

Improved Randomized Broadcast Protocols in Multi-hop Radio Networks

Chungki Lee James E. Burns* Mostafa H. Ammar

GIT-CC-93/14

February 1993

Abstract

This paper presents a suite of randomized broadcast protocols for the problem of broadcasting a message in multi-hop radio networks. The protocols are compared with the randomized broadcast protocol by Bar-Yehuda et al. The time complexity of one of the randomized broadcast protocols presented in this paper is shown, by simulation, to be much better than those of other protocols in most of the typical cases.

College of Computing
Georgia Institute of Technology
Atlanta, Georgia 30332-0280

*Author's current address: Bellcore, P.O. Box 7040, Red Bank, NJ 07701-5699.

1 Introduction

A radio network is a collection of radios that communicate with each other over radio channels. If all units in a network can hear each other, the radio network is single-hop; otherwise it is multi-hop and repeaters must be used to provide network connectivity. Radio networks have recently received significant attention due to the growing interest in cellular telephones and wireless communication networks. In this paper, the problem of broadcast in multi-hop radio networks is considered. Broadcast is initiated by a single node, called the *source*, which sends a message to all nodes in the network. Broadcast communication is an essential ingredient in many distributed network applications. Despite the broadcast nature of the radio medium, broadcast in multi-hop radio networks requires careful consideration due to the increased potential for collisions.

The broadcast problem in radio networks has been studied extensively in the literature [CK85, CK87, CW87, BII89, BGI91, BGI92, ABLP91, ABLP92]. Chlamtac and Kutten [CK85] showed that, given a radio network and a designated source, finding an optimal broadcast schedule that uses the minimum number of timeslots is NP-hard. Also they [CK85, CK87] presented broadcast protocols based on using a spanning tree. Chlamtac and Weinstein [CW87] presented a polynomial-time (centralized) algorithm for constructing a broadcast schedule that uses $O(D(\log N)^2)$ timeslots, where N is the number of nodes in the network and D is its diameter. Bar-Yehuda et al. [BGI92] presented a randomized broadcast protocol that runs in expected $O((D + \log(\frac{N}{\epsilon})) \log \Delta)$ timeslots to ensure that with probability $1 - \epsilon$ all nodes receive the message, where Δ is the maximum degree. For $D = O(1)$, they also showed a $\Omega(N)$ lower bound for deterministic broadcast protocols. Thus, for the broadcast problem, they claimed that there exist randomized protocols that are much more efficient than any deterministic one. Alon et al. [ABLP91] presented radio networks with diameter $D = O(1)$ in which every broadcast requires $\Omega((\log N)^2)$ timeslots, using a probabilistic argument.

In this paper, we present a suite of randomized broadcast protocols. Unlike deterministic broadcast protocols in which each node is assumed to know either the complete network topology or its neighbors, randomized broadcast protocols need no topological knowledge of the network except for some upper bounds on its size and its maximum degree. This property makes randomized protocols adaptive to changes in topology which occur throughout the execution, and resilient to non-malicious faults. Also randomized broadcast protocols are conceptually simple and require a minor amount of local computation. The (average) time complexity of one of the randomized broadcast protocols in this paper is shown, by simulation, to be much better than that of the randomized broadcast protocol by Bar-Yehuda et al. [BGI92] in many typical topologies.

This paper is organized as follows. In section 2, a model of radio networks with some necessary assumptions is presented. The randomized broadcast protocol by Bar-Yehuda et al. is summarized in section 3. A suite of randomized broadcast protocols is presented in section 4. Section 5 discusses the issue of termination of the protocols. Section 6 compares the various randomized protocols. Section 7 concludes the paper.

2 The Model

A radio network is modeled by a connected, undirected graph whose vertices represent nodes (radios) and whose edges represent two-way communication channels between their incident vertices. Thus, we assume that adjacent radios are of comparable power and are within range of each other. Nodes communicate in synchronous timeslots using radio transmissions. All nodes agree on the beginning of each timeslot (using, say, a satellite to provide timing signals). The length of a timeslot is assumed to be at least as long as the transmission time of the longest message plus the maximum propagation delay of messages between any pair of nodes.

The properties of radio communication are described by the following rules. In each timeslot, a node can act either as a transmitter or as a receiver, but not both. Thus, a transmitting node cannot directly detect whether or not its transmission is successful. A node receives a message in a timeslot if it acts as a receiver and exactly one of its neighbors transmits. If more than one neighbor of a node transmits, a collision occurs. If a collision occurs, receiving nodes cannot determine reliably that a collision has occurred because collisions are indistinguishable from background noise. Except for collisions, channels are assumed to be error-free.

A broadcast message is sent by a node, called the source, to all nodes in the network. Our randomized protocols make use of the following quantities:

- N , the number of nodes in the network (It is sufficient to know an upper bound on N).
- Δ , the maximum degree over all nodes in the network (It is sufficient to know an upper bound on Δ).

These quantities are assumed to be globally known and be constant throughout the execution of the protocols. They are used to parameterize the protocols. (It is sufficient for only the source to know the quantities since they could be sent along with the broadcast message.) Also, we assume as in [BGI92] that only a single broadcast by a single source is in progress at any point in time. The protocols require no other knowledge about the network, which allows them to tolerate changes in the network over time.

The performance measures considered are the *expected reception time* (the time until all nodes receive the broadcast message), and the *expected termination time* (the time until the protocol terminates). Another measure of interest in a randomized broadcast protocol is the *success probability* defined as the probability that all nodes receive a copy of the broadcast message before the protocol terminates.

3 The Original Randomized Broadcast Protocol

For completeness we first summarize the randomized broadcast protocol by Bar-Yehuda et al. and its properties [BGI92]. The basis for the protocol is a randomized procedure

called *Decay*. The main idea of the procedure *Decay* is to resolve conflicts among the transmitting neighbors of a receiver by randomly eliminating half of them at each timeslot. Suppose that $d \leq \Delta$ nodes are competing. Simultaneously, they all start a game of coin flips. At each timeslot half of the remaining nodes remove their candidacy on the average. It is expected that, with constant probability, before all nodes remove their candidacy there exists a timeslot with exactly one candidate. The procedure is described in Figure 1. In the description k is a parameter and m is the message to be sent. *Decay* is a randomized protocol with the property that if $k = 2\lceil \log \Delta \rceil$ and several neighbors of a node v use *Decay* to send messages, then with probability greater than $\frac{1}{2}$, node v receives one of the messages. Throughout the paper all logarithms are to base 2.

```

procedure Decay( $k, m$ );
repeat at most  $k$  times
    send  $m$  to all neighbors;
    flip (binary) coin
until coin = 0

```

Figure 1: Procedure Decay

The randomized broadcast protocol by Bar-Yehuda et al., called **Protocol 0**, makes use of *Decay*. The protocol executed by each node except the source is described in Figure 2. A network is said to execute a broadcast if the source transmits a broadcast message at timeslot 0 and every other node executes the protocol in Figure 2. Note that *Time* in the protocol is the current timeslot. The protocol has the following properties:

1. It terminates within $O((D + (\log \frac{N}{\epsilon})) \log \Delta)$ timeslots.
2. If $threshold = \lceil \log \frac{N}{\epsilon} \rceil$ and the network executes a broadcast then each node has received the broadcast message with probability $> 1 - \epsilon$, before termination.

Note that all participating nodes start executing *Decay* at the same timeslot (i.e., only at one plus integer multiples of $2\lceil \log \Delta \rceil$). Define the i th phase, or *phase* i , to be the duration of time from timeslot $1 + ((i - 1) * (2\lceil \log \Delta \rceil))$ to timeslot $i * (2\lceil \log \Delta \rceil)$. We observe the following features of **Protocol 0**.

1. Whenever a node receives a message during a phase, the node should wait to forward the message until the beginning of the following phase. Thus upon receiving a message for the first time, each node waits for $\lceil \log \Delta \rceil$ timeslots on the average before it starts to forward the message. This feature will delay propagation of the message.

```

procedure Protocol 0;
   $k := 2\lceil \log \Delta \rceil$ ;
  wait until receiving a message, say  $m$ ;
  do  $threshold$  times
    wait until  $(Time \bmod k) = 1$ ;
     $Decay(k, m)$ ;
  od

```

Figure 2: Protocol 0

2. During a phase a node can receive the same message several times from a neighbor because each neighbor keeps sending it until $coin = 0$. This continuous transmission by a neighbor helps to resolve collisions at nodes with more than one sending neighbor during the phase. However, there is some redundancy involved in the continuous transmission. We define a set X to be a *success set* of a node y in a timeslot if every node in X receives a message from y in that timeslot. Note that during a phase the size of the maximal success set is monotonically increasing with each timeslot until y stops sending. Thus, all timeslots of a phase before the final sending timeslot are redundant.

We conjecture that these features degrade the performance of the protocol. The protocols in the following section attempt to eliminate these features in one way or another to improve the performance.

4 New Randomized Broadcast Protocols

In this section a suite of new randomized broadcast protocols is presented.

4.1 Protocol 1

The first protocol, **Protocol 1**, is the same as **Protocol 0** except that after a node receives a message for the first time, the node tries to forward the message from the next timeslot without waiting. Thus the protocol described in Figure 3 also uses the procedure *Decay*. However, unlike **Protocol 0** the protocol does not require that all participants start executing *Decay* simultaneously. So the concept of a phase does not apply here. It seems that immediate transmission of the message by the node upon receiving it will speed up propagation of the message and allow more concurrent transmission of messages (i.e., we can make better use of spatial reuse of transmission timeslots). Thus we conjecture that **Protocol 1** is better than **Protocol 0** in most typical cases.

```

procedure Protocol 1;
   $k := 2\lceil \log \Delta \rceil$ ;
  wait until receiving a message, say  $m$ ;
   $t := \text{Time}$ ;
   $t_0 := (t + 1) \bmod k$ ;
  do threshold times
    wait until  $(\text{Time} \bmod k) = t_0$ ;
     $\text{Decay}(k, m)$ ;
  od

```

Figure 3: Protocol 1

4.2 Protocol 2

The second protocol we consider, **Protocol 2**, is also a variation of **Protocol 0**. The protocol tries to avoid an undesirable feature of **Protocol 0**, i.e., (possibly) continuous transmission by a node during a phase until it tosses 0. That is, unlike **Protocol 0**, a node flips a binary coin before it sends a message. If the value of the coin is 1, the node sends the message and keeps quiet during the remaining timeslots of the phase. Thus each node sends a message at most once during a phase. Note that in **Protocol 0** each node sends a message at least once during a phase. The protocol uses a new procedure *Decay2* described in Figure 4. As proved in the appendix, *Decay2* is a randomized protocol with the property that if $k = 2\lceil \log \Delta \rceil$ and several neighbors of a node execute *Decay2* simultaneously to send messages, the node receives one of the messages with probability $\geq \frac{1}{2}$.

```

procedure  $\text{Decay2}(k, m)$ ;
  do at most  $k$  times
    flip (binary) coin
    if  $\text{coin} = 1$  then send  $m$  to all neighbors; exit
  od

```

Figure 4: Procedure Decay2

The protocol executed by each node except the source is described in Figure 5. If the source transmits a broadcast message at timeslot 0 and every other node executes the protocol, the following properties (proved in the appendix) hold:

1. A broadcast terminates within $O((D + (\log \frac{N}{\epsilon})) \log \Delta)$ timeslots.

2. If $threshold = \lceil \log \frac{N}{\epsilon} \rceil$, each node has received the message with probability $> 1 - \epsilon$, before termination.

Note that all participating nodes start executing *Decay2* at the same timeslot.

```

procedure Protocol 2;
   $k := 2 \lceil \log \Delta \rceil$ ;
  wait until receiving a message, say  $m$ ;
  do  $threshold$  times
    wait until  $(Time \bmod k) = 1$ ;
     $Decay2(k, m)$ ;
  od

```

Figure 5: Protocol 2

Unlike **Protocol 0**, if $d \leq \Delta$ neighbors of a node start executing *Decay2* simultaneously, the node can receive a message exactly once from a neighbor during a phase. Also every node is given an opportunity to send until it actually sends during a phase. Notice that a node can receive the same message several times from different neighbors during a phase. **Protocol 2** requires that after a node has received the message during a phase, the node waits until the next phase to start sending. Again it seems that immediate transmission of the message by the node will speed up propagation of the message and allow multiple simultaneous transmissions of messages to be received. The protocol in the following subsection eliminates the wait.

4.3 Protocol 3

The third protocol, **Protocol 3**, is the same as **Protocol 2** except that after a node receives a message for the first time, the node tries to forward the message from the next timeslot without waiting. Thus the protocol uses the procedure *Decay2* again and is described in Figure 6. However, unlike **Protocol 2** the protocol does not require that all participants start executing *Decay2* simultaneously. We conjecture that this protocol is the best among the protocols in arbitrary networks.

4.4 Protocol 4

The fourth protocol, **Protocol 4**, uses a simple strategy. That is, after a node receives a message for the first time, the node sends the message with probability $\frac{1}{2^{\lceil \log \Delta \rceil}}$ for a certain number of times from the following timeslot. To guarantee the termination of the protocol at non-source nodes, we limit the number of times to a *threshold* value. Due to the random nature of the protocol, only if the threshold value is infinite is the protocol

```

procedure Protocol 3;
   $k := 2\lceil \log \Delta \rceil$ ;
  wait until receiving a message, say  $m$ ;
   $t := \text{Time}$ ;
   $t_0 := (t + 1) \bmod k$ ;
  do  $\text{threshold}$  times
    wait until  $(\text{Time} \bmod k) = t_0$ ;
     $\text{Decay2}(k, m)$ ;
  od

```

Figure 6: Protocol 3

```

procedure Protocol 4;
   $k := 2\lceil \log \Delta \rceil$ ;
  wait until receiving a message, say  $m$ ;
  do  $\text{threshold}$  times
    generate a uniform random number,  $r$ ;
    if  $r < \frac{1}{k}$ 
      then send  $m$  to all neighbors
    else keep quiet;
  od

```

Figure 7: Protocol 4

guaranteed to broadcast a message before termination. The protocol is described in Figure 7. In the protocol at most $\frac{\Delta}{2\lceil \log \Delta \rceil}$ neighbors of a node on the average can send a message during a timeslot, which helps to reduce the number of collisions. Note that $\frac{\Delta}{2\lceil \log \Delta \rceil}$ is relatively very small compared with Δ .

4.5 Determining Appropriate Threshold Values

The *success probability* of each protocol is strongly dependent on the value of the *threshold* parameter. For **Protocol 0** it has been shown that if $\text{threshold} = \lceil \log \frac{N}{\epsilon} \rceil$ then *success probability* $> 1 - \epsilon$. With the exception of **Protocol 2**, we have not yet been able to analytically determine *threshold* values required to achieve given success probabilities for our protocols.

For **Protocol 2** it can be shown (see Appendix) that with $\text{threshold} = \lceil \log \frac{N}{\epsilon} \rceil$,

success probability $> 1 - \epsilon$. For the other protocols (**Protocol 1, 3, and 4**) and after considerable experimentation we use the following thresholds

- **Protocols 1 & 3:** $threshold = \lceil \log \frac{N}{\delta} \rceil$
- **Protocol 4:** $threshold = 4 \lceil \log \frac{N}{\delta} \rceil$

where δ is a small constant. Our simulation results indicate that for $\delta = 0.01$ and with the above thresholds, success probabilities are > 0.99 .

5 Termination Time Analysis of the Protocols

We prove that if the reception time of any protocol is known, the termination time can be computed easily. Our results in the next section show the average reception time for the various protocols in several network environments.

Theorem 1. *Let t_0 be the reception time of **Protocol 1 (Protocol 3)** in a given network with N nodes and at most Δ degree. Assume that the source initiates a broadcast at timeslot 0 in **Protocol 1 (Protocol 3)**. Let k be $2 \lceil \log \Delta \rceil$. Then the termination time of **Protocol 1 (Protocol 3)** is $t_0 + k \lceil \log(\frac{N}{\epsilon}) \rceil$.*

Proof: Let x be a node that received a broadcast message last at timeslot $t_0 - 1$. Then node x will terminate last because every other node starts **Protocol 1 (Protocol 3)** not later than node x . The execution time of **Protocol 1 (Protocol 3)** by each node is $k \lceil \log(\frac{N}{\epsilon}) \rceil$ because each node executes the procedure *Decay (Decay2)* $\lceil \log(\frac{N}{\epsilon}) \rceil$ times. This implies the theorem. \square

Theorem 2. *Let t_0 be the reception time of **Protocol 2** in a given network with N nodes and at most Δ degree. Assume that the source initiates a broadcast at timeslot 0 in **Protocol 2**. Let k be $2 \lceil \log \Delta \rceil$. Then the termination time of **Protocol 2** is $t_0 + (k - (t_0 - 1) \bmod k) \bmod k + k \lceil \log(\frac{N}{\epsilon}) \rceil$.*

Proof: Let x be a node that received a broadcast message last at timeslot $t_0 - 1$. Then node x will terminate last because every other node starts **Protocol 2** not later than node x . After it received the message during a phase, it waits for the beginning of the next phase. The number of the timeslots to wait is the difference between the first timeslot of the next phase and the timeslot when it received a message. Since phases always start at timeslot one plus integer multiples of k , it is $(k - (t_0 - 1) \bmod k) \bmod k$. Then node x will execute the procedure *Decay2* $\lceil \log(\frac{N}{\epsilon}) \rceil$ times. This implies the theorem. \square

Theorem 3. *Let t_0 be the reception time of **Protocol 4** in a given network with N nodes and at most Δ degree. Assume that the source initiates a broadcast at timeslot 0 in **Protocol 4**. Then the termination time of **Protocol 4** is $t_0 + 4\lceil\log(\frac{N}{\epsilon})\rceil$.*

Proof: Let x be a node that received a broadcast message last at timeslot $t_0 - 1$. Then node x will terminate last because every other node starts **Protocol 4** not later than node x . The execution time of **Protocol 4** by each node is $4\lceil\log(\frac{N}{\epsilon})\rceil$ after it received the message. This implies the theorem. \square

6 Reception Time Analysis of the Protocols

Determining reception times of our protocols is very difficult for arbitrary graphs, so we first consider some graphs with simple structure that lend themselves to analysis. While these are not very realistic models of radio networks, they do help us to understand the behavior of the protocols in extreme conditions. We then consider graphs that are more representative of realistic radio networks.

The network topologies considered are lines, complete binary trees, meshes, and geometric graphs. These topologies give a broad selection of different diameters and degrees of nodes. Note that the time complexity of the protocols strongly depends on the diameter of the network and degrees of nodes. The maximum degrees and diameters of these topologies with N nodes are summarized in Table 1. The maximum degree and diameter of geometric graphs (defined below) are obtained empirically.

<i>Topology</i>	<i>Maximum Degree</i>	<i>Diameter</i>
line	2	$N - 1$
complete binary tree	3	$\lceil\log N\rceil$
mesh	4	$O(\sqrt{N})$
geometric graph	$O(\log N)$	$O(\sqrt{\frac{N}{\log N}})$

Table 1: Maximum Degrees and Diameters of Network Topologies

The reception time of the protocols is the time (the number of timeslots) until every node receives the message from the source. Where possible we derive average reception times analytically. In other situations we resort to simulation. For a network of a given size except geometric graphs, 100 runs with different random number seeds are made and the performance measures of those runs are averaged. In geometric graphs, for a given size, 100 different networks are generated, a run is made per network, and the performance measures of those runs are averaged.

6.1 Graphs of Small Degree

For graphs of small degree, we consider

- a line of N nodes with the source at one end,
- a complete binary tree with the source being the root,
- a mesh of N (even) nodes with the source at the upper left corner.

6.1.1 Lines

Line networks are not very realistic for radio networks, but provide a simple structure that lends itself to precise analysis.

Theorem 4. *Assume that the graph corresponding to a network is a line with N nodes with the source at one end, where $N > 2$. Also suppose that the source initiates a broadcast at timeslot 0. Then all nodes using **Protocol 0** receive the message with probability one by time $2(N - 2)$.*

Proof: A neighbor of the source receives a message at timeslot 0. After that, it takes $k = 2\lceil \log \Delta \rceil = 2$ timeslots to forward the message one distance down the line network, because the message is received at the first timeslot of a phase but the receiver has to wait until the beginning of the next phase. It takes only one timeslot for the node at the other end to receive a message after its neighbor sends a message. Therefore the total time needed is $1 + 2(N - 3) + 1 = 2(N - 2)$. \square

Theorem 5. *Assume that the graph corresponding to a network is a line with N nodes with the source at one end, where $N > 1$. Also suppose that the source initiates a broadcast at timeslot 0. Then all nodes using **Protocol 1** receive the message with probability one by time $N - 1$.*

Proof: In **Protocol 1** every node sends a message in the following timeslot after it receives the message. Also when a node sends a message, there is no collision at its intended receiver. Therefore every node receives the message by time $D = N - 1$. \square

Theorem 6. *Assume that the graph corresponding to a network is a line with N nodes with the source at one end, where $N > 1$. Also suppose that the source initiates a broadcast at timeslot 0. Then all nodes using **Protocol 2** receive the message on the average by time $\frac{1}{3}(8N - 13)$.*

Proof: A neighbor of the source receives a message at timeslot 0. Let us calculate how many phases are needed to forward a message one distance down the line network. At the start of the phase each node with a message sends the message with probability $\frac{1}{2}$.

If it did not send, it sends the message with probability $\frac{1}{2}$ in the second (last) timeslot of the phase. So the probability that a node sends a message during a phase is

$$\frac{1}{2} + \left(\frac{1}{2}\right)\left(\frac{1}{2}\right) = \frac{3}{4}.$$

Thus the probability that $i \geq 1$ phases are needed for a node to send a message successfully is $(\frac{1}{4})^{i-1}(\frac{3}{4})$. Therefore the expected number of phases needed for a node to send a message successfully is

$$\sum_{i=1}^{\infty} i \left(\frac{1}{4}\right)^{i-1} \left(\frac{3}{4}\right) = \frac{4}{3}.$$

Since the duration of a phase is 2, the expected time that all nodes receive a message is

$$1 + \left(\frac{4}{3}\right)2(N-2) = \frac{1}{3}(8N-13).$$

□

Theorem 7. Assume that the graph corresponding to a network is a line with N nodes with the source at one end, where $N > 1$. Also suppose that the source initiates a broadcast at timeslot 0. Then all nodes using **Protocol 3** receive the message on the average by time $2N - 3$.

Proof: A neighbor of the source receives a message at timeslot 0. Upon receiving a message, a node sends the message with probability $\frac{1}{2}$ in the following timeslot. If it did not send, it sends the message with probability $\frac{1}{2}$ in the second following timeslot. Then it repeats this sending process. Thus the probability that $i \geq 1$ timeslots are needed for a node to send a message successfully is $(\frac{1}{2})^i$. Therefore the expected number of timeslots needed for a node to send a message successfully is

$$\sum_{i=1}^{\infty} i \left(\frac{1}{2}\right)^i = 2.$$

The expected time that all nodes receive a message is

$$1 + 2(N-2) = 2N - 3.$$

□

Theorem 8. Assume that the graph corresponding to a network is a line with N nodes with the source at one end, where $N > 1$. Also suppose that the source initiates a broadcast at timeslot 0. Then all nodes using **Protocol 4** receive the message on the average by time $2N - 3$.