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 grid (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 [6]. 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.

Categories and Subject Descriptors

C.2.4 [**Computer-Communication Networks**]: Distributed Systems; C.4 [**Performance of Systems**]: Fault Tolerance

General Terms

Algorithms, Reliability

Keywords

Byzantine failure, Crash-stop failure, Broadcast, Fault Tolerance, Radio Network

1. INTRODUCTION

Reliable broadcast in the presence of Byzantine and crash-stop 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 r (these are 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. Two distance metrics, L_∞ and L_2 (further discussed in Section 2), are considered. The considered fault model (first introduced in [6]) allows an adversary to place faults so long as the number of faults in any single *neighborhood* (to be formally defined later) do not exceed some value t. The results also hold for a finite toroidal network, as boundary anomalies are eliminated. We present bounds on the maximum number of failures t that may occur in any given *neighborhood* without rendering reliable broadcast impossible. For the case of Byzantine failures, we improve on bounds presented in [6]. We present a protocol (utilizing a notion of indirect reports) that allows reliable broadcast to be achieved in the L_∞ metric whenever $t < \frac{1}{2}r(2r+1)$. This exactly matches the impossibility bound proved in [6], and thus establishes an exact threshold for Byzantine agreement under this network 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_∞ metric). We present arguments suggesting that in L_2, i.e., Euclidean metric, when r is sufficiently large, similar thresholds must hold. We argue that for sufficiently large r, Byzantine agreement is possible in Euclidean metric if slightly less than one-fourth (more precisely, a 0.23 fraction) of the nodes in any given neighborhood may be faulty, while it is possible to tolerate crash-stop failures that are somewhat less than half (more precisely, a 0.46 fraction) of the neighborhood population. Finally, we consider the issue of tolerable faults with a simple protocol that does not use indirect reports (i.e. the protocol of [6]). We present an asymptotically tighter bound (than in [6]) for achievability with Byzantine failures by proving that reliable broadcast in L_∞ metric is achievable for $t \leq \frac{2}{3}r^2$ using this simple protocol.

2. NETWORK MODEL

We consider the network model described in [6]. Nodes are located on an infinite grid (each grid unit is a 1×1 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

*This research is supported in part by Motorola, Inc., and a Verizon Fellowship.

PODC'05, July 17–20, 2005, Las Vegas, Nevada, USA.

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. Note that this idealized shared radio channel intrinsically preserves ordering of messages sent by a node, i.e., if a node transmits messages m_1 and m_2 respectively in order, they will be received in that same order by all neighbors. We call this idealized behavior the *reliable local broadcast* assumption. While this assumption does not hold *per se* in real wireless networks, it may be possible to implement a local broadcast primitive that can provide probabilistic guarantees (given that transmissions are successfully received with a certain probability).

In this paper, we consider two distance metrics: L_∞ and L_2. The L_∞ metric is essentially the metric induced by the L_∞ norm [8]. 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 the this metric. Thus $nbd(a, b)$ comprises a square of side $2r$ with its centroid at (a, b). The L_2 metric is induced by the L_2 norm [8], 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 $nbd(a, b)$ comprises nodes within a circle of radius r centered at (a, b). The L_∞ metric enables more tractable analysis, and allows us to establish exact fault tolerance thresholds. It also provides valuable intuition, on which we base an approximate argument for the L_2 metric (which is the metric of practical significance).

The adversary is allowed to place faults as long as no single neighborhood contains more than t faults. When we refer to the neighborhood of a node v, it includes v itself. Thus a correct node may have upto t faulty neighbors, while a faulty node may have upto $(t - 1)$ neighbors that are also faulty.

As in [6], 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 (e.g., the schedule mentioned in [6]) so long as time-optimality is not a concern. We shall further discuss the impact of relaxing these assumptions in Section 10. However, note that accidental collisions (not deliberately caused by the adversary) may be handled to some extent by a probabilistic primitive (as they can be treated akin to transmission errors). A designated source (that is assumed located at the origin of the grid coordinate system, without loss of generality) 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 any single neighborhood without rendering reliable broadcast impossible.

3. RELATED WORK

Reliable broadcast has been extensively studied for networks with point-to-point communication under various connectivity conditions [1]. The classic result of Pease, Shostak and Lamport [11], [9] 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 [5]. Byzantine agreement in k-cast channels has been considered in [4]. However this does not capture the spatially dependent connectivity that characterizes radio networks. Reliable broadcast in radio networks has also been studied in [7] and [6]. Crash-stop failures are considered in [7] for finite networks comprising nodes located in a regular grid pattern. The focus is on obtaining algorithms for efficient broadcast to the part of the network that is reachable from the source, and not on quantifying the number of faults that render some nodes unreachable. In [6], the infinite grid network and neighborhood fault model described in Section 2 were considered. It was shown that under a Byzantine failure model, reliable broadcast is not achievable for $t \geq \lceil \frac{1}{2} r(2r + 1) \rceil$ (in both L_∞ and L_2 metrics). Besides a protocol was described that was able to achieve reliable broadcast under the following conditions:

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

This 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.

In very recent work [12], further study of the locally bounded fault model has been undertaken. The focus is on arbitrary graphs instead of using a specific network model. While the discussion mentions both radio and message-passing networks, 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 [6] that is referred 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 6. 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. The work described in this paper differs substantially, 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.

4. NOTATION AND 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 $nbd(x, y)$. For succint description, 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. 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. The source of the broadcast is deemed to be situated at $(0, 0)$, without affecting generality of the results.

5. 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 10), this significantly simplifies the problem, and relaxes the requirements for agreement. Under our assumptions (also in [6]), 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, beause if it were to attempt sending contradicting messages, they would be heard by all its neighbors, and its duplicity would be 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.

In a fully connected network, it is thus 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 [5] 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 [5] 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 as follows: 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.

6. FAULT THRESHOLD FOR BYZANTINE FAILURES IN L_∞ METRIC

As discussed in Section 3, it was proved in [6] that reliable broadcast is impossible with Byzantine failures, in L_∞ as well as L_2 metrics, if $t \geq \lceil \frac{1}{2} r(2r + 1) \rceil$. We prove the following:

THEOREM 1. *Under a Byzantine failure model, if* $t < \frac{1}{2} r(2r+1)$*, reliable broadcast is achievable in the* L_∞ *metric.*

This is an exact match to the impossibility bound for L_∞, and establishes the *exact* threshold for achieving reliable broadcast in the square grid network under consideration. Since an L_∞ neighborhood comprises $(2r + 1)^2$ nodes, this threshold implies that slightly less than one-fourth of the nodes in a neighborhood may be faulty. We present a protocol that achieves this objective [1]. 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 direct/indirect reports for nodes within its four-hop neighborhood. This state may be reduced further by stipulating exact messages that a node should look out for, and this shall become clear from our constructive proof for the viability of reliable broadcast with $t < \frac{1}{2} r(2r + 1)$. The protocol operates as follows:

- Initially, the source 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.
- 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 locally broadcast a $HEARD(j, i, v)$ message.

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

  On receipt of a $HEARD(l, k, i, v)$ message, record the message, and locally 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.

  On committing to a value v, do a one-time local broadcast of $COMMITTED(j, v)$.

  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. j is said to have reliably determined the value committed to by node i if one of the following conditions holds:

  - i is its neighbor, and so j heard $COMMITTED(i, v)$ directly. In this case, there is no cause for doubt as to the value committed to by node i, since no other node is capable of spoofing i's address, and collisions are ruled out.

[1] We have since obtained results that allow the same fault threshold to be tolerated using a simpler protocol, with a corresponding simpler proof. Section 6.2 provides a brief discussion.

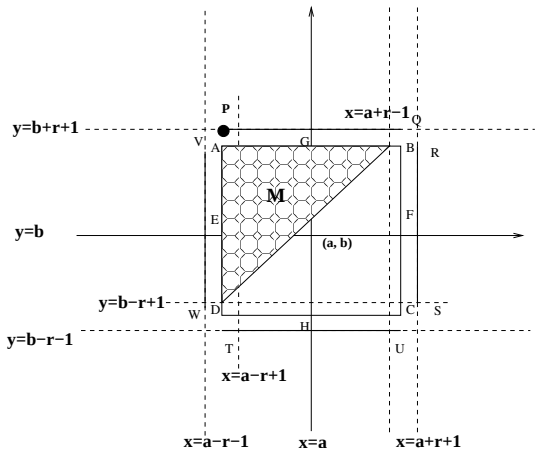

Figure 1: Nodes in $nbd(a,b)$ whose committed values P can reliably determine

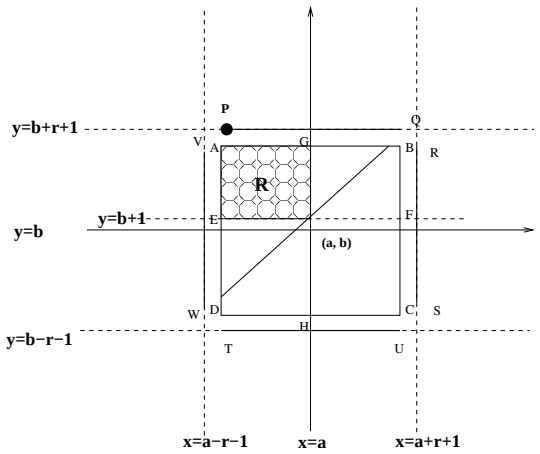

Figure 2: Nodes in $nbd(a,b)$ that are immediate neighbors of P

– 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 through the $HEARD$ messages that propagate via upto three intermediate nodes (i.e. upto four hops from the node that sent the $COMMITTED$ message), and the path information is obtained from these messages (as each forwarding node affixes its identifier to the message).

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

PROOF. The proof is by contradiction. Consider the first node, say j, that makes a wrong decision to commit to a 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 $pnbd(a,b) - nbd(a,b)$ 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 prove 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.

We show this for a corner node in $pnbd(a,b) - nbd(a,b)$ i.e. the node marked P (which is located at $(a-r, b+r+1)$) in Fig. 1. This represents the worst case. For all other

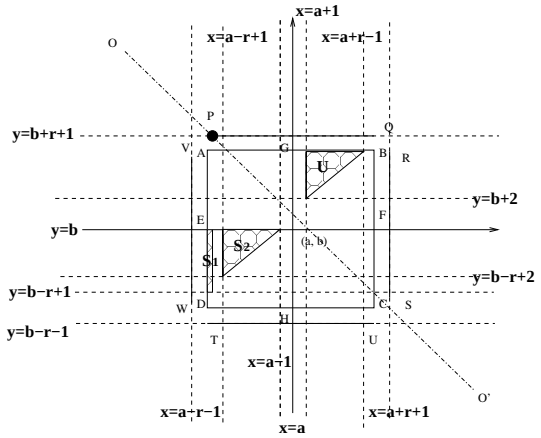

Figure 3: Nodes in $nbd(a, b)$ to which P has sufficient connectivity

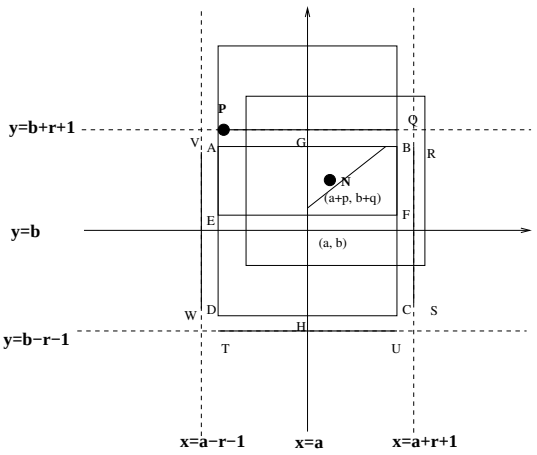

Figure 4: A node N in Region U

Region	x-extent	y-extent
A	$(a+p-r) \leq x \leq a$	$(b+1) \leq y \leq (b+q+r)$
B_1	$(a+1) \leq x \leq (a+p-1)$	$(b+1) \leq y \leq (b+q+r)$
B_2	$(a+1-r) \leq x \leq (a+p-1-r)$	$(b+1) \leq y \leq (b+q+r)$
C_1	$(a+p+1) \leq x \leq (a+r)$	$(b+q+1) \leq y \leq (b+r+1)$
C_2	$(a+p+1-r) \leq x \leq a$	$(b+q+1+r) \leq y \leq (b+1+2r)$
D_1	$(a+p) \leq x \leq (a+p+r-q)$	$(b+r+q-p+1) \leq y \leq (b+r+q)$
D_2	$(a+1) \leq x \leq (a+p)$	$(b+1+r+q) \leq y \leq (b+1+2r)$
D_3	$(a+1-r) \leq x \leq (a+p-r)$	$(b+1+r+q) \leq y \leq (b+1+2r)$
J	$(a-2r) \leq x \leq a$	$(b+1) \leq y \leq (b-p+r)$
K_1	$(a-2r) \leq x \leq a$	$(b-p+1) \leq y \leq b$
K_2	$(a-2r) \leq x \leq a$	$(b-p+r+1) \leq y \leq (b+r)$

Table 1: Spatial Extents of Various Regions

nodes in $pnbd(a, b) - nbd(a, b)$, the condition can be seen to be achieved via a similar argument. We briefly discuss this in Section 6.1.

We show that node P is able to reliably determine the values committed to by the nodes in the shaded region M in Fig. 1. Region M comprises $\{(a-r+p, b-r+q)|2r \geq q > p \geq 0\}$ and hence has $r(2r+1)$ nodes. The first observation is that P can directly hear the nodes in the shaded sub-region R in Fig. 2, comprising $\{(x, y)|(a-r) \leq x \leq a; (b+1) \leq y \leq (b+r)\}$ (this constitutes $r(r+1)$ nodes), and so is certain of the value they committed to. The remaining sub-regions are depicted in Fig. 3 as U (comprising $\frac{1}{2}r(r-1)$ nodes), S_1 (comprising r nodes), and S_2 (comprising $\frac{1}{2}r(r-1)$ nodes). We now explicitly prove existence of suitable node-disjoint paths for nodes that lie in the upper triangular region U in Fig. 3. Any node N in this region may be considered located at $(a+p, b+q)$ (Fig. 4), such that $r \geq q > p \geq 1$ 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). For greater clarity, the spatial extents of various demarcated regions used in the following argument are tabulated in Table 1.

Consider 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 $nbd(N)$ (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 $nbd(P)$. 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 correpondence between a point (x, y) in B_1 and a point $(x-r, y)$ 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 $\{(x, y)|(a+p) \leq x \leq (a+p+r-q), (b+$

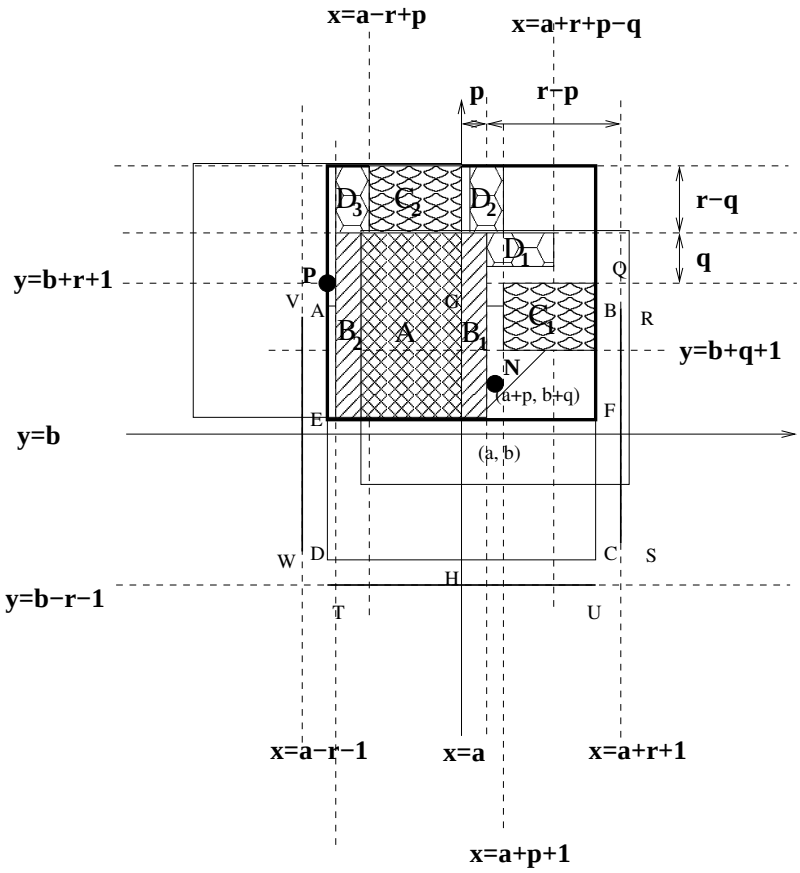

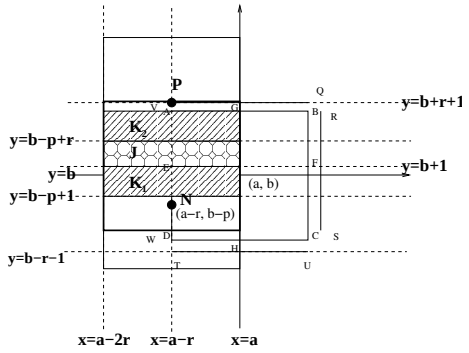

Figure 6: Connectivity between P and nodes in S_1

Figure 5: Construction depicting node-disjoint paths between N and P

$r+q-p+1) \leq y \leq (b+r+q)\}$, 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 $\leq r$). 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 correpondence with a node $(x-r,y)$ in D_3. Hence, 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.

We now consider nodes in regions S_1 and S_2 depicted in Fig. 3. Then: $S_1 = \{(a-r, b-p)|0 \leq p \leq (r-1)\}$. It can be shown that P has $r(2r+1)$ disjoint paths to each node in S_1, as depicted in Fig. 6. Any node N in S_1 is located at $(a-r, b-p)$. Consider region J comprising $\{(x,y)|(a-2r) \leq x \leq a; (b+1) \leq y \leq (b-p+r)\}$. All nodes in J are common neighbors of N and P, and provide $(r-p)(2r+1)$ paths of the form $N \rightarrow J \rightarrow P$. Region K_1 comprises $\{(x,y)|(a-2r) \leq x \leq a; (b-p+1) \leq y \leq b\}$, and falls enirely within $nbd(N)$. Region K_2 is $\{(x,y)|(a-2r) \leq x \leq a; (b-p+r+1) \leq y \leq (b+r)\}$, and falls in $nbd(P)$. For each node (x,y) falling in K_1, there is a one-to-one correspondence with a node $(x,y+r)$ in K_2, and thus we obtain $p(2r+1)$ paths of the form $N \rightarrow K_1 \rightarrow K_2 \rightarrow P$. This yields a total of $r(2r+1)$ paths (all lying entirely within $nbd(a-r, b+1)$), as depicted in Fig. 6.

Region S_2 comprises $\{(a-q, b-p)|(r-1) \geq q > p \geq 0\}$. For the nodes in S_2, observe that each node $(a-q+1, b-p+1)$ in S_2 possesses the same relative position w.r.t. P as the node $(a+p, b+q)$ in region U of Fig. 3 (note the axial symmetry about axis OO'), and due to the symmetric structure of the network, shall enjoy exactly the same connectivity properties to P as the node $(a+p, b+q)$ in region U. Since we have already shown existence of sufficient connectivity for those nodes, the same holds for nodes in S_2.

The inductive hypothesis, along with the base case, suffices to show that every honest node 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.

Note that the connectivity condition proved above suffices to prove that upto $2t < r(2r+1)$ crash-stop failures are tolerable in L_∞ metric. We elaborate further in Section 7. □

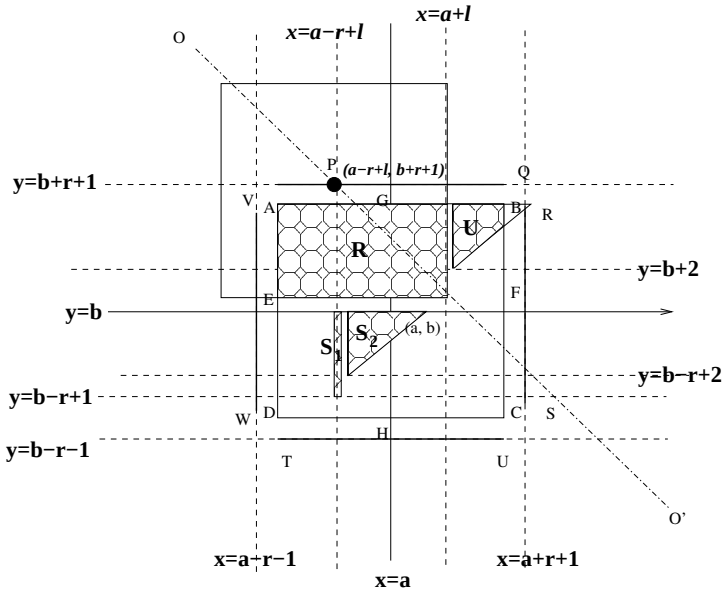

Figure 7: Non-worst Case Location of P

6.1 Non-worst Case Location of P

We briefly discuss how the connectivity argument holds for all $P \in pnbd(a,b) - nbd(a,b)$. We consider non-worst case locations of $P \in \{(a-r+l, b+r+1) | 1 \leq l \leq r\}$. For all other locations, the argument holds by symmetry. The situation is depicted in Fig. 7. One may consider P to be translated to the right by l units from its worst case location at $(a-r, b+r+1)$. Then, region R that lies in direct range of P (recall from Fig. 2) now comprises $r(r+l+1)$ nodes. If we also translate regions U, S_1, and S_2 by l units each to the right, they preserve their relative positions and hence connectivity to P. However, now $\frac{1}{2}l(l-1)$ nodes from U fall out of $nbd(a,b)$, but this is more than compensated by the increase of rl nodes in region R. Thus, if we count the number of nodes in $nbd(a,b) \cap U$, $nbd(a,b) \cap S_1$, and $nbd(a,b) \cap S_2$, it can be shown that they are $\geq r(r-l)$ in number. Together with the $r(r+l+1)$ nodes in region R, they provide at least $r(2r+1)$ nodes to which P is connected either directly or via $2t+1$ node-disjoint paths all lying within some single neighborhood.

6.2 A Simpler Protocol

We have formulated a new protocol [3] in which only the immediate neighbors of a node that sent a $COMMITTED$ message, send out a $HEARD$ message reporting it. Thus, information about the value committed to by a node propagates only upto its two hop neighborhood. This suffices to achieve reliable broadcast. The correctness of this protocol proceeds from the observation that a much simpler connectivity condition is sufficient to ensure reliable broadcast. Given that all honest nodes in $nbd(a,b)$ have been able to correctly determine the broadcast value, any node P in $pnbd(a,b) - nbd(a,b)$ should be connected to $2t+1$ nodes N in $nbd(a,b)$ via a single path each, such that collectively these $2t+1$ paths are node-disjoint and they all (the endpoints N, as well as any intermediate nodes) lie in some single neighborhood. Details are presented in [3].

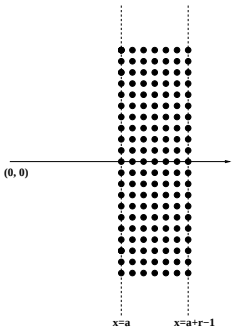

Figure 8: Network Partition due to Crash Stop Failures

7. CRASH-STOP FAILURES

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_∞ metric.

THEOREM 4. *Under a crash-stop failure model, if $t \geq r(2r+1)$, it is impossible to achieve reliable broadcast in L_∞ metric.*

PROOF. We present a construction with $t = r(2r+1)$ that renders reliable broadcast impossible. Consider the network in Fig. 8. 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. *Under a crash-stop failure model, if $t < r(2r+1)$, it is possible to achieve reliable broadcast in L_∞ metric.*

PROOF. A proof proceeds from the proof of Theorem 1. 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. We also describe simpler proof(s) in [2] and [3]. Note that in L_∞, this threshold corresponds to slightly less than half the nodes in a neighborhood. □

8. BROADCAST IN EUCLIDEAN METRIC

We 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. However, we present arguments to suggest that reliable broadcast in L_2 is achievable if the fraction of nodes in any neighborhood that exhibit Byzantine faults is slightly less than one-fourth. We work with the value $t < 0.23\pi r^2$.

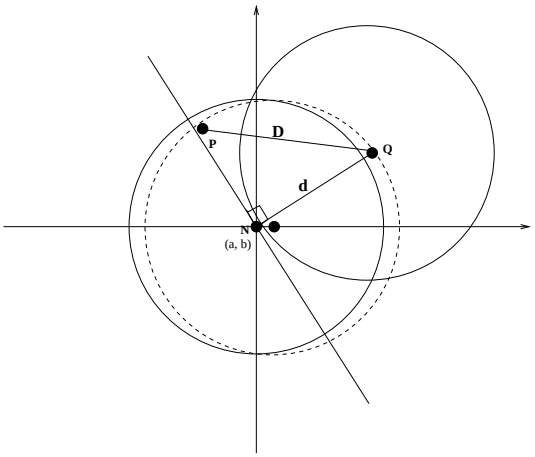

Figure 9: Illustrating an Approximate Argument for Euclidean Metric

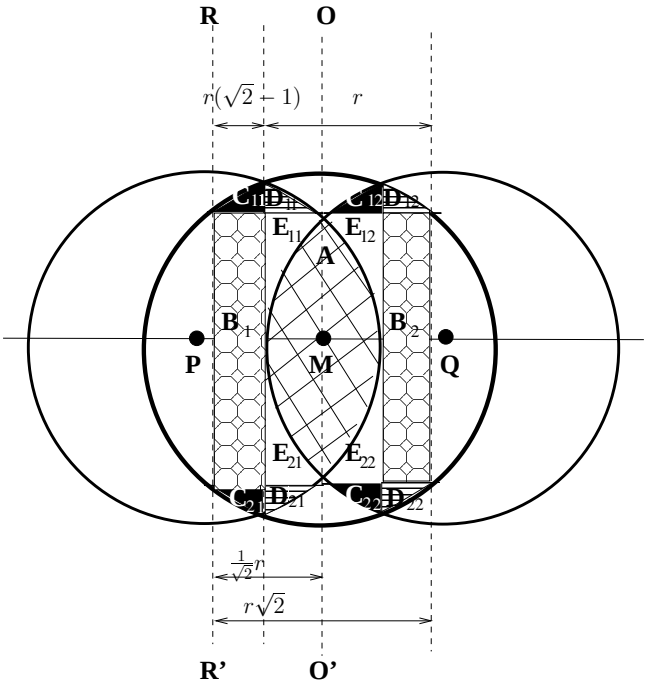

Figure 10: Approximate Construction depicting Node-Disjoint Paths (PQ from Fig. 9 rotated to horizontal axis)

The basis for the argument is that given a closed simple region of area A, and perimeter p, bounded by upto k straight line segments and circular arcs of radius r, where k is a small constant, the number of lattice point lying within it, N_l, is given by $N_l = A \pm O(p)$, and the constant hidden in the $O(p)$ term is small. The justification for this claim is based on Pick's Theorem [13], and is presented in the Appendix. For our argument, we consider sub-regions of a circle bounded by circular arcs of radius r, and straight line segments (each of length $\leq 2r$), with $k \leq 4$. These regions have a perimeter $p = O(r)$, and area $A = O(r^2)$. Hence, the number of nodes that lie in a subregion having area A is approximately $A \pm O(r)$. Thus, for sufficiently large r, the quadratic area term dominates, and the area is a good approximation for the number of lattice points lying within these regions. The argument proceeds by induction, as in Section 6.

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 $pnbd(a,b)$ are able to commit to the correct value.

Justification of Inductive Hypothesis:

We show that each node in $pnbd(a,b) - nbd(a,b)$ 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. As in Section 6, we proceed by showing that each node in $pnd(a,b) - nbd(a,b)$ 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. 9. Let d be the distance between the node at (a,b) (we call it N) and any node in $(pnbd(a,b) - nbd(a,b))$ (we call it node Q). Then $d \leq r+1$ (from the triangle inequality). Consider the half-neighborhood of (a,b) demarcated by the axis perpendicular to NQ (not counting the points lying on the axis). The number of lattice points lying in it may be approximated by $0.5\pi r^2 - O(r)$. Then, as the number of faults $t < 0.23\pi r^2$, it implies that there are at least $2t+1$ nodes 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 worst case situation. Fig. 10 depicts PQ rotated to the horizontal axis for purpose of clarity. The corresponding rotated coordinates shall be referred to as x', y'. We denote location of P in this rotated system by (P_x, P_y). Use of these coordinates enables simpler description of the spatial extents of various regions relative to the locations of P and Q. Note that the integer coordinates in this system are not coincident with the lattice points in the actual grid. This does not affect the argument, as we use these coordinates solely to define placement of the regions, and then rely on the area-based argument to count the number of actual grid points falling in each.

We attempt to construct node-disjoint paths between P and Q that all lie within the neighborhood centered at the midpoint of PQ (we call it M). If M is not an actual lattice

point, the neighborhood center is perturbed to the nearest lattice point. This can only affect the argument by $O(perimeter) = O(r)$. Similarly, when we state one-to-one correspondences between points in the following argument, any deviation from the calculated numbers (due to the discrete nature of the grid) can only be by $O(r)$.

The set of nodes marked A in Fig. 10 are common neighbors of P and Q (i.e., fall in $nbd(P) \cap nbd(Q)$), and lie within $nbd(M)$. They constitute two-hop PQ paths ($P \rightarrow A \rightarrow Q$). A set of three-hop paths $P \rightarrow B_1 \rightarrow B_2 \rightarrow Q$ is also formed. Region B_1 comprises $\{(x', y') | (x', y') \in (nbd(P) - nbd(Q)) \cap nbd(M), P_x < x' < P_x + r\sqrt{2} - r, P_y - \frac{r}{\sqrt{2}} \leq y' \leq P_y + \frac{r}{\sqrt{2}}\}$. B_2 is the translate of B_1 by r units along the direction of PQ, and $B_2 \in (nbd(Q) - nbd(P)) \cap nbd(M)$ (since PQ $\approx r\sqrt{2}$). This yields a one-to-one correspondence between points in B_1 and points in B_2.

Region $C_{11} \cup C_{21} = \{(x', y') | (x', y') \in (nbd(P) - nbd(Q)) \cap nbd(M), P_x < x' < P_x + \frac{r}{2\sqrt{2}}, (y' < P_y - \frac{r}{\sqrt{2}}) \vee (y' > P_y + \frac{r}{\sqrt{2}})\}$. Region $D_{11} \cup D_{21} = \{(x', y') | (x', y') \in (nbd(P) - nbd(Q)) \cap nbd(M), P_x + \frac{r}{2\sqrt{2}} \leq x', (y' < P_y - \frac{r}{\sqrt{2}}) \vee (y' > P_y + \frac{r}{\sqrt{2}})\}$. Region $E_{11} \cup E_{21} = \{(x', y') | (x', y') \in (nbd(P) - nbd(Q)) \cap nbd(M), P_x + r\sqrt{2} - r \leq x', P_y - \frac{r}{\sqrt{2}} \leq y' \leq P_y + \frac{r}{\sqrt{2}}\}$. The translated image(s) of $C_{11}, C_{21}, D_{11}, D_{21}$ by $\frac{1}{\sqrt{2}}r$ units in the x' direction, fall within $(nbd(Q) - nbd(P)) \cap nbd(M)$. Thus, there is a set of three-hop paths $P \rightarrow C_{11}(C_{21}) \rightarrow C_{12}(C_{22}) \rightarrow Q$ and $P \rightarrow D_{11}(D_{21}) \rightarrow D_{12}(D_{22}) \rightarrow Q$, since each point (x', y') in C_{i1} (D_{i1}) has a corresponding point (its image under translation along the direction of PQ) in C_{i2} (D_{i2}). Finally, there exist paths $P \rightarrow E_{11}(E_{21}) \rightarrow E_{12}(E_{22}) \rightarrow Q$ such that each point in E_{11} (E_{21}) has a one-to-one correspondence with its mirror image with respect to axis OO', lying in $E_{12}(E_{22})$. The number of such paths is approximately equal to the sum of the areas A, B_1, C_{11}, C_{21} D_{11}, D_{21}, E_{11} and E_{21} which turns out to be approximately $1.47r^2 = 0.47\pi r^2 > (2(0.23\pi r^2) + 1)$ for sufficiently large r. Thus approximately $0.23\pi r^2$ Byzantine faults may be tolerated. Note that all considered paths comprise upto three hops. Thus the protocol of Section 6 suffices for L_2.

We also argue 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 [6] for L_∞, which is depicted in Fig. 11. As proved in [6], 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. 11). The relevant area is approximately $0.6\pi r^2$, and we expect approximately $0.6\pi r^2 \pm O(r)$ nodes to lie in it. Half of these, i.e., around $0.3\pi r^2 \pm O(r)$ are to be faulty. This yields 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 would lie between a 0.23 and a 0.3 fraction i.e. close to a one-fourth fraction of faults.

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 seems to be in the range of half the neighborhood population.

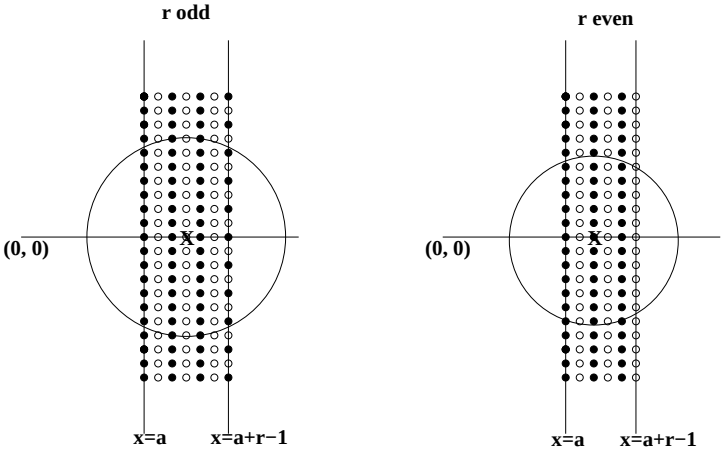

Figure 11: Impossibility Construction for Byzantine Failures in Euclidean metric

9. A SIMPLE BYZANTINE PROTOCOL: IMPROVED BOUNDS

We have obtained improved bounds for tolerable faults when an extremely simple Byzantine protocol (described in [6]) is used. In this protocol, initially the source transmits the value, and its immediate neighbors are able to commit to that value instantly. They 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 no further progress is possible. We present an asymptotically tighter bound for the number of tolerable Byzantine faults in the L_∞ metric (using this protocol) than that presented in [6].

THEOREM 6. *Under a Byzantine failure model, if $t \leq \frac{2}{3}r^2$, it is possible to achieve reliable broadcast in the L_∞ metric, with the described simple protocol.*

PROOF. The proof is omitted due to paucity of space. It is described in [2]. □

10. A STRONGER ADVERSARY

The presence of a broadcast channel introduces numerous possibilities for stronger adversarial behavior. A malicious node can potentially spoof another node's address and send spurious messages under guise. There is also 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 non-faulty node in its vicinity. When the number of collisions is bounded, it may be possible to come up with protocols that achieve reliable broadcast. As discussed

briefly in [6], the situation may not simply be remediable via retransmissions, if the adversary leverages collisions to send contradicting messages to different parts of the network. A protocol that involves consultation between the neighbors of a node as to the value(s) they heard it transmit, as well as any detected collisions, can potentially resolve the problem, and requires further investigation.

11. CONCLUSIONS

We have presented results regarding the number of Byzantine and crash-stop failures that may be tolerated in an idealized radio network without rendering reliable broadcast impossible. We have considered an adversarial model where the adversary is free to choose faulty nodes, so long as the placement satisfies the constraint that no neighborhood has more than t faults. However, in the presence of channel errors etc., the *reliable local broadcast* assumption, that underlies these results, is not trivial to realize. Thus, implementation of a reliable broadcast service based on the radio network model would require efficient implementation of a reliable local broadcast primitive that operates under realistic network conditions. Proposed mechanisms for reliable broadcast in multi-hop mobile networks (e.g., that described in [10]) have typically not focused on Byzantine node failures. Besides, these mechanisms do not leverage the broadcast nature of the shared wireless channel, relying instead on construction of clustering or backbone structures for reliable dissemination via unicast messages. 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.

12. REFERENCES

[1] H. Attiya and J. Welch. *Distributed Computing*. McGraw-Hill, 1998.

[2] V. Bhandari and N. H. Vaidya. On reliable broadcast in a radio network. Technical Report, CSL, UIUC, May 2005.

[3] V. Bhandari and N. H. Vaidya. On reliable broadcast in a radio network: A simplified characterization. Technical Report, CSL, UIUC, 2005.

[4] J. Considine, L. A. Levin, and D. Metcalf. Byzantine agreement with faulty majority using bounded broadcast. *CoRR*, cs.DC/0012024, 2000.

[5] D. Dolev. The byzantine generals strike again. *J. Algorithms*, 3(1):14–30, 1982.

[6] 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*, pages 275–282. ACM Press, 2004.

[7] E. Kranakis, D. Krizanc, and A. Pelc. Fault-tolerant broadcasting in radio networks. *J. Algorithms*, 39(1):47–67, 2001.

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

[9] L. Lamport, R. Shostak, and M. Pease. The byzantine generals problem. *ACM Trans. Program. Lang. Syst.*, 4(3):382–401, 1982.

[10] E. Pagani and G. P. Rossi. Providing reliable and fault tolerant broadcast delivery in mobile ad-hoc networks. *MONET*, 4(3):175–192, 1999.

[11] M. Pease, R. Shostak, and L. Lamport. Reaching agreement in the presence of faults. *J. ACM*, 27(2):228–234, 1980.

[12] A. Pelc and D. Peleg. Broadcasting with locally bounded byzantine faults. *Information Processing Letters*, 93(3):109–115, Feb 2005.

[13] D. E. Varberg. Pick's theorem revisited. *The American Mathematical Monthly*, 92(8):584–587, October 1985.

APPENDIX

We claimed in Section 8 that given a simple closed region S of area A bounded by upto k line segments or circular arcs, where k is a small constant, the number of lattice points in S is $A \pm O(p)$. We justify this by bounding S, within and without, by lattice polygons, and applying Pick's Theorem [13]. For any such region S, consider the lattice polygon comprising grid squares that lie completely within S (Fig. 12). In certain cases, instead of a single lattice polygon, we obtain a number of simple polygons that may share a common vertex, or are disconnected (if S has narrow constrictions or *necks* (Fig. 13)). In rare instances, no such polygon may be obtained, if S is extremely narrow, and has no grid square lying completely within it ($A = O(p)$ for such regions). We call the polygon(s) thus obtained P_{in} (in case of multiple polygons, P_{in} refers to their union). Note that $S - P_{in}$ comprises the grid squares that are partially in S, i.e., those traversed by the boundary of S. Since the boundary of S comprises upto k line segments and arcs of radius r, the number of grid squares traversed by the boundary $\leq 2p + ck$, where c is a constant. The area of P_{in} must thus be at least $A-(2p+ck)$. Let n_1 denote the number of lattice points falling in P_{in}. Similarly, consider the lattice polygon P_{out} obtained by taking the union of all grid squares that lie fully or partially in S. P_{out} is simple, fully contains S, and its area can be no more than $A + (2p + ck)$ (it can at most have an additional area comprising the grid squares traversed by the boundary of S). Let the number of lattice points falling in P_{out} be n_2. Then $n_1 \leq N_l \leq n_2$. By invoking Pick's Theorem [2], it can be shown that $n_1 \geq A - O(p)$, and $n_2 \leq A + O(p)$. Thus $N_l = A \pm O(p)$.

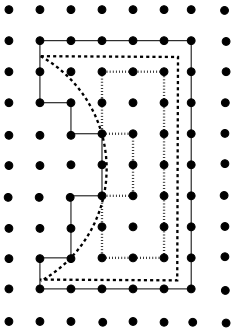

Figure 12: Bounding a Simple Closed Region via Lattice Polygons

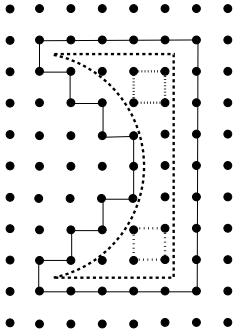

Figure 13: Region with *Neck*: Multiple Simple Polygons in Interior

[2]Pick's Theorem: Let A be the area of a simple closed lattice polygon. Let B denote the number of lattice points on the polygon boundary, and I the number of points in the polygon interior. Then: $A = I + \frac{1}{2}B - 1$.