JOURNAL OF COMPUTER AND SYSTEM SCIENCES **43**, 290–298 (1991)

A Lower Bound for Radio Broadcast

NOGA ALON*

Tel-Aviv University, Ramat-Aviv, Tel-Aviv 69978, Israel; and Bell Communications Research, Morristown, New Jersey 07060

AMOTZ BAR-NOY†

Stanford University, Stanford, California 94305

NATHAN LINIAL

IBM Almaden Research Center, San Jose, California 95120; and the Hebrew University, Jerusalem 91904, Israel

AND

DAVID PELEG‡

Stanford University, Stanford, California 94305

Received March 2, 1988; revised October 30, 1988

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 $\Omega(\log^2 n)$ rounds of transmissions. This matches an upper bound of $O(\log^2 n)$ rounds for networks of radius 2 proved earlier by Bar-Yehuda, Goldreich, and Itai, *in* "Proceedings of the 4th ACM Symposium on Principles of Distributed Computing, 1986," pp. 98–107. © 1991 Academic Press, Inc.

1. INTRODUCTION

Packet radio networks have received considerable attention during the last decade [BGI, CK, GVF, K, KGBK, SC]. A *radio network* is an undirected (multi-hop) network of processors that communicate in synchronous time-slots in the

* Supported in part by an Allon Fellowship, by a Bat-Sheva de Rothschild Grant, and by a Bergmann Memorial Grant.

† Supported in part by a Weizmann Fellowship and by Contract ONR N00014-85-C-0731. Current address: IBM, T. J. Watson Research Center, Yorktown Heights, NY 10598.

‡ Supported in part by Contract ONR N00014-85-C-0731. Current address: The Weizmann Institute, Rehovot 76100, Israel.

0022-0000/91 $3.00

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) transmits, 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 lower bound is proved for the number of rounds that are required to broadcast in certain radio networks. The complexity of broadcasting may change significantly, depending on whether or not the processors know the network and how their actions are coordinated. We assume complete knowledge of the network and no restrictions on the coordination mechanism. These assumptions are the most advantageous for the algorithm (and so the most severe for the lower bound.)

Having made these assumptions, there is no loss of generality in assuming that broadcasting proceeds according to a *schedule* S, which is a list $(T_1, ..., T_t)$ of *transmissions*. Each transmission is a set of processors, and 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.

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

$$\Omega(\log^2 n).$$

As already mentioned, the lower bound applies even at the harder case where G is known. For radius-2 graphs the methods of [BGI] yield a broadcast schedule of $O(\log^2 n)$ rounds. The algorithm of [BGI] is probabilistic and does not assume knowledge of the graph. Thus, 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 the same order of time complexity which requires no knowledge of the graph at hand.

The situation gets less satisfactory, however, when the radius grows. For general graphs of diameter D the probabilistic algorithm of [BGI] yields a schedule of $O(\log^2 n + D \log n)$ rounds. The deterministic centerlized (polynomial time) algorithm of [CW] provides a schedule of length $O(D \log^2 n)$. We cannot rule out the possibility that an $O(D + \log^2 n)$ schedule always exists. This is a problem on efficient mechanisms for pipelining message passing: Let V_i be the set vertices at distance i from the sender s; the network may be engaged in passing M from V_i to V_{i+1} while dealing with V_j and V_{j+1} for some $j > i+1$. How efficiently this may be done we do not know, and the question is quite intriguing.

2. Preliminaries

Let $G = (V, E)$ be a radius-2 graph; denote by V_1 (respectively, V_2) the set of all the vertices at distance one (respectively, two) from the sender s. For convenience, we assume $V_1 = N = \{1, 2, ..., n\}$. After the first round of any broadcasting schedule in G all the processors in V_1 have the message M. Therefore, the remaining rounds of any schedule only need to guarantee the arrival of M to all processors in V_2. The graphs considered here have no edges between vertices in V_1, and so existence of a t-round schedule can be cast in combinatorial terms as follows.

Let $\mathscr{H}$ and $\mathscr{F}$ be families of nonempty subsets of N. (Any $H \in \mathscr{H}$ is the set of neighbors of some vertex in V_2. Members of $\mathscr{F}$ are transmissions in the schedule.) We say that $F \in \mathscr{F}$ *hits* $H \in \mathscr{H}$ if $|F \cap H| = 1$. (This means that the vertex in V_2 corresponding to H got the message on the transmission corresponding to F.) Also $\mathscr{F}$ *hits* H if some $F \in \mathscr{F}$ does, and $\mathscr{F}$ hits $\mathscr{H}$ if it hits every $H \in \mathscr{H}$. Let $t(\mathscr{H})$ be the minimum cardinality of $\mathscr{F}$ such that every $H \in \mathscr{H}$ is hit by some $F \in \mathscr{F}$. (The shortest broadcast schedule.) Define $t(n) = \max\{t(\mathscr{H})\}$ over all $\mathscr{H}$ of n subsets of N. The problem is to determine or estimate $t(n)$.

In the present note we determine $t(n)$ up to a constant factor.

THEOREM 2.1. *There are two positive constants c_1, c_2 such that*

$$c_1 \log^2 n \leqslant t(n) \leqslant c_2 \log^2 n \text{ for all } n \geqslant 2.$$

The upper bound was established in [BGI] and we show the lower bound

$$t(n) = \Omega(\log^2 n). \tag{1}$$

Theorem 2.1 implies the desired corollary

COROLLARY 2.1. *There is a family of order n graphs with radius 2 for which any schedule for the broadcast problem requires $\Omega(\log^2 n)$ rounds.*

The rest of the paper is devoted to the proof of (1).

3. The Lower Bound

3.1. *Outline of the Proof*

The proof is done by a probabilistic method. We show the existence of a family $\mathscr{H}$ of subsets of N which cannot be hit by any $\mathscr{F}$ of size $(\log^2 n)/100$.

The lower bound of $O(\log^2 n)$ changes only by a constant factor, as long as the cardinality of the family $\mathscr{H}$ is polynomial in n, and in fact, the constructed family, $\mathscr{H}$, is composed of $0.2 \log n$ subfamilies $\mathscr{H}_l$, each of cardinality n^7. For each l, $0.4 \log n \leqslant l \leqslant 0.6 \log n$, let $\mathscr{H}_l$ be a random family of n^7 (not necessarily distinct) subsets H of N chosen as follows: for each $i \in N$, independently, $\Pr(i \in H = 1/2^l)$. It is shown that for any fixed family $\mathscr{F}$ of at most $\log^2 n/100$ sets there is only a small probability for $\mathscr{F}$ to hit $\mathscr{H}$. The sum of these probabilities over all such $\mathscr{F}$ is less than 1 so there is an $\mathscr{H}$ which is hit by no $\mathscr{F}$.

As observed in [BGI2] every $\mathscr{H}_l$ may be hit by an $\mathscr{F}$ of size $O(\log n)$. The proof essentially shows that each $\mathscr{H}_l$ requires an $\mathscr{F}_l$ of size $\Omega(\log n)$ in order to be hit and $\mathscr{F}_l$ "does not help" in hitting $\mathscr{H}_j$ for $j \neq l$.

It is easy to check that for a set A of a elements in N and a random set B of b elements in N, the probability of A hitting B is $(1+o(1))xe^{-x}$ where $x = ab/n$. Now, the cardinalities of the sets in $\mathscr{H}_l$ are almost surely very close to $n/2^l$; so, for a fixed F and $H \in \mathscr{H}_l$ the following three cases may occur: If $|F| \ll 2^l$ then, with high probability, $F \cap H = \varnothing$. If $|F| \gg 2^l$, then with high probability, $|F \cap H| \geqslant 2$. If $|F|$ is close to 2^l, then with a constant probability F hits H. Consequently, a set $F \in \mathscr{F}$ which is of the "right" size for some $\mathscr{H}_l$ is either too small or too large for other $\mathscr{H}_j$.

If we associate each $F \in \mathscr{F}$ with the appropriate $\mathscr{H}_l$, then there is an $\mathscr{H}_l$ with less than $\log n/20$ associated F's (since $|\mathscr{F}| < \log^2 n/100$ and $\mathscr{H}$ consists of $0.2 \log n$ subfamilies $\mathscr{H}_l$). A simple argument can be made to yield some lower bound, but not quite the correct one, mainly because of two difficulties. First, the function xe^{-x} mentioned above does not decay sufficiently fast as we move from its maximum at $x = 1$. Consequently, an F of size 2^l may "help" also in hitting $\mathscr{H}_j$ for j "close" to l. Second, estimating the probability of $\mathscr{F}$ hitting $\mathscr{H}$ by summing over all $F \in \mathscr{F}$ and $H \in \mathscr{H}$ gives bounds which are too crude and we need some independence. These difficulties are overcome by a refinement of the above pigeonhole argument (Lemma 3.1).

Let $\mathscr{F}$ be a fixed family of $t \leqslant (\log^2 n/100)$ subsets of N; Lemma 3.1 shows that there are an index l $(0.4 \log n \leqslant l \leqslant 0.6 \log n)$ and a subfamily $\mathscr{G}$ of $\mathscr{F}$ satisfying the following conditions:

1. $|\bigcup_{A \in \mathscr{G}} A|$ is "small."

2. For each $B \in \mathscr{F} \backslash \mathscr{G}$, $B' = B \backslash (\bigcup_{A \in \mathscr{G}} A)$ is "large."

This index l indicates which subfamily $\mathscr{H}_l$ is not hit by $\mathscr{F}$.

Let H be a member in $\mathscr{H}_l$. The claim is that the probability that $\mathscr{F}$ does not hit H is more than $1/n^5$. By condition (1) the union of all sets in $\mathscr{G}$ is sufficiently small; thus, the probability that $H \cap A = \varnothing$ for *all* $A \in \mathscr{G}$ is at least $1/n^2$ (Lemma 3.2). The probability that $|H \cap B'| \geqslant 2$ for a fixed $B \in \mathscr{F} \backslash \mathscr{G}$ is at least $1 - O(2^l/|B'|)$. The dependencies between the events $|B'_i \cap H| \geqslant 2$ and $|B'_j \cap H| \geqslant 2$ only help by the FKG inequality, and so the probability that $|H \cap B'| \geqslant 2$ for all $B \notin \mathscr{G}$ is greater than $\prod (1 - O(2^l/|B'|))$. Condition (2) implies that this is at least $1/n^3$

(Lemma 3.3). Since $\bigcup_{\mathscr{G}} A$ and the B^r are disjoint these two events are independent and the $1/n^5$ bound is proved.

The size of $\mathscr{H}_l(n^7)$ implies that the probability that for all $H \in \mathscr{H}_l$ there is an $F \in \mathscr{F}$ with $|F \cap H| = 1$ is at most e^{-n^2}. As there are less than $2^{n \log^2 n} \ll e^{n^2}$ such possible families $\mathscr{F}$, the lower bound follows.

3.2. *A Combinatorial Lemma*

Here is the lemma we have just discussed. We assume that n is large enough whenever needed, also the numerical constants certainly can be improved upon.

LEMMA 3.1. *Suppose $t \leqslant \log^2 n/100$ and let $\mathscr{F}$ be a family of t subsets of $N = \{1, 2, ..., n\}$. Then, there is an index l and a subfamily $\mathscr{G}$ of $\mathscr{F}$ such that the following four conditions hold*:

(i) $0.4 \log n \leqslant l \leqslant 0.6 \log n$.

(ii) $|\bigcup_{A \in \mathscr{G}} A| \leqslant 2^l \log n$.

(iii) *For each $B \in \mathscr{F} \backslash \mathscr{G}$ define $B^r = B \backslash (\bigcup_{A \in \mathscr{G}} A)$ then $|B^r| \geqslant 2^l$ for all $B \in \mathscr{F} \backslash \mathscr{G}$.*

(iv) *For each $k \geqslant 0$ let f_k denote the number of sets $B \in \mathscr{F} \backslash \mathscr{G}$ such that $2^{l+k} \leqslant |B^r| < 2^{l+k+1}$ then*

$$\sum_{k \geqslant 0} \frac{f_k}{2^k} \leqslant \log n. \tag{2}$$

Proof. Define a permutation $A_1, A_2, ..., A_t$ of the members of $\mathscr{F}$ as follows. Let A_1 be a set of minimum cardinality in F. Assuming $A_1, ..., A_i$ have already been chosen, $(1 \leqslant i \leqslant t)$, let A_{i+1} be a set in $\mathscr{F} \backslash \{A_1, ..., A_i\}$ such that $|A_{i+1} \backslash (\bigcup_{j=1}^{i} A_j)|$ is minimum. Define, also, $x_i = |A_i \backslash (\bigcup_{j<i} A_j)|$ for $1 \leqslant i \leqslant t$. For each l, $0.4 \log n \leqslant l \leqslant 0.6 \log n$, let $j = j(l)$ be the smallest j such that $x_j \geqslant 2^l$. (If there is no such j, put $j(l) = t+1$). Note that by the definition of the permutation $A_1, A_2, ..., A_t$, for every l and for every $j' \geqslant j(l)$

$$|A_{j'} \backslash \{A_1 \cup \cdots \cup A_{j(l)-1}\}| \geqslant 2^l. \tag{3}$$

For each l put $d_l = |\{i : 1 \leqslant i \leqslant t, 2^l \leqslant x_i < 2^{l+1}\}|$, and $d'_l = d_{l-1} + d_{l-2}/2 + d_{l-3}/4 + d_{l-4}/8 + \cdots$. Clearly

$$\sum \{d'_l : 0.4 \log n \leqslant l \leqslant 0.6 \log n\} \leqslant 2 \sum_{l \geqslant 0} d_l \leqslant 2t \leqslant \frac{\log^2 n}{50}. \tag{4}$$

Call an index l *good* if $0.4 \log n \leqslant l \leqslant 0.6 \log n$ and $d'_l \leqslant \log n$. By (4) the average value of d'_l over all $0.4 \log n \leqslant l \leqslant 0.6 \log n$ is at most $(\log n)/10$ and hence at least

90% of the indices l, $0.4\log n \leqslant l \leqslant 0.6\log n$, are good. Notice that if l is good and $j=j(l)$ then for $\mathcal{G}=\{A_1, ..., A_{j-1}\}$ we have

$$\left|\bigcup_{A\in\mathcal{G}} A\right| = x_1 + \cdots + x_{j-1} \leqslant \sum_{p<l} 2^{p+1} d_p$$
$$= 2^l\left(d_{l-1} + \frac{d_{l-2}}{2} + \cdots\right) = 2^l d_l' \leqslant 2^l \log n.$$

Hence $\mathcal{G}$ and l satisfy conditions (i) and (ii). Moreover, by (3), condition (iii) holds as well. To complete the proof we show that for at least one (and in fact for many) good l condition (iv) holds too. For each good l and each set A_k, with $k \geqslant j(l)$, define $s(l, k) = r$ if

$$2^{l+r} \leqslant \left|A_k \Big\backslash \bigcup_{i=1}^{j(l)-1} A_i\right| < 2^{l+r+1}.$$

Note that if $l' > l$ are both good then $j(l') \geqslant j(l)$ and hence if $k \geqslant j(l')$ then

$$\left|A_k \Big\backslash \bigcup_{i=1}^{j(l')-1} A_i\right| \leqslant \left|A_k \Big\backslash \bigcup_{i=1}^{j(l)-1} A_i\right|.$$

Consequently, in this case $s(l', k)$ is strictly smaller than $s(l, k)$. Therefore, for every fixed k,

$$\sum \left\{\frac{1}{2^{s(l,k)}} : l \text{ is good, } j(l) \leqslant k\right\} \leqslant 2. \tag{5}$$

For each good l define $y_l = \sum_{k \geqslant j(l)} (1/2^{s(l,k)})$. By (5)

$$\sum \{y_l : l \text{ is good}\} \leqslant \sum_{k=1}^{t} \left\{\frac{1}{2^{s(l,k)}} : l \text{ is good, } j(l) \leqslant k\right\} \leqslant 2t \leqslant \frac{\log^2 n}{50}.$$

Since there are at least $0.9 \cdot 0.2 \log n > (1/10) \log n$ good indices l, there is at least one such l with $y_l \leqslant \log n/5 < \log n$. Define $j = j(l)$ and $\mathcal{G} = \{A_1, A_2, ..., A_{j-1}\}$. Clearly these $\mathcal{G}$ and l satisfy conditions (i), (ii), and (iii). Moreover, if f_k denotes the number of sets $B \in \mathcal{F}\backslash\mathcal{G}$ such that $2^{l+k} \leqslant |B'| < 2^{l+k+1}$ then

$$\sum_{k \geqslant 0} \frac{f_k}{2^k} = y_l < \log n;$$

i.e., condition (iv) holds too. This completes the proof of the lemma. ∎

3.3. *The Proof*

Now we show our probabilistic construction. For each l, $0.4 \log n \leqslant l \leqslant 0.6 \log n$, let $\mathcal{H}_l = \{H_{l_1}, ..., H_{l_{n^7}}\}$ be a random family of n^7 (not necessarily distinct) subsets of

$N=\{1,2,\ldots,n\}$ chosen as follows, for each $i\in N$ and $1\leqslant j\leqslant n^7$, independently, $\Pr(i\in H_{l_j})=1/2^l$. Put

$$\mathscr{H}=\bigcup\{\mathscr{H}_l : 0.4\log n\leqslant l\leqslant 0.6\log n\}.$$

We show that with positive probability $t(\mathscr{H})>(\log^2 n/100)$. Since $\mathscr{H}$ has less than $n^7\log n$ sets all of which can be considered as subsets of an $[n^7\log n]$-element set this shows that $t(n^7\log n)\geqslant\Omega(\log^2 n)$ and hence $t(m)\geqslant\Omega(\log^2 m)$, completing the proof of Theorem 2.1. It thus remains to show that with positive probability

$$t(\mathscr{H})>\frac{\log^2 n}{100}.$$

Let $\mathscr{F}$ be a fixed family of $t\leqslant(\log^2 n/100)$ subsets of N. By Lemma 3.1 there are an index l and a subfamily $\mathscr{G}$ of $\mathscr{F}$ satisfying the conclusions (i)–(iv) of the lemma. Consider the subfamily $\mathscr{H}_l$ of $\mathscr{H}$ and let $H=H_{l_j}$ be one of the subsets in that subfamily. We claim that the probability that $|H\cap F|\neq 1$ for each $F\in\mathscr{F}$ is more than $1/n^5$ (for all sufficiently large n). To prove this claim we need the following two lemmas.

LEMMA 3.2. *The probability that $H\cap A=\emptyset$ for all $A\in\mathscr{G}$ is at least $1/n^2$.*

Proof. This probability is precisely

$$\left(1-\frac{1}{2^l}\right)^{|\bigcup_{A\in\mathscr{G}}A|}\geqslant\left(1-\frac{1}{2^l}\right)^{2^l\log n}=\frac{1}{n^{1+o(1)}}>\frac{1}{n^2}. \quad\blacksquare$$

LEMMA 3.3. *The probability that $|H\cap B'|\geqslant 2$ for all $B\in\mathscr{F}\backslash\mathscr{G}$ is at least $1/n^3$.*

Proof. We first note that by the well known FKG inequality (see, e.g., [Bo, Thm. 19.5]) the above probability is at least the product of the probabilities that $|H\cap B'|\geqslant 2$, as B ranges over all sets in $\mathscr{F}\backslash\mathscr{G}$. Fix a set B in $\mathscr{F}\backslash\mathscr{G}$ and let $k\geqslant 0$ be an integer so that $2^{l+k}\leqslant|B'|<2^{l+k+1}$. Put $y=|B'|$. Clearly if n is large enough:

$$\begin{aligned}\Pr(|H\cap B'|\geqslant 2)&=1-\left(1-\frac{1}{2^l}\right)^y-y\frac{1}{2^l}\left(1-\frac{1}{2^l}\right)^{y-1}\\&\geqslant 1-\left(1-\frac{1}{2^l}\right)^{2^{l+k}}-2^k\left(1-\frac{1}{2^l}\right)^{2^{l+k}-1}\\&\geqslant 1-e^{-2^k}-(1+o(1))2^ke^{-2^k}\\&=1-(1+o(1))\frac{2^k+1}{e^{2^k}}\geqslant 1-\frac{0.9}{2^k}.\end{aligned}$$

Consequently, by Lemma 3.1 (iv) and by the FKG inequality mentioned above we conclude that if f_k is the number of sets $B \in \mathscr{F} \backslash \mathscr{G}$ such that $2^{l+k} \leqslant |B^r| < 2^{l+k+1}$ then the probability that H contains at least two elements from each B^r is at least

$$\prod_{k \geqslant 0} \left(1 - \frac{0.9}{2^k}\right)^{f_k} > \prod_{k \geqslant 0} (e^{-3/2^k})^{f_k} \geqslant e^{-3 \log n} = \frac{1}{n^3}.$$

This completes the proof of the lemma. ∎

The event considered in Lemma 3.2 and that considered in Lemma 3.3 are clearly independent (since $\bigcup_{\mathscr{G}} A$ and $\bigcup_{B \in \mathscr{F} \backslash \mathscr{G}} B^r$ are disjoint) and hence

$$\begin{aligned}
&\Pr(|H \cap F| \neq 1 \text{ for all } F \in \mathscr{F}) \\
&\quad \geqslant \Pr(|H \cap A| = 0 \text{ for all } A \in \mathscr{G} \text{ and } |H \cap B^r| \geqslant 2 \text{ for all } B \in \mathscr{F} \backslash \mathscr{G}) \\
&\quad = \Pr(|H \cap A| = 0 \text{ for all } A \in \mathscr{G}) \cdot \Pr(|H \cap B^r| \geqslant 2 \text{ for all } B \in \mathscr{F} \backslash \mathscr{G}) \\
&\quad \geqslant \frac{1}{n^2} \cdot \frac{1}{n^3} = \frac{1}{n^5}.
\end{aligned}$$

We have thus proved that for each fixed family of at most $(\log^2 n)/100$ sets $\mathscr{F}$ there is some index l such that for each of the random sets H in $\mathscr{H}_l$, the probability that $|H \cap F| = 1$ for some $F \in \mathscr{F}$ is at most $1 - 1/n^5$. As the members of $\mathscr{H}_l$ are independent this implies that the probability that for all $H \in \mathscr{H}_l$ there is an $F \in \mathscr{F}$ with $|F \cap H| = 1$ is at most

$$\left(1 - \frac{1}{n^5}\right)^{n^7} \leqslant e^{-n^2}.$$

Therefore, for each fixed family $\mathscr{F}$ of at most $(\log^2 n)/100$ subsets of N, the probability that for all $H \in \mathscr{H}$ there is an $F \in \mathscr{F}$ with $|F \cap H| = 1$ is at most e^{-n^2}. As there are less than $2^{n \log^2 n} \ll e^{n^2}$ such possible families $\mathscr{F}$ this implies that for most families $\mathscr{H}$ constructed as above

$$t(\mathscr{H}) \geqslant \frac{\log^2 n}{100}.$$

(It is also easy to check that most of these families contains no empty sets). This completes the proof of Theorem 2.1.

REFERENCES

[BGI] R. BAR-YEHUDA, O. GOLDREICH, AND A. ITAI, On the time-complexity of broadcast in radio networks: An exponential gap between determinism and randomization, *in* "Proceedings of the 4th ACM Symposium on Principles of Distributed Computing, 1986," pp. 98–107.

[BGI2] R. BAR-YEHUDA, O. GOLDREICH, AND A. ITAI, private communication.

[Bo] B. BOLLOBÁS, "Combinatorics," Cambridge University Press, Cambridge, 1986.

[C] H. CHERNOFF, A measure of asymptotic efficiency for tests of hypothesis based on the sum of observations, *Ann. Math. Statist.* **23** (1952), 493–507.

[CK] I. CHLAMTAC AND S. KUTTEN, On broadcasting in radio networks—Problem analysis and protocol design, *IEEE Trans. Comm.* **33** (1985).

[CW] I. CHLAMTAC AND O. WEINSTEIN, The wave expansion approach 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. Comm.* (1976), 926–930.

[K] R. C. KUNZELMAN, Overview of the ARPA packet radio experimental network, *Proc. COMPCON* (1978), 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).