Broadcasting Algorithms in Radio Networks with Unknown Topology *

Artur Czumaj
Department of Computer Science
New Jersey Institute of Technology
E-Mail: czumaj@cis.njit.edu

Wojciech Rytter
Department of Computer Science
New Jersey Institute of Technology *and*
Institute of Informatics, Warsaw University
E-Mail: rytter@cis.njit.edu

Abstract

In this paper we present new randomized and deterministic algorithms for the classical problem of broadcasting in radio networks with unknown topology. We consider directed n-node radio networks with specified eccentricity D (maximum distance from the source node to any other node). In a seminal work on randomized broadcasting, Bar-Yehuda et al. presented an algorithm that for any n-node radio network with eccentricity D completes the broadcasting in $\mathcal{O}(D \log n + \log^2 n)$ time, with high probability. This result is almost *optimal, since as it has been shown by Kushilevitz and Mansour and Alon et al., every randomized algorithm requires $\Omega(D \log(n/D) + \log^2 n)$ expected time to complete broadcasting.*

Our first main result closes the gap between the lower and upper bound: we describe an optimal randomized broadcasting algorithm whose running time complexity is $\mathcal{O}(D \log(n/D) + \log^2 n)$, with high probability. In particular, we obtain a randomized algorithm that completes broadcasting in any n-node radio network in time $\mathcal{O}(n)$, with high probability; the best previously existing algorithm achieved the running time $\mathcal{O}(n \log n)$.

The main source of our improvement is a better "selecting sequence" used by the algorithm that brings some stronger property and improves the broadcasting time. Two types of "selecting sequence" are considered: randomized and deterministic ones. The algorithm with a randomized sequence is easier (more intuitive) to analyze but both randomized and deterministic sequences give algorithms of the same asymptotic complexity.

Next, we demonstrate how to apply our approach to deterministic broadcasting, and describe a deterministic oblivious algorithm that completes broadcasting in almost optimal time $\mathcal{O}(n \log^2 D)$, which improves upon best known algorithms in this case. The fastest previously known algorithm had the broadcasting time of $\mathcal{O}(n \log n \log D)$, it was non-oblivious and it was significantly more complicated; our algorithm can be seen as a natural extension of our randomized algorithm.

Finally, we show how our randomized broadcasting algorithm can be used to improve the randomized complexity of the gossiping problem.

* Research supported in part by NSF grant CCR-0105701.

1. Introduction

In this paper we consider the fundamental problems in distributed computing of broadcasting in networks. In the *broadcasting problem*, one distinguished source node has a message that needs to be sent to all other nodes in the network. We consider *directed* networks with *unknown structure*. A *radio network* is modeled by a network (directed graph) $\mathcal{N} = (V, E)$, where the set of nodes correspond to the set of transmitter-receiver stations. The nodes of the network are assigned different identifiers and we assume that $V = [n]$ with 1 being distinguished as the *source* node (where we use the standard notation $[n] = \{1, \dots, n\}$). A directed edge $(v, u) \in E$ means that node v can send a message to node u.

We consider the standard model of unknown radio networks, called also sometimes the *ad-hoc network model* (for more elaborate discussion about our model, see, e.g., [1, 4, 5, 7, 8, 10, 12, 17, 19, 20]). We assume that a node does not have any a prior knowledge about the topology of the network, its in-degree and out-degree, and the set of its neighbors. We assume that the only knowledge of each node is its own identifier and the size of the network, n. (We notice that one can relax the assumption about the knowledge about n to be either a linear upper bound for the number of nodes, or even to be unknown for the nodes in the network, see, e.g., [17] for more details.) Additionally, in a part of the paper, we shall assume that each node knows the *eccentricity* of the network, which we denote by D, that is equal to the maximum distance from the source to any other node in

$\mathcal{N}$. (Again, this knowledge is not necessary for our analysis if n is known, and we use it for simplicity only.)

The nodes are sending messages according to a *selecting sequence*, at each round this sequence specifies the probability with which all active nodes will transmit the message in this round, or in the deterministic algorithm, specifies which active node will transmit. It is convenient to assume that all nodes know the selecting sequence in advance, if not, then the source can include it together with its original starting message. (This is especially important if we consider a *randomized* selecting sequence, where one should understand that initially the source chooses the sequence at random and then distributes it to all other nodes together with the original starting message.) Our algorithms are *"oblivious,"* all nodes perform in the same way, according to the same selecting sequence. The nodes do not use any extra local memory.

To make the broadcasting problem feasible, we assume that every node in $\mathcal{N}$ is reachable by a path from the source node. We assume that all nodes have access to a global clock and work synchronously in discrete time steps called *rounds*. A node is *active* if it has already received a message from the source. If a node received the message at round r, then we say it is active at the end of round r; if it received the message at round $r' < r$, then we say it is active at the beginning of round r. The source message can be sent only through the edges of network. In each round each node can transmit the message to all its out-neighbors at once and can receive the messages from its in-neighbors. A node will receive a message at a given round if and only if *exactly one of its in-neighbors transmits* at that round. If more than one in-neighbor transmits simultaneously in a given round, then a *collision* occurs and none of the messages is received by the node. In that case, we assume that the node cannot distinguish such a collision from the situation when none of its neighbors is transmitting. Furthermore, we do not consider the possibility of spontaneous transmissions and we allow only active nodes to transmit. We say an algorithm *completes broadcasting in* T *rounds* if at the end of round T all nodes received the source message, or in other words, if all nodes are active at the end of round T.

1.1. Previous results

There has been a vast amount of research on broadcasting in unknown radio network models (see, e.g., the survey [7], and the recent papers [16, 17] and the references therein). In some of the work cited below, the network is modeled as an undirected graph, which is equivalent to the assumption that the directed graph is symmetric (and in which case, D, the eccentricity, is of the order of the diameter of the network). Therefore, the model of directed networks is more general.

For the randomized version of the problem, it has been shown by Alon et al. [1] that there exists a network of eccentricity $\mathcal{O}(1)$ for which broadcasting needs $\Omega(\log^2 n)$ expected time. Kushilevitz and Mansour [19] showed that any randomized broadcasting algorithm requires $\Omega(D \cdot \log(n/D))$ time for n-node networks of eccentricity D (see also [20]). Bar-Yehuda et al. [4] designed an almost optimal broadcasting algorithm achieving the running time of $\mathcal{O}((D + \log n) \cdot \log n)$. By the lower bounds from [1, 19], this algorithm is optimal for all $D \le n^{1-\epsilon}$, but it is by a logarithmic factor off from optimal for D close to n. In general n-node networks, when the bound on D is unknown, the algorithm due to Bar-Yehuda et al. [4] requires $\mathcal{O}(n \log n)$ time and no asymptotically faster algorithm has been known before. If all nodes have full knowledge of the network, then the lower bound of $\Omega(\log^2 n)$ is still valid and one can do broadcasting in $\mathcal{O}(D + \log^5 n)$ expected time [13].

The problem of deterministic broadcasting has been also intensively studied, but only very recently almost optimal algorithms have been designed. [8] designed the first subquadratic running time of $\mathcal{O}(n^{11/6})$, [9] gave the running time of $\mathcal{O}(n^{1.5})$ and the same bound was obtained by Peleg using a probabilistic construction. Chrobak et al. [10] were the first who designed an almost optimal algorithm that completes the broadcasting in $\mathcal{O}(n \log^2 n)$ time; very recently, Kowalski and Pelc [17] improved this bound to obtain a non-oblivious algorithm of complexity $\mathcal{O}(n \log n \log D)$. All known $\mathcal{O}(n\, \text{poly-log}(n))$ algorithms (including those in [10, 17]) are probabilistic and non-constructive. The best constructive algorithm known up to date is by Indyk [15]. The best known lower bound is $\Omega(n \log D)$ due to Clementi et al. [12].

In this paper, we study also the problem of gossiping in unknown radio networks (for more details, see [7, 10, 11, 20, 21]). The fastest known randomized algorithms achieve the running time of $\mathcal{O}(n \log^4 n)$ [11] and $\mathcal{O}(n \log^3 n)$ [20]; the fastest deterministic algorithm has the running time of $\widetilde{\mathcal{O}}(n^{1.5})$ [10, 21].

1.2. New contributions

In this paper we present several new algorithms which improve upon the complexity of the best known algorithms for the problems listed below. Let $\mathcal{N} = (V, E)$ be a directed n-node unknown radio network.

- We design a randomized broadcasting algorithm which completes broadcasting in $\mathcal{N}$ in $\mathcal{O}(n)$ time with high probability. This is the *first optimal algorithm* for this problem. (Using standard doubling technique for estimating the value of n, the algorithm does not have to know n.) We show algorithms using randomized and deterministic *selecting sequences*.

- Let D be the eccentricity of $\mathcal{N}$. We extend our $\mathcal{O}(n)$-time algorithm and present a randomized broadcasting algorithm which completes broadcasting in $\mathcal{N}$ in $\mathcal{O}(D \log(n/D) + \log^2 n)$ time with high probability. Here we also show algorithms using randomized and deterministic selecting sequences. Our upper bound matches the lower bound given by Kushilevitz and Mansour [19] and Alon et al. [1] for all values of the eccentricity D.
- We describe a deterministic oblivious broadcasting algorithm which completes broadcasting in $\mathcal{O}(n \log^2 D)$ time. The main source of our improvement is the use of new selecting sequences, extending our randomized algorithm.
- We give a randomized Las Vegas algorithm that performs gossiping in $\mathcal{N}$ in expected $\mathcal{O}(n \log^2 n)$ time. This improves upon the best previously known bound [20] by a logarithmic factor.

2. Preliminaries

We assume, without of loss generality, that n is a power of 2; otherwise, one should use $\lceil \log n \rceil$ instead of $\log n$. Similarly, we assume that D is a power of 2; otherwise, one should use $\lceil \log D \rceil$ instead of $\log D$. Moreover, each time we write an expression of the type $\log(N/K)$, we mean $\log(1 + N/K)$; this is to avoid the case $N = K$, when we want "$\log(N/K)$" to be $\mathcal{O}(1)$ rather than 0.

2.1. Previous approach to randomized broadcasting

Bar-Yehuda et al. [4] presented a randomized broadcasting algorithm that runs in almost optimal time. The following Simple Randomized Broadcasting Algorithm(T) is essentially identical to the algorithm presented in [4], and as it is proven in [4], this algorithm completes the broadcasting in $\mathcal{O}((D+\log n) \log n)$ rounds with probability at least $1 - n^{-1}$, where D is the eccentricity of the input network.

> **Simple Randomized Broadcasting Algorithm(T)**
>
> **Input:** Network $\mathcal{N} = (V, E)$.
> $M_r = \log n - (r \bmod \log n)$ for every $r \in \mathbb{N}$
>
> **for** $r = 1$ **to** T **do** { *round number* r }
> **for** each *active* node $v \in V$ *independently* **do**
> node v transmits with probability 2^{-M_r}

The main idea of this algorithm is that, informally, by choosing the "selecting sequence" $\langle M_1, M_2, \ldots \rangle$ we ensure that for any r, in the next $\mathcal{O}(\log n)$ rounds we expect to make one new node active, and thus, $\mathcal{O}(n \log n)$ rounds suffice to complete broadcasting.

In this paper we improve this algorithm and obtain asymptotically optimal running time for randomized broadcasting of $\mathcal{O}(D \log(n/D) + \log^2 n)$. The main source of our improvement is a better "selecting sequence" used by the algorithm that brings some stronger property and improves the broadcasting time.

3. Randomized broadcasting in linear time

In our new randomized algorithm we replace the selecting sequence $\langle M_1, M_2, \ldots \rangle$ by two types of selecting sequences $\mathfrak{I} = \langle I_1, I_2, \ldots \rangle$: a randomized sequence (each element is generated independently with some probability α_k), and a deterministic sequence defined constructively. All analyzes in this paper could be performed using either of the sequences, but for simplicity of the presentation we focus our attention on the randomized sequence.

3.1. Selective sequences and broadcasting algorithms

Let us remind that n is a power of 2 and let $\mathcal{LL}(n) = \lceil \log \log n \rceil$. For any $k \in \{0, 1, \ldots, \log n\}$, define

$$\alpha_k = \begin{cases} 2^{-(k+1)} & \text{for } 1 \le k \le \mathcal{LL}(n) \ , \\ \frac{1}{2 \log n} & \text{for } \mathcal{LL}(n) \le k \le \log n \ , \\ 1 - \sum_{i=1}^{\log n} \alpha_i & \text{for } k = 0 \ . \end{cases}$$

In our first algorithm we will define the randomized sequence $\mathfrak{I} = \langle I_1, I_2, \ldots \rangle$ such that $\mathbf{Pr}[I_r = k] = \alpha_k$.

> **Linear Randomized Broadcasting Algorithm** (T)
>
> **Input:** Network $\mathcal{N} = (V, E)$.
> Randomized sequence $\mathfrak{I} = \langle I_1, I_2, \ldots \rangle$ such that
> $\mathbf{Pr}[I_r = k] = \alpha'_k \ \forall r \in \mathbb{N}, \forall k \in \{0, 1, 2, \ldots, \log n\}$
>
> **for** $r = 1$ **to** T **do** { *round number* r }
> **for** each *active* node $v \in V$ *independently* **do**
> node v transmits with probability 2^{-I_r}

Our first main result is the following theorem.

Theorem 1 *Let $\mathcal{N} = (V, E)$ be an n-node radio network. There is a constant c such that if $T \ge c\, n$, then the Linear Randomized Broadcasting Algorithm(T) completes broadcasting in $\mathcal{N}$ with probability at least $1 - n^{-1}$.*

Remark 3.1 *It is tempting to run the Linear Randomized Broadcasting Algorithm with a simpler selection sequence, in which each I_i is chosen independently at random according to the following simple distribution:* for any integer $k \ge 1$, $\mathbf{Pr}[I_i = k] = 2^{-k}$. *However, for such a sequence it is* not true *that the Linear Randomized Broadcasting Algorithm completes broadcasting in $\mathcal{O}(n)$ rounds with probability at least $1 - n^{-1}$. Consider complete layered n-node network with three layers L_0, L_1, L_2, with $|L_0| = |L_2| = 1$*

and $|L_1| = n-2$*, where the only edges are from all nodes in* L_i *to all nodes in* L_{i+1} *for* $i = 0, 1$. *The Linear Randomized Broadcasting Algorithm with the modified sequence for such networks requires* $\Omega(n \log(1/\beta))$ *rounds to complete the broadcasting with the probability of at least* $1-\beta$. *With our setting of* $\beta = n^{-1}$*, this implies* $\Omega(n \log n)$ *time.*

3.1.1. Deterministic sequence

Some users may find the use of a randomized sequence too complicated or not necessary, but actually, one can use also a deterministic definition of the sequence $\mathfrak{I}$ which is (in the stochastic sense) almost identical to the previous definition. Let us consider a sequence $\mathfrak{I} = \langle I_1, I_2, \dots \rangle$ of integers from $\{1, 2, \dots \log n\}$. For an integer $1 \le k \le \log n$ define $\mathtt{density}(\mathfrak{I}, k)$ to be the smallest integer ℓ such that in every subsequence of ℓ consecutive elements of $\mathfrak{I}$ there is at least one with value k. Let $\mathfrak{I}$ be any infinite sequence satisfying the following *deterministic density property*:

- for all $k \in \{1, 2, \dots, \log n\}$: $\mathtt{density}(\mathfrak{I}, k) = \mathcal{O}(\min\{2^k, \log n\})$, and
- every subsequence of $\mathcal{O}(\log n)$ consecutive elements contains a subsequence $\langle \mathcal{LL}(n) + 1, \dots, \log n \rangle$.

Now, we define a simple sequence satisfying the deterministic density property. We first define recursively the auxiliary sequence $\mathsf{T}(s)$ for $s \ge 1$:

$$\mathsf{T}(1) = 1 \ ; \qquad \mathsf{T}(s) = \mathsf{T}(s-1)\ s\ \mathsf{T}(s-1) \text{ for } s \ge 2 \ .$$

For example, $\mathsf{T}(2) = \langle 1, 2, 1 \rangle$, $\mathsf{T}(3) = \langle 1, 2, 1, 3, 1, 2, 1 \rangle$, and $\mathsf{T}(4) = \langle 1, 2, 1, 3, 1, 2, 1, 4, 1, 2, 1, 3, 1, 2, 1 \rangle$.

Let us take two infinite sequences $\mathfrak{I}^{\langle 1 \rangle}$ and $\mathfrak{I}^{\langle 2 \rangle}$ which result by iterating the sequence $\mathsf{T}(\mathcal{LL}(n))$ and the sequence $\mathcal{LL}(n) + 1, \mathcal{LL}(n) + 2, \dots \log n$, respectively. The deterministic selective sequence $\mathfrak{I}$ is constructed by interleaving the sequences $\mathfrak{I}^{\langle 1 \rangle}$ and $\mathfrak{I}^{\langle 2 \rangle}$. For example, if $n = 2^7$ (and hence, $\mathcal{LL}(n) = 3$ and $\log n = 7$), then

$$\mathfrak{I} = \langle 1, 4, 2, 5, 1, 6, 3, 7, 1, 4, 2, 5, 1, 6, 1, 7, 2, 4, 1, 5, 3 \cdots \rangle$$

Lemma 3.2 *The deterministic sequence* $\mathfrak{I}$ *constructed above satisfies the deterministic density property.* □

Assuming we have a sequence $\mathfrak{I}$ satisfying the deterministic density property, we can replace it in the Linear Randomized Broadcasting Algorithm to obtain the following theorem.

Theorem 2 *There exists* $\mathsf{T} = \Theta(n)$ *such that if one replaces the randomized sequence* $\mathfrak{I}$ *by any deterministic sequence satisfying the deterministic density property, then the Linear Randomized Broadcasting Algorithm* (T) *completes broadcasting in any* n*-node radio network with probability at least* $1 - n^{-1}$.

3.2. Notational conventions and basic properties

Before we provide a formal analysis of the Linear Randomized Broadcasting Algorithm(T), let us first introduce some further notation and give some properties of the algorithm.

Let $\Gamma_v[r]$ be the set of nodes $u \in V$ that are active at the end of round r and for which $(u, v) \in E$. Let $\gamma_v[r] = |\Gamma_v[r]|$. Let Act_v be the random variable denoting the round in which node v becomes active (for the source node, which has index 1, we define $\mathsf{Act}_1 = 0$). Notice that a node v will make its first attempt to transmit at round $\mathsf{Act}_v + 1$, one round after it becomes active. Observe also that the broadcasting time of the Linear Randomized Broadcasting Algorithm(T) is equal to $\max\{\mathsf{Act}_v : v \in V\}$.

Eager nodes and successful rounds. For any node $v \in V$, we say v is *eager at round* r if $\gamma_v[r-1] > 0$ and $\frac{1}{2}\gamma_v[r-1] < 2^{I_r} \le \gamma_v[r-1]$, where I_r is the rth element in sequence $\mathfrak{I}$.

Lemma 3.3 *For any node* $v \in V$ *and any round* $r \in \mathbb{N}$*, if* $\gamma_v[r-1] > 0$*, then the probability that* v *is eager at round* r *is at least* $\frac{1}{2} \max\{1/\gamma_v[r-1], 1/\log n\}$.

Proof : Follows immediately from the definition of the α_k's that implies that for any $k \in \{1, \dots, \log n\}$ and any $r \in \mathbb{N}$, we have $\mathbf{Pr}[I_r = k] = \frac{1}{2} \max\{2^{-k}, 1/\log n\}$. □

Lemma 3.4 *For any node* $v \in V$ *and any round* $r \in \mathbb{N}$*, if* v *is eager at round* r*, then the probability that* v *is active at the end of round* r *is at least* 0.1.

Proof : Since v is eager at round r, we have $0 < \frac{1}{2}\gamma_v[r-1] < 2^{I_r} \le \gamma_v[r-1]$.

Let us consider round r in the Linear Randomized Broadcasting Algorithm(T). At the beginning of that round there are $s = \gamma_v[r-1] \ge 1$ nodes in $\mathcal{N}$ that are active and that can send a message to node v. In order the message to be successfully received, there must be exactly one node in $\Gamma_v[r-1]$ that will transmit in round r. Since each active node transmits the message in round r independently with probability 2^{-I_r}, and since $\frac{1}{2}s < 2^{I_r} \le s$, the probability that one message will be successfully transmitted to v is equal to $\binom{s}{1} \cdot 2^{-I_r} \cdot (1 - 2^{-I_r})^{s-1}$, which is at least 0.1 for all $s \ge 2$ (for $s = 1$ this probability is 1). This clearly means that v is active at the end of round r with probability at least 0.1. □

3.3. Proof of Theorem 1: analysis of the Linear Randomized Broadcasting Algorithm

In our analysis we concentrate ourselves on a single path from the source to an arbitrary node and show that with high probability the end node of the path will be active after

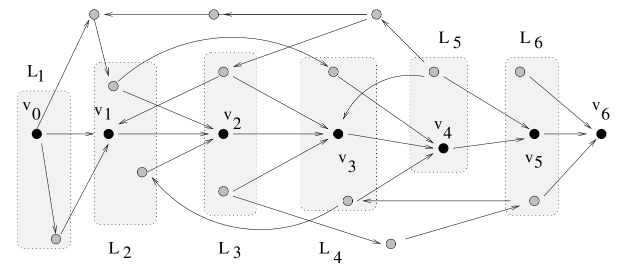

Figure 1. Let $P = (v_0, v_1, \ldots v_6)$. Then $L_i = \mathsf{Layer}_P(i)$ for the example graph.

$\mathcal{O}(n)$ rounds. Since there are n paths whose union ends at all the vertices in the network, this will suffice to prove that the algorithm will complete broadcasting in $\mathcal{O}(n)$ rounds, with high probability.

Let $P = \langle v_0, v_1, \ldots, v_\ell \rangle$ be any path that begins at the source. For any node $v \in V$, let us define

$$\mathrm{last}_P(v) = \begin{cases} v_i & \text{if } (v, v_i) \in E \text{ and } \forall_j((v, v_j) \in E \Rightarrow j \leq i) \\ \emptyset & \text{if } \forall_j (v, v_j) \notin E \ . \end{cases}$$

For any node i, $1 \leq i \leq \ell$, let $\mathsf{Layer}_P(i) = \{v \in V : \mathrm{last}_P(v) = v_i\}$, see Figure 1. The set $\mathsf{Layer}_P(i)$ is called the *layer* of rank i with respect to P. We say a layer $\mathsf{Layer}_P(i)$ is *leading* (with respect to P) at round r of the algorithm if i is the highest rank layer containing an active node at the beginning of round r (that is, $i = \max\{j : \exists v \in \mathsf{Layer}_P(j)$ with $\mathsf{Act}_v < r\}$) *and* node v_ℓ is not active at the beginning of round r. Define:

$$Z_P(i) \ = \ |\{r : \mathsf{Layer}_P(i) \text{ is leading at round } r\}| \ .$$

The following lemma describes some basic and key properties of the definitions above.

Lemma 3.5 *Consider the layers corresponding to some shortest path $P = \langle v_0, v_1, \ldots, v_\ell \rangle$, where $v_0 = 1$.*

- *The layers are disjoint, consequently $\sum_{i=0}^{\ell} |\mathsf{Layer}_P(i)| \leq n$.*
- *The last node of P, v_ℓ, receives the message after $\sum_{i=1}^{\ell} Z_P(i)$ rounds, that is, $\mathsf{Act}_{v_\ell} = \sum_{i=1}^{\ell} Z_P(i)$.*
- *For every $r \in \mathbb{N}$ and any i, $1 \leq i \leq \ell$, if $\mathsf{Layer}_P(i)$ is leading at round r, then $\gamma_{v_i}[r-1] \leq |\mathsf{Layer}_P(i)|$.*

Proof : The first two claims follow immediately from the definitions above. We show the third claim by contradiction. Let us suppose that $\gamma_{v_i}[r-1] > |\mathsf{Layer}_P(i)|$. Since $\gamma_{v_i}[r-1] = |\Gamma_{v_i}[r-1]| = |\{v \in V : (v, v_i) \in E$ and $\mathsf{Act}_v \leq r-1\}|$, $\gamma_{v_i}[r-1] > |\mathsf{Layer}_P(i)|$ implies that there must be some node x with $(x, v_i) \in E$ that is active in round $r-1$ and which is not in $\mathsf{Layer}_P(i)$. Since it is easy to see that $\{v \in V : (v, v_i) \in E\} \subseteq \bigcup_{j=i}^{\ell} \mathsf{Layer}_P(j)$, this yields $x \in \mathsf{Layer}_P(j)$ for some $j > i$. Therefore, if $\gamma_{v_i}[r-1] > |\mathsf{Layer}_P(i)|$ then $\mathsf{Layer}_P(i)$ is not leading at round r, what implies the third claim. □

Our next **key** lemma gives a majorization result for the round in which the last node from P becomes active. (Let us recall that $\mathfrak{Y}$ has *geometric distribution with parameter* ρ if $\mathbf{Pr}[\mathfrak{Y} = k] = (1-\rho)^{k-1} \cdot \rho$ for any $k \in \mathbb{N}$.)

Lemma 3.6 (Key majorization lemma) *Let $P = \langle v_0, v_1, \ldots, v_\ell \rangle$ be any shortest path from the source to a node v_ℓ. Let $Y_1, \ldots, Y_\ell$ be a sequence of independent random variable having geometric distribution with parameter ρ_i each, where $\rho_i = \frac{1}{20 \cdot \min\{\log n, |\mathsf{Layer}_P(i)|\}}$. Then the random variable Act_{v_ℓ} is stochastically majorized by the sum $\sum_{i=1}^{\ell} Y_i$, that is, for any $r \in \mathbb{N}$,*

$$\mathbf{Pr}\big[v_\ell \textit{ is active at the end of round } r\big] \ \geq \ \mathbf{Pr}\Big[\sum_{i=1}^{\ell} Y_i \leq r\Big] \ .$$

Proof : Since from Lemma 3.5 we know that $\mathsf{Act}_{v_\ell} = \sum_{i=1}^{\ell} Z_P(i)$, it is enough to prove that for every i, $1 \leq i \leq \ell$, the random variable $Z_P(i)$ is majorized by Y_i.

Let us fix any i and we show that $Z_P(i)$ is majorized by Y_i. Observe first that if $\mathsf{Layer}_P(i)$ is never a leading layer, then $Z_P(i)$ is majorized by Y_i. Otherwise, let $r_{\langle i \rangle}$ be the first round in which $\mathsf{Layer}_P(i)$ is leading. Let us consider any round $r \geq r_{\langle i \rangle}$ in which $\mathsf{Layer}_P(i)$ is leading. Then, there is at least one node in $\Gamma_{v_i}[r-1]$, and therefore $\gamma_{v_i}[r-1]$ nodes will try to transmit in round r to node v_i. By Lemma 3.3, the probability that node v_i is eager is at least $\frac{1}{2} \max\{1/\gamma_{v_i}[r-1], 1/\log n\}$, and by Lemma 3.4, if v_i is eager in round r, then the probability that it will become active (and hence $\mathsf{Layer}_P(i)$ will be not leading at round $r+1$) is at least 0.1. Therefore, since by Lemma 3.5 we know that $\gamma_{v_i}[r-1] \leq |\mathsf{Layer}_P(i)|$, we can conclude that the probability that node v_i will become active at round r is at least $\frac{1}{20 \cdot \min\{|\mathsf{Layer}_P(i)|, \log n\}}$. Since this bound is true in every round r in which $\mathsf{Layer}_P(i)$ is leading, we clearly have

$$\begin{aligned} \mathbf{Pr}[Z_P(i) = k] \ &\geq \ \frac{1}{20 \cdot \min\{|\mathsf{Layer}_P(i)|, \log n\}} \cdot \\ &\cdot \ \left(1 - \frac{1}{20 \cdot \min\{|\mathsf{Layer}_P(i)|, \log n\}}\right)^{k-1} \\ &= \ \mathbf{Pr}[Y_i = k] \ , \end{aligned}$$

what concludes the proof of the lemma. □

Next, we prove a concentration result for sums of geometric random variables (proof is deferred to Appendix A).

Lemma 3.7 *Let $X_1, \ldots, X_\ell$ be a sequence of independent integer-valued random variables, each X_i being geometrically distributed with a parameter p_i, $0 < p_i < 1$. For every $1 \leq i \leq \ell$, let $\mu_i = 1/p_i$, and assume that all μ_i are from a set Δ, that is, $\Delta = \{\mu_i : 1 \leq i \leq \ell\}$. If $\sum_{i=1}^{\ell} \mu_i \leq N$,* *then for every positive real number* β,

$$\mathbf{Pr}\Big[\sum_{i=1}^{\ell} X_i \;\le\; 2\cdot N + 8\cdot \ln(|\Delta|/\beta)\cdot \sum_{z\in\Delta} z\Big] \ge 1-\beta \ .$$

Equipped with Lemmas 3.6 and 3.7, we are now ready to complete the proof of Theorem 1.

Proof of Theorem 1: Our goal is to show that for every node $v \in V$ we have $\mathsf{Act}_v \le c\,n$ with probability at least $1-n^{-2}$. Since there are n nodes to be considered, the union bound will imply the theorem:

$$\begin{aligned}\mathbf{Pr}[\forall_{v\in V}\ \mathsf{Act}_v \le c\,n] &\ge 1-\sum_{v\in V}\mathbf{Pr}[\mathsf{Act}_v > c\,n]\\ &\ge 1-n\cdot n^{-2} \;=\; 1-n^{-1} \ .\end{aligned}$$

Let us fix any node $v \in V$. Let us consider a shortest path from the source to v and denote it by $P = \langle v_0, v_1, \dots, v_\ell\rangle$. By analyzing the algorithm on this path we show that $\mathbf{Pr}[\mathsf{Act}_v \le c\,n] \ge 1-n^{-2}$.

We apply Lemma 3.6 to estimate the probability that $\mathbf{Pr}[\mathsf{Act}_v \le c\,n] \ge 1-n^{-2}$. Let us define the random variable $Y_1, Y_2, \dots, Y_\ell$ as in Lemma 3.6. Now, we can estimate the expected value Act_v as follows:

$$\begin{aligned}\mathbf{E}[\mathsf{Act}_v] &= \sum_{i=1}^{\ell}\mathbf{E}[Y_i] \;=\; \sum_{i=1}^{\ell}\frac{1}{\rho_i} \;=\\ &= \sum_{i=1}^{\ell} 20\cdot\min\{\log n, |\mathsf{Layer}_P(i)|\}\\ &\le 20\cdot\sum_{i=1}^{\ell}|\mathsf{Layer}_P(i)| \;\le\; 20\cdot n \ ,\end{aligned}$$

where the last inequality follows from Lemma 3.5.

Next, we want to prove that the sum of random variables $\sum_{i=1}^{\ell} Y_i$ is highly concentrated around its mean. For this, we use Lemma 3.7, from which we obtain (with $X_i \equiv Y_i$, $\Delta = \{1, 2, \dots, 20 \log n\}$, $\beta = n^{-2}$, and $N = 20\cdot n$):

$$\mathbf{Pr}\Big[\sum_{i=1}^{\ell} Y_i \le 40n + 8\ln(20n^2\log n)(20\log n)^2\Big] \ge 1-n^{-2} \ .$$

This immediately implies that for a certain positive constant c we obtain

$$\mathbf{Pr}\big[\mathsf{Act}_v \le c\,n\big] \;=\; \mathbf{Pr}\Big[\sum_{i=1}^{\ell} Y_i \le c\,n\Big] \;\ge\; 1-n^{-2} \ .$$

This inequality, when combined with our arguments at the beginning of the proof, concludes the proof. □

3.4. Analysis of the Linear Randomized Broadcasting Algorithm with deterministic sequences

In this section, we sketch the proof of Theorem 2, of how to modify the analysis of the Linear Randomized Broadcasting Algorithm(T) to show that the randomized sequence $\mathfrak{I}$ can be replaced by any deterministic sequence satisfying the deterministic density property. Our main observation is that even though for the deterministic selecting sequence (and satisfying the deterministic density property) defined in Section 3.1.1 the key lemma in the proof of Theorem 1, which is Lemma 3.6, is false, we can replace it by a relaxed version of that lemma, Lemma 3.8, which is still sufficient to conduct the analysis from Section 3. Before we state that lemma, let us introduce some definitions.

We say an algorithm is in *stage* s if it is in a round r, where $(s-1)\cdot\sigma\log n + 1 \le r \le s\cdot\sigma\log n$, for certain constant σ. Let $P = \langle v_0, v_1, \dots, v_\ell\rangle$ be any path in $\mathcal{N}$ that begins at the source node. We use the same definitions of layers and related variables as in Section 3 and observe that Lemma 3.5 holds. If $\mathsf{Layer}_P(i)$ is leading at round r, then we define

$$\mathsf{Leftover}_P(r) \;=\; \sum_{j=i}^{\ell}|\mathsf{Layer}_P(j)| \ .$$

For any $s \in \mathbb{N}$, we define a function $\Delta_P(s)$, which represent the "gain" in stage s:

$$\Delta_P(s) \;=\; \mathsf{Leftover}_P((s-1)\cdot\sigma\log n) - \mathsf{Leftover}_P(s\cdot\sigma\log n).$$

The following lemma plays the role similar to the role of Lemma 3.6 in the proof of Theorem 1.

Lemma 3.8 *Let* $P = \langle v_0, v_1, \dots, v_\ell\rangle$ *be any shortest path from the source node to a node* v_ℓ*. Then, for any* $s \in \mathbb{N}$*, with probability at least* 0.1*, either*

- $\Delta_P(s) \ge \log n$*, or*
- $\mathsf{Leftover}_P(s\cdot\sigma\log n) = 0$*; in which case the node* v_ℓ *is active at the end of stage* s*.* □

The way one should understand this lemma is that for the deterministic sequence in every $\sigma\log n$ consecutive rounds (a stage) we have at least a "logarithmic" progress with probability at least 0.1.

Once we have Lemma 3.8, similar approach as that used in the proof of Theorem 1 together with the fact that $\mathsf{Leftover}_P(0) = n$, can be used to obtain the proof of Theorem 2. □

4. Improved randomized algorithm for shallow networks

The analysis in Section 3 gives asymptotically optimal broadcasting time in n-node radio networks. However, if we consider n-node networks of eccentricity D, then we can still hope to obtain a better broadcasting algorithm. In this section we present the first algorithm that achieves the optimal broadcasting time in "shallow networks": we extend our approach from Section 3 to a broadcasting requiring $\mathcal{O}(D \log(n/D) + \log^2 n)$ rounds.

4.1. Randomized selecting sequence for shallow networks

Similarly as in Section 3, we consider the algorithm using a randomized selecting sequence; later, in Section 4.3, we briefly discuss deterministic selecting sequences.

We consider the Randomized Broadcasting Algorithm for Shallow Networks (T), which is identical to the Linear Randomized Broadcasting Algorithm(T) with the sequence $\mathfrak{I}$ replaced by $\mathfrak{I}_D$. The sequence $\mathfrak{I}_D$ is defined below using a new distribution α'.

Let us fix D and n. Let $\lambda = \log(n/D)$. Let us remind that $\mathcal{LL}(n) = \lceil \log \log n \rceil$ and that we assumed that both, n and D are power of 2 (hence, λ is an integer). In the previous section, while discussing a randomized $\mathcal{O}(n)$-time broadcasting algorithm we were using a distribution α depending on n. Now, we use a similar distribution, denoted by α', that is defined for a given D as follows (notice that $\alpha_k = \alpha'_k$ for $D = n$),

$$\alpha'_k = \begin{cases} \frac{1}{2\lambda} & \text{for } 1 \le k \le \lambda, \\ \frac{1}{2\lambda} \cdot 2^{-(k-\lambda)} & \text{for } \lambda < k \le \lambda + \mathcal{LL}(n), k \le \log n, \\ \frac{1}{2\lambda} \cdot \frac{1}{\log n} & \text{for } \lambda + \mathcal{LL}(n) < k \le \log n, \\ 1 - \sum_{i=1}^{\log n} \alpha'_i & \text{for } k = 0 \enspace . \end{cases}$$

With these definitions, we can define a randomized sequence $\mathfrak{I}_D = \langle J_1, J_2, \ldots \rangle$ as the sequence of elements chosen independently at random according to the distribution α', that is, $\mathbf{Pr}[J_r = k] = \alpha'_k$.

Randomized Broadcasting Algorithm for Shallow Networks(T)

Input: Network $\mathcal{N} = (V, E)$, $n = |V|$, and the eccentricity D
Randomized sequence $\mathfrak{I}_D = \langle J_1, J_2, \ldots \rangle$ such that
$\mathbf{Pr}[J_r = k] = \alpha'_k \ \forall r \in \mathbb{N}, \forall k \in \{0, 1, 2, \ldots, \log n\}$

for $r = 1$ **to** T **do** { *round number* r }
  for each *active* node $v \in V$ *independently* **do**
    node v transmits with probability 2^{-J_r}

4.2. Optimal algorithm for shallow networks

In this section we analyze the Randomized Broadcasting Algorithm for Shallow Networks (T) for arbitrary n-node networks of eccentricity D. The analysis is similar to the analysis in Section 3 and below we only sketch the proof. (We could apply our analysis for any D to achieve the broadcasting time of $\mathcal{O}(D \log(n/D) + \log^2 n)$ with high probability. However, since such a result is already known for small $D = \mathcal{O}(\log^3 n)$, we concentrate our analysis on the most interesting case.)

Theorem 3 *Let $\mathcal{N} = (V, E)$ be an n-node radio network of eccentricity $D = \Omega(\log^3 n)$. There exists a constant c such that if $T \ge c\, D \log(n/D)$, then the Randomized Broadcasting Algorithm for Shallow Networks (T) completes broadcasting in $\mathcal{N}$ with probability at least $1 - n^{-1}$.*

Sketch of the proof : The proof is similar to the proof of Theorem 1. We can consider the layers with respect to a specified path starting at the source node. We separately analyze the time spent in *small* layers ($|\mathsf{Layer}_P(i)| \le n/D$) and *large* layers. We can show that the expected number of rounds spent in any small layer is $\mathcal{O}(\lambda) \equiv \mathcal{O}(\log(n/D))$, and hence, since there are at most D layers and by applying an appropriate concentration bound, we can show that the total number of rounds spent in small layers is $\mathcal{O}(D \cdot \log(n/D))$, with high probability. For large layers, we can show that the expected number of rounds spent in a large layer of size $s \cdot (n/D)$ is $\mathcal{O}(s \cdot \lambda) \equiv \mathcal{O}(s \cdot \log(n/D))$. Consequently, the time spent in all larger layers is upper bounded by $\sum_{i=1}^{\ell} \mathcal{O}(|\mathsf{Layer}_P(i)| \cdot (D/n) \cdot \log(n/D))$, with high probability, which is $\mathcal{O}(D \cdot \log(n/D))$.

Summarizing, the algorithm completes broadcasting in $\mathcal{O}(D \cdot \log(n/D))$ rounds, with high probability. $\square$

4.3. Deterministic selecting sequence for randomized algorithm for shallow networks

Now, we briefly discuss how our analysis can be applied to deterministic selecting sequences. We say an infinite sequence $\mathfrak{I}_D$ satisfies the D-modified deterministic density property if the following conditions are satisfied:

- every subsequence of $\mathcal{O}(\log(n/D))$ consecutive elements contains a subsequence $\langle 1, 2, \ldots, \log(n/D) \rangle$, and
- for every k, $\log(n/D) < k \le \log(n/D) + \mathcal{LL}(n)$, $\mathsf{density}(\mathfrak{I}_{\mathfrak{D}}, k) = \mathcal{O}(\log(n/D) \cdot 2^k \cdot (D/n))$, and
- every subsequence of $\mathcal{O}(\log n \cdot \log(n/D))$ consecutive elements contains a subsequence $\langle \log(n/D) + \mathcal{LL}(n) + 1, \ldots, \log n \rangle$.

It is easy to construct such a sequence deterministically, and we leave the easy details of the construction to the readers.

Using modification of the proof of Theorem 3 similar to those done in the proof of Theorem 2, we can show:

Theorem 4 *Let $\mathcal{N} = (V, E)$ be an n-node radio network of eccentricity $D = \Omega(\log^3 n)$. There exists $T = \Theta(D \log(n/D))$, such that if one replaces the randomized sequence $\mathfrak{I}_D$ by any deterministic sequence satisfying the*

D-*Modified deterministic density property, then the Randomized Broadcasting Algorithm for Shallow Networks* (T) *completes broadcasting in any* n*-node radio network of eccentricity* D *with probability at least* $1 - n^{-1}$. □

5. $\mathcal{O}(n \log^2 \mathsf{D})$-time oblivious deterministic broadcasting in shallow networks

We consider an unknown n-node radio network $\mathcal{N}$ with eccentricity D. It has been shown nonconstructively, using very complicated counting arguments, that there is a deterministic protocol for such networks that completes broadcasting in time $\mathcal{O}(n \log n \log \mathsf{D})$, see [17]. In this paper, we present a more natural approach to this problem and obtain an improvement to $\mathcal{O}(n \log^2 \mathsf{D})$. Moreover our algorithm is significantly simpler.

Our improvement comes from two main sources:

- appropriately chosen *structure* of the selecting sequence and
- in the selecting sequence only indices corresponding to powers of two between n/D and n are considered; there are only $\log \mathsf{D}$ such powers.

We consider an infinite sequence $\mathfrak{I}$ of elements from $\{0, 1, 2, \ldots, \log \mathsf{D}\}$. Define *sparseness*$(\mathfrak{I}, k)$ to be the minimal distance between two distinct positions in $\mathfrak{I}$ containing k. Let us define two properties of sequences $\mathfrak{I}$.

Sparseness property: $\mathsf{sparseness}(\mathfrak{I}, k) = \Omega(2^k)$ for each $0 \le k \le \log \mathsf{D}$.

Modified density property: $\mathsf{density}(\mathfrak{I}, k) = O(2^k)$ for each $0 \le k \le \log \mathsf{D}$.

Lemma 5.1 *There exists a sequence* $\mathfrak{I}$ *satisfying both the sparseness property and the modified density property.*

Proof : We use a modification of the finite sequence T from Section 3.1.1. Let $\mathsf{T}^*_i = \mathsf{T}(\log(\mathsf{D}) + 1)_i - 1$ for each index $1 \le i \le |\mathsf{T}(\log(\mathsf{D}) + 1)|$. Then $\mathfrak{I}$ results by iterating the sequence T^*. □

Following [10], for a given n, a family $\mathcal{F}$ of subsets of $[n]$ is called a j-*selector* if for any pair of disjoint sets $\mathsf{X} \subseteq [n]$ and $\mathsf{Y} \subseteq [n]$ with $j/2 \le |\mathsf{X}| \le j$ and $|\mathsf{Y}| \le j$ there exists a set $\mathsf{F} \in \mathcal{F}$ such that $|\mathsf{X} \cap \mathsf{F}| = 1$ and $|\mathsf{Y} \cap \mathsf{F}| = 0$

In [12], using probabilistic arguments, j-selectors of asymptotically optimal size have been constructed.

Lemma 5.2 **[12]** *For every* $1 \le j \le n$*, there exists a* j*-selector* $\mathcal{F}^j$ *of size* $\mathcal{O}(j \log(n/j))$.

For any k, $0 \le k \le \log \mathsf{D}$, let

$$\mathcal{F}^k = \{\mathcal{F}^k_1, \mathcal{F}^k_2, \mathcal{F}^k_3, \ldots, \mathcal{F}^k_{s_k}\}$$

be a $(2^k n/\mathsf{D})$-selector of size $s_k = \mathcal{O}(2^k (n/\mathsf{D}) \cdot \log(\mathsf{D}/2^k))$, as promised by Lemma 5.2.

Oblivious Deterministic Broadcasting Algorithm(M)

Input: Network $\mathcal{N} = (\mathsf{V}, \mathsf{E})$ of eccentricity D.
Sequence $\mathfrak{I} = \langle \mathfrak{I}_1, \mathfrak{I}_2, \ldots \rangle$ satisfying the sparseness and modified density properties.
$(2^k (n/\mathsf{D}))$-selectors $\mathcal{F}^k$, for all $0 \le k \le \log \mathsf{D}$.

for $t = 1$ **to** M **do** { *phase number* t of order $\mathfrak{I}_t$ }
  $k = \mathfrak{I}_t$;
  for $j = 2$ **to** k **do**
    for $r = 1$ **to** s_j **do**
      for each *active* node $v \in \mathsf{V}$ **do**
        if $v \in \mathcal{F}^j_r$ **then** v transmits

Lemma 5.3 *There exists a constant* c *such that if* $\mathsf{M} \ge c\,\mathsf{D}$*, then the Oblivious Deterministic Broadcasting Algorithm*(M) *completes broadcasting in* $\mathcal{N}$.

Proof : The proof is similar to the proof of Theorem 1. We consider layers with respect to a specified path P. The following fact follows from the definition of selector families, using the same arguments as in [10].
Claim If $|\mathsf{Layer}_\mathsf{P}(i)| \le 2^k \cdot (n/\mathsf{D})$ and $\mathsf{Layer}_\mathsf{P}(i)$ is *leading*, then after a phase of order k it is no longer *leading*. □

For each layer we can estimate the number of *phases* in which $\mathsf{Layer}_\mathsf{P}(i)$ is *leading*. Due to the modified density property this number is $\mathcal{O}(|\mathsf{Layer}_\mathsf{P}(i)| \cdot (\mathsf{D}/n) + 1)$, independently in which position of the selective sequence $\mathfrak{I}$ we start. Consequently the number of phases is $\mathcal{O}\left(\sum |\mathsf{Layer}_\mathsf{P}(i)| \cdot (\mathsf{D}/n) + 1\right) = \mathcal{O}(\mathsf{D})$. □

Theorem 5 *Let* $\mathcal{N} = (\mathsf{V}, \mathsf{E})$ *be an* n*-node radio network of eccentricity* D*. There exists a constant* c *such that if* $\mathsf{T} \ge c\, n \log^2 \mathsf{D}$*, then the Oblivious Deterministic Broadcasting Algorithm completes broadcasting in* $\mathcal{N}$ *after* T *rounds.*

Proof : Due to Lemma 5.2, the number of rounds performed by each phase of order k is

$$s_2 + s_3 + \ldots + s_k = \mathcal{O}(s_k) = \mathcal{O}(2^k \cdot (n/\mathsf{D}) \cdot \log \mathsf{D}) \ .$$

Due to the sparseness property of the sequence $\mathfrak{I}$ and Lemma 5.3 we perform $\tau_k = \mathcal{O}(\mathsf{D}/2^k)$ phases of order k. Hence the total number of rounds is

$$\mathcal{O}\Big(\sum_{k=0}^{\log \mathsf{D}} s_k \cdot \tau_k\Big) = \mathcal{O}\left(\sum_{k=0}^{\log \mathsf{D}} (2^k \cdot (n/\mathsf{D}) \cdot \log \mathsf{D}) \cdot (\mathsf{D}/2^k)\right) = \mathcal{O}\left(\sum_{k=0}^{\log \mathsf{D}} n \log \mathsf{D}\right) = \mathcal{O}(n \log^2 \mathsf{D})$$ □

6. Application to randomized gossiping

Our randomized algorithm from Section 3 can be used as a building block to improve the complexity of randomized broadcasting in n-node radio networks. We base on the previous works [11, 20]. Comparing to the algorithm from [11], we first begin with $\mathcal{O}(\log^2 n)$ rounds of the round robin algorithm to ensure that every message is received by at least $\mathcal{O}(\log^2 n)$ nodes. Then, we run algorithm RANDGOSSIP, but only with the values of $i = 2\log\log n + \Theta(1), \dots, \log n - 1$, and each execution of algorithm $\text{LTDBROADCAST}_v(2^{i+1})$ is replaced by our Linear Randomized Broadcasting Algorithm (T) that takes v as the source node and set $T = c \cdot 2^{i+1}$ for certain large constant c. With this modification, we can prove the following result. (More details are deferred to the full version of the paper.)

Theorem 6 *Let $\mathcal{N} = (V, E)$ be an n-node radio network. There exists a randomized Las Vegas algorithm that completes gossiping in $\mathcal{N}$ in $\mathcal{O}(n \log^2 n)$ rounds with probability at least $1 - n^{-1}$.* □

7. Final remarks

We have shown an optimal randomized algorithms for the broadcasting problem in radio networks. Our randomized broadcasting algorithm are *oblivious*, all nodes perform in the same way, according to a same selecting sequence.

For two other problems, deterministic broadcasting in shallow networks and randomized gossiping we presented *almost* optimal algorithms and, using our approach we improved the complexity of previously known algorithm for these problems.

Our algorithm for deterministic broadcasting in shallow networks is not only faster than previously known algorithms, it is much simpler and it is oblivious, while the algorithm from [17] is not.

Final note. After submitting this paper, we have learned that very recently a similar *randomized* broadcasting algorithm was obtained independently by Kowalski and Pelc [18]. The main model their consider was that of *undirected* networks, in which often the broadcasting can be performed more efficiently, but as it is claimed in [18], their randomized algorithm can be applied to directed networks as well. Their algorithm seems to be very similar to our randomized $\mathcal{O}(D \log(n/D) + \log^2 n)$-time algorithm.

Kowalski and Pelc [18] gave also a deterministic algorithm for broadcasting in *undirected* networks with the running time of $\mathcal{O}(n \min\{\log n, D\})$. Their algorithm can be combined with our deterministic algorithm to obtain an $\mathcal{O}(n \min\{\log n, \log^2 D\})$-time deterministic broadcasting algorithm in *undirected* networks.

References

[1] 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.

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

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

[4] 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(1): 104–126, 1992.

[5] I. Chlamtac and S. Kutten. On broadcasting in radio networks — problem analysis and protocol design. *IEEE Transaction on Communications*, 33(12): 1240–1246, December 1985.

[6] I. Chlamtac and O. Weinstein. The wave expansion approach to broadcasting in multihop radio networks. *IEEE Transaction on Communications*, 39(3): 426–433, March 1991.

[7] B. S. Chlebus. Randomized communication in radio networks. In S. Rajasekaran, P. M. Pardalos, J. H. Reif, and J. D. P. Rolim, eds., *Handobook of Randomized Algorithms*. Kluwer Academic Publishers, 2001, vol. I, pages 401–456.

[8] B. S. Chlebus, L. Gąsieniec, A. Gibbons, A. Pelc, and W. Rytter. Deterministic broadcasting in unknown radio networks. In *Proceedings of the 11th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 861–870, San Francisco, CA, January 9–11, 2000. SIAM, Philadelphia, PA.

[9] B. S. Chlebus, L. Gąsieniec, A. Östlin, and J. M. Robson. Deterministic Radio Broadcasting. In *Proceedings of the 27th Annual International Colloquium on Automata, Languages and Programming (ICALP)*, volume 1853 of *Lecture Notes in Computer Science*, pages 717–728, Geneva, Switzerland, July 9–15, 2000. Springer-Verlag, Berlin.

[10] M. Chrobak, L. Gąsieniec, and W. Rytter. Fast broadcasting and gossiping in radio networks. In *Proceedings of the 41st IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 575–581, Redondo Beach, CA, November 12–14, 2000. IEEE Computer Society Press, Los Alamitos, CA.

[11] M. Chrobak, L. Gąsieniec, and W. Rytter. A randomized algorithm for gossiping in radio networks. In *Proceedings of the 7th Annual International Computing Combinatorics Conference (COCOON)*, volume 2108 *Lecture Notes in Computer Science*, pages 483–492, Guilin, China, August 20–23, 2001. Springer-Verlag, Berlin.

[12] A. E. F. Clementi, A. Monti, and R. Silvestri. Distributed broadcasting in radio networks of unknown topology. *Theoretical Computer Science*, 302: 337-364, 2003. A preliminary version appeared in *Proceedings of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 709–718, Washington, DC, January 7–9, 2001. SIAM, Philadelphia, PA.

[13] I. Gaber and Y. Mansour. Broadcast in radio networks. *Journal of Algorithms*, 46(1): 1–20, January 2003. A preliminary version appeared in *Proceedings of the 6th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages

577–585, San Francisco, CA, January 22–24, 1995. SIAM, Philadelphia, PA.

[14] R. G. Gallager. A perspective on multiaccess channels. *IEEE Transactions on Information Theory*, 31(2): 124–142, March 1985.

[15] P. Indyk. Explicit constructions of selectors and related combinatorial structures, with applications. In *Proceedings of the 13th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 697–704, San Francisco, CA, January 6–8, 2002. SIAM, Philadelphia, PA.

[16] D. Kowalski and A. Pelc. Deterministic broadcasting time in radio networks of unknown topology. In *Proceedings of the 43rd IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 63–72, Vancouver, BC, Canada, November 16–19, 2002. IEEE Computer Society Press, Los Alamitos, CA.

[17] D. Kowalski and A. Pelc. Faster deterministic broadcasting in ad hoc radio networks. In *Proceedings of the 20th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 2607 *Lecture Notes in Computer Science*, pages 109–120, Berlin, Germany, February 27 – March 1, 2003. Springer-Verlag, Berlin.

[18] D. Kowalski and A. Pelc. Deterministic broadcasting time in radio networks of unknown topology. To appear in *Proceedings of the 22nd ACM Symposium on Principles of Distributed Computing (PODC)*, Boston, MA, July 13–16, 2003. ACM Press, New York, NY.

[19] E. Kushilevitz and Y. Mansour. An $\Omega(D\log(N/D))$ lower bound for broadcast in radio networks. *SIAM Journal on Computing*, 27(3): 702–712, 1998.

[20] D. Liu and M. Prabhakaran. On randomized broadcasting and gossiping in radio networks. In *Proceedings of the 8th Annual International Computing Combinatorics Conference (COCOON)*, volume 2387 *Lecture Notes in Computer Science*, pages 340–349, Singapore, August 15–17, 2002. Springer-Verlag, Berlin.

[21] Y. Xu. An $\mathcal{O}(n^{1.5})$ determinstic gossiping algorithm for radio networks. *Algorithmica*, 36(1):93–96, 2003.

A. Appendix: Proof of Lemma 3.7

Since $\mathbf{E}[X_i] = 1/p_i = \mu_i$, we have $\mathbf{E}\big[\sum_{i=1}^{\ell} X_i\big] = \sum_{i=1}^{\ell} 1/p_i = \sum_{i=1}^{\ell} \mu_i \le N$.

Next, for any $z \in \Delta$, let $K_z = \{i : \mu_i = z\}$. Then, clearly, we have

$$\sum_{i=1}^{\ell} X_i \;=\; \sum_{z\in\Delta}\sum_{i\in K_z} X_i \enspace .$$

Next, our goal is to study $\sum_{i\in K_z} X_i$ for all $z \in \Delta$. Since all X_i with $i \in K_z$ are identically distributed with geometric distribution with the parameter $1/z$, the standard relation between geometric and binomial distributions implies the following inequality that holds for every $z \in \Delta$ and for every $M \in \mathbb{N}$:

$$\mathbf{Pr}\Big[\sum_{i\in K_z} X_i > M\Big] \;\le\; \mathbf{Pr}\big[\mathbb{B}(M, 1/z) \le |K_z|\big] \enspace . \qquad (1)$$

Therefore, we focus our analysis on the study of the binomial distribution. Let us recall a standard variant of Chernoff bound: for any $M \in \mathbb{N}$ and any $0 \le p \le 1$: $\mathbf{Pr}\big[\mathbb{B}(M,p) \le \frac{1}{2}\mathbf{E}[\mathbb{B}(M,p)]\big] \le \exp(-\mathbf{E}[\mathbb{B}(M,p)]/8) = \exp(-Mp/8)$.

For simplicity, let us consider two cases separately:

- $|K_z| \ge 4\ln(|\Delta|/\beta)$: we apply the Chernoff bound above with $M = 2z|K_z|$ to obtain
$$\begin{aligned}\mathbf{Pr}\big[\mathbb{B}(2z|K_z|, 1/z) \le |K_z|\big] &\le \exp(-2|K_z|/8)\\ &\le \exp(-\ln(|\Delta|/\beta))\\ &= \beta/|\Delta| \enspace .\end{aligned}$$
- $|K_z| < 4\ln(|\Delta|/\beta)$: we apply the Chernoff bound above with $M = 8z\ln(|\Delta|/\beta)$ to obtain
$$\begin{aligned}&\mathbf{Pr}\big[\mathbb{B}(8z\ln(|\Delta|/\beta), 1/z) \le |K_z|\big]\\ &\quad\le \mathbf{Pr}\big[\mathbb{B}(8z\ln(|\Delta|/\beta), 1/z) \le 4\ln(|\Delta|/\beta)\big]\\ &\quad\le \exp(-\ln(|\Delta|/\beta)) \;=\; \beta/|\Delta| \enspace .\end{aligned}$$

Thus, we can summarize these two cases in a single bound:

$$\mathbf{Pr}\big[\mathbb{B}(\max\{2z|K_z|, 8z\ln(|\Delta|/\beta)\}, 1/z) \le |K_z|\big] \;\le\; \beta/|\Delta| \enspace .$$

Now, we can combine this inequality with (1) to obtain for every $z \in \Delta$:

$$\begin{aligned}&\mathbf{Pr}\Big[\sum_{i\in K_z} X_i > \max\{2z|K_z|, 8z\ln(|\Delta|/\beta)\}\Big]\\ &\quad\le \mathbf{Pr}\big[\mathbb{B}(\max\{2z|K_z|, 8z\ln(|\Delta|/\beta)\}, 1/z) < |K_z|\big]\\ &\quad\le \beta/|\Delta| \enspace .\end{aligned}$$

Therefore, we can apply the union bound to obtain

$$\begin{aligned}&\mathbf{Pr}\Big[\exists z \in \Delta \text{ such that } \Big(\sum_{i\in K_z} X_i > \max\{2z|K_z|, 8z\ln(|\Delta|/\beta)\}\Big)\Big]\\ &\quad\le \sum_{z\in\Delta}\mathbf{Pr}\Big[\sum_{i\in K_z} X_i < \max\{2z|K_z|, 8z\ln(|\Delta|/\beta)\}\Big]\\ &\quad\le |\Delta|\cdot(\beta/|\Delta|) \;=\; \beta \enspace .\end{aligned}$$

From this, we see that for all $z \in \Delta$, $\sum_{i\in K_z} X_i \le \max\{2z|K_z|, 8z\ln(|\Delta|/\beta)\}$ with probability at least $1-\beta$. Therefore, with probability at least $1-\beta$ we obtain

$$\begin{aligned}\sum_{i=1}^{\ell} X_i &= \sum_{z\in\Delta}\sum_{i\in K_z} X_i \;\le\; \sum_{z\in\Delta}\max\{2z|K_z|, 8z\ln(|\Delta|/\beta)\}\\ &\le \sum_{z\in\Delta}(2z|K_z| + 8z\ln(|\Delta|/\beta))\\ &= 2\cdot\sum_{z\in\Delta} z\cdot|K_z| + 8\cdot\ln(|\Delta|/\beta)\cdot\sum_{z\in\Delta} z\\ &= 2\cdot\sum_{i=1}^{\ell}\mu_i + 8\cdot\ln(|\Delta|/\beta)\cdot\sum_{z\in\Delta} z\\ &\le 2\cdot N + 8\cdot\ln(|\Delta|/\beta)\cdot\sum_{z\in\Delta} z \enspace . \qquad \Box\end{aligned}$$