

Inondation dans les réseaux dynamiques[‡]

Hervé Baumann¹, Pierluigi Crescenzi² et Pierre Fraigniaud³

¹ *University Paris Diderot. Email : herve.baumann@liafa.jussieu.fr*

² *University of Florence.*

³ *CNRS and Univ. Paris Diderot. Email : pierre.fraigniaud@liafa.jussieu.fr*

Cette note résume nos travaux sur l'inondation dans les réseaux dynamiques. Ces derniers sont définis à partir d'un processus Markovien de paramètres p et q générant des séquences de graphes (G_0, G_1, G_2, \dots) sur un même ensemble $[n]$ de sommets, et tels que G_t est obtenu à partir de G_{t-1} comme suit : si $e \notin E(G_{t-1})$ alors $e \in E(G_t)$ avec probabilité p , et si $e \in E(G_{t-1})$ alors $e \notin E(G_t)$ avec probabilité q . Clementi et al. (PODC 2008) ont analysé différents processus de diffusion de l'information dans de tels réseaux, et ont en particulier établi un ensemble de bornes sur les performances de l'inondation. L'inondation consiste en un protocole élémentaire où chaque nœud apprenant une information à un temps t la retransmet à tous ses voisins à toutes les étapes suivantes. Evidemment, en dépit de ses avantages en terme de simplicité et de robustesse, le protocole d'inondation souffre d'une utilisation abusive des ressources en bande passante. Dans cette note, nous montrons que l'inondation dans les réseaux dynamiques peut être mis en œuvre de façon à limiter le nombre de retransmissions d'une même information, tout en préservant les performances en termes du temps mis par une information pour atteindre tous les nœuds du réseau. La principale difficulté de notre étude réside dans les dépendances temporelles entre les connexions du réseau à différentes étapes de temps.

Keywords: Gossip protocol, epidemic protocol, evolving graphs, broadcasting

1 Introduction

Gossip protocols have been identified as simple, efficient, and robust mechanisms for various network and system tasks, such as, e.g., multicast, resource location, and distributed databases management. In parallel, the epidemiology community has considered several models for the analysis of the spreading of an infection in a population. These models include the famous SIR (Susceptible-Infective-Removed) and SIS (Susceptible-Infective-Susceptible) models, aiming at capturing the way a virus disseminates in a population as a function of the reaction to the virus experienced by every person. Gossip (and epidemic) protocols tolerate a high degree of dynamism in their running environment. It is therefore of premier importance to evaluate the precise impact of this dynamism on the efficiency of the gossip protocols. It is indeed known that network dynamics can have a tremendous impact in certain circumstances. This impact can be quite positive whenever the network evolution is ergodic (e.g., when measuring the global bandwidth of an ad hoc radio network [GT02]), but also quite negative whenever the network evolution is arbitrary (e.g., when measuring the cover time of random walks [AKL08]).

In the framework of gossip protocols and epidemiology, one way to handle dynamism is to assume that the network evolves with time as a sequence (G_0, G_1, G_2, \dots) of graphs on the same set of vertices, where the graph G_t considered at time t is an Erdős-Renyi random graph drawn in $\mathcal{G}_{n,p}$. In this model, several investigations have been recently performed to measure the impact of the network evolution on the performances of algorithms. For instance, it was proved that radio broadcasting performs efficiently, even for p below the connectivity threshold of $\mathcal{G}_{n,p}$ (see [CMPS07]). In [AKL08] it is proved that the cover time of a random walk remains polynomial for any $p > 0$. Threshold phenomena have also been identified ; for

[†] A part of this work was done during the stay of the second author at University Paris Diderot as Invited Professor. All authors were supported by COST Action 295 "DYNAMO". The first and third authors received additional supports from ANR project ALADDIN, and from INRIA project GANG. The second author received additional supports from the EU under the EU/IST Project 15964 AEOLUS.

[‡]The full version of this paper will appear in *Proc. 28th ACM Symp. on Principles of Distributed Computing (PODC 2009)*

instance, it was proved that SIR epidemic only contaminates a constant number of nodes if $p = c/n$ with $c < 1$, but contaminate a constant fraction of the nodes if $c > 1$ (see [DGM06]).

In term of modeling real world networks, the elementary random process (G_0, G_1, G_2, \dots) with $G_t \in \mathcal{G}_{n,p}$ suffers from the absence of time dependencies. The network does evolve but the structure of the network at time t is independent from its structure at time $t' < t$. This does not precisely reflect what is observed in many contexts such as wireless networks (the connection between two nodes is highly correlated to the previous existence of this connection) and P2P networks (the occurrence of an information exchange between two participants is highly correlated to the existence of previous exchanges). A more evolved model capturing time-dependencies has been recently considered in [CMM⁺08, CMPS09]. Similarly to the elementary random process, the model, called *edge-Markovian* process, and denoted $\mathcal{M}_{n,p,q}$, generates a random sequence of graphs (G_0, G_1, G_2, \dots) on the same node set $[n]$. This sequence is set based on a birth-rate p and death-rate q as follows[§]: G_0 is an Erdős-Renyi random graph in $\mathcal{G}_{n,\hat{p}}$ where $\hat{p} = p/(p+q)$, and, for any $t > 0$, a non-existing edge $e \notin E(G_{t-1})$ exists in $E(G_t)$ with probability p , while an existing edge $e \in E(G_{t-1})$ disappears from $E(G_t)$ with probability q .

In their companion papers [CMM⁺08, CMPS09], Clementi et al. analyzed the flooding protocol in edge-Markovian dynamic graphs, i.e., in graph sequences generated by the edge-Markovian process. Flooding in dynamic graphs is the gossip mechanism in which every node becoming aware of an information at step t forwards this information to all its neighbors at all forthcoming steps $t' > t$. Flooding is a core mechanism for information dissemination in contexts in which the network topology is highly dynamic and unknown, such as P2P networks, mobile networks, or any networks susceptible to faults, and several variants of flooding designed to limit the bandwidth consumption have been proposed [CL07, LCC⁺02]. Clementi et al. produced several bounds on the flooding time in edge-Markovian dynamic graphs. In particular they proved that there is a wide class of dynamic graphs in which the flooding time does not (asymptotically) depends on the edge death-rate q .

Despite the interests of flooding in term of simplicity and robustness, this protocol suffers from a severe drawback in dynamic graphs : it requires every node, upon reception of a source message, to forward this message during *all* forthcoming time steps. This results in a waste of resources in terms both of link bandwidth and node computation. Of course, if one knows that flooding completes in T time steps, then all nodes can be bounded to be active for only this amount of time. Nevertheless, T is typically growing with the size n of the network, hence been active T steps still results in a significant waste of resources. Our objective is to force flooding to perform more parsimoniously, by limiting the number of steps during which every node is active in forwarding a message to its neighbors, yet allowing the message to eventually reach all nodes in short time. Parsimonious flooding enables to save bandwidth and computational resources, and potentially energy as well, the latter parameter being known to be crucial for ad hoc and sensor networks.

Previous work.

There is a vast body of literature on broadcast and gossip protocols in static networks. Broadcasting in random graphs $\mathcal{G}_{n,p}$ has been analyzed in [FPRU90, FG85, Pit87], and randomized gossip protocols in specific metrics have been analyzed in [KKD01, KK02]. In all these cases, there are no time-dependencies induced by any evolution of the network structure.

Several papers tackle information spreading problems in the context of wireless networks (see, e.g., [CLF⁺07, SD09]) but the time-dependencies are either ignored, or overcome by assuming sufficiently long time slots for these dependencies to become negligible. In fact, the only information spreading work we are aware of, dealing explicitly with time-dependencies regarding the presence of the links, is [CMM⁺08, CMPS09]. In the former, the model is the one considered in this paper, i.e., edge-Markovian dynamic graphs. In the latter, the authors consider other types of Markovian dynamics, including the *geometric* Markovian evolving graphs, and give general bounds on the flooding time in Markovian dynamic graphs satisfying certain expansion properties.

More specifically, it is proved in [CMM⁺08] that, for any initial graph G_0 (i.e., not necessarily $G_0 \in \mathcal{G}_{n,\hat{p}}$), and any birth-rate and death-rate $0 < p, q < 1$, the flooding time in edge-Markovian graphs is at most

[§] The general definition in [CMM⁺08] assumes that G_0 can be arbitrary, and the definition of edge-Markovian evolving graphs given here is called *stationary* edge-Markovian evolving graphs in [CMPS09].

$O(\log n / \log(1 + np))$. For the stable graph G_0 (i.e., $E(G_0) = \emptyset$), it is proved that, for any $0 < p, q < 1$, the flooding time is at least $\Omega(\log n / np)$, and if $p \geq c \log n / n$ for $c > 1$ then the flooding time is at least $\Omega(\log n / \log(1 + np))$. The case $G_0 \in \mathcal{G}_{n, \hat{p}}$ is considered in [CMPS09], where it is proved that for $\hat{p} \geq c \log n / n$ with c large enough, the flooding time is at most $O(\frac{\log n}{\log n \hat{p}} + \log \log n \hat{p})$ and at least $\Omega(\frac{\log n}{\log n \hat{p}})$.

2 Parsimonious flooding

For a positive integer k , we say that a flooding protocol (in dynamic graphs) is k -active if each node forwards a source message only during the k time steps immediately following the step at which the node receives that message for the first time. For instance, the 1-active flooding protocol is the standard flooding protocol for static networks : a message is forwarded only once, at the step immediately following its reception. However, in dynamic networks, the flooding protocol may have to be active for $k > 1$ steps in order for the message to reach all nodes. The smaller the parameter k , the lesser the resource consumption by the protocol.

Our objective is to determine the minimum k for which the k -active flooding protocol completes correctly, i.e., the message eventually reaches all nodes. We define the *reachability threshold* for the flooding protocol in $\mathcal{M}_{n,p,q}$ as the smallest integer k such that, for any source $s \in [n]$, the k -active flooding protocol from s completes correctly almost surely[¶]. Clearly if flooding completes in T steps in $\mathcal{M}_{n,p,q}$, then the reachability threshold of flooding in $\mathcal{M}_{n,p,q}$ is at most T . But in fact, we will show that, for a large spectrum of parameters p and q , the reachability threshold is just $o(T)$, and often even just constant. Moreover, we will also show that being active for a number of steps equal to the reachability threshold (up to a multiplicative constant) is sufficient for the flooding protocol to complete in *optimal* time, i.e., in asymptotically the same number of steps as when being perpetually active.

3 Our results

In this paper, we first revisit the results in [CMM⁺08, CMPS09], and we give tight bounds on the flooding time (without any bound on the activity constraints) for all possible values of $p, q \in (0, 1)$ whenever $G_0 \in \mathcal{G}_{n, \hat{p}}$. These bounds are summarized in Table 1, where $\hat{p} = \frac{p}{p+q}$. For $\hat{p} \geq \frac{c \log n}{n}$ with $c > 1$, flooding performs, a.s., in $\Theta(\frac{\log n}{\log(n\hat{p})})$ rounds. For $0 < \hat{p} \leq \frac{c}{n}$ with $c > 0$, flooding a.s. performs in $\Theta(\frac{\log n}{np})$ rounds. In between, the situation is more complex, and depends on the relative values of \hat{p} and p (see Table 1). If $np \geq \log n \hat{p}$ then the flooding time is $\Theta(\frac{\log n}{\log(n\hat{p})})$, whereas if $np \leq \log n \hat{p}$ then the flooding time is $\Theta(\frac{\log n}{np})$.

In parallel to the computation of the bounds on the flooding time, we have established tight bounds on the reachability threshold. If $\hat{p} \geq \frac{c \log n}{n}$ with $c > 1$, then this parameter is equal to 1. That is, just one step of activity is enough to make sure that, a.s., the message reaches all nodes. If $\hat{p} \leq \frac{c \log n}{n}$ with $c < 1$, then the reachability threshold is $\Theta(\frac{\log n}{np})$. (Note that the condition is on \hat{p} while the value for the reachability threshold depends on p).

Interestingly enough, for any $p, q \in (0, 1)$, we also prove that if k is the activity threshold, then, a.s., an $O(k)$ -active flooding protocol completes in the same time as flooding without constraints on the activity, up to a multiplicative constant. In other words, the reachability threshold for the flooding protocol is essentially of the same order of magnitude as the activity threshold for this protocol to complete in optimal time, up to multiplicative constants. In particular, for $\hat{p} \geq \frac{c \log n}{n}$ with $c > 1$, one step of activity is sufficient for flooding to complete in optimal time. Similarly, for $\frac{1}{n} \ll \hat{p} \leq \frac{c \log n}{n}$ with $c < 1$, and $np \geq \log n \hat{p}$, the reachability threshold $\Theta(\frac{\log n}{np})$ is significantly smaller than the optimal flooding time $\Theta(\frac{\log n}{\log(n\hat{p})})$, yet being active for $O(\frac{\log n}{np})$ steps is sufficient for flooding to complete in asymptotically optimal time $O(\frac{\log n}{\log(n\hat{p})})$. For all the remaining cases, the thresholds for completion and for optimality coincide, up to multiplicative constants.

[¶] A series of events \mathcal{E}_n holds almost surely (a.s.) if $\Pr[\mathcal{E}_n] \rightarrow 1$ when $n \rightarrow \infty$, i.e., $\Pr[\mathcal{E}_n] = 1 - o(1)$. These events holds with high probability (w.h.p.) if $\Pr[\mathcal{E}_n] \geq 1 - O(\frac{1}{n^\alpha})$ for some $\alpha > 0$.

	$0 < \hat{p} \leq \frac{c}{n}, c > 0$	$\frac{1}{n} \ll \hat{p} \leq \frac{c \log n}{n}, c < 1$ $np \leq \log n \hat{p}$ $np \geq \log n \hat{p}$	$\hat{p} \geq \frac{c \log n}{n}, c > 1$
Flooding time	$\Theta\left(\frac{\log n}{np}\right)$	$\Theta\left(\frac{\log n}{np}\right)$ $\Theta\left(\frac{\log n}{\log(n\hat{p})}\right)$	$\Theta\left(\frac{\log n}{\log(n\hat{p})}\right)$
Reachability threshold	$\Theta\left(\frac{\log n}{np}\right)$	$\Theta\left(\frac{\log n}{np}\right)$ $\Theta\left(\frac{\log n}{np}\right)$	1

TABLE 1: Summary of results

Références

- [AKL08] C. Avin, M. Koucký, and Z. Lotker. How to Explore a Fast-Changing World. In 35th International Colloquium on Automata, Languages and Programming (ICALP), pages 121–132, 2008.
- [CL07] N. Chang and M. Liu. Controlled flooding search in a large network. *IEEE/ACM Transactions on Networking* 15(2) :436–449, 2007.
- [CLF⁺07] D. Chakrabarti, J. Leskovec, C. Faloutsos, S. Madden, C. Guestrin, and M. Faloutsos. Information survival threshold in sensor and P2P networks. In 26th IEEE INFOCOM, pages 1316–1324, 2007.
- [CMM⁺08] A. Clementi, C. Macci, A. Monti, F. Pasquale, and R. Silvestri. Flooding Time in Edge-Markovian Dynamic Graphs. In 27th ACM Symp. on Principles of Dist. Comp. (PODC), pages 213–222, 2008.
- [CMPS07] A. Clementi, A. Monti, F. Pasquale, and R. Silvestri. Communication in Dynamic Radio Networks. In 26th ACM Symp. on Principles of Distributed Computing (PODC), pages 205–214, 2007.
- [CMPS09] A. Clementi, A. Monti, F. Pasquale, and R. Silvestri. Information Spreading in Stationary Markovian Evolving Graphs. In 23rd IEEE Int. Parallel and Dist. Processing Symp. (IPDPS), to appear.
- [DGM06] M. Draief, A. Ganesh, and L. Massoulié. Thresholds for virus spread on networks. *Ann. Appl. Probab.* 18(2) :359–378, 2008.
- [FG85] A. Frieze and G. Grimmett. The shortest-path problem for graphs with random arc-lengths. *Discrete Applied Math.* 10 :57–77, 1985.
- [FPRU90] U. Feige, D. Peleg, P. Raghavan, and E. Upfal. Randomized broadcast in networks. In Int. Symposium on Algorithms (SIGAL), LNCS 450, pages 128–137, Springer, 1990.
- [GT02] M. Grossglauser and D. Tse. Mobility increases the capacity of ad-hoc wireless networks. *IEEE/ACM Transactions on Networking* 10(4) :477–486 (2002).
- [KK02] D. Kempe and J. Kleinberg. Protocols and impossibility results for gossip-based communication mechanisms. In 43st IEEE Symp. on Foundations of Comp. Sc. (FOCS), pages 471–480, 2002.
- [KKD01] D. Kempe, J. Kleinberg, and A. Demers. Spatial gossip and resource location protocols. In 33rd ACM Symposium on Theory of Computing (STOC), pages 163–172, 2001.
- [LCC⁺02] Q. Lv, P. Cao, E. Cohen, K. Li, S. Shenker. Search and replication in unstructured peer-to-peer networks. In 16th ACM Int. Conference on Supercomputing (ICS), pages 84–95, 2002.
- [Pit87] B. Pittel. On spreading a rumour. *SIAM J. Applied Math.*, 47 :213–223, 1987.
- [SD09] A. Sarwate and A. Dimakis. The Impact of Mobility on Gossip Algorithms. In 28th Conference on Computer Communications (INFOCOM), 2009.