On the Speed of Information Spreading in Dynamic Networks

Francesco Pasquale

December 15, 2008

Abstract

There is a growing interest in the study of graphs that evolve over time. Communication networks, peer-to-peer systems, social networks, ad-hoc radio networks and the Internet are only few examples of networks that are intrinsically *dynamic*.

In order to analyse real world phenomena from a mathematical point of view, we have to define abstract models of concrete scenarios. This introduces a natural trade-off: the more a model is *realistic*, the more it is hard to analyse. Finding the right balance between these two poles is one of the hardest part of the work.

In the thesis we study the speed of information spreading in dynamic networks. Our aim is to introduce a general framework where simple communication primitives can be analysed, to achieve a better understanding of the main features of dynamic networks. The thesis provides a revised version of the results presented in [6, 5, 7]. We here give an informal overall description of such results.

One of the simplest communication tasks that seems worth studying in dynamic networks is the *broadcast* problem: one distinguished node of the network (the *source* node) aims to send a message to all the other nodes of the network.

We model a dynamic network as an *evolving graph*. An evolving graph is a sequence of graphs (i.e. the *snapshots*) with the same set of nodes. To analyse the speed of information spreading in evolving graphs, we use a communication process that is usually called *flooding* (or flooding mechanism). The flooding process can be described as follows: start with one single *informed* node (the source), and assume that when a non-informed node v has an informed neighbour then node v becomes informed itself, at the next time step. The question is how long does it take to get all the nodes informed?

Given an evolving graph and a source node, the *completion time* of the flooding process is the first time step in which all the nodes are informed. The *flooding time* of an evolving graph is the maximum completion time over all possible sources.

The flooding time of a *static* graph equals the diameter of the graph; in particular it is finite if and only if the graph is connected. Can we say something like that for evolving graphs? Is it true that, for example, if every snapshot of an evolving graph $\mathcal{G}$ has *small* diameter then the flooding time of $\mathcal{G}$ is *small*? In order to investigate such issue in a clean way, let us introduce the concept of *adversarial* evolving graph.

Adversarial evolving graphs. How can a graph evolve in a *worst-case* way, with respect to a communication task? Suppose to have a set of nodes V and assume there is an *adversary* that, time step by time step, chooses the set of edges E_t, yielding an evolving graph $\mathcal{G} = \{G_t : t \in \mathbb{N}\}$ where $G_t = (V, E_t)$ is the *snapshot* of the evolving graph at time step t. Fix some constraints for the adversary, for example: all the snapshots G_t must be connected and must have *small* diameter. What is the flooding time of such an adversarial evolving graph? For how long can the adversary slow down the flooding process?

On one hand, if all the snapshots are connected then at every time step there is at least one new informed node, so the flooding time is upper bounded by the number of nodes. On the other hand, it is easy to define an adversarial strategy such that at every time step the graph has *small* diameter, but the flooding time is still *big*. Hence concepts like *connectivity* and *diameter*, that are strongly related to communication tasks in static graphs, should be revised in evolving graphs. In a static graph, the spread of information is faster in graphs with small diameter, whereas if the graph is not connected, then there is no way to communicate. Both those observations turns out to be false in evolving graphs.

Random evolving graphs. Graphs that evolve randomly are not a new subject in the networking area. For instance, a lot of effort has been devoted to define random mobility models that well approximate realistic mobile networks [3]. Typically, the network-evolution model is suitably chosen according to the particular application we are studying. Most of these models are very suitable for analysing concrete communication problems, performing accurate and meaningful simulations; on the other hand, their are often too complicated to allow a general and deep mathematical analysis of basic communication tasks. In the thesis, we adopt a different approach: starting from a random graph model that is well established in the static case, we try to extend it to the dynamic case.

One of the most elegant and widely studied model of random graphs is the Erdös-Rényi model $G_{n,p}$. Informally speaking, it can be described as follows: given a set of n nodes, every edge between two nodes exists with probability p, independently from the other edges. The properties of random graphs $G_{n,p}$ have been extensively studied [2]. In order to analyse communication problems on random evolving graphs, our first purpose is to define a natural evolving version of the $G_{n,p}$ model.

We introduce the following model of random evolving graph: start with an arbitrary initial graph, and at every time step, if an edge exists then it will *die* at next time step with probability q; while, if an edge does not exist, then it will *appear* at next time step with probability p. We call this model *edge-Markovian evolving graph* (edge-MEG, in short). It is easy to see that, whatever the initial graph is, an edge-MEG eventually converge to an Erdös-Rényi random graph.

We study the flooding time of edge-Markovian evolving graphs. An edge-MEG is defined by four parameters: the number n of nodes, the edge *birth-rate* p, *death-rate* q, and the initial graph E_0. The flooding time of edge-MEGs will depend on all those four parameters, but exploiting the dependence on the initial graph is quite a hopeless task. So we look for an upper bound of the following form: for every initial graph, the flooding time of edge-MEG $\mathcal{G}(n, p, q, E_0)$ is $\mathcal{O}(f)$ for some function $f = f(n, p, q)$. The function f that we find does not depend on the death-rate q. In order to understand how far away from the actual flooding time our upper bound is, we look for a lower bound independent on the death rate of this form: there exists an initial graph such that for every death rate q the flooding time is $\Omega(g)$ for some function $g = g(n,p)$. Quite surprisingly, we find function g that asymptotically matches our upper bound f, except when the birth-rate p is in a small range I_n. It turns out that, when p is outside I_n, the influence of the death-rate q on the flooding time is only for a constant factor that is hidden in the asymptotic formula; when p is inside I_n, the role of the death-rate q is no longer asymptotically negligible. We don't know the actual asymptotic formula for the flooding time when p is in that range, but we can prove that it *must* depend on the death-rate q too.

The stationary case. In the previous paragraph we have seen that there is an explicit, almost tight, asymptotic formula for the flooding time of edge-MEGs with respect to the *worst-case* initial graph. One natural question is: what happens if the initial graph is random according to the stationary distribution of the Markov chain that defines the process? For such an initial condition, all the snapshots of the evolving graph have the same marginal distribution. A random process with this property is called *stationary*. In experimental papers, this situation is often referred to as *perfect simulation*.

By studying stationary edge-MEGs, we realized that our techniques can be used to give upper bounds on the flooding time of a very general class of models called Markovian evolving graphs [1]. Roughly, a Markovian evolving graph (MEG, in short) is a Markov chain with a finite set of graphs as state space. Clearly, an edge-MEG is a special case of MEG. However, Markovian evolving graphs are much more general than this. Consider, for example, the following model that we here call *geometric-MEG*. Assume to have a two-dimensional grid and a set of n nodes, each node performing a random walk on the grid. Consider the random evolving graph where there is an edge between two nodes if their Euclidean distance is less than some fixed parameter r. Clearly, a geometric-MEG is a special case of MEG as well. Moreover it is a sort of prototype for a lot of similar *mobility* models where the *physical* random node movements over some metric space determine the evolution of the *overlay* random graph.

We prove a general theorem that establishes an upper bound on the flooding time of stationary Markovian evolving graphs in terms of their *expansion* properties. We have seen that the diameter of the snapshots of an evolving graph does not mean so much for the flooding time. Instead, expansion properties of the snapshots *do* mean. Our theorem says that if the snapshots are *good* expanders, the information spreading in the evolving graph is *fast*. More importantly, the theorem gives an explicit asymptotic formula that bounds the flooding time as a function of the expansion parameters of the evolving graph. We then show how to use the general theorem to obtain upper bounds on the flooding time of stationary edge-MEGs and geometric MEGs. In both cases the resulting upper bounds are nearly tight.

Radio broadcasting in dynamic networks. The radio network model [4] received great attention during the last two decades. Informally speaking, a radio network is modeled as a graph with the following communication constraints: 1. *Non directionality*: when a node sends a message, the message is sent along all its outgoing edges; 2. *Collisions*: if two or more neighbors of a node u send a message in the same time step, then a collision occurs and node u doesn't receive any message.

It is reasonable to claim that almost all major issues about distributed broadcast in radio networks have been solved when the network topology does not change. However, theoretical results for *dynamic* radio networks are still very few.

In order to understand the difficulties introduced by collisions, let us consider again adversarial evolving graphs. Given the set of nodes V, at every time step an adversary chooses the sets of edges E_t thus yielding the evolving graph $\mathcal{G} = \{G_t = (V, E_t) : t \in \mathbb{N}\}$. We impose that, for example, every snapshot G_t must be connected. Can we define a protocol that performs the broadcast task against such an adversary in the radio model? If we consider only deterministic protocols, then it's easy to see that we can't. Interestingly enough, we prove that there are randomized protocols that not only complete the broadcasting against any meaningful adversary, but have also *good* performances.

Finally, what can we say about radio broadcasting in random evolving graphs? Unfortunately, we are still far away from a complete understanding of this problem in full generality. However, we have some interesting results that highlight most of the important features. Here is a little taste. In [8] the authors study the radio broadcast problem in static $G_{n,p}$. They prove that if p is over the connectivity threshold, then there is a protocol that completes the broadcasting in $\mathcal{O}(\log n)$ time steps. We consider the random evolving graph where the snapshots are independent Erdös-Rényi random graphs. We prove that, for this *dynamic* $G_{n,p}$, there is a protocol that completes the broadcasting in $\mathcal{O}(\log n)$ time steps even if p is *under* the connectivity threshold.

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, 2008.

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

[3] T. Camp, J. Boleng, and V. Davies. A survey of mobility models for ad hoc network research. *Wireless Communication and Mobile Computing*, 2(5):483–502, 2002.

[4] I. Chlamtac and S. Kutten. On broadcasting in radio networks - problem analysis and protocol design. *IEEE Transactions on Communications*, 33(12):1240–1246, 1985.

[5] A. Clementi, C. Macci, A. Monti, F. Pasquale, and R. Silvestri. Flooding time in edge-markovian dynamic graphs. In *Proc. of 27th PODC*, pages 213–222. ACM, 2008.

[6] A. Clementi, A. Monti, F. Pasquale, and R. Silvestri. Communication in dynamic radio networks. In *Proc. of 26th PODC*, pages 205–214. ACM, 2007.

[7] A. Clementi, A. Monti, F. Pasquale, and R. Silvestri. Information spreading in stationary markovian evolving graphs. In *Proc. of 23rd IPDPS*. IEEE, 2009. To appear.

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