Rumor Spreading in Random Evolving Graphs

Andrea Clementi[1], Pierluigi Crescenzi[2], Carola Doerr[*3], Pierre Fraigniaud[†4], Francesco Pasquale[‡5], and Riccardo Silvestri[5]

[1] Università Tor Vergata di Roma
[2] Università di Firenze
[3] Université Paris Diderot and Max Planck Institute Saarbrücken
[4] CNRS and Université Paris Diderot,
[5] Sapienza Università di Roma

Abstract

We analyze randomized broadcast in dynamic networks modeled as edge-Markovian evolving graphs. We prove that, for the practically relevant case in which the birth rate $p = \Omega(\frac{1}{n})$, and the dead rate q is constant, randomized broadcast completes w.h.p. in $O(\log n)$ time steps. Our result provides the first formal argument demonstrating the robustness of the "push" protocol against network changes.

[*]Supported by a Feodor Lynen postdoctoral research fellowship of the Alexander von Humboldt Foundation and by the Agence Nationale de la Recherche under the project ANR-09-JCJC-0067-01.

[†]Additional support from the ANR project DISPLEXITY, and from the INRIA project GANG.

[‡]Partially supported by EU FET project MULTIPLEX 317532.

1 Introduction

1.1 Context and Objective

Rumor spreading is a well-known gossip-based distributed algorithm for disseminating information in large networks. According to the synchronous `Push` version of this algorithm, an arbitrary source node is initially informed, and, at each time step (a.k.a. round), an informed node u chooses one of its neighbors v uniformly at random, and this node becomes informed at the next time step.

Rumor spreading (originally called *rumor mongering*) was first introduced in [13], in the context of replicated databases, as a solution to the problem of distributing updates and driving replicas towards consistency. Successively, it has been proposed in several other application areas, such as failure detection in distributed systems [36], peer-sampling [29], adaptive machine discovery [26], and distributed averaging in sensor networks [5] (for a nice survey of gossip-based algorithm applications, see also [31]). Apart from its applications, rumor spreading has also been deeply analyzed from a theoretical and mathematical point of view. Indeed, as already observed in [13], rumor spreading is just an example of an epidemic process: hence, its analysis "benefits greatly from the existing mathematical theory of epidemiology" (even if its application in the field of distributed systems has almost opposite goals). In particular, the *completion time* of rumor spreading, that is, the number of steps required in order to have all nodes informed with high probability[1] (w.h.p.), has been investigated in the case of several different network topologies, such as complete graphs [21, 34, 30], hypercubes [16], random graphs [16, 18, 19], preferential attachment graphs [6, 14], and some power-law degree graphs [20]. Besides obtaining bounds on the completion time of rumor spreading, most of these works also derive deep connections between the completion time itself and some classic measures of graph spectral theory, such as, for example, the *conductance* of a graph (as far as we know, the most recent results of this kind are the ones presented in [7, 8, 22]) or its *vertex expansion* (see [35, 23]).

It is important to observe that the techniques and the arguments adopted in these studies strongly rely on the fact that the underlying graph is *static* and does not change over time. For instance, most of these analyses exploit the crucial fact that the degree of every node (no matter whether this is a random variable or a deterministic value) never changes during the entire execution of the rumor spreading algorithm. It is then natural to ask ourselves what is the speed of rumor spreading in the case of *dynamic* networks, where nodes and edges can appear and disappear over time (several emerging networking technologies such as ad hoc wireless, sensor, mobile networks, and peer-to-peer networks are indeed inherently dynamic).

In order to investigate the behavior of distributed protocols in the case of dynamic networks, the concept of evolving graph has been introduced in the literature. An *evolving graph* is a sequence of graphs $(G_t)_{t\geq 0}$ where $t \in \mathbb{N}$ (to indicate that we consider the graph *snapshots* at discrete time steps t, although it may evolve in a continuous manner) with the same set of n nodes.[2] This concept is general enough for allowing us to model basically any kind of network evolution, ranging from *adversarial* evolving graphs (see, for example, [11, 32]) to *random* evolving graphs (see, for example, [4]).

Indeed, although only the edges are subject to changes, a node whose all incident edges are not present at a given step t can be considered as having left the network at time t, where the network is viewed as the giant component of G_t. Hence, the concept of evolving graph also captures some essence of the node dynamics. In the case of *random* evolving graphs, at each time step, the graph G_t is chosen randomly according to some probability distribution over a specified family of graphs. One very well-known and deeply studied example of such a family is the set $\mathcal{G}_{n,p}$ of *Erdős-Rényi* random graphs [1, 15, 24]. In the evolving graph setting, at every time step t, each possible edge exists with probability p (independently of the previous

[1] An event holds with high probability if it holds with probability at least $1 - 1/n^c$ for some constant $c > 0$.

[2] As far as we know, this definition has been formally introduced for the first time in [17].

graphs $G_{t'}$, $t' < t$, and independently of the other edges in G_t).

Random evolving graphs can exhibit communication properties which are much stronger than static networks having the same expected edge density (for a recent survey on computing over dynamic networks, see [33]). This has been proved in the case of the simplest communication protocol that implements the broadcast operation, that is, the `Flooding` protocol, according to which a source node is initially informed, and, whenever an uninformed node has an informed neighbor, it becomes informed itself at the next time step. It has been shown [3, 10, 12] that the `Flooding` completion time may be very fast (typically poly-logarithmic in the number of nodes) even when the network topology is, w.h.p., sparse, or even highly disconnected at every time step. Therefore, such previous results provide analytical evidences of the fact that random network dynamics not only do not hurt, but can actually help data communication, which is of the utmost importance in several contexts, such as, e.g., delay-tolerant networking [37, 38].

The same observation has been made when the model includes some sort of *temporal* dependency, as it is in the case of the random *edge-Markovian* model. According to this model, the evolving graph starts with an arbitrary initial graph G_0, and, at every time step t,

- if an edge does not exist in G_t, then it will appear in the next graph G_{t+1} with probability p, and
- if an edge exists in G_t, then it will disappear in the next graph G_{t+1} with probability q.

For every initial graph G_0, an edge-Markovian evolving graph will eventually converge to a (random) graph in $\mathcal{G}_{n,\tilde{p}}$ with stationary edge-probability $\tilde{p} = \frac{p}{p+q}$. However, there is a Markovian dependence between graphs at two consecutive time steps, hence, given G_t, the next graph G_{t+1} is not necessarily a random graph in $\mathcal{G}_{n,\tilde{p}}$. Interestingly enough, the edge-Markovian model has been recently subject to experimental validations, in the context of sparse opportunistic mobile networks [38], and of dynamic peer-to-peer systems [37]. These validations demonstrate a good fitting of the model with some real-world data traces. The completion time of the `Flooding` protocol has been recently analyzed in this model, for all possible values of $\tilde{p}$ (see [3, 12]). A variant of the model, in which the "birth" and "death" probabilities p and q depend not only on the number of nodes but also on some sort of distance between the nodes, has been investigated in [25].

The `Flooding` protocol however generates high message complexity. Moreover, although its completion time is an interesting analog for dynamic graphs of the diameter for static graphs, it is not reflecting the kinds of gossip protocols mentioned at the beginning of this introduction, used for practical applications. Hence the main objective of this paper is to analyze the more practical `Push` protocol, in edge-Markovian evolving graphs.

1.2 Framework

We focus our attention on dynamic network topologies yielded by the edge-Markovian evolving graphs for parameters p (*b*irth) and q (*d*eath) that correspond to a good fitting with real-world data traces, as observed in [37, 38]. These traces describe networks with relatively high dynamics, for which the death probability q is at least one order of magnitude greater than the birth probability p. In order to set parameters p and q fitting with these observations, let us consider the expected number of edges $\bar{m}$, and the expected node-degree $\bar{d}$ at the stationary regime, governed by $\tilde{p} = \frac{p}{p+q}$. We have $\bar{m} = \frac{p}{p+q}\binom{n}{2}$, and $\bar{d} = \frac{2\bar{m}}{n} = (n-1)\frac{p}{p+q}$. Thus, at the stationary regime, the expected number of edges ν that switch their state (from non existing to existing, or vice versa) in one time step satisfies

$$\nu = \bar{m}q + (\binom{n}{2} - \bar{m})p = \frac{n(n-1)}{2}\left(\frac{pq}{p+q} + \left(1 - \frac{p}{p+q}\right)p\right) = n(n-1)\frac{pq}{p+q} = nq\bar{d}.$$

Hence, in order to fit with the high dynamics observed in real-world data traces, we set q constant, so that a constant fraction of the edges disappear at every step, while a fraction p of the non-existing edges appear. We consider an arbitrary range for p, with the unique assumption that $p \geq \frac{1}{n}$. (For smaller p's, the completion time of any communication protocol is subject to the expected time $\frac{1}{np} \gg 1$ required for a node to acquire just one link connected to another node). To sum up, we essentially focus on the following range of parameters:

$$\frac{1}{n} \leqslant p < 1 \text{ and } q = \Omega(1). \tag{1}$$

This range includes network topologies for a wide interval of expected edge density (from very sparse and disconnected graphs, to almost-complete ones), and with an expected number of switching edges per time step equal to some constant fraction of the expected total number of edges.

Remark. It is worth noticing that analyzing the Push protocol in edge-Markovian graphs is not only subject to temporal dependencies, but also to *spatial* dependencies, thus making the analysis of the Push protocol more challenging. This holds even in the simpler random evolving graph model, i.e., the sequence of independent random graphs $G_t \in \mathcal{G}_{n,p}$. Indeed, even if this case does not include temporal dependencies, the Push protocol introduces spatial dependences that has to be carefully handled. To see why, consider a time step of the Push protocol, where we have k informed nodes, and let us try to evaluate how many new informed nodes there will be in the next time step. Given an informed node u, let $\delta(u)$ be the neighboring node selected by u according to the Push protocol (i.e., $\delta(u)$ is chosen uniformly at random among the current neighbors of u). By conditioning on the degree of u, it is not hard to calculate the probability that $\delta(u) = v$, for any non informed node v. However, the events "$\delta(u_1) = v_1$" and "$\delta(u_2) = v_2$" are not necessarily independent. Indeed, the event "$\delta(u_1) = v_1$" decreases the probability of the existence of an edge between u_1 and u_2, and so it affects the value of the random variable $\delta(u_2)$. This positive dependency prevents us from using the classical methods for analyzing the Push protocol in static graphs, or makes the use of these methods far more complex.

1.3 Our results

For the parameter range in Eq. (1), we show that, w.h.p., starting from any n-node graph G_0, the Push protocol informs all n nodes in $\Theta(\log n)$ time steps. Hence, in particular, even if the graph G_t is w.h.p. disconnected at every time step (this is the case for $p \ll \frac{\log n}{n}$), the completion time of the Push protocol is as small as it could be (the Push protocol cannot perform faster than $\Omega(\log n)$ steps in any static or dynamic graph since the number of informed nodes can at most double at every step). It is also interesting to compare the performances of the Push protocol with the one of Flooding . The known lower bound for Flooding on edge-Markovian graphs [12] (which is clearly a lower bound for Push , too) demonstrates that for $p = \Theta(1/n)$, the two protocols have the same asymptotic completion time. Moreover it is clear that, for $p = \Omega(1/n)$, the completion-time slowdown factor of the Push protocol is at most logarithmic. This property is a remarkable one, since the expected number of exchanged messages per node in Push may be exponentially smaller than the one in Flooding (for instance, consider the case $p = \Theta(1/\sqrt{n})$ which corresponds to an expected node degree $\Theta(\sqrt{n})$).

We also address another range of parameters p and q. Although it does not precisely fit with the measures in [37, 38], they can be of independent interest for other settings. This second case is the sequence of independent $\mathcal{G}_{n,p}$ graphs, that is, the case where $p + q = 1$. Actually, the analysis of this special case will allow us to focus on the first important probabilistic issue that needs to be solved: spatial dependencies. Indeed, even in this case, as already mentioned, the Push protocol induces a positive correlation among some crucial events that determine the number of new informed nodes at the next time step. This holds despite the fact that every edge is set independently from the others. For a sequence of independent $\mathcal{G}_{n,p}$

graphs, we prove that for every p (i.e., also for $p = o(\frac{1}{n})$) and $q = 1 - p$ the completion time of the Push protocol is, w.h.p., $\mathcal{O}(\log n/(\hat{p}n))$, where $\hat{p} = \min\{p,\, 1/n\}$. By comparing the lower bound for Flooding in [12], it turns out that this bound is tight, even for very sparse graphs.

Remark. Notice that, in [9, 27], by using a different approach, Isopi and Panconesi show a logarithmic bound for more "static" network topologies, i.e., for the range $p = \frac{c}{n}$ where $c > 0$ is a constant, and for any $q \in (0, 1)$. This parameter range includes edge-Markovian graphs with a small expected number of switching edges (this happens when $q = o(1)$). In this case, too, Push completes, w.h.p., in $O(\log n)$ rounds. This gives yet another evidence that dynamism helps.

Structure of the paper. In Section 2, we present our terminology and some preliminary concepts that will be used throughout the paper. In Section 3, we consider the independent dynamic Erdős-Rényi graphs, while Section 4 provides the analysis of the Push protocol in the the case of the edge-Markovian evolving graph model. In Section 5, finally, we summarize our results and propose some future research.

2 Preliminaries

The number of vertices in the graph will always be denoted by n. We abbreviate $[n] := \{1, \ldots, n\}$ and $\binom{[n]}{2} := \{\{i, j\} \mid i, j \in [n]\}$. For any subset $E \subseteq \binom{[n]}{2}$ and any two subsets $A, B \subseteq [n]$, define

$$E(A) = \{ \text{ edges of } E \text{ incident to } A \ \} \ \text{ and } \ E(A, B) = \{\{u, v\} \in E \mid u \in A, v \in B\}.$$

We consider the edge-Markovian evolving graph model $\mathcal{G}(n, p, q; E_0)$ [10] where E_0 is the starting set of edges.
The Push Protocol over $\mathcal{G}(n, p, q; E_0)$ can be represented as a random process over the set $\mathcal{S}$ of all possible pairs (E, I) where E is a subset of edges and I is a subset of nodes. In particular, the combined Markov process works as follows

$$\ldots \rightarrow \ (E_t, I_t) \ \stackrel{\text{edge-Markovian}}{\longrightarrow} \ (E_{t+1}, I_t) \ \stackrel{\texttt{Push}\text{ protocol}}{\longrightarrow} \ (E_{t+1}, I_{t+1}) \ \stackrel{\text{edge-Markovian}}{\longrightarrow} \ \ldots$$

where E_t and I_t represent the set of existing edges and the set of informed nodes at time t, respectively. All events, probabilities and random variables are defined over the above random process. Given a graph $G = ([n], E)$, a node $v \in [n]$, and a subset of nodes $A \subseteq [n]$ we define $\deg_G(v, A) = |\{(v, a) \in E \mid a \in A\}|$. When we have a sequence of graphs $\{G_t = ([n], E_t) \ : \ t \in \mathbb{N}\}$ we write $\deg_t(v, A)$ instead of $\deg_{G_t}(v, A)$. Given a graph G and an informed node $u \in I$, we define $\delta_G(u)$ as the random variable indicating the node selected by u in graph G according to the Push protocol. When G and/or t are clear from the context, they will be omitted.

3 Warm up: the time-independent case

In this section we analyze the special case of a sequence of independent $G_{n,p}$ (observe that a sequence of independent $G_{n,p}$ is a special case of edge-Markovian evolving graph with $q = 1 - p$). We show that the completion time of the Push protocol is $\mathcal{O}(\log n/(\hat{p}n))$ w.h.p., where $\hat{p} = \min\{p,\, 1/n\}$. In Theorem 1 we prove the result for $p \geqslant 1/n$ and in Theorem 2 for $p \leqslant 1/n$. From the lower bound on the flooding time for edge-Markovian evolving graphs [12], it turns out that our bound is optimal.

As mentioned in the introduction, even though in this case there is no time-dependency in the sequence of graphs, the `Push` protocol introduces a kind of dependence that has to be carefully handled. The key challenge is to evaluate the probability that a non-informed node v receives the information from at least one of the informed nodes; i.e., $1 - \mathbf{P}\left(\cap_{u \in I}\{\delta(u) \neq v\}\right)$. In order to overcome that dependency issue we consider the `Push` operation on a *modified* random graph where the corresponding events are independent and we prove that the number of new informed nodes in the original random graph is stochastically at least as large as in the modified version.

Definition 3.1 ((I, b)-modified graph) *Let G be a graph with node-set $[n]$, let $I \subseteq [n]$ be a subset of nodes, and let b be an integer with $1 \leqslant b \leqslant n$. The (I, b)-modified G is the graph H with node-set $[n] \cup \{v_1, \ldots, v_b\}$, where $\{v_1, \ldots, v_b\}$ is a set of extra* virtual *nodes, obtained from G by the following operations:*

1. *For every node $u \in I$ with $\deg_G(u) > b$, remove all edges incident to u;*
2. *For every node $u \in I$ with $\deg_G(u) \leqslant b$, add all edges $\{u, v_1\}, \ldots, \{u, v_b\}$ between u and the virtual nodes;*
3. *Remove all edges between any pair of nodes that are both in I.*

In the next lemma we prove that, if the informed nodes perform a `Push` operation both in a graph and in its modified version, then the number of new informed nodes in the original graph is (stochastically) larger than the number of informed nodes in the modified one. We will then apply this result to $G_{n,p}$ random graphs.

Lemma 3.2 (Virtual nodes) *Let $G = ([n], E)$ be a graph and let b be an integer with $1 \leqslant b \leqslant n$. Let $I \subseteq [n]$ be a set of nodes performing a `Push` operation in graphs G and H, where H is the (I, b)-modified G according to Definition 3.1. Let X and Y be the random variables counting the numbers of new informed nodes in G and H respectively. Then for every $h \in [0, n]$ it holds that $\mathbf{P}(X \leqslant h) \leqslant \mathbf{P}(Y \leqslant h)$.*

Proof. Let $u \in I$ be an informed node and consider the following coupling of random variables $\delta_G(u)$ and $\delta_H(u)$ indicating the nodes selected by u according to the `Push` operation in graphs G and H respectively.

If $\deg_G(u) \leqslant b$ then let $\delta_H(u)$ be uniform over the neighbors of u in H and let $\delta_G(u)$ be chosen in the following way: If $\delta_H(u) \in [n] \setminus I$ then $\delta_G(u) = \delta_H(u)$; otherwise (i.e., when $\delta_H(u)$ is a virtual node) let $\delta_G(u)$ be uniform over the informed neighbors of u in G with probability $1 - x$, and uniform over the non-informed ones with probability x, where $x = \frac{k(b-h)}{(h+k)b}$ and h and k are the numbers of informed and non-informed neighbors of u respectively.

If $\deg_G(u) > b$ then u performs the `Push` operations independently in G and H (notice that when $\deg_G(u) > b$ node u has no neighbors in H).

By construction we have that the set of new (non-virtual) informed nodes in H is a subset of the set of new informed nodes in G. Moreover, it is easy to check that $\delta_G(u)$ is uniform over the set of neighbors of u in G. □

Let I be the set of informed nodes performing a `Push` operation on a $G_{n,p}$ random graph. As previously observed, if $v \in [n] \setminus I$ is a non-informed node, then events $\{\{\delta_G(u) = v\} \,:\, u \in I\}$ are not independent, but events $\{\{\delta_H(u) = v\} \,:\, u \in I\}$ on the (I, b)-modified graph H are independent because of Operation 3 in Definition 3.1. In the next lemma we use this fact and Lemma 3.2 to evaluate the increasing rate of the number of informed nodes over a sequence of independent $G_{n,p}$.

Lemma 3.3 (The increasing rate of informed nodes) *Let $I \subseteq [n]$ be the set of informed nodes, let G be a new $G_{n,p}$ random graph with node-set $[n]$, and let X be the random variable counting the number of nodes in $[n] \setminus I$ that get informed after the* `Push` *operation performed in G by nodes in I. It holds that $\mathbf{P}\left(X \geqslant \lambda \cdot \min\{|I|,\, n-|I|\}\right) \geqslant \lambda$, where λ is a positive constant.*

Proof. Let H be the $(I, 3np)$-modified version of $G = ([n], E)$ according to Definition 3.1. Now we show that the number of nodes that gets informed in H is at least $\lambda \cdot \min\{|I|, n-|I|\}$ with probability at least λ, for a suitable constant λ.

Let $u \in I$ be an informed node and let $v \in [n] \setminus I$ be a non-informed one. Observe that by the definition of H, u cannot choose v in H if the edge $\{u,v\} \notin E$ or if the degree of u in G is larger than $3np$ (see Operation 3 in Definition 3.1). Thus the probability that node u chooses node v in random graph H according to the `Push` protocol is

$$\mathbf{P}\left(\delta_H(u) = v\right) = \mathbf{P}\left(\delta_H(u) = v \mid \{u,v\} \in E \,\wedge\, \deg_G(u) \leqslant 3np\right) \mathbf{P}\left(\{u,v\} \in E \,\wedge\, \deg_G(u) \leqslant 3np\right). \tag{2}$$

If $\deg_G(u) \leqslant 3np$ then node u in H has exactly $3np$ virtual neighbors plus at most other $3np$ non-informed neighbors. It follows that

$$\mathbf{P}\left(\delta_H(u) = v \,|\, \{u,v\} \in E \,\wedge\, \deg_G(u) \leqslant 3np\right) \geqslant 1/(6np). \tag{3}$$

We also have that

$$\begin{aligned} \mathbf{P}\left(\{u,v\} \in E \,\wedge\, \deg_G(u) \leqslant 3np\right) &= \mathbf{P}\left(\{u,v\} \in E\right) \mathbf{P}\left(\deg_G(u) \leqslant 3np \mid \{u,v\} \in E\right) \\ &= p \cdot \mathbf{P}\left(\deg_G(u) \leqslant 3np \mid \{u,v\} \in E\right). \end{aligned}$$

Since $\mathbf{E}\left[\deg_G(u) \,|\, \{u,v\} \in E\right] \leqslant np+1$ with $np \geqslant 1$, from Chernoff bound (Lemma A.1 in Appendix A) positive constants c and $\beta < 1$ exist such that

$$\mathbf{P}\left(\deg_G(u) > 3np \,|\, \{u,v\} \in E\right) \leqslant \mathbf{P}\left(\deg_G(u) > 2np+1 \,|\, \{u,v\} \in E\right) \leqslant e^{-cnp} = \beta < 1. \tag{4}$$

By using (3) and (4) in (2) we get $\mathbf{P}\left(\delta_H(u) = v\right) \geqslant \frac{\alpha}{n}$, for some constant $\alpha > 0$.

Since the events $\{\{\delta_H(u) = v\}, v \in I\}$ are independent, the probability that node v is not informed in H is thus

$$\mathbf{P}\left(\cap_{u \in I} \delta_H(u) \neq v\right) \leqslant (1-\alpha/n)^{|I|} \leqslant e^{-\alpha|I|/n}.$$

Let Y be the random variable counting the number of new informed nodes in H. The expectation of Y is

$$\mathbf{E}[Y] \geqslant (n-|I|)\left(1 - e^{-\alpha|I|/n}\right) \geqslant (\alpha/2)(n-|I|)|I|/n.$$

Hence we get

$$\mathbf{E}[Y] \geqslant \begin{cases} (\alpha/4)|I| & \text{if } |I| \leqslant n/2\,, \\ (\alpha/4)(n-|I|) & \text{if } |I| \geqslant n/2\,. \end{cases}$$

Since $Y \leqslant \min\{|I|,\, n-|I|\}$, from Observation B.2 in Appendix B it follows that

$$\mathbf{P}\left(Y \geqslant (\alpha/8) \cdot \min\{|I|,\, n-|I|\}\right) \geqslant \alpha/8\,.$$

Finally the thesis follows from Lemma 3.2. □

We can now derive the upper bound on the completion time of the `Push` protocol over a sequence of independent $G_{n,p}$ random graphs.

Theorem 1 *Let $\mathcal{G} = \{G_t \,:\, t \in \mathbb{N}\}$ be a sequence of independent $G_{n,p}$ with $p \geqslant 1/n$. The completion time of the* `Push` *protocol over $\mathcal{G}$ is $\mathcal{O}(\log n)$ w.h.p.*

Proof. Consider a generic time step t of the execution of the `Push` protocol where $I_t \subseteq [n]$ is the set of informed nodes and $m_t = |I_t|$ is its size. For any t such that $m_t \leqslant n/2$, Lemma 3.3 implies that $\mathbf{P}\left(m_{t+1} \geqslant (1+\lambda)m_t\right) \geqslant \lambda$, where λ is a positive constant. Let us define event $\mathcal{E}_t = \{m_t \geqslant (1+\lambda)m_{t-1}\} \vee \{m_{t-1} \geqslant n/2\}$ and let $Y_t = Y_t((E_1, I_1), \ldots, (E_t, I_t))$ be the indicator random variable of that event. Observe that if $t = \frac{\log n}{\log(1+\lambda)}$ then $(1+\lambda)^t \geqslant n/2$. Hence, if we set $T_1 = \frac{2}{\lambda}\frac{\log n}{\log(1+\lambda)}$, we get

$$\mathbf{P}\left(m_{T_1} \leqslant n/2\right) \leqslant \mathbf{P}\left(\sum_{t=1}^{T_1} Y_t \leqslant (\lambda/2)T_1\right).$$

The above probability is at most as large as the probability that in a sequence of T_1 independent coin tosses, each one giving `head` with probability λ, we see less than $(\lambda/2)T_1$ `heads` (see e.g. Lemma 3.1 in [2]). A direct application of Chernoff bound (Lemma A.1 in Appendix A) shows that this probability is smaller than $e^{-(1/4)\lambda T_1} \leqslant n^{-c}$, for a suitable constant $c > 0$. We can thus state that, after $\mathcal{O}(\log n)$ time steps, there at least $n/2$ informed nodes w.h.p.
If $m_{T_1} \geqslant n/2$, then, for every $t \geqslant T_1$, Lemma 3.3 implies that $\mathbf{P}\left(n - m_{t+1} \leqslant (1-\lambda)(n-m_t)\right) \geqslant \lambda$. Observe that if $t = \frac{\log n}{\lambda}$ then $(1-\lambda)^t \leqslant 1/n$, so that for $T_2 = \frac{2}{\lambda} \cdot \frac{\log n}{\lambda} + T_1$ the probability that the `Push` protocol has not completed at time T_2 is

$$\mathbf{P}\left(m_{T_2} < n\right) \leqslant \mathbf{P}\left(m_{T_2} < n \,|\, m_{T_1} \geqslant \frac{n}{2}\right) + \mathbf{P}\left(m_{T_1} < \frac{n}{2}\right).$$

As we argued in the analysis of the spreading till $n/2$, the probability $\mathbf{P}\left(m_{T_2} < n \,|\, m_{T_1} \geqslant \frac{n}{2}\right)$ is not larger than the probability that in a sequence of $\frac{2}{\lambda} \cdot \frac{\log n}{\lambda}$ independent coin tosses, each one giving `head` with probability λ, there are less than $\frac{\log n}{\lambda}$ `heads`. Again, by applying the Chernoff bound (Lemma A.1 in Appendix A), the latter is not larger than n^{-c} for a suitable positive constant c. □

In order to prove the bound for $p \leqslant 1/n$, we first show that one single `Push` operation over the union of a sequence of graphs informs (stochastically) less nodes than the sequence of `Push` operations performed in every single graph (this fact will also be used in Section 4 to analyse the edge-MEG).

Lemma 3.4 (Time windows) *Let $\{G_t = ([n], E_t) \,:\, t = 1, \ldots, T\}$ be a finite sequence of graphs with the same set of nodes $[n]$. Let $I \subseteq [n]$ be the set of informed nodes in the initial graph G_1. Suppose that at every time step every informed node performs a* `Push` *operation, and let X be the random variable counting the number of informed nodes at time step T. Let $H = ([n], F)$ be such that $F = \cup_{t=1}^{T} E_t$ and let Y be the random variable counting the number of informed nodes when the nodes in I perform one single* `Push` *operation in graph H. Then for every $\ell = 0, 1, \ldots, n$ it holds that $\mathbf{P}(X \leqslant \ell) \leqslant \mathbf{P}(Y \leqslant \ell)$.*

Proof. Consider the sequence of graphs $\{H_t = ([n], F_t) \,:\, t = 1, \ldots, T\}$ where graph H_t is the union of graphs $G_1, \ldots, G_t$, i.e. for every t we set $F_t = \bigcup_{i=1}^{t} E_i$. We inductively construct one single `Push` operation in $H \equiv H_T$, building it on the probability space of the `Push` protocol in $(G_1, \ldots, G_T)$, in a way that the set of informed nodes in H is a subset of the set of informed nodes in G_T.

For every node u that is informed at the beginning of the process, i.e. $u \in I$, and for every $t = 1, \ldots, T$, let N_t be the set of neighbors of u in graph G_t, let $d_t = |N_t|$ be its size, let $h_t = |\bigcup_{i=1}^{t} N_i|$ be the number of neighbors of u in graph H_t, and let $\delta_{G_t}(u)$ be the random variable indicating the neighbor chosen by u

u.a.r. in N_t. Finally, let $\{C_t : t = 2, \ldots, T\}$ be a sequence of independent Bernoulli random variables with $\mathbf{P}(C_t = 1) = d_t/h_t$. Now we recursively define random variables $\delta_{H_1}(u), \ldots, \delta_{H_T}(u)$:
Define $\delta_{H_1}(u) = \delta_{G_1}(u)$. For $t = 2, \ldots, T$ define

$$\delta_{H_t}(u) = \begin{cases} \delta_{G_t}(u) & \text{if } \delta_{G_t}(u) \in N_t \setminus \left(\bigcup_{i=1}^{t-1} N_i \right) \text{ and } C_t = 1 \\ \delta_{H_{t-1}}(u) & \text{otherwise} \end{cases} \tag{5}$$

By construction, it holds that $\delta_{H_T}(u) \in \{\delta_{G_1}(u), \ldots, \delta_{G_T}(u)\}$, hence the set of informed nodes in H_T is a subset of the set of informed nodes in G_T. Now we show that for every t node u chooses one of its neighbors uniformly at random in H_t, i.e. for every $v \in \bigcup_{i=1}^{t} N_i$ it holds that $\mathbf{P}(\delta_{H_t}(u) = v) = 1/h_t$.

We proceed by induction on t. The base of the induction directly follows from the choice $\delta_{H_1}(u) = \delta_{G_1}(u)$. Now assume that for every $v \in \bigcup_{i=1}^{t-1} N_i$ it holds that $\mathbf{P}\left(\delta_{H_{t-1}}(u) = v\right) = 1/h_{t-1}$ and let $v \in \bigcup_{i=1}^{t} N_i$. We distinguish two cases:

- If $v \in N_t \setminus \left(\bigcup_{i=1}^{t-1} N_i \right)$ then, according to (5) we have that $\delta_{H_t}(u) = v$ if and only if $\delta_{G_t}(u) = v$ and $C_t = 1$, hence

$$\mathbf{P}(\delta_{H_t}(u) = v) = \mathbf{P}(\delta_{G_t}(u) = v \wedge C_t = 1) = \frac{1}{d_t} \cdot \frac{d_t}{h_t} = \frac{1}{h_t}$$

- If $v \in \bigcup_{i=1}^{t-1} N_i$ then we have that $\delta_{H_t}(u) = v$ if and only if $\delta_{H_{t-1}}(u) = v$ and at least one of the two conditions in (5) does not hold (that is $C_t = 0$ or $\delta_{G_t}(u) \in N_t \cap \left(\bigcup_{i=1}^{t-1} N_i \right)$). Hence,

$$\mathbf{P}(\delta_{H_t}(u) = v) = \mathbf{P}\left(\delta_{H_{t-1}}(u) = v\right) \left[\mathbf{P}(C_t = 0) + \mathbf{P}\left(\delta_{G_t}(u) \in N_t \cap \left(\bigcup_{i=1}^{t-1} N_i \right) \wedge C_t = 1 \right) \right]$$

By the induction hypothesis we have that $\mathbf{P}\left(\delta_{H_{t-1}}(u) = v\right) = 1/h_{t-1}$, and by observing that the size of $N_t \cap \left(\bigcup_{i=1}^{t-1} N_i \right)$ is $d_t + h_{t-1} - h_t$ it follows that

$$\mathbf{P}(\delta_{H_t}(u) = v) = \frac{1}{h_{t-1}} \left(\frac{h_t - d_t}{h_t} + \frac{d_t + h_{t-1} - h_t}{d_t} \cdot \frac{d_t}{h_t} \right) = \frac{1}{h_t}$$

□

Observe that if we look at a sequence of independent $G_{n,p}$ with $p \leqslant 1/n$ for a time-window of approximately $1/(np)$ time steps, then every edge appears at least once in the sequence with probability at least $1/n$. The above lemma thus allows us to reduce the case $p \leqslant 1/n$ to the case $p \geqslant 1/n$.

Theorem 2 *Let $\mathcal{G} = \{G_t : t \in \mathbb{N}\}$ be a sequence of independent $G_{n,p}$ with $p \leqslant 1/n$ and let $s \in [n]$. The* `Push` *protocol with source s over $\mathcal{G}$ completes the broadcast in $\mathcal{O}(\log n/(np))$ time steps w.h.p.*

Proof. Consider the sequence of random graphs $\mathcal{H} = \{H_s : s \in \mathbb{N}\}$ where H_s is the union of random graphs

$$H_s = ([n], F_s) \text{ such that } F_s = E_{sT} \cup E_{sT+1} \cup \cdots \cup E_{sT+T-1} \text{ with } T = 2/(np).$$

Observe that every H_s is a $G_{n,\hat{p}}$ with $\hat{p} \geqslant 1/n$. Indeed, the probability that an edge does not exist in F_s is

$$(1-p)^T \leqslant e^{-pT} = e^{-2/n}.$$

Hence the probability that the edge exists is $1 - e^{-2/n} \geqslant 1/n$.
Let $\tau_{\mathcal{G}}$ and $\tau_{\mathcal{H}}$ be the random variables indicating the completion time of the Push protocol over sequences $\mathcal{G}$ and $\mathcal{H}$ respectively. From Theorem 1 it follows that $\tau_{\mathcal{H}} = \mathcal{O}(\log n)$ w.h.p. and from Lemma 3.4 it follows that for every t it holds that

$$\mathbf{P}(\tau_{\mathcal{G}} \geqslant Tt) \leqslant \mathbf{P}(\tau_{\mathcal{H}} \geqslant t).$$

Hence, it holds that

$$\tau_{\mathcal{G}} = \mathcal{O}(T \log n) = \mathcal{O}\left(\frac{\log n}{np}\right) \text{ w.h.p.}$$

□

4 Edge-Markovian graphs with high dynamics

In this section we prove that the Push protocol over an edge-Markovian graph $\mathcal{G}(n, p, q; E_0)$ with $p \geqslant 1/n$ and $q = \Omega(1)$ has completion time $\mathcal{O}(\log n)$ w.h.p.

As observed in the Introduction, the stationary random graph is an Erdős-Rényi $G_{n,\tilde{p}}$ where $\tilde{p} = \frac{p}{p+q}$ and the mixing time of the edge Markov chain is $\Theta\left(\frac{1}{p+q}\right)$. Thus, if p and q fall into the range defined in (1), we get that the stationary random graph can be sparse and disconnected (when $p = o\left(\frac{\log n}{n}\right)$) and that the mixing time of the edge Markov chain is $O(1)$. Thus, we can omit the term E_0 and assume it is random according to the stationary distribution.

The time-dependency between consecutive snapshots of the dynamic graph does not allow us to obtain directly the *increasing rate* of the number of informed nodes that we got for the independent-$G_{n,p}$ model. In order to get a result like Lemma 3.3 for the edge-Markovian case, we need in fact a *bounded-degree* condition on the current set of informed nodes (see Definition 4.2) that does not apply when the number of informed nodes is *small* (i.e., smaller than $\log n$). However, in order to reach a state where at least $\log n$ nodes are informed, we can use a different ad-hoc technique that analyzes the spreading rate yielded by the source only.

Lemma 4.1 (The Bootstrap) *Let $\mathcal{G} = \mathcal{G}(n, p, q)$ be an edge-Markovian graph with $p \geqslant 1/n$ and $q = \Omega(1)$, and consider the* Push *protocol in $\mathcal{G}$ starting with one informed node. For any positive constant γ, after $\mathcal{O}(\log n)$ time steps there are at least $\gamma \log n$ informed nodes w.h.p.*

Proof. We consider the message-spreading process yielded by the source node only and, instead of directly analyzing this process on the edge-Markovian sequence $\{G_t = ([n], E_t) : t \in \mathbb{N}\}$, we consider it in the sequence $\{H_t = ([n], E_{2t} \cup E_{2t+1})\}$. Thanks to Lemma 3.4, this is feasible since the number of informed nodes in H_t is stochastically smaller than the number of informed nodes in G_{2t}. We split the analysis in two cases: $p \leqslant \log n/n$ and $p \geqslant \log n/n$.

Case $p \geqslant \log n/n$: Consider an arbitrary time step t during the execution of the protocol and for convenience' sake let us rename it $t = 0$. Let I_0 be the set of informed nodes in that time step with $|I_0| \leqslant \gamma \log n$. Let $H = ([n], E_1 \cup E_2)$ be the random graph obtained by taking the edges that are present in at least one of the next two time steps and consider the Push operation of the source node in H. Every edge has probability at least p in H (see Observation B.1 in Appendix B). In particular, every node v is connected to the source node s in H with probability at least p. Thus, if we name X the random variable counting the number of non-informed nodes connected to the source node in H, we have that the expectation of X is

$$\mathbf{E}[X] = \sum_{v \in [n] \setminus I_0} \mathbf{P}(\{s,v\} \in E_1 \cup E_2) \geqslant (n - |I_0|)p \geqslant 2\alpha np$$

for a suitable positive constant α. Since edges are independent, from Chernoff bound (Lemma A.1 in Appendix A) it follows that

$$\mathbf{P}(X \leqslant \alpha np) \leqslant e^{-\varepsilon np}$$

for a suitable positive constant ε. Hence, since $p \geqslant \log n/n$, it follows that there are at least $\alpha \log n$ nodes in $[n] \setminus I_0$ that are connected to s in H w.h.p. The probability that the source s, by applying the `Push` operation in H, sends the message to one of those nodes is

$$\begin{aligned}\mathbf{P}(\delta_H(s) \in [n] \setminus I_0) &\geqslant \mathbf{P}(\delta_H(s) \in [n] \setminus I_0 \mid X \geqslant \alpha \log n)\, \mathbf{P}(X \geqslant \alpha \log n) \\ &\geqslant \frac{\alpha \log n}{|I_0| + \alpha \log n} \mathbf{P}(X \geqslant \alpha \log n) \geqslant \lambda\end{aligned}$$

for a suitable positive constant λ.
From Lemma 3.4, the probability that there are no new informed nodes after two time steps is at most as large as the probability that the source node does not inform a new node in H; i.e.,

$$\mathbf{P}(I_2 = I_0) \leqslant \mathbf{P}(\delta_H(s) \notin [n] \setminus I_0) \leqslant 1 - \lambda\,.$$

Thus for every time step t during the bootstrap, if $p \geqslant \log n/n$, after two time steps there is at least one new informed node with probability at least λ; i.e.,

$$\mathbf{P}(|I_{t+2}| \geqslant |I_t| + 1) \geqslant \lambda\,.$$

Hence, after $(4\gamma/\lambda)\log n$ time steps, there are at least $\gamma \log n$ informed nodes w.h.p.

Case $p \leqslant \log n/n$: In order to analyze the bootstrap phase on the sequence $\{H_t = ([n], E_{2t} \cup E_{2t+1})\}$, we first condition on the event $\overline{F}$ that in the first $T = (4\gamma/\lambda)\log n$ time steps it never happens that a new edge appears between the source node and a node that is already informed. Formally, $\overline{F}$ is the complementary event of $F := \cup_{t=1}^{T} F_t$ where F_t denotes the event "In H_{t+1} at least one edge will appear between the source node and a previously informed node". As we will see below, we have $\mathbf{P}(F) = \mathcal{O}(\log^3 n/n)$ and $\mathbf{P}\left(|I_T| \leqslant \gamma \log n \,|\, \overline{F}\,|\right) \leq n^{-\varepsilon}$ for a suitable positive constant ε.
Observe that if an edge does not exist in H_t then it will appear in H_{t+1} with probability $1-(1-p)^2$. Since $p \leqslant \log n/n \leqslant 1/4$, by applying the standard inequalities

$$e^{-2x} \leqslant 1 - x \leqslant e^{-x}, \text{ for any } 0 \leqslant x \leqslant \frac{1}{2}$$

we get

$$2p \leqslant 1 - (1-p)^2 \leqslant 4p$$

For F_t as defined above we have

$$\mathbf{P}(F_t) \leqslant 4p|I_t| \leqslant 4\gamma \frac{\log^2 n}{n}\,, \tag{6}$$

where in the last inequality we used the facts that $p \leqslant \log n/n$ and that, during the bootstrap, $|I_t| \leqslant \gamma \log n$.

Now consider the two following events: S_1^t is the event "The source informs a new node in H_{t+1}" and S_2^t is the event "The number of edges between the source node and the set of informed nodes decreases in H_{t+1}"; i.e.,

$$S_1^t = \{|I_{t+1}| = |I_t| + 1\} \text{ and } S_2^t = \{\deg_{t+1}(s, I_{t+1}) \leqslant \deg_t(s, I_t) - 1\}$$

Now we show that, at every time step, at least one of the two events above holds with constant probability if event F_t does not hold. Indeed, in that case, if the number of informed nodes connected to the source node is zero, then if some non-informed node will be connected to the source node at the following time step we will have at least a new informed node (event S_1^t) and this happens with constant probability. If there is at least one informed node connected to the source, then if one of those edges will disappear then $\deg(s, I_t)$ will decrease (event S_2^t). More formally, if $\deg_t(s, I_t) = 0$ we have that

$$\mathbf{P}\left(S_1^t \,|\, \overline{F_t}\right) \geqslant 1 - (1-2p)^{n-|I_t|} \geqslant 1 - e^{-2p(n-|I_t|)} \geqslant 1 - e^{-(2/n)(n-|I_t|)} \geqslant 1 - e^{-1}\,.$$

If $\deg_t(s, I_t) \geqslant 1$, we get $\mathbf{P}\left(S_2^t \,|\, \overline{F_t}\right) \geqslant q$. Hence for $\lambda = \min\{q,\, 1-e^{-1}\}$, we have that

$$\mathbf{P}\left(S_1^t \vee S_2^t \,|\, \overline{F_t}\right) \;\geqslant\; \lambda\,. \tag{7}$$

If we define $T = (4\gamma/\lambda)\log n$ then we can show that after T time steps there are at least $\gamma \log n$ informed nodes w.h.p. Indeed, let X_1 and X_2 be the random variables indicating the number of time steps that events S_1 and S_2 hold, respectively. Remind that its complement $\overline{F}$ is the event "In the first T time steps it never happens that a new edge appears between the source node and a node that is already informed". Since $T = \mathcal{O}(\log n)$, from Eq. 6 it follows that $\mathbf{P}(F) = \mathcal{O}(\log^3 n/n)$. Moreover, observe that if event $\overline{F}$ holds then $X_1 \geqslant X_2$. Indeed, if no edge between the source and any previously informed node appears, then, when an edge between the source node and an informed node disappears (event of S_2 type), the source must have previously informed that node (S_1 event). Thus the probability that the bootstrap is not completed at time T is

$$\mathbf{P}\left(|I_T| \leqslant \gamma \log n\right) \;\leqslant\; \mathbf{P}\left(X_1 \leqslant \gamma \log n \,|\, \overline{F}\right) + \mathbf{P}(F) \;\leqslant\; \mathbf{P}\left(X_1 + X_2 \leqslant 2\gamma \log n \,|\, \overline{F}\right) + \mathbf{P}(F)\,.$$

Since from Eq. 7 we have that, at every time step, the event $S_1 \vee S_2$ holds with probability at least λ, then $\mathbf{P}\left(X_1 + X_2 \leqslant 2\gamma \log n \,|\, \overline{F}\right)$ is smaller than the probability that in a sequence of $T = (4\gamma/\lambda)\log n$ independent coin tosses, each one giving `head` with probability λ, we see less than $2\gamma \log n$ `heads`: this is smaller than $n^{-\varepsilon}$ for a suitable positive constant ε. □

We can now start the second part of our analysis where the `Push` operation of all informed nodes (forming the subset I) will be considered and, thanks to the bootstrap, we can assume that $|I| = \Omega(\log n)$.
As mentioned at the beginning of the section, we need to introduce the concept of *bounded-degree state* (E, I) of the Markovian process describing the information-spreading process over the dynamic graph, where E is the set of edges and I is the set of informed nodes.

Definition 4.2 (Bounded-Degree State) *A state (E, I) such that $|E(I)| \leqslant (8/q)n\tilde{p}|I|$ (where $\tilde{p} = \frac{p}{p+q}$ is the stationary edge probability) will be called a* bounded-degree *state.*

In the next lemma we show that, if I is the set of informed nodes with $|I| \geqslant \log n$, if in the starting random graph G_0 every edge exists with probability approximately $(1 \pm \varepsilon)p$, and if it evolves according to the edge-Markovian model and the informed nodes perform the `Push` protocol, then for a long sequence of time steps the random process is in a bounded-degree state. We will use this property in Theorem 3 by observing that, for every initial state, after $\mathcal{O}(\log n)$ time steps an edge-Markovian graph with $p \geqslant 1/n$ and $q \in \Omega(1)$ is in a state where every edge $\{u, v\}$ exists with probability $p_{\{u,v\}} \in [(1-\varepsilon)\tilde{p},\, (1+\varepsilon)\tilde{p}]$.

Lemma 4.3 *Let $\mathcal{G} = \mathcal{G}(n, p, q, E_0)$ be an edge-Markovian graph starting with G_0 and consider the* `Push` *protocol in $\mathcal{G}$ where I_0 is the set of informed nodes at time $t = 0$. Then, for any constant $c > 0$, for a sequence of $c \log n$ time steps every state is a bounded-degree one w.h.p.*

Proof. Let us fix $c = 8/q$ as in Definition 4.2. We show that (E_0, I_0) is a bounded-degree state w.h.p. and that if (E_t, I_t) is a bounded-degree state, then (E_{t+1}, I_{t+1}) is a bounded-degree state as well w.h.p. Let us name $X_t = |E_t(I_t)|$. The expected size of $E_0(I_0)$ is

$$\mathbf{E}\left[X_0\right] \leqslant \left[\binom{|I_0|}{2} + |I_0|(n - |I_0|)\right](1+\varepsilon)\tilde{p} \leqslant (1+\varepsilon)n\tilde{p}|I_0|\,.$$

Since edges are independent, $c \geqslant 8$, and $n\tilde{p}|I_0| = \Omega(\log n)$, from Chernoff bound (Lemma A.1 in Appendix A) it follows that $|E_0(I_0)| \leqslant cn\tilde{p}|I_0|$ w.h.p. Now let $t \geqslant 0$ and assume that $X_t \leqslant cn\tilde{p}|I_0|$. Observe that the size of $E_{t+1}(I_{t+1})$ satisfies

$$X_{t+1} \;=\; |E_{t+1}(I_t)| + |E_{t+1}(\hat{I}_{t+1}, [n] \setminus I_t)|\,, \tag{8}$$

where $\hat{I}_{t+1} := I_{t+1} \setminus I_t$. As for the first addend, we have that

$$\begin{aligned}\mathbf{E}\left[|E_{t+1}(I_t)| \mid X_t\right] &= (1-q)X_t + p\left[\binom{|I_t|}{2} + |I_t|(n-|I_t|) - X_t\right] \\ &= (1-(p+q))\,X_t + p\left[\binom{|I_t|}{2} + |I_t|(n-|I_t|)\right]\end{aligned}$$

because all the X_t edges existing at time t are still there at time $t+1$ with probability $1-q$ and all the edges that do not exist at time t appear with probability p. Since $p = \tilde{p}(p+q) \leqslant 2\tilde{p}$, if $p + q \geqslant 1$ then

$$\mathbf{E}\left[|E_{t+1}(I_t)|\right] \leqslant 2n\tilde{p}|I_t| \leqslant \frac{q}{4}cn\tilde{p}|I_t|\,,$$

regardless of the value of X_t. If instead $p + q \leqslant 1$ then, if $X_t \leqslant cn\tilde{p}|I_t|$ we have that

$$\begin{aligned}\mathbf{E}\left[|E_{t+1}(I_t)| \mid X_t \leqslant cn\tilde{p}|I_t|\right] &\leqslant (1-p-q)\,cn\tilde{p}|I_t| + np|I_t| \\ &= cn\tilde{p}|I_t|\left(1 - p - q + \frac{(p+q)}{c}\right) \\ &\leqslant \left(1 - \frac{q}{2}\right)cn\tilde{p}|I_t|\,,\end{aligned} \tag{9}$$

where in the last inequality we used that $p \geqslant 0$ and $(p+q)/c \leqslant q/2$.

As for the second addend, we observe that every pair $e = \{u, v\}$ with $u \in \hat{I}_{t+1}$, $v \in [n] \setminus I_t$, and $u \neq v$ exists in $E_{t+1}(\hat{I}_{t+1}, [n] \setminus I_t)$ with probability $p_e \in [(1-\varepsilon)\tilde{p},\, (1+\varepsilon)\tilde{p}]$ since it has never been observed before time $t+1$. Hence

$$\mathbf{E}\left[|E_{t+1}(\hat{I}_{t+1}, [n] \setminus I_t)|\right] \leqslant |\hat{I}_{t+1}|(n - |I_t|)(1+\varepsilon)\tilde{p} \leqslant \frac{q}{4}cn\tilde{p}|I_t|\,. \tag{10}$$

From (9) and (10) in (8) we get

$$\mathbf{E}\left[X_{t+1} \mid X_t \leqslant cn\tilde{p}|I_t|\right] \leqslant \left(1 - \frac{q}{4}\right)cn\tilde{p}|I_t| \leqslant \left(1 - \frac{q}{4}\right)cn\tilde{p}|I_{t+1}|\,.$$

Since edges are independent, $q = \Omega(1)$, and $n\tilde{p}|I_{t+1}| = \Omega(\log n)$, from Chernoff bound (Lemma A.1 in Appendix A) it follows that $X_{t+1} \leqslant cn\tilde{p}|I_{t+1}|$ w.h.p. □

Now we can bound the *increasing rate* of the number of informed nodes in an edge-Markovian graph. The proof of the following lemma combines the analysis adopted in the proof of Lemma 3.3 with some further ingredients required to manage the time-dependency of the edge-Markovian model.

Lemma 4.4 (The increasing rate of new informed nodes) *Let (E, I) be a bounded-degree state and let X be the random variable counting the number of non-informed nodes that get informed after two steps of the* `Push` *operation in the edge-Markovian graph model. It holds that $\mathbf{P}\left(X \geqslant \varepsilon \cdot \min\{|I|,\, n - |I|\}\right) \geqslant \lambda$, where ε and λ are positive constants.*

Proof. Let $G_0 = ([n], E_0)$ be the current graph and let $G_1 = ([n], E_1)$ and $G_2 = ([n], E_2)$ be the next two random graphs obtained according to the edge-Markovian process starting from G_0. Let $H = ([n], E_H)$ be such that $E_H = E_1 \cup E_2$ and let $\hat{H}$ be the $(I, 3cn\tilde{p})$-modified version of H according to Definition 3.1, where c is a sufficiently large constant (from what follows, it will be clear that it is sufficient to fix any $c \geqslant 32/q$). From Lemmas 3.2 and 3.4, we have that the number of informed nodes in $\hat{H}$ is stochastically smaller than the number of informed nodes in G_2. In what follows we evaluate the number of new informed nodes in $\hat{H}$ and we show that with positive constant probability it is at least a constant fraction of $\min\{|I|,\, n - |I|\}$.

Let I_A be the set of informed nodes that have degree at most $cn\tilde{p}$, i.e.,

$$I_A = \{u \in I \,:\, \deg_{G_0}(u) \leqslant cn\tilde{p}\}\,.$$

In what follows, I_A will denote the set of *active* informed nodes. Observe that

$$\sum_{u \in I} \deg_{G_0}(u) \leqslant 2|E(I)|.$$

Since (E, I) is a bounded-degree state, we have $2|E(I)| \leqslant (16/q)n\tilde{p}|I|$. Thus, if $c \geqslant 32/q$ then we have that $|I_A| \geqslant |I|/2$.

Consider an active informed node $u \in I_A$ and let $v \in [n] \setminus I$ be a non-informed one. The probability that node u selects node v in $\hat{H}$ according to the `Push` protocol is

$$\begin{aligned}\mathbf{P}\left(\delta_{\hat{H}}(u) = v\right) \;&=\; \mathbf{P}\left(\delta_{\hat{H}}(u) = v \mid \{u, v\} \in E_H,\, \deg_H(u) \leqslant 3cn\tilde{p}\right) \cdot \\ &\quad \cdot \mathbf{P}\left(\deg_H(u) \leqslant 3cn\tilde{p} \mid \{u, v\} \in E_H\right) \mathbf{P}\left(\{u, v\} \in E_H\right)\,. \qquad (11)\end{aligned}$$

Indeed, by the definition of $\hat{H}$, u cannot select v in $\hat{H}$ if the edge $\{u, v\}$ does not exist in H or if the degree of u in H is larger than $3cn\tilde{p}$.

Now observe that

$$\mathbf{P}\left(\delta_{\hat{H}}(u) = v \,|\, \{u, v\} \in E_H,\, \deg_H(u) \leqslant 3cn\tilde{p}\right) \geqslant 1/(6cn\tilde{p})\,. \qquad (12)$$

Indeed, node u has $3cn\tilde{p}$ virtual neighbors in $\hat{H}$ plus up to $3cn\tilde{p}$ non-informed neighbors. As for $\mathbf{P}\left(\{u, v\} \in E_H\right)$, from Observation B.1 (see Appendix B), it follows that

$$\mathbf{P}\left(\{u, v\} \in E_H\right) \geqslant p = \tilde{p}(p + q) \geqslant q \cdot \tilde{p}\,. \qquad (13)$$

We now show that $\mathbf{P}\left(\deg_H(u) \leqslant 3cn\tilde{p} \mid \{u, v\} \in E_H\right)$ is larger than a positive constant. Observe that we can write

$$\deg_H(u) = \sum_{w \in [n] \setminus \{u\}} X_w\,,$$

where X_w is the indicator random variable of the event $\{u, w\} \in E_H$. Thus,

$$\mathbf{E}\left[\deg_H(u) \,|\, \{u, v\} \in E_H\right] = \sum_{w \in [n] \setminus \{u\}} \mathbf{P}\left(X_w = 1 \,|\, \{u, v\} \in E_H\right) . \tag{14}$$

Now observe that, for $w \neq v$, $\mathbf{P}\left(X_w = 1 \,|\, \{u, v\} \in E_H\right) = \mathbf{P}\left(X_w = 1\right)$ and it can have two values, depending on whether or not edge $\{u, w\}$ existed in G_0,

$$\begin{aligned} \mathbf{P}\left(X_w = 1 \,|\, \{u, w\} \notin E_0\right) &= p + (1-p)p \,, \\ \mathbf{P}\left(X_w = 1 \,|\, \{u, w\} \in E_0\right) &= 1 - q + qp \,. \end{aligned}$$

Hence, if we split the sum in (14) in the w's that were neighbors of u in E_0 and those that were not, we get

$$\begin{aligned} \mathbf{E}\left[\deg_H(u) \,|\, \{u, v\} \in E_H\right] &\leqslant 1 + (1 - q + qp)\deg_{G_0}(u) + (n - \deg_{G_0}(u))(p + (1-p)p) \\ &\leqslant 1 + \deg_{G_0}(u) + (n - \deg_{G_0}(u))2p \\ &\leqslant cn\tilde{p} + 3np \\ &\leqslant 2cn\tilde{p} \,, \end{aligned}$$

where, from the first line to the second one we used that $p + (1-p)p \leqslant 2p$ and $1 - q + qp \leqslant 1$, from the second to the third line we used that $1 \leqslant np$ and that $\deg_{G_0}(u) \leqslant cn\tilde{p}$, because $u \in I_A$, and from the third line to the fourth one we used that $p = (p+q)\tilde{p} \leqslant 2\tilde{p}$ and $c \geqslant 6$. From Markov's inequality it thus follows that

$$\mathbf{P}\left(\deg_H(u) \geqslant 3n\tilde{p} \,|\, \{u, v\} \in E_H\right) \leqslant 2/3 \,. \tag{15}$$

By combining (12), (13), and (15) in (11) we get

$$\mathbf{P}\left(\delta_{\hat{H}}(u) = v\right) \geqslant \frac{\alpha}{n}$$

for a suitable positive constant α.

Since the events $\{\delta_{\hat{H}}(u) \neq v \,:\, u \in I_A\}$ are independent, the probability that node v is not informed in $\hat{H}$ is

$$\mathbf{P}\left(\bigcap_{u \in I_A} \delta_{\hat{H}}(u) \neq v\right) \leqslant (1 - \alpha/n)^{|I_A|} \leqslant e^{-\alpha |I_A|/n} \leqslant e^{-(\alpha/2)|I|/n} \,.$$

Let X be the random variable counting the number of new informed nodes in $\hat{H}$. The expectation of X is thus

$$\mathbf{E}\left[X\right] \geqslant (n - |I|)\left(1 - e^{-(\alpha/2)|I|/n}\right) \geqslant (\alpha/4)(n - |I|)|I|/n \,.$$

Hence we have that

$$\mathbf{E}\left[X\right] \geqslant \begin{cases} (\alpha/8)|I| & \text{if } |I| \leqslant n/2 \,, \\ (\alpha/8)(n - |I|) & \text{if } |I| \geqslant n/2 \,. \end{cases}$$

Since $X \leqslant \min\{|I|,\, n - |I|\}$ the thesis then follows from Observation B.2 (see Appendix B). □

Now we can prove that in $\mathcal{O}(\log n)$ time steps the `Push` protocol informs all nodes in an edge-Markovian graph, w.h.p.

Theorem 3 *Let $\mathcal{G} = \mathcal{G}(n, p, q, E_0)$ be an edge-Markovian graph with $p \geqslant 1/n$ and $q = \Omega(1)$ and let $s \in [n]$ be a node. The* `Push` *protocol with source s completes the broadcast over $\mathcal{G}$ in $\mathcal{O}(\log n)$ time steps w.h.p.*

Proof. Lemma 4.1 implies that after $\mathcal{O}(\log n)$ time steps there are $\Omega(\log n)$ informed nodes w.h.p. From Observation B.1 (see Appendix B) and Lemma 4.3, it follows that, after further $\mathcal{O}(\log n)$ time steps, the edge-Markovian graph reaches a bounded-degree state and remains so for further $\Omega(\log n)$ time steps. Let us rename $t = 0$ the time step where there are $\Omega(\log n)$ informed nodes and every edge $e \in \binom{[n]}{2}$ exists with probability $p_e \in [(1-\varepsilon)\tilde{p}\,,\ (1+\varepsilon)\tilde{p}]$. We again abbreviate $m_t := |I_t|$. Observe that if recurrence $m_{2(t+1)} \geqslant (1+\varepsilon)m_{2t}$ holds $\log n/\log(1+\varepsilon)$ times, then there are $n/2$ informed nodes. Let us thus name

$$T = \frac{2}{\lambda}\frac{\log n}{\log(1+\varepsilon)}$$

If at time $2T$ there are less than $n/2$ informed nodes, then recurrence $m_{2(t+1)} \geqslant (1+\varepsilon)m_{2t}$ held less than $\lambda T/2$ times. Since, at each time step, the recurrence holds with probability at least λ (there are less than $n/2$ informed nodes and the state is a bounded-degree one w.h.p.), the above probability is at most as large as the probability that in a sequence of T independent coin tosses, each one giving `head` with probability λ, we see less than $(\lambda/2)T$ `heads` (see, e.g., Lemma 3.1 in [2]). By using Chernoff bound (Lemma A.1 in Appendix A) such a probability is smaller than $e^{-\gamma\lambda T}$, for a suitable positive constant γ. Since γ and λ are constants and $T = \Theta(\log n)$ we have that

$$\mathbf{P}\left(m_{2T} \leqslant n/2\right) \leqslant n^{-\delta} \tag{16}$$

for a suitable positive constant δ. When m_t is larger than $n/2$ and the edge-Markovian graph is in a bounded-degree state, from Lemma 4.4 it follows that the recurrence

$$n - m_{t+1} \leqslant (1-\varepsilon)(n - m_t)$$

holds with probability at least λ. If this recurrence holds $\log n/\log\left(1/(1-\varepsilon)\right)$ times then the number of informed nodes cannot be smaller than n. Hence, if we name $\tilde{T} := (2/\lambda)\log n/\log\left(1/(1-\varepsilon)\right)$, thanks to the same argument we used to get (16), we obtain that, after $2T + 2\tilde{T}$ time steps, all nodes are informed w.h.p. □

5 Conclusions

In this paper we studied the `Push` protocol over edge-MEGs. We first analyzed the independent $G_{n,p}$ case (i.e. the edge-MEG with $q = 1 - p$) and we showed that the completion time is $\mathcal{O}(\log n/n\hat{p})$ w.h.p., where $\hat{p} = \min\{p,\, 1/n\}$. Then we studied the general edge-MEG model with $p \geqslant 1/n$ and $q = \Omega(1)$ and we showed that the completion time is logarithmic. This bound is obviously tight because the `Push` protocol cannot inform n nodes in less than $\log_2 n$ time steps.

Our results can be extended to the case of "more static" sparse dynamic graphs. Indeed, in [9, 27], Isopi and Panconesi show a logarithmic bound on the completion time of the `Push` protocol over the $\mathcal{G}(n, p, q)$ model when $p = \Theta(1/n)$ and $q = o(1)$.

We believe that the most challenging question is to analyze rumor spreading over more general classes of evolving graphs where edges may be not independent: for instance, it would be interesting to analyze the `Push` protocol over geometric models of mobile networks [12, 28].

References

[1] C. Avin, M. Koucky, and Z. Lotker. How to explore a fast-changing world. In *Proc. of 35th ICALP*, volume 5125 of *LNCS*, pages 121–132. Springer-Verlag, 2008.

[2] Y. Azar, A.Z. Broder, A.R. Karlin, and E. Upfal. Balanced allocations. *SIAM J. on Computing*, 29(1):180–200, 1999.

[3] H. Baumann, P. Crescenzi, and P. Fraigniaud. Parsimonious flooding in dynamic graphs. *Distributed Computing*, 24(1):31–44, 2011. (Ext. Abs. in *ACM PODC 2009*).

[4] B. Bollobás. *Random Graphs*. Cambridge University Press, 2001.

[5] S. Boyd, G. Arpita, P. Balaji, and S. Devavrat. Gossip algorithms: Design, analysis and applications. In *Proc. of 24th INFOCOM*, pages 1653–1664. IEEE, 2005.

[6] F. Chierichetti, S. Lattanzi, and A. Panconesi. Rumor Spreading in Social Networks. In *Proc. of 36th ICALP*, volume 5556 of *LNCS*, pages 375–386. Springer-Verlag, 2009.

[7] F. Chierichetti, S. Lattanzi, and A. Panconesi. Almost tight bounds on rumour spreading by conductance. In *Proc. of 42nd ACM STOC*, pages 399–408. ACM New York, 2010.

[8] F. Chierichetti, S. Lattanzi, and A. Panconesi. Rumour spreading and graph conductance. In *Proc. of 21th ACM-SIAM SODA*, pages 1657–1663. SIAM, 2010.

[9] A. Clementi, P. Crescenzi, C. Doerr, P. Fraigniaud, M. Isopi, A. Panconesi, F. Pasquale, and R. Silvestri. Rumor spreading in random evolving graphs. In *Proc. of 21st ESA*, volume 8125 of *LNCS*, pages 325–336. Springer, 2013.

[10] A. Clementi, C. Macci, A. Monti, F. Pasquale, and R. Silvestri. Flooding time of edge-markovian evolving graphs. *SIAM J. Discrete Math.*, 24(4):1694–1712, 2010. (Ext. Abs. in *ACM PODC 2008*).

[11] A. Clementi, A. Monti, F. Pasquale, and R. Silvestri. Broadcasting in dynamic radio networks. *J. Comput. Syst. Sci.*, 75(4):213–230, 2009. (Ext. Abs. in *ACM PODC 2007*).

[12] A. Clementi, A. Monti, F. Pasquale, and R. Silvestri. Information spreading in stationary markovian evolving graph. *IEEE Trans. Parallel Distrib. Syst.*, 22(9):1425–1432, 2011. (Ext. Abs. in *IEEE IPDPS 2009*).

[13] A. Demers, D. Greene, C. Hauser, W. Irish, J. Larson, S. Shenker, H. Sturgis, D. Swinehart, and D. Terry. Epidemic algorithms for replicated database maintenance. In *Proc. of 6th ACM PODC*, pages 1–12. ACM New York, 1987.

[14] B. Doerr, M. Fouz, and T. Friedrich. Social networks spread rumors in sublogarithmic time. In *Proc. of 43rd ACM STOC*, pages 21–30. ACM New York, 2011.

[15] P. Erdős and A. Rényi. On Random Graphs. *Publ. Math.*, 6:290–297, 1959.

[16] U. Feige, D. Peleg, P. Raghavan, and E. Upfal. Randomized broadcast in networks. *Random Structures and Algorithms*, 1(4):447–460, 1990.

[17] A. Ferreira. On models and algorithms for dynamic communication networks: The case for evolving graphs. In *Proc. of 4th ALGOTEL*, pages 155–161, 2002.

[18] N. Fountoulakis, A. Huber, and K. Panagiotou. Reliable broadcasting in random networks and the effect of density. In *Proc. of 29th IEEE INFOCOM*, pages 2552–2560. IEEE, 2010.

[19] N. Fountoulakis and K. Panagiotou. Rumor spreading on random regular graphs and expanders. In *Proc. of 13th APPROX/RANDOM*, volume 6302 of *Lecture Notes in Computer Science*, pages 560–573. Springer, 2010.

[20] N. Fountoulakis, K. Panagiotou, and T. Sauerwald. Ultra-fast rumor spreading in social networks. In *Proc. of 23rd ACM-SIAM SODA*, pages 1642–1660. SIAM, 2012.

[21] A. Frieze and G. Grimmett. The shortest-path problem for graphs with random arc-lengths. *Discrete Applied Mathematics*, 10(1):57–77, 1985.

[22] G. Giakkoupis. Tight bounds for rumor spreading in graphs of a given conductance. In *Proc. of 28th STACS*, volume 9 of *LIPIcs*, pages 57–68. Schloss Dagstuhl, 2011.

[23] G. Giakkoupis and T. Sauerwald. Rumor spreading and vertex expansion. In *Proc. of 23rd ACM-SIAM SODA*, pages 1623–1641. SIAM, 2012.

[24] E.N. Gilbert. Random graphs. *Annals of Mathematical Statistics*, 30(4):1141–1144, 1959.

[25] P. Grindrod and D.J. Higham. Evolving graphs: dynamical models, inverse problems and propagation. In *Proc. R. Soc. A, 466(2115)*, pages 753–770, 2010.

[26] M. Harchol-Balter, T. Leighton, and D. Lewin. Resource discovery in distributed networks. In *Proc. of 18th PODC*, pages 229–237. ACM New York, 1999.

[27] M. Isopi and A. Panconesi. Personal communication.

[28] P. Jacquet, B. Mans, and G. Rodolakis. Information Propagation Speed in Mobile and Delay Tolerant Networks. *IEEE Transaction on Information Theory*, 56:5001–5015, 2010.

[29] M. Jelasity, S. Voulgaris, R. Guerraoui, A.-M. Kermarrec, and M. van Steen. Gossip-based Peer Sampling. *ACM Trans. Comp. Syst.*, 25(3):Article 8, 2007.

[30] R. Karp, C. Schindelhauer, S. Shenker, and B. Vocking. Randomized rumor spreading. In *Proc. of 41st IEEE FOCS*, pages 565–574. IEEE, 2000.

[31] A.-M. Kermarrec and M. van Steen. Gossiping in distributed systems. *SIGOPS Oper. Syst. Rev.*, 41(5):2–7, 2007.

[32] F. Kuhn, N. Linch, and R. Oshman. Distributed Computation in Dynamic Networks. In *Proc. 42nd ACM STOC*, pages 513–522. ACM New York, 2010.

[33] F. Kuhn and R. Oshman. Dynamic networks: models and algorithms. *ACM SIGACT News*, 42(1):82–96, 2011.

[34] B. Pittel. On spreading a rumor. *SIAM Journal on Applied Mathematics*, 47(1):213–223, 1987.

[35] T. Sauerwald and A. Stauffer. Rumor spreading and vertex expansion on regular graphs. In *Proc. of 22nd ACM-SIAM SODA*, pages 462–475. SIAM, 2011.

[36] R. Van Renesse, Y. Minsky, and M. Hayden. A Gossip-Style Failure Detection Service. In *Proc. of Middleware*, pages 55–70, 1998.

[37] M. Vojnovic and A. Proutier. Hop limited flooding over dynamic networks. In *Proc. of 30th IEEE INFOCOM*, pages 685–693. IEEE, 2011.

[38] J. Whitbeck, V. Conan, and M.D. de Amorim. Performance of Opportunistic Epidemic Routing on Edge-Markovian Dynamic Graphs. *IEEE Transactions on Communications*, 59(5):1259–1263, 2011.

Appendix

A Useful inequalities

Lemma A.1 (Chernoff bound) *Let* $X = \sum_{i=1}^{n} X_i$ *where* X_i*'s are independent Bernoulli random variables and let* $\mu = \mathbf{E}[X]$*. Then for every* $0 < \delta < 1$ *it holds that*

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

$$\mathbf{P}(X \geqslant (1+\delta)\mu) \leqslant e^{-(\delta^2/3)\mu}$$

B Some observations

Observation B.1 *Consider the general two state Markov chain*

$$\left(\begin{array}{c|cc} & 0 & 1 \\ \hline 0 & 1-p & p \\ 1 & q & 1-q \end{array}\right)$$

Then

- *For every initial state* $x \in \{0,1\}$*, the probability that the chain is is state* 1 *in at least one of the first two time steps is*
$$\mathbf{P}(X_2 = 1 \text{ or } X_1 = 1 \,|\, X_0 = x) \geqslant p$$
- *Let* $p_t = \mathbf{P}(X_t = 1)$ *be the probability that the chain is in state* 1 *at time* t*. Then*
$$p_t = \frac{p}{p+q} + \left(p_0 - \frac{p}{p+q}\right)(1-p-q)^t$$

Observation B.2 *Let* X *be a random variable taking values between* 0 *and* m*, for some positive real* m*. If* $\mathbf{E}[X] \geqslant \lambda m$ *for some* $0 \leqslant \lambda \leqslant 1$*, then*

$$\mathbf{P}\left(X \geqslant \frac{\lambda}{2} m\right) \geqslant \lambda/2$$