

Oblivious k -shot Broadcasting in Ad Hoc Radio Networks *

Paraschos Koutris

Aris Pagourtzis

School of Elec. & Comp. Eng., National Technical University of Athens
Polytechnioulou Zografou, 157 80 Athens, Greece,

Email: paris.koutris@gmail.com, pagour@cs.ntua.gr

Abstract

We study distributed broadcasting protocols with few transmissions ('shots') in radio networks of unknown topology. In particular, we examine the case in which a bound k is given and a node may transmit at most k times during the broadcasting protocol. We focus on oblivious algorithms, that is, algorithms where each node decides whether to transmit or not with no consideration of the transmission history. Our main contributions are (a) a lower bound of $\Omega(\frac{n^2}{k})$ on the broadcasting time of any oblivious k -shot broadcasting algorithm and (b) an oblivious broadcasting protocol that achieves a matching upper bound, namely $O(\frac{n^2}{k})$, for every $k \leq \sqrt{n}$ and an upper bound of $O(n^{3/2})$ for every $k > \sqrt{n}$. We also initiate the study of the behavior of general broadcasting protocols by showing an $\Omega(n^2)$ lower bound for any adaptive 1-shot broadcasting protocol.

Keywords: distributed algorithms, k -shot broadcasting, oblivious protocols, ad hoc radio networks.

1 Introduction

Energy efficiency has become a central issue in wireless networks, due to the constantly increasing use of autonomous devices with limited power resources. A lot of recent research focuses on how to accomplish communication tasks in an energy-efficient manner without compromising the system performance too much. Much of the work so far has been devoted to the problem of adjusting the transmission ranges of nodes so that the energy cost is minimized.

However, if nodes transmit at a fixed power level it makes sense to consider the number of transmissions as an energy consumption measure. Such a study was initiated by Gasieniec *et al.* [10], where broadcasting protocols with few transmissions ('shots') per node were considered for radio networks with known topology. Here, we study the problem in radio networks of in which nodes have no knowledge of the topology of the network.

We assume that a bound k is given and a node may transmit at most k times during the broadcasting

protocol (*k-shot broadcasting*); note that the bound k may well represent the number of transmissions that the power supply of a node can handle. We also assume that the communication is *synchronous*, that is, nodes may transmit or receive only at discrete time slots; each such simultaneous transmission is called a communication *step*. At each step a node may decide to act either as a *transmitter* or a *receiver*. Whenever a node transmits all its neighbors receive the message. If, however, two neighbors of a node v transmit simultaneously then a *collision* occurs and v hears nothing but noise, which is indistinguishable from background noise (therefore we assume that collision detection is not available).

We examine in particular the task of *broadcasting*. In the beginning, there is a unique node (the *source* node) which holds a message m , and the goal of a broadcasting protocol is to disseminate m to every node of the graph in a minimum number of steps.

We consider two types of protocols: *adaptive* and *oblivious* protocols; the former refers to protocols where a node may decide whether to transmit or not by taking into account any information it has received during the previous steps, while the latter term refers to protocols where a node makes transmission decisions with no consideration of the transmission history. Even though adaptive protocols are more powerful, oblivious algorithms are much easier to implement and demand minimal processing time for each node.

In this paper we mainly focus on oblivious protocols for k -shot broadcasting and study the way in which the limitation on the number of transmissions interacts with the time complexity the protocol. We also initiate a framework for extending our results to adaptive protocols.

Related work. Distributed broadcasting in radio networks of unknown topology with no limitation in the number of shots has been extensively studied in the literature.

The problem was first introduced by Chlamtac and Kutten [4]. Bar-Yehuda, Goldreich and Itai [1] gave the first randomized protocol, which completes broadcasting in $O(D \log n + \log^2 n)$ expected time when applied to graphs with n nodes and diameter D . Several papers followed [9, 14] that led to a tight upper bound of $O(D \log(n/D) + \log^2 n)$.

As for the deterministic case, a lower bound of $\Omega(n \log n)$ for general networks was given by Brusci and Del Pinto in [3], improved (for small D) to $\Omega(n \log D)$ by Clementi *et al.* [8]. Chlebus *et al.* [5] gave the first broadcasting protocol of sub-quadratic

*A short version of this paper appeared as a Brief Announcement in PODC 2010 [13]
Copyright ©2011, Australian Computer Society, Inc. This paper appeared at the 17th Computing: The Australasian Theory Symposium (CATS 2011), Perth, Australia, January 2011. Conferences in Research and Practice in Information Technology (CRPIT), Vol. 119, Alex Potanin and Taso Viglas, Ed. Reproduction for academic, not-for profit purposes permitted provided this text is included.

time complexity $O(n^{11/6})$. This bound was later improved to $O(n^{5/3} \log^3 n)$ by De Marco and Pelc [17] and then by Chlebus *et al.* [6], who gave an algorithm with time complexity $O(n^{3/2})$ based on finite geometries. Chrobak, Gąsieniec and Rytter [7] further improved the bound to $O(n \log^2 n)$. Finally, De Marco [16] gave the best currently known upper bound of $O(n \log n \log \log n)$, thus leaving a sub-logarithmic gap between the upper and lower bound.

It should be noted that all the aforementioned algorithms with time complexity better than $O(n^{3/2})$ are non-constructive; more specifically, the algorithms make use of combinatorial structures whose existence is shown via the probabilistic method. The best constructive bound so far is that of Indyk [11] who presented a somewhat slower constructive version of the protocol of [7], achieving $O(n \log^{O(1)} n)$ time complexity. It is also noteworthy that all algorithms proposed so far for deterministic distributed broadcasting in directed graphs are oblivious; this seems to be inherently related to the fact that the nodes have no knowledge about the graph topology.

For undirected networks, Chlebus *et al.* [5] gave a deterministic $O(n)$ -time broadcasting algorithm with spontaneous wake-up, that is, nodes may transmit even before receiving the source message. In case the nodes do not use spontaneous wake-up, an optimal $O(n \log n)$ -time broadcasting algorithm was presented in [14].

As mentioned above, broadcasting with a limited number of shots was first proposed in [10]. It has also been considered in [12], where randomized algorithms were proposed; in both cases, only broadcasting in known networks was studied. Another approach to limiting the number of shots was presented in [2], where the authors propose randomized algorithms which use few shots for each node and achieve nearly optimal broadcasting time.

Our contribution. To the best of our knowledge, the present paper is the first that addresses the issue of deterministic k -shot broadcasting in general radio networks of unknown topology.

We show (a) a lower bound of $\Omega(\frac{n^2}{k})$ on the broadcasting time of any oblivious k -shot broadcasting algorithm and (b) an oblivious broadcasting protocol that achieves a matching upper bound, namely $O(\frac{n^2}{k})$, for every $k \leq \sqrt{n}$ and an upper bound of $O(n^{3/2})$ for every $k > \sqrt{n}$. This bound implies the following tradeoff between time complexity and the number of maximum transmissions per node:

$$\#\text{shots} \times \#\text{steps} = \begin{cases} \Theta(n^2) & \text{for } k \leq \sqrt{n}, \\ \Omega(n^2) & \text{and} \\ O(k \cdot n^{3/2}) & \text{for } k > \sqrt{n}. \end{cases}$$

In order to prove the lower bound, we develop a technique which, given an oblivious protocol, constructs a graph (a *collision* graph) which succeeds in blocking the progress of broadcasting for sufficiently many steps. It should also be noted that the lower bound holds even in the case of oblivious broadcasting in symmetric (undirected) networks.

Our algorithm which matches the lower bound is based on the $O(n^{3/2})$ -time algorithm of [6]. An unexpected consequence of our results is that we may

impose a \sqrt{n} -shot restriction on the algorithm of [6] without affecting its performance.

Finally, for the case of 1-shot broadcasting, we prove that the lower bound $\Omega(n^2)$ holds even for adaptive algorithms that are as strong as possible (in the sense that they can make use not only of their own history but also of the complete history of other nodes).

2 Model and Preliminaries

Following [6], [5] we model a radio network as a directed graph. This means that if between two nodes u and v there exists an edge (u, v) but not the opposite edge (v, u) , then node u can transmit to node v , but not vice versa. Furthermore, we assume that the nodes have unique labels from the set $V = \{1, 2, \dots, n\}$, where n is the number of nodes in the network. Initially, a node is aware only of its own label and whether it is the source node or not. This means that it has no knowledge, full or partial, about the topology of the underlying graph. We also assume that every node knows the size n of the network.

We consider protocols under the assumption that a node may transmit only after it has received the source message, i.e. there are no *spontaneous transmissions*. Moreover, we assume that the nodes are not capable of *detecting collisions*, that is, if an attempt to transmit to a node v was unsuccessful, then v is not able to sense it.

We say that a broadcasting algorithm (or protocol) completes broadcasting when all nodes of the network have received the source message. For our purposes, the *running time* or *broadcasting time* of an algorithm is a function of the size of the network n and is defined as the worst-case number of steps needed to complete broadcasting over all possible network topologies of size n .

We now define the notion of *oblivious k -shot protocols*. As mentioned earlier, a protocol is oblivious if nodes do not take into account any information that may be gained during the execution of the protocol. Formally, an oblivious protocol can be succinctly described as a sequence of *transmission sets*, which are subsets of the node set V . We call such a sequence a *schedule*. Once a node receives the message at step t , it wakes up and transmits at the first k steps after t in which it appears in the transmission set. This model (in the unlimited shot sense) captures an important class of broadcasting algorithms, since most known algorithms for deterministic broadcasting in networks with unknown topology fall into this class ([6], [5], [7], [9], [16]).

Finally, we introduce some useful notation. We refer to a path graph S as a *chain* and denote by $V(S)$ the set of its nodes. For simplicity, we denote by $|S|$ the number of nodes in S . We also say that a graph G *starts* with a chain S when S is a subgraph of G , no node in S but the last is connected to nodes in $V \setminus V(S)$ and the source is the first node of the chain. We define the concatenation of two chains S_1 and S_2 , denoted by $S_1 \circ S_2$, as the graph consisting of S_1 and S_2 with the last node of S_1 connected to the first node of S_2 , and no other edge between S_1 and S_2 . We will also denote by $S \circ w$ the concatenation of chain S with the chain consisting of a single node w .

3 A Lower Bound for Oblivious k -shot Broadcasting

In this section we prove an $\Omega(\frac{n^2}{k})$ lower bound for any oblivious k -shot broadcasting protocol. We first need to introduce some more notation.

We denote by $\text{shots}(v, T)$ the minimum between k and the number of times v appears in a transmission set after step T . Note that $\text{shots}(v, T)$ is the number of times node v will transmit if it receives the message at step T . Let also $t_i(v, T)$, where $i \leq \text{shots}(v, T)$, be the step where node v appears for the i -th time in a transmission set after step T . Moreover, we define

$$t_{\leq i}(v, T) = \begin{cases} t_i(v, T) & \text{if } i \leq \text{shots}(v, T), \\ t_{\text{shots}(v, T)}(v, T) & \text{otherwise.} \end{cases}$$

We say that a sequence of nodes $S = \langle v_1, \dots, v_{|S|} \rangle$ occurs in a schedule \mathcal{S} if there is a subsequence $\mathcal{S}' = \langle \mathcal{T}_1, \dots, \mathcal{T}_{|S|} \rangle$ of transmission sets in \mathcal{S} such that for all $i = 1, \dots, |S|$, it holds that $v_i \in \mathcal{T}_i$. The first occurrence of sequence S after some step T is defined in a similar way, where in addition we ask for \mathcal{T}_1 to appear after step T and for $\mathcal{T}_{|S|}$ to appear as early as possible in the schedule. We denote by $t_1(S, T)$ the step where $\mathcal{T}_{|S|}$ appears.

Now, let us consider an oblivious k -shot broadcasting protocol \mathcal{P} . We will show that for any such protocol there is a graph $\mathcal{S}_{\mathcal{P}}$, which we explicitly construct, that delays the progress of broadcasting as much as the claimed bound. In order to show this, we will prove two lemmata.

Lemma 3.1. *Consider a sequence S and let $T = t_1(S, 0)$. Consider also any set $Q \subseteq V \setminus V(S)$. Then, there exists a node $w \in Q$ such that*

$$t_{\leq k}(w, T) \geq T + |Q| - 1$$

Proof. We first define a bipartite graph $G = (A, B, E)$ as follows. The upper set A corresponds to the nodes in the set $Q' = Q \setminus \{v_i\}$, where v_i is an arbitrary node in Q . Let $L = \max_{u \in Q'} t_{\leq k}(u, T)$, that is, the last step in which some node from Q' transmits if it receives the message at step T . We set node w to be the node that maximizes L . The lower set B corresponds to steps from $\{T + 1, \dots, L\}$ in which some node from Q' transmits. If node u transmits at step t , we add an edge between $u \in A$ and $t \in B$ (see Figure 1 for an example of the construction).




Figure 1: Example of the construction of the bipartite graph for a 2-shot protocol: nodes 1, 2, 3 are on the upper set A and the schedule after step T is $\{2, 4\}, \{3, 1\}, \{5\}, \{2, 1\}, \{3, 5, 4\}, \dots$. We have that $L = T + 5$.

We say that an induced subgraph H of G is *conflicting* if for any node $w \in V(H)$ the following two properties hold:

1. If $w \in A$, then all neighbors of w in G also belong to $V(H)$
2. If $w \in B$, then $\deg_H(w) > 1$

Note that in H , no vertex of B has only one neighbour, which means that all transmission sets included in H contain at least two nodes.

Let us now state and prove the following property: *G contains no conflicting subgraphs.*

Indeed, suppose that G has some conflicting subgraph $H = (A', B', E')$. Consider the graph G_t with the following topology: graph G_t starts with the chain corresponding to S and the last node v_S of S is connected to all nodes in A' . Moreover, v_t is connected to every node in A' and has no other neighbors. At step T , v_S transmits and all nodes in A' get the message. However, nodes in A' transmit only at steps in B' , according to the first property of H . Since every step in B' has at least two neighboring nodes in A' , the corresponding transmission set contains at least two nodes possessing the source message and therefore a conflict occurs at every such step. Moreover, by the end of step L , every node in A' either has transmitted k times or has no more transmissions available. Thus, v_t never gets the message, which is a contradiction.

Based on this property, it is easy to see that G has at least one vertex $u \in B$ with $\deg_G(u) = 1$ (otherwise G would be a conflicting subgraph itself). Suppose now that we remove u along with its *only* neighbor to obtain graph G' . Notice that G' is an induced subgraph of G and for any node $u \in A \cap G'$, all its neighbors belong to $V(G')$. Consequently, G' cannot be a conflicting subgraph, thus there exists some vertex $u' \in B$ such that $\deg_{G'}(u') = 1$. This process may continue $|Q| - 1$ times (since each time we remove at most one node from A), until all nodes in A are chosen. Intuitively, this process maps each node to a unique step in B . Thus, $|B| \geq |Q| - 1$, therefore $L \geq T + |Q| - 1$. \square

Lemma 3.2. *Consider a sequence S and let $T = t_1(S, 0)$. Then, there exists a sequence R of length at most k such that*

$$t_1(R, T) \geq T + n - |S| - k$$

Proof. Let $Q = V \setminus V(S)$. We first construct a set $W = \{w_1, w_2, \dots, w_k\} \subseteq Q$ of size k as follows. We apply lemma 3.1 to S and Q ; thus, there exists a node w_1 such that $t_{\leq k}(w_1, T) \geq T + |Q| - 1$. Generally, for any $i \leq k$, we apply lemma 3.1 to S and set $Q_i = Q \setminus \{w_1, \dots, w_{i-1}\}$, in order to obtain a node w_i such that $t_{\leq k}(w_i, T) \geq T + |Q_i| - 1 = T + |Q| - i$. Thus, the set W we obtain has the following property: for any $w \in W$, it holds that $t_{\leq k}(w, T) \geq T + |Q| - k$.

In the sequel, we will show that we can extend S using nodes in W as follows: we first choose the node r_1 which maximizes $t_{\leq 1}(r_1, T)$, then we choose the node r_2 which maximizes $t_{\leq 2}(r_2, T)$, and so on. Of course we need to take care of extreme cases, for example when all nodes except r_1 have $t_{\leq 2}(u, T) \leq t_1(r_1, T)$. We next show formally this claim.

We use the nodes from W to construct a sequence $R = \langle r_1, \dots, r_k \rangle$ of length k . We will allow that some places in R are empty by setting, e.g.

for the i -th place, $r_i = \varepsilon$. The construction proceeds as follows: in order to compute r_i , we calculate the value $M_i = \max_{v \in W \setminus \{r_1, \dots, r_{i-1}\}} \{t_{\leq i}(v, T)\}$. If $M_i \leq \max_{j < i} \{M_j\}$, then we set $r_i = \varepsilon$. Otherwise, we set r_i to be the node which maximizes M_i . Intuitively, r_i node with the latest i -th transmission among the remaining nodes of W .

We will show that $t_1(R, T) = \max_i \{M_i\}$. In order to prove this, we will prove by induction that $t_1(\langle r_1, \dots, r_i \rangle, T) = \max_{j \leq i} \{M_j\}$.

For the induction base, we observe that r_1 first appears at step M_1 , thus the proposition trivially holds. Now, let us examine the sequence $\langle r_1, \dots, r_{i+1} \rangle$. By the induction hypothesis, we know that $t_1(\langle r_1, \dots, r_i \rangle, T) = \max_{j \leq i} \{M_j\}$. We now distinguish two cases.

- $r_{i+1} = \varepsilon$: Then, we have $t_1(\langle r_1, \dots, r_{i+1} \rangle, T) = t_1(\langle r_1, \dots, r_i \rangle, T)$. Moreover, due to the construction, it holds that $M_{i+1} \leq \max_{j < i+1} \{M_j\}$. Thus, $\max_{j \leq i+1} \{M_j\} = \max_{j \leq i} \{M_j\}$ and, using the induction hypothesis, the proposition holds.
- $r_{i+1} \neq \varepsilon$: In this case, it holds that $t_{i+1}(r_{i+1}, T) = M_{i+1}$ and that $M_{i+1} > \max_{j < i+1} \{M_j\}$. Thus, $\max_{j \leq i+1} \{M_j\} = M_{i+1}$. Furthermore, the previous i occurrences of r_{i+1} are not later than step $\max_{j \leq i} \{M_j\}$ by construction. Consequently, $t_1(\langle r_1, \dots, r_{i+1} \rangle, T) = M_{i+1} = \max_{j \leq i+1} \{M_j\}$.

Now, let us consider the last node r_j of R such that $r_j \neq \varepsilon$. Clearly, it holds that $\forall i \neq j : M_j \geq M_i$ and that $t_1(R, T) = \max_i \{M_i\} = M_j = t_{\leq j}(r_j, T)$. In the case where $j = k$, we have that $t_{\leq j}(r_j, T) = t_{\leq k}(r_k, T)$. Otherwise, since $r_i = \varepsilon$ for $i > j$, no node in $\{w_{j+1}, \dots, w_k\}$ occurs after step $t_{\leq j}(r_j, T)$, thus $t_{\leq j}(r_j, T) \geq t_{\leq k}(p, T)$ for any $p \in \{w_{i+1}, \dots, w_k\}$.

In any case, there exists a node $q \in W$ such that $t_1(R, T) \geq t_{\leq k}(q, T)$. However, $t_{\leq k}(q, T) \geq T + |Q| - k$ by construction of set W . Consequently, $t_1(R, T) \geq T + |Q| - k = T + n - |S| - k$. \square

We now use repeatedly lemma 3.2 in order to prove the main result of this section.

Theorem 3.3. *For any oblivious k -shot broadcasting protocol \mathcal{P} , there exists a graph $S_{\mathcal{P}}$ where \mathcal{P} needs $\Omega(\frac{n^2}{k})$ steps to complete broadcasting.*

Proof. Consider the source node v_0 . W.l.o.g. we may assume that v_0 appears in the first transmission set and thus $t_1(\langle v_0 \rangle, 0) = 1$. Then, from lemma 3.2 with $S = \langle v_0 \rangle$, there exists a sequence S_1 of length at most k such that $T_1 = t_1(\langle v_0 \rangle \circ S_1, 0) \geq n - k$.

Now, we may apply lemma 3.2 with $S = \langle v_0 \rangle \circ S_1$ so as to find a sequence S_2 , where $T_2 = t_1(\langle v_0 \rangle \circ S_1 \circ S_2, 0) \geq T_1 + n - 2k - 1$ and S_2 has length at most k . We may continue this process until all nodes have been chosen. Thus, we can construct a sequence $S_{\mathcal{P}}$ of nodes which occurs for the first time at step

$$\sum_{j=1}^{\lfloor \frac{n}{k} \rfloor} (n - jk - 1) + \Omega(1) = \Omega(\frac{n^2}{k})$$

The chain $S_{\mathcal{P}}$ corresponding to this sequence S is the claimed graph. \square

4 An Oblivious Algorithm for k -shot Broadcasting

We will present an oblivious algorithm (OBLIVIOUS k -SHOT) which is based on the algorithm presented in [6] and performs optimal k -shot broadcasting in time $O(\frac{n^2}{k})$ for any $k \leq \sqrt{n}$. For $k \geq \sqrt{n}$, the algorithm completes broadcasting in $O(n^{3/2})$ steps, matching the time performance of the algorithm in [6].

Let p be the smallest prime greater than or equal to $\lceil \sqrt{n} \rceil$. We map a node with label i to the point $\langle i \text{ div } p, i \text{ mod } p \rangle$. A line $L_{a,b}$ with direction a ($a = 0, \dots, p$) and offset b ($b = 0, \dots, p-1$) is defined as the following set of points:

$$L_{a,b} = \begin{cases} \{\langle x, y \rangle : x \equiv b \pmod{p}\} & \text{if } a = p, \\ \{\langle x, y \rangle : y \equiv a \cdot x + b \pmod{p}\} & \text{otherwise.} \end{cases}$$

It is easy to observe that the sets defined have the following useful properties, which will be crucial in analyzing the running time of the algorithm.

- There are p disjoint lines in each direction, one for each offset.
- The total number of distinct lines is $p \cdot (p+1)$.
- Each line contains exactly p nodes.
- Each node belongs to $p+1$ lines, one in each direction.
- Two lines of different directions have exactly one node in common.
- For any two different nodes, there is exactly one line that contains both of them.

The algorithm multiplexes two different procedures, the classic ROUND-ROBIN procedure where nodes transmit alone, one after the other (transmission sets as singletons), and the LINE-TRANSMIT procedure, where lines are used as transmission sets. Note that we do not need to set specific termination conditions, because the k -shot restriction guarantees termination.

Procedure ROUND-ROBIN

```

while not finished do
  for  $v = 1, 2, \dots, p^2$  do
    | node  $v$  transmits
  end
end

```

Procedure LINE-TRANSMIT

```

while not finished do
  for  $a = 0, \dots, p$  do
    /* STAGE */
    for  $b = 0, \dots, p-1$  do
      | all nodes in  $L_{a,b}$  transmit
    end
  end
end
end

```

We define $K = \lceil \frac{p}{k-2} \rceil$ and the procedures are multiplexed such that a step of the LINE-TRANSMIT is followed by K steps of the ROUND-ROBIN procedure (see Figure 2).

Algorithm 3: OBLIVIOUS k -SHOT

```
 $K = \lceil \frac{p}{k-2} \rceil ;$   
while not finished do  
  | perform a LINE-TRANSMIT step ;  
  | perform  $K$  ROUND-ROBIN steps ;  
end
```




Figure 2: The multiplexing of the procedures LINE-TRANSMIT and ROUND-ROBIN.

Before we analyze the time complexity of the algorithm, we need to examine whether the OBLIVIOUS k -SHOT algorithm will accomplish broadcasting in any graph. Notice that it suffices to prove that every node transmits alone at least once.

Lemma 4.1. *In OBLIVIOUS k -SHOT, every node transmits at least once as the only node in the network.*

Proof. Assume that node v receives the message at some step t . We will show that v transmits at least once at some step of the ROUND-ROBIN procedure. In order to obtain a contradiction, let us assume that v transmits all k times during the LINE-TRANSMIT procedure. Since each node transmits exactly once during a stage, this means that v transmits at k consecutive stages.

We will now compute the number of ROUND-ROBIN steps during these k stages. Note that each stage has p line transmissions and after each line transmission K ROUND-ROBIN steps are performed. Thus, we have at least $(p \cdot (k-2) + 1) \cdot K > p \cdot (k-2) \cdot \lceil \frac{p}{k-2} \rceil \geq p^2$ ROUND-ROBIN steps. Clearly, this means that v would appear at least once at a ROUND-ROBIN step during this period, which yields a contradiction. \square

Let us now examine the time complexity of this broadcasting protocol. We say that we make one *unit of progress* in two cases: (a) each time a node receives the message (in this case we say that the node becomes *active*) and (b) when all neighbors of a node receive the message (in this case that the node becomes *passive*). Achieving progress $2n-1$ means that every node has received the message and thus broadcasting is complete.

We next prove a useful lemma. Let $r < k$ be a constant which we will fix later so as to optimize the analysis.

Lemma 4.2. *Let F be the set of active nodes at the beginning of stage c . If $|F| \leq k/r$, then the average progress per stage during the next $r \cdot |F|$ stages is constant.*

Proof. Let us consider the set of consecutive stages $C = \{c, c+1, \dots, c+r \cdot |F| - 1\}$. If every node in

F becomes passive during C , then the progress is at least $|F|$, and thus the average progress per stage is at least $\frac{|F|}{r \cdot |F|} = \frac{1}{r}$.

Otherwise, there exists some node $v \in F$ that remains active after C . Note first that v has not completed its k shots during C , for otherwise by lemma 4.1, it would have broadcasted alone and thus would have become passive. Therefore, node v transmits during all stages of C .

Then, there must exist $(r-1) \cdot |F|$ stages during which v transmits as the only node from F . This is because $r \cdot |F| \leq k \leq \sqrt{n} \leq p$ and hence any node from F may collide at most once with v during C (due to line properties). This means that there must exist at least $(r-1) \cdot |F|$ nodes outside F that collide with v during C . All these nodes have obtained the message during C , which implies an average progress per stage of at least $\frac{(r-1) \cdot |F|}{r \cdot |F|} = 1 - \frac{1}{r}$.

Thus, the average progress per stage is at least $\min\{\frac{1}{r}, 1 - \frac{1}{r}\}$, which is a constant. \square

It is easy to see that the average progress is maximized when $r = 2$. Then, we have that if $|F| \leq k/2$, the average progress per stage is at least $1/2$.

Theorem 4.3. *OBLIVIOUS k -SHOT completes k -shot broadcasting in $O(\frac{n^2}{k})$ steps for $k \leq \sqrt{n}$ and $O(n^{3/2})$ steps for $k > \sqrt{n}$.*

Proof. We will calculate the time needed to make progress $2n-1$. Clearly, if at any stage the number of active nodes is more than $k/2$, then the ROUND-ROBIN procedure guarantees that the progress over the next $O(n)$ steps is at least $k/2$. Thus, the average progress per step is $\Omega(k/n)$. Otherwise, the number of active nodes will be at most $k/2$ and thus the progress per stage will be at least $1/2$ by using lemma 4.2 with $r = 2$. Since each stage, together with the corresponding ROUND-ROBIN transmissions, has $p \cdot (K+1) = O(n/k)$ steps, the average progress per step is again $\Omega(k/n)$. Consequently, we have that the total time complexity will be $O(\frac{n^2}{k})$.

For any number of shots greater than \sqrt{n} , one may use OBLIVIOUS k -SHOT restricted to $k = \sqrt{n}$ shots, thus obtaining an $O(n^{3/2})$ broadcasting algorithm. \square

5 Beyond Oblivious Broadcasting

In this final section, we present an $\Omega(n^2)$ lower bound which holds for any adaptive algorithm for the special case of 1-shot broadcasting. The bound holds not only for directed, but also for symmetric graphs.

We first need a formal definition of an *adaptive* broadcasting protocol. We will use the model proposed by Kowalski and Pelc [15]. We denote by $H_t(v)$ the *message history* of node v until the end of step t . Specifically, $H_t(v)$ is a sequence of pairs (M_1, M_2, \dots, M_t) . M_1 is either the pair $(0, \emptyset)$ (we call this *empty history*) or the pair (s, m) , where m is the message of the source node and s is the label of the source node. If node v does not receive a message at step i , then $M_i = (0, \emptyset)$. Otherwise, M_i is the pair $(w, H_{i-1}(w))$, where w is the label of the node from which node v received a message at step i .

A broadcasting protocol can now be defined by a function $\pi(v, t, H_{t-1}(v))$, which takes values in the

set $\{\text{receive}, \text{send}\}$. The function decides whether node v with message history $H_{t-1}(v)$ acts as a *receiver* (**receive**) or as a *transmitter* (**send**) at step t . If v acts as a transmitter, it sends the message $(v, H_{t-1}(v))$. Note that it is enough to restrict our attention to sending only the entire history of the transmitter, since the receiver can always deduce any information from the received history.

Let us consider any deterministic 1-shot broadcasting protocol \mathcal{P} which completes broadcasting in any graph with n nodes. We construct a graph $G_{\mathcal{P}}$ such that broadcasting is slowed down as much as possible. We start by considering the family \mathcal{G}_0 of all possible connected graphs with n nodes. The construction proceeds by considering the steps of protocol \mathcal{P} , at each step refining the family of graphs. We will show that by the end of the construction, we are left with a graph in which \mathcal{P} completes broadcasting in $\Omega(n^2)$ steps.

We divide the construction into $n - 3$ phases and denote by \mathcal{G}_i the family of graphs by the end of phase i . The construction is based on the following lemma.

Lemma 5.1. *Assume that the family \mathcal{G}_{i-1} includes only graphs which start with a chain S with i nodes ($i \leq n - 3$) and the last node v_S of S transmits no earlier than step T . Then, there exists a node $w \in \bar{S} = V \setminus V(S)$ such that the family \mathcal{G}_i includes only graphs which start with chain $S \circ w$ and w transmits no earlier than step $T + n - i$.*

Proof. For simplicity, we set $\mathcal{H}_0 = \mathcal{G}_{i-1}$. We also denote by $\mathcal{H}_0^w \subseteq \mathcal{H}_0$ the family of graphs in \mathcal{H}_0 where node w is connected to v_S . Without loss of generality, let us assume that node v_S transmits its history $H_{T-1}(v_S)$ at step T . In any graph $G \in \mathcal{H}_0^w$, node w receives the same history $H_{T-1}(v_S)$ at step T and has received only empty history before step T . Thus, for any graphs $G_a, G_b \in \mathcal{H}_0^w$, it holds that $H_T(w, G_a) = H_T(w, G_b) = H_T(w)$. Since protocol \mathcal{P} determines the action of any w at step $T + 1$ from $\pi(w, T + 1, H_T(w))$, the action of any node w is the same for any graph in \mathcal{H}_0^w and we denote it by $\pi_S(w)$.

Assume that for $w_i, w_j \in \bar{S}$ it holds that $\pi_S(w_i) = \pi_S(w_j) = \text{send}$. Then, consider any graph $G_{ij} \in \mathcal{H}_0^{w_i} \cap \mathcal{H}_0^{w_j}$ such that w_i and w_j are the only nodes connected to a node v_t . At step $T + 1$, the nodes w_i and w_j transmit simultaneously, a collision occurs and thus v_t never gets the message, a contradiction (see Figure 3). Thus, there exists at most one node $w_1 \in \bar{S}$ such that $\pi_S(w_1) = \text{send}$. In this case, we refine the family of graphs to the family $\mathcal{H}_1 = \mathcal{H}_0 \setminus \mathcal{H}_0^{w_1}$ and set $\bar{S}_1 = \bar{S} \setminus \{w_1\}$. Otherwise, when no node from \bar{S} transmits, we set $\mathcal{H}_1 = \mathcal{H}_0$ and $\bar{S}_1 = \bar{S}$.




Figure 3: A collision graph: a conflict occurs when nodes w_i and w_j transmit simultaneously. Since v_t has no other neighbors and the protocol is 1-shot, v_t never gets the message.

Clearly, for every graph in \mathcal{H}_1 , no node from \bar{S} transmits at step $T + 1$. This means that only nodes

from $V(S)$ transmit at step $T + 1$. Consequently, all the nodes in \bar{S}_1 may receive a message only from v_S and thus the history (and action) of any node $w \in \bar{S}_1$ is the same for any graph in \mathcal{H}_1^w . Using the same argument as before, there exists at most one node w_2 such that $\pi(w_2, T + 2, H_{T+1}(w_2)) = \text{send}$. Then, we further refine the family of graphs to $\mathcal{H}_2 = \mathcal{H}_1 \setminus \mathcal{H}_1^{w_2}$ and set $\bar{S}_2 = \bar{S}_1 \setminus \{w_2\}$ (or $\mathcal{H}_2 = \mathcal{H}_1$, $\bar{S}_2 = \bar{S}_1$ if no node transmits at $T + 2$). We may apply this argument repeatedly at least $n - i - 1$ times, each time further refining the family of graphs. The final family $\mathcal{G}_i = \mathcal{H}_i$ ($i \geq n - i - 1$) consists of graphs where the only node of $\bar{S}_i = \{w_{i+1}\}$ is connected to v_S ; thus, every graph in \mathcal{G}_i starts with the chain $S \circ w_{i+1}$. Moreover, w_{i+1} does not transmit earlier than step $T + i + 1$ in any graph of \mathcal{G}_i . \square

Theorem 5.2. *For any 1-shot broadcasting protocol \mathcal{P} , there exists a graph $G_{\mathcal{P}}$ where \mathcal{P} needs $\Omega(n^2)$ steps to complete broadcasting.*

Proof. We will prove the lower bound using induction on the number of phases. Specifically, we will show that family \mathcal{G}_i includes only graphs starting with a chain of $i + 1$ nodes and that the last node of the chain does not transmit before step $1 + \sum_{j=1}^i (n - j)$.

At the end of phase 0, the chain consists only of the source node, which transmits at step 1. Thus, the claim holds trivially. Using the induction hypothesis, we know that the family \mathcal{G}_i includes only graphs which start with a chain of $i + 1$ nodes and the last node transmits no earlier than step $1 + \sum_{j=1}^i (n - j)$. Applying lemma 5.1, the construction refines the family \mathcal{G}_i to \mathcal{G}_{i+1} , where every graph of \mathcal{G}_{i+1} starts with a chain of $(i + 1) + 1$ nodes and the last node transmits no earlier than step $1 + \sum_{j=1}^i (n - j) + (n - i - 1) = 1 + \sum_{j=1}^{i+1} (n - j)$.

After phase $n - 3$, the family \mathcal{G}_{n-3} includes only three graphs (for each configuration of the remaining two nodes). In one of the graphs (graph $G_{\mathcal{P}}$), the protocol needs one more step to complete broadcasting. Thus, in $G_{\mathcal{P}}$ the protocol completes broadcasting no earlier than step

$$1 + \sum_{i=1}^{n-3} (n - i) + 1 = \frac{n(n-1)}{2} - 1$$

\square

The proof of this theorem is constructive. Thus, for any 1-shot broadcasting protocol \mathcal{P} , we can actually construct a graph G where \mathcal{P} needs at least $\Omega(n^2)$ steps to complete broadcasting.

6 Conclusions

In this paper, we initiate the study of deterministic k -shot broadcasting in radio networks with unknown topology. We manage to show an exact energy-time tradeoff for values of $k \leq \sqrt{n}$. It remains an open question of whether it is possible to match the lower bound of $\Omega(\frac{n^2}{k})$ for $k > \sqrt{n}$.

It is also interesting to examine whether the lower bound can also be generalized to hold for adaptive k -shot protocols for any value of k . Finally, it would be desirable to drop the requirement that the number of nodes n is known to the nodes. It seems that the

standard doubling technique cannot work in this case, because of the k -shot restriction.

References

- [1] Reuven Bar-Yehuda, Oded Goldreich, and Alon Itai. On the time-complexity of broadcast in radio networks: An exponential gap between determinism and randomization. In *PODC*, pages 98–108, 1987.
- [2] Petra Berenbrink, Colin Cooper, and Zengjian Hu. Energy efficient randomised communication in unknown adhoc networks. *Theor. Comput. Sci.*, 410(27-29):2549–2561, 2009.
- [3] Danilo Bruschi and Massimiliano Del Pinto. Lower bounds for the broadcast problem in mobile radio networks. *Distributed Computing*, 10(3):129–135, 1997.
- [4] I. Chlamtac and S. Kutten. On broadcasting in radio networks—problem analysis and protocol design. *IEEE Transactions on Communications*, 33(12):1240–1246, dec. 1985.
- [5] Bogdan S. Chlebus, Leszek Gasieniec, Alan Gibbons, Andrzej Pelc, and Wojciech Rytter. Deterministic broadcasting in unknown radio networks. In *SODA '00: Proceedings of the eleventh annual ACM-SIAM symposium on Discrete algorithms*, pages 861–870, Philadelphia, PA, USA, 2000. Society for Industrial and Applied Mathematics.
- [6] Bogdan S. Chlebus, Leszek Gasieniec, Anna Östlin, and John Michael Robson. Deterministic radio broadcasting. In Ugo Montanari, José D. P. Rolim, and Emo Welzl, editors, *ICALP*, volume 1853 of *Lecture Notes in Computer Science*, pages 717–728. Springer, 2000.
- [7] Marek Chrobak, Leszek Gasieniec, and Wojciech Rytter. Fast broadcasting and gossiping in radio networks. In *FOCS*, pages 575–581, 2000.
- [8] Andrea E. F. Clementi, Angelo Monti, and Riccardo Silvestri. Distributed broadcast in radio networks of unknown topology. *Theor. Comput. Sci.*, 302(1-3):337–364, 2003.
- [9] Artur Czumaj and Wojciech Rytter. Broadcasting algorithms in radio networks with unknown topology. In *FOCS*, pages 492–501. IEEE Computer Society, 2003.
- [10] Leszek Gasieniec, Erez Kantor, Dariusz R. Kowalski, David Peleg, and Chang Su. Time efficient k -shot broadcasting in known topology radio networks. *Distributed Computing*, 21(2):117–127, 2008.
- [11] Piotr Indyk. Explicit constructions of selectors and related combinatorial structures, with applications. In *SODA '02: Proceedings of the thirteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 697–704, Philadelphia, PA, USA, 2002. Society for Industrial and Applied Mathematics.
- [12] Erez Kantor and David Peleg. Efficient k -shot broadcasting in radio networks. In Idit Keidar, editor, *DISC*, volume 5805 of *Lecture Notes in Computer Science*, pages 481–495. Springer, 2009.
- [13] Paraschos Koutris and Aris Pagourtzis. Brief announcement: k -shot distributed broadcasting in radio networks. In Andréa W. Richa and Rachid Guerraoui, editors, *PODC*, pages 77–78. ACM, 2010.
- [14] Dariusz R. Kowalski and Andrzej Pelc. Broadcasting in undirected ad hoc radio networks. In *PODC '03: Proceedings of the twenty-second annual symposium on Principles of distributed computing*, pages 73–82, New York, NY, USA, 2003. ACM.
- [15] Dariusz R. Kowalski and Andrzej Pelc. Time of deterministic broadcasting in radio networks with local knowledge. *SIAM J. Comput.*, 33(4):870–891, 2004.
- [16] Gianluca De Marco. Distributed broadcast in unknown radio networks. In Shang-Hua Teng, editor, *SODA*, pages 208–217. SIAM, 2008.
- [17] Gianluca De Marco and Andrzej Pelc. Faster broadcasting in unknown radio networks. *Inf. Process. Lett.*, 79(2):53–56, 2001.