

# On Reliable Broadcast in a Radio Network

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 an infinite (or finite toroidal) radio network under Byzantine and crash-stop failures. We present bounds on the maximum number of failures that may occur in any given neighborhood without rendering reliable broadcast impossible. We improve on previously proved bounds for the number of tolerable Byzantine faults (presented in a PODC 2004 paper [1]). Our results indicate that it is possible to achieve reliable broadcast if slightly less than one-fourth fraction of nodes in any neighborhood are faulty, and impossible otherwise. We also show that reliable broadcast is achievable with crash-stop failures if slightly less than half the nodes in any given neighborhood may be faulty. In particular, we establish *exact thresholds* under a specific distance metric.

**Index Terms**—Byzantine faults, Crash-stop faults, Broadcast, Fault Tolerance, Radio Network, Broadcast Channel, Possibility/Impossibility

## I. INTRODUCTION

Reliable broadcast in the presence of crash-stop and Byzantine failures is a well-studied problem with numerous practical implications. With the proliferation of wireless networks, there has been interest in the achievability of reliable broadcast in radio networks, which are characterized by a shared wireless medium where every node can talk to all nodes within its transmission radius (deemed as neighbors) and a sent message is heard by all the neighbors. We consider the problem of reliable broadcast in an infinite radio network (with nodes situated on a unit square grid) under Byzantine and crash-stop failures. The results also hold for a finite toroidal network, as boundary anomalies are eliminated. We present bounds on the maximum number of failures that may occur in any given neighborhood (to be formally defined later) without rendering reliable broadcast impossible. For the case of Byzantine failures, we improve on bounds presented in a PODC '04 paper

[1]. We also prove that reliable broadcast is possible with crash-stop failures if slightly less than half the nodes in any neighborhood are faulty. In particular, we establish *exact thresholds* under a specific distance metric.

## II. NETWORK MODEL

We consider the network model described in [2] and [1]. Nodes are located on an infinite grid (each grid unit is a  $1 \times 1$  square). Nodes can be uniquely identified by their grid location  $(x, y)$ . All nodes have a transmission radius  $r$ . A message broadcast by a node  $(x, y)$  is heard by all nodes within distance  $r$  from it (where distance is defined in terms of the particular metric under consideration, and  $r$  is assumed to be an integer). The set of these nodes is termed the neighborhood of  $(x, y)$ . Thus there is an assumption that the channel is perfectly reliable, and a local broadcast is correctly received by all neighbors. We call this the *reliable local broadcast* assumption. In this paper, we consider two distance metrics viz.  $L_\infty$  and  $L_2$ . The  $L_\infty$  metric is essentially the metric induced by the  $L_\infty$  norm [3], 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  $nbu(a, b)$  comprises a square of side  $2r$  with its centroid at  $(a, b)$ . The  $L_2$  metric is induced by the  $L_2$  norm [3], 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  $nbu(a, b)$  comprises nodes within a circle of radius  $r$  centered at  $(a, b)$ .

As in [1], we assume that a node may not spoof another node's identity, and that no collisions are possible, i.e., there exists a pre-determined TDMA schedule that all nodes follow. Such schedules are easily determined for the grid network under consideration [1] (so long as time-optimality is not a concern). We shall further discuss the impact of relaxing these assumptions in Section X. A designated source (that is assumed located at the origin of the grid coordinate system

w.l.o.g.) broadcasts a message with a binary value. The aim is to propagate the correct value to all nodes in the network. We seek to determine the maximum number of faulty nodes  $t$  that may be present in the neighborhood of any given node without rendering reliable broadcast impossible.

### III. RELATED WORK

Reliable broadcast has been extensively studied for networks with point-to-point communication under various connectivity conditions [4]. The classic result of Pease, Shostak and Lamport [5], [6] states that in case of full connectivity, Byzantine agreement with  $f$  faulty nodes is possible if and only if  $n \geq 3f + 1$ . Under more general communication graphs, the requirements for Byzantine agreement are that  $n \geq 3f + 1$ , and the network be at least  $(2f + 1)$ -connected [7]. Byzantine agreement in  $k$ -cast channels has been considered in [8]. However this does not capture the spatially dependent connectivity that characterizes radio networks. Reliable broadcast in radio networks has been studied in [2] and [1]. Crash-stop failures are considered in [2] for finite networks comprising nodes located in a regular grid pattern and algorithms are described for efficient broadcast to the part of the network that is reachable from the source. However this work does not attempt to quantify the number of faults that render some nodes unreachable. In [1], it is shown that for a network of nodes located on an infinite grid of unit squares and having transmission radius  $r$ , reliable broadcast is not achievable for  $t \geq \lceil \frac{1}{2}r(2r + 1) \rceil$  (in both  $L_\infty$  and  $L_2$  metrics). Besides a protocol is described and it is proved that it allows reliable broadcast to be achieved under the following conditions:

- If  $t < \frac{1}{2}(r(r + \sqrt{\frac{r}{2}} + 1)) - 2$ , then reliable broadcast is achieved in the  $L_\infty$  metric.
- If  $t < \frac{1}{4}(r(r + \sqrt{\frac{r}{2}} + 1)) - 2$ , then reliable broadcast is achievable in the  $L_2$  metric.

The considered protocol stipulates that nodes wait till they hear the same value from  $t + 1$  neighbors before they commit to it, and re-broadcast it exactly once for the benefit of other neighbors. Under this protocol, no non-faulty node will ever accept the wrong value. However, there is a possibility of some nodes never being able to decide, and the achievability bounds do not match the impossibility bound, leaving a region of uncertainty.

We consider the network model described above,

and examine the possibility of achieving reliable broadcast under Byzantine and crash-stop failures. For a Byzantine failure model, we present a protocol (utilizing a notion of indirect reports) that allows reliable broadcast to be achieved under the same network model in the  $L_\infty$  metric whenever  $t < \frac{1}{2}r(2r + 1)$ . This exactly matches the impossibility bound of [1], and thus establishes an exact threshold for Byzantine agreement under this model. We also prove that reliable broadcast is achievable under the crash-stop model iff the number of faulty nodes  $t$  in any neighborhood is governed by  $t < r(2r + 1)$  (in the  $L_\infty$  metric). We present informal arguments suggesting that in  $L_2$  i.e. Euclidean distance metric, Byzantine agreement is possible if slightly less than one-fourth of the nodes in any given neighborhood may be faulty, while it is possible to tolerate crash-stop failures that are slightly less than half the neighborhood population. Finally, we consider the issue of tolerable faults when using a simple protocol that does not use indirect reports (i.e. the protocol of [1]). We present an asymptotically tighter bound (than that in [1]) for achievability with Byzantine failures by proving that reliable broadcast is achievable for  $t \leq \frac{2}{3}r^2$  using the simple protocol.

In a very recent work [9], further study of the locally bounded fault model has been undertaken. The paper focuses on arbitrary graphs instead of using a specific network model. It also claims to hold generally for both radio and message-passing networks. However there is an assumption that duplicity (sending different messages to different neighbors) is impossible, which seems to stem from the radio network model. Upper and lower bounds for achievability of reliable broadcast are presented based on graph-theoretic parameters, for arbitrary graphs. However, no exact thresholds are established. The paper considers two algorithms for broadcast. One is the simple algorithm of [1] that they refer to as the Certified Propagation Algorithm (CPA). Another algorithm, termed as the Relaxed Propagation Algorithm (RPA), is informally described and involves a notion of indirect reports similar to the protocol we describe in Section VI. It is shown that RPA is a more powerful algorithm, as there exist graphs for which RPA succeeds but CPA does not. It is also shown that there exist certain graphs in which algorithms that work with knowledge of topology succeed in achieving reliable broadcast, while those that lack this knowledge fail to do so. Our work differs substantially from theirs, in that we focus on a specific network model and obtain

an exact threshold for byzantine as well as crash-stop fault-tolerance. We also present a specific algorithm for byzantine agreement in the considered model, which localizes the circulation of indirect reports, and thus reduces communication overhead.

#### IV. NOTATION/TERMINOLOGY

We briefly describe here notation and terminology that shall be used in this paper. Nodes are identified by their grid 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  $nb(x, y)$ . For succinct description, we define a term  $pnb(x, y)$  where  $pnb(x, y) = nb(x - 1, y) \cup nb(x + 1, y) \cup nb(x, y - 1) \cup nb(x, y + 1)$ . Intuitively  $pnb(x, y)$  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. Besides, throughout this paper, a non-faulty node shall be variously alluded to as an honest or correct node, while a node exhibiting byzantine failure shall occasionally be referred to as a malicious node.

#### V. BYZANTINE AGREEMENT IN A RADIO NETWORK

Radio networks present a special case for the Byzantine agreement problem due to the broadcast nature of the channel. In the absence of address-spoofing and deliberate collisions (discussed further in Section X), this significantly simplifies the problem, and relaxes the requirements for agreement. Under our assumptions (also in [1]), if a node transmits a value, all its neighbors hear the transmission, and are certain of the identity of the sender. The transmitting node is thus incapable of duplicity, because if it were to attempt sending contradicting messages, they would be heard by all its neighbors, and its duplicity would stand detected. Thus any protocol could stipulate that if the neighbors of a node hear it transmitting multiple contradictory versions of a message, they should accept only the first message, and ignore the rest. Thus, in a fully connected network, it is possible to tolerate an arbitrary number of Byzantine faults. In a more general network, the absence of duplicity implies a relaxation of the requirements proved in [7] in that it is no longer required that  $n \geq 3f + 1$  for tolerating  $f$  faults. If only  $f$  Byzantine faults were allowed in the whole network, the necessary and sufficient condition for reliable broadcast would be exactly the same as the connectivity condition of [7] viz. that the graph be  $(2f + 1)$ -connected. Since we consider a model in which an adversary may place upto  $t$  faults in any single neighborhood, a general *sufficient* condition that

may be stated for an arbitrary network graph  $G = (V, E)$  is that for each pair of nodes  $(v_1, v_2)$  s.t.  $v_1, v_2 \in V$ , either  $(v_1, v_2) \in E$ , else  $\exists S \subseteq V$  such that the adversary may place at most  $f$  faults in  $S$  without violating the constraint, and  $v_1$  be connected to  $v_2$  via  $2f + 1$  node-disjoint paths that lie entirely within  $S$ . Note that this requires knowledge of network topology. The protocol we present in this paper is based on a localized variant of this sufficient condition.

#### VI. RELIABLE BROADCAST WITH BYZANTINE FAILURES

As discussed in Section III, it was proved in [1] that reliable broadcast is impossible in  $L_\infty$  as well as  $L_2$  metrics if  $t \geq \lceil \frac{1}{2}r(2r + 1) \rceil$ . We prove the following:

**THEOREM 1:** If  $t < \frac{1}{2}r(2r + 1)$ , reliable broadcast is achievable in the  $L_\infty$  metric.

This is an exact match to the impossibility bound for  $L_\infty$ , and thus establishes the threshold for achieving reliable broadcast in the square grid network under consideration. We present a protocol that achieves this objective. Without loss of generality we assume the message to comprise a binary value (say 0 or 1). A node that is not the source is said to *commit* to a value when it becomes certain that it is indeed the value originated by the source. The protocol requires maintenance of state by each node pertaining to nodes within its three-hop neighborhood. This state may be reduced further by earmarking exact messages that a node should lookout for, and this shall become clear from our constructive proof for the viability of reliable broadcast with  $t < \frac{1}{2}r(2r + 1)$ . However, at a basic level, the protocol operates as follows:

- Initially, the source broadcasts the message.
- Each neighbor  $i$  of the source re-broadcasts the first value it heard from the source (and committed to) once in a *COMMITTED*( $i, v$ ) message.
- Hereafter, the following protocol is followed by each node  $j$  (including those involved in the previous two steps):

On receipt of a *COMMITTED*( $i, v$ ) message from neighbor  $i$ , record the message, and broadcast a *HEARD*( $j, i, v$ ) message.

On receipt of a *HEARD*( $k, i, v$ ) message from a neighbor  $k$ , record the message, and broadcast a *HEARD*( $j, k, i, v$ ) message.

On receipt of a *HEARD*( $l, k, i, v$ ) message, record

the message, and broadcast a  $HEARD(j, l, k, i, v)$  message.

On receipt of a  $HEARD(g, l, k, i, v)$  message, record the message, but do not re-propagate.

A node  $j$  commits to a value  $v$  if it reliably determines that at least  $t + 1$  nodes lying in some single neighborhood have committed to  $v$ . A node is said to have reliably determined the value committed to by node  $i$  if:

- $i$  is its neighbor, and so  $j$  heard  $COMMIT(i, v)$  directly. In this case, there is no cause for doubt as to what value was committed to by node  $i$ , since no other node is capable of spoofing  $i$ 's address, and collisions are ruled out.
- $j$  heard indirect reports of  $i$  having committed to a particular value  $v$  through  $t + 1$  node-disjoint paths that all lie within *some single neighborhood*. The indirect reports are obtained via the  $HEARD$  messages that propagate via upto three intermediate nodes, and the path information is obtained from these messages (as each forwarding node affixes its identifier to the message). Observe that as the  $t + 1$  node-disjoint paths all lie within a single neighborhood, and as no more than  $t$  nodes in the neighborhood may be faulty, all the  $(t + 1)$  paths cannot have a faulty node each, and it is therefore impossible for the node to arrive at a wrong conclusion by following this rule.

**THEOREM 2: (Correctness)** No node shall commit to a wrong value by following the above rule.

*Proof:* The proof is by contradiction. Consider the first node, say  $j$ , that makes a wrong decision to commit to value  $v$ . This implies it reliably determined that  $t + 1$  already committed nodes lying in some single neighborhood  $N_1$  had committed to  $v$ . Since reliable determination of a node  $i$  having committed to a value  $v$  involves hearing  $i$  directly or hearing indirect reports (that  $i$  committed to  $v$ ) via at least  $t + 1$  node-disjoint paths lying in some single neighborhood  $N_2$ , and since the number of faults in  $N_2$  may be at most  $t$ , it implies that all these paths cannot have relayed the wrong value, and so  $v$  must indeed be the value committed to by  $i$ . Thus no node can make a wrong determination of what value each of the  $t + 1$  nodes in  $N_1$  committed to; they must all indeed have committed to  $v$ . Since  $j$  is the first node to make a wrong decision, the  $t + 1$

nodes could not have made a wrong decision. Also, all of these nodes cannot be faulty, as no more than  $t$  nodes in any neighborhood may exhibit Byzantine failure. Thus  $v$  must indeed be the correct value. ■

**THEOREM 3: (Completeness)** Each node is eventually able to commit to the correct value.

*Proof:* We prove that each node will be able to meet the conditions stipulated by the protocol for committing to the correct value. The proof also clarifies the operation of the protocol, and in fact would allow one to stipulate exactly which messages each node should act upon (given that the topology of the network is completely known), thereby reducing the state maintained at each node. The essence of the proof lies in showing that each node  $j$  (except the direct neighbors of  $(0, 0)$ ) is connected to at least  $2t + 1$  nodes that lie in some single neighborhood  $N_1$ , such that the connectivity to each such node is through  $2t + 1$  node-disjoint paths that all lie in some neighborhood  $N_2$ , and the nodes in  $N_1$  are able to commit to the correct value before node  $j$  has done so.

The proof proceeds by induction.

*Base Case:*

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

*Proof of Inductive Hypothesis:*

We show that each node in  $pnbd(a, b)$  is able to reliably determine the value committed to by  $2t + 1$  nodes in  $nbd(a, b)$ . Since no more than  $t$  of these can be faulty, this guarantees that the node will become aware of  $t + 1$  nodes in  $nbd(a, b)$  having committed to a (the correct) value, and will also commit to it. In order to show this, we show that each node is connected to at least  $2t + 1$  nodes in  $nbd(a, b)$  either directly, or through  $2t + 1$  node disjoint paths that all lie entirely within some single neighborhood. Thus at least  $t + 1$  of these




Fig. 1. Nodes in  $nb_d(a, b)$  whose committed values  $P$  can reliably determine




Fig. 2. Nodes in  $nb_d(a, b)$   $P$  can hear directly




Fig. 3. Nodes in  $nb_d(a, b)$  to which  $P$  has sufficient connectivity




Fig. 4. Nodes that are neighbors of both  $P$  and  $N$

paths are guaranteed to be fault-free and shall allow communication of the correct value.

We show this for a corner node in  $pnbd(a, b)$  i.e. the node marked  $P$  (which is located at  $(a-r, b+r+1)$ ) in Fig. 1, which represents the worst case. For all other nodes in  $pnbd(a, b)$ , the condition can be seen to be achieved via a similar argument, but even more easily. We omit the proof for the sake of brevity.

We show that node  $P$  is able to reliably determine the values committed to by the nodes in the shaded region in Fig. 1 which comprises  $r(2r+1) + 2r > r(2r+1)$  nodes. The first observation is that  $P$  can directly hear the nodes in the shaded region in Fig. 2, and so is certain of the value they committed to. We now explicitly prove existence of suitable node-disjoint paths for nodes that lie in the upper triangular region shaded in Fig. 3. Consider

a node  $N$  located at  $(a+p, b+q)$  (Fig. 4). Observe that  $q \geq p \geq 1$  in this region. We show the existence of  $r(2r+1)$  node-disjoint paths between  $N$  and  $P$ , that all lie within the same single neighborhood (centered at  $(a, b+r+1)$ , and indicated by the square with dark outline in Fig. 5). The region marked  $A$  comprises  $\{(x, y) | (a+p-r) \leq x \leq a; (b+1) \leq y \leq (b+q+r)\}$ , and nodes in this region are neighbors of both  $N$  and  $P$ . Thus, there are  $(r-p+1)(r+q)$  paths of the form  $N \rightarrow A \rightarrow P$  that comprise one intermediate node each. The region  $B_1$  comprises  $\{(x, y) | (a+1) \leq x \leq (a+p-1); (b+1) \leq y \leq (b+q+r)\}$ , and falls in  $nb_d(N)$  (recall that  $N$  is located at  $(a+p, b+q)$ ). The region  $B_2$  comprises  $\{(x, y) | (a+1-r) \leq x \leq (a+p-1-r); (b+1) \leq y \leq (b+q+r)\}$ , and falls in  $nb_d(P)$ . As may be seen,  $B_2$  is obtained by a translation of  $B_1$  to the left by  $r$  units. Thus there is a one-to-one correspondence between a point  $(x, y)$  in  $B_1$  and a point




Fig. 5. Construction depicting node-disjoint paths between N and P

$(x-r, y)$  in  $B_2$ , such that the points in each pair are neighbors. This yields  $(p-1)(r+q)$  paths of the form  $N \rightarrow B_1 \rightarrow B_2 \rightarrow P$ .

Region  $C_1$  comprises  $\{(x, y) | (a+p+1) \leq x \leq (a+r); (b+q+1) \leq y \leq (b+r+1)\}$  and thus falls within  $nbd(N)$ . Region  $C_2$  comprises  $\{(x, y) | (a+p+1-r) \leq x \leq a; (b+q+1+r) \leq y \leq (b+1+2r)\}$  and falls within  $nbd(P)$ . It may be seen that there is a one-to-one correspondence between any point  $(x, y)$  in  $C_1$  and point  $(x-r, y+r)$  in  $C_2$ , with the paired points being neighbors. Hence there exist  $(r-p)(r-q+1)$  paths of the form  $N \rightarrow C_1 \rightarrow C_2 \rightarrow P$  that comprise two intermediate nodes each. Region  $D_1$  comprises  $\{(x, y) | (a+p) \leq x \leq (a+r+p-q); (b+r+q-p+2) \leq y \leq (b+r+q+1)\}$ , and falls in  $nbd(N)$ . Region  $D_2$  comprises  $\{(x, y) | (a+1) \leq x \leq (a+p); (b+1+r+q) \leq y \leq (b+1+2r)\}$ . Region  $D_3$  comprises  $\{(x, y) | (a+1-r) \leq x \leq (a+p-r); (b+1+r+q) \leq y \leq (b+1+2r)\}$ , and

falls in  $nbd(P)$ . We note that regions  $D_1, D_2$  and  $D_3$  have exactly the same number of nodes each. Besides, the regions  $D_1$  and  $D_2$  are mutually located in a manner that each node in  $D_2$  is a neighbor of each node in  $D_1$  (maximum distance between any two nodes  $< r$ ). Hence, any one-to-one pairing of nodes in  $D_1$  with nodes in  $D_2$  is valid. Further, a node located at  $(x, y)$  in  $D_2$  has a one-to-one correspondence with a node  $(x-r, y)$  in  $D_3$ . Thus, there are  $p(r-q+1)$  paths of the form  $N \rightarrow D_1 \rightarrow D_2 \rightarrow D_3 \rightarrow P$  that comprise three intermediate nodes each (Fig. 5). Thus the  $r(2r+1)$  node-disjoint paths are obtained.

For the nodes in the lower region of Fig. 3, a similar construction will yield the required paths, as indicated in Fig. 6.

Observe that the inductive hypothesis along with the base case suffice to show that every honest node




Fig. 6. Indication of how construction shall proceed for lower region

will eventually commit to the correct message, since starting at  $(0,0)$ , one can cover the entire infinite grid by moving up, down, left and right. Thus the neighborhood of every grid point can be shown to have decided i.e. every honest node will have decided on the correct value.

We note that the connectivity condition proved above is also sufficient to prove that upto  $2t < r(2r+1)$  crash-stop failures are tolerable in  $L_\infty$  metric. We shall elaborate further in Section VII.

## VII. CRASH-STOP FAILURES

We first note that when only crash-stop failures are admissible, no special protocol is required. Each node that receives a value, commits to it, re-broadcasts it once for the benefit of others, and then may terminate local execution of the protocol. Thus the sole criterion for achievability is reachability. In this failure mode, we establish an exact threshold for tolerable faults in  $L_\infty$  metric.

**THEOREM 4:** If  $t \geq r(2r+1)$ , it is impossible to achieve reliable broadcast in  $L_\infty$  metric.

*Proof:* We present a construction with  $t = r(2r+1)$  that renders reliable broadcast impossible. Consider the network in Fig. 7. The nodes in the designated region  $\{(x,y) | a \leq x < a+r\}$  are all faulty while all other nodes are correct. As may be seen, the maximum number of faulty nodes in any given neighborhood is  $\leq r(2r+1)$ . However this configuration partitions all nodes in the half-plane  $x \geq a+r$  from the source and they are unable to receive the broadcast.

**THEOREM 5:** If  $t < r(2r+1)$ , it is possible to achieve reliable broadcast in  $L_\infty$  metric.

*Proof:* One possible proof proceeds from the proof of Theorem 1, as was noted earlier. Since, we showed that each node is connected to each of  $r(2r+1)$  already committed nodes lying in some single neighborhood, via  $r(2r+1)$  node-disjoint paths that all lie within some single neighborhood, it follows that upto  $t < r(2r+1)$  crash-stop faults may be tolerated, as each node would still be connected to at least one non-faulty committed node, via at least one fault-free path. However, we also present a simpler proof that indicates achievability of reliable broadcast. This proof presents a clearer picture of the progress of the broadcast in the network. The proof is by induction, similar to the inductive argument for Byzantine agreement.




Fig. 7. Network Partition due to Crash Stop Failures

*Base Case:*

When  $(0,0)$  initially broadcasts the message, all correct nodes in  $nbd(0,0)$  hear it directly, and thereby receive the broadcast.

*Inductive Hypothesis:*

If all correct nodes in  $nbd(a,b)$  have received the broadcast, then all correct nodes in  $pnbd(a,b)$  will also receive the broadcast.

*Proof of Inductive Hypothesis:*

Consider the situation as in Fig. 8. All correct nodes in  $nbd(a,b)$  (depicted by square ABCD) have received the broadcast. We consider the partition of ABCD into two rectangles by the horizontal axis through  $(a,b)$ . These regions are depicted as ABFE and EFCD in Fig. 8. The nodes on the partitioning axis i.e. on EF may be included in any one region or split between the two. It does not affect the proof, as these nodes do not play a role in the proof argument. A similar partitioning by the vertical axis through  $(a,b)$  yields AGHD and GBCH, with nodes along GH assigned arbitrarily to either region. Since the number of faulty nodes in ABCD  $< r(2r+1)$ , one of the regions ABFE and EFCD has  $\leq \frac{1}{2}r(2r+1) = r^2 + \frac{r}{2}$  faults i.e. strictly less than  $r(r+1)$  faults<sup>1</sup>. Similarly one of the regions AGHD and GBCH has  $\leq \frac{1}{2}r(2r+1) = r^2 + \frac{r}{2}$  faults. Without loss of generality we assume that the

<sup>1</sup>If  $t$  items are split between two regions, one will get  $\leq \frac{t}{2}$  and the other will get  $\geq \frac{t}{2}$ .




Fig. 8. Reliable Broadcast Propagation: Stage 1




Fig. 9. Reliable Broadcast Propagation: Stage 2

regions satisfying the condition are ABFE and AGHD. Then every node in  $\{(x, a+r+1) | (a-r \leq x \leq a+r)\}$  i.e. along line segment PQ in the figure has at least  $r(r+1)$  neighbors in  $nbnd(a, b)$  and these neighbors fall entirely in region ABFE. Given that the number of faults in ABFE is strictly less than  $r(r+1)$  each node on PQ is able to hear the broadcast from at least one correct neighbor in  $nbnd(x, y)$ . By a similar argument, every node in  $\{(a-r-1, y) | (b-r \leq y \leq b+r)\}$ , i.e., along segment VW, has at least  $r(r+1)$  neighbors in AGHD, and is thus able to receive the broadcast from at least one correct neighbor in  $nbnd(a, b)$ .

Given that ABFE has strictly less than  $r(r+1)$  faulty nodes, it follows that GBFO (being a subset of ABFE) also has strictly less than  $r(r+1)$  faults. Thus each node in  $\{(a+r+1, y) | (b \leq y \leq b+r)\}$  (segment RR') has at least one correct neighbor belonging to  $nbnd(a, b)$  and can receive the broadcast. By a similar argument, every node in  $\{(x, b-r-1) | (a-r \leq x \leq a)\}$  (segment TT') is able to receive the broadcast. Therefore all those nodes belonging to  $pnbd(a, b) - nbnd(a, b)$  that lie along in these regions (depicted by the dark line segments in Fig. 8) receive the broadcast. We know need to show that the remaining nodes will also be able to do so. These remaining nodes are the ones along line segment U'U and segment S'S. We explicitly consider the nodes along segment U'U. The same argument holds for S'S.




Fig. 10. Illustrating an Approximate Argument for Euclidean Metric

Now consider the nodes in the shaded region  $\{(x, y) | a \leq x \leq a+r, b-r \leq y < b\}$ . If even one of these nodes is correct, then the nodes along U'U are guaranteed to receive the broadcast. If all these nodes are faulty then these faulty nodes number  $r^2+r$ . Therefore, if we consider the neighborhood centred at  $(a, b-r-1)$  (Fig. 9), only  $r^2$  more nodes can be faulty in this entire neighborhood apart from those in the shaded region. This number of faulty nodes is not sufficient to completely partition a correct node in WH'T'T from all correct nodes in TT'J'J. Then at least one correct node in region TT'J'J should be able to hear from at least one correct node in region WH'T'T, and in turn all other correct nodes in TT'J'J should be able to receive the broadcast. Similarly, at least one correct node in region U'UK'K should be able to hear from at least one correct node in TT'J'J, and in turn all correct nodes along U'U should be able to receive the broadcast. A symmetric argument holds for the nodes along S'S.

Thus, if all nodes in  $nbnd(a, b)$  receive the broadcast, then all nodes in  $pnbd(a, b)$  also receive the broadcast. Since the considered failures are only of crash-stop kind, the received value is guaranteed to be correct. ■

## VIII. RELIABLE BROADCAST IN EUCLIDEAN METRIC

We now briefly consider the issue of reliable broadcast in the  $L_2$  i.e. Euclidean metric. We refrain from establishing exact thresholds as it is difficult to precisely determine lattice points falling in areas bounded by circular arcs. We do however present






Fig. 11. Approximate Construction depicting Node-Disjoint Paths (PQ from Fig. 10 rotated to x-axis)

informal arguments that suggest that reliable broadcast in  $L_2$  is definitely achievable if slightly less than one-fourth fraction of nodes in any neighborhood exhibit Byzantine faults. We work with the value  $t < 0.23\pi r^2$ . The basis for the argument is that for sufficiently large  $r$ , the number of nodes that lie in various subregions (having area  $A$ ) of a circle of radius  $r$  (elaborated later) are approximately  $A \pm O(r)$ . Thus, we expect the argument to hold well for large values of  $r$ . The argument proceeds by induction, as in the previous section.

#### Base Case:

All honest nodes in  $nbd(0,0)$  are able to commit to the correct value. This follows trivially since they hear the origin directly.

#### Inductive Hypothesis:

If all honest neighbors of a node located at  $(a,b)$  are able to commit to the correct value, then all honest nodes in  $pnbnd(a,b)$  are able to commit to the correct value.

#### Justification of Inductive Hypothesis:

We show that each node in  $pnbnd(a,b)$  should be able to reliably determine the value committed to by  $2t + 1$  nodes in  $nbd(a,b)$ . Since no more than  $t$  of these can be faulty, this would guarantee that the node will become aware of  $t + 1$  nodes in  $nbd(a,b)$  having committed to a (the correct) value, and will also commit to it. In order

to show this, we show that each node is connected to at least  $2t + 1$  nodes in  $nbd(a,b)$  either directly, or through  $2t + 1$  node disjoint paths that all lie entirely within some single neighborhood. Thus at least  $t + 1$  of these paths are guaranteed to be fault-free and shall allow communication of the correct value.

Consider the node at  $(a,b)$ , as in Fig. 10. Let  $d$  be the distance between the node at  $(a,b)$  (we call it node  $N$ ) and any node in  $(pnbnd(a,b) - nbd(a,b))$  (we call it node  $Q$ ). Then  $d \leq r + 1$ . Consider the half-neighborhood of  $(a,b)$  demarcated by the medial axis perpendicular to  $NQ$  (not counting the points falling on the medial axis). Then, as the number of faults  $t < 0.23\pi r^2$ , it implies that there must be at least  $2t + 1$  nodes lying in this half-neighborhood. We attempt to quantify the number of node-disjoint paths between any node  $P$  in this half-neighborhood, and the node  $Q$ . Observe that in the worst case, the distance  $D$  between  $P$  and  $Q$  is  $\approx r\sqrt{2}$ . We consider the situation in Fig. 11 with  $PQ$  rotated to the horizontal axis. The distance  $PQ$  is  $r\sqrt{2}$ . We attempt to construct node-disjoint paths that all lie within the neighborhood centered at  $M$  (the midpoint of  $PQ$ ). The set of nodes marked  $A$  are common neighbors of  $P$  and  $Q$  and constitute two-hop  $PQ$  paths ( $P \rightarrow A \rightarrow Q$ ). A set of three-hop paths  $P \rightarrow B_1 \rightarrow B_2 \rightarrow Q$  and  $P \rightarrow D_1 \rightarrow D_2 \rightarrow Q$  is also formed where each point  $(x,y)$  in region  $B_1$  has a corresponding point  $(x+r,y)$  in  $B_2$ . Similarly there is a set of three-hop paths  $P \rightarrow C_1 \rightarrow C_2 \rightarrow Q$  and  $P \rightarrow D_1 \rightarrow D_2 \rightarrow Q$ , since each point  $(x,y)$  in  $C_1$  ( $D_1$ ) has a corresponding point  $(x + \frac{1}{\sqrt{2}}r, y)$  in  $C_2$  ( $D_2$ ). Finally, there is a set of paths  $P \rightarrow E_1 \rightarrow E_2 \rightarrow Q$  such that each point in  $E_1$  has a one-to-one correspondence with its mirror image with respect to axis  $OO'$  which lies in  $E_2$ . The number of such paths is approximately equal to the sum of the areas  $A, B_1, C_1, D_1$ , and  $E_1$  which turns out to be approximately  $1.47r^2 = 0.47\pi r^2 > (2(0.23\pi r^2) + 1)$ . Thus approximately  $0.23\pi r^2$  Byzantine faults may be tolerated.

We also argue similarly that reliable broadcast is not possible if  $t \geq 0.3\pi r^2$ . The argument is based on a construction identical to that presented in [1] for  $L_\infty$ , which is depicted in Fig. 12. As already argued in [1], this arrangement of faults renders reliable broadcast impossible. Note that the maximum number of faults lying in any single neighborhood is given by the number of faulty nodes in the circled region (Fig. 12). The relevant area under the circle is approximately




Fig. 12. Impossibility Construction for Byzantine Failures in Euclidean metric

$0.6\pi r^2$ , and we expect approximately  $0.6\pi r^2 \pm O(r)$  nodes to lie in it. Of these around  $0.3\pi r^2 \pm O(r)$  are expected to be faulty. This concludes the argument that if  $t \geq 0.3\pi r^2$  (approximately), reliable broadcast would be unachievable. Thus the critical threshold for  $L_2$  metric seems to lie between a 0.23 and a 0.3 fraction i.e. close to one-fourth fraction of faults.

Observe that the above argument also leads to the conclusion that upto  $2t = 0.46\pi r^2$  crash-stop failures may be tolerated, while around  $0.6\pi r^2$  failures would render reliable broadcast impossible. Thus, for crash-stop failures, the threshold is expected to be somewhere around half the number of nodes in a neighborhood.

## IX. RELIABLE BROADCAST WITH A SIMPLER BYZANTINE PROTOCOL

We present bounds for tolerable faults when an extremely simple protocol (described in [1]) is used. In this protocol, initially the source transmits the value, and its immediate neighbors are able to commit to that value instantly. They then re-broadcast the value committed to and terminate protocol operation. Any other node that has heard the same value reported by at least  $t + 1$  neighbors, commits to it, re-broadcasts it, and then terminates. Thus the protocol proceeds till either all nodes have terminated, or a situation is reached where no further progress is possible. We present an asymptotically tighter bound for the number of tolerable Byzantine faults in the  $L_\infty$  metric (using this protocol) than that presented in [1] viz. we claim and prove that reliable broadcast is always possible for  $t \leq \frac{2}{3}r^2$  which dominates the bound of  $t < \frac{1}{2}(r(r + \sqrt{\frac{r}{2}} + 1)) - 2$ , proved in [1], for all sufficiently large  $r$ .

**THEOREM 6:** If  $t \leq \frac{2}{3}r^2$ , it is possible to achieve reliable broadcast, in the  $L_\infty$  metric, with the given protocol.

The proof is presented in the Appendix.

## X. IMPACT OF ADDRESS-SPOOFING AND COLLISIONS

The presence of a broadcast channel introduces numerous difficulties by way of the possibility of a malicious node spoofing another node's address and sending spurious messages under guise, as well as the possibility of disruption of communication via deliberate collisions. The results presented in this paper assume that neither problem exists.

When the adversary has control over low-level networking functions, reliable broadcast is extremely difficult to achieve. If address spoofing is allowed, any malicious node may attempt to impersonate any honest node. Similarly, reliable broadcast is rendered impossible if the adversary can cause an unbounded number of collisions, since a faulty node can cause collision with any transmission made by a good node in its vicinity. When the number of collisions is bounded, it may be possible to come up with protocols that achieve reliable broadcast. If the adversary uses collisions to merely disrupt communication, the problem is trivially solved by re-transmitting messages a sufficient number of times. However, the adversary might use it to send contradicting messages to different parts of the network (a situation briefly discussed in [1]). This situation might be remediable via a protocol that involves consultation between the neighbors of a node as to the value they heard it transmit, as well as any detected collisions, and requires further investigation.

## XI. CONCLUSIONS

We have presented results regarding the number of Byzantine and crash-stop failures that may be tolerable in a radio network without rendering reliable broadcast impossible. We have considered an adversarial model where the adversary is free to choose faulty nodes as long as the placement satisfies the constraint that no neighborhood has more than  $t$  faults. Another useful model to consider would be that of random failure, whereby each node has a probability of failure  $p_f$ , and nodes fail independently of each other. Observe that in case of crash-stop failures, the problem is similar to the problem of site percolation [10].

Another aspect that requires further attention is

that of efficient implementation of a reliable broadcast service in a real wireless network. In the presence of channel errors etc., the basic reliable local broadcast requirement is by itself not trivial to achieve. A mechanism for reliable broadcast in a multi-hop mobile network is described in [11]. However, only temporary and non-Byzantine node failures are taken into account, and the mechanism primarily relies on a clustering scheme with unicast messages (where link errors are handled via retransmissions). There is need for further work on efficient Byzantine fault-tolerant protocols for multi-hop wireless networks, in order to bridge the gap between theory and practice.

#### REFERENCES

- [1] C.-Y. Koo, "Broadcast in radio networks tolerating byzantine adversarial behavior," in *Proceedings of the twenty-third annual ACM symposium on Principles of distributed computing*. ACM Press, 2004, pp. 275–282.
- [2] E. Kranakis, D. Krizanc, and A. Pelc, "Fault-tolerant broadcasting in radio networks," *J. Algorithms*, vol. 39, no. 1, pp. 47–67, 2001.
- [3] E. Kreyszig, *Advanced Engineering Mathematics*, 7th ed. John Wiley & Sons, 1993.
- [4] H. Attiya and J. Welch, *Distributed Computing*. McGraw-Hill, 1998.
- [5] M. Pease, R. Shostak, and L. Lamport, "Reaching agreement in the presence of faults," *J. ACM*, vol. 27, no. 2, pp. 228–234, 1980.
- [6] L. Lamport, R. Shostak, and M. Pease, "The byzantine generals problem," *ACM Trans. Program. Lang. Syst.*, vol. 4, no. 3, pp. 382–401, 1982.
- [7] D. Dolev, "The byzantine generals strike again," *J. Algorithms*, vol. 3, no. 1, pp. 14–30, 1982.
- [8] J. Considine, L. A. Levin, and D. Metcalf, "Byzantine agreement with faulty majority using bounded broadcast," *CoRR*, vol. cs.DC/0012024, 2000.
- [9] A. Pelc and D. Peleg, "Broadcasting with locally bounded byzantine faults," *Information Processing Letters*, vol. 93, no. 3, pp. 109–115, Feb 2005.
- [10] G. Grimmett, *Percolation*. Springer-Verlag, New York, 1989.
- [11] E. Pagani and G. P. Rossi, "Providing reliable and fault tolerant broadcast delivery in mobile ad-hoc networks," *MONET*, vol. 4, no. 3, pp. 175–192, 1999.

#### APPENDIX

##### *Proof of Theorem 6 (by induction)*

###### *Base Case:*

All honest nodes in  $nbd(0,0)$  are able to commit to the correct value. This follows trivially since they hear the origin directly.

###### *Inductive Hypothesis:*

If all honest neighbors of a node located at  $(a,b)$

i.e. all honest nodes in  $nbd(a,b)$  are able to commit, then all honest nodes in  $pnb(a,b)$  are able to commit.

###### *Proof of Inductive Hypothesis:*

A *sufficient* condition for a node to be able to commit to the correct message value is that at least  $2t + 1 = \frac{4}{3}r^2 + 1$  of its neighbors must have committed and broadcast their committed value before it. Assume that all honest neighbors of node  $(a,b)$  have arrived at a decision. Then after all these nodes have broadcast their committed value, a certain number of other nodes will *definitely* be able to commit as the sufficient condition is satisfied for them. We consider a subset of these nodes which are indicated in Fig. 13 i.e.  $2\lceil \frac{r}{2} \rceil + 1$  such nodes along each edge of the central square are definitely able to commit, for all  $r > 1$ . That these nodes are able to commit is evident by observing that the number of committed neighbors of these nodes is  $\geq (r + 1 + \frac{r}{2})r > \frac{3}{2}r^2 + r > \frac{4}{3}r^2 + 1 = 2 \cdot \frac{2}{3}r^2 + 1$  (shaded region in Fig. 13). Once these nodes have broadcast their committed value, the adjacent row of  $2\lceil \frac{r}{2} \rceil + 1$  nodes (Fig. 14) will be able to commit and so on, till the *stack* of committed nodes adjoining each edge of the central square reaches a size of  $\lfloor \frac{r}{3} \rfloor$  rows. This may be seen as follows: we have already argued that row 1 will be able to commit. Given that *row* $(i - 1)$  has committed, *row*  $i$  can commit if  $(\lceil \frac{3}{2}r \rceil + 1)(r + 1 - i) + (i - 1)(2\lceil \frac{r}{2} \rceil + 1) + (i - 1)(\lceil \frac{r}{2} \rceil - i + 1) \geq \frac{4}{3}r^2 + 1$ . This condition holds for all  $i \leq \lfloor \frac{r}{\sqrt{6}} \rfloor$ , when  $r \geq 2$ . This implies that the stack can grow to at least  $\frac{r}{3}$  rows, since  $\sqrt{6} < 3$  (Fig. 15).

Once this first stage is over, we show that the remaining nodes would be able to commit. As Fig. 16 depicts, after the first stage completes, there are 8 more nodes which will definitely be able to commit since their committed neighbors  $\geq (r + 1 + \lceil \frac{r}{2} \rceil)r + 2\lceil \frac{r}{2} \rceil \lfloor \frac{r}{3} \rfloor \geq \frac{11r^2}{6} \geq \frac{4r^2}{3}$  (for all  $r \geq 2$ ). Thereafter all the other remaining nodes will be in a position to commit since the minimum number of committed neighbors that any of these nodes has is  $\geq (r + 1)r + 2\lceil \frac{r}{2} \rceil \lfloor \frac{r}{3} \rfloor + 4 > \frac{4r^2}{3}$  (see shaded region in Fig. 16). Thus the inductive hypothesis stands proven.

The inductive hypothesis along with the base case suffice to show that every honest node will eventually commit to the correct message.




Fig. 13. First Stage

Fig. 14. Progress of First Stage

Fig. 15. Completion of First Stage




Fig. 16. Second Stage: Step 1

Fig. 17. Progress of Second Stage

Fig. 18. Completion of Second Stage