{2p}\ln\frac{1}{2(1-p)}})(\frac{1}{2}\ln\frac{1}{2p}+\frac{1}{2}\ln\frac{1}{2(1-p)})-4\ln\ln n} & \end{aligned} \tag{14}$$

(since n large enough to ensure $\ln\frac{1}{2p} + \ln\frac{1}{2(1-p)} \geq \frac{1}{\ln n}$, and $c < 1$, leading to $d < c(\ln n)^2 < (\ln n)^2$)

$$= \frac{2}{3} e^{-\frac{c}{2}\ln n - 4\ln\ln n} \geq \frac{2(\ln n)^3}{3n} \text{ from our choice of } c$$

$$Pr[\ j \text{ alive; at least half } nbd(j) \text{ faulty }] \geq (1-p)q > \frac{1}{2}\frac{2(\ln n)^3}{3n} = \frac{(\ln n)^3}{3n} \tag{15}$$

Let us mark out a subset of nodes j such that the neighborhoods of these nodes are all disjoint, as in Fig. 4. Then, as noted earlier, the number of such nodes that we may obtain is $k \geq \left(\lfloor\frac{\lfloor\sqrt{n}\rfloor}{2r+1}\rfloor\right)^2 \geq \frac{n}{2d}$ for large n. Let I_j be an indicator variable that takes value 1 if j is non-faulty but has at least half faulty neighbors. Then $Pr[I_j = 1] \geq \frac{(\ln n)^3}{3n}$, and all I_j's are independent. We have chosen n large enough to ensure that $\ln\frac{1}{2p} + \ln\frac{1}{2(1-p)} \geq \frac{1}{\ln n}$, i.e. $d \leq c(\ln n)^2$. Let I'_j be an indicator variable that takes value 1 if j is non-faulty but commits to a wrong value. From Theorem 5, we know that if a non-faulty node has half or more faulty neighbors, it will commit to the wrong value with probability at least $\frac{1}{2}$. Thus $Pr[I'_j = 1] \geq \frac{1}{2}Pr[I_j = 1] \geq \frac{(\ln n)^3}{6n}$. Let X be a random variable indicating the number of non-faulty nodes with half or more faulty neighbors that commit to the wrong value. Then $X = \sum I'_j$, and $E[X] = \sum Pr[I'_j = 1] \geq \frac{(\ln n)^3}{6n}\left(\frac{n}{2d}\right) = \frac{(\ln n)^3}{12d} > \frac{\ln n}{12} \to \infty$ (as $d < (\ln n)^2$). Thus we can choose any constant $0 < \beta < 1$ (e.g., set $\beta = \frac{1}{2}$) and apply the Chernoff bound in Lemma 2 to obtain:

$$\lim_{n\to\infty} Pr[X > (1-\beta)E[X]] \geq \lim_{n\to\infty} 1 - e^{-\frac{\beta^2 E[X]}{2}} = 1 \because E[X] \to \infty \tag{16}$$

Thus, as $n \to \infty$, the probability that some non-faulty node(s) fail to commit to the correct value tends towards 1:

$$\lim_{n\to\infty} Pr[\text{ reliable broadcast fails}] \to 1$$

■

VII. Sufficient Condition for Random Networks

We obtain a sufficient condition for a network of n nodes deployed uniformly at random, based on the sufficient condition for the grid network model. To maintain consistency with the grid network formulation, we assume a toroidal region of area $\sqrt{n}$ x $\sqrt{n}$, with n nodes located uniformly at random. The average degree of a node is the average number of the remaining $n-1$ nodes that fall within its neighborhood (recall we are using L_∞ distance metric), i.e., $d_{avg}(n,p) = \frac{(n-1)(2r(n,p))^2}{n} \approx 4r^2(n,p)$ for large n.

Theorem 8: When failure probability $p < \frac{1}{2}$, and $r(n,p) \geq \sqrt{\frac{100\ln n}{\frac{1}{2}-p+\frac{1}{2}\ln\frac{1}{2(1-p)}}}$, reliable broadcast is asymptotically achievable in the random network model with high probability.

Proof: At the outset, we make the observation that if $r(n,p) = \sqrt{n}$, all nodes are neighbors, and trivially broadcast is achievable. Thus this result is of interest only so long as $r(n,p) < \sqrt{n}$.

In light of Fact 1:

$$\begin{aligned} D(Q_{\frac{1}{2}}||p) &= \frac{1}{2}\ln\frac{1}{2p} + \frac{1}{2}\ln\frac{1}{2(1-p)} \\ \geq \frac{1}{2}(1-2p) + \frac{1}{2}\ln\frac{1}{2(1-p)} &= \frac{1}{2} - p + \frac{1}{2}\ln\frac{1}{2(1-p)} \end{aligned} \tag{17}$$

Also, since $0 \leq p < \frac{1}{2}$:

$$0 < \frac{1}{2} - p + \frac{1}{2}\ln\frac{1}{2(1-p)} \leq \frac{1}{2}(1-\ln 2) < \frac{1}{2} \tag{18}$$

Similar to grid networks, we use a notion of quarter-neighborhoods. For a given broadcast instance, we again use relative coordinates by treating the source's coordinates as $(0,0)$. With some abuse of the grid network notation introduced in Section III, we can extend the notion of $nbd(x,y)$, to include all nodes within distance r of point (x,y) (regardless of whether or not there is a node at (x,y)), where x and y are real numbers. The notion of $pnd(x,y)$ is also similarly extended to imply $nbd(x-1,y) \cup nbd(x+1,y) \cup nbd(x,y-1) \cup nbd(x,y+1)$ for all points (x,y).

Note again that in this model, a node's (or point's) coordinates are real numbers. We thus associate eight quarter-neighborhoods with each *node*, with spatial extents as in Table I, except that now x and y must be treated as real numbers. Also, now it is not possible to assert that there are only $2n$ distinct quarter-neighborhoods. Thus, all eight quarter-neighborhoods of a node must be treated as distinct[3], yielding $8n$ quarter-neighborhoods in all.

The quarter-neigborhoods are axis-parallel rectangles of area $r(n,p)(r(n,p)-1) \geq \frac{r^2(n,p)}{2}$ (for $r(n,p) \geq 2$). Then, if $4r^2(n,p) \geq \frac{400\ln n}{\frac{1}{2}-p+\frac{1}{2}\ln\frac{1}{2(1-p)}}$, then we can apply Lemma 6 for all axis-parallel rectangles of area $r(n,p)(r(n,p)-1) \geq \frac{50\ln n}{\frac{1}{2}-p+\frac{1}{2}\ln\frac{1}{2(1-p)}} \geq \frac{100\ln n}{1-\ln 2}$, to obtain that they all have at least $\frac{50\ln n}{\frac{1}{2}-p+\frac{1}{2}\ln\frac{1}{2(1-p)}} - 50\ln n > \frac{25\ln n}{\frac{1}{2}-p+\frac{1}{2}\ln\frac{1}{2(1-p)}} > \frac{50\ln n}{1-\ln 2}$ nodes, with probability at least $1 - \frac{50\ln n}{n} \to 1$.

Thus all such rectangles are *non-empty*. Also:

$$\frac{25\ln n}{\frac{1}{2}-p+\frac{1}{2}\ln\frac{1}{2(1-p)}} \geq \frac{25\ln n}{D(Q_{\frac{1}{2}}||p)} > \frac{8\ln n}{D(Q_{\frac{1}{2}}||p)} \tag{19}$$

[3]Note that distinct does not mean disjoint.

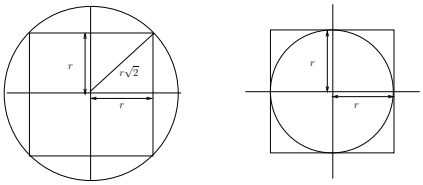

Fig. 5. Relationship between L_∞ and L_2 neighborhoods

Hence all the quarter-neighborhoods have at least $\frac{8\ln n}{D(Q_{\frac{1}{2}}||p)}$ nodes (which is the quarter-neighborhood population in the grid network case). Then using a proof argument similar to Theorem 2, one can prove the following theorem:

Theorem 9: If $p < \frac{1}{2}$, and $r(n,p) \geq \sqrt{\frac{100\ln n}{\frac{1}{2}-p+\frac{1}{2}\ln\frac{1}{2(1-p)}}}$, then

$$\lim_{n\to\infty} Pr[\text{ all } 8n \textit{ qnbds} \text{ have non-faulty majority}] \to 1$$

Thus, one can use a broadcast protocol similar to that for grid networks (a node commits to a value if it is received from half or more nodes in some quarter-neighborhood), and, for all broadcast sources, and instances, the correctness and completeness continue to hold, as follows:

Correctness: Relying on Theorem 9, we can apply a proof argument similar to Theorem 3.

Completeness: The proof uses the an inductive argument similar to the proof of Theorem 4, except that the terms $nbd(x,y), pnd(x,y)$ and quarter-neighborhood must be interpreted as per their re-definition in this section. In the base case, all neighbors of the source (which is at $(0,0)$) commit to the correct value trivially. In the inductive step, one can show that if all nodes in $nbd(x,y)$ (as per extended notation) have comitted to the correct value, all nodes in $pnd(x,y) - nbd(x,y)$ have some *qnbd* contained in $nbd(x,y)$, and can thus commit to the value received from a majority of nodes in this *qnbd*. ■

Since the area within range of a node is $(2r)^2 = 4r^2$ (for the valid domain of r values) in the L_∞ metric, the result indicates that an average node degree d_{avg} of $\frac{400\ln n}{\frac{1}{2}-p+\frac{1}{2}\ln\frac{1}{2(1-p)}}$ suffices for reliable broadcast. Hence the *critical* average node degree $d^{avg}_{critical}$ is $O(\frac{\ln n}{\frac{1}{2}-p+\frac{1}{2}\ln\frac{1}{2(1-p)}})$. [4]

VIII. Conditions in Euclidean Metric

Our results derived for L_∞ metric continue to hold for L_2 metric, with only the constants in the asymptotic notation changing. A similar argument was used in [3].

Lemma 7: If reliable broadcast is achievable in L_∞ for all $r \geq r_{min}$, then it is achievable in L_2 for all $r \geq r_{min}\sqrt{2}$.

Proof: The proof is by contradiction. Suppose that, for a given failure configuration, broadcast is achievable in L_∞ for all $r \geq r_{min}$ but is not achievable for all $r \geq r_{min}\sqrt{2}$ in L_2. Observe that it is possible to circumscribe a L_∞ neighborhood of range r by a L_2 neighborhood of range $r\sqrt{2}$ (Fig. 5). Hence the non-faulty nodes in an L_2 network of transmission range $r\sqrt{2}$ can be made to simulate the operation of nodes in a L_∞ network with range r (as the L_∞ neighborhood is fully contained within the L_2 neighborhood). Also, given that this is a network of known topology, with no address spoofing allowed, the faulty nodes cannot gain any unfair advantage by not simulating the the L_∞ network. This implies that if broadcast is achievable in the L_∞ network of range r , so must it be in the L_2 network of range $r\sqrt{2}$. If there is some $r \geq r_{min}$ for which we can achieve broadcast in the L_∞ network asymptotically, but not in the the L_2 network of range $r\sqrt{2}$, we obtain a contradiction, as achievability in the L_∞ network would imply achievability in the L_2 network. ■

Lemma 8: If reliable broadcast fails in L_∞ for all $r \leq r_{min}$, then it fails in L_2 for all $r \leq r_{min}$.

Proof: The proof is by contradiction. Suppose there is a failure configuration in which broadcast fails in L_∞ for range r, but does not fail in L_2 for range r. Observe that an L_∞ neighborhood of transmission range r circumscribes an L_2 neighborhood of range r (Fig. 5). Thus, for any given failure configuration, if broadcast succeeds in the the L_2 network of range r, so can it in the L_∞ network of radius r, as we could simply make the fault-free nodes in the L_∞ network simulate the behavior of nodes in the L_2 network. This yields a contradiction. ■

IX. Non-Toroidal Networks

We used the assumption that the network is toroidal to avoid edge effects. However, our Byzantine failure results continue to hold even if the network were spread over a non-toroidal rectilinear domain. The necessary conditions continue to hold since the degree of nodes at the edges can be no more more than the degree of nodes towards the center; if reliable broadcast is impossible even with the assumption of equal degree for all nodes, it must certainly be impossible when some nodes (those at the edges) have a smaller degree. The sufficient condition for Byzantine failures continues to hold since the described protocol relies on information from quarter-neighborhoods, and it can be seen that even the nodes at the edges have at least one quarter-neighborhood within the network region.

X. Crash-Stop Failures

As mentioned in Section IV, broadcast in presence of crash-stop failures is equivalent to connectivity under failure. We have also derived results for crash-stop failures in a grid network. These are closely related to the results of [1], discussed in Section IV. However, we prove that, given a *failure* probability p, the critical degree is $\Theta(d_{min} + \frac{\ln n}{\ln\frac{1}{p}})$ for connectivity (please see [15] for details). Our results are more accurate than those of [1] when $p \to 0$, and our necessary condition holds in a certain subdomain where that of [1] does not. To illustrate, when $p = 0$, their condition yields

[4] A more intuitive way of viewing the result is that *critical* degree is $O(\max\{\ln n, \frac{\ln n}{D(Q_{\frac{1}{2}}||P)}\})$.

a degree of the order of $\log n$ while our condition yields a degree d_{min}, which is accurate, since in absence of failure, the minimum transmission range for a non-empty neighborhood suffices. However, both our necessary conditions and those of [1] will fail to apply when p becomes very close to 1. Besides, the necessary condition derived in [1] was actually a necessary condition for coverage, and hence a joint condition for connectivity and coverage. Also relevant to the connectivity issue is analysis presented in [16] regarding the feasible rate in a sensor network, which may potentially be adapted to yield similar connectivity results.

XI. Discussion

An interesting observation is that the form of the results for Byzantine failures is very similar to the results for crash-stop failures/connectivity. For Byzantine failures, we have obtained that the critical node degree for grid networks is $\Theta(d_{min} + \frac{\ln n}{\ln \frac{1}{2p} + \ln \frac{1}{2(1-p)}})$, which may be re-stated as $\Theta(d_{min} + \frac{\ln n}{D(Q_{\frac{1}{2}}||P)})$ where $Q_{\frac{1}{2}}$ denotes the $Bernoulli(\frac{1}{2})$ distribution, P denotes the $Bernoulli(p)$ distribution, and $D(Q||P)$ denotes the *relative entropy* (or Kullback-Leibler distance) between distributions Q and P. Similarly, the node degree for crash-stop failures/connectivity is $\Theta(d_{min} + \frac{\ln n}{\ln \frac{1}{p}})$, and may be viewed as as $\Theta(d_{min} + \frac{\ln n}{\lim\limits_{q \to 1} D(Q||P)})$, where Q is the $Bernoulli(q)$ distribution, and P is the $Bernoulli(p)$ distribution.

Recall that we derive the necessary condition from isolated failure events, and this is found to match the sufficient condition within a constant factor. Thus, possibly failure events involving isolated nodes not receiving correct broadcast may be the dominant failure events [5]. Focusing on these isolated failure events, the obtained expressions for node degree can be explained in the light of Sanov's Theorem [14]. As per Sanov's Theorem, the probability of occurrence of the event-set $\mathcal{E} = \{$ half or more neighbors faulty$\}$ is dominated by the probability of the event in $\mathcal{E}$ closest in relative entropy to the governing fault distribution P. Since we are considering the regime $p < \frac{1}{2}$, the closest event is that of exactly half the neighbors faulty, corresponding to $Q_{\frac{1}{2}}$. In light of this, the critical degree expression for Byzantine failures is quite intuitive. One can similarly explain the crash-stop results.

The necessary and sufficient condition for connectivity in a sensor network where nodes sleep with probability p was shown in [11] to be $\Theta(\frac{\ln(n(1-p))}{1-p})$ (when expressed in our notation) for the case of a randomly deployed network. This problem is equivalent to that of crash-stop failures in random networks. Our sufficient condition for random networks with Byzantine failure probability $p < \frac{1}{2}$ is $O(\frac{\ln n}{\frac{1}{2}-p+\frac{1}{2}\ln \frac{1}{2(1-p)}})$. There is a similarity of form in the two results, and one may interpret the critical node degree as being $O(\max\{\ln n(1-p), \frac{\ln n(1-p)}{D(Q||P)}\})$ where $q = 1$ for the sleeping/crash-stop case in [11], and $q = \frac{1}{2}$ for the Byzantine failure case.

Also note that both our grid network and random network results (for Byzantine failure) have similar structural form, involving a minimum term required for connectivity without disruptive (Byzantine) behavior, and a second term required to ensure broadcast even in presence of failure.

[5] Note that in [17], it was found that the primary disconnection events in non-faulty *random* networks are those involving single isolated nodes.

XII. Conclusions

We considered the problem of reliable broadcast in wireless networks with permanent probabilistic failures, and obtained tight bounds for asymptotic achievability of broadcast in a grid deployment. We also presented a sufficient condition for Byzantine failures in a random network. In recent work, we have also obtained a necessary condition for random networks.

Acknowledgement

We acknowledge an anonymous reviewer of a prior manuscript version whose useful remarks suggesting the extensibility of our grid network sufficiency result to other network models motivated us to work out the sufficient condition for random networks described in Section VII.

References

[1] S. Shakkottai, R. Srikant, and N. Shroff, "Unreliable sensor grids: Coverage, connectivity, and diameter," in *Proc. of Infocom 2003*, 2003.

[2] E. Kreyszig, *Advanced Engineering Mathematics*, 7th ed. John Wiley & Sons, 1993.

[3] C.-Y. Koo, "Broadcast in radio networks tolerating byzantine adversarial behavior," in *Proc. of ACM PODC '04*. ACM Press, 2004, pp. 275–282.

[4] V. Bhandari and N. H. Vaidya, "On reliable broadcast in a radio network," in *Proc. of ACM PODC '05*. ACM Press, 2005, pp. 138–147.

[5] C.-Y. Koo, V. Bhandari, J. Katz, and N. H. Vaidya, "Reliable broadcast in radio networks: The bounded collision case," in *Proc. of ACM PODC '06*. ACM Press, 2006.

[6] T. H. Cormen, C. E. Leiserson, and R. L. Rivest, *Introduction to Algorithms*. MIT Press, 1990.

[7] A. Pelc and D. Peleg, "Broadcasting with locally bounded byzantine faults," *Information Proc. Letters*, vol. 93, no. 3, pp. 109–115, Feb 2005.

[8] ——, "Feasibility and complexity of broadcasting with random transmission failures," in *Proc. of ACM PODC '05*. ACM Press, 2005, pp. 334–341.

[9] S. Kumar, T. H. Lai, and J. Balogh, "On k-coverage in a mostly sleeping sensor network," in *Proc. of MobiCom '04*. New York, NY, USA: ACM Press, 2004, pp. 144–158.

[10] S. Shakkottai, R. Srikant, and N. Shroff, "Correction to unreliable sensor grids: Coverage, connectivity, and diameter," Personal Comm., 2005.

[11] D. Kim, C. Hsin, and M. Liu, "Asymptotic connectivity of low duty-cycled wireless sensor networks," in *Proc. MILCOM*, 2005.

[12] K. Jogdeo and S. M. Samuels, "Monotone convergence of binomial probabilities and a generalization of ramanujan's equation," *The Annals of Mathematical Statistics*, vol. 39, no. 4, pp. 1191–1195, August 1968.

[13] W. Hoeffding, "Probability inequalities for sums of bounded random variables," *Journal of the American Statistical Association*, vol. 58, no. 301, pp. 13–30, Mar. 1963.

[14] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. John Wiley & Sons, Inc., 1991.

[15] V. Bhandari and N. H. Vaidya, "Reliable broadcast in wireless networks with probabilistic failures," Technical Report, CSL, UIUC, Jan. 2007.

[16] X. Liu and R. Srikant, "An information-theoretic view of connectivity in wireless sensor networks," in *Proc. of SECON '04*. Santa Clara, CA: IEEE, Oct. 4-7 2004.

[17] P. Gupta and P. R. Kumar, "Critical power for asymptotic connectivity in wireless networks," in *Stochastic Analysis, Control, Optimization and Applications: A Volume in Honor of W.H. Fleming*, W. M. McEneany, G. Yin, and Q. Zhang, Eds. Boston: Birkhauser, 1998, pp. 547–566.