Journal of Computer and System Sciences 81 (2015) 219–233

Contents lists available at ScienceDirect

Journal of Computer and System Sciences

www.elsevier.com/locate/jcss

Parsimonious flooding in geometric random-walks ☆

Andrea Clementi [a,*], Riccardo Silvestri [b]

[a] Dipartimento di Ingegneria dell'Impresa, Università Tor Vergata di Roma, Italy
[b] Dipartimento di Informatica, Sapienza Università di Roma, Italy

A R T I C L E   I N F O

Article history:
Received 3 June 2014
Accepted 5 June 2014
Available online 18 June 2014

Keywords:
Information/virus spreading
Broadcasting protocols
Geometric random-walks
Mobile ad-hoc networks

A B S T R A C T

We study the epidemic process yielded by the k-Flooding Protocol in geometric Mobile Ad-Hoc Networks. We consider n agents on a square performing independent random walks. At any time step, every active agent informs every non-informed agent which is within distance R from it. An informed agent is active only for k time steps. Initially, a source agent is informed and we look at the completion time of the protocol. We prove optimal bounds on the completion time of the process. Our method of analysis provides a clear picture of the geometric shape of the information spreading over the time.

1. Introduction

In the *Geometric Random-Walk Model* [7,10,12], n agents perform independent random walks over a square of side length L. We consider the following *epidemic* process performed by the agents. Every agent can be in three different states: non-informed (white), informed–active (red), informed-removed (black). During a time step, every agent performs one step of the random walk and every red agent informs (infects) all white agents lying within distance R. A white agent, informed (for the first time) at time step t, becomes red at time step $t+1$. When an agent becomes red, he stays that for the next k time steps and, then, he becomes black and stays that forever.

At the initial time step, a source agent is in the red state. The completion time of the above infection process is the first time step in which every agent gets into the black state. If this time step does not exist then we say that the infection process does not complete.

Such a random process is inspired by two main scenarios: The *Susceptible–Infected–Removed (SIR)* process which is widely studied in Mathematical Epidemiology [1,2,4,14,20] and the *(Parsimonious) k-Flooding Protocol* [3] in opportunistic networks such as *geometric Mobile Ad-Hoc Networks (MANET)*.

While the standard flooding is inefficient in terms of agent's energy consumption and message complexity, the k-flooding protocol (for small values of k) strongly reduces the agent's energy consumption and the overall message complexity. However, as we will discuss later, the infection process yielded by the k-flooding (for small k) over a dynamic network is much more complex than that yielded by the standard flooding.

The k-flooding has been studied in [3] on *Edge-Markovian Evolving Graphs* (*Edge-MEG*). An Edge-MEG [6,7] is a Markovian random process that generates an infinite sequence of graphs over the same set of n agents. If an edge exists at time t then, at time $t+1$, it dies with probability q. If instead the edge does not exist at time t, then it will come into existence at

☆ An extended abstract of this work has been presented at *25th International Symposium*, *DISC 2011*, Rome, Italy, September 2011. Research partially supported by the Italian MIUR 2010N5K7EB_002 under the PRIN-COFIN 2010-11 Project *ARS-TechnoMedia*.

* Corresponding author.
E-mail addresses: clementi@mat.uniroma2.it (A. Clementi), silvestri@di.uniroma1.it (R. Silvestri).

http://dx.doi.org/10.1016/j.jcss.2014.06.002
0022-0000/

time $t+1$ with probability p. The stationary distribution of an Edge-MEG with parameter p and q is the classic Erdös–Rényi random graph $\mathcal{G}(n, \tilde{p})$ where $\tilde{p} = \frac{p}{p+q}$. The work [3] gives tight bounds on the k-flooding on stationary Edge-MEG for arbitrary values of p, q and k. In particular, it derives the *reachability threshold* for the k-flooding, i.e., the smallest $k = k(n, p, q)$ over which the protocol completes.

Edge-MEG is an analytical model of dynamic networks capturing time-dependencies, an important feature observed in real scenarios such as (faulty) wireless networks and P2P networks. However, it does not model important features of MANET. Indeed, in Edge-MEG, edges are independent Markov chains while, in several real opportunistic networks, the correlation among edges cannot be neglected: agents act over a geometric space [7,10,12,16].

1.1. Our contribution

We study the k-flooding protocol in the geometric random-walk model (the resulting MANET will be called *geometric-MANET*). The move radius ρ determines the maximal distance an agent can travel in one time step. Even though network modeling issues are out of the aims of this work, both the transmission radius and the move radius play a crucial role in our analysis and, thus, a short discussion about them is needed. Both parameters depend on several factors. In mathematical epidemiology, they depend on the agent's mobility, on the nature of infection and on the agent's social behavior that all together determine the average rate of "positive" contacts. For instance, in some particular human viruses (such as HIV), the nature of contacts required for getting an effective infection yields a very low average infection-rate: this can be modeled by setting the move radius to be significantly larger than the transmission radius. In MANET, the move radius depends (besides on the agent's mobility) on the adopted protocol that tunes the transmission rate of the agents: the larger is the time between two consecutive transmissions the larger is ρ. A larger ρ could yield a better message complexity at the cost of a larger completion time (as for the standard flooding, the correct trade-offs are derived in [8]).

It turns out that the issue of setting the "most suitable" move radius concerns both the network modeling and the protocol design. For these reasons, we investigate the k-flooding for a wide range of parameters R and ρ.

If the transmission radius R is below the *connectivity threshold* (this threshold is here computed with respect to the uniform distribution of n (static) agents over the square), we prove that, *with high probability (w.h.p.)*,[1] for any $\rho \geqslant 0$ and for any $k = O(1)$, the k-flooding does not complete.

We thus study the k-flooding when R is above the connectivity threshold, emphasizing that the "static" connectivity among (generic) agents says nothing about the dynamic *expansion*-properties yielded by the red agents over the time: this is the crucial issue in the k-flooding process.

For any $\rho \lesssim R$ and for any $k \geqslant 1$, we prove that the information spreads at "optimal" speed $\Theta(R)$, i.e., the k-flooding protocol w.h.p. completes within $O(L/R)$ time steps. Observe that, since $\rho \lesssim R$, this bound is asymptotically optimal.

Then, we consider a slightly-different version of the random-walk model with move radius ρ that can be up to any polynomial of R, i.e. $\rho \leqslant \text{POLY}(R)$. We prove that the information spreads at "optimal" speed $\Theta(\rho)$. So, for any $k \geqslant 1$, the k-flooding w.h.p. completes in time $O(L/\rho)$ which is optimal for any $\rho \geqslant R$. Notice that this optimal information speed makes the 1-flooding time smaller than the static diameter $O(L/R)$ of the stationary graph: our bound is thus the first analytical evidence that agent's mobility actually speeds-up this infection process. Finally, we observe that, in both cases, the energy-efficient 1-flooding protocol is as fast as the standard flooding [7,8].

1.2. Adopted techniques and further results

The presence of black agents in the infection process makes the previous analysis techniques – those adopted for the flooding – almost useless. In particular, percolation theory [12,10,18], meeting and cover time of random walks on graphs [19], and the expansion/bootstrap arguments [7,8] strongly rely on the fact that an informed agent will be active for all the flooding process. Furthermore, the analysis of k-flooding over the Edge-MEG model [3] strongly relies on the stochastic independence among the edges and the consequent high node-expansion property of $\mathcal{G}(n, p)$: properties that clearly do not hold in geometric-MANET.

Our method of analysis significantly departs from all those mentioned above. Besides the optimal bounds on the completion time, our analysis provides a clear characterization of the geometric evolution of the infection process. We make use of a grid partition of the square into cells of size $\Theta(R)$ and define a set of possible states a cell can assume over the time depending on the number of red, white and black agents inside it.[2] We then derive the local state-evolution law of any cell. Thanks to the regularity of this law, we can characterize the evolution of the geometric *wave* formed by the *red cells* (i.e. cells containing some red agent). A crucial property we prove is that, at any time step, white cells (i.e. cells containing white agents only) will never be adjacent to black cells, so there is always a red wave "working" between the black region and the white one. Furthermore, we show that the red wave eventually spans the entire region before all agents become black.

[1] As usual, an event is said to hold *with high probability (w.h.p.)*, if its probability is at least $1 - 1/n^{\Theta(1)}$.

[2] When $\rho \gg R$, we also need a further grid partition into *supercells* of size $\Theta(\rho)$ and a more complex argument.