

# The Cost of Global Broadcast in Dynamic Radio Networks\*

Mohamad Ahmadi<sup>1</sup>, Abdolhamid Ghodselahi<sup>2</sup>, Fabian Kuhn<sup>3</sup>, and Anisur Rahaman Molla<sup>4</sup>

- 1 Department of Computer Science, University of Freiburg, Freiburg, Germany  
mahmadi@cs.uni-freiburg.de
- 2 Department of Computer Science, University of Freiburg, Freiburg, Germany  
hghods@cs.uni-freiburg.de
- 3 Department of Computer Science, University of Freiburg, Freiburg, Germany  
kuhn@cs.uni-freiburg.de
- 4 Department of Computer Science, University of Freiburg, Freiburg, Germany  
armolla@cs.uni-freiburg.de

---

## Abstract

We study the single-message broadcast problem in dynamic radio networks. We show that the time complexity of the problem depends on the amount of stability and connectivity of the dynamic network topology and on the adaptiveness of the adversary providing the dynamic topology. More formally, we model communication using the standard graph-based radio network model. To model the dynamic network, we use a variant of the synchronous dynamic graph model introduced in [Kuhn et al., STOC 2010]. For integer parameters  $T \geq 1$  and  $k \geq 1$ , we call a dynamic graph  $T$ -interval  $k$ -connected if for every interval of  $T$  consecutive rounds, there exists a  $k$ -vertex-connected stable subgraph. Further, for an integer parameter  $\tau \geq 0$ , we say that the adversary providing the dynamic network is  $\tau$ -oblivious if for constructing the graph of some round  $t$ , the adversary has access to all the randomness (and states) of the algorithm up to round  $t - \tau$ .

As our main result, we show that for any  $T \geq 1$ , any  $k \geq 1$ , and any  $\tau \geq 1$ , for a  $\tau$ -oblivious adversary, there is a distributed algorithm to broadcast a single message in time  $O\left(\left(1 + \frac{n}{k \cdot \min\{\tau, T\}}\right) \cdot n \log^3 n\right)$ . We further show that even for large interval  $k$ -connectivity, efficient broadcast is not possible for the usual adaptive adversaries. For a 1-oblivious adversary, we show that even for any  $T \leq (n/k)^{1-\varepsilon}$  (for any constant  $\varepsilon > 0$ ) and for any  $k \geq 1$ , global broadcast in  $T$ -interval  $k$ -connected networks requires at least  $\Omega(n^2/k^2 \log n)$  time. Further, for a 0-oblivious adversary, broadcast cannot be solved in  $T$ -interval  $k$ -connected networks as long as  $T < n - k$ .

**1998 ACM Subject Classification** F.2.2 Analysis of Algorithms and Problem Complexity – Non-numerical Algorithms and Problems, G.2.2 Discrete Mathematics – Graph Theory

**Keywords and phrases** radio network, dynamic network, global broadcast, interval connectivity, hitting game

**Digital Object Identifier** 10.4230/LIPIcs.OPODIS.2015.7

## 1 Introduction

By now, a rich theory on algorithms for large-scale wireless networks exists and we have a rather precise understanding of the complexity of many basic computation and communication

---

\* Research supported by ERC Grant No. 336495 (ACDC).



tasks for a variety of wireless network models. While many wireless communication models and modeling assumptions have been studied, to a large part, the considered models all share one basic property. Most of the existing work is based on static networks and on communication models where wireless signal reception is modeled in a completely deterministic way. For example, in the classic radio network model, a wireless network is modeled as a graph and a node in the graph can receive a message transmitted by some neighbor if and only if no other neighbor transmits at the same time, e.g., [5, 21]. In the SINR (or physical) model, nodes have fixed coordinates in some geometric space and a transmitted signal can be successfully received if and only if the signal-to-noise-and-interference ratio at the receiver is above a certain fixed threshold, e.g., [14, 22].

The situation in actual networks however is quite different and wireless signal reception might behave in a rather unpredictable way. There can be multiple sources for interference which cannot be controlled by a distributed algorithm and signal propagation depends on various properties of the environment. As a result, we often obtain wireless communication links with unreliable behavior [16, 24, 26, 27, 28]. In addition, wireless devices might be mobile leading to a potentially completely dynamic network topology.

As a consequence, in recent years, researchers in the wireless algorithms community have also started to consider radio network models which exhibit nondeterministic behavior and sometimes general dynamic topologies, e.g., [9, 11, 12, 2, 17, 13]. In the present paper, we continue this line of research and study the global broadcast problem in dynamic radio networks for a range of modeling assumptions. Note that in ordinary, static radio networks, albeit appearingly simple, global broadcast is one of the best studied problems in the area, (see, e.g., [4, 5, 10, 15, 23] and many others). We model a dynamic network by applying the dynamic network model introduced in [18]. Time is divided into synchronous rounds and a wireless network is modeled as a dynamic graph with a fixed set of  $n$  nodes and a set of edges which can change from round to round. For two parameters  $T \geq 1$  and  $k \geq 1$ , a dynamic graph is called  $T$ -interval  $k$ -connected if for any interval of  $T$  consecutive rounds, the set of edges which are present throughout these  $T$  rounds induces a graph with vertex connectivity at least  $k$  (in [18], the model was only introduced for  $k = 1$ ). We refer to [19] for a more thorough discussion of the model of [18] and of several earlier related dynamic network models (e.g., [3, 6, 11, 25]).

Communication is modeled by using the standard radio network model. In each round, each node can either transmit a message or listen. A listening node successfully receives a message transmitted by a neighbor in the current graph if and only if no other neighbor transmits in the same round. We assume that nodes cannot detect collisions, i.e., whether 0 or more than 1 neighbors transmit is indistinguishable for a listening node. Note that the described dynamic network model does not only allow to model topology changes due to arbitrary node mobility. It also allows to model unreliable links where the presence/availability can change for various reasons.

We assume that the dynamic graph is provided by a worst-case adversary. As we study randomized distributed protocols, we need to specify to what extent the adversary can adapt to the random decisions of the nodes when determining the sequence of network topologies. For the adaptiveness of the adversary, we use a more fine-grained classification than what is usually done. For an integer parameter  $\tau \geq 0$ , we say that the adversary is  $\tau$ -oblivious if for determining the graph in round  $r$ , the adversary knows the randomness of all nodes of all the rounds up to round  $r - \tau$ . Typically, only the extreme cases are studied. An adversary which does not have access to the random decisions of the algorithm (i.e.,  $\tau = \infty$ ) is called an *oblivious* adversary, whereas an adversary which has access to the randomness of the

algorithm is called an *adaptive* adversary. If the adversary even has access to the randomness of the current round ( $\tau = 0$ ), it is called *strongly adaptive*, otherwise ( $\tau = 1$ ), it is called *weakly adaptive*. For more precise formal definitions of the modeling assumptions, we refer to Section 2.

In our paper, we consider the problem of broadcasting a single message from a source node to all the nodes of a dynamic network. The most relevant previous work in the context of the present work appeared in [11, 13, 17]. In [11], it is shown that in 1-interval 1-connected networks (i.e., the graph is connected in every round)<sup>1</sup>, the complexity of global broadcast for a 1-oblivious adversary is  $\Theta(n^2/\log n)$ . In [17] and [13],  $\infty$ -interval 1-connected graphs are considered (i.e., there is a stable connected subgraph which is present throughout the whole execution). In [17], it is shown that even for a 0-oblivious adversary, it is possible to solve broadcast in  $O(n \log^2 n)$  rounds and it is shown that  $\Omega(n)$  rounds are necessary even if the stable connected subgraph has diameter 2. In [13], it is shown that when only assuming an  $\infty$ -oblivious adversary, the running time can be improved to  $O((D + \log n) \log n)$ , where  $D$  is the diameter of the stable connected subgraph. Note that in this case, the algorithm in [13] achieves essentially the same time complexity as is possible in static graphs of diameter  $D$  [5, 20, 23]. In [13], it is also shown that for a 1-oblivious adversary,  $\Omega(n/\log n)$  rounds are necessary even for  $D = 2$ .

## 1.1 Contributions

In the following, we state the results of the paper. For formal details regarding problem statement and modeling not specified in the introduction, we refer to Section 2. Our main result is a randomized broadcast algorithm for the described dynamic radio network model. The algorithm (and also partly its analysis) is based on a combination of the techniques used in [11] and [17]. We prove the following main theorem.

► **Theorem 1.** *Let  $T \geq 1$ ,  $\tau \geq 1$ , and  $k \geq 1$  be positive integer parameters. Assume that the adversary is  $\tau$ -oblivious. Then, in a dynamic  $T$ -interval  $k$ -connected  $n$ -node radio network, with high probability, single message broadcast can be solved in time*

$$O\left(\left(1 + \frac{n}{k \cdot \min\{\tau, T\}}\right) \cdot n \log^3 n\right).$$

► **Remark.** Note that for small and for large values of  $\min\{\tau, T\}$ , one can do slightly better. It is straightforward to generalize the broadcast algorithm of [11] to complete single message broadcast in time  $O(n^2/k \log n)$  in 1-interval  $k$ -connected radio networks against a 1-oblivious adversary. Using the result from [17], we also know that for a sufficiently large constant  $c$  and  $T \geq cn \log^2 n$ , single-message broadcast can be solved in  $O(n \log^2 n)$  rounds even for  $\tau = 0$ . Our upper bound therefore beats previous results for  $\min\{\tau, T\} = \omega(\log^4 n)$  and  $T = O(n \log^2 n)$ .

In addition to the upper bound of Theorem 1, we also prove a lower bound which essentially shows that even for very large values of  $T$ , some relaxation on the standard adaptive adversaries is necessary in order to get an upper bound which improves with  $T$ . For  $\tau = 1$ , we show that at least for small  $k$ , the generalized upper bound of [11] is essentially optimal. The lower bound can be seen as a generalization of the simple  $\Omega(n^2/\log n)$  lower bound for  $k = 1$  and  $T = 1$  proven in [11].

<sup>1</sup> In [11], the connectivity condition on the dynamic network is phrased differently and slightly more general.

■ **Table 1** An overview over the existing bounds on global broadcast in the dynamic radio network model. The results marked in bold are the results of the present paper. For the  $T = \infty$  results,  $D$  refers to the diameter of the stable subgraph.

interval conn.	vertex conn.	adversary	complexity
$T = 1$	$k = 1$	$\tau = 1$	$\Theta(n^2/\log n)$ [11]
$T = \infty$	$k = 1$	$\tau = 0$	$O(n \log^2 n) / \Omega(n)$ , $D=2$ [17]
$T = \infty$	$k = 1$	$\tau = \infty$	$O((D + \log n) \cdot \log n)$ [13]
$T = \infty$	$k = 1$	$\tau = 1$	$\Omega(n/\log n)$ , $D=2$ [13]
$T \geq 1$	$k \geq 1$	$\tau \geq 1$	$O\left(\left(1 + \frac{n}{k \cdot \min\{\tau, T\}}\right) \cdot n \log^3 n\right)$
$T \leq (n/k)^{1-\varepsilon}$	$k \geq 1$	$\tau = 1$	$\Omega(n^2/(k^2 \log n))$
$T < n - k$	$k \geq 1$	$\tau = 0$	<i>impossible</i>

► **Theorem 2.** For every constant  $\varepsilon > 0$  and every  $T \leq (n/k)^{1-\varepsilon}$ , the expected time to solve single-message broadcast in  $T$ -interval  $k$ -connected radio networks against a 1-oblivious adversary is at least

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

In addition, we show that unless the interval connectivity is very large, single-message broadcast cannot be solved in the presence of a strongly adaptive (0-oblivious) adversary.

► **Theorem 3.** For any  $k \geq 1$  and any  $T < n - k$ , it is not possible to solve single-message broadcast in  $T$ -interval  $k$ -connected radio networks against a 0-oblivious adversary.

The discussion of the above result appears in Section 5. We note that the above theorem is tight in the following sense. As soon as  $T \geq n - k$ , global broadcast can be solved (with potentially exponential time complexity) and as soon as  $T = cn \log^2 n$  for a sufficiently large constant  $c$ , we know from [17] that it can even be solved in time  $O(n \log^2 n)$ . All results, as well as a comparison with previous work are summarized in Table 1.

► **Remark.** In [18], interval connectivity was introduced to (in particular) study the problem of broadcasting multiple messages in a dynamic network in a standard message passing model. It is shown that interval connectivity  $T$  allows to essentially speed up multi-message broadcast by a factor of  $T$ . We find it interesting that when considering a radio network model, interval connectivity seems to provide a similar speed-up, even for broadcasting a single message. Something similar also holds for graphs with large vertex connectivity. In [7, 8], it is shown that even on static graphs, vertex connectivity  $k$  allows to speed up multi message broadcast by essentially a factor  $k$ . Here, we show that a similar speed up can be obtained in radio networks even for broadcasting a single message.

## 2 Model and Problem Definition

**Dynamic Network.** As described in Section 1, we adapt the synchronous dynamic network model of [18] to model dynamic networks.<sup>2</sup> Time is divided into rounds such that for all  $r \geq 1$ , round  $r$  starts at time  $r - 1$  and ends at time  $r$ . A dynamic network is given by a

<sup>2</sup> Similar dynamic network models have also been used before [18], for example in [3, 11, 25]. For additional references and a thorough discussion, we refer to [19].

sequence of undirected graphs  $\langle G_1, G_2, \dots \rangle$ , where  $G_r = (V, E_r)$  is a static graph representing the network topology in round  $r$ . The node set  $V$  is a set of  $n$  nodes corresponding to the wireless devices in the network and the edge set  $E_r$  is the set of active communication links in round  $r$ . A dynamic graph  $\langle G_1, G_2, \dots \rangle$  is called  $T$ -interval  $k$ -connected for integer parameters  $T \geq 1$  and  $k \geq 1$  if and only if for all  $r \geq 1$ , the graph

$$\tilde{G}_{r,T} = (V, \bar{E}_{r,T}), \quad \bar{E}_{r,T} := \bigcap_{r'=r}^{r+T-1} E_{r'}$$

is a graph with vertex connectivity at least  $k$ .

**Communication Model.** An  $n$ -node distributed algorithm  $\mathcal{A}$  is defined by  $n$  randomized processes which are assigned to the nodes of the dynamic graph by an adversary. For simplicity we use the term *node*  $u$  to also refer to the process which is assigned to node  $u$ . In each round, each node decides either to transmit a message or to listen to the wireless channel. The behavior of the wireless channel is modeled by using the standard radio network model first used in [5, 10]. When node  $u$  decides to transmit in round  $r$ , its message reaches all of its neighbours in  $G_r$ . A node  $v$  which listens in round  $r$  receives a message transmitted by a neighbor  $u$  if and only if  $u$  is the only neighbor of  $v$  in  $G_r$  which is transmitting in round  $r$ . If no message reaches  $v$  (no neighbor is transmitting),  $v$  receives silence, indicated by  $\perp$ . If two or more messages reach  $v$ ,  $v$  also receives  $\perp$ , i.e.,  $v$  cannot distinguish 2 or more transmitting neighbors from silence.

**Adversary.** We assume that the network changes under the control of an adversary. For any round  $r$  the adversary has to determine  $G_r$  based on the knowledge it has. For an integer  $\tau \geq 0$ , we call an adversary  $\tau$ -oblivious if for any  $r \geq 1$ , the adversary constructs  $G_r$  based on the knowledge of: (1) the algorithm description, (2) the network topologies of rounds  $1, \dots, r-1$ , and (3) the nodes' random choices of the first  $r-\tau$  rounds.

**Global Broadcast.** A distributed algorithm solving the global broadcast problem needs to disseminate a single message  $\mathcal{M}$  from a distinguished source node to all the processes in the network. We assume that in a distributed broadcast algorithm, non-source nodes are activated (and can start to actively transmit) when they first receive the broadcast message  $\mathcal{M}$ . Nodes that do not yet know  $\mathcal{M}$  remain silent.

**Mathematical Notation.** For two integers  $a \leq b$ ,  $[a, b]$  denotes the set of all integers between  $a$  and  $b$  (including  $a$  and  $b$ ). Further, for an integer  $a \geq 1$ , we use  $[a]$  as a short form to denote  $[a] := [1, a]$ . We say that a probability event happens with high probability (w.h.p.) if it happens with probability at least  $1 - 1/n^c$ , where  $n$  is the number of nodes and  $c > 0$  is a constant which can be chosen arbitrarily large by adjusting other constants.

## 3 Upper Bound

### 3.1 Randomized Broadcasting Algorithm

We now describe our randomized algorithm which solves broadcast in a  $T$ -interval  $k$ -connected radio network against a  $\tau$ -oblivious adversary. As stated in Section 1.1, the algorithm has a time complexity of  $O((1 + n/(k\psi)) \cdot n \log^3 n)$  with high probability where  $\psi := \min\{\tau, T, n/2k\}$ . In light of the comment following Theorem 1 in Section 1.1, throughout

## 7:6 The Cost of Global Broadcast in Dynamic Radio Networks

Section 3, we assume that  $\psi = \Omega(\log^3 n)$  as otherwise, one can achieve a stronger upper bound by just using an adapted version of [11].

In the first round, the source node transmits the message to its neighbors. Because we assume that each graph is  $k$ -vertex connected, after one round, at least  $k + 1$  nodes know the message. From there on, our randomized algorithm works in *phases*. To simplify notation, in the following, we ignore the first round and assume that at time 0, the algorithm starts with at least  $k + 1$  nodes which know the broadcast message  $\mathcal{M}$ . The phases of the algorithm are defined as follows.

► **Definition 4 (Phase).** The  $j^{\text{th}}$  time interval of  $\psi$  consecutive rounds is called phase  $j$ , where  $j$  is a positive integer. Hence, phase  $j$  starts at time  $(j - 1)\psi$  and ends at time  $j\psi$  and it consists of rounds  $(j - 1)\psi + 1, \dots, j\psi$ .

Let  $t_v$  denote the round in which  $\mathcal{M}$  is received by node  $v$  for the first time. In each round  $t$  the set  $V$  is partitioned into following three subsets. The *previously informed nodes*  $I(t)$  are the nodes that have received  $\mathcal{M}$  in some phase before the current phase. Note that in the first phase,  $I(t)$  consists of at least  $k + 1$  informed nodes. The nodes that have received  $\mathcal{M}$  for the first time in the current phase in some round before time  $t$  are called *newly informed nodes*, and they are denoted by  $N(t)$ . Finally, the set of *uninformed nodes* at time  $t$  is denoted by  $U(t) := V \setminus \{I(t) \cup N(t)\}$ .

The algorithm can be seen as a combination of two existing protocols which appeared in [11] and [17]. The protocol of [11] is a very basic one where all informed nodes always try to transmit the message independently with the same uniform probability. In the *harmonic broadcast* protocol of [17], informed nodes use harmonically decaying probabilities to forward the message. In each phase of our algorithm, in the first  $\lceil \psi/2 \rceil$  rounds, a variant of the protocol of [11] is applied and in the second  $\lfloor \psi/2 \rfloor$  rounds, the idea of the protocol of [17] is applied. In the following, the algorithm is described in detail.

**First half of a phase.** In the first  $\lceil \psi/2 \rceil$  rounds of a phase, all informed nodes, i.e., all  $v \in \{I(t) \cup N(t)\}$ , transmit the message with probability  $1/n$ .

**Second half of a phase.** The nodes in  $U(t) \cup I(t)$  remain silent throughout the second half of a phase. However, in each round  $t$ , any node  $v \in N(t)$  transmits the message with probability  $p_v(t)$ , given by


$$\forall t > \left\lfloor \frac{t}{\psi} \right\rfloor \cdot \psi + \left\lceil \frac{\psi}{2} \right\rceil : \forall v \in N(t) : p_v(t) := \frac{1}{1 + \left\lfloor \frac{t - \hat{t}_v - 1}{\mathcal{T}} \right\rfloor}, \quad (1)$$

where

$$\hat{t}_v := \begin{cases} \left\lfloor \frac{t}{\psi} \right\rfloor \cdot \psi + \left\lceil \frac{\psi}{2} \right\rceil, & \text{if } \lfloor t/\psi \rfloor \cdot \psi < t_v < \lfloor t/\psi \rfloor \cdot \psi + \lceil \psi/2 \rceil \\ t_v, & \text{otherwise} \end{cases}$$

and  $\mathcal{T}$  will be fixed in Lemma 5.

Thus, in the second half of a phase, only nodes participate which for the first time receive  $\mathcal{M}$  in the current phase. Each node  $v$  which gets newly informed in the phase executes the following protocol. As soon as  $v$  knows  $\mathcal{M}$  and as soon as the second half of the phase has started,  $v$  starts transmitting  $\mathcal{M}$  to its neighbors. For the first  $\mathcal{T}$  rounds,  $v$  transmits the  $\mathcal{M}$  with probability 1, for the next  $\mathcal{T}$  rounds  $v$  transmits  $\mathcal{M}$  with probability  $1/2$ , and the probability for the next time intervals of  $\mathcal{T}$  rounds becomes  $1/3, 1/4$ , etc.



■ **Figure 1** Time intervals  $[\theta_{i-1}, \theta_i]$  with equal number of free and busy rounds where  $i > 1$ .

### 3.2 Analysis

Recall that by the definition of  $\psi$ , we have  $k\psi \leq n/2$ ,  $\psi \leq T$ , and  $\psi \leq \tau$ . The  $T$ -interval  $k$ -connectivity of the dynamic network guarantees the existence of a stable spanning subgraph with vertex connectivity of at least  $k$  throughout the whole duration of every phase. We call this reliable spanning subgraph the *backbone* of the phase. Note that we may have different backbones in different phases. Let  $P(t)$  denote the sum of transmitting probabilities of all the nodes in round  $t$ , i.e.,  $P(t) := \sum_{v \in V} p_v(t)$ . For the analysis of our algorithm, we say that round  $t$  is *busy* if  $P(t) \geq 1$  and otherwise we say that round  $t$  is *free*. If the node  $v$  is the only node transmitting in a round, we say that node  $v$  gets *isolated* in that round.

For any phase  $j$ , let  $\theta_0 := j\psi - \lfloor \psi/2 \rfloor$ , i.e.,  $\theta_0$  is the time when the second half of the phase starts. For  $i > 0$ , we define  $\theta_i > \theta_{i-1}$  to be the first time such that in the time interval  $[\theta_{i-1}, \theta_i]$  (i.e., in rounds  $\theta_{i-1} + 1, \dots, \theta_i$ ) the number of busy rounds equals the number of free rounds (see Figure 1).

We further define  $m \geq 0$  such that  $\theta_m$  is the last such time defined for a given phase. The case  $m = 0$  implies that throughout the second half of the phase, the number of busy rounds is always larger than the number of free rounds.


We use the following lemma adapted from Lemma 13 of [17].

► **Lemma 5** ([17]). *Consider a node  $v$ . Let  $t > \hat{t}_v$  be such that at least half of the rounds  $\hat{t}_v + 1, \dots, t$  are free. If  $\mathcal{T} \geq \lceil 12 \ln(n/\epsilon) \rceil$  for some  $\epsilon > 0$ , then with probability larger than  $1 - \epsilon/n$  there exists a round  $t' \in \{\hat{t}_v + 1, \dots, t\}$  such that  $v$  is isolated in round  $t'$ .*

► **Lemma 6.** *For all phases, in each time interval  $[\theta_{i-1}, \theta_i]$ , where  $i \in [m]$ , if round  $\theta_{i-1} + 1$  is busy then any node  $v$  with  $\hat{t}_v \in \{\theta_{i-1}, \dots, \theta_i - 1\}$  gets isolated in some round  $t' \in \{\hat{t}_v + 1, \dots, \theta_i\}$  with high probability.*

**Proof.** Let  $\bar{t}$  denote the first round that the number of free rounds equals the number of busy rounds starting from round  $\hat{t}_v + 1$ . For the sake of contradiction, assume that  $\bar{t} > \theta_i$ , that is, the number of free rounds is less than the number of busy rounds in  $\{\hat{t}_v + 1, \dots, \theta_i\}$  and we also know that the number of busy rounds is greater than the number of free rounds in  $\{\theta_{i-1} + 1, \dots, \hat{t}_v\}$  (because of minimality of  $\theta_i$  and the fact that round  $\theta_{i-1} + 1$  is busy). It follows that the number of busy rounds is greater than the number of free rounds in  $\{\theta_{i-1} + 1, \dots, \theta_i\}$  contradicting our assumption on the equality of free and busy rounds in  $\{\theta_{i-1} + 1, \dots, \theta_i\}$ . Therefore,  $\bar{t} \leq \theta_i$  and according to Lemma 5 the claim holds. ◀

As one can see in Figure 2, at the beginning of each phase, the uninformed nodes in the backbone form one or several connected subgraphs which we call the *uninformed connected components*. For each uninformed connected component there must exist some edge in the backbone (within a phase) connecting an informed node to a node in that component. Note that because the adversary is  $\tau$ -oblivious and thus also oblivious to the last  $\psi \leq \tau$  rounds, the adversary has to determine the dynamic graph throughout a phase before the phase starts. The backbone graph of a phase can therefore not change depending on the randomness of the algorithm during the phase.



■ **Figure 2** Backbone of a phase. Available components are identified by thick circles.

► **Definition 7** (Available Components and Available Nodes). At the end of the first  $\lceil \psi/2 \rceil$  rounds of each phase, any uninformed connected component that includes at least one newly informed node is called an available component. All the nodes in an available component are called available nodes.

► **Lemma 8.** Consider an arbitrary phase and an arbitrary  $i \geq 1$ . If at the beginning of round  $\theta_i$  of the phase there exists at least one uninformed available node, then w.h.p. at least one available node gets informed in round  $\theta_i$ .

**Proof.** We will show that for every  $i \geq 1$ , w.h.p., if there is some node available  $u$  with  $\hat{t}_u = \theta_{i-1}$  and at the beginning of round  $\theta_i$ , there is at least one uninformed available node, then at least one available node  $v$  gets informed in round  $\theta_i$ . The claim of the lemma then follows by induction on  $i$ . If there are no available nodes, there is nothing to prove. If there are available nodes, there is at least one node  $u$  which gets newly informed in the first half of the phase and we therefore have  $\hat{t}_u = \theta_0$ . Using the above claim, it then w.h.p. follows that if there still is an uninformed available node at time  $\theta_1 - 1$ , some uninformed available node  $u'$  gets informed in round  $\theta_1$  and thus  $\hat{t}_{u'} = \theta_1$ . For  $i > 1$ , the induction step now follows in the same way. It therefore remains to show that w.h.p., if there is some node available  $u$  with  $\hat{t}_u = \theta_{i-1}$  and at the beginning of round  $\theta_i$ , there is at least one uninformed available node, then at least one available node  $v$  gets informed in round  $\theta_i$ .

By Lemma 6 we know that w.h.p., all the nodes  $u$  with  $\hat{t}_u \in \{\theta_{i-1}, \dots, \theta_i - 1\}$  get isolated in some round  $t' \in \{\theta_{i-1} + 1, \dots, \theta_i\}$ . Hence, by induction on  $j$ , w.h.p., for all  $j \leq i - 1$  there is some node  $u'$  with  $\hat{t}_{u'} = \theta_j$  and therefore all nodes  $u$  with  $\hat{t}_u \in \{\theta_0, \dots, \theta_i - 1\}$  get isolated in some round  $t' \in \{\theta_0 + 1, \dots, \theta_i\}$ . Consequently, w.h.p., all newly informed nodes  $N(\theta_i - 1)$  at time  $\theta_i - 1$  get isolated in some round  $t' \in \{\theta_0 + 1, \dots, \theta_i\}$ . Let  $v$  be an uninformed available node before round  $\theta_i$  (i.e., at time  $\theta_i - 1$ ). Because  $v$  is available, at time  $\theta_i - 1$ , there is a informed available neighbor  $u$  in the backbone graph of the current phase. We clearly have  $u \in N(\theta_i - 1)$  and thus w.h.p.,  $u$  gets isolated in some round  $t' \in \{\theta_0 + 1, \dots, \theta_i\}$ . As soon as  $u$  gets isolated,  $v$  gets informed and we can therefore conclude that  $u$  gets isolated in round  $\theta_i$  and thus  $v$  gets informed in round  $\theta_i$ . ◀

The proof of the following lemma appears in the full version of this paper [1].

► **Lemma 9.** Consider an arbitrary phase and assume that at the beginning of the second  $\lceil \psi/2 \rceil$  rounds of the phase there are  $z$  available nodes. Then, w.h.p., for some constant  $c > 0$ , at the end of the phase we have at least  $\min\{z, c\psi/\ln^2 n\}$  newly informed nodes.

Using the established technical lemmas, we can now proof our upper bound theorem.



► **Theorem 1 (restated).** *Let  $T \geq 1$ ,  $\tau \geq 1$ , and  $k \geq 1$  be positive integer parameters. Assume that the adversary is  $\tau$ -oblivious. Then, in a dynamic  $T$ -interval  $k$ -connected  $n$ -node radio network, with high probability, single message broadcast can be solved in time*

$$O\left(\left(1 + \frac{n}{k \cdot \min\{\tau, T\}}\right) \cdot n \log^3 n\right).$$

**Proof.** Consider some phase  $j$  and let  $\mathcal{B}_j$  be the backbone of phase  $j$ , i.e.,  $\mathcal{B}_j$  is the stable  $k$ -connected subgraph of phase  $j$ . Consider the subgraph  $\mathcal{B}_j[U_j]$  of  $\mathcal{B}_j$  induced by the uninformed nodes  $U_j$ . This induced subgraph might consist of several connected components. However, each of the components is connected to at least  $k$  nodes in  $I_j$  as it is shown in Figure 2 (recall that we can assume that  $|I(t)| \geq k + 1$ ). Note that if one of these at least  $k$  nodes gets isolated in the first half of the phase, all nodes in the component become available for the second half of the phase.

In the first half of phase  $j$ , in any round  $t$ , each node in  $I(t)$  transmits the message with probability  $1/n$ . Therefore, for every node  $u \in I(t)$ , the probability that  $u$  gets isolated in round  $t + 1$  (in the first half of a phase) is at least

$$\Pr(u \text{ gets isolated in round } t + 1) \geq \frac{1}{n} \left(1 - \frac{1}{n}\right)^{n-1} > \frac{1}{en}. \quad (2)$$

In following, we analyze the progress in the first half of some phase  $j$ . Consider an uninformed node  $v \in U_j$  (at the beginning of phase  $j$ ). Let  $\mathcal{A}_v$  be the event that  $v$  becomes available in the first half of phase  $j$ . Event  $\mathcal{A}_v$  definitely occurs if one of the at least  $k$  initially informed neighbors of  $v$ 's component in  $\mathcal{B}_j[U_j]$  gets isolated in one of the at least  $\psi/2$  rounds of the first half of the phase. The probability for this is

$$\Pr(\mathcal{A}_v) \geq 1 - \left(1 - \frac{k}{en}\right)^{\psi/2} > 1 - e^{-k\psi/2en} \geq \frac{k\psi}{4en}.$$

The last inequality follows from the fact that for all  $0 \leq x \leq 1$ ,  $e^{-x} \leq 1 - x/2$ . Let  $X$  be the number of nodes in  $U_j$  that get available in phase  $j$ . For convenience, we define  $\lambda := |U_j|/n \leq 1$ . We then have

$$\mathbb{E}[X] = \sum_{v \in U_j} \Pr(\mathcal{A}_v) \geq |U_j| \cdot \frac{k\psi}{4en} = \frac{\lambda k\psi}{4e}. \quad (3)$$

We define  $F := \min\left\{\frac{\psi}{16e}, \frac{c\psi}{\ln^2 n}\right\}$ , where  $c > 0$  is the constant that is used in Lemma 9. Note that by Lemma 9, in phase  $j$ , w.h.p., at least  $\min\{X, F\}$  uninformed nodes become informed.

We define a phase to be successful if  $X \geq \lambda F$ . Let  $\mathcal{S}$  be the event that phase  $j$  is successful and let  $\bar{\mathcal{S}}$  be the complementary event. We can upper bound the expected value of  $X$  as follows:

$$\mathbb{E}[X] < \Pr(\bar{\mathcal{S}}) \cdot \lambda F + (1 - \Pr(\bar{\mathcal{S}})) \cdot \lambda n.$$

Combining with the upper bound in (3), we obtain (recall that we assume that  $k\psi \leq n/2$ ).

$$\Pr(\bar{\mathcal{S}}) < \frac{n - \frac{k\psi}{4e}}{n - F} \leq \frac{n - \frac{k\psi}{4e}}{n - \frac{\psi}{16e}} \leq \left(1 - \frac{k\psi}{4en}\right) \left(1 + \frac{\psi}{8en}\right)^{\binom{k\geq 1}{\leq}} \leq \left(1 - \frac{k\psi}{8en}\right). \quad (4)$$

By Lemma 9, in a successful phase, w.h.p., at least  $\lambda F$  new nodes get informed. Hence, in a successful phase, w.h.p., we get rid of at least an  $(F/n)$ -fraction of the remaining uninformed nodes. In order to inform all nodes, w.h.p., we therefore need at most  $O(n \log(n)/F) = O(n \log^3(n)/\psi)$  successful phases. Using (4) and a standard Chernoff argument, we can thus w.h.p. upper bound the total number of phases by  $O(n^2 \log^3(n)/(k\psi^2))$ . As each phase takes  $\psi$  round, this concludes the proof. ◀

## 4 Lower Bound

In this section we prove a lower bound for global broadcast in  $T$ -interval  $k$ -connected radio networks against a 1-oblivious adversary. Furthermore, we show impossibility of solving the same problem against a strongly adaptive adversary (0-oblivious adversary).

Our lower bound is based on a general technique for proving lower bounds for communication problems in radio networks, introduced by Newport in [23]. Using this technique, one first defines a combinatorial game for which a lower bound can be proved directly. It is then shown how to reduce the game to the problem in order to leverage the game's lower bound to obtain the desired lower bound for the problem.


To prove Theorem 2 using this technique, we first introduce an abstract hitting game, called the  $(\beta, \ell, \varphi)$ -periodic hitting game and directly prove a lower bound for winning this game. We note that this game is more involved than the games used in previous work, e.g., [13, 23]. Based on a lower bound for the hitting game, for a given instance of the game we instantiate an  $n$ -node target network. By instantiation of an  $n$ -node network, we mean assigning  $n$  processes with unique IDs to the nodes of the network. For the instantiation, one needs to also have information which is not available to the player in the game. However, we show that by playing the game, the player can still simulate the execution of a given broadcast algorithm on the corresponding target network to the given instance of the game. We show that this simulation of a broadcast algorithm allows to win the hitting game and the lower bound on the hitting game and the simulation together imply a lower bound for solving global broadcast.

**$(\beta, \ell, \varphi)$ -periodic hitting game.** The game is defined for three integers  $\beta, \ell, \varphi > 0$  and proceeds in rounds. Time is divided into phases of  $\varphi$  rounds, where the  $j^{\text{th}}$  phase of the game is called phase  $\pi_j$ . That is, phase  $\pi_1$  consists of rounds 1 to  $\varphi$ , phase  $\pi_2$  consists of rounds  $\varphi + 1$  to  $2\varphi$ , etc. The first round of any phase  $\pi_j$  is called  $t_j$ . The player of the game is represented by a probabilistic automaton  $\mathcal{P}$  and plays the game against a referee. Two sets are defined for this game, a *selection set*  $S := [\beta]$  which is fixed during the game, and a *target set* which might change from round to round. The target set of round  $t$  is denoted by  $X(t)$ . In each round  $t$ ,  $\mathcal{P}$  chooses one element from  $S$  and outputs it as the guess  $\gamma(t)$  for round  $t$ . Round  $t$  is called a *successful round*, if and only if  $\gamma(t) \in X(t)$ .

At the beginning of each phase  $\pi_j$  ( $j \geq 1$ ), the referee determines a set  $Y_j$  consisting of  $\ell$  elements chosen uniformly at random from  $S$ . We define the target set as follows. For convenience, assume that  $Y_0, Y_{-1}$  and  $X(0)$  are empty sets.

$$\forall j, \forall t \in [t_j, t_j + \varphi - 1] : X(t) := \begin{cases} Y_j \cup [X(t-1) \setminus (\{\gamma(t-1)\} \cup (Y_{j-2} \setminus Y_{j-1}))] & \text{if } t = t_j, \\ X(t-1) \setminus \{\gamma(t-1)\} & \text{if } t \neq t_j. \end{cases}$$

That is, at the beginning of each phase  $\pi_j$ , the referee chooses  $\ell$  elements  $Y_j$  from  $S$  uniformly at random and adds them to the target set. Two phases ( $2\varphi$  rounds) later, each of these  $\ell$  elements which still remains in the target set (and which is not in  $Y_{j+1}$ ) is removed from the target set by the referee. Moreover, after each successful round, the referee removes the correct guess from the target set. Player  $\mathcal{P}$  wins the game in  $r$  rounds if and only if either round  $r$  is the  $\beta^{\text{th}}$  successful round for  $\mathcal{P}$ , or before round  $r + 1$  (in phase  $j$ ),  $X(r+1) \cap Y_{j-1} = \emptyset$  or  $X(r+1) \cap Y_j = \emptyset$ . The second condition will be used to ensure sufficiently large interval connectivity of the target network as long as the game is not won. The only information that the player receives at the end of each round is whether the round was successful or not. The player is also notified if it wins the game.



■ **Figure 3** A snapshot of the dynamic network used in the hitting game simulation.

Intuitively, as long as the target set changes sufficiently often, it should always appear essentially random to the player. Therefore, the best strategy for hitting the target set is to always choose an almost uniformly random guess, leading to roughly  $\beta/\ell$  rounds to get a single successful round. The following lemma states this intuition formally. The proof of the lemma appears in the full version of this paper [1].

► **Lemma 10.** *For any  $\varphi \leq \beta/3$  and for  $\ell \geq \ell_0$  for a sufficiently large constant  $\ell_0 > 0$ , the expected number of rounds for a player to win the  $(\beta, \ell, \varphi)$ -periodic hitting game is at least  $\Omega(\beta^2/\ell)$ .*

► **Lemma 11.** *If algorithm  $\mathcal{A}$  solves the global broadcast problem in any  $T$ -interval 1-connected dynamic  $n$ -nodes network against a 1-oblivious adversary in  $f(n) = n^{O(1)}$  rounds in expectation for a sufficiently large value of  $T$ , then we can construct a player  $\mathcal{P}$  to win the  $(\lfloor n/2 \rfloor - \ell, \ell, cT \ln n)$ -periodic hitting game in expected  $O(f(n) \log n)$  rounds, for some positive constants  $c$  and  $\ell$ .*

**Proof.** We construct a player  $\mathcal{P}$  to simulate the execution of  $\mathcal{A}$  on a particular  $T$ -interval 1-connected dynamic  $n$ -node network (the target network). Then the player uses the transmitting behavior of the nodes in the simulation to generate guesses for playing the game. We start by defining the target network for a given instance of the  $(\lfloor n/2 \rfloor - \ell, \ell, cT \ln n)$ -periodic hitting game.

**The Target Network.** For the following discussion, we set  $\beta := \lfloor n/2 \rfloor - \ell$  and  $\varphi := cT \log n$  to denote the size of the selection set and the length of a phase of the hitting game. We assume that we are given an instance of the  $(\beta, \ell, \varphi)$ -periodic hitting game. Based on how the hitting game develops, we define an  $n$ -node dynamic target network. We first describe the core (backbone) part of the network. The nodes of the dynamic network are defined as  $V := \{0, \dots, n-1\}$ . We assume that node 0 is the source and we identify the next  $\beta$  nodes (i.e., the set  $[\beta]$ ) with the selection set  $S$  of the hitting game. Throughout the execution, node 0 is connected to all nodes in  $[\beta]$  and it is not connected to any other node. Throughout the simulation of the broadcast algorithm, we use  $I$  and  $U$  to denote the set of informed and uninformed nodes, respectively (a node is informed iff it knows the broadcast message  $\mathcal{M}$ ). Clearly, as soon as the source node broadcasts  $\mathcal{M}$ , the set of informed nodes is  $I = \{0, \dots, \beta\}$  and we thus have  $U = \{\beta+1, \dots, n-1\}$ . To simplify notation, we assume that already at

the start of the simulation, all nodes in  $\{0, \dots, \beta\}$  know  $\mathcal{M}$  and thus, we start round 1 with  $I = \{0, \dots, \beta\}$  and  $U = \{\beta + 1, \dots, n - 1\}$ . We will assume that the number of uninformed nodes is always at least  $2\ell$ . As soon as it drops below, we stop carrying out the simulation.

Throughout the simulation, we always assume that all nodes in  $I$  form a clique and all nodes in  $U$  form a clique. Apart from this, the topology of the core network is determined by the target set of the hitting game that we are trying to win by simulating  $\mathcal{A}$ . Assume that in some round  $r$  of the hitting game, the target set is  $X(r) \subset [\beta]$ . During the simulation, in the backbone network we then use the nodes in  $X(r)$  as bridge nodes to connect the informed nodes to the uninformed nodes. Each node  $x \in X(r)$  is connected to exactly one node  $n(x) \in U$  such that each node in  $U$  is connected to at most one node in  $X(r)$ . We assume that whenever a new node is added to  $X(r)$ , its neighbor in  $U$  is chosen uniformly at random among all nodes in  $U$  which are not already connected to a bridge node in  $X(r)$ . Note that the size of  $X(r)$  is always at most  $2\ell$  and because we assumed that  $|U| \geq 2\ell$ , we can always do such an assignment of bridge nodes. Whenever the player makes a successful guess  $x \in X(r)$ , we move  $x$  to the set of informed nodes  $I$  and we connect  $x$  with all nodes in  $I$  and disconnect it with all nodes in the remaining set  $U$  of uninformed nodes. Note that in the hitting game, after a successful guess  $x \in X(r)$ ,  $x$  is also removed from the target set. The target network at any time is either the described core network (backbone) or the complete graph  $K_n$ , for which the choice will be explained later.

**The Simulation.** The simulation of the broadcast algorithm  $\mathcal{A}$  is done in a round-by-round manner. As the dynamic topology used in the simulation depends on the target set of the hitting game, the player  $\mathcal{P}$  of the hitting game does not know the dynamic topology. We need to show that  $\mathcal{P}$  can still correctly simulate the behavior of the broadcast algorithm.

As discussed above, we assume that at the beginning of the simulation, the set of informed nodes is  $I = \{0, \dots, \beta\}$ . Each round of  $\mathcal{A}$  is simulated by  $\mathcal{P}$  by making at most  $c \ln n$  guesses in the hitting game. More specifically, a given round  $r$  of  $\mathcal{A}$  is simulated as follows.

First, note that because we assume that the adversary is 1-oblivious,  $\mathcal{P}$  can base the graph of round  $r$  on the states of all nodes at the beginning of the round. Hence, in particular, the graph of round  $r$  can depend on the probability  $p_v(r)$  with which each node  $v \in I$  transmits in the given round. We define a round  $r$  of  $\mathcal{A}$  to be *busy* if  $\sum_{v \in I} p_v(r) > \frac{c}{2} \ln n$ , otherwise a round  $r$  is called *free*. In a busy round, the network graph is assumed to be the complete graph  $K_n$  and in a free round, the network graph is assumed to be exactly the backbone graph as described above. We assume that  $\mathcal{P}$  always knows the set of informed nodes and because only informed nodes are allowed to transmit,  $\mathcal{P}$  can determine all messages which are transmitted in a round by simulating the random decisions of the nodes in  $I$ . We say that the simulated execution of  $\mathcal{A}$  is *bad* if either there is a free round in which more than  $c \ln n$  nodes in  $I$  decide to transmit a message or if there is a busy round in which exactly one node in  $I$  decides to transmit. Otherwise, the simulated execution is called *good*. If an execution turns out to be bad,  $\mathcal{P}$  stops the simulation of  $\mathcal{A}$  and simply continues making random guesses until it wins the hitting game. Note that the expected time to win the hitting game in this way is at most  $O(\beta^2/\ell)$  as unless there have been at least  $\Omega(\ell)$  successes during the last  $\phi = O(\beta)$  rounds of the hitting game, the probability for a successful guess is always  $\Omega(\ell/\beta)$ . As long as the number of simulated rounds  $f(n)$  of  $\mathcal{A}$  is polynomial in  $n$  (and thus also in  $\beta$ ), for an arbitrary given constant  $d > 0$  and sufficiently large constant  $c$ , the probability to obtain a good execution is at least  $1 - 1/n^d$ . For sufficiently large constant  $c$ , the expected time to win the hitting game is therefore dominated by the expected time to win the game conditioned on the event that the simulation of  $\mathcal{A}$  creates a good execution.

In the following, we therefore assume that the generated broadcast execution is good. In the following, we also assume that in the current phase (of length  $\varphi$ ) of the hitting game, there are still at least  $c \ln n$  guesses that can be made. If this is not the case, player  $\mathcal{P}$  first makes a sequence of unsuccessful guesses to finish the phase ( $\mathcal{P}$  can for example repeat the last guess it has made before to make sure it is not successful). As we assumed that  $T$  is sufficiently large, we can assume that  $\varphi \gg c \ln n$  and therefore we only waste a small fraction of all guesses by doing this.

Let us first assume that a simulated round  $r$  of  $\mathcal{A}$  is busy. As in this case, the communication network is a complete graph and since we assume that in a good execution, no node gets isolated (transmits alone), every node receives silence and we therefore do not need to simulate any receive behavior. In this case, we also do not make any guesses in the hitting game. If round  $r$  is free, the number of nodes that transmit is between 0 and  $c \ln n$ . First recall that the nodes in  $I$  are fully connected and  $\mathcal{P}$  can therefore clearly simulate their receive behavior. Further, let  $Z(r) \subseteq [\beta]$  be the set of nodes in  $\beta$  which are transmitting in round  $r$ . For each  $z \in Z(r)$ , player  $\mathcal{P}$  uses  $z$  as a guess in the hitting game. Note that because there are at most  $c \ln n$  guesses to be made and because we assumed that there are still at least  $c \ln n$  guesses in the current phase of the hitting game, during making the guesses for all  $z \in Z(r)$ , we do not change the phase (and thus the target set) in the hitting game. The node  $z$  therefore is a bridge node connecting  $I$  to a node  $n(z) \in U$  in round  $r$  of the broadcast algorithm if and only if  $z$  is a successful guess. In that case,  $n(z)$  is a uniformly random node in  $U$ . Hence, if  $z$  is a successful guess,  $\mathcal{P}$  chooses  $n(z)$  uniformly at random in  $U$  and it moves  $n(z)$  from  $U$  to  $I$ . Note that  $z$  is also removed from the target set, and all connections of  $n(z)$  to nodes in the remaining set  $U$  are removed. Note also that by choosing  $n(z)$  uniformly at random in  $U$ , player  $\mathcal{P}$  does not only simulate the randomness of the broadcast algorithm, but it also simulates the randomness of the adversary. As long as the execution is good, in the given dynamic network, the broadcast algorithm informs a new node if and only if one of the bridge nodes  $v$  transmits in a free round. For any bridge node  $v$  which transmits in a free round, the corresponding uninformed bridge node  $n(v)$  gets informed. The described simulation therefore correctly simulates the broadcast algorithm and it informs a new node if and only if it makes a correct guess. As we only stop the simulation once the number of uninformed nodes drops below  $2\ell$ , we need at least  $n - (\beta + 1) - 2\ell + 1 = \lceil n/2 \rceil - \ell \geq \beta$  successful guesses and thus win the game to stop the simulation.

It remains to show that the dynamic network used in the simulation is  $T$ -interval connected. Every  $\varphi$  guesses and thus after at least  $\varphi/(c \ln n)$  rounds of the simulation, we add  $\ell$  new edges connecting some  $v \in I$  and  $n(v) \in U$ . As long as  $v$  is not used as a guess, such an edge remains for  $2\varphi$  guesses. As long as it is always guaranteed that one of these edges survives the next  $2\varphi$  guesses (and thus at least  $2\varphi/(c \ln n)$  rounds), the network is at least  $T = \varphi/(c \ln n)$ -interval connected. Hence, the network is not guaranteed to be  $T$ -interval connected if there is a phase  $j$  in the hitting game such that all the elements of the set  $Y_j$  added to the target set at the beginning of phase  $j$  are successfully guessed by the end of phase  $j + 1$ . Recall that in this case, the player also wins the game and therefore the claim of the lemma follows.  $\blacktriangleleft$

For  $k = 1$ , the statement of our main lower bound theorem (Theorem 2) now directly follows by combining Lemmas 10 and 11.

**► Lemma 12.** *For every constant  $\varepsilon > 0$  and every  $T \leq n^{1-\varepsilon}$ , the expected time to solve single-message broadcast in  $T$ -interval 1-connected radio networks against a 1-oblivious adversary is at least  $\Omega\left(\frac{n^2}{\log n}\right)$ .*

**Proof.** For the sake of contradiction let us assume that  $\mathcal{A}$  can solve broadcast for any  $T$ -interval 1-connected network in  $f(n) = o(n^2/\ell \log n)$  rounds, then based on Lemma 11 a player can solve any instance of the  $(n/2, \ell, cT \log n)$ -periodic hitting game in  $o(n^2/\ell)$  rounds which contradicts Lemma 10 and this proves the necessity of  $\Omega(n^2/\ell \log n)$  rounds to solve broadcast in any  $T$ -interval 1-connected network. Based on Lemma 10, by choosing sufficiently large constant value for  $\ell$  the claimed lower bound follows.  $\blacktriangleleft$

In order to obtain Theorem 2, we need to generalize the above result from  $T$ -interval 1-connected networks to  $T$ -interval  $k$ -connected networks for arbitrary  $k \geq 1$ . As shown below, this can be achieved by using a simple generic reduction.

► **Theorem 2 (restated).** *For every constant  $\varepsilon > 0$  and every  $T \leq (n/k)^{1-\varepsilon}$ , the expected time to solve single-message broadcast in  $T$ -interval  $k$ -connected radio networks against a 1-oblivious adversary is at least  $\Omega(\frac{n^2}{k^2 \log n})$ .*

**Proof.** Given an  $n$ -node graph  $G$ , let  $H_k(G)$  be the graph which is obtained by replacing each node of  $G$  by a clique of size  $k$  and by replacing each edge  $\{u, v\}$  of  $G$  by a complete bipartite subgraph  $K_{k,k}$  between the two  $k$ -cliques representing  $u$  and  $v$ . If  $G$  is connected, in order to disconnect  $H_k(G)$  by deleting some nodes, we need to completely remove at least one of the  $k$ -cliques representing the nodes of  $G$ . Hence, if  $G$  is connected,  $H_k(G)$  is  $k$ -vertex connected. It follows in the same way that if we have a dynamic graph  $G_1, \dots, G_t$  which is  $T$ -interval 1-connected, the dynamic graph  $H_k(G_1), \dots, H_k(G_t)$  is  $T$ -interval  $k$ -connected. Even if all nodes of such a graph  $H_k(G_i)$  know to which of the cliques representing the nodes of  $G_i$  they belong, solving broadcast in the dynamic graph  $H_k(G_1), \dots, H_k(G_t)$  cannot be easier than solving broadcast in  $G_1, \dots, G_t$ . If each graph  $H_k(G_i)$  has  $N$  nodes, the graphs  $G_i$  have  $N/k$  nodes and the claimed lower bound directly follows by applying Lemma 12.  $\blacktriangleleft$

## 5 Impossibility of Broadcast Against a 0-Oblivious Adversary

In this section, we prove the impossibility result that we stated in Section 1.1. We show that unless  $T$  is almost equal to  $n$ , the global broadcast problem cannot be solved in the presence of a 0-oblivious adversary, even for very large vertex connectivity  $k$ .

► **Theorem 3 (restated).** *For any  $k \geq 1$  and any  $T < n - k$ , it is not possible to solve single-message broadcast in  $T$ -interval  $k$ -connected radio networks against a 0-oblivious adversary.*

**Proof.** We show that a strongly adaptive adversary (i.e., a 0-oblivious adversary) can apply a simple strategy to prevent any algorithm from solving the global broadcast problem in a  $T$ -interval  $k$ -connected network, where  $T < n - k$ . Consider the following adversary strategy to determine the sequence of network topologies.

The adversary partitions the  $n$  nodes into two distinct sets  $A$  and  $B$ , such that  $A$  includes the source node and is of size  $T + k$ , and  $B$  is of size  $n - (T + k)$ . Since  $T < n - k$ , there exists at least one node in  $B$ . Note that at the beginning, no node in  $B$  knows the broadcast message  $\mathcal{M}$  (or anything about  $\mathcal{M}$ ). Because the adversary is 0-oblivious, in each round  $r$ , it can determine the graph after all nodes have made their random decisions. It can therefore determine the graph based on which nodes transmit.

If in a round  $r$ , either 0 nodes transmit or at least 2 nodes transmit, the network graph is chosen to be the complete graph. Note that in such a round, there is either silence or all nodes experience a collision. In both cases, all listening nodes receive  $\perp$  and therefore no node in  $B$  can learn something about  $\mathcal{M}$ .

If in a round, exactly one node  $v$  in  $A$  transmits, the network graph consists of all edges except the edges connecting  $v$  to nodes in  $B$ . Like this, also in this case all nodes in  $B$  receive  $\perp$  and they therefore cannot learn something about  $\mathcal{M}$ .

It remains to show that the given dynamic graph is  $T$ -interval  $k$ -connected. During the whole execution,  $A$  is a clique consisting of  $T+k$  nodes. Hence,  $A$  is a  $T$ -interval  $k$ -connected network. To show that whole  $n$ -node network is also a  $T$ -interval  $k$ -connected network, it is sufficient to show that for any node  $v \in B$ , in any  $T$  consecutive rounds, there exist at least  $k$  fixed edges from  $v$  to the nodes in  $A$ . To do so, fix some arbitrary time interval of  $T$  consecutive rounds. During the time interval, there are at most  $T$  rounds in which exactly one node transmits. Therefore, because  $|A| = T+k$ , there are at least  $k$  nodes in  $A$  which do not transmit alone during the given time interval. The edges from these  $k$  nodes to all nodes in  $B$  are therefore available throughout the  $T$  rounds. Therefore throughout any interval of  $T$  rounds, each node in  $B$  is connected to a set of at least  $k$  nodes in  $A$ . Consequently, the constructed dynamic network is  $T$ -interval  $k$ -connected. ◀

Notice that at least for store-and-forward algorithms even collision detection does not help to overcome the impossibility result. The 0-oblivious adaptive adversary knows the random choices of the algorithm in the current round and it can thus prevent any progress.

We also note that the above result turns out to be tight in the following sense. If  $T \geq n-k$ , global broadcast can be solved. If in each round, every node independently tries to broadcast with some probability (say  $1/n$ ), if  $T+k \geq n$  there is a non-zero probability (it may be very small) that  $T$  different nodes are isolated in  $T$  consecutive rounds. Consider an interval of  $T$  rounds and let  $I$  and  $U$  be the sets of informed and uninformed nodes at the beginning of this interval. From  $T$ -interval  $k$ -connectivity, we get that there are at least  $k$  nodes in  $I$  which are stably connected to nodes in  $I$  throughout the  $T$  rounds. Before broadcast is solved, we have  $|I| \leq n-1$  and if in the  $T$  rounds,  $T$  different nodes in  $I$  are isolated, at least one of the  $k$  nodes stably connected to  $U$  gets isolated and we can therefore make progress. Note that for  $T = n-k$ , the probability for making progress might be exponentially small, resulting in an exponential running time for the broadcast problem. Note however also that once  $T \geq cn \log n$  for a sufficiently large constant  $c$ , it is not hard to show that broadcast can be solved in polynomial time against a 0-oblivious adversary and if  $T$  is larger than  $cn \log^2 n$  for a sufficiently large constant  $c$ , it is shown in [17], that it can be solved in time  $O(n \log^2 n)$ .

---

## References

- 1 Mohamad Ahmadi, Abdolhamid Ghodselahi, Fabian Kuhn, and Anisur R. Molla. The cost of global broadcast in dynamic radio networks. *CoRR*, abs/1601.01912, 2016.
- 2 Antonio F. Anta, Alessia Milani, Miguel A. Mosteiro, and Shmuel Zaks. Opportunistic information dissemination in mobile ad-hoc networks: the profit of global synchrony. *Distributed Computing*, 25(4):279–296, 2012.
- 3 Chen Avin, Michal Koucký, and Zvi Lotker. How to explore a fast-changing world (cover time of a simple random walk on evolving graphs). In *Proc. 5th Coll. on Automata, Languages and Programming (ICALP)*, pages 121–132, 2008.
- 4 Reuven Bar-Yehuda, Oded Goldreich, and Alon Itai. Efficient emulation of single-hop radio network with collision detection on multi-hop radio network with no collision detection. *Distributed Computing*, 5:67–71, 1991.
- 5 Reuven Bar-Yehuda, Oded Goldreich, and Alon 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.

- 6 Hervé Baumann, Pierluigi Crescenzi, and Pierre Fraigniaud. Parsimonious flooding in dynamic graphs. In *Proc. of 28th ACM Symp. on Principles of Distributed Computing (PODC)*, pages 260–269, 2009.
- 7 Keren Censor-Hillel, Mohsen Ghaffari, and Fabian Kuhn. Distributed connectivity decomposition. In *Proc. 33rd Symp. on Principles of Distributed Computing (PODC)*, 2014.
- 8 Keren Censor-Hillel, Mohsen Ghaffari, and Fabian Kuhn. A new perspective on vertex connectivity. In *Proc. 25th ACM-SIAM Symp. on Discrete Algorithms (SODA)*, pages 546–561, 2014.
- 9 Keren Censor-Hillel, Seth Gilbert, Fabian Kuhn, Nancy Lynch, and Calvin Newport. Structuring unreliable radio networks. *Distributed Computing*, 27(1):1–19, 2014.
- 10 Imrich Chlamtac and Shay Kutten. On broadcasting in radio networks—problem analysis and protocol design. *IEEE Transactions on Communications*, 33(12):1240–1246, 1985.
- 11 Andrea Clementi, Angelo Monti, Francesco Pasquale, and Riccardo Silvestri. Broadcasting in dynamic radio networks. *J. Comput. Syst. Sci.*, 75(4):213–230, 2009.
- 12 Andrea Clementi, Angelo Monti, Francesco Pasquale, and Riccardo Silvestri. Optimal gossiping in geometric radio networks in the presence of dynamical faults. *Networks*, 59(3):289–298, 2012.
- 13 Mohsen Ghaffari, Nancy Lynch, and Calvin Newport. The cost of radio network broadcast for different models of unreliable links. In *Proc. 32nd Symp. on Principles of Distributed Computing (PODC)*, pages 345–354, 2013.
- 14 Piyush Gupta and Panganmala R. Kumar. The Capacity of Wireless Networks. *IEEE Transactions on Information Theory*, 46(2):388–404, 2000.
- 15 Tomasz Jurdzinski, Dariusz R. Kowalski, Michal Rozanski, and Grzegorz Stachowiak. On the impact of geometry on ad hoc communication in wireless networks. In *Proc. 33rd Symp. on Principles of Distributed Computing (PODC)*, pages 357–366, 2014.
- 16 Kyu-Han Kim and Kang G. Shin. On accurate measurement of link quality in multi-hop wireless mesh networks. In *Proc. Conf. on Mobile Computing and Networking (MOBICOM)*, pages 38–49, 2006.
- 17 Fabian Kuhn, Nancy Lynch, Calvin Newport, Rotem Oshman, and Andréa W. Richa. Broadcasting in unreliable radio networks. In *Proc. 29th Symp. on Principles of Distributed Computing (PODC)*, pages 336–345, 2010.
- 18 Fabian Kuhn, Nancy Lynch, and Rotem Oshman. Distributed computation in dynamic networks. In *Proc. 42nd Symp. on Theory of Computing (STOC)*, pages 513–522, 2010.
- 19 Fabian Kuhn and Rotem Oshman. Dynamic Networks: Models and Algorithms. *ACM SIGACT News*, 42(1):82–96, 2011.
- 20 Eyal Kushilevitz and Yishay Mansour. An  $\omega(d \log(n/d))$  lower bound for broadcast in radio networks. *SIAM journal on Computing*, 27(3):702–712, 1998.
- 21 Thomas Moscibroda and Roger Wattenhofer. Maximal independent sets in radio networks. In *Proc. 24th Symp. on Principles of Distributed Computing (PODC)*, pages 148–157, 2005.
- 22 Thomas Moscibroda and Roger Wattenhofer. The complexity of connectivity in wireless networks. In *Proc. 25th Conf. on Computer Communications (INFOCOM)*, pages 1–13, 2006.
- 23 Calvin Newport. Radio network lower bounds made easy. In *Distributed Computing*, pages 258–272. 2014.
- 24 Calvin Newport, David Kotz, Yougu Yuan, Robert S. Gray, Jason Liu, and Chip Elliott. Experimental evaluation of wireless simulation assumptions. *Simulation*, 83(9):643–661, 2007.
- 25 Regina O’Dell and Roger Wattenhofer. Information dissemination in highly dynamic graphs. In *Proc. of Workshop on Foundations of Mobile Computing (DIALM-POMC)*, pages 104–110, 2005.



- 26 Krishna Ramachandran, Irfan Sheriff, Elizabeth Belding, and Kevin Almeroth. Routing stability in static wireless mesh networks. In *Proc. Conf. on Passive and Active Network Measurement*, pages 73–82, 2007.
- 27 Kannan Srinivasan, Maria A. Kazandjieva, Saatvik Agarwal, and Philip Levis. The  $\beta$ -factor: Measuring wireless link burstiness. In *Proc. 6th Conf. on Embedded Networked Sensor System*, pages 29–42, 2008.
- 28 Mark D. Yarvis, Steven W. Conner, Lakshman Krishnamurthy, Jasmeet Chhabra, Brent Elliott, and Alan Mainwaring. Real-world experiences with an interactive ad hoc sensor network. In *Proc. Conf. of Parallel Processing*, pages 143–151, 2002.