A Lowe Bound for Broadcasting in Mobile Ad Hoc Networks*

Ravi Prakash* André Schiper† Mansoor Mohsin* David Cavin† Yoav Sasson†
* {ravip,mmohsin}@utdallas.edu
† {andre.schiper,david.cavin,yoav.sasson}@epfl.ch

*Department of Computer Science, The University of Texas at Dallas, USA
†École Polytechnique Fédérale de Lausanne (EPFL), Switzerland

EPFL Technical Report IC/2004/37

Abstract

We give a lower time complexity bound of $\Omega(n)$ for broadcasting in a mobile ad hoc network, where n is the number of nodes in the network. This is a tighter lower bound over the $\Omega(D \log\ n)$ bound by Bruschi and Pinto [1997], for a network of diameter D. We show that when node mobility is considered, the dominating factor in the complexity of an algorithm is the number of nodes in the network and not its diameter. We consider a synchronous model, in which time is divided into rounds. In each round a node receives messages, performs some computation, sends messages, and possibly moves to a new location.

1 Introduction

A Mobile Ad Hoc Network (MANET) is a set of *nodes* communicating with each other via multihop wireless links. Each node can directly communicate with only those nodes that are in its communication range. Intermediate nodes forward messages to the nodes that are more than one hop distance from the source. Since the nodes are mobile, the topology of the network is constantly changing.

In this paper, we study *broadcasting*, the problem of sending a message from a *source* node to all the other nodes of the network. The following simplified model has been widely considered to analyze the complexity of the broadcasting problem: nodes send messages in synchronous *rounds* or *time-slots*. In every round a node acts either as a *transmitter* or a *receiver*. A node acting as a receiver successfully receives a message in a round if exactly one of its neighbors transmits in that round. If two or more neighbors transmit in a round then there is a *collision* and none of the transmitted messages is received successfully.

All the broadcasting protocols that have been published did not consider node mobility when proving the upper and lower bounds.[1] In this paper, we investigate broadcasting for one-dimensional and two-dimensional networks with explicit node mobility. In particular, we give a lower bound of

*The work presented in this paper was supported by the National Competence Center in Research on Mobile Information and Communication Systems (NCCR-MICS), a center supported by the Swiss National Science Foundation under grant 5005-67322.

[1] In Basagni et al. [1999] the authors claim that their scheduling algorithm is mobility independent. However, this means that the scheduling algorithm is independent of nodes' neighbors.

$\Omega(n)$ rounds for deterministic distributed broadcasting algorithms for MANETs. This shows that when node mobility is considered, the dominating factor in the complexity of an algorithm is the number of nodes in the network and not its diameter. This contradicts many published results that we now discuss.

Since Chlamtac and Kutten [1985] have proved the broadcast problem to be $\mathcal{NP}$-hard, the lower bounds are useful in comparing the performance of proposed solutions. Many researchers have proposed different lower bounds for the broadcast problem in MANETs under varying system models.

When each node has full knowledge of the network topology, Alon et al. [1991] prove that any broadcast protocol require $\Omega(D + \log^2 n)$ rounds of transmission. For this centralized model, Chlamtac and Weinstein [1991] give a broadcast protocol that works in $O(D \log^2 n)$, which is optimal for a constant D. An $O(D + \log^5 n)$ broadcast protocol is given by Gaber and Mansour [2003], which is optimal when $D = \Omega(\log^5 n)$.

Bar-Yehuda et al. [1992] were the first to give a randomized distributed broadcast protocol. Their protocol works in $O((D+\log\ n/\epsilon)\cdot\log\ n)$ time-slots, with probability $1-\epsilon$. Hence, depending on the value of ϵ, the protocol can obtain near optimal results but it does not guarantee a reliable broadcast within this bound. A lower bound of $\Omega(D \log(n/D))$ is given by Kushilevitz and Mansour [1998] for any randomized broadcast protocol.

The first lower bound on deterministic distributed broadcast protocol was given by Bruschi and Pinto [1997]. They prove a lower bound of $\Omega(D \log\ n)$, for $D \leq n/2$. A broadcast protocol for this model was given by Basagni et al. [1999]. The total number of rounds required by this broadcast is $O(D\Delta\ \log^{\lfloor \log\ \Delta \rfloor} n)$, where Δ is the maximum degree of the network. When $\Delta = 3$, the protocol completes in $O(2D\ \log\ n)$ rounds, which equals the lower bound proved by Bruschi and Pinto [1997]. Furthermore, when $\Delta = 7$, the protocol is optimal for constant diameter networks [Alon et al., 1991]. Other deterministic broadcasting algorithms are given by Chlebus et al. [2002], which consider networks with/without collision detection mechanisms.

2 System Model

We introduce a model that incorporates node mobility. We assume a grid-based system, with nodes located at the grid points. Any number of nodes could be co-located at a grid point at any instant of time. Messages can be sent reliably by a node to other nodes that are no more than R inter-grid point distance away, where $R \geq 1$ is the communication range of each node. Time is divided in synchronous rounds: messages are sent at the beginning of a round, and we assume that each node within the communication range of the sender receives the message by the end of the same round. Once a node has received all the messages meant for it in a round, it performs some computation which may result in a change of its state. Then the node can move to a neighboring grid point, before starting the next round.

Formally, we define our model by (1) a *message generation* function, (2) a *state transition* function, and (3) a *mobility* function:

- *Message generation function* in round r:

$$(s_x^r) \rightarrow msent_x^r$$

- *State transition function* in round r:

$$(s_x^r,\ msent_x^r,\ mreceived_x^r) \rightarrow {s'}_x^r$$

- *Mobility function* before starting round $r+1$:

$$(location_x^r) \rightarrow (location_x^{r+1})$$

where:

- s_x^r state of node x during round r.
- $location_x^r$ location of node x during round r.
- $msent_x^r$: messages sent by node x during round r.
  These messages are received by all nodes y such that $distance(location_y^r,\ location_x^r) \leq R$.
- $mreceived_x^r$: messages received by node x during round r.

It should be noted that this round model is different from the round model of the papers mentioned in Section 1, where in each round a node acts either as transmitter or as receiver, and a message is received by some node only if exactly one of its neighbors transmits in that round. Let us call the rounds of Section 1 *micro-rounds* and the rounds of this section *macro-rounds*. It is obvious that, if we want to emulate our macro-rounds using micro-rounds, *at least* one micro-round is needed to emulate one macro-round. So a lower bound expressed in terms of macro-rounds is also a lower bound in terms of micro-rounds. This allows us to compare our lower bound to the lower bounds mentioned in Section 1.

3 Bounds with One- and Two-Dimensional Network

In static networks (no mobility), the added knowledge of network diameter helps in determining the tight-bound on the time complexity for broadcasts. Our goal is to investigate if such knowledge is helpful in the context of MANETs (with mobility). Specifically, we are given an upper bound D on the diameter of the MANET. Given this information, is it possible to assert anything about completion time of broadcast? Of course, while broadcast is taking place, from round to round, the diameter of the MANET may change, but never exceed D. We also assume that the nodes stay connected all the time. In other words, we consider an adversary that initially places nodes on grid points and moves them in order to delay as much as possible the termination of broadcast, with the constraints that (1) nodes remain connected all the time and (2) the diameter of the MANET never exceeds D.

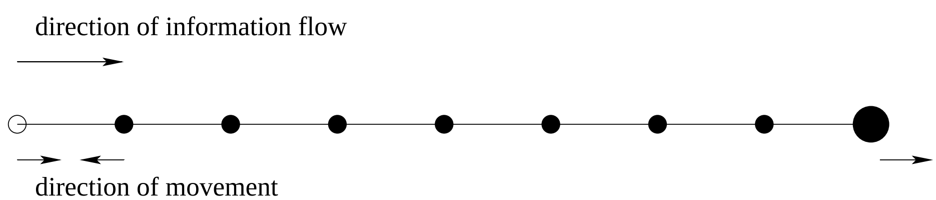

Figure 1: Mobility adversary in a one-dimensional network.

Lemma 1 *In a one-dimensional MANET network it is possible to complete broadcasting in $O(D)$ rounds provided $R \geq 3$.*

Proof: Consider the network topology shown in Figure 1. The nodes form a graph of diameter D. The node to the extreme left represented by an unfilled circle has some message m to broadcast. A total of $D-1$ individual nodes along the chain are represented as small filled circles. All the remaining nodes are co-located at a single grid-point and represented as a large filled circle to the extreme right. In round 0, the node to the extreme left has message m to broadcast. Obviously, the best that the broadcasting algorithm can do to reduce the broadcast completion time is for each node, once it has received m for the first time in round r, to send m to all its neighbors in all successive rounds (not only in round $r+1$, since its neighbors can change, due to mobility). With no mobility, the algorithm obviously terminates in D rounds.

With node mobility, the best that the adversary can do in order to delay the completion of the broadcast is the following:

1. Before the beginning of each round, the nodes at the extreme right (large filled circle) move one grid point to the right, if this *does not increase the diameter of the network*. Otherwise these nodes do not move.

2. A node that receives m for the first time in round r, moves one grid point to the left before the beginning of round $r+1$.

Note that, even though this is not necessary for our proof, this adversary is not stronger than the one introduced in the beginning of this section (i.e., the network can remain connected). Consider the following additional mobility controlled by the adversary. Before starting round $r+1$, (a) all the nodes that have received m before round r move one grid unit to the right, (b) the node that received m for the first timed in round r and all nodes that have not yet received m (except those at the extreme right) move one grid unit to the left, and (c) among the nodes at the extreme right (in the large filled circle) one moves to the left, while all others move one grid unit to the right provided this move to the right does not result in the diameter exceeding D. It is easy to show, by induction on the round number, that this mobility function keeps the network connected.

In order to evaluate the impact of this adversary, let x denote the distance (in grid points) between two neighbor nodes in Figure 1. The parameter x is such that $R/2 < x \leq R$, in order for the graph to be connected and so that in each round at most one new node receives m. The message m, initially at position 0, propagates to the right at the speed of $x-1$ (grid points per round). The rightmost nodes, initially at position xD move to the right at most at the speed of 1. So m reaches the rightmost node at latest at round r such that $r(x-1) = xD + r$, i.e., at round $r = xD/(x-2)$. Hence, for $x > 2$, the number of rounds is $O(D)$. Since $R/2 < x$, $R > 3$ ensures $x > 2$. □

Regarding the case where $R \leq 3$, it is easy to see that in this case the above mobility adversary leads to a broadcast in $O(n)$ rounds.

Lemma 2 *In a two-dimensional MANET the lower-bound on time to complete broadcasting is* $\Omega(n)$.

Proof: The result can be proved by providing one example where the adversary controls the movement of nodes and it takes in the order of n rounds to broadcast a message in a network of diameter no greater than D.

Consider the example network shown in Figure 2(a), where the distance between neighboring nodes is R grid points. Node a, in the right hand corner of the third row from the bottom has message m to broadcast. Nodes b, c and d serve as visual points of reference. Without affecting the demonstration, the two lowest rows, represented by unfilled circles, can be ignored. The diameter

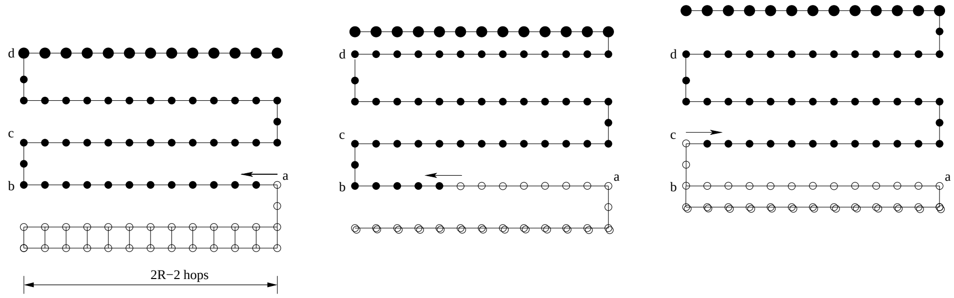

(a) Initial configuration.

(b) After R rounds.

(c) After $2R$ rounds.

Figure 2: Mobility adversary in a two-dimensional network.

of the network is equal to $10R - 1$. The total number of nodes in the network, n, is significantly greater than the diameter. Each row, except for the topmost, has $2R - 1$ nodes. The rows are $2R$ grid distance apart from each other. Successive rows are connected by one node, which is R grid distance from both the rows. Each large circle in the top row represents multiple nodes — the *reserve*. Again, the best that the broadcasting algorithm can do to reduce the broadcast completion time, upon receiving m for the first time in round r, is to send m to all its neighbors in *all* successive rounds.

We consider a *phase* containing $2R$ rounds. In the first round, node a initiates the broadcast. While the message is propagated, nodes move in the following fashion:

Lowest row: All nodes in the lowest row move up by one grid distance every round of the phase, covering a total of $2R$ grid distance.

Second lowest row: The nodes remain in place during the first R rounds and move up by one grid distance during the remaining R rounds of the phase.

Top most row: One node stays at each of the occupied positions in the top most row. All other nodes in these positions move up by one grid distance in each of the $2R$ rounds. The only exception is one node in the top row at the right most position which moves only during the first R rounds.

Figure 2(b) represents the situation after R rounds (unfilled circles represent nodes that have received m). Both the lowest and the top most rows moved up by R grid points. The two lowest rows now coincide. During this time, m has been propagated halfway through the (initially) third lowest row.

After the full $2R$ rounds, represented in Figure 2(c), the broadcast message m has reached node c and a new row has been constructed at the top.

The following assertions can be made:

1. By definition, the transformation took place in a phase of $2R$ rounds.

2. The resulting Figure 2(c) is identical (by symmetry and translation) to the initial Figure 2(a). Node c in Figure 2(c) is the starting point for the broadcast in the next phase.

3. The diameter of the graph stays unchanged during the transformation from Figures 2(a) to 2(b) to 2(c).

4. Graph connectivity is maintained throughout the transformation.

5. The remaining distance to cover after the transformation remains the same to before the transformation.

6. $2R$ nodes have been covered in the phase and $2R$ nodes have been consumed from the *reserve* in order to construct the new top-most row.

The same schema can be repeated as long as there are sufficient nodes in reserve in the top-most row, alternating between left most and right most position for the exception regarding the top most row. Since in Figure 2(a) there are $n - 4R$ nodes that have not received m, the schema can be repeated for a total of $\lfloor \frac{n-4R}{2R} \rfloor$ phases, each of duration $2R$ rounds. Hence, the number of rounds required for m to be received by all nodes is in the order of n. The lower bound for broadcasting is thus $\Omega(n)$. □

We complete this lower-bound result with an upper-bound:

Lemma 3 *In a two-dimensional MANET the upper-bound on time to complete broadcast is $O(n)$.*

Proof: Assume that in round 0 some node has message m to broadcast and sends m to its neighbors. Each node, once it has received m for the first time, sends m to all its neighbors in all successive rounds. We prove the result by showing that in each round – until all nodes have received m – at least one new node receives m for the first time.
Assume for contradiction that there is a round r during which no new node receives m for the first time. Let $Covered(r)$ denote the set of nodes that have received m before round r and $NonCovered(r) \neq \emptyset$ the set of nodes that have not received m at the beginning of round r. If no nodes in $NonCovered(r)$ receives m in round r, this means that, at the beginning of round r, no node in $Covered(r)$ is neighbor to a node in $NonCovered(r)$. A contradiction with the assumption that the nodes stay all the time connected. Thus in each round at least one node receives m for the first time. Hence, in n rounds every node has received m. □

Theorem 1 *The time complexity for broadcasting in two-dimensional MANETs is $\Theta(n)$.*

Proof: The proof follows directly from Lemmas 2 and 3.

4 Discussion

The preceding results show that knowledge of upper bound on the the diameter of a MANET does not help in limiting the time to complete a broadcast. Hence, the assumed trivial lower bound of $\Omega(D)$ is not true when node mobility is considered. Instead, $\Omega(n)$ rounds are required by *any* broadcasting protocol when the network nodes are mobile, which contradicts previous results, e.g., Bar-Yehuda et al. [1992] ($O((D + \log\ n/\epsilon) \cdot \log\ n)$) and Basagni et al. [1999] ($O(D\Delta \log^{\lfloor \log\ \Delta \rfloor} n)$, where Δ is the maximum degree of the network, i.e., $O(2D \log n)$ rounds when $\Delta = 3$).

A question to ponder is whether it would be possible to complete broadcasting in a MANET in $O(D)$ rounds if additional restrictions were imposed on the mobility of nodes and/or changes in MANET topology between rounds. One possible restriction worth investigating is:

- All MANET nodes are restricted to move only within a region, say a rectangle of dimensions $D1 \times D2$, such that the diameter of the region is possibly greater than $D \times R$. Can broadcasting be achieved in time that is a function of D, $D1$ and $D2$, regardless of n?

Other possible restrictions, worth investigating, to constrain the mobility of nodes and possibly complete broadcasting in time less than $O(n)$ are:

1. Only a subset of nodes are allowed to move in each round.
2. If a link is formed, nodes that are incident on the link move in a coordinated fashion such that the link survives some minimum number of rounds.
3. There is an upper bound on the number of links that can be added and removed in each round.

References

Noga Alon, Amotz Bar-Noy, Nathan Linial, and David Peleg. A lower bound for radio broadcast. *Journal of Computer and System Sciences*, 43(2):290–298, October 1991.

Reuven Bar-Yehuda, Oded Goldreich, and Alton Itai. On the time-complexity of reliable broadcast in multi-hop radio networks: An exponentional gap between determinism and randomization. *Journal of Computer and System Sciences*, 45(1):104–126, August 1992.

Stefano Basagni, Danilo Bruschi, and Imrich Chlamtac. A mobility-transparent deterministic broadcast mechanism for ad hoc networks. *IEEE/ACM Transactions on Networking (TON)*, 7(6): 799–807, December 1999.

Danilo Bruschi and Massimiliano Del Pinto. Lower bounds for the broadcast problem in mobile radio networks. *Distributed Computing*, 10(3):129–135, 1997.

Imrich Chlamtac and Shay Kutten. On broadcasting in radio networks — problem analysis and protocol design. *IEEE Transactions on Communications*, 33(12):1240–1246, December 1985.

Imrich Chlamtac and Orly Weinstein. The wave expansion approach to broadcasting in multihop radio networks. *IEEE Transactions on Communications*, 39(3):426–433, March 1991.

Bogdan S. Chlebus, Leszek Gąsieniec, Alan Gibbsons, Andrzej Pelc, and Wojciech Rytter. Deterministic broadcasting in ad hoc radio networks. *Distributed Computing*, 15(1):27–38, 2002.

Iris Gaber and Yishay Mansour. Centralized broadcast in multihop radio networks. *Journal of Algorithms*, 46(1):1–20, January 2003.

Eyal Kushilevitz and Yishay Mansour. An $\omega(d \log(n/d))$ lower bound for broadcast in radio networks. *SIAM Journal on Computing*, 27(3):702–712, June 1998.