An $\Omega(D \log(N/D))$ Lower Bound for Broadcast in Radio Networks

Eyal Kushilevitz* Yishay Mansour†

Abstract

We show that for any *randomized* broadcast protocol for radio networks, there exists a network in which the expected time to broadcast a message is $\Omega(D \log(N/D))$, where D is the diameter of the network and N is the number of nodes. This implies a tight lower bound of $\Omega(D \log N)$ for all $D \leq N^{1-\varepsilon}$, where $\varepsilon > 0$ is any constant.

1 Introduction

Traditionally, radio networks received a considerable attention due to their military significance. The growing interest in cellular telephones and wireless communication networks has reinforced the interest in radio networks. The basic feature of radio networks, that distinguishes them from other networks, is that a processor can receive a message only from a *single* neighbor at a certain time. If two (or more) neighbors of a processor transmit concurrently, then the processor would not receive either messages.

In many applications, the users of the radio network are mobile, and therefore the topology is instable. For this reason, it is desirable for radio-networks algorithms to refrain from making assumptions about the network topology, or about the information that processors have concerning the topology. In this work we assume that none of the processors have initially any topological information, except for the size of the network and its diameter.[1] See [Ten81, Gal85, BGI92, BGI91] for a discussion on this model and related models.

We study *broadcast* protocols; those protocols are initiated by a single processor (the *originator*) that has a message M he wishes to propagate to all the other processors in the network. In many of the radio networks applications (e.g. cellular phones) broadcast is a central primitive which is frequently used, for example to perform a network-wide search for a user.

Bar-Yehuda et al. [BGI92] present a *random-*

*Aiken Computation Lab., Harvard University. e-mail: eyalk@das.harvard.edu . Research supported by research contracts ONR-N0001491-J-1981 and NSF-CCR-90-07677.

†Computer science Dept., Tel-Aviv University and IBM – T. J. Watson Research Center. e-mail: mansour@cs.tau.ac.il

12th ACM Symposium on Principles on Distributed Computing, Ithaca NY

[1] Usually, when the topology is instable, the diameter is unknown to the processors and only a bound on the size of the network is available. However, since we are proving a *lower bound* this assumption only makes the result stronger.

ized broadcast algorithm, that runs in expected $O(D \log N + \log^2 N)$ time slots, where N is the number of processors in the network and D is its diameter. In contrast they show that for any *deterministic* broadcast algorithm there are networks of constant diameter on which the algorithm needs $\Omega(N)$ time slots.

Alon et al. [ABLP91] made the first step towards proving the optimality of the upper bound of [BGI92]. Their result can be viewed as a graph-theoretic result; they show that there exist networks of diameter $D = 3$ on which any schedule needs at least $\Omega(\log^2 N)$ time slots. This lower bound shows that there are networks on which broadcast requires this many time-slots, and it matches the known upper bounds [BGI92, CW87], in the case of constant-diameter networks.

In this work we complete the picture by proving an $\Omega(D \log(N/D))$ lower bound. Our result is of a different nature; we show that for any *randomized* broadcast algorithm and parameters N and D there is an ordering of the N processors in a network of diameter D such that the expected number of time slots, used by the algorithm, is $\Omega(D \log(N/D))$. For $D \leq N^{1-\epsilon}$ this gives an $\Omega(D \log N)$ lower bound. Hence, it proves the tightness of the upper bound of [BGI92] for all N and $D \leq N^{1-\epsilon}$. Moreover, the lower bound holds even if each of the N processors is allowed to use a *different* program (e.g. the processors can use their IDs). On the other hand, it is worth pointing out that it remains an open problem whether there exists a network for which *any* algorithm requires $\Omega(D \log(N/D))$ time-slots. Namely, it may be the case that for any network there is a schedule that completes in $O(\log^2 N + D)$ time slots. In fact, the networks constructed in our lower-bound proof have this property, which hints that the lower bound heavily relies on the lack of topological knowledge at the processors.

Broadcast in radio networks has received considerable attention in previous works. [CW87] present a deterministic sequential algorithm that given the network finds, in polynomial time, a legal schedule that requires at most $O(D \log^2 N)$ time slots. Broadcast that is based on using a spanning tree was suggested in [CK85a, CK87]. In [BII89] it is shown how to reduce the amortized cost per broadcast by using a BFS tree. Simulation of point to point networks on radio networks is found in [CK85b, ABLP92, BGI91].

An important issue in the study of radio networks is whether collisions can be detected; namely, whether a listener can distinguish between the case that none of its neighbors transmit and the case that two or more of them transmit. In our model it is assumed that the listener cannot distinguish between the two cases (say, it hears noise in both cases). There is another common model in which it is assumed that the two cases are distinguishable (say, if no neighbor transmits the listener hears silence while if two or more neighbors transmit the listener hears noise). A discussion justifying both models can be found in [Gal85, BGI92]. Willard [Wil86] studies a broadcast problem, in a single multi-access channel under this second model (i.e., when collision-detection is available). He shows matching upper and lower bounds of $\Theta(\log \log n)$ expected time slots[2] in this model.

[2] Willard shows an $\Omega(\log \log n)$ lower bound in the single multi-access channel model. Although this bound applies to a different model, it should be noted that his bound is also significantly restricted by the types of algorithms for which it applies. In particular, he requires independence between the decision whether to transmit in a certain time slot and the decisions made in previous time slots. In our case such a restriction is unacceptable as the upper bound of [BGI92] has such dependencies. Also he does not handle the case where each processor may use different program.

for the same problem in our model. Again, this lower bound holds even if the processors may use different programs. Hence, we demonstrate a provable exponential gap between these two models.

The rest of this paper is organized as follows: Section 2 contains some necessary definitions. Section 3 contains the proof of the main lemma in the *uniform* case, where all the processors use the same program. Section 4 contains the proof of the main lemma in the *non-uniform* case, where processors may use different programs. The proof for this case is based on a probabilistic reduction to the uniform case. Finally in Section 5 we prove the main theorem. The proof involves constructing a "difficult" network in a probabilistic way.

2 Preliminaries

A *radio network* is described by an undirected graph $G(V, E)$.[3] The nodes of the graph represent processors of the network, and an edge between nodes v and u implies that v can send messages to u (and vice-versa). The *neighborhood* of a node u includes all the nodes v such that there is an edge (u, v) in E.

The time is viewed as divided into slots (or *rounds*), and a node can either transmit or not in a given slot. A radio network has the property that if two or more nodes in the neighborhood of a node u transmit at the same time slot then none of the messages is received at u. I.e. a node receives successfully a transmission at time t, if exactly one of its neighbors broadcasts at time t.

Each processor in the radio network uses a *probabilistic* program. This program defines whether not. As we are not concerned with the computational power of the processors we can simply view this program as a probability distribution, which may depend on the history. A protocol is *uniform* if all processors use the same program. Otherwise, if each processor may have a different program, the protocol is *non-uniform.*

A broadcast protocol is a protocol that is initiated by a single processor, called *originator*, that holds a message M (any other processor is inactive until receiving a message for the first time). The aim of the protocol is that each processor in the network will receive a copy of the message M.

[3]None of the results presented in this work will be changed if the network is a directed one. However, it is common in this area to assume that the network is undirected.

3 Uniform Processors

In this section we prove the main lemma for the uniform case, where all processors use the same program. It shows that if there are n processors[4] arranged in a clique, then there exists a t $(2 \leq t \leq n)$ such that if t processors wish to transmit (we call these t processors the *participants*) then the expected number of rounds (time slots) until a round in which exactly one of them transmits is $\Omega(\log n)$. In fact, we show that this is the case for most of the t's of the form $t = 2^i$. Note that the assumption that the topology is not known to the processors, in the context of this lemma, means that t, the number of processors that are trying to transmit, is not known to any processor. We can view the scenario as having a family of networks with $n + 1$ nodes, composed from a clique of size n and an originator which is connected to t of the nodes in the clique. The (unknown) topology is chosen to be one of these networks.

[4]Note that we use here n (and not N) as the number of processors. This will be convenient while using the lemma in the proof of the theorem.

Lemma 1 *Let Π be a broadcast protocol, let the network be as above, and let n be an upper bound on the number of participants. We call a round* successful *if exactly one processor transmits. Let $E(T_\ell^\Pi)$ denotes the expected number of rounds until the* first *successful round, given that the number of participants is 2^ℓ (the expectation is taken over the probabilistic choices of the processors). Then,*

$$E_\ell[E(T_\ell^\Pi)] = \Omega(\log n),$$

where ℓ is chosen uniformly from the range $1 \le \ell \le \log n$.

Proof: The first observation that we make is that the lemma deals only with the *first* success. Therefore, we can assume that the (probabilistic) decision in which rounds a processor tries to transmit is made at the beginning of the protocol. This is done by letting each of the 2^ℓ processors to choose whether to transmit in round s or not in the same way as it chooses in the original protocol, in the case that all previous rounds were unsuccessful. Clearly, as far as the *first* success is concerned this modification has no effect on the protocol. Hence, the decision of a processor whether to transmit in round s may depend on the round number, s, and the probabilistic choices of the processor in first $s-1$ rounds but it does *not* depend on choices made by other processors.[5] Therefore, we can think about the processors as if they choose *in advance*, for every round $s = 1, 2, \ldots$, whether they will try to transmit.[6]

[5] The message that the processors need to transmit also has influence on their decisions. However, it can be thought as part of the program used by the processors.

[6] To avoid measurability concerns it is convenient to assume that the protocol is such that s is in the range $1, \ldots, F$, for some *finite* F. If this is not the case, we can always choose F such that the probability of choosing only in the range $1, \ldots, F$ is arbitrarily close to 1. This will cause minor changes in our proof.

For simplicity of notations, we assume that n is a power of 2. Define,

$$p_{s,\ell} \triangleq Pr\ (\ \text{failure in rounds } 1, \ldots, s-1 \textbf{ and } \text{success in round } s|2^\ell \text{ participants}).$$

As the events described in the definition are disjoint (for *fixed* ℓ and different s's), and assuming that the protocol succeeds with probability 1 (no matter what ℓ is), we have for all ℓ

$$\sum_{s=1}^{\infty} p_{s,\ell} = 1. \tag{1}$$

At some point in the proof below, it will be inconvenient that $p_{s,\ell}$ depends on events happen in previous rounds. However, we can get rid of this dependency simply by writing

$$p_{s,\ell} \le Pr(\text{ success in round } s|2^\ell \text{ participants}). \tag{2}$$

Now, let λ be a parameter (to be fixed later). By the definition of $E(T_\ell^\Pi)$, we have

$$\begin{aligned} E(T_\ell^\Pi) &= \sum_{s=1}^{\infty} s \cdot p_{s,\ell} \\ &= \sum_{s=1}^{\lambda\cdot\log n-1} s \cdot p_{s,\ell} + \sum_{s=\lambda\cdot\log n}^{\infty} s \cdot p_{s,\ell} \\ &\ge \sum_{s=1}^{\lambda\cdot\log n-1} (s - \lambda\cdot\log n + \lambda\cdot\log n)\cdot p_{s,\ell} \\ &\quad + \sum_{s=\lambda\cdot\log n}^{\infty} \lambda\cdot\log n\cdot p_{s,\ell} \end{aligned}$$

By equation (1),

$$\begin{aligned} &= \lambda\cdot\log n + \sum_{s=1}^{\lambda\cdot\log n-1} (s-\lambda\cdot\log n)\cdot p_{s,\ell} \\ &= \lambda\cdot\log n - \underbrace{\sum_{s=1}^{\lambda\cdot\log n-1} (\lambda\cdot\log n - s)\cdot p_{s,\ell}}_{\triangleq\alpha(\ell)} \end{aligned} \tag{3}$$

To complete the proof of the lemma it is enough to show that, for a random ℓ, $\alpha(\ell) \leq c \log n$, where $c < \lambda$ is a constant. To this end we consider the sum of the $\alpha(\ell)$'s and show that this sum is bounded by $c \log^2 n$. Formally, define

$$\alpha^* \triangleq \sum_{\ell=1}^{\log n} \alpha(\ell) = \sum_{\ell=1}^{\log n} \sum_{s=1}^{\lambda \cdot \log n - 1} (\lambda \cdot \log n - s) \cdot p_{s,\ell}$$

We now use equation (2) and change the order of summation and we get

$$\alpha^* \leq \sum_{s=1}^{\lambda \cdot \log n - 1} (\lambda \cdot \log n - s) \sum_{\ell=1}^{\log n} Pr(\text{success in round } s \mid 2^\ell \text{ participants}) \quad (4)$$

The next claim gives a bound on the last sum. Intuitively it says that you cannot have high probability of success in (a fixed) round s, for more than a few values of 2^ℓ. Formally,

Claim 2 *For any s,*

$$\sum_{\ell=1}^{\log n} Pr(\textit{success in round } s \mid 2^\ell \textit{ participants}) < 2.$$

Proof: Fix s. As already discussed, we assume that the processors make all their choices in advance. The *history* of choices of a processor is a string in $\{0,1\}^{s-1}$, where the value of the ith bit means trying ("1") or not trying ("0"). Define

$$q(s) \triangleq Pr(\text{trying in round } s) = \sum_{\text{history } h} Pr(h) \cdot Pr(\text{trying in round } s|h).$$

Note that $q(s)$ does *not* depend on ℓ. We assume, without loss of generality, that $q(s) > 0$ (rounds with $q(s) = 0$ can be omitted from the protocol).

Recall that a successful round is one in which exactly one processor is trying to transmit. Therefore,

$$Pr(\text{success in round } s \mid 2^\ell \text{ participants}) = 2^\ell \cdot q(s) \cdot (1 - q(s))^{2^\ell - 1}.$$

We get

$$\begin{aligned} &\sum_{\ell=1}^{\log n} Pr(\text{success in round } s|2^\ell \text{ participants}) \\ &= \sum_{\ell=1}^{\log n} 2^\ell q(s)(1-q(s))^{2^\ell-1} \\ &= q(s) \sum_{\ell=1}^{\log n} 2^\ell (1-q(s))^{2^\ell-1} \\ &\leq 2 \cdot q(s) \sum_{j=1}^{n-1} (1-q(s))^j \\ &= 2 \cdot q(s) \cdot \frac{1-(1-q(s))^n}{q(s)} < 2 \end{aligned}$$

which completes the proof of the claim.

Using this claim and equation (4) we have

$$\sum_{\ell=1}^{\log n} \alpha(\ell) = \alpha^* < 2 \sum_{s=1}^{\lambda \cdot \log n - 1} (\lambda \cdot \log n - s) < \lambda^2 \cdot \log^2 n. \quad (5)$$

By the definitions, and equation (3) we have,

$$E_\ell[E(T_\ell^\Pi)] = \sum_{\ell=1}^{\log n} \frac{[E(T_\ell^\Pi)]}{\log n} \geq \sum_{\ell=1}^{\log n} \lambda - \frac{\alpha(\ell)}{\log n} = \lambda \log n - \frac{\sum_{\ell=1}^{\log n} \alpha(\ell)}{\log n}.$$

By equation (5), this is greater than

$$\lambda \log n - \frac{\lambda^2 \log^2 n}{\log n} = (\lambda - \lambda^2) \cdot \log n.$$

By choosing $\lambda = 1/2$ we get $E_\ell[E(T_\ell^\Pi)] > 1/4 \cdot \log n$ as desired. □

4 Non-Uniform Processors

In this section we prove the main lemma for the non-uniform case, where the n processors may use *different* programs. The main idea of the proof is to "reduce" the non-uniform case to the uniform one, and use the result of the previous section (Lemma 1).

Lemma 3 *Let Π be a protocol for n distinct processors $P_1, \ldots, P_n$ that run (possibly) different programs. Let $E(T_\ell^\Pi)$ denotes the expected number of rounds until the* first *successful round, given that a* random *set of 2^ℓ processors participates (the expectation is taken over the choice of the set and the probabilistic choices made by the processors). Then,*

$$E_\ell[E(T_\ell^\Pi)] = \Omega(\log n),$$

where ℓ is chosen uniformly from the range $1 \leq \ell \leq \log n$.

Proof: As argued in the previous section each program can be thought of as a "schedule" – a choice of a subset of rounds in which the processor will transmit. Processor P_i chooses its schedule from a distribution μ_i.

We now define, based on the (possibly different) programs used by $P_1, \ldots, P_n$, a new program that will be used by each of L *uniform* processors $Q_1, \ldots, Q_L$: Processor Q_j chooses at random $1 \leq i \leq n$ and simulates the program of processor P_i. Namely, it chooses a schedule s with probability $\frac{1}{n}\sum_{i=1}^{n}\mu_i(s)$, where $\mu_i(s)$ is the probability that processor P_i chooses the schedule s. We denote by $c(Q_j)$ the processor P_i that Q_j chose to simulate. We emphasize that all the Q_j's run the *same* program (i.e. they are uniform), and that different Q_j's may choose to simulate the same processor P_i (we will choose L "small enough" so that this will happen only with a "small" probability).

The following claim says that given that for $Q_1, \ldots, Q_{2^\ell}$ all the corresponding $c(Q_j)$'s are distinct, then the probability distribution of the schedules chosen by the Q_j's is the same as that of a *random* set of 2^ℓ processors P_i.

Claim 4 *Let $Q = \{Q_1, \ldots, Q_{2^\ell}\}$. For every $Q_j \in Q$, let $c(Q_j)$ be a random processor P_i. If $\forall j_1, j_2 : c(Q_{j_1}) \neq c(Q_{j_2})$, then $P = \{c(Q_j) | Q_j \in Q\}$ is a random set of 2^ℓ processors (in $P_1, \ldots, P_n$), and the following holds: for every choice of 2^ℓ schedules $\vec{s}_{2^\ell} = (s_1, \ldots, s_{2^\ell})$*

$$Pr[\vec{s}_{2^\ell} | \textit{processors } Q \textit{ run }] = Pr[\vec{s}_{2^\ell} | \textit{processors } P \textit{ run }].$$

The following claim is the main tool in the reduction from the non-uniform case to the uniform case.

Claim 5 *Let Q be as above, and let $Q' = \{Q'_1, \ldots, Q'_{2^\ell}\}$ be a set of 2^ℓ processors. Each processor Q'_j run the program of Q_j at the odd steps and the [BGI92] program at the even steps (note that the [BGI92] program is also a uniform program, and therefore so is the program run by the processors Q'). Let β_ℓ be the probability that $\forall j_1, j_2 : c(Q_{j_1}) \neq c(Q_{j_2})$. Let $T_\ell^{Q'}$ be the random variable indicating the time of first success when the 2^ℓ identical programs in Q' run, and recall that T_ℓ^Π is a random variable indicating the time of first success when a random subset of 2^ℓ distinct programs $P_{i_1}, \ldots, P_{i_{2^\ell}}$ run. Then,*

$$E[T_\ell^{Q'}] \leq 2\beta_\ell E[T_\ell^\Pi] + 8(1 - \beta_\ell)\log n.$$

In the above claim we mixed the given (unknown) protocol with the [BGI92] protocol. This is because we have no guarantee about the running time of the simulation, in case that some Q_j's choose to simulate the same P_i. For example, a protocol that lets processor P_i transmit at time slot i would not terminate if all the Q_j simulate the same processor P_i.

Proof: Let *unique* be the event that $\forall Q_{j_1}, Q_{j_2} \in Q' : c(Q_{j_1}) \neq c(Q_{j_2})$. Then,

$$\begin{aligned} E[T_\ell^{Q'}] &= E[T_\ell^{Q'} | unique] \cdot Pr[unique] \\ &\quad + E[T_\ell^{Q'} | not\ unique] \cdot Pr[not\ unique] \end{aligned}$$

By definition

$$Pr[unique] = \beta_\ell.$$

By claim 4,

$$E[T_\ell^{Q'} | unique] \leq 2E[T_\ell^{\Pi}].$$

In the case that the choices of $c(Q_j)$ are not unique we cannot use the properties of the original protocol. However, we can use the fact that the [BGI92] protocol has expected time until first success of at most $4 \log n$. Therefore,

$$E[T_\ell^{Q'} | not\ unique] \leq 8 \log n$$

which completes the proof of the claim. •

The next claim says that with "high probability" the choices $c(Q_j)$ are unique.

Claim 6 *Let β_ℓ be the probability that $\forall j_1 \neq j_2 : c(Q_{j_1}) \neq c(Q_{j_2})$ and assume that $2^\ell \leq n^{1/4}$. Then,*

$$\beta_\ell > 1 - \frac{1}{\sqrt{n}}$$

Proof: Note that

$$Pr[j_1 \neq j_2 \text{ and } c(Q_{j_1}) = c(Q_{j_2})] = \frac{1}{n}.$$

Therefore,

$$\begin{aligned} \beta_\ell &= Pr[\forall j_1 \neq j_2 : c(Q_{j_1}) \neq c(Q_{j_2})] \\ &\geq 1 - \binom{2^\ell}{2} \frac{1}{n} \end{aligned}$$

Since $2^\ell \leq n^{1/4}$ the lemma follows.

Let $L = n^{1/4}$. By Claims 5 and 6,

$$\begin{aligned} E[T_\ell^{Q'}] &\leq 2\beta_\ell E[T_\ell^{\Pi}] + (1 - \beta_\ell) 8 \log n \\ &\leq 2E[T_\ell^{\Pi}] + \frac{8 \log n}{\sqrt{n}} \end{aligned}$$

or

$$E[T_\ell^{\Pi}] \geq \frac{1}{2} E[T_\ell^{Q'}] - \frac{4 \log n}{\sqrt{n}}.$$

We now take the expectation over all values $1 \leq \ell \leq \log L$ and get

$$E_\ell[E[T_\ell^{\Pi}]] \geq \frac{1}{2} E_\ell[E[T_\ell^{Q'}]] - \frac{4 \log n}{\sqrt{n}}.$$

By Lemma 1

$$E_\ell[E[T_\ell^{Q'}]] = \Omega(\log L) = \Omega(\log n),$$

which implies that

$$E_\ell[E[T_\ell^{\Pi}]] = \Omega(\log n)$$

as desired. □

5 Main Theorem

In this section we prove the main theorem. We show that for every broadcast algorithm that does not know the topology of the network, for every N, and every D, there exist networks of N processors and diameter D such that the expected running time of the algorithm (until all processors receive the message) is $\Omega(D \log(N/D))$. This implies a similar lower bound for the case where the *worst case* running time is considered and a small probability of error is allowed (which is the scenario in which the upper bound of [BGI92] is described).

Given an algorithm, and the values N and D we construct a network as follows. Let $n = N/D$, and assume for simplicity that n is a power of 2. We construct a complete layered network of $D+2$ layers. The first layer (layer 0) contains one node, s, which will be the originator of the broadcast. Each of the next D layers (layers $1, 2, ..., D$) consists of $n_i = 2^{\ell_i} \leq n$ nodes, where ℓ_i is chosen uniformly (and independently for each layer i) in the range $1, ..., \log n$. The last layer contains all the other nodes (so that the total number of nodes will be N). Each node in layer i is connected to all nodes in layers $i-1$ and $i+1$.

Recall that the topology of the network is not known to the processors (If the topology was known, then an efficient uniform protocol would be to let a processor at layer i broadcast with probability $1/n_i$, the expected time is $O(D)$. A non-uniform protocol that knows the topology simply lets one node in each layer to transmit). The algorithm can depend however on other information that the processors have, in particular – the clock, the content of the message, etc.

We discuss the uniform case, in the sense that all the processors at layer i have the same protocol. The extension to the non-uniform case employs the techniques of the previous section, and the proof is the same but the notation becomes cumbersome. (In particular, in the non-uniform case, at each layer i we will choose not only n_i but also a random set of n_i processors.) The main property that this construction has is the following: For all i, and all runs of the protocol, all the processors in layer i have the same view; every message received at one of these processors, is received by all other processors at the same time. Therefore the broadcast progresses in a layer-by-layer fashion. Moreover, this implies that all the processors in layer i choose schedules according to the same distribution μ (the choice of μ depends on the history but all the processors of layer i share the same history), which allows us using Lemma 1.

Finally, before going into the details, we make one more assumption that makes our argument simpler. We give the processors of layer i, at the time they get the first message from a processor in layer $i-1$, all the other messages they will get from layer $i-1$ at the future as well as the actual values of $\ell_1, \ldots, \ell_{i-1}$. As this extra information can only help the processors to make the broadcast faster we are allowed to make this assumption.

Let t_i be the random variable indicating the number of rounds since the processors of layer i get the message (and become active) until their success (the first time that a single processor in layer i transmits). We need to show that for some choice of $\ell_1, ..., \ell_D$ we get $E_\Pi(\sum_{i=1}^{D} t_i) = \Omega(D \log(N/D))$, where the expectation is taken over the random choices of the algorithm Π. Certainly, it is enough to show that $E_{\ell_1,...,\ell_D,\Pi}(\sum_{i=1}^{D} t_i) = \Omega(D \log(N/D))$. By linearity of expectation we get

$$E_{\ell_1,...,\ell_D,\Pi}(\sum_{i=1}^{D} t_i) = \sum_{i=1}^{D} E_{\ell_1,...,\ell_D,\Pi}(t_i).$$

So all we have to bound now is $E_{\ell_1,...,\ell_D,\Pi}(t_i)$. Clearly, the choice of $\ell_{i+1}, ..., \ell_D$ has no influence on the expectation of t_i, i.e.

$$E_{\ell_1,...,\ell_D,\Pi}(t_i) = E_{\ell_1,...,\ell_i,\Pi}(t_i).$$

Also, by the discussion above, with every history (which depends on the random choices made in the first $i-1$ layers, including the choice of $\ell_1, ..., \ell_{i-1}$) we can associate a probability distribution μ used by the processors in layer i to choose their schedules (note that since we assume that the processors of layer i get all the future information with the first message, they can make all their random choices at this time). Therefore, we can write

$$\begin{aligned} & E_{\ell_1,...,\ell_i,\Pi}(t_i) \\ &= \sum_{b_1,\cdots,b_{i-1}} E_{\ell_i,\Pi}(t_i | \ell_1 = b_1, \cdots, \ell_{i-1} = b_{i-1}) \\ &\quad \cdot Pr[\ell_1 = b_1, \cdots, \ell_{i-1} = b_{i-1}] \end{aligned}$$

$$= \sum_{b_1,\cdots,b_{i-1}} E_{\ell_i,\Pi}(t_i|\ell_1 = b_1,\cdots,\ell_{i-1} = b_{i-1})$$
$$\cdot \prod_{j=1}^{i-1} Pr[\ell_j = b_j]$$
$$= \frac{1}{(\log n)^{i-1}}$$
$$\cdot \sum_{b_1,\cdots,b_{i-1}} E_{\ell_i,\Pi}(t_i|\ell_1 = b_1,\cdots,\ell_{i-1} = b_{i-1})$$

It remains to bound the expression $E_{\ell_i,\Pi}(t_i|\ell_1 = b_1,\cdots,\ell_{i-1} = b_{i-1})$. As mentioned, we allow the processors at layer i to have access to $b_1,\cdots,b_{i-1}$ (the actual values of $\ell_1,\ldots,\ell_{i-1}$). Therefore we need to evaluate $E_{\ell_i,\Pi_i}(t_i)$, where Π_i is the protocol at layer i with the additional information about the lower layers. By Lemma 1, for each such Π_i,

$$E_{\ell_i,\Pi_i}(t_i) \geq c\log n,$$

for some constant c. We get,

$$E_{\ell_1,\ldots,\ell_i,\Pi}(t_i) \geq \frac{1}{(\log n)^{i-1}} \sum_{b_1,\cdots,b_{i-1}} c\log n$$
$$= c\log n$$

This implies

$$E_{\ell_1,\ldots,\ell_D,\Pi}(\sum_{i=1}^{D} t_i) = \Omega(D\log n) = \Omega(D\log(N/D)),$$

which completes the proof of our main theorem:

Theorem 7 *For any non-uniform broadcast protocol, there exists a network in which the expected time to complete a broadcast is $\Omega(D\log(N/D))$, where N is the number of processors, and D is the diameter.*

In case that $D \leq N^{1-\varepsilon}$, the above proof shows a lower bound of $\Omega(D\log N)$. Combining our result with the results of Alon et al. [ABLP91] and Bar-Yehuda et al. [BGI92] we have the following result.

Corollary 8 *For any non-uniform broadcast protocol, there exists a network in which the time to complete a broadcast is $\Omega(\log^2 N + D\log(N/D))$, where N is the number of processors in the network and D is the diameter. Furthermore, there is a (uniform) protocol that requires only $O(\log^2 N + D\log N)$ expected time (which is tight for all $D \leq N^{1-\varepsilon}$).*

Note that unlike [ABLP91] we show that for any protocol there exists a network for which the lower bound holds, while they prove that there exists a network on which any protocol requires the lower bound.

6 Acknowledgments

We wish to thank Oded Goldreich for his very useful comments.

References

[ABLP91] N. Alon, A. Bar-Noy, N. Linial, and D. Peleg. A lower bound for radio broadcast. *Journal of Computer and System Sciences*, 43:290–298, 1991.

[ABLP92] N. Alon, A. Bar-Noy, N. Linial, and D. Peleg. Single round simulation on radio networks. *Journal of Algorithms*, 13:188–210, 1992.

[BGI91] R. Bar-Yehuda, O. Goldreich, and A. Itai. Efficient emulation of single-hop radio network with collision detection on multi-hop radio network with no collision detection. *Distributed Computing*, 5:67–71, 1991.

[BGI92] R. Bar-Yehuda, O. Goldreich, and A. Itai. On the time-complexity of broadcast in multi-hop radio networks: An exponential gap between determinism and randomization. *Journal of Computer and System Sciences*, 45:104–126, 1992.

[BII89] R. Bar-Yehuda, A. Israeli, and A. Itai. Multiple communication in multi-hop radio networks. In *Proceedings of the 8th Annual ACM Symposium on Principles of Distributed Computing*, pages 329–338, 1989.

[CK85a] I. Chlamtac and S. Kutten. On broadcasting in radio networks – problem analysis and protocol design. *IEEE Trans. Comm.*, COM-33, 1985.

[CK85b] I. Chlamtac and S. Kutten. A spatial reuse tdma/fdma for mobile multi-hop radio networks. In *INFOCOM*, pages 389–394, 1985.

[CK87] I. Chlamtac and S. Kutten. Tree-based broadcasting in multihop radio networks. *IEEECOM*, C-36:1209–1223, 1987.

[CW87] I. Chlamtac and O. Weinstein. The wave expansion approach to broadcasting in multihop radio networks. In *INFOCOM*, pages 874–881, 1987.

[Gal85] R. Gallager. A perspective on multiaccess channels. *IEEE Transactions on Information Theory*, 31:124–142, 1985.

[Ten81] A. S. Tenenbaum. *Computer Networks*. Prentice-Hall, Englewood Cliffs, NJ, 1981.

[Wil86] D. E. Willard. Log-logarithmic selection resolution protocols in a multiple access channel. *SIAM Journal on Computing*, 15(2):468–477, May 1986.