

Tight Bounds on Information Dissemination in Sparse Mobile Networks*

Alberto Pettarin
Department of Information
Engineering
University of Padova
Padova, 35131, Italy
pettarin@dei.unipd.it

Andrea Pietracaprina
Department of Information
Engineering
University of Padova
Padova, 35131, Italy
capri@dei.unipd.it

Geppino Pucci
Department of Information
Engineering
University of Padova
Padova, 35131, Italy
geppo@dei.unipd.it

Eli Upfal
Department of Computer
Science
Brown University
Providence, RI 02912, USA
eli@cs.brown.edu

ABSTRACT

Motivated by the growing interest in mobile systems, we study the dynamics of information dissemination between agents moving independently on a plane. Formally, we consider k mobile agents performing independent random walks on an n -node grid. At time 0, each agent is located at a random node of the grid and one agent has a rumor. The spread of the rumor is governed by a dynamic communication graph process $\{G_t(r) \mid t \geq 0\}$, where two agents are connected by an edge in $G_t(r)$ iff their distance at time t is within their transmission radius r . Modeling the physical reality that the speed of radio transmission is much faster than the motion of the agents, we assume that the rumor can travel throughout a connected component of G_t before the graph is altered by the motion. We study the *broadcast time* T_B of the system, which is the time it takes for all agents to know the rumor. We focus on the sparse case (below the percolation point $r_c \approx \sqrt{n/k}$) where, with high probability, no connected component in G_t has more than a logarithmic number of agents and the broadcast time is dominated by the time it takes for many independent random walks to meet one other. Quite surprisingly, we show that for a system below the percolation point, the broadcast time does not depend on the transmission radius. In fact,

*Support for the first three authors was provided, in part, by MIUR of Italy under project AlgoDEEP, and by the University of Padova under the Strategic Project STPD08JA32 and Project CPDA099949/09. This work was done while the first author was visiting the Department of Computer Science of Brown University, partially supported by “Fondazione Ing. Aldo Gini”, Padova, Italy.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

PODC'11, June 6–8, 2011, San Jose, California, USA.
Copyright 2011 ACM 978-1-4503-0719-2/11/06 ...\$10.00.

we prove that $T_B = \tilde{\Theta}(n/\sqrt{k})$ for any $0 \leq r < r_c$, even when the transmission range is significantly larger than the mobility range in one step, giving a tight characterization up to logarithmic factors. Our result complements a recent result of Peres et al. (SODA 2011) who showed that above the percolation point the broadcast time is polylogarithmic in k .

Categories and Subject Descriptors

G.3 [Probability and Statistics]: Probabilistic algorithms;
F.2.m [Analysis of Algorithms and Problem Complexity]: Miscellaneous; G.2.2 [Discrete Mathematics]: Graph Theory—*network problems*

General Terms

Algorithms, Theory

Keywords

Information Dissemination, Mobile Networks, Multiple Random Walks, Broadcast

1. INTRODUCTION

The emergence of mobile computing devices has added a new intriguing component, *mobility*, to the study of distributed systems. In fully mobile systems, such as wireless mobile ad-hoc networks (MANETs), information is generated, transmitted and consumed within the mobile nodes, and communication is carried out without the support of static structures such as cell towers. These systems have been implemented in vehicular networks and sensor networks attached to soldiers on a battlefield or animals in a nature reserve [23, 14, 17, 26]. Characterizing the power and limitations of mobile networks requires new models and analytical tools that address the unique properties of these systems [15, 8], which include:

- *Small transmission radius*: the transmission range of individual agents is restricted by limitations on energy consumption and interference from other agents;

- *Planarity*: agents reside, move and transmit on (subsets of) a plane. Low diameter graphs that are often used to model static communication networks are not useful here;
- *Dynamic communication graphs*: communication channels between agents are changing dynamically as mobile agents move in and out of the transmission radius of other agents;
- *Relative speeds*: transmission speed is significantly faster than the physical movement of the agents. A message can execute several hops before the network is altered by motion.

In this work we study the dynamics of information dissemination between agents moving independently on a plane. We consider a system of k mobile agents performing independent random walks on an n -node grid, starting at time 0 in a uniform distribution over the grid nodes. We focus on the fundamental communication primitive of broadcasting a rumor originating at one arbitrary agent to all other agents in the system. We characterize the *broadcast time* T_B of the system, which is the time it takes for all agents to receive the rumor.

We model the spreading of information in a mobile system by a dynamic communication graph process $\{G_t(r) \mid t \geq 0\}$, where the nodes of $G_t(r)$ are the mobile agents, and two agents are connected by an edge iff their distance at time t is within their transmission radius r . We are interested in *sparse systems* in which the transmission radius is below the percolation point $r_c \approx \sqrt{n/k}$ (i.e., the minimum radius which guarantees that $G_t(r_c)$ features a giant connected component), and where, with high probability, no connected component of $G_t(r_c)$ has more than a logarithmic number of agents [24, 25]. The broadcast time in sparse systems is dominated by the time it takes for many independent random walks to meet one another. Incorporating the fact that radio transmission is much faster than the motion of the agents, we assume that the rumor can travel throughout a connected component of G_t within one step, before the graph is altered by the motion.

Our main result is quite surprising: we show that below the percolation point the broadcast time does not depend on the transmission radius. We prove that $T_B = \tilde{\Theta}(n/\sqrt{k})$ for any r below r_c , giving a tight characterization up to logarithmic factors¹. Our bound holds both when the transmission radius is significantly larger than the mobility range (i.e., the distance an agent can travel in one step), and when, in contrast to previous work [7, 8], the transmission radius as well as the mobility range are very small. Our work complements a recent result by Peres et al. [25] who proved an upper bound polylogarithmic in k for the broadcast time in a system of k mobile agents which follow independent Brownian motions in \mathbb{R}^d , with transmission radius above the percolation point.

Our analysis techniques are applicable to a number of interesting related problems such as covering the grid with many random walks and bounding the extinction time in random predator-prey systems.

¹The tilde notation hides polylogarithmic factors, e.g. $\tilde{O}(f(n)) = O(f(n) \log^c n)$ for some constant c .

1.1 Related Work

Information dissemination has been extensively studied in the literature under a variety of scenarios and objectives. Here we restrict our attention to the results more directly related to our work.

A prolific line of research has addressed broadcasting and gossiping in static graphs, where the nodes of the graph represent active entities which exchange messages along incident edges according to specific protocols (e.g., *push*, *pull*, *push-pull*). The most recent results in this area relate the performance of the protocols to expansion properties of the underlying topology, with particular attention to the case of social networks, where broadcasting is often referred to as *rumor spreading* [6]. (For a relatively recent, comprehensive survey on this subject, see [16].) Unfortunately, mobile networks do not feature properties similar to those of social networks, mostly because of the physical limitations of both the movement and the radio transmission processes. Indeed, as noted in [20], the short range of communication attainable by low-power antennas enforces the same local dynamics typical of disease epidemics [11], which require physical proximity to propagate. Indeed, the analysis of opportunistic networks, where nodes relay messages as they come close to one another, employs models from the study of human mobility [5, 4].

In the theory community there has been growing interest in modeling and analyzing information dissemination in dynamic scenarios, where a number of agents move either in a continuous space or along the nodes of some underlying graph and exchange information when their positions satisfy a specified proximity constraint. In [7, 8] the authors study the time it takes to broadcast information from one of k mobile agents to all others. The agents move on a square grid of n nodes and in each time step an agent can (a) exchange information with all agents at distance at most r from it, and (b) move to any random node at distance at most ρ from its current position. The results in these papers only apply to a very dense scenario where the number of agents is linear in the number of grid nodes (i.e., $k = \Theta(n)$). They show that the broadcast time is $\Theta(\sqrt{n}/r)$ w.h.p., when $\rho = O(r)$ and $r = \Omega(\sqrt{\log n})$ [7], and it is $O((\sqrt{n}/\rho) + \log n)$ w.h.p., when $\rho = \Omega(\max\{r, \sqrt{\log n}\})$ [8]. These results crucially rely on $r + \rho = \Omega(\sqrt{\log n})$, which implies that the range of agents' communications or movements at each step defines a connected graph.

In more realistic scenarios, like the one adopted in this paper, the number of agents is decoupled from the number of locations (i.e., the graph nodes) and a smoother dynamics is enforced by limiting agents to move only between neighboring nodes. A reasonable model consists of a set of multiple, simple random walks on a graph, one for each agent, with communication between two agents occurring when they are nodes whose distance is at most $r \geq 0$. One variant of this setting is the so-called *Frog Model*, where initially one of k agents is active (i.e., is performing a random walk), while the remaining agents do not move. Whenever an active agent hits an inactive one, the latter is activated and starts its own random walk. This model was mostly studied in the infinite grid focusing on the asymptotic (in time) shape of the set of vertices containing all active agents [3, 18].

A model similar to our scenario is often employed to represent the spreading of computer viruses in networks, and the

spreading time is also referred to as *infection time*. Kesten and Sidoravicius [19] characterized the rate at which an infection spreads among particles performing continuous-time random walks with the same jump rate. In [10], the authors provide a general bound on the average infection time when k agents (one of them initially affected by the virus) move in an n -node graph. For general graphs, this bound is $O(t^* \log k)$, where t^* denotes the maximum average meeting time of two random walks on the graph, and the maximum is taken over all pairs of starting locations of the random walks. Also, in the paper tighter bounds are provided for the complete graph and for expanders. Observe that the $O(t^* \log k)$ bound specializes to $O(n \log n \log k)$ for the n -node grid by applying the known bound on t^* of [1]. A tight bound of $\Theta((n \log n \log k)/k)$ on the infection time on the grid is claimed in [28], based on a rather informal argument where some unwarranted independence assumptions are made. Our results show that this latter bound is incorrect.

Recent work by Peres et al. [25] studies a process in which agents follow independent Brownian motions in \mathbb{R}^d . They investigate several properties of the system, such as detection, coverage and percolation times, and characterize them as functions of the spatial density of the agents, which is assumed to be greater than the percolation point. Leveraging on these results, they show that the broadcast time of a message is polylogarithmic in the number of agents, under the assumption that a message spreads within a connected component of the communication graph instantaneously, before the graph is altered by agents' motion.

1.2 Organization of the Paper

The rest of the paper is organized as follows. In Section 2, we define the quantities of interest and establish some technical facts which are used in the analysis. Section 3 contains our main results: first, we prove the upper bound on the broadcast time in the most restricted case, that is, when the information exchange occurs through physical contact of the agents (i.e., $r = 0$), and then we provide a matching lower bound, which holds for every value of the transmission radius r below the percolation point. Finally, in Section 4 we briefly discuss the connection between our result and other interesting related problems and devise some future research directions.

2. PRELIMINARIES

We study the dynamics of information exchange among a set A of k agents performing independent random walks on an n -node 2-dimensional square grid \mathcal{G}_n , which is commonly adopted as a discrete model for the domain where agents wander. We assume that $n \geq 2k$, since sparse scenarios are the most interesting from the point of view of applications; however, our analysis can be easily extended to denser scenarios. We suppose that the agents are initially placed uniformly and independently at random on the grid nodes. Time is discrete and agent moves are synchronized. At each step, an agent residing on a node v with n_v neighbors ($n_v = 2, 3, 4$), moves to any such neighbor with probability $1/5$ and stays on v with probability $1 - n_v/5$. With these probabilities it is easy to see that at any time step the agents are placed uniformly and independently at random on the grid nodes. The following two lemmas con-

tain important properties of random walks on \mathcal{G}_n , which will be employed for deriving our results².

LEMMA 1. *Consider a random walk on \mathcal{G}_n , starting at time $t = 0$ at node v_0 . There exists a positive constant c_1 such that for any node $v \neq v_0$, the probability $p(v, v_0)$ that v is visited within $(\|v - v_0\|)^2$ steps is*

$$p(v, v_0) \geq \frac{c_1}{\max\{1, \log(\|v - v_0\|)\}}.$$

PROOF. The Lemma is proved in [3, Theorem 2.2] for the infinite grid \mathbb{Z}^2 . By the ‘‘Reflection Principle’’ [13, Page 72], for each walk in \mathbb{Z}^2 that started in \mathcal{G}_n , crossed a boundary and then crossed the boundary back to \mathcal{G}_n , there is a walk with the same probability that does not cross the boundary and visits all the nodes in \mathcal{G}_n that were visited by the first walk. Thus, restricting the walks to \mathcal{G}_n can only change the bound by a constant factor. \square

LEMMA 2. *Consider the first ℓ steps of a random walk in \mathcal{G}_n which was at node v_0 at time 0.*

1. *The probability that at any given step $1 \leq i \leq \ell$ the random walk is at distance at least $\geq \lambda\sqrt{i}$ from v_0 is at most $2e^{-\lambda^2/2}$.*
2. *There is a constant c_2 such that, with probability greater than $1/2$, by time ℓ the walk has visited at least $c_2\ell/\log \ell$ distinct nodes in \mathcal{G}_n .*

PROOF. We observe that the distance from v_0 in each coordinate defines a martingale with bounded difference 1. Then, the first property follows from the Azuma-Hoeffding Inequality [22, Theorem 2.6]. As for the second property, let R_ℓ be the set of nodes reached by the walk in ℓ steps. By Lemma 1, $\mathbb{E}[R_\ell] = \Omega(\ell/\log \ell)$ (even when v_0 is near a boundary), while $\text{Var}(R_\ell) = \Theta(\ell^2/\log^4 \ell)$ (see [27]). The result follows by applying Chebyshev's inequality. \square

Let M be a set of messages, which will be referred to as *rumors* henceforth, such that for each $m \in M$ there is (at least) one agent *informed* of m at time $t = 0$. W.l.o.g., we can assume that the number of distinct rumors is at most equal to the number of agent. We denote by $M_a(t)$ the set of rumors that agent $a \in A$ is informed of at time t , for any $t \geq 0$; possibly, $M_a(0) = \emptyset$. We assume that each agent is equipped with a *transmission radius* $r \in \mathbb{N}$, representing the maximum distance at which the agent can send information in a single time step.

The spread of rumors can be represented by a dynamic communication graph process $\{G_t(r) \mid t \geq 0\}$, where $G_t(r)$, the *visibility graph at time t* , is a graph with vertex set A and such that there is an edge between two vertices iff the corresponding agents are within distance r at time t . Following a common assumption justified by the physical reality that the speed of radio transmission is much faster than the motion of the agents [25], we suppose that rumors can travel throughout a connected component of $G_t(r)$ before the graph is altered by the motion. We assume that within the same connected component agents exchange all rumors they are informed of. Formally, let C be a connected component of $G_t(r)$: for all $a \in C$, $M_a(t) = \bigcup_{a' \in C} M_{a'}(t-1)$.

²Throughout the paper, the distance between two grid nodes u and v , denoted by $\|u - v\|$, is defined to be the Manhattan distance. Also, all logarithms are taken to the base two.

Note that the sets $M_a(t)$ can only grow over time, that is, agents do not “forget” rumors. The following quantities will be studied in this paper.

DEFINITION 1 (BROADCAST TIME, GOSSIP TIME).

The broadcast time T_B^m of a rumor $m \in M$ is the first time at which every agent is informed of m , that is, for all $t \geq T_B^m$ and $a \in A$, $m \in M_a(t)$. The gossip time T_G of the system is the first time at which every agent is informed of every rumor, that is, for any $t \geq T_G$ and $a \in A$, $M_a(t) = M$.

Note that both T_B^m and T_G depend on the transmission radius r , but we will omit this dependence to simplify the notation. We will also write T_B instead of T_B^m when the message m is clearly identified by the context.

3. BROADCASTING BELOW THE PERCOLATION POINT

In this section we give bounds to the broadcast time T_B of a rumor when the transmission radius is below the percolation point $r_c \approx \sqrt{n}/k$, that is, when all the connected components of $G_t(r)$ comprise at most a logarithmic number of agents. In this regime, we show that quite surprisingly T_B does not depend on the transmission radius, the reason being that the broadcast time is dominated by the time it takes for many independent random walks to intersect one another. In Subsection 3.1 we prove an upper bound on the broadcast time T_B in the extreme case $r = 0$, that is, when agents can exchange information only when they meet on a grid node. The same upper bound clearly holds for any other $r > 0$. Then, in Subsection 3.2 we show that the upper bound is tight, within logarithmic factors, for all values of the transmission radius below the percolation point. We also argue that the bounds on T_B easily extend to gossip time T_G .

3.1 Upper Bound on the Broadcast Time

The main technical ingredient of the analysis carried out in this subsection is the following lower bound on the probability that two random walks \bar{a}, \bar{b} on the grid meet within a given time interval and not too far from their starting positions, which is a result of independent interest.

LEMMA 3. Consider two independent simple random walks on the grid $\bar{a} = \langle a_0, a_1, \dots \rangle$, and $\bar{b} = \langle b_0, b_1, \dots \rangle$, where a_t and b_t denote the locations of the walks at time $t \geq 0$. Let $d = \|a_0 - b_0\| \geq 1$ and define D to be the set of nodes at distance at most d from both a_0 and b_0 . For $T = d^2$, there exists a constant $c_3 > 0$ such that

$$\begin{aligned} P_{\bar{a}, \bar{b}}(T) &\triangleq \Pr(\exists t \leq T \text{ such that } a_t = b_t \in D) \\ &\geq c_3 / \max\{1, \log d\}. \end{aligned}$$

PROOF. The case $d = 1$ is immediate. Consider now the case $d > 1$. Let $P_t(w, x)$ denote the probability that a walk that started at node w at time 0 is at node x at time t , and let $R(w, u, D, s)$ be the expected number of times that two walks which started at nodes w and u at time 0 meet at nodes of D during the time interval $[0, s]$, then

$$R(w, u, D, s) = \sum_{t=0}^s \sum_{x \in D} P_t(w, x) P_t(u, x).$$

Let $\tau(a, b)$ be the first meeting time of the walks \bar{a} and \bar{b} at a node of D . Then

$$\begin{aligned} R(a_0, b_0, D, T) &= \sum_{t=0}^T \Pr(\tau(a, b) = t) R(a_t, a_t, D, T - t) \\ &\leq P_{\bar{a}, \bar{b}}(T) \max_x R(x, x, D, T). \end{aligned}$$

Thus,

$$P_{\bar{a}, \bar{b}}(T) \geq \frac{R(a_0, b_0, D, T)}{\max_x R(x, x, D, T)}.$$

It is easy to verify that $|D| \geq d^2/4$. Applying Theorem 1.2.1 in [21] we have:

$$\begin{aligned} R(a_0, b_0, D, T) &\geq \sum_{t=0}^T \sum_{x \in D} P_t(a_0, x) P_t(b_0, x) \\ &\geq \sum_{t=\frac{T}{2}+1}^T \sum_{x \in D} 4 \left(\frac{1}{\pi t}\right)^2 e^{-\frac{\|x-a_0\|^2 + \|x-b_0\|^2}{t}}. \end{aligned}$$

By bounding $\|x - a_0\|^2$ and $\|x - b_0\|^2$ from above with T in the formula, easy calculations show that $R(a_0, b_0, D, T) = \Omega(1)$. Similarly, using the fact that there are no more than $4i$ nodes at distance exactly i from x , we have:

$$\begin{aligned} \max_x R(x, x, D, T) &\leq 1 + \sum_{t=1}^T \sum_{i=1}^t 4i 4 \left(\frac{1}{\pi t}\right)^2 2e^{-\frac{i^2}{t}} \\ &\leq 1 + \left(\frac{4}{\pi}\right)^2 \sum_{t=1}^T \frac{1}{t^2} \left(\left(\sum_{i=1}^{\sqrt{t}} i\right) + \left(\sum_{i=1+\sqrt{t}}^t i e^{-i^2/t}\right) \right) \\ &\leq 1 + \left(\frac{4}{\pi}\right)^2 \sum_{t=1}^T \frac{1}{t^2} \left(\frac{t}{2} + \left(\sum_{i=1+\sqrt{t}}^t i^2 e^{-i^2/t}\right) \right) \\ &\leq 1 + \left(\frac{4}{\pi}\right)^2 \sum_{t=1}^T \frac{1}{t^2} \left(\frac{t}{2} + \frac{e}{(e-1)^2} t \right) = O(\log T). \end{aligned}$$

We conclude that there is a constant $c_3 > 0$ such that $P_{\bar{a}, \bar{b}}(T) \geq c_3 / \log d$. \square

Observe that considering the difference random walk $\bar{a} - \bar{b} = \langle a_0 - b_0, a_1 - b_1, \dots \rangle$ and computing the probability that it hits the origin in the prescribed number of steps does not provide any information about the place where the meeting occurs, hence it is not immediate to derive the above result through that approach.

The remainder of this section is devoted to proving the following upper bound on the broadcast time of a single rumor m in the case $r = 0$. We assume that $M_a(0) = \{m\}$ for some $a \in A$, and $M_{a'}(0) = \emptyset$ for any other $a' \neq a$.

THEOREM 1. Let $r = 0$. For any $k \geq 2$, with probability at least $1 - 1/n^2$,

$$T_B = \tilde{O}\left(\frac{n}{\sqrt{k}}\right).$$

We observe that since the diameter of \mathcal{G}_n is $2\sqrt{n} - 2$, we can use Lemma 3 to show that with probability at least $1 - 1/n^2$, at time $8n \log^2 n$ an agent has met all other agents walking in \mathcal{G}_n . Thus, the theorem trivially holds for $k = O(\text{poly} \log(n))$.

From now on we concentrate on the case $k = \Omega(\log^3 n)$. We tessellate \mathcal{G}_n into cells of side $\ell \triangleq \sqrt{14n \log^3 n / (c_3 k)}$, where c_3 is defined in Lemma 3. We say that a cell Q is reached at time t_Q if t_Q is the first time when a node of the cell hosts an agent informed of the rumor, and we call this first visitor the *explorer* of Q . We first show that, after a suitably chosen number $T_1 = O(\ell^2 \log^4 n)$ of steps past t_Q , there is a large number of informed agents within distance $O(\ell \log^{5/2} n)$ from Q . Furthermore, we show that while the rumor spreads to cells adjacent to Q , at any time $t \geq t_Q + T_1$ a large number of informed agents are at locations close to Q . These facts will imply that the exploration process proceeds smoothly and that all agents are informed of the rumor shortly after all cells are reached.

The above argument is made rigorous in the following sequence of lemmas.

LEMMA 4. Consider an arbitrary $\ell \times \ell$ cell Q of the tessellation. Let $T_1 = 16\beta\gamma\ell^2 \log^4 n$ and $c_4 = 8\sqrt{5}\beta\gamma$, where $\beta = 7/(2c_1)$ and $\gamma = 18/c_3$. By time $\tau_1 = t_Q + T_1$, at least $4\beta \log^2 n$ agents are informed and are at distance at most $2(1 + c_4 \log^{5/2} n)\ell$ from Q , with probability $1 - 1/n^8$, for sufficiently large n .

PROOF. Since at any given time the agents are at random and independent locations, by the Chernoff bound we have that the following *density condition* holds with probability at least $1 - 1/n^9$, for sufficiently large n : for any cell Q' and any time instant $t \in [0, n \log^4 n]$, the number of agents residing in cell Q' at time t is at least $(7 \log^3 n)/c_3$. In the rest of the proof, we assume that the density condition holds.

First, we prove that, by time τ_1 , there are at least $4\beta \log^2 n$ informed agents in the system. We assume that at every time step $t \in [t_Q, \tau_1]$ there is always an uninformed agent in the same cell where the explorer resides (otherwise the sought property follows immediately by the density condition). For $1 \leq i \leq 4\beta \log^2 n$, let $t_i \geq t_Q$ be the time at which the explorer of Q informs the i -th agent. For notational convenience, we let $t_0 = t_Q$. To upper bound t_i , for $i > 0$, we consider a sequence of $\gamma \log^2 n$ consecutive, non-overlapping time intervals of length $4\ell^2$ beginning from time t_{i-1} . By the previous assumption, at the beginning of each interval the cell where the explorer resides contains an uninformed agent a . Hence, by Lemma 3, the probability that the explorer fails to meet an uninformed agent during all of these intervals is

$$\begin{aligned} \Pr(t_i > t_{i-1} + 4\gamma\ell^2 \log^2 n) &\leq (1 - c_3/\log(2\ell))^{\gamma \log^2 n} \\ &\leq 1/n^9, \end{aligned}$$

where the last inequality holds for sufficiently large n by our choice of γ . By iterating the argument for every i , we conclude that with probability at least $1 - 4\beta \log^2 n/n^9$, there are at least $4\beta \log^2 n$ informed agents at time τ_1 . Let I denote the set of informed agents identified through the above argument, and observe that each agent of I was in the cell containing the explorer at some time step $t \in [t_Q, \tau_1]$.

To conclude the proof of the lemma, we note that, by Lemma 2, the probability that the explorer, during the interval $[t_Q, \tau_1]$, reaches a grid node at distance greater than $(c_4 \log^{5/2} n)\ell$ from its position at time t_Q is bounded by $2T_1/n^{10}$. Consider an arbitrary agent $a \in I$. As observed above, there must have been a time instant $\bar{t} \in [t_Q, \tau_1]$ when a and the explorer were in the same cell, hence at distance at

most $(2 + c_4 \log^{5/2} n)\ell$ from Q . From time \bar{t} until time τ_1 the random walk of agent a proceeds independently of the random walk of the explorer. By applying again Lemma 2, we can conclude that the probability that one of the agents of I is at distance greater than $2(1 + c_4 \log^{5/2} n)\ell$ from Q at time τ_1 is at most $8\beta \log^2 n/n^9$. By adding up the upper bounds to the probabilities that the event stated in the lemma does not hold, we get $1/n^9 + 4\beta \log^2 n/n^9 + 2T_1/n^{10} + 8\beta \log^2 n/n^9$, which is less than $1/n^8$ for sufficiently large n . \square

LEMMA 5. Consider an arbitrary $\ell \times \ell$ cell Q of the tessellation. Let T_1, τ_1, c_4 and β be defined as in Lemma 4, and let $T_2 = (2(2 + c_4 \log^{5/2} n)\ell)^2$, $\tau_2 = \tau_1 + T_2$, and $c_5 = (4\sqrt{\log 16})c_4$. Then, the following two properties hold with probability at least $1 - 1/n^6$ for n sufficiently large:

1. For Q and for each of its adjacent cells, there exists a time t , with $\tau_1 \leq t \leq \tau_2$, at which there is an informed agent in the cell;
2. At any time t , with $\tau_1 \leq t \leq \tau_2 + T_1$, there are at least $\beta \log^2 n$ informed agents at distance at most $(2 + (2c_4 + c_5) \log^{5/2} n)\ell$ from Q .

PROOF. We condition on the event stated in Lemma 4, which occurs with probability $1 - 1/n^8$. Hence, assume that by time τ_1 there are at least $4\beta \log^2 n$ informed agents at distance at most $d_4 \triangleq 2(1 + c_4 \log^{5/2} n)\ell$ from Q . Consider the center node v of Q (resp., Q' adjacent to Q), so that at τ_1 there are at least $4\beta \log^2 n$ informed agents at distance at most $d_4 + 2\ell$ from v . By Lemma 1 the probability that v is not touched by an informed agent between τ_1 and τ_2 is at most $(1 - (c_1/\log(d_4 + 2\ell)))^{4\beta \log^2 n}$, which is less than $1/n^7$, for sufficiently large n , by our choice of β . Thus, Point 1 follows.

As for Point 2, consider an informed agent a which, at time τ_1 , is at a node x at distance at most d_4 from Q . Fix a time $t \in [\tau_1, \tau_2 + T_1]$. By Lemma 2 the probability that at time t agent a is at distance greater than $(c_5 \log^{5/2} n)\ell$ from x is at most $1/2$. Hence, at time t the average number of informed agents at distance at most $d_4 + (c_5 \log^{5/2} n)\ell$ from Q is at least $2\beta \log^2 n$. Since agents move independently, Point 2 follows by applying the Chernoff bound to bound the probability that at time t there are less than $\beta \log^2 n$ informed agents at distance at most $d_4 + (c_5 \log^{5/2} n)\ell$ from Q , and by applying the union bound over all time steps of the interval $[\tau_1, \tau_2 + T_1]$. \square

We are now ready to prove the main theorem of this subsection:

PROOF OF THEOREM 1. As observed at the beginning of the subsection, we can limit ourselves to the case $k = \Omega(\log^3 n)$. Consider the tessellation of \mathcal{G}_n into $\ell \times \ell$ cells defined before, and focus on a cell Q reached for the first time at t_Q . By Lemma 5, we know that with probability at least $1 - 1/n^6$, in each time step $t \in [\tau_1, \tau_2 + T_1]$ there are at least $\beta \log^2 n$ informed agents at distance at most $d_5 \triangleq (2 + (2c_4 + c_5) \log^{5/2} n)\ell$ from Q and there exists a time $t' \in [\tau_1, \tau_2]$ such that an informed agent is again inside Q . By applying again the lemma, we can conclude that, with probability at least $(1 - 1/n^6)^2$, at any time step $t'' \in [t' + T_1, t' + 2T_1 + T_2]$ there are at least $\beta \log^2 n$ informed agents at distance at most d_5 from Q . Note that the two time intervals $[\tau_1, \tau_2 + T_1]$ and $[t' + T_1, t' + 2T_1 + T_2]$ overlap

and the latter one ends at least T_1 time steps later. Thus, by applying the lemma $n \log^4 n$ times, we ensure that, with probability at least $(1 - 1/n^6)^{n \log^4 n} \geq 1 - \log^4 n/n^5$, from time τ_1 until the end of the broadcast, there are always at least $\beta \log^2 n$ informed agents at distance at most d_5 from Q .

Lemma 5 shows that each of the neighboring cells of Q is reached within time $\tau_2 = t_Q + T_1 + T_2$ with probability $1 - 1/n^6$. Therefore, all cells are reached within time $T^* = (2\sqrt{n}/\ell)(T_1 + T_2)$ with probability at least $1 - 1/n^5$. Hence, by applying a union bound over all cells, we can conclude that with probability at least $(1 - 1/n^5)(1 - \log^4 n/n^4) \geq 1 - 1/n^3$ there are at least $\beta \log^2 n$ informed agents at distance at most d_5 from each cell of the tessellation, from time $T^* + T_1$ until the end of the broadcast.

Consider now an agent a which at time $T^* + T_1$ is uninformed and resides in a certain cell Q . By an argument similar to the one used to prove Lemma 4, we can prove that a meets at least one of the informed agents around Q within $O(\ell^2 \log^5 n)$ time steps with probability at least $1 - 1/n^6$. A union bound over all uninformed agents completes the proof. \square

Observe that the broadcast time is a non-increasing function of the transmission radius. Therefore, the upper bound developed for the case $r = 0$ holds for any $r > 0$, as stated in the following corollary.

COROLLARY 1. *For any $k \geq 2$ and $r > 0$, $T_B = \tilde{O}\left(n/\sqrt{k}\right)$ with probability at least $1 - 1/n^2$.*

As another immediate corollary of the above theorem, we can prove that the gossiping of multiple distinct rumors completes within the same time bound, with high probability.

COROLLARY 2. *For any $k \geq 2$ and $r > 0$, $T_G = \tilde{O}\left(n/\sqrt{k}\right)$ with probability at least $1 - 1/n$.*

3.2 Lower Bound on the Broadcast Time

In this subsection we prove that the result of Corollary 1 is indeed tight, up to logarithmic factors, for any value r of the transmission radius below the percolation point. Note that this result is also a lower bound on T_G if there are multiple rumors in the system. First observe that with probability at least $1 - 2^{-(k-1)}$, there exists an agent placed at distance at least $\sqrt{n}/2$ from the source of m . W.l.o.g., we assume that the x -coordinates of the positions occupied by such an agent and the source agent differ by at least $\sqrt{n}/4$ and that the latter is at the left of the former. (The other cases can be dealt with through an identical argument.) In the proof, we cannot solely rely on a distance-based argument since we need to take into account the presence of “many” agents which may act as relay to deliver the rumor.

We define the *informed area* $\mathcal{I}(t)$ at time t as the set of grid nodes visited by any informed agent up to time t , and let $x(t)$ to be the rightmost grid node in $\mathcal{I}(t)$. We will show that there is a sufficiently large value T such that, at time T , there is at least one uninformed agent right of $x(T)$. We need the following definition:

DEFINITION 2 (ISLAND). *Let A be the set of agents. For any $\gamma > 0$, let $G_t(\gamma)$ be the graph with vertex set A and such that there is an edge between two vertices iff the corresponding agents are within distance γ at time t . Then*

any connected component of $G_t(\gamma)$ is called an island of parameter γ at time t .

Next, we prove an upper bound on the size of the islands.

LEMMA 6. *Let $\gamma = \sqrt{n/(4e^6 k)}$. Then, the probability that there exists an island of parameter γ with more than $\log n$ agents at any time t , with $0 \leq t \leq 8n \log^2 n$, is at most $1/n^2$.*

PROOF. Since at any given time the agents are uniformly distributed in \mathcal{G}_n , the probability that a given agent is within distance γ of another given agent at time t_0 is bounded by $4\gamma^2/n$. Fix a time t_0 and let $\mathcal{B}_w(t_0)$ denote the event that there exists an island with at least $w > \log n$ elements at time t_0 . Then, recalling that w^{w-2} is the number of unrooted trees over w labeled nodes, we have that

$$\begin{aligned} \Pr(\mathcal{B}_w(t_0)) &\leq \binom{k}{w} w^{w-2} \left(\frac{4\gamma^2}{n}\right)^{w-1} \\ &\leq \left(\frac{ek}{w}\right)^w w^{w-2} \left(\frac{4\gamma^2}{n}\right)^{w-1}. \end{aligned}$$

Using definition of γ and the bound $w \geq 1 + \log n$ and $k \leq n$, we have

$$\Pr(\mathcal{B}_w(t_0)) \leq \frac{ek}{w^2} e^{-5(w-1)} \leq \frac{en}{w^2 n^5} \leq \frac{1}{n^4},$$

for a sufficiently large n . Applying the union bound over $O(n \log^2 n)$ time steps concludes the proof. \square

Next we show that, with high probability, for values of r below percolation, the informed area cannot expand to the right too fast.

LEMMA 7. *Suppose $r \leq \sqrt{n/(64e^6 k)}$. Let $\gamma = \sqrt{n/(4e^6 k)}$ and let t_0 and $t_1 = t_0 + \gamma^2/(144 \log n)$ be two time steps. Then, with probability $1 - 2/n^2$,*

$$\|x(t_1) - x(t_0)\| \leq \gamma \log n.$$

PROOF. By Lemma 2, with probability $1 - 2/n^3$ an agent cannot cover a distance of more than $(\gamma - r)/2$ in $\gamma^2/(144 \log n)$ time steps. Thus with probability $1 - 1/n^2$, any two agents belonging to distinct islands of $G_{t_0}(\gamma)$ cannot come within distance r of each other in the interval $[t_0, t_1]$. Therefore, in that time interval, the rumor can propagate exclusively among agents belonging to those islands of $G_{t_0}(\gamma)$ containing at least one informed agent. By Lemma 6 we conclude that with probability $1 - 2/n^2$, in the interval $[t_0, t_1]$, $x(t)$ can move to the right of at most $\gamma(\log n - 1) + (\gamma - r)/2 < \gamma \log n$ positions. \square

Finally, we can prove the main theorem of the subsection:

THEOREM 2. *Let $k \geq 2$ and suppose that $r \leq \sqrt{n/(64e^6 k)}$. Then, with probability $1 - (2^{-(k-1)} + 1/n + 2/n^2)$,*

$$T_B = \Omega\left(\frac{n}{\sqrt{k} \log^2 n}\right).$$

PROOF. As mentioned before, with probability at least $1 - 2^{-(k-1)}$ there exists an agent a placed at distance at least $\sqrt{n}/2$ from the source of the rumor; we may assume that their x -coordinates differ by at least $\sqrt{n}/4$ and that the uninformed agent is to the right of the source agent. Let $T =$

$n/(2304e^3\sqrt{k}\log^2 n)$ and $\gamma = \sqrt{n/(4e^6k)}$. By Lemma 7, with probability $1 - 1/n$ the frontier cannot move right in T steps more than $(\gamma \log n/2)T/(\gamma^2/(144\log n)) < \sqrt{n}/8$. By Lemma 2, with probability $1 - 2/n^2$, agent a cannot move left more than $2\sqrt{T\log n} < \sqrt{n}/8$, so that agent a cannot be informed by time T . Hence, the broadcast time is at least $T_B > T = \Omega\left(n/(\sqrt{k}\log^2 n)\right)$ with probability at least $1 - (2^{-(k-1)} + 1/n + 2/n^2)$. \square

4. FURTHER RESULTS AND FUTURE RESEARCH

In this work we took a step toward a better understanding of the dynamics of information spreading in mobile networks. We proved a tight bound (up to logarithmic factors) on the broadcast of a rumor in a mobile network where agents perform independent random walks on a grid and the transmission radius defines a system below the percolation point. Our results complement the work of Peres et al. [25], who studied the behavior of a similar system above the percolation point. A similar bound holds for the gossip problem in this model, where at time 0 each agent has a distinct rumor and all agents need to receive all rumors.

Our analysis techniques are applicable to some interesting related problems. For example, similar bounds on the broadcast time T_B can be obtained for the Frog Model [3], where only informed agents move and uninformed agents remain at their initial positions. In particular, we can show that the broadcast time in the Frog Model is upper bounded by $T_B = \tilde{O}\left(n/\sqrt{k}\right)$. The argument is similar to the proof of Theorem 1, where Lemma 3 is replaced with Lemma 1 and the analysis of the initial phase of the information dissemination process is carried out by using Point 2 of Lemma 2. Also, a closer look at Theorem 2 reveals that the same argument employed in our dynamic model to bound T_B from below applies to the Frog Model. Thus, we have tight bounds, up to logarithmic factors, in this latter model as well.

Another measure of interest in systems of mobile agents is the *coverage time* T_C , that is, the first time at which every grid node has been visited at least once by an informed agent [25]. While in the Frog Model the broadcast time is obviously upper bounded by the coverage time, this relation is not so obvious in our dynamic model, since the coverage of the grid nodes does not imply that all agents have been informed of the rumor. Nevertheless, one can verify that, in our model, $T_C \approx T_B = \tilde{O}\left(n/\sqrt{k}\right)$. Indeed, by Point 2 of Lemma 5 and by Lemma 1, after $O(\ell^2)$ steps from the first time at which an informed agent reached a given cell, all the nodes of that cell have been visited by some informed agent. Hence, by the cell-by-cell spreading process devised in the proof of Theorem 1, we can conclude that the coverage time is bounded by $\tilde{O}\left(n/\sqrt{k}\right)$. (In fact, the same tight relation between T_C and T_B can be proved in the Frog Model.)

Another by-product of our techniques is a high-probability upper bound $O\left((n\log^2 n)/k + n\log n\right)$ on the *cover time* of k independent random walks on the n -grid (i.e., the time until each grid node has been touched by at least one such walk), improving on the previous results of [2, 12] which provide the same bound only for the expected value. Finally, in a closely related scenario, namely a random *predator-prey system* where $k = \Omega(\log n)$ predators are to catch moving

preys on an n -node grid by performing independent random walks [9], our techniques yield a high-probability upper bound $O\left((n\log^2 n)/k\right)$ on the extinction time of the preys.

In an effort to go beyond the pure mathematical contribution, we are now working on extending our analysis techniques to more complex scenarios that are interesting from the point of view of applications. Introducing mobility and communication barriers seems a first natural extension to be considered. Another dimension of mobile networks that we plan to further explore is the communication complexity of generic distributed computations among the moving agents of the system. Finally, the most challenging future direction leads to formulate sound analytical mobility models representing the dynamics of people travelling on road or subway networks, whereas, up to now, these types of systems have been studied by physicists, transportation scientists and engineers only by means of empirical or simulative techniques.

5. ACKNOWLEDGMENTS

Many thanks to Jeff Steif for referring us to some crucial references and to Andrea Clementi and Riccardo Silvestri for pointing out a claim not fully justified in an earlier draft of this paper.

6. REFERENCES

- [1] D. Aldous and J. Fill. *Reversible Markov Chains and Random Walks on Graphs*. Unpublished manuscript, 1998.
- [2] N. Alon, C. Avin, M. Koucký, G. Kozma, Z. Lotker, and M. R. Tuttle. Many random walks are faster than one. In *Proc. SPAA*, pages 119–128, 2008.
- [3] O. S. M. Alves, F. P. Machado, and S. Y. Popov. The shape theorem for the frog model. *The Annals of Applied Probability*, 12(2):533–546, 2002.
- [4] A. Chaintreau, P. Fraigniaud, and E. Lebar. Opportunistic spatial gossip over mobile social networks. In *Proc. WOSN*, pages 73–78, 2008.
- [5] A. Chaintreau, P. Hui, J. Crowcroft, C. Diot, R. Gass, and J. Scott. Impact of human mobility on opportunistic forwarding algorithms. *IEEE Trans. Mob. Comput.*, 6(6):606–620, 2007.
- [6] F. Chierichetti, S. Lattanzi, and A. Panconesi. Almost tight bounds for rumour spreading with conductance. In *Proc. STOC*, pages 399–408, 2010.
- [7] A. E. F. Clementi, A. Monti, F. Pasquale, and R. Silvestri. Information spreading in stationary markovian evolving graphs. In *Proc. IPDPS*, pages 1–12, 2009.
- [8] A. E. F. Clementi, F. Pasquale, and R. Silvestri. MANETS: High mobility can make up for low transmission power. In *Proc. ICALP*, pages 387–398, 2009.
- [9] C. Cooper, A. Frieze, and T. Radzik. Multiple random walks in random regular graphs. *SIAM Journal of Discrete Mathematics*, 23(4):1738–1761, 2009.
- [10] T. Dimitriou, S. Nikolettseas, and P. Spirakis. The infection time of graphs. *Discrete Applied Mathematics*, 154(18):2577–2589, 2006.
- [11] R. Durrett. Stochastic spatial models. *SIAM Review*, 41:677–718, 1999.
- [12] R. Elsässer and T. Sauerwald. Tight bounds for the

- cover time of multiple random walks. In *Proc. ICALP*, pages 415–426, 2009.
- [13] W. Feller. *An Introduction to Probability Theory and Its Applications, Vol. I*. Wiley, 3 edition, 1968.
- [14] M. Gerla. From battlefields to urban grids: New research challenges in ad hoc wireless networks. *Pervasive and Mobile Computing*, 1(1):77–93, 2005.
- [15] M. Grossglauser and D. N. C. Tse. Mobility increases the capacity of ad hoc wireless networks. *IEEE/ACM Transactions on Networking*, 10(4):477–486, 2002.
- [16] J. Hromkovic, R. Klasing, A. Pelc, P. Ruzicka, and W. Unger. *Dissemination of Information in Communication Networks*. Springer, Berlin, 2005.
- [17] P. Juang, H. Oki, Y. Wang, M. Martonosi, L.-S. Peh, and D. Rubenstein. Energy-efficient computing for wildlife tracking: Design tradeoffs and early experiences with zebranet. In *Proc. ASPLOS*, pages 96–107, 2002.
- [18] H. Kesten and V. Sidoravicius. A shape theorem for the spread of an infection. arXiv:math/0312511v1 [math.PR], 2003.
- [19] H. Kesten and V. Sidoravicius. The spread of a rumor or infection in a moving population. *The Annals of Probability*, 33(6):2402–2462, 2005.
- [20] J. Kleinberg. The wireless epidemic. *Nature*, 449(7160):287–288, 2007.
- [21] G. F. Lawler. *Intersections of random walks*. Birkhäuser, Boston, 1991.
- [22] M. Mitzenmacher and E. Upfal. *Probability and Computing*. Cambridge University Press, Cambridge, 2005.
- [23] S. Olariu and M. C. Weigle. *Vehicular Networks: From Theory to Practice*. Chapman and Hall/CRC, 2009.
- [24] M. Penrose. *Random Geometric Graphs*. Oxford University Press, Oxford, 2003.
- [25] Y. Peres, A. Sinclair, P. Sousi, and A. Stauffer. Mobile geometric graphs: Detection, coverage and percolation. In *Proc. SODA*, 2011.
- [26] I. Stojmenovic. *Handbook of Wireless Networks and Mobile Computing*. Wiley, 2002.
- [27] D. C. Torney. Variance of the range of a random walk. *Journal of Statistical Physics*, 44(1):49–66, 1986.
- [28] Y. Wang, S. Kapadia, and B. Krishnamachari. Infection spread in wireless networks with random and adversarial node mobilities. In *Proc. SIGMOBILE Workshop on Mobility Models*, pages 17–24, 2008.