Journal of Computer and System Sciences 75 (2009) 213–230

Contents lists available at ScienceDirect

Journal of Computer and System Sciences

www.elsevier.com/locate/jcss

Broadcasting in dynamic radio networks ☆,☆☆

Andrea E.F. Clementi [a,*], Angelo Monti [b], Francesco Pasquale [a,1], Riccardo Silvestri [b]

[a] *Dipartimento di Matematica, Università di Roma "Tor Vergata", I-00133 Rome, Italy*
[b] *Dipartimento di Informatica, "La Sapienza" Università di Roma, I-00133 Rome, Italy*

ARTICLE INFO

Article history:
Received 25 March 2008
Received in revised form 25 March 2008
Available online 14 December 2008

Keywords:
Randomized algorithms
Dynamic networks
Radio networks
Random graphs

ABSTRACT

It is reasonable to claim that almost all major questions related to radio broadcasting can be considered closed as far as *static* networks are considered: the network never changes during the entire protocol's execution. On the other hand, theoretical results on communication protocols in any scenario where the network topology may change during protocol's execution (i.e. a *dynamic* radio network) are very few.

In this paper, we present a theoretical study of broadcasting in radio networks having dynamic unknown topology. The dynamic network is modeled by means of adversaries: we consider two of them. We first analyze an oblivious, memoryless random adversary that can be seen as the dynamic version of the average-case study presented by Elsässer and Gasieniec in *JCSS*, 2006. We then consider the deterministic worst-case adversary that, at each time slot, can make any network change (thus the strongest adversary). This is the dynamic version of the worst-case study provided by Bar-Yehuda, Goldreich and Itai in *JCSS*, 1992.

In both cases we provide tight bounds on the completion time of randomized broadcast protocols.

1. Introduction

In a *radio* network, every node (station) can directly transmit to some subset of the nodes depending on the power of its transmitter and on the topological characteristics of the surrounding region. When a node u can directly transmit to a node v, we say that there is a (wireless) directed link (u, v). The set of nodes together with the set of these links form a directed communication graph which represents the radio network.

According to the series of previous important theoretical works [4,6,7,10,19,20], the communication is assumed to be synchronous. Synchronous communication allows us to focus on the impact of the *interference* phenomenon on the network performance. When a node sends a message, the latter is sent in parallel on all outgoing edges. On the other hand, a node can receive a message during a time slot iff there is exactly one of its in-coming neighbors that sends the message during that time slot. If two or more neighbors send a message during the same time slot, then a *collision* occurs and the node receives nothing because of the interference phenomenon.

The broadcast task consists of sending a message from a given *source* node to all nodes of the network. Broadcasting is a fundamental communication primitive in radio networks and it is the subject of a large number of research works in both

☆ Partially supported by the European Union under the EU/IST Project 15964 AEOLUS.
☆☆ A preliminary version of this paper appeared in *Proc. of 26th Annual ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing (PODC'07)*.
* Corresponding author.
E-mail addresses: clementi@mat.uniroma2.it (A.E.F. Clementi), monti@di.uniroma1.it (A. Monti), pasquale@mat.uniroma2.it (F. Pasquale), silvestri@di.uniroma1.it (R. Silvestri).
1 Partially supported by the European Union under the EU/IST Project 15964 AEOLUS.

0022-0000/$ – see front matter
doi:10.1016/j.jcss.2008.10.004

algorithmic and networking areas [4,6,7,19,20]. The *completion time* of a broadcast protocol in a synchronous network is the number of time slots required by the protocol to inform all nodes. A node is *informed* if it has received the source message.

It is reasonable to claim that almost all major questions related to radio broadcasting can be considered closed as far as *static* networks are considered: the network never changes during the entire protocol's execution. A series of theoretical works establishes tight bounds on the completion time of broadcasting that strongly depend on what nodes know about the graph and on the kind of the protocol (see Section 1.2).

On the other hand, theoretical results on communication protocols in any scenario where the network topology may change during the protocol's execution (i.e. a *dynamic* radio network) are very few [21] (see Section 1.2). It is not even known whether a randomized broadcast protocol exists that has *finite* expected completion time in arbitrary dynamic radio networks.

1.1. Our contribution

We follow a high-level approach to investigate broadcasting in dynamic radio networks by considering general *adversarial networks* [1,3,21]. We study networks where edges change during each time slot according to some adversarial strategy. We investigate two somewhat extremal adversaries. A *weak random* adversary where dynamic changes are fully random and memoryless (thus oblivious), and a *strong worst-case* adversary where arbitrary dynamic changes are deterministically and adaptively chosen at each time slot. Such two extremal scenarios do not find immediate applications on real radio networks. However, a tight analysis of them allows us to draw the range spanned by broadcast completion time against *any* dynamic adversary strategy. Moreover, such extremal choices about the adversary aim to answer two fundamental questions: *1.* Do dynamic scenarios *always* constitute a *hurdle* for radio communication? *2.* How hard can radio communication be against *worst-case* adversaries?

We will consider general randomized protocols as well as *non-spontaneous oblivious* protocols. The latter are easy-to-implement and energy-efficient so they are very suitable for radio networks. In such protocols only informed nodes are active and any action of an informed node i, at time slot t, depends only on i and t. So, in oblivious protocols, the actions of an informed node do not depend on any information received during the execution of the protocol. An even more restricted class of protocols is that of homogeneous ones: A protocol is said to be *homogeneous* if it is non-spontaneous and the transmission probability of every informed node i at time slot t depends only on t. Observe that when decisions must be oblivious and the topology is unknown there seems to be no reason to a priori distinguish the strategy of two nodes.

The weak random adversary. The dynamic network is modeled by an oblivious random process defined as follows. At each time slot t of the execution of the protocol, a (new) graph G_t is selected according to the well-known random graph model $\mathcal{G}_{n,p}$ where n is the number of nodes and p is the edge probability [2,5]. This adversarial strategy will be simply denoted as *dynamic* $\mathcal{G}_{n,p}$. This model can be considered as the dynamic version of random networks studied by Elsässer and Gasieniec in [11] (see also Section 1.2).

We first assume that protocols know p. For any probability $p \geqslant 1/n$, we provide a randomized oblivious protocol that, with high probability (in short *w.h.p.*), completes radio broadcasting in a dynamic $\mathcal{G}_{n,p}$ in $O(\log n)$ time slots (we say that an event occurs *with high probability* if it happens with probability at least $1-n^{-\Theta(1)}$). This bound is tight: we indeed prove that, for any $p < 1-\epsilon$ (where $\epsilon < 1$ is any positive constant), *any* randomized protocol completes radio broadcasting in a dynamic $\mathcal{G}_{n,p}$ in $\Omega(\log n)$ expected time. So, the lower bound holds for spontaneous, non-oblivious randomized protocols too.

We then consider the case when protocols do not know p: the adversary, based on the protocol strategy, can choose p in order to minimize the probability of successful communications. Clearly, the above logarithmic lower bound holds. We first show that a simple, homogeneous version of the Bar-Yehuda–Goldreich–Itai's (BGI's) protocol [4] has $O(\log^2 n)$ completion time w.h.p., for any probability $p \geqslant 1/n$. Then, we prove that, for any homogeneous randomized protocol, there exists p, with $\frac{\ln^2 n}{n} \leqslant p \leqslant \frac{1}{\ln^5 n}$, so that the protocol completes broadcasting in a dynamic $\mathcal{G}_{n,p}$ in $\Omega(\log^2 n/\log\log n)$ expected time.

Let us observe that the above protocols work in logarithmic time even when, at every time slot, the expected node degree is 1 and the radio network is w.h.p. disconnected (the latter happens whenever $p = o(\log n/n)$ [5]). This makes our upper bounds significantly different from the logarithmic upper bound for static random graphs [11] that holds only for $p = \Omega(\log^{1+\varepsilon} n/n)$ (see Section 1.2).

We thus answer to Question *1* above by providing the first rigorous proof of the fact that oblivious fully-random network changes, instead of working as a hurdle, *help* information propagation. This is rather surprising due to the unpredictable collisions yielded by dynamic radio networks.

The strong worst-case adversary. We investigate adversaries that can make any network change and that are adaptive, i.e., their actions at time slot t depend on the execution of the protocol and on the state of the network till time slot $t-1$. However, the adversary must be meaningful. An adversary is *meaningful* if, at any time slot, it keeps at least one link on from an informed node to a non-informed one. This condition is a *minimal* one: the completion time of any protocol against non-meaningful worst-case adversaries is clearly infinite. Observe that "meaningfulness" is much weaker than (global) graph connectivity, a condition commonly adopted in all previous works on this topic. In the sequel, meaningful worst-case adversaries will be simply called worst-case adversaries.

Table 1

	Random graphs	Worst-case graphs
Static	$\Theta(\log n)$, $\forall p \geqslant \frac{\ln^{1+\varepsilon} n}{n}$ [11]	$\Theta(n)$ [10,13]
Dynamic	$\Theta(\log n)$, $\forall p \geqslant \frac{1}{n}$ [this paper]	$\Theta(\frac{n^2}{\log n})$ [this paper]

An alternative view of our worst-case model is an adversary that can make *any* change but the completion time of a protocol counts only useful time slots: in a *useful* time slot, at least one link must exist from an informed node to a non-informed one.

It is important to observe that, for any deterministic protocol, there is a worst-case adversary that, at each time slot, yields a connected graph (so the adversary is meaningful) over which the protocol never completes broadcasting. Indeed, consider a network scenario where two nodes are informed and one is still not informed. It is then easy to see that, for any deterministic protocol, there is an adversary strategy that keeps always at least one link on from an informed node to the non-informed one and, exploiting collisions, it avoids the message to arrive at the non-informed node.

We instead show that the use of randomness makes broadcasting (against the worst-case adversary) feasible and relatively efficient. We present a simple oblivious randomized protocol that, for any worst-case adversary, completes broadcasting in $O(n^2/\log n)$ time, w.h.p. Then we prove this upper bound to be optimal for *any* randomized protocol (so, again, for spontaneous and non-oblivious ones too). Such results thus provide the first rigorous answer to our second fundamental question. In particular, our quadratic upper bound implies that no meaningful adversary exists that forces *exponential* broadcast completion time.

A comparison between our results for dynamic networks and those known for static networks of unknown topology is summarized in Table 1 (all results concern randomized protocols).

Finally, we emphasize that our work significantly departs from all previous theoretical works on this topic in two important issues:

- In some theoretical studies [15–17], dynamic network models are considered where nodes and edges may change at any time slot. However, such changes are somewhat locally *declared* in the *previous* time slot. Instead, our work investigates highly-dynamic networks in which the next changes are completely unknown to the protocol.
- To the best of our knowledge, all previous theoretical studies on broadcasting in dynamic radio networks of unknown topology assume the network is connected during *all* time slots of the protocol. Our results show this assumption is too strong: information propagation can go on successfully even under much weaker conditions against both random and worst-case adversaries.

1.2. Related theoretical works

Static networks. For brevity's sake, we here consider only theoretical results on general networks of unknown topology. The best-known deterministic protocol for radio networks has $O(n\log^2 n)$ completion time, which is proved in [7]. Then, in [8] an $\Omega(n\log D)$ lower bound is shown on the completion time of any deterministic protocol, where D is the source eccentricity. When nodes know D then a protocol working in $O(n\log^2 D)$ time is presented in [10].

In [4], a randomized protocol is proposed, denoted here as BGI's protocol, that completes broadcasting in $O(D\log n + \log^2 n)$, w.h.p. Then, in [10] an improved version of the BGI's protocol is presented obtaining completion time $O(D\log(n/D) + \log^2 n)$, w.h.p. On the other hand, in [13] a lower bound $\Omega(D\log(n/D))$ is shown. Finally, broadcasting in *static random* graphs $\mathcal{G}_{n,p}$ has been recently studied in [11]. A $\Theta(\log n)$ bound is proved for oblivious randomized protocols. The upper bound holds for any choice of $p \geqslant \ln^{1+\epsilon} n/n$, so graphs are w.h.p. connected.

Dynamic networks. A theoretical study of broadcasting in a class of dynamic radio networks is presented in [9]. The results concern deterministic protocols and they are stated in terms of *fault-tolerance*. At each time slot, the deterministic *adversary* decides a fault pattern starting from an initial graph of *known* topology. The worst-case analysis is then made on the *residual* graph, i.e., the connected subgraph (containing the source) of the initial graph that has been *always* fault-free. It is proved that the *round robin* strategy is asymptotically optimal thus getting an optimal bound $\Theta(Dn)$. Then for graphs of maximal in-degree Δ, a deterministic protocol is presented having completion time $O(D\Delta\log^3 n)$. Deterministic broadcasting in faulty radio networks of *known* topology is studied in [18]. An initial graph is given and, at each time slot, every node is faulty with probability p, where p is a fixed positive *constant* such that $0 < p < 1$. A completion time of $O(opt\log n)$ is shown where OPT is the optimal completion time in the fault-free case.

Broadcasting on *highly-dynamic graphs* is studied in [17]. The adversary can arbitrarily change the edges of the graph at each time slot but the graph must be always connected. A further critical assumption is that each node is somewhat *previously* informed about any change in its neighborhood and it can act accordingly. The main result is the existence of deterministic protocols that complete broadcasting in $O(n^2)$ (worst-case) completion time.

Finally, reliable broadcasting over mobile *grid* networks is studied in [15,16]. At each time slot, a node can move from one grid point to an arbitrary adjacent one. A lower bound $\Omega(D\log n)$ for the line grid and an $\Omega(n\log n)$ lower bound for the square grid are proved in [15]. Then, a protocol is provided in [16] that completes broadcasting on the line grid within

$O(D \log n)$ time slots. We emphasize that the local node-mobility is somewhat previously "known" by every node in this model too. Time slots are not homogeneous: during a *control slot*, the nodes declare their next moves.

2. Random adversary

For any n and for any probability parameter p, the dynamic random graph, denoted as *dynamic* $\mathcal{G}_{n,p}$, is an infinite sequence of random graphs

$$G_0, G_1, \ldots, G_t, \ldots$$

where each G_t is independently selected according to the *random graph model* $\mathcal{G}_{n,p}$ [5]. A *random graph* $\mathcal{G}_{n,p}$ is an undirected graph $G(V,E)$ where V is the set of n nodes and the probability that $(i,j) \in E$ is equal to p. In the sequel p will denote the edge probability of random graphs. A broadcast protocol running on a dynamic $\mathcal{G}_{n,p}$, at any time slot t, works in graph G_t.

We distinguish two cases depending on whether or not the protocol knows the probability p.

2.1. Case p known

We now present an oblivious randomized protocol that makes use of an oblivious version (the third loop below) of the BGI's *Decay* procedure [4].

```
DynBroad(n,p)
```

for $\lceil c \ln n \rceil$ time slots (where c is a suitable constant)
    The source sends the message;
for $\lceil c \ln n \rceil$ time slots
    Each informed node sends the message;
for $k = 0, 1, \ldots \lceil \ln n \rceil$
    Each informed node sends the message with probability $q = e^{-k}$
for $\lceil c \ln n \rceil$ time slots
    Each informed node sends the message with probability $q = 1/(np)$

The protocol clearly terminates within $O(\log n)$ time slots. In what follows we will show that, for $p \geqslant 1/n$, `DynBroad(n,p)` completes broadcasting in a dynamic $\mathcal{G}_{n,p}$, w.h.p. The proof evaluates the number of informed nodes after each of the four loops of the protocol. Note that the analysis significantly departs from those in [4] and [11] for static unknown graphs.

Lemma 2.1 *(First loop). Assume that the source sends the message for $c \ln n$ time slots, with $c > 1$.*

- *If $p \geqslant 1/n$ then at least $\ln n$ nodes will be informed w.h.p.*
- *If $p \geqslant 1/\ln n$ then at least $n/2$ nodes will be informed w.h.p.*

Proof. For each node $i = 1, 2, \ldots, n$ other than the source, let X_i be the random variable whose value is 1 if node i is informed within $c \ln n$ time slots and 0 otherwise. It holds that

$$\Pr\{X_i = 0\} = (1-p)^{c \ln n} \leqslant e^{-cp \ln n}.$$

Consider the random variable $X = \sum_{i=1}^{n} X_i$ counting the number of informed nodes after $c \ln n$ time slots.

If $p \geqslant 1/n$ we have

$$e^{-cp \ln n} \leqslant e^{-c \frac{\ln n}{n}} \leqslant 1 - \frac{c \ln n}{2n}.$$

Hence $\Pr\{X_i = 1\} \geqslant \frac{c \ln n}{2n}$ for each $i = 1, 2, \ldots, n$. The expected value of X is $\mathbf{E}[X] \geqslant \frac{c}{2} \ln n$. Since X_i's are independent, by Chernoff's bound (see (A.2) in Appendix A) with $\mu = \frac{c}{2} \ln n$ and $\delta = 1 - \frac{2}{c}$, it holds that

$$\Pr\{X \leqslant \ln n\} \leqslant e^{-\alpha \ln n} = \frac{1}{n^{\alpha}}$$

where α is a positive constant.

If $p \geqslant 1/\ln n$, we have $e^{-cp \ln n} \leqslant e^{-c}$. Hence $\Pr\{X_i = 1\} \geqslant 1 - e^{-c}$ for each $i = 1, 2, \ldots, n$. The expected value of X is $\mathbf{E}[X] \geqslant (1 - e^{-c})n$. Since X_i's are independent, by Chernoff's bound (see (A.2) in Appendix A) with $\mu = (1 - e^{-c})n$ and $\delta = \frac{1}{2}(1 - \frac{1}{e^c - 1})$, it holds that $\Pr\{X \leqslant n/2\} \leqslant e^{-\alpha n}$ where α is a positive constant. □

Lemma 2.2 *(Second loop). Let p be such that $1/n \leqslant p \leqslant 1/\ln n$. Assume we start with at least $\frac{1}{4}\ln n$ informed nodes and that at each time slot every informed node sends the message. Then, after $c\ln n$ time slots, at least $\frac{1}{2p}$ nodes are informed w.h.p.*

Proof. Let m_k be the number of informed nodes at time slot k of the second loop. By hypothesis we have $m_0 \geqslant \frac{1}{4}\ln n$ and so $m_k \geqslant \frac{1}{4}\ln n$ for every k. Consider the events

$$\mathcal{E}_k = \text{“}m_k \geqslant (1+a)m_{k-1}\text{”},$$
$$\mathcal{Z}_k = \text{“}m_{k-1} \leqslant \frac{1}{2p}\text{”},$$
$$\mathcal{F}_k = \text{“}m_k < \frac{1}{4}(1+a)^k \ln n\text{”}$$

where a is a positive constant that will be given later. Let us observe how the events are mutually related. First of all it holds that

$$\mathcal{Z}_{k+1} \subseteq \mathcal{Z}_k. \tag{1}$$

In fact if the informed nodes are less then $\frac{1}{2p}$ in time slot k they were less than $\frac{1}{2p}$ in time slot $k-1$. Moreover

$$\bigcap_{i=1}^{k} \mathcal{E}_i \subseteq \overline{\mathcal{F}_k}.$$

Indeed, if inequality $m_i \geqslant (1+a)m_{i-1}$ holds for every time slot $i = 1, \dots, k$ and, since $m_0 \geqslant \frac{1}{4}\ln n$, then $m_k \geqslant \frac{1}{4}(1+a)^k \ln n$. By looking at complementary sets we have

$$\mathcal{F}_k \subseteq \overline{\left(\bigcap_{i=1}^{k} \mathcal{E}_i\right)} = \bigcup_{i=1}^{k} \overline{\mathcal{E}_i}. \tag{2}$$

Finally, observe that for every $k \leqslant \frac{\ln \frac{2}{p\ln n}}{\ln(1+a)}$ (i.e. when $\frac{1}{4}(1+a)^k \ln n \leqslant \frac{1}{2p}$) it holds that $\mathcal{F}_k \subseteq \mathcal{Z}_{k+1}$.

Set $\hat{k} = \lceil \frac{\ln \frac{2}{p\ln n}}{\ln(1+a)} \rceil$ and observe that $\hat{k} \leqslant c\ln n$ for a suitable constant $c > 0$. Then it holds that

$$\begin{aligned}
\Pr\left\{m_{\lceil c\ln n\rceil} < \frac{1}{2p}\right\} &\leqslant \Pr\left\{m_{\hat{k}} \leqslant \frac{1}{2p}\right\} = \Pr\{\mathcal{F}_{\hat{k}} \cap \mathcal{Z}_{\hat{k}+1}\} \leqslant \Pr\{\mathcal{F}_{\hat{k}} \cap \mathcal{Z}_{\hat{k}}\} \\
&\leqslant \Pr\left\{\left(\bigcup_{i=1}^{\hat{k}} \overline{\mathcal{E}_i}\right) \cap \mathcal{Z}_{\hat{k}}\right\} = \Pr\left\{\bigcup_{i=1}^{\hat{k}} (\overline{\mathcal{E}_i} \cap \mathcal{Z}_{\hat{k}})\right\} \leqslant \sum_{i=1}^{\hat{k}} \Pr\{\overline{\mathcal{E}_i} \cap \mathcal{Z}_{\hat{k}}\} \\
&\leqslant \sum_{i=1}^{\hat{k}} \Pr\{\overline{\mathcal{E}_i} \cap \mathcal{Z}_i\} = \sum_{i=1}^{\hat{k}} \Pr\{\overline{\mathcal{E}_i} \mid \mathcal{Z}_i\} \Pr\{\mathcal{Z}_i\} \leqslant \sum_{i=1}^{\hat{k}} \Pr\{\overline{\mathcal{E}_i} \mid \mathcal{Z}_i\}.
\end{aligned}$$

Where from the first to the second line we used (2), from the second to the third line we used (1). The next claim implies that

$$\Pr\{\overline{\mathcal{E}_i} \mid \mathcal{Z}_i\} \leqslant \frac{1}{n^\gamma} \quad \text{for } i = 1, 2, \dots. \tag{3}$$

Claim 1. *If the number m of informed nodes is such that $\frac{1}{4}\ln n \leqslant m \leqslant \frac{1}{2p}$, then a positive constant α exists such that at least αm new nodes will be informed in one time step w.h.p.*

Proof. For every non-informed node $i = 1, \dots, n-m$ let X_i be the random variable whose value is 1 if node i gets informed and 0 otherwise. It holds

$$\Pr\{X_i = 1\} = mp(1-p)^{m-1} \geqslant \frac{m}{n}(1-p)^{m-1} \geqslant \frac{m}{3n}$$

where we used the hypothesis on m and p. Let $X = \sum_{i=1}^{n-m} X_i$, we have that

$$\mathbf{E}[X] \geqslant (n-m)\frac{m}{3n} \geqslant \frac{m}{6}$$

where we used that $m \leqslant \frac{1}{2p} \leqslant \frac{n}{2}$. By using Chernoff's bound (see (A.2)) with $\delta = 1/2$ and hypothesis $m \geqslant \frac{1}{4}\ln n$ we get

$$\Pr\{X \leqslant \alpha m\} \leqslant \frac{1}{n^{\gamma}}$$

with α and γ positive constants. □

Finally, the thesis follows from (3). □

In the sequel we will use the following result.

Fact 2. *Let $m < n$ be the number of informed nodes and let u be an uninformed one. If every informed node sends the message with probability q, then in $\mathcal{G}_{n,p}$ the probability that node u receives the message is*

$$mpq(1-pq)^{m-1}.$$

Proof. Let T be the random variable counting the number of transmitting nodes and let X be the random variable whose value is 1 if node u gets the message and 0 otherwise. Then

$$\begin{aligned}
\Pr\{X=1\} &= \sum_{j=0}^{m} \Pr\{X=1 \mid T=j\} \Pr\{T=j\} \\
&= \sum_{j=1}^{m} jp(1-p)^{j-1} \binom{m}{j} q^j (1-q)^{m-j} \\
&= pq \sum_{j=1}^{m} j \binom{m}{j} \big(q(1-p)\big)^{j-1} (1-q)^{m-j} \\
&= mpq \sum_{j=1}^{m} \binom{m-1}{j-1} \big(q(1-p)\big)^{j-1} (1-q)^{m-j} \\
&= mpq\big(q(1-p)+1-q\big)^{m-1} = mpq(1-pq)^{m-1}. \qquad \square
\end{aligned}$$

Lemma 2.3 *(Third loop). Let p be such that $1/n \leqslant p \leqslant 1/\ln n$. Assume we start with m informed nodes, with $m \geqslant \frac{1}{2p}$. If, for $k = 0, 1, 2, \ldots, \lceil \ln n \rceil$, every informed node sends the message with probability e^{-k}, then at least γn nodes are informed w.h.p., for some positive constant γ.*

Proof. Let m_k be the number of informed nodes at time slot k of this phase. From lemma's hypothesis it holds that $m_0 \geqslant 1/(2p)$. Let $h = \lceil \ln(np) \rceil$ and consider the first $k = 1, \ldots, h$ time slots. If for each of them it holds $k < \ln(2m_k p)$, then, in the last one of them, it holds $\lceil \ln(np) \rceil < \ln(2m_h p)$. So $m_h \geqslant n/2$. Otherwise a time slot k must exist such that $k = \lceil \ln(2m_k p) \rceil$. In what follows, we only care about time slot k and we simply denote the number of informed nodes during this time slot as m. First of all, note that the transmission probability $q = e^{-k}$ of the informed nodes satisfies

$$\frac{1}{2emp} \leqslant q \leqslant \frac{1}{2mp}. \tag{4}$$

Consider the $n-m$ non-informed nodes and let X_i, $i = 1, \ldots, n-m$, be the random variable whose value is 1 if node i is informed in time slot k and 0 otherwise. From Fact 2 and (4) it holds that

$$\Pr\{X_i = 1\} = mpq(1-pq)^{m-1} \geqslant \frac{1}{2e}\left(1 - \frac{1}{2m}\right)^{m-1} \geqslant \frac{1}{2e} e^{-\frac{m-1}{2m-1}} \geqslant \frac{1}{2e} e^{-1/2}.$$

Now consider $X = \sum_{i=1}^{n-m} X_i$. If there are $m \geqslant n/2$ informed nodes the lemma is proved, otherwise the expected value of X is

$$\mathbf{E}[X] = \frac{n-m}{2e^{3/2}} \geqslant \frac{n}{4e^{3/2}}.$$

In order to prove that, after time slot k, the total number of informed nodes is a constant fraction of n (w.h.p.), we cannot apply Chernoff's bound on X since $X_1, X_2, \ldots, X_{n-m}$ are not independent. We thus need to introduce the random variable T counting the number of nodes that send the source message. Since there are m informed nodes, each one is sending the message independently with probability q, it holds that T is a binomial random variable, i.e.

$$\Pr\{T = j\} = \binom{m}{j} q^j (1-q)^{m-j}, \quad j = 0, 1, \ldots, m \text{ and } \mathbf{E}[T] = mq.$$

From Chernoff's bound (see (A.3) in Appendix A) with $\mu = mq$ and $\delta = 1/2$ we have

$$\Pr\left\{T \notin \left[\frac{1}{2}mq, \frac{3}{2}mq\right]\right\} \leqslant 2e^{-\frac{1}{12}mq} \leqslant 2e^{-\frac{1}{24e}\frac{1}{p}} \leqslant \frac{2}{n^{\frac{1}{24e}}}$$

where in the last two inequalities we used $q \geqslant \frac{1}{2emp}$ and $p \leqslant \frac{1}{\ln n}$. For $j = 0, 1, \ldots, m$ and for $i = 1 \ldots, n-m$, define the conditioned random variables X_i^j that equals X_i under the event $T = j$. Observe that

$$\Pr\{X_i^j = 1\} = \Pr\{X_i = 1 \mid T = j\} = jp(1-p)^{j-1}.$$

Moreover, for each fixed $j = 0, 1, \ldots, m$, we note that $X_1^j, \ldots, X_{n-m}^j$ are independent. When j is close to the expected value of T, i.e. when $\frac{1}{2}mq \leqslant j \leqslant \frac{3}{2}mq$, it holds that

$$\Pr\{X_i^j = 1\} = jp(1-p)^{j-1} \geqslant \frac{1}{2}mqp(1-p)^{\frac{3}{2}mq-1} \geqslant \frac{1}{4e}e^{-\frac{p}{1-p}(\frac{3}{4p}-1)} \geqslant \beta$$

where β is a positive constant. Hence for each $j \in [\frac{1}{2}mq, \frac{3}{2}mq]$ the random variable $X^j = \sum_{i=1}^{n-m} X_i^j$ has expectation

$$\mathbf{E}[X^j] \geqslant \beta(n-m) \geqslant \frac{\beta}{2}n.$$

From Chernoff's bound (see (A.2) in Appendix A) with $\mu = \frac{\beta}{2}n$ and $\delta = 1/2$, it follows that

$$\Pr\left\{X^j \leqslant \frac{\beta}{4}n\right\} \leqslant e^{-\frac{\beta}{16}n}.$$

We can now go back to the random variable X and obtain

$$\begin{aligned}
\Pr\left\{X \geqslant \frac{\beta}{4}n\right\} &= \sum_{j=0}^{m} \Pr\left\{X^j \geqslant \frac{\beta}{4}n\right\} \Pr\{T = j\} \\
&\geqslant \sum_{j \in [\frac{1}{2}mq, \frac{3}{2}mq]} \Pr\left\{X^j \geqslant \frac{\beta}{4}n\right\} \Pr\{T = j\} \\
&\geqslant \left(1 - e^{-\frac{\beta}{16}n}\right) \sum_{j \in [\frac{1}{2}mq, \frac{3}{2}mq]} \Pr\{T = j\} \\
&= \left(1 - e^{-\frac{\beta}{16}n}\right) \Pr\left\{T \in \left[\frac{1}{2}mq, \frac{3}{2}mq\right]\right\} \\
&\geqslant \left(1 - e^{-\frac{\beta}{16}n}\right)\left(1 - \frac{2}{n^{1/(24e)}}\right) \geqslant 1 - \frac{1}{n^{\varepsilon}}
\end{aligned}$$

for a suitable positive constant ε. □

The next lemma is used for the analysis of the fourth protocol's loop. In next subsection, it will be used to analyze another protocol for the case that p is unknown as well.

Lemma 2.4 *(Fourth loop). Let $p \geqslant 1/n$ and let γ be a constant such that $0 < \gamma < 1$. Assume we start with at least γn informed nodes and, at each time slot, every informed node sends the message with probability q such that $\frac{1}{enp} \leqslant q \leqslant \frac{1}{np}$. Then, after $c \ln n$ of such time slots, all nodes are informed w.h.p.*

Proof. Let m_0 be the number of informed nodes at the beginning. From lemma's hypothesis we have that $m_0 \geqslant \gamma n$. At any successive time slot, the number m of informed nodes satisfies $\gamma n \leqslant m_k \leqslant n$. From Fact 2, the probability that a non-informed node receives the message during the time slot is

$$mpq(1-pq)^{m-1} \geqslant \frac{\gamma}{e}\left(1 - \frac{1}{n}\right)^{n-1} \geqslant \frac{\gamma}{e^2} \tag{5}$$

where we used $\frac{1}{enp} \leqslant q \leqslant \frac{1}{np}$. At the beginning we have $n - m_0$ non-informed nodes. For $i = 1, 2, \ldots, n - m_0$, consider the event

$$\mathcal{E}_i = \{\text{node } i \text{ does not receive the message within } c \ln n \text{ time slots}\}.$$

From (5), it holds that

$$\Pr\{\mathcal{E}_i\} \leqslant \left(1 - \frac{\gamma}{e^2}\right)^{c \ln n} \leqslant e^{\gamma c/e^2 \ln n} = \frac{1}{n^{\gamma c/e^2}}.$$

Then, the probability that, after $c \ln n$ time slots, there exists a non-informed node is

$$\Pr\{\exists i\colon\ \mathcal{E}_i\} \leqslant \sum_{i=1}^{n-m_0} \Pr\{\mathcal{E}_i\} \leqslant \frac{\gamma}{n^{\gamma c/e^2 - 1}}.$$

Finally, by setting $c > e^2/\gamma$, it holds that all nodes are informed w.h.p. $\square$

Theorem 2.5. *Let $p \geqslant 1/n$. Protocol `DynBroad(n,p)` completes broadcasting in a dynamic $\mathcal{G}_{n,p}$, w.h.p.*

Proof. Two cases may arise. If $p \geqslant 1/\ln n$ then, from Lemma 2.1, after the first loop there are at least $n/2$ informed nodes w.h.p. Otherwise, if $1/n \leqslant p \leqslant 1/\ln n$, from Lemmas 2.1 and 2.2, after the second loop there are at least $1/(2p)$ informed nodes w.h.p. Then, from Lemma 2.3, after the third loop, there are at least γn informed nodes w.h.p.

So, in both cases, after the third loop, there are at least γn informed nodes w.h.p. where γ is a positive constant, $0 < \gamma < 1$. Then, from Lemma 2.4, after the last loop, all the n nodes are informed w.h.p. $\square$

We observe that if p is $1 - o(1)$ the broadcast task can be completed in $o(\log n)$ time by considering the simple protocol where only the source transmits with probability 1 (e.g. if $p = 1 - 1/n^2$ broadcasting is completed w.h.p. in one time slot). Instead when p does not tend to 1 as n goes to infinity, we now show a lower bound matching the previous logarithmic upper bound.

To this aim, we need the following lemma that will be used for the case that p is unknown as well.

Lemma 2.6. *Let ε be any positive constant and let $p \leqslant 1 - \varepsilon$. Consider any broadcast protocol in a dynamic $\mathcal{G}_{n,p}$. Let m be the number of informed nodes at a given time slot and let m' be the number of informed nodes at the successive time slot. If $m \leqslant n - \frac{18 \ln n}{1-p}$ then it holds that*

$$\Pr\left\{m' - m > \left(1 - \frac{1-p}{2e}\right)(n-m)\right\} \leqslant \frac{1}{n}.$$

Proof. Consider any time slot $t \geqslant 1$ of the protocol's execution and let m and m' be the number of informed nodes at time slot t and $t+1$, respectively. For any k, under the condition that exactly k nodes transmit at time slot t, we define, for each node j, the 0–1 random variable X_j^k that is equal to 1 if node j is not informed at time slot t and it is informed at time slot $t+1$. It is easy to verify that

$$\Pr\{X_j^k = 1\} = \begin{cases} kp(1-p)^{k-1} & \text{if } j \text{ is not informed at time } t, \\ 0 & \text{otherwise.} \end{cases}$$

As for $\tilde{p}_k = kp(1-p)^{k-1}$, we get, for all k,

$$\tilde{p}_k \leqslant \max\{p, 1 - 1/e\} \leqslant 1 - \frac{1-p}{e}. \tag{6}$$

Let $X^k = \sum_{j=1}^n X_j^k$, then the expected value $\mathbf{E}[X^k] = (n-m)\tilde{p}_k$. Observe that, for any fixed k, the random variables $X_1^k, \dots, X_n^k$ are independent. From (6), Corollary 3 (in Appendix A) and lemma's hypothesis it follows that

$$\Pr\left\{X^k \geqslant \left(1 - \frac{1-p}{2e}\right)(n-m)\right\} \leqslant \Pr\left\{X^k \geqslant \frac{1+\tilde{p}_k}{2}(n-m)\right\}$$
$$\leqslant e^{-\lambda(1-\tilde{p}_k)(n-m)} \leqslant e^{-\lambda \frac{1-p}{e}(n-m)} \leqslant e^{-\frac{18\lambda}{e} \ln n} \leqslant \frac{1}{n}$$

where $\lambda = (1 - \ln 2)/2$ as in Corollary 3. Let T be the random variable counting the number of nodes that transmit at time slot t. We then get

$$\Pr\left\{m' - m \geqslant \left(1 - \frac{1-p}{2e}\right)(n-m)\right\}$$
$$= \sum_{k=0}^{n} \Pr\left\{X^k \geqslant \left(1 - \frac{1-p}{2e}\right)(n-m)\right\} \Pr\{T = k\}$$
$$\leqslant \frac{1}{n} \sum_{k=0}^{n} \Pr\{T = k\} = \frac{1}{n}. \qquad \square$$

Theorem 2.7. *Let ε be any positive constant and let $p \leqslant 1-\varepsilon$. Any broadcast protocol in a dynamic $\mathcal{G}_{n,p}$ has expected completion time $\Omega(\log n)$.*

Proof. Let m_t be the random variable counting the number of informed nodes at time slot $t \geqslant 0$ and define the event $\mathcal{E}_t$ as

$$\mathcal{E}_t = \left\{ m_{t+1} - m_t < \left(1 - \frac{1-p}{2e}\right)(n - m_t) \right\}.$$

Let us assume that the events $\mathcal{E}_0, \ldots, \mathcal{E}_{t-1}$ hold. Then, from $m_0 = 1$, we get

$$m_t \leqslant n - (n-1)\left(\frac{1-p}{2e}\right)^t. \tag{7}$$

A constant $\alpha > 0$ (depending only on ϵ) exists such that if $t \leqslant \alpha \ln n$,

$$n - (n-1)\left(\frac{1-p}{2e}\right)^t \leqslant n - \frac{18 \ln n}{1-p}.$$

Hence, for any $t \leqslant \alpha \ln n$, Lemma 2.6 implies that

$$\Pr\left\{ \mathcal{E}_t \,\Big|\, \bigcap_{i=1}^{t-1} \mathcal{E}_i \right\} > 1 - \frac{1}{n}.$$

It follows that, for any $t \leqslant \alpha \ln n$,

$$\Pr\left\{ \bigcap_{i=1}^{t} \mathcal{E}_i \right\} = \prod_{i=1}^{t} \Pr\left\{ \mathcal{E}_i \,\Big|\, \bigcap_{j=1}^{i-1} \mathcal{E}_j \right\} \geqslant \left(1 - \frac{1}{n}\right)^t \geqslant \left(1 - \frac{1}{n}\right)^{\alpha \ln n} \geqslant e^{-\frac{\alpha \ln n}{n-1}} \geqslant e^{-2\alpha}.$$

Hence, there is a positive constant probability that broadcasting is not completed within $\alpha \ln n$ time. □

2.2. Case p unknown

Let us consider the following homogeneous variant of the BGI's Decay procedure [4] denoted as `BGI(n)`.

```
BGI(n)
for ⌈c ln n⌉ time slots (where c is a suitable constant)
    for k = 0, 1, ..., ⌈ln n⌉
        Each informed node sends the message with probability q = e^{-k}
```

Protocol `BGI(n)` terminates within $O(\log^2 n)$ time slots. Now we show that it completes broadcasting in a dynamic $\mathcal{G}_{n,p}$, w.h.p.

In the following, we call a *phase* any execution of the inner **for**-loop of the protocol `BGI(n)`.

Lemma 2.8. *If $p \geqslant 1/n$ then there is a constant α such that after the execution of the first $\alpha \ln n$ phases of the protocol `BGI(n)` at least $\ln n$ nodes get informed w.h.p.*

Proof. If $p \geqslant \frac{\ln n}{n}$, by a straightforward application of Chernoff's bound A.1, the thesis follows (in one time slot). On the other hand if $1/n \leqslant p < \frac{\ln n}{n}$ the proof is much harder. Indeed, if we try to evaluate the number of new informed nodes in every step, we cannot get any *with high probability* bound, since the number of informed nodes is too small. We thus need to consider the number of new informed nodes after a logarithmic number of steps. Unfortunately, a straightforward analysis does not work since the involved random variables are not independent.

Thus let $\mathcal{E}$ be the event that occurs if the number of informed nodes after the first $\alpha \ln n$ phases is less than $\ln n$. The constant $\alpha > 1$ will be determined later. We will prove that $\Pr\{\mathcal{E}\} \leqslant 1/n$.

The phases are numbered starting from 1. For any $\ell = 1, 2, \ldots$, let $\mathcal{Z}_\ell$ be the event that occurs if during the phase ℓ no new nodes get informed. It is immediate to see that if $\mathcal{E}$ occurs then there are more than $\alpha \ln n - \ln n$ phases, among the first $\alpha \ln n$ ones, for which $\mathcal{Z}_\ell$ occurs. That is, there exists $B \in \mathbf{B}$ where

$$\mathbf{B} = \left\{ A \subseteq \left\{1, \ldots, \lceil \alpha \ln n \rceil\right\} \colon\ |A| > \alpha \ln n - \ln n \right\}$$

and for every $\ell \in B$, $\mathcal{Z}_\ell$ occurs. This implies that

$$\Pr\{\mathcal{E}\} \leqslant \sum_{B\in\mathbf{B}} \Pr\left\{\bigcap_{\ell\in B} (\mathcal{Z}_\ell \cap \mathcal{E})\right\}.$$

For any $\ell = 1, 2, \ldots,$ let $\mathcal{E}_\ell$ be the event that occurs if the number of informed nodes at the beginning of phase ℓ is less than $\ln n$. Since for every $\ell \in \{1, \ldots, \lceil \alpha \ln n \rceil\}$ $\mathcal{E} \subseteq \mathcal{E}_\ell$, it holds that

$$\Pr\{\mathcal{E}\} \leqslant \sum_{B\in\mathbf{B}} \Pr\left\{\bigcap_{\ell\in B} (\mathcal{Z}_\ell \cap \mathcal{E}_\ell)\right\}.$$

By applying the identity $\Pr(\bigcap_{i=1}^k \mathcal{A}_i) = \prod_{j=1}^k \Pr(\mathcal{A}_j \mid \bigcap_{i=1}^{j-1} \mathcal{A}_i)$, we obtain

$$\Pr\{\mathcal{E}\} \leqslant \sum_{B\in\mathbf{B}} \prod_{\ell\in B} \Pr\left\{\mathcal{Z}_\ell \cap \mathcal{E}_\ell \;\middle|\; \bigcap_{i\in B_\ell} (\mathcal{Z}_i \cap \mathcal{E}_i)\right\}$$

where, for any $B \subseteq \mathbf{B}$ and for any $\ell \in B$, $B_\ell = \{i \in B \mid i < \ell\}$. By applying the identity $\Pr(\mathcal{A} \cap \mathcal{B} \mid \mathcal{H}) = \Pr(\mathcal{A} \mid \mathcal{H} \cap \mathcal{B}) \cdot \Pr(\mathcal{B} \mid \mathcal{H})$, we obtain, for any B and for any ℓ,

$$\begin{aligned}\Pr\left\{\mathcal{Z}_\ell \cap \mathcal{E}_\ell \;\middle|\; \bigcap_{i\in B_\ell} (\mathcal{Z}_i \cap \mathcal{E}_i)\right\} &= \Pr\left\{\mathcal{Z}_\ell \;\middle|\; \bigcap_{i\in B_\ell} (\mathcal{Z}_i \cap \mathcal{E}_i) \cap \mathcal{E}_\ell\right\} \cdot \Pr\left\{\mathcal{E}_\ell \;\middle|\; \bigcap_{i\in B_\ell} (\mathcal{Z}_i \cap \mathcal{E}_i)\right\} \\ &\leqslant \Pr\left\{\mathcal{Z}_\ell \;\middle|\; \bigcap_{i\in B_\ell} (\mathcal{Z}_i \cap \mathcal{E}_i) \cap \mathcal{E}_\ell\right\}.\end{aligned}$$

For any subset of the nodes $F \subseteq \{1, \ldots, n\}$ and for any $\ell = 1, 2, \ldots,$ let $\mathcal{I}_{\ell,F}$ be the event that occurs if the set of informed nodes is F at the beginning of the phase ℓ. Since the events $\mathcal{I}_{\ell,F}$, as F varies over all the subsets, form a partition, it holds that, for any B and for any ℓ,

$$\begin{aligned}&\Pr\left\{\mathcal{Z}_\ell \;\middle|\; \bigcap_{i\in B_\ell} (\mathcal{Z}_i \cap \mathcal{E}_i) \cap \mathcal{E}_\ell\right\} \\ &\quad = \sum_{F\subseteq[n]} \Pr\left\{\mathcal{Z}_\ell \;\middle|\; \bigcap_{i\in B_\ell} (\mathcal{Z}_i \cap \mathcal{E}_i) \cap \mathcal{E}_\ell \cap \mathcal{I}_{\ell,F}\right\} \cdot \Pr\left\{\mathcal{I}_{\ell,F} \;\middle|\; \bigcap_{i\in B_\ell} (\mathcal{Z}_i \cap \mathcal{E}_i) \cap \mathcal{E}_\ell\right\}.\end{aligned}$$

Observe that if either $s \notin F$ or $|F| \geqslant \ln n$ then $\Pr(\mathcal{I}_{\ell,F} \mid \bigcap_{i\in B_\ell} (\mathcal{Z}_i \cap \mathcal{E}_i) \cap \mathcal{E}_\ell) = 0$. Moreover, if $|F| < \ln n$ then $\mathcal{E}_\ell \cap \mathcal{I}_{\ell,F} = \mathcal{I}_{\ell,F}$. It follows that if $F \in \mathbf{F}$ where $\mathbf{F} = \{A \subseteq \{1, \ldots, n\}\colon |A| < \ln n,\ s \in F\}$ then

$$\begin{aligned}&\Pr\left\{\mathcal{Z}_\ell \;\middle|\; \bigcap_{i\in B_\ell} (\mathcal{Z}_i \cap \mathcal{E}_i) \cap \mathcal{E}_\ell\right\} \\ &\quad = \sum_{F\in\mathbf{F}} \Pr\left\{\mathcal{Z}_\ell \;\middle|\; \bigcap_{i\in B_\ell} (\mathcal{Z}_i \cap \mathcal{E}_i) \cap \mathcal{I}_{\ell,F}\right\} \cdot \Pr\left\{\mathcal{I}_{\ell,F} \;\middle|\; \bigcap_{i\in B_\ell} (\mathcal{Z}_i \cap \mathcal{E}_i) \cap \mathcal{E}_\ell\right\}.\end{aligned}$$

Now, a crucial observation comes:

$$\Pr\left\{\mathcal{Z}_\ell \;\middle|\; \bigcap_{i\in B_\ell} (\mathcal{Z}_i \cap \mathcal{E}_i) \cap \mathcal{I}_{\ell,F}\right\} = \Pr\{\mathcal{Z}_\ell \mid \mathcal{I}_{\ell,F}\}.$$

Indeed, the behavior of the execution of the protocol, from the beginning of the phase ℓ onward, only depends on the set F of informed nodes at the beginning of that phase. It does not depend on how F was formed by the previous phases and, in particular, it does not depend on $\bigcap_{i\in B_\ell} (\mathcal{Z}_i \cap \mathcal{E}_i)$, provided that F is given.

For any $\ell = 1, 2, \ldots,$ let $\mathcal{Z}_\ell^0$ be the event that occurs if during the time-slot 0 of the phase ℓ no new nodes get informed. Recall that in time-slot 0 all the informed nodes transmit with probability 1. Since $\mathcal{Z}_\ell \subseteq \mathcal{Z}_\ell^0$, it holds that $\Pr(\mathcal{Z}_\ell \mid \mathcal{I}_{\ell,F}) \leqslant \Pr(\mathcal{Z}_\ell^0 \mid \mathcal{I}_{\ell,F})$. Taking into account all these observations, we obtain

$$\begin{aligned}&\Pr\left\{\mathcal{Z}_\ell \;\middle|\; \bigcap_{i\in B_\ell} (\mathcal{Z}_i \cap \mathcal{E}_i) \cap \mathcal{E}_\ell\right\} \\ &\quad \leqslant \sum_{F\in\mathbf{F}} \Pr\bigl(\mathcal{Z}_\ell^0 \mid \mathcal{I}_{\ell,F}\bigr) \cdot \Pr\left\{\mathcal{I}_{\ell,F} \;\middle|\; \bigcap_{i\in B_\ell} (\mathcal{Z}_i \cap \mathcal{E}_i) \cap \mathcal{E}_\ell\right\}.\end{aligned}$$

And, recalling that the events $\mathcal{I}_{\ell,F}$, as F varies over all the subsets, form a partition, we see that the following inequality holds

$$\Pr\left\{\mathcal{Z}_\ell \;\Big|\; \bigcap_{i\in B_\ell}(\mathcal{Z}_i \cap \mathcal{E}_i)\cap \mathcal{E}_\ell\right\} \leqslant \max_{F\in\mathbf{F}} \Pr\big(\mathcal{Z}_\ell^0 \;\big|\; \mathcal{I}_{\ell,F}\big).$$

It follows that

$$\Pr\{\mathcal{E}\} \leqslant \sum_{B\in\mathbf{B}}\prod_{\ell\in B}\max_{F\in\mathbf{F}} \Pr\big(\mathcal{Z}_\ell^0 \;\big|\; \mathcal{I}_{\ell,F}\big).$$

At this point, we evaluate $\Pr(\mathcal{Z}_\ell^0 \mid \mathcal{I}_{\ell,F})$. It is easy to verify that

$$\Pr\big(\mathcal{Z}_\ell^0 \;\big|\; \mathcal{I}_{\ell,F}\big) = \big(1-|F|p(1-p)^{|F|-1}\big)^{n-|F|}.$$

It is not hard to see that there exists a constant $\gamma < 1$ such that, for any $1/n \leqslant p \leqslant \ln n/n$ and for any $1 \leqslant m < \ln n$, it holds that

$$\big(1-mp(1-p)^{m-1}\big)^{n-m} \leqslant \gamma.$$

It follows that

$$\Pr\{\mathcal{E}\} \leqslant \sum_{B\in\mathbf{B}}\gamma^{|B|} \leqslant \gamma^{(\alpha-1)\ln n}\sum_{k=(\alpha-1)\ln n+1}^{\alpha\ln n}\binom{\alpha\ln n}{k}.$$

Finally, by exploiting standard upper bounds for the tail of the sum of binomial coefficients, it is not hard to find a constant α (depending upon γ) such that

$$\sum_{k=(\alpha-1)\ln n+1}^{\alpha\ln n}\binom{\alpha\ln n}{k} \leqslant \left(\frac{1}{2\gamma}\right)^{(\alpha-1)\ln n}.$$

So, we conclude that

$$\Pr\{\mathcal{E}\} \leqslant \left(\frac{1}{2}\right)^{(\alpha-1)\ln n} \leqslant \frac{1}{n}. \qquad \square$$

Lemma 2.9. *If $p \geqslant 1/\ln n$ then there exists a constant α such that after the execution of the first $\alpha \ln n$ phases of the protocol* `BGI(n)` *all nodes are informed w.h.p.*

Proof. After the first time slot of the protocol, when the source node sends the message, the expected number of informed nodes is $np \geqslant n/\ln n$. By using Chernoff's bound we have that there are at least $\frac{1}{2}\frac{n}{\ln n}$ informed nodes w.h.p.

For every non-informed node i, and for every phase $j = 1, \dots, \lceil\alpha\ln n\rceil$ define event $\mathcal{F}_i^j =$ "Node i is not informed at the end of phase j," and observe that the probability that node i is not informed after $\alpha\ln n$ phases is

$$\Pr\big\{\mathcal{F}_i^{\lceil\alpha\ln n\rceil}\big\} = \Pr\left\{\bigcap_{j=1}^{\lceil\alpha\ln n\rceil}\mathcal{F}_i^j\right\} = \prod_{j=1}^{\lceil\alpha\ln n\rceil}\Pr\big\{\mathcal{F}_i^j \;\big|\; \mathcal{F}_i^{j-1}\big\}. \tag{8}$$

Now we show that $\Pr\{\mathcal{F}_i^j \mid \mathcal{F}_i^{j-1}\}$ is upper bounded by a constant $c < 1$. Let m_k^j be the number of informed nodes in time slot k of phase j, and observe that $\frac{1}{2}\frac{n}{\ln n} \leqslant m_k^j \leqslant n$. As k grows from 1 to $\ln n$, there must exist a time slot such that $k = \lceil\ln m_k^j p\rceil$. Let m be the number of informed nodes in that time slot and note that the transmission probability q is $\frac{1}{emp} \leqslant q \leqslant \frac{1}{mp}$. From Fact 2, for every non-informed node i, the probability to be informed in that time slot is

$$mpq(1-pq)^{m-1} \geqslant \frac{1}{e}\left(1-\frac{1}{m}\right)^{m-1} \geqslant \frac{1}{e^2}.$$

If we define event $\mathcal{H}_i^j =$ "Node i is not yet informed before time slot $k = \lceil\ln m_k^j p\rceil$ of phase j," then it holds that

$$\begin{aligned}\Pr\big\{\mathcal{F}_i^j \;\big|\; \mathcal{F}_i^{j-1}\big\} &= \Pr\big\{\mathcal{F}_i^j\cap\mathcal{H}_i^j \;\big|\; \mathcal{F}_i^{j-1}\big\} = \Pr\big\{\mathcal{F}_i^j \;\big|\; \mathcal{F}_i^{j-1}\cap\mathcal{H}_i^j\big\}\Pr\big\{\mathcal{H}_i^j \;\big|\; \mathcal{F}_i^{j-1}\big\}\\ &\leqslant \Pr\big\{\mathcal{F}_i^j \;\big|\; \mathcal{F}_i^{j-1}\cap\mathcal{H}_i^j\big\} \leqslant 1-\frac{1}{e^2}.\end{aligned}$$

And filling it in (8), the probability that node i is not yet informed after $\alpha\ln n$ time slots is

$$\Pr\big\{\mathcal{F}_i^{\lceil\alpha\ln n\rceil}\big\} \leqslant \left(1-\frac{1}{e^2}\right)^{\lceil\alpha\ln n\rceil} \leqslant e^{-\frac{\alpha}{e^2}\ln n} = \frac{1}{n^{\alpha/e^2}}.$$

Finally, from the union bound, the probability that, after $\alpha \ln n$ time slots, there exists a non-informed node is

$$\Pr\{\exists i\colon\ \mathcal{F}_i^{\lceil \alpha \ln n \rceil}\} = \Pr\left\{\bigcup_i \mathcal{F}_i^{\lceil \alpha \ln n \rceil}\right\} \leqslant n \frac{1}{n^{\alpha/e^2}} = \frac{1}{n^{\alpha/e^2 - 1}}.$$

And the thesis follows by choosing $\alpha > e^2$. □

Theorem 2.10. *Protocol* `BGI(n)` *completes broadcasting in a dynamic $\mathcal{G}_{n,p}$ w.h.p. for any $p \geqslant 1/n$.*

Proof. If $p \geqslant \frac{1}{\ln n}$ then from Lemma 2.9 after $O(\log n)$ phases all nodes are informed w.h.p.

If $1/n \leqslant p \leqslant 1/\ln n$ then from Lemma 2.8, after $O(\log n)$ phases, there are at least $\ln n$ informed nodes. In the next $O(\log n)$ phases of the protocol, consider only the first time slot of the phase, i.e. the time slot where all informed nodes send the message. From Lemma 2.2 it holds that, after such $O(\log n)$ phases, there are at least $1/(2p)$ informed nodes w.h.p. Now, from Lemma 2.3, in the next phase at least γn nodes will be informed w.h.p., with $\gamma > 0$ constant. Finally, in the remaining phases, consider only the time slots where $k = \lceil \ln(np) \rceil$, i.e. the time slots where the transmission probability q is $1/(enp) \leqslant q \leqslant 1/(np)$. Then, thanks to Lemma 2.4 all nodes will be informed w.h.p. within $O(\log n)$ phases. □

When a homogeneous randomized protocol does not know p, the adversary can choose it in order to force the protocol to run for $\Omega(\log^2 n/\log\log n)$ expected time. The main technical step is Lemma 2.11 below, which states that, for any fixed edge probability p, there exists an interval of transmission probabilities such that if the protocol's transmission probability is out of this interval then the number of new informed nodes is small. This line of reasoning is similar to that used in Theorem 6.2 in [12] for lower bounding the wake up time in single-hop static radio networks. However, the technical issues to be solved in our framework significantly depart from that static case because of the presence of an unknown dynamic random topology, i.e., dynamic $\mathcal{G}_{n,p}$.

Lemma 2.11. *Let p such that $\frac{\ln^2 n}{n} \leqslant p \leqslant \frac{1}{\ln^5 n}$ and consider any homogeneous broadcast protocol in a dynamic $\mathcal{G}_{n,p}$. Let m be the number of informed nodes at a given time slot t and let m' be the number of informed nodes at time slot $t+1$. If $\frac{n}{4} \leqslant m \leqslant n - \ln^3 n$ and the Protocol's transmission probability $q \notin [\frac{1}{e^4 mp \ln^2 n}, \frac{e^4 \ln^2 n}{mp}]$ at time slot t, then, for sufficiently large n, it holds that*

$$\Pr\left\{m' - m > \frac{1}{\ln^2 n}(n-m)\right\} \leqslant \frac{1}{\ln^2 n}.$$

Proof. Consider any time slot $t \geqslant 1$ of the protocol's execution and let m and m' be the number of informed nodes at time slot t and $t+1$ respectively. For any k, $k \leqslant m$, under the condition that exactly k nodes transmit at time slot t, we define, for each node j, the 0–1 random variable X_j^k that is equal to 1 if node j is not informed at time slot t and it is informed at time slot $t+1$. It is easy to show that

$$\Pr\{X_j^k = 1\} = \begin{cases} kp(1-p)^{k-1} & \text{if } j \text{ is not informed at time } t, \\ 0 & \text{otherwise.} \end{cases}$$

Define

$$I = \left[\frac{1}{e^2 p \ln^2 n}, \frac{e^2 \ln n}{p}\right] \quad \text{and} \quad \tau = \frac{n-m}{\ln^2 n}.$$

Observe that, for any fixed k, random variables $X_1^k, \ldots, X_n^k$ are independent. Let $X^k = \sum_{j=1}^n X_j^k$ and let T be the random variable counting the number of nodes that transmit at time slot t. It thus follows

$$\begin{aligned} \Pr\left\{m' - m > \frac{1}{\ln^2 n}(n-m)\right\} &= \sum_{i=1}^m \Pr\{X^i > \tau\} \Pr\{T = i\} \\ &= \sum_{i \notin I} \Pr\{X^i > \tau\} \Pr\{T = i\} + \sum_{i \in I} \Pr\{X^i > \tau\} \Pr\{T = i\}. \end{aligned} \tag{9}$$

Our goal is to get an upper bound on each of the two sums in the right-hand side of the above equation. We denote $\tilde{p}_i = ip(1-p)^{i-1}$ so that the expected value of X^i is $\mu_i = \mathbf{E}[X^i] = (n-m)\tilde{p}_i$. Note that, for $i < \frac{1}{e^2 p \ln^2 n}$, it holds that

$$\tilde{p}_i = ip(1-p)^{i-1} < \frac{1}{e^2 p \ln^2 n} p = \frac{1}{e^2 \ln^2 n}$$

and for $i > \frac{e^2 \ln n}{p}$, it holds that

$$\tilde{p_i} = ip(1-p)^{i-1} < n\frac{1}{\ln^5 n}e^{-p\frac{e^2\ln n}{p}+p} < \frac{e}{n^{e^2-1}\ln^5 n} < \frac{1}{e^2\ln^2 n}.$$

Hence

$$\mu_i = (n-m)\tilde{p}_i \leqslant \frac{\tau}{e^2}, \quad \text{i.e. } \tau \geqslant e^2\mu_i, \text{ for any } i \notin I. \tag{10}$$

For any $i \notin I$ and $\delta_i = \tau/\mu_i - 1$ (note that (10) implies $\delta_i > 0$), from Chernoff's bound (see (A.1) in Appendix A) we get

$$\Pr\{X^i \geqslant \tau\} \leqslant \frac{e^{\tau-\mu_i}}{(\tau/\mu_i)^\tau} \leqslant e^{-\tau\ln\frac{\tau}{e\mu_i}} < e^{-\tau\ln e} < e^{-\frac{n-m}{\ln^2 n}} < e^{-\frac{\ln^3 n}{\ln^2 n}} < \frac{1}{2\ln^2 n}.$$

We use this upper bound to obtain

$$\sum_{i\notin I}\Pr\{X^i > \tau\}\Pr\{T=i\} < \frac{1}{2\ln^2 n}. \tag{11}$$

Now we get an upper bound on the second sum

$$\sum_{i\in I}\Pr\{X^i > \tau\}\Pr\{T=i\} \leqslant \sum_{i\in I}\Pr\{T=i\} = \sum_{i\in I}\binom{m}{i}q^i(1-q)^{m_t-i}. \tag{12}$$

We consider two cases:

Case ($q < \frac{1}{e^4 mp\ln^2 n}$). Define $\bar{q} = \frac{1}{e^4 mp\ln^2 n}$. Note that $\binom{m}{i}q^i(1-q)^{m-i}$ is decreasing for increasing $i \geqslant m\bar{q}$ whereas $i \geqslant e^2 m\bar{q}$ for each $i \in I$. Thus we have

$$\binom{m}{i}q^i(1-q)^{m-i} \leqslant \binom{m}{\dot{i}}q^{\dot{i}}(1-q)^{m-\dot{i}} \quad \text{for each } i \in I$$

where $\dot{i} = \lceil e^2 m\bar{q}\rceil$. From (12), using $\binom{x}{y} \leqslant (\frac{ex}{y})^y$ and $p \leqslant \frac{1}{\ln^5 n}$ we have

$$\begin{aligned}\sum_{i\in I}\Pr\{X^i > \tau\}\Pr\{T=i\} &\leqslant n\binom{m}{\dot{i}}q^{\dot{i}}(1-q)^{m-\dot{i}} \\ &< n\left(\frac{em}{\dot{i}}q\right)^{\dot{i}} \leqslant n\left(\frac{emq}{e^2 m\bar{q}}\right)^{\frac{1}{e^2 p\ln^2 n}} < \frac{1}{4\ln^2 n}.\end{aligned} \tag{13}$$

Case ($q > \frac{e^4\ln^2 n}{mp}$). Define $\bar{q} = \frac{e^4\ln^2 n}{mp}$. Note that $\binom{m}{i}q^i(1-q)^{m-i}$ is increasing for increasing $i \leqslant m\bar{q}$ whereas $i \leqslant \frac{m\bar{q}}{e^2}$ for each $i \in I$. Thus we have

$$\binom{m}{i}q^i(1-q)^{m-i} \leqslant \binom{m}{\dot{i}}q^{\dot{i}}(1-q)^{m-\dot{i}} \quad \text{for each } i \in I$$

where $\dot{i} = \lfloor\frac{m\bar{q}}{e^2}\rfloor$. From (12), using $\binom{x}{y} \leqslant (\frac{ex}{y})^y$ and $p \leqslant \frac{1}{\ln^5 n}$ we have

$$\begin{aligned}\sum_{i\in I}\Pr\{X^i > \tau\}\Pr\{T=i\} &\leqslant n\binom{m}{\dot{i}}q^{\dot{i}}(1-q)^{m-\dot{i}} \\ &< n\left(\frac{em}{\dot{i}}\right)^{\dot{i}}e^{-q(m-\dot{i})} \leqslant e^{\ln n+\dot{i}\ln(\frac{em}{\dot{i}})-qm+\dot{i}} \\ &\leqslant \left(\frac{1}{n}\right)^{(e^4-2e^2)\ln^6 n-1} < \frac{1}{4\ln^2 n}.\end{aligned} \tag{14}$$

By combining (13) and (14), we get

$$\sum_{i\in I}\Pr\{X^i > \tau\}\Pr\{T=i\} < \frac{1}{2\ln^2 n}. \tag{15}$$

The lemma thus follows from (9), (11) and (15). □

We now show that there exist $\Omega(\frac{\log n}{\log\log n})$ edge probabilities such that their corresponding intervals are pairwise disjoint. A homogeneous broadcast protocol that does not know the probability p of the dynamic $\mathcal{G}_{n,p}$ cannot avoid that at least one of these intervals (and the corresponding edge probability $\tilde{p}$) does exist that contains at most $O(\log n)$ transmission probabilities of the protocol. Hence, for most of the time slots in a dynamic $\mathcal{G}_{n,\tilde{p}}$, the number of new informed nodes will be small.

Theorem 2.12. *Given any homogeneous broadcast protocol $\mathcal{P}$, the adversary can choose a probability p, with $\frac{\ln^2 n}{n} \leqslant p \leqslant \frac{1}{\ln^5 n}$, so that $\mathcal{P}$ has expected completion time $\Omega(\frac{\log^2 n}{\log\log n})$ in a dynamic $\mathcal{G}_{n,p}$.*

Proof. Let q_t be the probability transmission of the homogeneous protocol $\mathcal{P}$ at time slot t, with $1 \leqslant t \leqslant \bar{t}$, where $\bar{t} = \frac{\ln^2 n}{24 \ln\ln n}$. Define the intervals

$$Q_k = \left[\frac{1}{e^4 n p_k \ln^2 n}, \frac{4e^4 \ln^2 n}{n p_k}\right]$$

where $p_k = \frac{1}{\ln^{5k} n}$ and $1 \leqslant k \leqslant \frac{\ln n}{6 \ln\ln n}$ and consider the distribution of the $\bar{t}$ values q_t in the $\frac{\ln n}{6\ln\ln n}$ intervals. It must exist an interval $Q_{\tilde{k}}$ containing at most $\frac{\ln n}{4}$ values.

Consider now an execution of the protocol $\mathcal{P}$ in $\mathcal{G}_{n,p_{\tilde{k}}}$. Let m_t be the number of informed nodes at time step t and consider the events $\mathcal{E}_i$, $1 \leqslant i \leqslant \bar{t}$ where

$$\mathcal{E}_i = \begin{cases} m_{i+1} - m_i \leqslant \frac{1}{\ln^2 n}(n - m_i) & \text{if } m_i \geqslant \frac{n}{4},\ q_i \notin Q_{\tilde{k}}, \\ m_{i+1} - m_i \leqslant (1 - \frac{1}{2e^2})(n - m_i) & \text{otherwise.} \end{cases}$$

Claim 2.13. *If the sequence of events $\mathcal{E}_1, \mathcal{E}_2, \ldots \mathcal{E}_{\bar{t}}$ occurs, then it holds that $m_{\bar{t}} \leqslant n - \frac{n^{1/4}}{24e}$.*

Proof. Let $\hat{t}$ be the first time step such that $m_{\hat{t}} \geqslant \frac{n}{4}$ and assume w.l.o.g. that $\hat{t} < \bar{t}$. We get $m_{\hat{t}-1} < \frac{n}{4}$ and thus, by $\mathcal{E}_{\hat{t}-1}$

$$m_{\hat{t}} \leqslant \left(1 - \frac{1}{2e^2}\right)(n - m_{\hat{t}-1}) + m_{\hat{t}-1} = \left(1 - \frac{1}{2e^2}\right)n + \frac{m_{\hat{t}-1}}{2e^2}$$
$$< \left(1 - \frac{1}{2e^2}\right)n + \frac{n}{8e^2} = n - \frac{3}{8e^2}n < \frac{23}{24}n.$$

Now let r_i be the number of not informed nodes at time step i and define

$$\lambda_i = \begin{cases} \frac{1}{\ln^2 n} & \text{if } q_i \notin Q_{\tilde{k}}, \\ 1 - \frac{1}{2e^2} & \text{otherwise.} \end{cases}$$

We have $r_{\hat{t}} \geqslant \frac{n}{24}$ and $r_{i+1} = n - m_{i+1} \geqslant (1 - \lambda_i)(n - m_i)$ for $\hat{t} \leqslant i \leqslant \bar{t}$. Thus

$$r_{\bar{t}} \geqslant r_{\hat{t}} \prod_{i=\hat{t}}^{\bar{t}} (1 - \lambda_i) \geqslant \frac{n}{24}\left(\frac{1}{2e^2}\right)^s \left(1 - \frac{1}{\ln^2 n}\right)^{\bar{t}-\hat{t}-s} \geqslant \frac{n}{24}\left(\frac{1}{e^3}\right)^s e^{-\frac{\bar{t}}{\ln^2 n - 1}}$$

where s is the number of time slots such that $q_i \in Q_{\tilde{k}}$, $\hat{t} \leqslant i \leqslant \bar{t}$.

Note that $s \leqslant \frac{\ln n}{4}$ and $\bar{t} < \ln^2 n - 1$. Hence $r_{\bar{t}} > \frac{n^{1/4}}{24e}$ and the claim follows since $m_{\bar{t}} = n - r_{\bar{t}}$. □

From the claim, we have that when the sequence of events $\mathcal{E}_1, \mathcal{E}_2, \ldots, \mathcal{E}_{\bar{t}}$ occurs, the broadcast requires more than $\bar{t} = \Omega(\frac{\log^2 n}{\log\log n})$ time slots. To prove the theorem we now show that the sequence $\mathcal{E}_1, \mathcal{E}_2, \ldots, \mathcal{E}_{\bar{t}}$ occurs with constant probability.

Whenever $m_i \geqslant \frac{n}{4}$ and $q_i \notin Q_{\tilde{k}}$ with $i \leqslant \bar{t}$, consider the event $\mathcal{E}'_i$ = "$m_{i+1} - m_i \leqslant \frac{1}{\ln^2 n}(n - m_i)$." From the claim we get $m_i \leqslant n - \ln^3 n$. From $1 \leqslant \tilde{k} \leqslant \frac{\ln n}{6\ln\ln n}$ and the definition of $p_{\tilde{k}}$, we get $\frac{\ln^2 n}{n} \leqslant p_{\tilde{k}} \leqslant \frac{1}{\ln^5 n}$. From $m_i \leqslant n$, we get $\frac{1}{e^4 m_t p_{\tilde{k}} \ln^2 n} < \frac{1}{e^4 n p_{\tilde{k}} \ln^2 n}$ and from $m_i \geqslant \frac{n}{4}$, we get $\frac{e^4 \ln^2 n}{m_i p_{\tilde{k}}} \leqslant \frac{4e^4 \ln^2 n}{n p_{\tilde{k}}}$. Thus, interval

$$Q'_{\tilde{k}} = \left[\frac{1}{e^4 m_i p_{\tilde{k}} \ln^2 n}, \frac{e^4 \ln^2 n}{m_i p_{\tilde{k}}}\right]$$

is a sub-interval of $Q_{\tilde{k}}$. By hypothesis $q_i \notin Q_{\tilde{k}}$, hence we get $q_i \notin Q'_{\tilde{k}}$. Thus applying Lemma 2.6, we have

$$\Pr\{\mathcal{E}'_i\} \geqslant 1 - \frac{1}{\ln^2 n}. \tag{16}$$

Consider event $\mathcal{E}''_i$ = "$m_{i+1} - m_i \leqslant (1 - \frac{1}{2e})(n - m_i)$" where $1 \leqslant i \leqslant \bar{t}$. Note that from the claim, since $p_{\tilde{k}} \leqslant \frac{1}{\ln^5 n}$, we get $m_i \leqslant n - \frac{n^{\frac{1}{4}}}{24e} < n - \frac{18 \ln n}{1 - p_{\tilde{k}}}$. Since $p_{\tilde{k}} \leqslant \frac{1}{\ln^5 n}$, we get $1 - \frac{1 - p_{\tilde{k}}}{2e} < 1 - \frac{1}{2e^2}$. Thus, by applying Lemma 2.11 we have

$$\Pr\left\{m_{t+1} - m_t > \left(1 - \frac{1}{2e^2}\right)(n - m_t)\right\}$$
$$\leqslant \Pr\left\{m_{t+1} - m_t > \left(1 - \frac{1 - p_{\tilde{k}}}{2e}\right)(n - m_t)\right\} \leqslant \frac{1}{n}$$

thus

$$\Pr\{\mathcal{E}_i''\} \geqslant 1 - \frac{1}{n} \geqslant 1 - \frac{1}{\ln^2 n}. \tag{17}$$

From (16) and (17), we get

$$\Pr\left\{\bigcap_{i=1}^{\bar{t}} \mathcal{E}_i\right\} = \prod_{i=1}^{\bar{t}} \Pr\left\{\mathcal{E}_i \,\middle|\, \bigcap_{j=1}^{i-1} \mathcal{E}_j\right\} \geqslant \left(1 - \frac{1}{\ln^2 n}\right)^{\bar{t}} \geqslant e^{-\frac{\bar{t}}{\ln^2 n - 1}} \geqslant e^{-1}$$

where the last step follows since $\bar{t} < \ln^2 n$.

We can thus claim that probability that the broadcast on $\mathcal{G}_{n,p_{\tilde{k}}}$ is not completed within the first $\Omega(\frac{\log^2 n}{\log\log n})$ time slots is a positive constant. □

3. Deterministic adversary

In this section we consider broadcasting against the worst-case adversary. At each time slot t, the adaptive adversary chooses the set E_t of edges, thus yielding an infinite sequence of graphs $G_1, G_2, \ldots, G_t, \ldots$. As stated in the introduction, we consider only meaningful adversaries.

It is interesting to observe that the BGI's procedure fails to complete broadcasting against the adaptive worst-case adversary. However, we now show that a very simple oblivious protocol works efficiently.

Theorem 3.1. *There exists a homogeneous randomized protocol that, for any adaptive worst-case adversary, completes broadcasting within $O(\frac{n^2}{\log n})$ time slots, w.h.p.*

Proof. Let us consider the following homogeneous protocol:

At every time slot all the informed nodes transmit with probability $q = \frac{\ln n}{n}$.

Consider a non-informed node u that has $k \geqslant 1$ informed neighbors in a given time slot. Then the probability that u gets the message in this time slot is $kq(1-q)^{k-1}$. Consider the function

$$f(x) = xq(1-q)^{x-1} \quad \text{with } x \in [1, n].$$

If $q \leqslant 1 - (1/n)^{\frac{1}{n-1}} \approx \frac{\ln n}{n}$ then the minimum of f lies in $x = 1$. If we choose $q = \ln n/n$, we have that $f(x) \geqslant \ln n/n$ for each $x \in [1, n]$. Hence, at each time slot, there exists a non-informed node that has probability at least $\ln n/n$ to get informed. The expected time to get a new informed node is thus at most $n/\ln n$ and, so, the expected completion time of the broadcasting is $O(\frac{n^2}{\log n})$.

In order to show that this upper bound holds with high probability we need a more careful argument. Let us fix an adversary strategy $\mathcal{A}$. Note that, being adaptive, this strategy considers all possible protocol's actions and all possible network's configurations at run time. From the previous discussion on the expected completion time, we set $q = \frac{\ln n}{n}$. Then, it is easy to verify that for each $k \geqslant 1$

$$kq(1-q)^{k-1} \geqslant \bar{q} \quad \text{where } \bar{q} = \frac{q}{2}.$$

Consider the probability $p_{t,k}$ that, at time slot t, there are at least k informed nodes. It holds that

$$p_{t,k} \geqslant (1 - \bar{q})p_{t-1,k} + \bar{q}p_{t-1,k-1}.$$

Indeed, the inequality is obtained by summing up the probabilities of two disjoint events: either there are at least $k-1$ informed nodes at time slot $t-1$ and a new node gets informed, or there are at least k informed nodes at time slot $t-1$ and no new node gets informed.

By solving the above inequality with respect to the first term of the right side, we obtain

$$p_{t,k} \geqslant (1-\bar{q})^{t-k}\bar{q}^{k-1} + \bar{q}\sum_{s=k}^{t-1}(1-\bar{q})^{t-1-s}p_{s,k-1}.$$

By induction on k it is possible to verify that

$$p_{t,k} \geqslant 1 - (1-\bar{q})^{t-k} \sum_{s=0}^{k-1} \frac{(t\bar{q})^s}{s!}.$$

By evaluating the value of $p_{t,k}$ when $t = T = 10\frac{n^2}{\ln n}$ and $k = n$, we get $p_{T,n} \geqslant 1 - e^{-n}$. □

Theorem 3.2. *Given any randomized broadcast protocol, there is an adaptive worst-case adversary that forces the protocol to have $\Omega(\frac{n^2}{\log n})$ expected completion time.*

Proof. Consider homogeneous protocols first, i.e. protocols in which at every time slot, every informed node transmits with the same probability. Let m be the number of informed nodes at time slot t and let q be the transmission probability. The adversary adopts the following strategy: If $q \leqslant \ln m/m$ then the adversary connects only one informed node with a non-informed one and all remaining nodes are kept isolated; otherwise, it connects *all* the m informed nodes to a non-informed one. In both cases, when there are $m > 3$ informed nodes, the probability that a new node gets informed is less than $2\ln m/m$. This is trivially true if $q \leqslant \ln m/m$, otherwise observe that function $f(q) = mq(1-q)^{m-1}$ is decreasing when $q > 1/m$. So if $q > \ln m/m$ the probability that a new node will be informed is

$$f(q) \leqslant f(\ln m/m) = \ln m\left(1 - \frac{\ln m}{m}\right)^{m-1} \leqslant 2\frac{\ln m}{m}.$$

Now let X_m be the random variable counting the time slots needed to inform a new node when there are m informed nodes. Then $\mathbf{E}[X_m] \geqslant \frac{m}{2\ln m}$ and so the expected time to complete broadcasting is

$$\mathbf{E}\left[\sum_{m=1}^{n-1} X_m\right] = \sum_{m=1}^{n-1} \mathbf{E}[X_m] \geqslant \frac{1}{2}\sum_{m=2}^{n-1} \frac{m}{\ln m} \in \Theta\left(\frac{n^2}{\log n}\right).$$

The extension of the lower bound to general protocols is quite straightforward. Let m be the number of informed nodes at time slot t, and let q_i, $i = 1, \dots, m$ be the transmission probability of informed node i. The adversary adopts the following strategy:

(1) If a node exists such that its transmission probability at time slot t is less than $\ln m/m$, then the adversary connects this node with a non-informed node and all remaining nodes are kept isolated;
(2) Otherwise, it connects *all* the m informed nodes to a non-informed one.

When there are $m > 3$ informed nodes the probability that a new node gets informed is less than $2\ln m/m$. Indeed, this is trivial if we are in Case 1. As for Case 2, first observe that, when $q_i \geqslant 1/m$ for any $i = 1, \dots, m$, function

$$g(q_1, \dots, q_m) = \sum_{i=1}^{m} q_i \prod_{j \neq i} (1 - q_j)$$

satisfies the following monotone property. For any $q'_i \geqslant q_i$, $i = 1, \dots, m$, it holds that $g(q'_1, \dots, q'_m) \leqslant g(q_1, \dots, q_m)$. So, the probability that a new node gets informed is

$$\sum_{i=1}^{m} q_i \prod_{j \neq i} (1 - q_j) \leqslant m\frac{\ln m}{m}\left(1 - \frac{\ln m}{m}\right)^{m-1} \leqslant 2\frac{\ln m}{m}. \quad \square$$

4. Conclusions

Concerning the weak random adversary, an interesting open question is whether the lower bound can be extended to oblivious protocols when p is unknown.

As for the worst-case adversary, note that the adversary in the proof of Theorem 3.2 is adaptive since it needs to know the informed nodes at any time slot. Finding a good lower bound for oblivious adversaries is an open question.

We studied two extremal adversaries aiming to establish the broadcast complexity against the somewhat most favorable, natural dynamic scenario and against the worst-case one, respectively. Our tight results on these two adversaries set up a framework that aims to stimulate future studies on more realistic adversaries “lying” between the two above. An interesting approach would be that of introducing *time dependencies* in our random adversary: the random topology at a given time slot is somewhat related to the topology at the previous time slot. For instance, the case where only a fixed fraction of (unknown) edges are subject to random changes. Another case is where any pair of nodes has a fixed probability of keeping the previous state: connected or not.

The challenging ultimate goal of this line of research is to provide analytical results about *geometric* dynamical models [21].

Acknowledgments

We wish to thank Paolo Penna and Luca Trevisan for comments and very helpful suggestions.

Appendix A. Useful inequalities

Lemma A.1 *(Chernoff's bounds). Let $X = \sum_{i=1}^{n} X_i$ where $X_1, \ldots, X_n$ are independent Bernoulli random variables.*

- *If $\mathbf{E}[X] \leqslant \mu$, than for each $\delta > 0$ it holds that*

$$\mathbf{P}\{X \geqslant (1+\delta)\mu\} \leqslant \left[\frac{e^{\delta}}{(1+\delta)^{(1+\delta)}}\right]^{\mu}. \tag{A.1}$$

- *If $\mathbf{E}[X] \geqslant \mu$, than for each $0 < \delta < 1$ it holds that*

$$\mathbf{P}\{X \leqslant (1-\delta)\mu\} \leqslant e^{-\frac{\delta^2}{2}\mu}. \tag{A.2}$$

- *If $\mathbf{E}[X] = \mu$, than for each $0 < \delta < 1$ it holds that*

$$\Pr\{X \notin [(1-\delta)\mu, (1+\delta)\mu]\} \leqslant 2e^{-\frac{\delta^2}{3}\mu}. \tag{A.3}$$

Lemma A.2. *(See Exercise 4.13 in [14].) Let $0 < p < 1$ and let $X_1, \ldots, X_n$ be independent random variables such that for each $i = 1, \ldots, n$*

$$\Pr\{X_i = 1\} = p, \qquad \Pr\{X_i = 0\} = 1 - p.$$

Let $X = \sum_{i=1}^{n} X_i$ so that $\mathbf{E}[X] = np$. Than for each $p < x \leqslant 1$ it holds

$$\Pr\{X \geqslant xn\} \leqslant e^{-nF(x,p)}$$

where

$$F(x, p) = x \ln \frac{x}{p} + (1-x) \ln \frac{1-x}{1-p}. \tag{A.4}$$

Corollary 3. *Let $X_1, \ldots, X_n$ and X be as in the previous lemma. Then*

$$\Pr\left\{X \geqslant \frac{1+p}{2} n\right\} \leqslant e^{-\lambda(1-p)n}$$

where $\lambda = (1 - \ln 2)/2$.

References

[1] M. Adler, C. Scheideler, Efficient communication strategies for ad hoc wireless networks, in: Proc. of 10th Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA), ACM Press, 1998, pp. 259–268.
[2] N. Alon, J. Spencer, The Probabilistic Method, Wiley and Sons, 1992.
[3] B. Awerbuch, P. Berenbrink, A. Brinkmann, C. Scheideler, Simple routing strategies for adversarial systems, in: Proc. of 42th IEEE Symposium on Foundations of Computer Science (FOCS), IEEE Computer Society, 2001, pp. 158–167.
[4] R. Bar-Yehuda, O. Goldreich, A. Itai, On the time-complexity of broadcast in multi-hop radio networks: An exponential gap between determinism and randomization, J. Comput. System Sci. 45 (1) (1992) 104–126.
[5] B. Bollobàs, Random Graphs, Cambridge University Press, 2001.
[6] B. Chlebus, L. Gasieniec, A. Gibbons, A. Pelc, W. Rytter, Deterministic broadcasting in unknown radio networks, in: Proc. of the 11th Annual ACM–SIAM Symposium on Discrete Algorithms (SODA), ACM Press, 2000, pp. 861–870.
[7] M. Chrobak, L. Gasieniec, W. Rytter, Fast broadcasting and gossiping in radio networks, J. Algorithms 43 (2) (2002) 177–189.
[8] A. Clementi, A. Monti, R. Silvestri, Distributed broadcast in radio networks of unknown topology, Theoret. Comput. Sci. 302 (1–3) (2003) 337–364.
[9] A. Clementi, A. Monti, R. Silvestri, Round robin is optimal for faulttolerant broadcasting on wireless networks, J. Parallel Distrib. Comput. 64 (1) (2004) 89–96.
[10] A. Czumaj, W. Rytter, Broadcasting algorithms in radio networks with unknown topology, J. Algorithms 60 (2) (2006) 115–143.
[11] R. Elsässer, L. Gasieniec, Radio communication in random graphs, J. Comput. System Sci. 72 (3) (2006) 490–506.
[12] T. Jurdzinski, G. Stachowiak, Probabilistic algorithms for the wakeup problem in single-hop radio networks, in: Proc. of 13th Annual International Symposium on Algorithms and Computation (ISAAC02), in: Lecture Notes in Comput. Sci., vol. 2518, Springer, 2002, pp. 535–549.
[13] E. Kushilevitz, Y. Mansour, An $\Omega(d \log(n/d))$ lower bound for broadcast in radio networks, SIAM J. Comput. 27 (3) (1998) 702–712.
[14] M. Mitzenmacher, E. Upfal, Probability and Computing, Cambridge University Press, 2005.
[15] M. Moshin, D. Cavin, Y. Sasson, R. Prakash, A. Schiper, A lower bound for broadcasting in mobile ad hoc networks, Ecole Polytechnique Federale de Lausanne, Switzerland, Tech. Rep. IC/2004, 2004.
[16] M. Moshin, D. Cavin, Y. Sasson, R. Prakash, A. Schiper, Reliable broadcast in wireless mobile ad hoc networks, in: Proc. of 39th IEEE Hawaii International Conference on System Science, vol. 09, IEEE Computer Society, 2006, p. 233.1.

[17] R. O'Dell, R. Wattenhofer, Information dissemination in highly dynamic graphs, in: Proc. of ACM DIALM-POMC, ACM Press, 2005, pp. 104–110.
[18] A. Pelc, D. Peleg, Feasibility and complexity of broadcasting with random transmission failures, in: Proc. of 24th Annual ACM Symposium on Principles of Distributed Computing (PODC), ACM Press, 2005, pp. 334–341.
[19] S. Ramanathan, E. Lloyd, Scheduling algorithms for multihop radio networks, IEEE/ACM Trans. Netw. 1 (2) (1993) 166–177.
[20] R. Ramaswami, K. Parhi, Distributed scheduling of broadcasts in radio networks, in: Proc. of IEEE INFOCOM, IEEE Computer Society, 1989, pp. 497–504.
[21] C. Scheideler, Models and techniques for communication in dynamic networks, in: Proc. of 19th Symposium on Theoretical Aspects of Computer Science (STACS), in: Lecture Notes in Comput. Sci., vol. 2285, Springer, 2002, pp. 27–49.