February 1988 Report No. STAN-CS-88-1195

A Lower Bound for Radio Broadcast

by

A. Bar-Soy, N. Linial, and **D. Peleg**

Department of Computer Science

Stanford **University**
Stanford, California 94302

A Lower Bound for Radio Broadcast

Amotz Bar-Noy * Nathan Linial † David Peleg ‡

February 16, 1988

Abstract

A ***radio network*** is a synchronous network of processors that communicate by transmitting messages to their neighbors, where a processor receives a message in a given step if and only if it is silent in this step and precisely one of its neighbors transmits. In this paper we prove the existence of a family of radius-2 networks on n vertices for which any broadcast schedule requires at least $\Omega((\log n/\log\log n)^2)$ rounds of transmissions. This almost matches an upper bound of $O(\log^2 n)$ rounds for networks of radius 2 proved earlier by Bar-Yehuda, Goldreich and Itai [BGI].

Keywords: radio networks, radio broadcast, lower bound.

*Stanford University. Supported in part by a Weizmann fellowship and by contract ONR N00014-85-C-0731.

† IBM Almaden Research Center and the Hebrew University of Jerusalem.

*Stanford University. Supported in part by contract ONR N00014-85-C-0731.

1 Introduction

Packet radio networks have received considerable attention during the last decade [BGI, CK, GVF, K,K GBK, SC]. A *radio network* is an undirected (multi-hop) network of processors that communicate among themselves in synchronous time-slots in the following manner. In each step a processor can either transmit or keep silent. A processor receives a message in a given step if and only if it keeps silent and precisely one of its neighbors transmits in this step. If none of its neighbors transmits, it hears nothing. If more than one neighbor (including itself) has transmitted, a collision occurs and the processor hears only noise.

In this paper we consider the *broadcast* operation in radio networks [BGI, CK, CW, GVF]. Broadcast is a process by which a message M, initiated by a processor s (the sender) is delivered to all other processors in the network.

A *schedule* S is a list $(T_1, \ldots, T_t)$ of *transmissions*, specifying for each step i, $i = 1, 2, \ldots, t$, the set T_i of processors that have to transmit in step i. The schedule is applied as a broadcast procedure as follows. In step i, every processor $v \in T_i$ which already holds a copy of M transmits it (a processor $v \in T_i$ that does not have a copy yet, remains silent). The schedule S is a *broadcast schedule* for the sender s in G if after applying S, every processor in the network has a copy of M.

To formalize this notion, let us define for a vertex x the set $\Gamma(x)$ of its neighbors in G and let $deg(x) = |\Gamma(x)|$. For $T \subseteq V$ we let

$$H(T) = \{x \in V \setminus T : |\Gamma(x) \cap T| = 1\}.$$

Let $S = (T_1, \ldots, T_t)$ be a sequence of subsets of V with $T_1 = \{s\}$. Define $U_1 = \{Y\}$ and for $j \geq 1$

$$U_{j+1} = U_j \cup H(U_j \cap T_j).$$

Informally U_j is the set of vertices to whom the message M was relayed before round j. The processors transmitting at round j are therefore those in $U_j \cap T_j$ and $H(U_j \cap T_j)$ is the set of those processors receiving the message on round j. We say that S is a broadcast schedule for (G, s) if $U_t = V$.

We are interested in the existence of *short* broadcast schedules. Clearly, the radius of a network G from s (i.e., the largest distance between s and any other vertex in G) serves as a lower bound for the length of any broadcast schedule. Also, examples have been shown of radius-2 graphs of order n where every broadcast schedule requires $\Omega(\log n)$ rounds [BGI2]. In this paper we demonstrate the existence of a family of radius-2 networks on n vertices for which the number of rounds required by any broadcast schedule is

$$\Omega\left(\left(\frac{\log n}{\log\log n}\right)^2\right).$$

Note that the problem of *finding* efficient broadcast schedules depends on whether the graph G is known to the designer of the procedure. Our lower bound applies even at the harder case where G is known. For radius-2 graphs this is close to optimal in that the methods of [BGI] yield a broadcast schedule of $O(\log^2 n)$ rounds for them. The algorithm of [BGI] is probabilistic and does not assume knowledge of the graph. Thus, except for the log log n terms the upper and lower bounds match in a satisfactory way. Namely, for the family of graphs we construct there is a lower bound for the length of any schedule, even if G is known. On the other hand there is a probabilistic algorithm of almost the same complexity which requires no knowledge of the graph at hand. We think that $\Theta(\log^2 n)$ is the correct answer and the (log log $n)^2$ term remains only because of some weakness in our proof technique.

The situation is however less satisfactory when the radius grows. For general graphs of diameter D the probabilistic algorithm of [BGI] yields a schedule of 0($D \log n + \log^2 n$) rounds. The deterministic centerlized (polynomial time) algorithm of [CW] constructs a schedule of length $O(D \log^2 n)$. We cannot rule out the possibility that an 0($D + \log^2 n$) schedule always exists. The main difficulty is in understanding how efficient one may get in pipelining the message passing: Letting V_i be all the vertices at distance i from the sender s, the network may be engaged in passing A4 from V_i to V_{i+1} at the same time that it deals with V_j and V_{j+1} for some $j > i + 1$. How efficiently this may be done we do not know. It is an intriguing question to try and understand if this is indeed possible.

2 Preliminaries

We prove the lower bound in the case where all the processors are at distance at most two from the sender. The networks we consider are of the following form. We start with a bipartite graph $G(U, V, E)$ where $U = \{u_1, \ldots, u_r\}$, $V = \{v_1, \ldots, v_w\}$ and $n = r + w$. Further, we add a sender vertex s, adjacent to all the nodes of U. After the first round $(T_1 = \{s\})$, the message arrives at all neighbors of the sender in U. The rest of the schedule therefore needs only ensure the arrival of the message at all the processors of V. Consequently in a schedule S only the *set* of transmissions is relevant while the order in which they are performed may be ignored. Note also that each set T_i of those processors who transmit at round $i > 1$ is, without loss of generality, a subset of U.

We represent each bipartite graph $G(U, V, E)$ by its adjacency matrix A,,, where $A_{i,j} = 1$ if $(u_i, v_j) \in E$ and zero otherwise. For a given matrix A, a

transmission T is a set of rows $T \subseteq \{1, \ldots, n\}$. The size of a transmission T is denoted $|T|$. A *schedule* S is a set of transmissions. We say that a transmission T *hits* a column j in A if there is exactly one row $i \in T$ s.t. $A(i,j) = 1$. The transmission T misses column j if $A(i, j) = 0$ for all $i \in T$. We denote by $H(T, A)$ (respectively, $M(T, A)$) the set of columns hit (respectively, missed) by T. For a schedule S, $H(S, A) = \bigcup_{T \in S} H(T, A)$ and $M(S, A) = \bigcap_{T \in S} M(T, A)$. The schedule S *exhausts* A if all the columns of A are in $H(S, A)$. Note that for networks in the particular form considered here, our notation of $H(T)$ is consistent with the one from the previous section.

In the proofs of the lemmas we use the following propositions. The first states the Chernoff estimates [C] for the tails of the binomial distribution.

Proposition 2.1 (Chernoff) *Let X be a random variable with a binomial distribution and expectation $E = E(X)$. Then for $\alpha > 0$,*

1. $Pr(X > (1+\alpha)E) < e^{-\frac{1}{2}\alpha^2 E}$.

2. $Pr(X < (1-\alpha)E) < e^{-\frac{1}{3}\alpha^2 E}$.

A slight modification of Proposition 2.1 yields

Proposition 2.2 *For X, E and α as in Proposition 2.1,*

1. *If* $F > E$ *then* $Pr(X > (1+\alpha)F) < e^{-\frac{1}{2}\alpha^2 F}$.

2. *If* $F < E$ *then* $Pr(X < (1-\alpha)F) < e^{-\frac{1}{3}\alpha^2 F}$.

(Note that only Part 1 needs verifying while Part 2 is immediate.)

We recall two more easy facts, which we will use freely.

Proposition 2.3 *For* $0 < \alpha \leq \frac{1}{2}$ *and* $\beta > 0$,

$$e^{-3\alpha\beta} < (1 - \alpha)^{\beta} < e^{-\alpha\beta}.$$

Proposition 2.4 *The function* xe^{-x} *increases when* $0 < x \leq 1$ *and decreases when* $x > 1$.

Unless specified otherwise, all logarithms are to base 2. For simplicity we omit all floor and ceiling roundings throughout. Also, we assume without further notice that all our parameters are sufficiently large whenever needed.

3 The Bound

Let us first motivate our lower bound technique. Consider a radio network of the form described in the previous section. Suppose the network is based on a bipartite graph $G(U, V, E)$ in which each vertex $v \in V$ has degree d. In our matrix notation this corresponds to having exactly d 1's in every column. For such a network, a simple probabilistic argument shows the existence of a broadcast schedule of length $O(\log n)$. This is proved by considering a collection of $O(\log n)$ random transmissions, in which each row (U vertex) is chosen with probability $\frac{1}{d}$. For the general case, decompose V into $\log |U|$ sets

$$V_i = \{v \in V : 2^{i-1} \leq deg(v) < 2^i\} \qquad (1 \leq i \leq \log |U|)$$

and let G_i be the subgraph of G induced by $V_i \cup U$. Each G_i has an $O(\log n)$ schedule and the concatenation of these schedules is an $O(\log^2 n)$ schedule for all of G. Our lower bound shows that for suitably chosen G's this procedure is essentially the best possible. Each V_i is set to have size polynomial in $|U|$ and the edges at G_i are drawn at random from an appropriate distribution.

In simple terms our proof may be described as making the following claims: In dealing with G_i a transmission is almost useless unless its size is close enough to $\frac{|U|}{2^i}$ (Lemmas 3.1 and 3.2). Consequently we may think of each V_i as being handled separately, with almost no help from transmissions destined at other V_j's. For each individual V_i we show that if the edges at G_i are chosen at random from a proper distribution then the sets hit by various transmissions cannot combine in a way that is more efficient than in the situation where they are chosen randomly and independently (Lemmas 3.3 and 3.4). This is because a transmission of the right size not only *hits* a fraction of the target vertices, but also misses a fraction of them. These two arguments are put together by a straightforward pigeonhole argument to yield the theorem.

Let us now transform the above informal description into a more precise framework. Recall that our bipartite graphs are represented by an adjacency matrix $A_{r\times w}$. Throughout, we will have m = r^c for some constant c and w = *m log r*.

An *i-block* is a random *r x* m O-l matrix whose entries are chosen to be 1 with probability $p_i = 2^i/r$. We should, in fact speak about i-blocks of dimensions r *x m*, but the parameters r and *m* will always be clear from the context and we simplify our notations by dropping these indices. Denote by $\mathcal{A}_{i,m}$ the probability space of all i-blocks (again r will be fixed and so we omit this index). There is an induced probability distribution on $\mathcal{A}$, the set of all r x w O-l matrices: Independently choose a O-block, a l-block,..., a (logr $-$ l)-block and concatenate them to obtain a member in d.

For a given transmission *T* with *k* rows, its range, *R(T),* is the interval of integers i such that:

$$\log \frac{r}{k} - a \le i \le \log \frac{r}{k} + b$$

where

$$a = 2\log\log r + 2,$$
$$b = \log\log\log r + 2.$$

We are now ready to prove the lower bound itself.

Lemma *3.1 Given a transmission T and an integer i not in the range of T,*

$$Pr\left(|H(T,A)| > \frac{m}{\log^2 r}\right) < e^{-\frac{m}{4\log^2 r}} = o(1),$$

where $A \in \mathcal{A}_{i,m}$, the space of i-blocks.

Proof: Fix T and i and let $|T| = k$. Let $X_{T,i,j}$ denote the event in $\mathcal{A}_{i,m}$ "T hits the j-th column in an i-block". Recall that in an i-block the entries are chosen to be 1 with probability $p = 2^i/r$. *We* thus have

$$\Pr(X_{T,i,j}) = kp(1-p)^{k-1} \le \frac{kpe^{-kp}}{1-p} \le 2kpe^{-kp} .$$

Therefore the expected number of columns that T hits over $A \in \mathcal{A}_{i,m}$ is

(*) $$E(|H(T,A)|) \le 2kpme^{-kp} .$$

There are two cases for which $i \notin R(T)$. If

$$i > \log \frac{r}{k} + b = \log \frac{r}{k} + \log\log\log r + 2,$$

then $kp = \frac{k2^i}{r} > 4\log\log r$, so by (*) and by Proposition *2.4,*

$$E(|H(T,A)|) \le 8m \log\log r e^{-4\log\log r} < \frac{m}{2\log^2 r} .$$

In the other case

$$i < \log \frac{r}{k} - a = \log \frac{r}{k} - 2 \log \log r - 2,$$

so $kp < \frac{1}{4 \log^2 r} < 1$, which by $(*)$and by Proposition 2.4 implies that

$$E(|H(T,A)|) \leq 2kpm \leq \frac{m}{2 \log^2 r} \ .$$

Using Proposition 2.2 (Part 1) with $\alpha = 1$ and $F = \frac{m}{2 \log^2 r}$ *we* get

$$Pr\left(|H(T,A)| > \frac{m}{\log^2 r}\right) \leq e^{-\frac{m}{4 \log^2 r}}.$$

∎

Lemma 3.2 *Let* $0 \leq i \leq \log r - 1$, *let* $m \geq r^4$ *and let S be a schedule consisting of* $t = o(\log^2 r)$ *transmissions. Assume that i is not in the range of any transmission in S. Then*

$$Pr\left(|H(S, A)| \geq \frac{m}{2}\right) < \delta_m = (\log^2 r) e^{-\frac{m}{4 \log^2 r}} = o(1),$$

where $A \in \mathcal{A}_{i,m}$.

Proof: By Lemma 3.1 the probability that any transmission in S hits more than $\frac{m}{\log^2 r}$ columns is at most $te^{-\frac{m}{4 \log^2 r}}$. Therefore with probability $1 - te^{-\frac{m}{4 \log^2 r}} \leq 1 - \delta_m$ the whole schedule hits at most $\frac{mt}{\log^2 r} < \frac{m}{2}$ columns. ∎

Now we show that a transmission of the proper size not only hits but also misses a fraction of the vertices in V_i.

Lemma 3.3 *Given a transmission T and an integer i in the range of T,*

$$Pr\left(|M(T,A)| > \frac{m}{2 \log^{18} r}\right) > 1 - \varepsilon_m = 1 - e^{-\frac{m}{12 \log^{18} r}} = 1 - o(1),$$

where $A \in \mathcal{A}_{i,m}$.

Proof: Assume $|T| = k$. Let $Y_{T,i,j}$ denote the event in $\mathcal{A}_{i,m}$ "T misses the j-th column of an i-block".

$$\Pr(Y_{T,i,j}) = (1-p)^k \geq e^{-3kp}.$$

By the assumptions of the lemma, $i \in R(T)$, hence

$$i \leq \log \frac{r}{k} + \log\log\log r + 2,$$

so

$$kp = \frac{k2^i}{r} \leq 4\log\log r.$$

Therefore the expected number of columns that T misses is

$$E(|M(T,A)|) \geq me^{-3kp} \geq \frac{m}{\log^{18} r}.$$

Using Proposition 2.2 (Part 2) with $\alpha = \frac{1}{2}$ and $F = \frac{m}{\log^{18} r}$ *we* get

$$Pr\left(|M(T,A)| < \frac{m}{2\log^{18} r}\right) \leq e^{-\frac{m}{12\log^{18} r}} = \varepsilon_m.$$

∎

Define $\varepsilon = e^{-r}$.

Lemma 3.4 *Define* $\eta = 2\log^{18} r$ and *let* $m \geq r^3$. *Let S be a schedule consisting of* $t = O\left(\log_\eta \frac{m}{r^3}\right)$ *transmissions, and let i be an integer in the range of all the transmissions in S. Then*

$$Pr\left(|M(S,A)| \geq \frac{m}{\eta^t}\right) \geq \gamma_m = (1-\varepsilon)^{\log_\eta m},$$

where $A \in \mathcal{A}_{i,m}$.

Proof: By induction on m. For $m = r^3$ there are no transmissions so all the columns are missed. For simplicity assume that $m = \eta^j r^3$ for $j \geq 1$. Apply one of the transmissions T in S. Lemma 3.3 implies that with probability $1 - \varepsilon_m$ there are at least $\frac{m}{\eta} \geq r^3$ missed columns. This probability is at least $1 - \varepsilon$, since ε_m is a decreasing function of m and $\varepsilon_{r^3} \leq e^{-r} = \varepsilon$. Using the induction hypothesis we get that for the remaining $t - 1$ transmissions the schedule S \ {T} misses at least $\frac{m/\eta}{\eta^{t-1}} = \frac{m}{\eta^t}$ columns with probability greater than $\gamma_{m/\eta}$.

Combining the probabilities we get that the probability that at least $\frac{m}{\eta^t}$ columns are missed is

$$\gamma_{m/\eta}(1-\varepsilon) = (1-\varepsilon)^{\log_\eta \frac{m}{\eta}}(1-\varepsilon) = (1-\varepsilon)^{\log_\eta m} = \gamma_m$$

∎

Corollary 3.1 *For $m \geq r^4$ and $t = o\left(\frac{\log r}{\log\log r}\right)$ (and hence also $t = o(\log_\eta r)$ and $t = o\left(\log_\eta \frac{m}{r^3}\right)$)*

$$Pr\left(M(S,A)| \geq \frac{m}{\eta^{\log_\eta n}} = \frac{m}{r} \geq 1\right) \geq \gamma_m,$$

where $A \in \mathcal{A}_{i,m}$.

Lemma 3.5 *Let $m \geq r^5$ and let $0 \leq i \leq \log r - 1$. Then there exists a block B_i such that for every schedule $S = S' \bigcup S''$ where S' is a schedule consisting of $o(\frac{\log r}{\log\log r})$ transmissions T such that $i \in R(T)$ and S'' is a schedule consisting of $o(\log^2 r)$ transmissions T such that $i \notin R(T)$, S does not exhaust the block B_i.*

Proof: Consider an arbitrary schedule $S = S' \bigcup S''$ as in the lemma. From Corollary 3.1 it follows that with probability $\gamma_m > 0$, S' misses at least $m' = m/\imath$ columns of an i-block. By Lemma 3.2, with probability $1 - \delta_{m'} \geq 1 - e^{-r^2}$, S'' hits less than $m'/2$ columns out of any m' columns in an i-block.

Combining these two results together we get that with probability at most $1-\gamma_m(1-\delta_m) < \delta_m \leq e^{-r^2}$, S exhausts an i-block. There are at most 2' distinct transmissions, hence the number of possible schedules is at most $2^{r\log^2 r}$. Summing the probabilities we get that $2^{r\log^2 r}e^{-r^2} < 1$. Hence there exists a block B_i as desired. ∎

Lemma 3.6 *For every $m \geq r^5$ there exists a matrix $A_{r\times w}$ such that no schedule with $o\left(\left(\frac{\log r}{\log\log r}\right)^2\right)$ transmissions exhausts A.*

Proof: Construct A as follows. For every $0 \leq i \leq \log r - 1$, choose the i'th block in A to be the block B_i whose existence is asserted by Lemma 3.5. Assume that S is a schedule that exhausts A and S consists of $o\left(\left(\frac{\log r}{\log\log r}\right)^2\right)$ transmissions.

The range of any transmission covers at most O(loglog r) blocks. Therefore, by a simple pigeonhole argument, there exists an i-block where i is covered by the range of at most $o(\frac{\log r}{\log\log r})$ transmissions in S. Obviously there are at most $o(\log^2 r)$ transmissions whose range does not cover i. Applying Lemma 3.5 we get a contradiction. ∎

Theorem 3.1 *There exists an $O(n^4)$-vertex radius-2 network for which the number of rounds required by any broadcast schedule is at least*

$$\Omega\left(\left(\frac{\log n}{\log\log n}\right)^2\right).$$

∎

References

[BGI] R. Bar-Yehuda, 0. Goldreich and A. Itai, On the Time-Complexity of Broadcast in Radio Networks: An Exponential Gap between Determinism and Randomization, *Proc.* 4th ACM Symp. on Principles of Distributed Computing, 1986, pp. 98–107.

[BGI2] R. Bar-Yehuda, 0. Goldreich and A. Itai, Private Communication.

[C] H. Chernoff, A Measure of Asymptotic Efficiency for Tests of Hypothesis Based on the Sum of Observations, Ann. Math. Stat. 23, (1952), 493-507.

[CK] I. Chlamtac and S. Kutten, On Broadcasting in Radio Networks – Problem Analysis and Protocol Design, IEEE Trans. on Commun. 33, (1985).

[CW] I. Chlamtac and 0. Weinstein, The Wave Expansion Approch to Broadcasting in Multihop Radio Networks, *Proc. INFOCOM*, 1987.

[GVF] I. Gitman, R. M. Van Slyke and H. Frank, Routing in Packet-Switching Broadcast Radio Networks, *IEEE* Trans. on Commun., (1976), 926-930.

[K] R. C. Kunzelman, Overview of the ARPA Packet Radio Experimental Network, *Proc. COMPCON*, 1978, pp. 157-160.

[KGBK] R. E. Kahn, S. A. Gronemeyer, J. Burchfiel and R. C. Kunzelman, Advances in Packet Radio Technology, *Proc. IEEE 66,* (1978).

[SC] N. Shacham and E. J. Craighill, Dynamic Routing for Real-Time Transport in Packet Radio Networks, *Proc. INFOCOM*, 1982.