Communication in Dynamic Radio Networks *

Andrea E. F. Clementi † Angelo Monti ‡ Francesco Pasquale†*
Riccardo Silvestri‡

Abstract

We study the completion time of distributed broadcast protocols in *dynamic* radio networks. The dynamic network is modelled by means of adversaries: we consider two of them that somewhat are the extremal cases.

We first analyze the weakest one, i.e., an oblivious, memoryless random adversary. At each time slot t, a graph G_t is selected according to the well-known random graph model $G_{n,p}$. We derive a randomized protocol that has $O(\log n)$ completion time. Then, we prove that any randomized protocol has $\Omega(\log n)$ completion time. This tight bound holds when the protocol knows p. When p is unknown, we present an oblivious homogeneous version of the Bar Yehuda-Goldreich-Itai's randomized protocol having $O(\log^2 n)$ completion time and we prove a lower bound $\Omega(\log^2 n/\log\log n)$ that holds for any randomized oblivious homogeneous protocol. We emphasize that the above (poly-)logarithmic upper bounds also hold when random graphs are sparse and disconnected, i.e., for $p = o(\ln n/n)$.

We then consider the deterministic worst-case adversary that, at each time slot, can make any network change (thus the strongest adversary). Up to now, it is not even known whether *finite* expected completion time is achievable against this adversary. We present a simple randomized protocol that works in $O(n^2/\log n)$ completion time. This bound is then shown to be optimal.

*Partially supported by the European Union under the Project AEOLUS.
†Dipartimento di Matematica, Università di Roma "Tor Vergata". {clementi,pasquale}@mat.uniroma2.it.
‡Dipartimento di Informatica, Università di Roma "La Sapienza". {monti,silvestri}@di.uniroma1.it.

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.

In the radio network model adopted in almost all previous theoretical works [5, 8, 7, 10, 20, 19], the communication is assumed to be synchronous. A possible way to get global timing information is to equip each node with a GPS-like receiver [16, 4]. From a theoretical point of view, assuming the synchronous mode of communication allows 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 presence of collisions make protocols and their analysis for radio networks significantly different from those working in wired network models.

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 algorithmic and networking areas [5, 8, 7, 20, 19]. The *completion time* of a broadcast protocol in a synchronous network is the number of time slots required by the protocol to inform all (reachable) 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 Subsection 1.2).

On the other hand, theoretical results on communication protocols in any scenario in which the network topology may change during the protocol's execution (i.e. a *dynamic* radio network) are very few [21]. Some dynamic models have been studied concerning the broadcast operation. However, for almost all of them, there are only experimental results. Analytical results are only known for dynamic networks with very restrictive characteristics (see Subsection 1.2). It is not even known whether a randomized broadcast protocol exists that has *finite* expected completion time against 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, 2, 21]. We study scenarios in which the edges of the network 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 the completion time of broadcasting against any dynamic adversary strategy. More importantly, such extremal choices about the adversary aim to answer the two following fundamental questions: *1.* Do (meaningful) dynamic scenarios *always* constitute a *hurdle* for radio communication? *2.* How much hard can radio communication be against (meaningful) *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. In oblivious protocols, the actions of an informed node do not depend on any information (but the source message) received during the execution of the protocol. An even more restricted kind 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 modelled 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 $G_{n,p}$ where n is the number of nodes and p is the edge probability [3, 6]. This adversarial strategy will be simply denoted as *dynamic* $\mathcal{G}(n,p)$.
Dynamic $\mathcal{G}(n,p)$ can be considered as the "random" version of *highly-dynamic* radio networks studied by O'Dell and Wattenhofer [17, 21] (see also Subsection 1.2). In the latter model, nodes (agents) and links (each link is here intended as a temporary connection between two agents) are subject to arbitrary changes (and/or faults). Such changes are so fast with respect to the protocol speed that the links of two consecutive time slots are completely unrelated.
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 dynamic $\mathcal{G}(n,p)$ in $O(\log n)$ time slots. 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 dynamic $\mathcal{G}(n,p)$ in $\Omega(\log n)$ expected time. We emphasize that the lower bound holds for any kind of randomized protocol.
The above upper bound assumes that the protocol knows p. We then consider the case in which the protocol does not know p: the adversary, based on the protocol strategy, can choose p in order to minimize the probability of successful communications. In this case, we first show that a simple, homogeneous version of the Bar Yehuda-Goldreich-Itai's (BGI's) protocol [5] 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 is p, with $\frac{\ln^2 n}{n} \leqslant p \leqslant \frac{1}{\ln^5 n}$, so that the protocol completes broadcasting in dynamic $\mathcal{G}(n,p)$ in $\Omega(\log^2 n/\log\log n)$ expected time.
Let us observe that the above protocols work in (poly-)logarithmic time even when, at every time slot, the expected maximal node degree is constant and the radio network is with high probability disconnected (the latter happens whenever $p = o(\log n/n)$ [6]). 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 Subsection 1.2).
We thus answer to Question (*1*) above by providing the first tight mathematical form of the fact that oblivious fully-random network changes, instead of working as an hurdle, *help* information propagation. This is far to be trivial due to the unpredictable collisions yielded by dynamic radio networks.

The strong worst-case adversary. We consider 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.
It is important to observe that, for any deterministic protocol, there is a worst-case adversary such that, at each time slot, the graph is connected (so the adversary is meaningful) and the protocol

never completes broadcasting.
Rather interestingly, 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 (oblivious or not) randomized protocol.
Such results thus provide the first tight mathematical answer to our second fundamental question. In particular, our quadratic upper bound implies that no meaningful adversary exists that yields *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).

Table 1

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

Another kind of deterministic adversary is the *oblivious worst-case* one. This adversary knows the protocol but it must decide all network changes a priori before the protocol's execution. Note that there is no difference between adaptive and oblivious adversaries when the protocol is deterministic while, against randomized protocols, adaptive adversaries may be much more powerful. As for the oblivious worst-case adversary, we will observe that the proof of the lower bound for the adaptive adversary can be modified in order to get the same lower bound when the randomized protocol is oblivious. On the other hand, it is an open question whether, against the oblivious worst-case adversary, adaptive randomized protocols can achieve better completion time.
Finally, we emphasize that our work significantly departs from all previous theoretical works on this topic in two important issues:

In some theoretical studies [17, 16, 15], 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 that the networks are connected during *all* time slots of the protocol. Our results show that 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 and it is proved in [8]. Then, [9] proved an $\Omega(n \log D)$ lower bound on the completion time of any deterministic protocol, where D is the source eccentricity.
As for randomized protocols, Bar-Yehuda, Goldreich, and Itai [5] proposed a protocol, denoted here as BGI's protocol, that completes broadcasting in $O\left(D \log n + \log^2 n\right)$, w.h.p. Then, [10] improved the BGI's protocol obtaining completion time $O(D \log(n/D) + \log^2 n)$, w.h.p. On the other hand, [13] proved a lower bound $\Omega(D \log(n/D))$.
Finally, broadcasting in *static random* graphs $G_{n,p}$ 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 $p \geqslant \log^{1+\epsilon} n/n$. Note that, in this case, the graph is w.h.p. connected. We emphasize that techniques and results (both for lower and upper bounds) in [11] do not work in our dynamic random graph

model. Roughly speaking, one of the crucial difference is that their techniques strongly rely on the static *layered* broadcast process yielded by the static graph: During the entire protocol execution, every node keeps the same distance from the source and it is possible to define layer L_i as the set of nodes at distance i from the source. The broadcast process is then analyzed by considering message propagation from one layer to the next one. This is clearly impossible in our dynamic model since the distance between two nodes may change arbitrarily from one time slot to the next one.

Dynamic networks. A theoretical study of broadcasting in general 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 which has been *always* fault-free. This assumption is necessary in the worst-case setting: any deterministic protocol fails to inform nodes on dynamical paths managed by an adversary that knows the protocol. In [9], it is proved that the *round robin* strategy is asymptotically optimal. Then for graphs of maximal in-degree Δ, a deterministic protocol is presented having completion time $O(D\Delta \log^2 n)$.
Deterministic broadcasting in faulty radio networks of *known* topology is studied in [18]. It their model, 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 $0 < p < 1$. [18] proves an $O(opt \log n)$ for the broadcasting completion time where *opt* is the optimal completion time in the fault-free case.
More recently, O'Dell and Wattenhofer [17] studies broadcasting on *highly-dynamic graphs*. Here, the adversary can arbitrarily change the edges of the graph at each time slot. The constraint is that 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 in [17] 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]. In their model, at each time slot, a node can move from one grid point to an arbitrary adjacent one. [15] proves a lower bound $\Omega(D \log n)$ for one-dimensional grids and a $\Omega(n \log n)$ lower bound for two-dimensional grids. Then, [16] provides a protocol which completes broadcasting on any one-dimensional grid working in $O(D \log n)$ time slots. We emphasize that local node mobility is somewhat previously "known" by every node in this model too. Time slots are not homogeneous: during a *control slot*, nodes declare their next moves. This clearly strongly helps the scheduling of transmissions during *application slots*.

2 The Random Adversary

In this section, we will consider the broadcast operation against *dynamic random graphs*. 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* $G_{n,p}$ [6]. A *random graph* $G_{n,p}$ 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 always denote the edge probability of random graphs.
A broadcast protocol in dynamic $\mathcal{G}(n,p)$, at any time slot t, acts 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 [5].

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

The protocol clearly terminates within $O(\log n)$ time slots.

Theorem 2.1 *Let $p \geqslant 1/n$. The protocol* `DynamicBroadcast(n,p)` *completes broadcasting in dynamic $\mathcal{G}(n,p)$, w.h.p.*

The full proof is given in Appendix B.1. The proof evaluates the number of informed nodes after every loop of the protocol. In particular, after the second loop, we prove that, for any $p \geqslant 1/n$, the number m of informed nodes is w.h.p. at least $\frac{1}{2p}$. Then, the main technical contribution here is the following lemma that evaluates the number of informed nodes after the third loop. Note that the analysis significantly departs from those in [5] and [11] for static unknown graphs.

Lemma 2.2 *Let p 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} (the third Protocol's loop), then at least γn nodes are informed w.h.p., where $\gamma = \frac{1}{8e^{7/4}}$.*

Proof. Let m_k be the number of informed nodes at time slot k of this phase. From Lemma's hypothesis it holds that $m_0 \geqslant 1/(2p)$. Consider the first $\lceil \ln(np) \rceil$ time slots. If for each of them it holds $k \leqslant \ln(2m_k p)$, then, in the last one of them, it holds $\lceil \ln(np) \rceil \leqslant \ln(2m_k p)$. So $m_k \geqslant n/2$. Otherwise a time slot k must exist for which $k = \lceil \ln(2m_k p) \rceil$. 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

$$\frac{1}{2emp} \leqslant q \leqslant \frac{1}{2mp}$$

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. 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 the random variable $X = \sum_{i=1}^{n-m} X_i$ counting the number of new informed nodes. 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 node is a constant fraction of n (w.h.p.), we cannot apply the Chernoff 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 independently sending the message with probability q, it holds that

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

From Chernoff's bound (Eq. (2) in the Appendix) on T with[1] $\mu = mq$ 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

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

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

$$\Pr\left\{X_i^j = 1\right\} = jp(1-p)^{j-1} \geqslant \frac{1}{2}mqp(1-p)^{\frac{3}{2}mq-1} \geqslant \frac{1}{4e}e^{-\frac{p}{1-p}\left(\frac{3}{4p}-1\right)} \geqslant \frac{1}{4e}e^{-3/4} = \frac{1}{4e^{7/4}}$$

Hence for each $j \in \left[\frac{1}{2}mq, \frac{3}{2}mq\right]$ the random variable $X^j = \sum_{i=1}^{n-m} X_i^j$ has expectation

$$\mathbf{E}\left[X^j\right] \geqslant \frac{n-m}{4e^{7/4}} \geqslant \frac{n}{8e^{7/4}}$$

From Chernoff's bound (Eq. (1) in the Appendix) with $\mu = \frac{n}{8e^{7/4}}$ and $\delta = 1/2$, it follows that

$$\Pr\left\{X^j \leqslant 4an\right\} \leqslant e^{-an} \quad \text{where } a = \frac{1}{64e^{7/4}}$$

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

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

for a suitable positive constant ε. □

We observe that if p is $1 - o(1)$ the broadcast task can be completed in $o(\log n)$ by considering the simple protocol in which only the source transmits the message with probability 1 (e.g. if $p = 1 - 1/n^2$ broadcasting is completed in one time slot). The proof of next theorem is given in Appendix B.2.

Theorem 2.3 *Let ε be any positive constant and let $p \leqslant 1-\varepsilon$. Any broadcast protocol in dynamic $\mathcal{G}(n,p)$ has expected completion time $\Omega(\log n)$.*

[1] Note that T can be viewed as $T = \sum_{l=1}^{m} T_l$ where the random variable T_l has value 1 if node l sends the message and 0 otherwise. Note that $T_1, \ldots, T_m$ are independent.

2.2 Case p unknown

Let us consider the following homogeneous variant of the BGI's Decay procedure [5], 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. A sketch of the proof of the following theorem is given in Appendix B.1.

Theorem 2.4 *Protocol* `BGI(n)` *completes broadcasting in dynamic* $\mathcal{G}(n,p)$ *w.h.p. for any* $p \geqslant 1/n$.

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.

Theorem 2.5 *For any homogeneous broadcast protocol, the adversary can choose a probability* p, *with* $\frac{\ln^2 n}{n} \leqslant p \leqslant \frac{1}{\ln^5 n}$, *so that the protocol in dynamic* $\mathcal{G}(n,p)$ *has expected completion time* $\Omega\left(\frac{\log^2 n}{\log\log n}\right)$.

The proof is rather complex and it is given in Appendix B.2 we here state the main technical lemma and provide a short outline of the overall proof.

Lemma 2.6 *Let* p *such that* $\frac{\ln^2 n}{n} \leqslant p \leqslant \frac{1}{\ln^5 n}$ *and consider any homogeneous broadcast protocol in 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 successive time slot. If* $\frac{n}{4} \leqslant m \leqslant n - \ln^3 n$ *and the Protocol's transmission probability* q *at time slot* t *does not belong to the interval* $\left[\frac{1}{e^4 mp \ln^2 n}, \frac{e^4 \ln^2 n}{mp}\right]$, *then, for* n *sufficiently large, it holds that*

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

The lemma 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. We then show that there exist $\Omega\left(\frac{\log n}{\log\log n}\right)$ 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 dynamic $\mathcal{G}(n,\tilde{p})$, the number of new informed nodes will be small.

3 Deterministic Worst-Case 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 an homogeneous randomized protocol that, for any adaptive worst-case adversary, completes broadcasting within* $O\left(\frac{n^2}{\log n}\right)$ *time slots, w.h.p.*

Sketch of the proof. Let us consider the following simple homogeneous protocol:

> *At every time slot all the informed nodes transmit with probability* q
> (the choice of q will be given later)

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

In order to show that this upper bound holds with high probability we need a more careful argument that is sketched below.

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's 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 is possible to prove that

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

Intuitively speaking this 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* $\Omega\left(\frac{n^2}{\log n}\right)$ *expected completion time.*

Sketch of the proof. Consider any protocol. Let m be number of informed nodes at time slot t. The adversary adopts the following strategy. 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. Otherwise, it connects *all* the m informed nodes to a non informed one. In both cases, it is possible to prove that, when there are $m \geqslant 2$ informed nodes, the probability that a new node gets informed is less than $2\ln m/m$. 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

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

□

Oblivious worst-case adversary. Note that the adversary in the above proof is adaptive since it needs to know the informed nodes at any time slot. However when the protocol is oblivious the adversary a priori knows the transmission probabilities $q(i,t)$ for any node i and time slot t. This allows us to transform the adversary's strategy of the above proof into an oblivious one. We can thus obtain the $\Omega\left(\frac{n^2}{\log n}\right)$ lower bound for oblivious randomized protocols against the oblivious worst-case adversary.

4 Open Questions

As for the weak random adversary, when p is unknown, it remains open the question whether the lower bound can be extended to oblivious protocols and whether it can be made tight.
As for the worst-case adversary, we don't know whether an adaptive randomized protocol against the oblivious adversary can beat our lower bound for oblivious protocols. Moreover, it is possible to modify the adaptive adversary yielding the $\Omega(n^2/\log n)$ lower bound in Theorem 3.2 in order to keep the network always connected. It would be interesting to know whether the same property can be guaranteed by the oblivious adversary.
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.

References

[1] M. Adler and C. Scheideler. Efficient communication strategies for ad hoc wireless networks. *Proc. of 10th Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA)*, pp. 259–268, 1998.

[2] B. Awerbuch, P. Berenbrink, A. Brinkmann, and C. Scheideler. Simple routing strategies for adversarial systems. *Proc. of 42th IEEE Symposium on Foundations of Computer Science (FOCS)*, pp. 158–167, 2001.

[3] N. Alon and J. Spencer. *The Probabilistic method*. Wiley and Sons, 1992.

[4] S.V. Bana and P. Varaiya. Real-time vehicle location with desired accuracy. *Proc. of IEEE Intelligent Transportation Systems Conference*, pp. 160–165, 2001.

[5] R. Bar-Yehuda, O. Goldreich, and A. Itai. On the time-complexity of broadcast in mult-hop radio networks: an exponential gap between determinism and randomization. *Journal of Computer and System Sciences*, 45(1), pp. 104–126, 1992.

[6] B. Bollobas. *Random Graphs*. Cambridge University Press, 2001.

[7] B. Chlebus, L. Gasieniec, A. Gibbons, A. Pelc, and W. Rytter. Deterministic broadcasting in unknown radio networks. *Proc. of the 11th Annual ACM-SIAM Symposium on Discrete algorithms (SODA)*, pp. 861–870, 2000.

[8] M. Chrobak, L. Gasieniec, and W. Rytter. Fast Broadcasting and Gossiping in Radio Networks. *Journal of Algorithms*, 43(2), pp. 177–189, 2002.

[9] A. Clementi, A. Monti, and R. Silvestri. Round Robin is optimal for fault-tolerant broadcasting on wireless networks. *Journal of Parallel and Distributed Computing*, 64(1), pp. 89–96, 2004.

[10] A. Czumaj and W. Rytter. Broadcasting algorithms in radio networks with unknown topology. *Journal of Algorithms*, 60(2), pp. 115–143, 2006.

[11] R. Elsasser and L. Gasieniec. Radio communication in random graphs. *Journal of Computer and System Sciences*, 72(3), pp. 490–506, 2006.

[12] E. Kranakis, D. Krizanc, and A. Pelc. Fault-tolerant broadcasting in radio networks. *Journal of Algorithms*, 39(1), pp. 47–67, 2001.

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

[14] M. Mitzenmacher and E. Upfal. *Probability and Computing*, Cambridge University Press, 2005.

[15] M. Moshin, D. Cavin, Y. Sasson, R. Prakash, and A. Schiper. A lower bound for broadcasting in mobile ad hoc networks. Ecole Polytechnique Federale de Lausanne, Tech. Rep. IC/2004/37, 2004.

[16] M. Moshin, D. Cavin, Y. Sasson, R. Prakash, and A. Schiper. Reliable Broadcast in Wireless Mobile Ad Hoc Networks. *Proc. of 39th IEEE Annual Hawaii International Conference on System Science(HICSS)*, p. 233.1, 2006.

[17] R. O'Dell and R. Wattenhofer. Information Dissemination in Highly Dynamic Graphs. *Proc. of ACM DIALM-POMC*, pp. 104–110, 2005.

[18] A. Pelc and D. Peleg. Feasibility and complexity of broadcasting with random transmission failures. *Proc. of 24th Annual ACM Symposium on Principles of Distributed Computing (PODC)*, pp. 334–341, 2005.

[19] S. Ramanathan and E. Lloyd. Scheduling algorithms for multihop radio networks. *IEEE/ACM Transactions on Networking*, 1(2) pp. 166–177, 1993.

[20] R. Ramaswami and K. Parhi. Distributed scheduling of broadcasts in radio networks. *Proc. of IEEE INFOCOM*, 1989.

[21] C. Scheideler. Models and Techniques for Communication in Dynamic Networks. *Proc. of 19th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, pp. 27–49, 2002.

Appendix

A Preliminaries

In what follows we remind some basic results that will be used in our proofs.

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

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

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

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

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

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*

$$\Pr\{X_i = 1\} = p \qquad \Pr\{X_i = 0\} = 1 - p \qquad \text{for each } i = 1, \ldots, n$$

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

$$F(x,p) = x \ln \frac{x}{p} + (1-x) \ln \frac{1-x}{1-p} \tag{3}$$

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

Proof. Put $x = (1+p)/2$, then $1 - x = (1-p)/2$. From Eq. (3) it holds

$$\begin{aligned}
F\left(\frac{1+p}{2}, p\right) &= \frac{1+p}{2} \ln \frac{1+p}{2p} + \frac{1-p}{2} \ln 1/2 \\
&= \frac{1+p}{2} \ln \left(1 + \frac{1-p}{2p}\right) - \frac{1-p}{2} \ln 2 \\
&\geqslant \frac{1+p}{2} \frac{\frac{1-p}{2p}}{1 + \frac{1-p}{2p}} - \frac{1-p}{2} \ln 2 \\
&= \frac{1-p}{2}(1 - \ln 2)
\end{aligned}$$

and the thesis follows by Lemma A.2. □

B The Random Adversary

B.1 Upper Bounds

Case p known. We first prove that protocol `DynamicBroadcast(n,p)` completes the broadcast w.h.p. The proof need the following lemmas.

Lemma B.1 *Assume that the source node 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

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

- 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 $i = 1, 2, \ldots, n$. Consider the random variable $X = \sum_{i=1}^{n} X_i$ counting the number of informed nodes after $c \ln n$ time slots. The expected value of X is

$$\mathbf{E}[X] = \sum_{i=1}^{n} \mathbf{E}[X_i] \geqslant \frac{c}{2} \ln n$$

  Since X_i's are independent, by using Eq. (1) 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, for $c > 2$, $\alpha = \frac{(c-2)^2}{4c}$.

- 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 $i = 1, 2, \ldots, n$. Let us consider the random variable $X = \sum_{i=1}^{n} X_i$ counting the number of informed nodes after the $c \ln n$ time slots. The expected value of X is

$$\mathbf{E}[X] = \sum_{i=1}^{n} \mathbf{E}[X_i] \geqslant (1 - e^{-c})n$$

  Since X_i's are independent, by using Eq. (1) with $\mu = (1 - e^{-c})n$ and $\delta = \frac{1}{2}\left(1 - \frac{1}{e^c - 1}\right)$, it holds that

$$\Pr\{X \leqslant n/2\} \leqslant e^{-\alpha n}$$

  where α is a positive constant for $c > 0$, $\alpha = \frac{1}{8}\left(1 - \frac{1}{e^c - 1}\right)^2 (1 - e^{-c})$.

□

Lemma B.2 *Let p such that $1/n \leqslant p \leqslant 1/\ln n$. Assume that we start with at least $\ln n$ informed nodes and 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, $k = 0, 1, \dots, \lceil c \ln n \rceil$. From the Lemma's hypothesis it holds that $m_0 \geqslant \ln n$. For each k, consider the $n - m_k$ non informed nodes and let X_i, $i = 1, \dots, n - m_k$, be the random variable whose value is 1 if node i is informed in time slot k and 0 otherwise. Then, either $m_k \geqslant 1/(2p)$ and the lemma is proved, or it holds that

$$\begin{aligned}\Pr\{X_i = 1\} &= m_k p(1-p)^{m_k - 1} \\ &\geqslant \frac{m_k}{n} e^{-\frac{p}{1-p}(m_k - 1)} \\ &\geqslant \frac{m_k}{n} e^{-\frac{p}{1-p}\left(\frac{1}{2p} - 1\right)} \\ &\geqslant \frac{m_k}{n} e^{-1/2}\end{aligned}$$

Consider the random variable $X = \sum_{i=1}^{n-m_k} X_i$ counting the number of new informed nodes in time slot k. Then,

$$\mathbf{E}[X] \geqslant (n - m_k)\frac{m_k}{e^{1/2} n} \geqslant \frac{m_k}{2e^{1/2}}$$

Where in the last inequality we use $m_k \leqslant 1/(2p) \leqslant n/2$. From Eq. (1) with $\mu = \frac{m_k}{2e^{1/2}}$ and $\delta = 1/2$, it holds that

$$\Pr\left\{X \leqslant \frac{m_k}{4e^{1/2}}\right\} \leqslant e^{-\frac{1}{8}\frac{m_k}{2e^{1/2}}} \leqslant \frac{1}{n^{\frac{1}{16e^{1/2}}}}$$

where in the last inequality we used $m_k \geqslant \ln n$. So we have that, at each time slot k, either $m_k \geqslant 1/(2p)$, or

$$\Pr\{X \leqslant 4am_k\} \leqslant \frac{1}{n^a}$$

where $a = \frac{1}{16e^{1/2}}$. Observe that, from the above inequality, w.h.p., it holds that $m_{k+1} \geqslant (1+a)m_k$ for any k. Then consider the following recursive equation holds for

$$\begin{cases} m_{k+1} \geqslant (1+a)m_k \\ m_0 \geqslant \ln n \end{cases}$$

By solving the recurrence, we get $m_k \geqslant (1+a)^k \ln n$ for any k. Let $\bar{k}$ be the smallest time slot such that $(1+a)^{\bar{k}} \ln n \geqslant \frac{1}{2p}$ and observe that $\bar{k} \leqslant c \ln n$. Hence, we obtain

$$\begin{aligned}\Pr\left\{m_{c\ln n} \geqslant \frac{1}{2p}\right\} &\geqslant \Pr\left\{m_{\bar{k}} \geqslant \frac{1}{2p}\right\} \\ &\geqslant \Pr\left\{\forall k < \bar{k},\ m_{k+1} > (1+a)m_k\right\} \\ &= 1 - \Pr\left\{\exists k < \bar{k}\ :\ m_{k+1} \leqslant (1+a)m_k\right\} \\ &\geqslant 1 - \sum_{k=0}^{\bar{k}-1} \Pr\{m_{k+1} \leqslant (1+a)m_k\} \\ &\geqslant 1 - \frac{\bar{k}}{n^{2a}} \\ &\geqslant 1 - \frac{c \ln n}{n^{2a}} \geqslant 1 - \frac{1}{n^{2a-\varepsilon}}\end{aligned}$$

□

Observation B.3 *Let $m < n$ be the number of informed nodes and let u be a non informed one. If every informed node sends the message with probability q, then in $G_{n,p}$ the probability that node u receive 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 get the message and 0 otherwise. Then

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

□

Lemma B.4 *Let $p \geqslant 1/n$ and let γ be a constant such that $0 < \gamma < 1$. Assume we start a phase with at least γn informed nodes and, at each time slot, every informed node sends the message with probability q with $\frac{1}{enp} \leqslant q \leqslant \frac{1}{np}$. Then, after $c \ln n$ of such time slots, all nodes are informed w.h.p.*

Proof. Let m_k be the number of informed nodes at time slot k of this phase. From Lemma's hypothesis we have that $m_0 \geqslant \gamma n$. The probability that a non informed node receives the message in time slot k is

$$m_k pq(1-pq)^{m_k-1} \geqslant \gamma npq(1-pq)^{n-1} \geqslant \frac{\gamma}{e}\left(1-\frac{1}{n}\right)^{n-1} \geqslant \frac{\gamma}{e^2} \tag{4}$$

where we used that $\gamma n \leqslant m_k \leqslant n$ for each k and $\frac{1}{enp} \leqslant q \leqslant \frac{1}{np}$.
Now consider the $n - m_k$ non informed nodes. For $i = 1, 2, \ldots, n - m_k$, consider the event

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

From Eq. (4), it holds that

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

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

$$\Pr\{\exists i \;:\; \mathcal{E}_i\} = \Pr\left\{\bigcup_{i=1}^{n-m_k} \mathcal{E}_i\right\} \leqslant \sum_{i=1}^{n-m_k} \Pr\{\mathcal{E}_i\} \leqslant \frac{n-m_k}{n^{\gamma c/e^2}} \leqslant \frac{\gamma}{n^{\gamma c/e^2-1}}$$

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

Proof of Theorem 2.1. Two cases may arise. If $p \geqslant 1/\ln n$ then, from Lemma B.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 B.1 and B.2, after the second loop there are at least $1/(2p)$ informed nodes w.h.p. Finally, from Lemma 2.2, 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 B.4, after the last loop all the n nodes are informed w.h.p. □

Case p unknown.
Sketch of the proof of Theorem 2.4. Consider a subsequence of time slots of `BGI(n)` in which all informed nodes send the message with probability close to $1/m$, where m is the current number of informed nodes. In this time slot, the expected number of nodes that send the message is 1. From a slightly different variant of Lemma B.1 after $O(\log n)$ of such time slots, there are at least $\ln n$ informed nodes. Consider a subsequence of time slots of `BGI(n)` in which all informed nodes send the message with probability 1 (this happens for $k = 0$ in the inner loop). From slightly different variants of Lemmas B.1 and B.2, after $O(\log n)$ of such time slots, there are at least $1/(2p)$ informed nodes. Moreover, from Lemma 2.2, after the first execution of the inner `for` loop, there are at least $n/2$ informed nodes. So, in both cases after $2c \ln n$ time slots there are at least $n/2$ informed nodes. Now for each execution of the inner `for` loop, consider the time slot when $k = \lceil \ln np \rceil$. In this time slot each node sends the message with probability $q = e^{-\lceil \ln np \rceil}$ so that $1/(enp) \leqslant q \leqslant 1/(np)$, and thanks to Lemma B.4, we have that, after the $\lceil c \ln n \rceil$ execution of the inner `for` loop, all nodes are informed w.h.p. □

B.2 Lower Bounds

Case p known. In order to prove the logarithmic lower bound, we need the following Lemma which will be used for the case p unknown as well.

Lemma B.5 *Let $p < 1$ and consider any broadcasting protocol in 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 show that

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

As for $\tilde{p_k} = kp(1-p)^{k-1}$, we can prove that

$$\tilde{p_i} \leqslant \max\{p,\ 1 - 1/e\} \leqslant 1 - \frac{1-p}{e} \tag{5}$$

Let $X^k = \sum_{j=1}^n X_j^k$ so that the expected value $\mathbf{E}\left[X^k\right] = (n-m)\tilde{p_k}$. Observe that, for any fixed k, the random variables $X_1^k, \ldots, X_n^k$ are independent. It follows that

$$
\begin{aligned}
\Pr\left\{X^k \geqslant \left(1 - \frac{1-p}{2e}\right)(n-m)\right\} &\leqslant \Pr\left\{X^k \geqslant \tfrac{1+\tilde{p_k}}{2}(n-m)\right\} && \text{from Ineq. 5} \\
&\leqslant e^{-\lambda(1-\tilde{p_k})(n-m)} && \text{from Coroll. A.3} \\
&\leqslant e^{-\lambda\frac{1-p}{e}(n-m)} && \text{from Ineq. 5} \\
&\leqslant e^{-\frac{18\lambda}{e}\ln n} && \text{from Lemma's hypothesis} \\
&\leqslant \frac{1}{n} && (6)
\end{aligned}
$$

Let T be the random variable counting the number of nodes that transmit at time slot t. We then get

$$
\begin{aligned}
\Pr\left\{m' - m \geqslant \left(1 - \frac{1-p}{2e}\right)(n-m)\right\} &= \sum_{k=0}^{n} \Pr\left\{X^k \geqslant \left(1 - \frac{1-p}{2e}\right)(n-m)\right\} \Pr\{T = k\} \\
&\leqslant \frac{1}{n}\sum_{k=0}^{n} \Pr\{T = k\} = \frac{1}{n}
\end{aligned}
$$

□

Proof of Theorem 2.3. 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

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

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

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

Hence, for any $t \leqslant \alpha \ln n$, Lemma B.5 implies that

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

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

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

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

Case p unknown.
Proof of Lemma 2.6. 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

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

Define the interval $I = \left[\frac{1}{e^2 p \ln^2 n}, \frac{e^2 \ln n}{p}\right]$ and $\tau = \frac{1}{\ln^2 n}(n-m)$. Observe that, for any fixed k, the 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

$$\begin{aligned} \Pr\left\{m' - m > \frac{1}{\ln^2 n}(n-m)\right\} &= \sum_{i=1}^{m} \Pr\left\{X^i > \tau\right\} \Pr\{T = i\} \\ &= \sum_{i \notin I} \Pr\left\{X^i > \tau\right\} \Pr\{T = i\} + \sum_{i \in I} \Pr\left\{X^i > \tau\right\} \Pr\{T = i\} \qquad (8) \end{aligned}$$

Our goal is get an upper bound on each of the two sums in the right-hand side of the above inequality.

- We denote $\tilde{p}_i = ip(1-p)^{i-1}$ so that the expected value $\mu_i = \mathbf{E}\left[X^i\right] = (n-m)\tilde{p}_i$.

  Note that, for $i < \frac{1}{e^2 p \ln^2 n}$, it holds $\tilde{p}_i = ip(1-p)^{i-1} < \frac{1}{e^2 p \ln^2 n} p = \frac{1}{e^2 \ln^2 n}$, and, for $i > \frac{e^2 \ln n}{p}$, it holds $\tilde{p}_i = ip(1-p)^{i-1} < n \frac{1}{\ln^5 n} e^{-p\frac{e^2 \ln n}{p} + p} < \frac{e}{n^{e^2 - 1} \ln^5 n} < \frac{1}{e^2 \ln^2 n}$. hence

$$\mu_i = (n-m)\tilde{p}_i \leqslant \frac{\tau}{e^2} \quad \text{i.e. } \tau \geqslant e^2 \mu_i, \text{ for any } i \notin I. \qquad (9)$$

  Chernoff bound say that, for any $\delta_i > 0$, it holds

$$\Pr\left\{X^i \geqslant (1+\delta_i)\mu_i\right\} \leqslant \left[\frac{e^{\delta_i}}{(1+\delta_i)^{(1+\delta_i)}}\right]^{\mu_i}$$

  for any $i \notin I$ we choose $\delta_i = \tau/\mu_i - 1$ (note that by (9) it holds $\delta_i > 0$), thus we have

$$\begin{aligned} \Pr\left\{X^i \geqslant \tau\right\} &\leqslant \left[\frac{e^{\frac{\tau}{\mu_i} - 1}}{(\tau/\mu_i)^{\tau/\mu_i}}\right]^{\mu_i} = \frac{e^{\tau - \mu_i}}{(\tau/\mu_i)^{\tau}} \leqslant \left(\frac{e}{\tau}\right)^{\tau} {\mu_i}^{\tau} = e^{-\tau \ln \frac{\tau}{e\mu_i}} \\ &< e^{-\tau \ln e} < e^{-\frac{n - m_t}{\ln^2 n}} < e^{-\frac{\ln^3 n}{\ln^2 n}} = \frac{1}{n} < \frac{1}{2\ln^2 n} \end{aligned}$$

  and we use this upper bound to obtain

$$\sum_{i \notin I} \Pr\left\{X^i > \tau\right\} \Pr\{T = i\} < \sum_{i \notin I} \frac{1}{2\ln^2 n} \Pr\{T = i\} < \frac{1}{2\ln^2 n} \qquad (10)$$

- Now we get an upper bound on the second sum.

$$\sum_{i \in I} \Pr\left\{X^i > \tau\right\} \Pr\{T = i\} \leqslant \sum_{i \in I} \Pr\{T = i\} = \sum_{i \in I} \binom{m}{i} q^i (1-q)^{m_t - i} \qquad (11)$$

  we consider two cases:

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

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

where $\bar{i} = \lceil e^2 m\bar{q} \rceil$. From (11), using $\binom{x}{y} \leqslant \left(\frac{ex}{y}\right)^y$ and $p \leqslant \frac{1}{\ln^5 n}$ we have

$$\begin{aligned}
\sum_{i \in I} \Pr\{X > \tau\} \Pr\{T = i\} &\leqslant m\binom{m}{\bar{i}} q^{\bar{i}} (1-q)^{m-\bar{i}} \\
&< n\left(\frac{em}{\bar{i}} q\right)^{\bar{i}} \\
&\leqslant n\left(\frac{emq}{e^2 m\bar{q}}\right)^{\frac{1}{e^2 p \ln^2 n}} \\
&\leqslant n\left(\frac{1}{e}\right)^{\frac{\lg^3 n}{e^2}} \\
&\leqslant \left(\frac{1}{n}\right)^{\frac{\ln^2 n}{e^2} - 1} \\
&< \frac{1}{4\ln^2 n} \qquad (12)
\end{aligned}$$

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

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

where $\bar{i} = \left\lfloor \frac{m\bar{q}}{e^2} \right\rfloor$. From (11), using $\binom{x}{y} \leqslant \left(\frac{ex}{y}\right)^y$ and $p \leqslant \frac{1}{\ln^5 n}$ we have

$$\begin{aligned}
\sum_{i \in I} \Pr\{X > \tau\} \Pr\{T = i\} &\leqslant m\binom{m}{\bar{i}} q^{\bar{i}} (1-q)^{m-\bar{i}} \\
&< n\left(\frac{em}{\bar{i}}\right)^{\bar{i}} e^{-q(m-\bar{i})} \\
&\leqslant e^{\ln n + \bar{i}\ln\left(\frac{em}{\bar{i}}\right) - qm + \bar{i}} \\
&\leqslant e^{\ln n + \frac{e^2 \ln^2 n}{p} - \frac{e^4 \ln^2 n}{p} + \frac{e^2 \ln^2 n}{p}} \\
&\leqslant e^{\ln n - (e^4 - 2e^2)\ln^7 n} \\
&\leqslant \left(\frac{1}{n}\right)^{(e^4 - 2e^2)\ln^6 n - 1} \\
&< \frac{1}{4\ln^2 n} \qquad (13)
\end{aligned}$$

From (12) and (13) we get

$$\sum_{i \in I} \Pr\left\{X^i > \tau\right\} \Pr\{T = i\} < \sum_{i \in I} \frac{1}{\ln^2 n} \Pr\{T = i\} < \frac{1}{2\ln^2 n} \qquad (14)$$

The Lemma follows from (8), (10) and (14). □

Proof of Theorem 2.5. Let q_t be the probability transmission of the homogeneous protocol at time slot t, $1 \leqslant t \leqslant \bar{t}$, where $\bar{t} = \frac{\ln^2 n}{24 \ln\ln n}$. Define the intervals

$$Q_k = \left[\frac{1}{e^4 n p_k \ln^2 n}, \frac{4e^4 \ln^2 n}{n p_k}\right] \text{ where } p_k = \frac{1}{\ln^{5k} n} \text{ and } 1 \leqslant k \leqslant \frac{\ln n}{6 \ln\ln n}.$$

and consider the distribution in the $\frac{\ln n}{6\ln\ln n}$ intervals of the $\bar{t}$ values q_t. It must be an interval $Q_{\tilde{k}}$ containing at most $\frac{\ln n}{4}$ values.

Consider now an execution of the protocol 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

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

Claim B.6 *If the sequence of events $\mathcal{E}_1, \mathcal{E}_2, \dots \mathcal{E}_{\bar{t}}$ occurs, then it holds $m_{\bar{t}} \leqslant n - \frac{n^{\frac{1}{4}}}{24e}$.*

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

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

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

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

we have
$r_{\hat{t}} \geqslant \frac{n}{24}$ e $r_{i+1} = n - m_{i+1} \geqslant n - \lambda_i(n - m_i) - m_i = (1 - \lambda_i)(n - m_i)$ 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 slot where $q_i \in Q_{\tilde{k}}$, $\hat{t} \leqslant i \leqslant \bar{t}$.
Note that $s \leqslant \frac{\ln n}{4}$ and $\bar{t} < \ln^2 n - 1$. Hence $r_{\bar{t}} > \frac{n^{\frac{1}{4}}}{24e}$ and the Claim follows since $m_{\bar{t}} = n - r_{\bar{t}}$. □

From the Claim we have that when the sequence of events $\mathcal{E}_1, \mathcal{E}_2, \dots \mathcal{E}_{\bar{t}}$ occurs, the broadcast require more than $\bar{t} = \Omega\left(\frac{\log^2 n}{\log \log n}\right)$ time slots. To prove the Theorem we show that the sequence $\mathcal{E}_1, \mathcal{E}_2, \dots \mathcal{E}_{\bar{t}}$ occurrs with a constant probability.

Consider the event $\mathcal{E}'_i = m_{i+1} - m_i \leqslant \frac{1}{\ln^2 n}(n - m_i)$ where $m_i \geqslant \frac{n}{4}$ and $q_i \notin Q_{\tilde{k}}$ con $i \leqslant \bar{t}$
Note that

- from the Claim we get $\frac{n}{4} \leqslant m_i \leqslant n - \ln^3 n$
- since $1 \leqslant \tilde{k} \leqslant \frac{\ln n}{6 \ln \ln n}$, and by 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 the 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 $Q_{\tilde{k}}$. Hence, since by hypothesis $q_i \notin Q_{\tilde{k}}$, we get $q_i \notin Q'_{\tilde{k}}$.

Thus applying Lemma B.5, we have

$$\Pr\{\mathcal{E}'_i\} \geqslant 1 - \frac{1}{\ln^2 n} \tag{15}$$

Consider the event $\mathcal{E}"_i = m_{i+1} - m_i \leqslant \left(1 - \frac{1}{2e}\right)(n - m_i)$ 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 applying Lemma 2.6 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

$$\Pr\left\{\mathcal{E}"_i\right\} \geqslant 1 - \frac{1}{n} \geqslant 1 - \frac{1}{\ln^2 n} \tag{16}$$

From 15 e 16, for the probability of the sequence $\mathcal{E}_1, \mathcal{E}_2, \ldots \mathcal{E}_{\bar{t}}$ 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 $\bar{t} < \log^2 n$.

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