

Sharing Multiple Messages over Mobile Networks

Yuxin Chen

Department of Electrical Engineering
Stanford University
Email: yxchen@stanford.edu

Sanjay Shakkottai and Jeffrey G. Andrews

Department of Electrical and Computer Engineering
The University of Texas at Austin
Email: {shakkott, jandrews}@ece.utexas.edu

Abstract—Information dissemination in a large network is typically achieved when each user shares its own information or resources with each other user. Consider n users randomly located over a fixed region, and k of them wish to flood their individual messages among all other users, where each user only has knowledge of its own contents and state information. The goal is to disseminate all messages using a *low-overhead* strategy that is *one-sided* and *distributed* while achieving an order-optimal spreading rate over a random geometric graph.

In this paper, we investigate the random-push gossip-based algorithm where message selection is based on the sender's own state in a random fashion. It is first shown that random-push is inefficient in static random geometric graphs. Specifically, it is $\Omega(\sqrt{n})$ times slower than optimal spreading. This gap can be closed if each user is mobile, and at each time moves “locally” using a random walk with velocity $v(n)$. We propose an efficient dissemination strategy that alternates between individual message flooding and random gossiping. We show that this scheme achieves the optimal spreading rate as long as the velocity satisfies $v(n) = \omega(\sqrt{\log n/k})$. The key insight is that the mixing introduced by this velocity-limited mobility approximately *uniformizes* the locations of all copies of each message within the optimal spreading time, which emulates a balanced geometry-free evolution over a complete graph.

I. INTRODUCTION

In wireless ad hoc or social networks, a variety of scenarios require agents to share their individual information or resources with each other for mutual benefit. A partial list includes file sharing and rumor spreading [1]–[3], distributed computation and parameter estimation [4]–[6], and scheduling and control [7], [8]. Due to the huge centralization overhead and unpredictable dynamics in large networks, it is usually more practical to disseminate information and exchange messages in a *decentralized* and *asynchronous* manner to combat unpredictable topology changes and the lack of global state information. Besides, it is often desirable to disseminate all messages efficiently in order to allow effective and prompt actions in various applications. These motivate the exploration of dissemination strategies that are inherently simple, distributed and asynchronous while achieving optimal spreading rates.

A. Motivation and Related Work

Among distributed asynchronous algorithms, gossip algorithms are a class of protocols which propagate messages according to rumor-style rules, initially proposed in [9]. Specifically, each agent in each round attempts to communicate with one of its neighbors in a random fashion to disseminate a limited number of messages. There are two types of *push-based* message selection strategies: (a) *one-sided protocols*

that are based only on the disseminator's own current state; and (b) *two-sided protocols* based on current states of both the sender and the receiver. Encouragingly, a simple one-sided push-only gossip algorithm with random message selection and peer selection is sufficient for efficient dissemination in some cases like a complete graph, with only a logarithmic gap between its spreading time and the optimal lower limit [10]. This type of one-sided gossiping has the advantages of being easily implementable and inherently distributed.

It has been pointed out, however, that the spreading rate of one-sided random gossip algorithms is frequently constrained by the network geometry, e.g. the conductance of the graph [10], [11] that captures the bottleneck ratio. For instance, one-sided rumor-style spreading is much more efficient in a complete graph than in a ring or a random geometric graph. Intuitively, since each user can only communicate with its nearest neighbors, the geometric constraints in these graphs limit the location distribution of all copies of each message during the evolution process, which largely limits the conductance of potential flow. In fact, one-sided push-based random gossiping can only complete spreading over static wireless networks $\Omega\left(\sqrt{\frac{n}{\text{poly}(\log n)}}\right)^1$ slower than the optimal lower limit, detailed later. Here n denotes the number of users.

This large gap can be closed if we take advantage of random linear coding where a random combination of all messages are transmitted instead of a specific message [12], or two-sided protocols which always disseminates innovative message if possible [13]. Otherwise optimal spreading time can be achieved in either complete graphs [12] or geometric graphs with constant maximum degree [14]. However, performing random network coding incurs very large computation overhead for each user, and is still impractical in practice. Two-sided protocols frequently require additional feedback that increases communication overhead. Also, the state information of the target may sometimes be unobtainable due to privacy or security concern. Furthermore, if there are $k \ll \sqrt{n}$ messages that need to be disseminated, neither network coding nor two-sided protocols can achieve optimal spreading time $\Theta(k)$ in a static unicast wireless network. This is because the dissemination time cannot be smaller than the diameter of the

¹The standard notion $f(n) = \omega(g(n))$ means $\lim_{n \rightarrow \infty} g(n)/f(n) = 0$; $f(n) = o(g(n))$ means $\lim_{n \rightarrow \infty} f(n)/g(n) = 0$; $f(n) = \Omega(g(n))$ means \exists a constant c such that $f(n) \geq cg(n)$; $f(n) = O(g(n))$ means \exists a constant c such that $f(n) \leq cg(n)$; $f(n) = \Theta(g(n))$ means \exists constants c_1 and c_2 such that $c_1 g(n) \leq f(n) \leq c_2 g(n)$.

underlying graph which scales as $\Theta(\sqrt{n})$.

Although a simple one-sided push-only random gossiping is not advocated for static geometric graphs, it may potentially achieve better performance if each user has some degree of mobility – an intrinsic feature of many wireless and social networks. For instance, full mobility² changes the geometric graph to a complete graph, which clearly allows optimal spreading time. However, how much benefit can be obtained from other *degrees* of mobility – which may be significantly lower than full mobility – is not clear. Most existing results on gossip algorithms center on evolutions associated with homogeneous graph structure or a fixed adjacency matrix, which can not be readily extended for dynamic topology changes. To the best of our knowledge, the first work to analyze gossiping with mobility was [15], which focused on *energy-efficient* distributed averaging instead of *time-efficient* message propagation.

Another line of work has studied broadcasting scaling laws using more sophisticated non-gossip schemes over static wireless networks, e.g. [16], [17]. Recently, Resta *et al.* [18] began investigating broadcast schemes for mobile networks with a *single static* source constantly broadcasting new data, while we focus on a different problem with multiple mobile sources each sharing distinct message. Besides, [18] analyzed how to combat the *adverse* effect of mobility to ensure the same pipelined broadcasting as in static networks, whereas we are interested in how to take advantage of mobility to overcome the geometric constraint. In fact, with mobility, simply performing random gossiping – which is simpler than most non-gossip schemes and does not require additional overhead – is sufficient to achieve optimality.

B. Problem Definition and Main Modeling Assumptions

We study the following problem. Suppose there are n users randomly located in a plane of fixed region. The task is to disseminate $k \leq n$ distinct messages (each contained in one user initially without being divided into pieces) among all users. The message spreading can be categorized into two types: (a) *single-message dissemination*: a single user (or a constant number of users) wishes to flood its message to all other users; (b) *multi-message dissemination*: a large number $k(k \gg 1)$ of users wish to spread individual messages to all other users. It should be noted that distinct messages may not be injected into the network simultaneously. They may arrive in the network (possibly in batches) sequentially, but the arrival time information is unknown in the network.

Our objective is to design a *gossip-style one-sided* algorithm in the absence of coding, such that it can take advantage of the intrinsic feature of mobility to accelerate dissemination. Only the “*push*” operation is considered in this paper, i.e. a sender determines which message to transmit solely based on its own

current state, and in particular not using the intended receiver’s state. We are interested in identifying the range of the degree of mobility within which our algorithm achieves near-optimal spreading time $O(k \text{ poly}(\log n))$ for each message regardless of message arrival patterns. (As an aside, we note that the poly-logarithmic gap from optimality cannot be closed with push-only random gossip even for complete graphs [12], [13]).

Our basic network model is as follows. Initially, there are n users uniformly distributed over a unit square. We ignore edge effects so that every node can be viewed as homogeneous. Our models and analysis are *mainly* based on the context of wireless ad hoc networks, but it can be easily applied to other network scenarios that can be modeled as a random geometric graph.

Physical-Layer Transmission Model. Each transmitter employs the same amount of power P , and the noise power density is assumed to be η . The path-loss model is used where the power density attenuates according to a power law of the Euclidean distance. That said, node j receives the signal from transmitter i with power $Pr_{ij}^{-\alpha}$, where r_{ij} denotes the Euclidean distance between i and j with α being the path loss exponent. Denote by $\mathcal{T}(t)$ the set of transmitters at time instance t . We assume that a packet from transmitter i is successfully received by a node j at time t if

$$\text{SINR}_{ij}(t) := \frac{Pr_{ij}^{-\alpha}}{\eta + \sum_{k \neq i, k \in \mathcal{T}(t)} Pr_{kj}^{-\alpha}} \geq \beta, \quad (1)$$

where SINR_{ij} is the signal-to-interference-plus-noise ratio (SINR) received at j , and β the SINR threshold required for successful reception. For simplicity, we suppose only one fixed-size message or packet can be transmitted for each transmission pair in each time instance.

Suppose that each node can move with velocity $v(n)$ in this mobile network. We provide a precise description of the mobility pattern as follows.

Mobility Model. We divide the entire square into $m := 1/v^2(n)$ subsquares each of area $v^2(n)$ (where $v(n)$ is the velocity of the mobile), which forms a $\sqrt{m} \times \sqrt{m}$ discrete torus. At each time instance, every node moves according to a *random walk* on the $\sqrt{m} \times \sqrt{m}$ discrete torus. More precisely, if a node resides in a subsquare $(i, j) \in \{1, \dots, \sqrt{m}\}^2$ at time t , it may choose to stay in (i, j) or move to any of the eight adjacent subsquares each with probability $1/9$ at time $t + 1$. The position inside the new subsquare is selected *uniformly at random*. See Fig. 1 for an illustration.

We note that when $v(n) = 1/3 = \Theta(1)$, the pattern reverts to the full mobility model. In this random-walk model, each node moves independently according to a uniform ergodic distribution. In fact, a variety of mobility patterns have been proposed to model mobile networks, including i.i.d. (full) mobility [19], random walk (discrete-time) model [20], Brownian motion (continuous-time) pattern [21], etc. For simplicity, we model it as a discrete-time pattern, which captures intrinsic features of mobile networks like uncontrolled placement and movement of nodes.

²By full mobility, we mean that the location of the mobile is distributed independently and uniformly random over the entire network over consecutive time-steps (i.e., the velocity of the mobile can be “arbitrarily large”). This is sometimes also referred to in literature as the i.i.d. mobility model. In this paper, we consider nodes with “velocity-limited” mobility capability.




Figure 1. The unit square is equally divided into $m = 1/v^2(n)$ subsquares. Each node can jump to one of its 8 neighboring subsquares or stay in current subsquare with equal probability $1/9$ at the beginning of each slot.

C. Summary of Contributions

The main contributions of this paper include the following.

- 1) **Single-message dissemination in mobile networks.** We derive an upper bound on single-message ($k = \Theta(1)$) spreading time using push-only random gossiping (called RANDOM PUSH) in mobile networks. A gain of $O(v(n)\sqrt{n})$ in the spreading rate can be obtained compared with static networks, which is still limited by the underlying geometry unless there is full mobility.
- 2) **Multi-message dissemination in static networks.** A lower bound on multi-message spreading time is derived using RANDOM PUSH for static networks. It turns out that there may exist a gap as large as $\Omega\left(\frac{\sqrt{n}}{\text{poly}(\log n)}\right)$ between its spreading time and the optimal lower limit, which coincides with the *diameter* of the graph. The key intuition is that the copies of each message M_i tend to cluster around the source i at all time instances, thus resulting in capacity loss and low conductance of flows.
- 3) **Multi-message dissemination in mobile networks.** We design a *one-sided uncoded* message-selection strategy called MOBILE PUSH that accelerates multi-message spreading ($k = \omega(1)$) with mobility. An upper bound on the spreading time is derived, which is the *main result* of this paper. Once $v(n) = \omega\left(\sqrt{\frac{\log n}{k}}\right)$ (which is still *significantly smaller* than full mobility), the optimal spreading time $O(k \text{ poly}(\log n))$ can be achieved. The underlying intuition is that if the mixing time associated with the mobility model is relatively smaller than optimal spreading time, the *mixing* property approximately *uniformizes* the location of all copies of each message, which allows the evolution to emulate the process in complete graphs and hence results in the speedup.

II. STRATEGIES AND MAIN RESULTS

The main results of this work are outlined in this section, where only unicast scenario is considered. The dissemination protocols for wireless networks are a class of scheduling algorithms that can be decoupled into (a) *physical-layer transmission* strategies (link scheduling) and (b) *message selection* strategies (message scheduling).

One physical-layer transmission strategy and two message selection strategies are described separately, along with the order-wise performance bounds.

A. Strategies

1) *Physical-Layer Transmission Strategy:* In order to achieve efficient spreading, it is natural to resort to a decentralized transmission strategy that supports order-wise largest number (i.e. $\Theta(n)$) of concurrent successful transmissions per time instance. The following strategy is a candidate that achieves this objective with local communication.

UNICAST Physical-Layer Transmission Strategy:

- At each time slot, each node i is designated as a sender independently with constant probability θ , and a potential receiver otherwise. Here, $\theta < 0.5$ is independent of n and k .
- Every sender i attempts to transmit one message to its *nearest* potential receiver $j(i)$.

This simple “link” scheduling strategy, when combined with appropriate push-based message selection strategies leads to the order-optimality results in this paper. We note that the authors in [19], by adopting a slightly different strategy in which θn nodes are randomly designated as senders (as opposed to link-by-link random selection as in our paper), have shown that the success probability for each unicast pair is a constant. In the same spirit as [19], we can prove (which we omit here) that there exists a fixed constant c such that

$$\mathbb{P}(\text{SINR}_{i,j(i)}(t) > \beta) \geq c, \quad (n \rightarrow \infty) \quad (2)$$

for our strategy. That said, $\Theta(n)$ concurrent transmissions can be successful, which is order-optimal. Due to the randomness in both our mobility model and transmission strategy, we further assume that physical-layer success events are *temporally independent* for simplicity of analysis and exposition. Indeed, we can show that accounting for the correlation yields the same scaling results.

Although this physical-layer transmission strategy supports $\Theta(n)$ concurrent local transmissions, it does not tell us how to take advantage of these resources to allow efficient propagation. This will be specified by the message-selection strategy, which potentially determines how each message is propagated and forwarded over the entire network.

2) *Message Selection Strategy:* We now turn to the objective of designing a *one-sided* message-selection strategy (only based on the transmitter’s current state) that is efficient in the absence of network coding. We are interested in a *decentralized* strategy in which no user has prior information

on the number of distinct messages existing in the network. One common strategy within this class is:

RANDOM PUSH Message Selection Strategy:

- In every time slot: each sender i randomly selects one of the messages it possesses for transmission.

This is a simple gossip algorithm solely based on random message selection, which is surprisingly efficient in many cases like a complete graph. It will be shown later, however, that this simple strategy is inefficient in a static wireless network / random geometric graph.

In order to take advantage of the mobility, we propose the following alternating strategy within this class:

MOBILE PUSH Message Selection Strategy:

- Denote by M_i the message that source i wants to spread.
- In every odd time slot: for each sender i , if it has an individual message M_i , then i selects M_i for transmission; otherwise i randomly selects one of the messages it possesses for transmission.
- In every even time slot: each sender i randomly selects one of the messages it possesses for transmission.

In the above strategy, each sender alternates between random gossiping and flooding of its own message. This alternating operation is crucial if we do not know *a priori* the number of distinct messages. Basically, random gossiping allows rapid spreading by taking advantage of all available throughput, and provides a non-degenerate efficient approach that ensures approximately “uniform” evolution for all distinct messages. On the other hand, individual message flooding step plays the role of self-advocating, which guarantees that a sufficiently large number of copies of each message can be forwarded with the assistance of mobility (which is not true in static networks). This is critical at the initial stage of the evolution.

B. Main Results

Now we proceed to state our main theorems, each of which characterizes the performance for one distinct scenario.

1) *Single-Message Dissemination in Mobile Networks with RANDOM PUSH:* The first theorem states the *limited* benefits of velocity on the spreading rate for single-message spreading using RANDOM PUSH. It can be shown using similar proof that MOBILE PUSH and RANDOM PUSH achieve the same scaling for single-message dissemination.

Theorem 1. Assume that the velocity constraint obeys $v(n) = \Omega\left(\sqrt{\frac{\log n}{n}}\right)$, and that the number k of distinct messages obeys $k = \Theta(1)$. RANDOM PUSH message selection strategy is used in the unicast scenario. Denote by $T_{sp}^{uc}(i)$ be the time taken for all users to receive message M_i after M_i is injected into the network, then with probability at least $1 - n^{-2}$ we have

$$\forall i, \quad T_{sp}^{uc}(i) = O\left(\frac{\log n}{v(n)}\right). \quad (3)$$

It can be observed that the gain of dissemination rate in mobile networks with respect to static networks is $\Theta(v(n)\sqrt{n})$. In fact, a random static network can usually be treated as equivalent (in order of magnitude) to a random geometric graph with transmission radius $r(n) = \Theta\left(\sqrt{\frac{\log n}{n}}\right)$. It is well known that such a random geometric graph has conductance $\Phi(n) = \Theta(r(n))$, which requires spreading time $\Theta\left(\frac{\log n}{\Phi(n)}\right) = \Theta\left(\frac{\log n}{r(n)}\right)$ [11]. Mobility plays the role of increasing the transmission radius, thus resulting in the speedup. It can be easily verified, however, that the universal lower bound on the spreading time is $\Theta(\log n)$ which can be achieved with full mobility. Therefore, the speedup is still limited by the *degree* of velocity $v(n)$ unless $v(n) = \Theta(1)$.

2) *Multi-Message Dissemination in Static Networks with RANDOM PUSH:* Now we turn to multi-message spreading ($k = \omega(\text{poly}(\log n))$) over static networks with random gossiping. It can be shown that the scaling performance of MOBILE PUSH is the same as RANDOM PUSH – the individual message flooding operation does not accelerate the spreading since each source has only $O(\log n)$ potential neighbors. The following theorem implies that simple RANDOM PUSH is inefficient in static wireless networks, under a message injection scenario where users start message dissemination sequentially. The setting is as follows: ($k - 1$) of the sources inject their messages into the network at some time prior to the k -th source. At a future time when each user in the network has at-least $w = \omega(\text{poly}(\log n))$ messages, the k -th message (denoted by M^*) is injected into the network.

Theorem 2. Assume that a new message M^* arrives in a static network later than other $k - 1$ messages, and suppose that M^* is first injected into the network from a state such that each node has received at least $w = \omega(\text{poly}(\log n))$ distinct messages. Denote by T^* the time until every user receives M^* using RANDOM PUSH, then for any constant $\epsilon > 0$ we have

$$T^* > w^{1-\epsilon} \sqrt{\frac{n}{128 \log n}} \quad (4)$$

with probability exceeding $1 - n^{-2}$.

Theorem 2 implies that if M^* is injected into the network when each user contains $\Omega\left(\frac{k^{1+2\epsilon}}{\sqrt{n}}\right)$ messages for any $\epsilon > 0$, then random push algorithm is unable to achieve the optimal spreading time $O(k \text{ poly}(\log n))$. In particular, if the message is first transmitted out when each user contains $\Omega(k/\text{poly}(\log n))$ messages, then it requires at least $\Theta(k^{1-\epsilon}\sqrt{n}/\text{poly}(\log n))$ time slots to complete spreading. That said, there may exist a gap as large as $\Omega\left(\frac{\sqrt{n}}{\text{poly}(\log n)}\right)$ between optimal performance and the one using RANDOM PUSH in statics networks. The reason is that the copies of each message tend to cluster around the source instead of spreading out at any time slot. A number of transmissions are wasted due to the blindness of the one-sided message selection, which results in capacity loss and hence the low conductance of information flow.

3) *Multi-Message Dissemination in Mobile Networks with MOBILE PUSH*: Although limited velocity *cannot* optimize propagation rate for single-message dissemination, it turns out to be remarkably helpful in multi-message case as stated in the following theorem.

Theorem 3. Assume that the velocity obeys: $v(n) = \omega\left(\sqrt{\frac{\log n}{k}}\right)$, where the number k of distinct messages obeys $k = \omega(\text{poly}(\log n))$. *MOBILE PUSH* message selection strategy is used along with unicast transmission strategy. Let $T_{mp}^{uc}(i)$ be the time taken for all users to receive message M_i after M_i is first injected into the network, then with probability at least $1 - n^{-2}$ we have

$$\forall i, \quad T_{mp}^{uc}(i) = O(k \log^2 n). \quad (5)$$

Since each node can receive at most one message in each time slot, the optimal spreading time is lower bounded by $\Theta(k)$ for any graph. Thus, our strategy with limited velocity spreads the information essentially as fast as possible. Intuitively, this is due to the fact that the velocity (even with *restricted* magnitude) helps *uniformize* the locations of all copies of each message, which significantly increases the *conductance* of the underlying graph in each slot. Although the velocity is significantly smaller than full mobility (which simply results in a complete graph), the relatively low mixing time helps achieve the same objective of *uniformization* approximately. On the other hand, the low spreading rates in static networks arise from the fact that the copies of each message tend to cluster around the source at any time instant, which decreases the number of flows going towards new users without this message.

III. ANALYSIS

The proofs of Theorem 2-3 are briefly outlined in this section. Interested readers are referred to [22] for complete proofs of all theorems and lemmas. Before continuing, we would like to state some preliminaries regarding the mixing time of a random walk on a 2-dimensional grid, and some related concentration results.

A. Preliminaries

1) *Mixing Time*: Define the probability of a typical node moving to subsquare A_i at time t as $\pi_i(t)$ starting from any subsquare. Define the *mixing time* of our random walk mobility model as $T_{\text{mix}}(\epsilon) := \min\{t : |\pi_i(t) - \frac{1}{m}| \leq \epsilon, \forall i\}$, which characterizes the time until the underlying Markov chain is close to its stationary distribution. It is well known that the mixing time of a random walk on a grid satisfies

$$T_{\text{mix}}(\epsilon) \leq \hat{c} m \log \frac{1}{\epsilon} \quad (6)$$

for some constant \hat{c} . We take $\epsilon = n^{-10}$ throughout this paper, so $T_{\text{mix}}(\epsilon) \leq c_0 m \log n$ holds with $c_0 = 10\hat{c}$. After $c_0 m \log n$ amount of time slots, all the nodes will reside in any subsquare *almost uniformly likely*. In fact, n^{-10} is very conservative and a much larger ϵ suffices for our purpose, but this gives us a good sense of the sharpness of the mixing time order.

2) *Concentration Results*: The following concentration result is also useful for our analysis.

Lemma 1. Assume that b ($b > 32m \log n$) nodes are thrown independently into m subsquares. Suppose for any subsquare A_i , the probability q_{A_i} of each node being thrown to A_i is bounded as

$$\left|q_{A_i} - \frac{1}{m}\right| \leq \frac{1}{3m}. \quad (7)$$

Then for any constant ϵ , the number of nodes $N_{A_i}(t)$ falling in any subsquare A_i ($1 \leq i \leq m < n$) at any time $t \in [1, n^2]$ satisfies

a) if $b = \Theta(m \log n)$, then

$$\mathbb{P}\left(\forall(i, t) : \frac{b}{6m} \leq N_{A(i)}(t) \leq \frac{7b}{3m}\right) \geq 1 - \frac{2}{n^3};$$

b) if $b = \omega(m \log n)$, then

$$\mathbb{P}\left(\forall(i, t) : \frac{(\frac{2}{3} - \epsilon)b}{m} \leq N_{A(i)}(t) \leq \frac{(\frac{4}{3} + \epsilon)b}{m}\right) \geq 1 - \frac{2}{n^3}.$$

This implies that the nodes residing in each subsquare at each time of interest will be reasonably close to the true mean. In fact, this concentration forms the basis for our analysis.

B. Multi-message Spreading in Static Networks with RANDOM PUSH

1) *The Lower Bound on the Spreading Time*: We sketch the proof of Theorem 2 in this subsection. To begin our analysis, we partition the entire unit square as follows:

- The unit square is divided into a set of nonoverlapping *tiles* $\{B_j\}$ each of side length $\sqrt{32 \log n / n}$ as illustrated in Fig. 2 (Note that this is a different partition from subsquares $\{A_j\}$ resulting from mobility model).




Figure 2. The plot illustrates that the number $N_{V_i}(t_X)$ of nodes containing M_i in vertical strip V_i by time t_X is decaying rapidly with geometric rate.

- The above partition also allows us to slice the network area into **vertical strips** each of width $\sqrt{32 \log n / n}$ and length 1. Label the vertical strips as $\{V_l\}$ ($1 \leq l \leq \sqrt{n / (32 \log n)}$) in increasing order from left to right, and denote by $N_{V_l}(t)$ and $\mathcal{N}_{V_l}(t)$ the number and the set of nodes in V_l that contains M^* by time t .
- The vertical strips are further grouped into **vertical blocks** $\{V_j^B\}$ each containing $\log n$ strips, i.e. $V_j^B = \{V_l : (j-1) \log n + 1 \leq l \leq j \log n\}$.

Remark 1. Since each tile has an area of $32 \log n / n$, concentration results imply that there are $\Theta(\log n)$ nodes residing in each tile with high probability. Since each sender only attempts to transmit to its nearest receiver, the communication process occurs only to nodes within the same tile or in adjacent tiles.

Without loss of generality, we assume that the source of M^* resides in the *leftmost* vertical strip V_1 . We aim at counting the time taken for M^* to cross each vertical block horizontally. In order to decouple the counting for different vertical blocks, we construct a new spreading process \mathcal{G}^* as follows.

Spreading in Process \mathcal{G}^* :

- 1) At $t = 0$, distribute M^* to all nodes residing in vertical strip V_1 .
- 2) Each node adopts RANDOM PUSH as the message selection strategy.
- 3) Define $T_l^B = \min \{t : N_{V_l^B}(t) > 0\}$ as the first time that M^* reaches vertical block V_l^B . For all $l \geq 2$, distribute M^* to all nodes residing in V_{l-1}^B at $t = T_l^B$.

It can be verified using coupling approach that \mathcal{G}^* evolves stochastically faster than the true process. By enforcing mandatory dissemination at every $t = T_l^B$, we enable separate counting for spreading time in different blocks – the spreading in V_{l+1}^B after T_{l+1}^B is independent of what has happened in V_l^B . Roughly speaking, since there are $\sqrt{\frac{n}{32 \log n}}$ blocks, the spreading time over the entire region is $\Theta\left(\sqrt{\frac{n}{\text{poly}(\log n)}}\right)$ times the spreading time over a typical block.

The following lemma characterizes the rate of propagation across different strips over a typical block.

Lemma 2. Consider the spreading of M^* over V_1^B in the constructed process \mathcal{G}^* . Suppose each node contains at least $w = \omega(\text{poly}(\log n))$ messages initially. Define $t_X := w^{1-\epsilon}$, and define $l^* = \min \{l : N_{V_l}^*(t_X) = O(w^\epsilon \log n)\}$. Then with probability at least $1 - 3n^{-3}$, we have

- (a) $l^* \leq \log n / 2$;
- (b) $\forall s$ ($1 \leq s < l^*$), there exists a constant c_{31} such that

$$N_{V_s}(t_X) \leq \left(\frac{\log n}{w^\epsilon}\right)^{s-1} (c_{31} \sqrt{n} \log n); \quad (8)$$

- (c) $N_{V_{\log n/2}}(t_X) \leq \log^2 n$.

Sketch of Proof of Lemma 2: The set of nodes $\mathcal{N}_{V_{l^*}}(t_X)$ may receive M^* either from nodes in adjacent strips or in V_{l^*} . Specifically, they can only receive M^* from the set of nodes $\mathcal{N}_{V_{l^*-1}}(t_X) \cup \mathcal{N}_{V_{l^*}}(t_X) \cup \mathcal{N}_{V_{l^*+1}}(t_X)$. Also, since every node contains $\Omega(w)$ messages, each transmitter containing M^* at time t will select M^* for transmission with probability at most $1/w$. Through fixed-point analysis we can derive

$$\begin{aligned} N_{V_{l^*}}(t_X) &\leq \frac{t_X}{w} (N_{V_{l^*-1}}(t_X) + N_{V_{l^*}}(t_X) + N_{V_{l^*+1}}(t_X)) \\ \implies N_{V_{l^*}}(t_X) &\leq \frac{2}{w^\epsilon} (N_{V_{l^*-1}}(t_X) + N_{V_{l^*+1}}(t_X)) \end{aligned}$$

with high probability. Repeatedly applying such comparisons between adjacent strips from $s = l^*$ to $s = 2$ yields

$$\begin{aligned} &N_{V_{l^*-s}}(t_X) + N_{V_{l^*+s}}(t_X) \\ &\leq \frac{\log n}{w^\epsilon} (N_{V_{l^*-s-1}}(t_X) + N_{V_{l^*+s+1}}(t_X)) \end{aligned}$$

That said, N_{V_l} decays as l increases with geometric rate ($1 \leq l < l^*$). Such rapid decay rate in turn guarantees that $l^* = O(\log n)$ with high probability. ■

The key observation from the above lemma is that the number of nodes in V_s containing M^* is decaying rapidly as s increases, which is illustrated in Fig. 2. After $N_{V_l}(t_X)$ decreases to $O(\log^2 n)$ for some l , we can show that $N_{V_l}(t_X)$ will vanish within no more than $O(\log n)$ further strips. This implies that M^* is unable to cross V_1^B by $t_X = w^{1-\epsilon}$.

Since there are $\Theta(\sqrt{n/\text{poly}(\log n)})$ vertical blocks in total, and crossing each block takes up $\Omega(w^{1-\epsilon})$ time slots, the time taken for M^* to cross all blocks can thus be bounded as

$$T^* = \Omega\left(w^{1-\epsilon} \sqrt{\frac{n}{\text{poly}(\log n)}}\right) \quad (9)$$

with high probability.

2) *Discussion:* This theorem implies that if a message M^* is injected into the network when each user contains $\Omega(k/\text{poly}(\log n))$ messages, the spreading time for M^* is $\Omega(k^{1-\epsilon} \sqrt{n/\text{poly}(\log n)})$ for arbitrarily small ϵ . That said, there exists a gap as large as $\Omega(\sqrt{n/\text{poly}(\log n)})$ from optimality. The tightness of this lower bound can be verified by deriving an *upper bound* using the conductance-based approach as follows.

We observe that the message selection probability for M^* is always lower bounded by $1/k$. Hence, we can couple a new process adopting a different message-selection strategy such that a transmitter containing M^* selects it for transmission with *state-independent* probability $1/k$ at each time. It can be verified that this process evolves stochastically slower than the original one. The conductance associated with the evolution for M^* is $\Phi(n) = \frac{1}{k} \Theta(r(n)) = O\left(\frac{1}{k} \sqrt{\frac{\log n}{n}}\right)$. Applying similar analysis as in [11] yields

$$T_i = O\left(\frac{\text{poly}(\log n)}{\Phi(n)}\right) = O(k \sqrt{n} \text{poly}(\log n)) \quad (10)$$

with probability exceeding $1 - n^{-2}$, which is only a poly-logarithmic gap from the lower bound we derived.

The tightness of this upper bound implies that the propagation bottleneck is captured by the conductance-based measure – the copies of each message tend to cluster around the source at any time instead of spreading out. That said, only the nodes lying around the boundary are likely to forward the message to new users. Capacity loss occurs to the users inside the cluster since many transmissions occur to receivers who have already received the message and are thus wasted. This bottleneck can be overcome with the assistance of mobility.

C. Multi-message Spreading in Mobile Networks with MOBILE PUSH

The proof of Theorem 3 is sketched in this subsection. We divide the entire evolution process into 3 phases. The duration of Phase 1 is chosen to allow each message to be forwarded to a sufficiently large number of users, and these copies be available almost “uniformly distributed” over the entire region. After this initial phase (which acts to “seed” the network with a sufficient number of all the messages), random gossiping in Phases 2 and 3 ensures the spread of all messages to all nodes.

1) *Phase 1:* This phase accounts for the first $c_6 \left(c_0 m \log n + \frac{m}{c_h \log n} \right) \log^2 n = O(m \log^3 n)$ time slots, where c_0 is the preconstant of the typical mixing time, and c_6 is a properly chosen constant. At the end of this phase, each message will be contained in at least $\Theta(m \log n)$ nodes. The time intended for this phase largely exceeds the mixing time of the random walk mobility model, which enables these copies to “uniformly” spread out over space.

We are interested in counting how many nodes will contain a particular message M_i by the end of Phase 1. Instead of counting all potential multi-hop relaying of M_i , we only look at the set of nodes that receive M_i directly from the source i in *odd* slots. This approach provides a lower bound on $N_i(t)$ at the end of Phase 1, but it suffices for our purpose.

Consider the following scenario: at time t_1 , node i attempts to transmit its message M_i to receiver j . Denote by $Z_i(t)$ ($1 \leq i \leq n$) the subsquare position of node i , and define the relative coordinate $Z_{ij}(t) := Z_i(t) - Z_j(t)$. Clearly, $Z_{ij}(t)$ forms another two-dimensional random walk on a discrete torus. The following lemma provides an upper bound on the expected number of time slots by time t during which the walk returns to $(0, 0)$.

Lemma 3. *For the random walk $Z_{ij}(t)$ defined above, there exist constants c_3 and c_h such that for any $t < \frac{m}{c_h \log n}$:*

$$\mathbb{E} \left(\sum_{k=1}^t \mathbb{1}(Z_{ij}(k) = (0, 0)) \middle| Z_{ij}(0) = (0, 0) \right) \leq c_3 \log t. \quad (11)$$

Here, $\mathbb{1}(\cdot)$ denotes the indicator function.

Sketch of Proof of Lemma 3: Denote by \mathcal{H}_{bd} the event that $Z_{ij}(t)$ hits the boundary \mathcal{A}_{bd} (as defined in Lemma 4)

before $t = m/(c_h \log n)$. The probability $q_{ij}^0(t)$ of $Z_{ij}(t)$ returning to $(0, 0)$ at time t can then be bounded as

$$q_{ij}^0(t) \leq \mathbb{P}(\mathcal{H}_{bd}) + \mathbb{P}(Z_{ij}(t) = (0, 0) \wedge \overline{\mathcal{H}_{bd}}) \quad (12)$$

Now, observe that when restricted to the set of sample paths where $Z_{ij}(t)$ does not reach the boundary by t , we can couple the sample paths of $Z_{ij}(t)$ to the sample paths of a random walk $\tilde{Z}_{ij}(t)$ over an infinite plane before the corresponding hitting time to the boundary. Denote by $\tilde{\mathcal{H}}_{bd}$ the event that $\tilde{Z}_{ij}(t)$ hits \mathcal{A}_{bd} by $t = m/(c_h \log n)$, then

$$\begin{aligned} \mathbb{P}(Z_{ij}(t) = (0, 0) \wedge \overline{\mathcal{H}_{bd}}) &= \mathbb{P}(\tilde{Z}_{ij}(t) = (0, 0) \wedge \overline{\tilde{\mathcal{H}}_{bd}}) \\ &\leq \mathbb{P}(\tilde{Z}_{ij}(t) = (0, 0)) \end{aligned}$$

The return probability obeys $\mathbb{P}(\tilde{Z}_{ij}(t) = (0, 0)) \sim t^{-1}$ for a random walk over an infinite plane, and $\mathbb{P}(\mathcal{H}_{bd})$ will be bounded in Lemma 4. Summing up all $q_{ij}^0(t)$ yields (11). ■

Lemma 4. *For the symmetric random walk $Z_{ij}(t)$ defined above, denote the set \mathcal{A}_{bd} of subsquares on the boundary as*

$$\mathcal{A}_{bd} = \left\{ A_i \mid A_i = \left(\pm \frac{\sqrt{m}}{2}, j \right) \text{ or } A_i = \left(j, \pm \frac{\sqrt{m}}{2} \right), \forall j \right\}.$$

Define the first hitting time to the boundary as $T_{hit} = \min \{t : Z_{ij}(t) \in \mathcal{A}_{bd}\}$, then there is a constant c_h such that

$$\mathbb{P} \left(T_{hit} < \frac{m}{c_h \log n} \right) \leq \frac{1}{n^4}. \quad (13)$$

In order to derive an estimate on the number of distinct nodes receiving M_i directly from source i , we need to calculate the number of slots where i fails to forward M_i to a new user. In addition to physical-layer outage events, some transmissions occur to users already possessing M_i , and hence are not successful. Recall that we are using one-sided push-only strategy, and hence we cannot always send an innovative message. Denote by $F_i(t)$ the number of *wasted transmissions* from i to some users already containing M_i by time t . This can be estimated as in the following lemma.

Lemma 5. *For $t_0 = \frac{m}{c_h \log n}$, the number of wasted transmissions $F_i(t)$ defined above obeys*

$$\mathbb{E}(F_i(t_0)) \leq c_5 \frac{m \log n}{n} t_0 \quad (14)$$

for some fixed constant c_5 with probability exceeding $1 - 3n^{-3}$.

Sketch of Proof of Lemma 5: Consider a particular pair of nodes i and j , where i is the source and j contains M_i . A wasted transmission occurs when (a) i and j meets in the same subsquare again, and (b) i is designated as a sender with j being the intended receiver. The probability of event (a) can be calculated using Lemma 3, while the probability of (b) is $\Theta(m/n)$ due to sharp concentration on N_{A_i} . ■

The above result is helpful in estimating the *expected number* of distinct users containing M_i . However, it is not obvious whether $F_i(t)$ exhibits desired sharp concentration. The difficulty is partly due to the dependence among $\{Z_{ij}(t)\}$

for different t arising from its Markov property. Due to their underlying relation with location of i , $Z_{ij_1}(t)$ and $Z_{ij_2}(t)$ are not independent either for $j_1 \neq j_2$. However, this difficulty can be circumvented by constructing different processes that exhibit approximate mutual independence as follows.

The time duration $[1, c_6(c_0 m \log n + m/(c_h \log n)) \log^2 n]$ of Phase 1 are divided into $c_6 \log^2 n$ non-overlapping subphases $P_{1,j}$ ($1 \leq j \leq \log^2 n$) for some constant c_6 . Each odd subphase accounts for $m/(c_h \log n)$ time slots, whereas each even subphase contains $c_0 m \log n$ slots. See Fig. 3 for an illustration. Instead of studying the true evolution, we consider different evolutions for each subphase. In each odd subphase, source i attempts to transmit message M_i to its intended receiver as in the original process. But in every even subphase, all new transmissions will be immediately deleted. The purpose for constructing these *clearance* or *relaxation* processes in even subphases is to allow for independent counting for odd subphases. The duration $c_0 m \log n$ of each even subphase, which is larger than the typical mixing time duration of the random walk, is sufficient to allow each user to move to everywhere almost uniformly likely.




Figure 3. Phase 1 is divided into $2c_6 \log^2 n$ subphases. Each odd subphase accounts for $m/(c_h \log n)$ slots, during which all nodes perform message spreading. Each even subphase contains $c_0 m \log n$ slots, during which no transmissions occur; it allows all nodes containing a typical message to be uniformly spread out.

Lemma 6. Set t to be $c_6 \left(c_0 m \log n + \frac{m}{c_h \log n} \right) \log^2 n$, which is the end time slot of Phase 1. The number of users containing each message M_i can be bounded below as

$$\forall i, \quad N_i(t) > 32m \log n \quad (15)$$

with probability at least $1 - c_7 n^{-2}$.

In fact, if $m \log^2 n \ll n$ holds, the above lemma can be further refined to $N_i(t) = \Theta(m \log^2 n)$. This implies that, by the end of Phase 1, each message has been flooded to $\Omega(m \log n)$ users. They are able to *cover* all subsquares (i.e., the messages' locations are roughly uniformly distributed over the unit square) after a further mixing time duration.

2) *Phase 2:* This phase starts from the end of Phase 1 and ends when $N_i(t) > n/8$ for all i . We use $t = 0$ to denote the starting slot of Phase 2 for convenience of presentation. Instead of directly looking at the original process, we generate a new process $\tilde{\mathcal{G}}$ which evolves slower than the original process \mathcal{G} . Define $\mathcal{S}_i(t)$ and $\tilde{\mathcal{S}}_i(t)$ as the set of messages that node i

contains at time t in \mathcal{G} and $\tilde{\mathcal{G}}$, with $S_i(t)$ and $\tilde{S}_i(t)$ denoting their cardinality, respectively. For more clear exposition, we divide the entire phase into several time blocks each of length $k + c_0 \log n/v^2(n)$, and use t_B to label different time blocks. We define $\tilde{N}_i^B(t_B)$ to denote $\tilde{N}_i(t)$ with t being the starting time of time block t_B . $\tilde{\mathcal{G}}$ is generated from \mathcal{G} : everything in these two processes remains the same (including locations, movements, physical-layer outage events, etc.) except message selection strategies, detailed below:

Message Selection Strategy in the Coupled Process $\tilde{\mathcal{G}}$:

- 1) Initialize: At $t = 0$, for all i , copy the set $\mathcal{S}_i(t)$ of all messages that i contains to $\tilde{\mathcal{S}}_i(t)$. Set $t_B = 0$.
- 2) In the next $c_0 \log n/v^2(n)$ time slots, all new messages received in this subphase are immediately deleted, i.e., no successful forwarding occurs in this subphase regardless of the locations and physical-layer conditions.
- 3) In the next k slots, for every sender i , each message it contains is randomly selected with probability $1/k$ for transmission.
- 4) For all i , if the number of nodes containing M_i is larger than $2\tilde{N}_i^B(t_B)$, delete M_i from some of these nodes so that $\tilde{N}_i(t) = 2\tilde{N}_i^B(t_B)$ by the end of this time block.
- 5) Set $t_B \leftarrow t_B + 1$. Repeat from (2) until $N_i > n/8$ for all i .

Thus, each time block consists of a relaxation period and a spreading period. The key idea is to simulate an *approximately spatially-uniform evolution*, which is summarized as follows:

- After each spreading subphase, we give the process a *relaxation* period to allow each node to move almost uniformly likely to all subsquares. This is similar to the relaxation period introduced in Phase 1.
- Trimming the messages alone does *not* necessarily generate a slower process, because it potentially increases the selection probability for each message. Therefore, we force the message selection probability to be a lower bound $1/k$, which is *state-independent*. Surprisingly, this conservative bound suffices for our purpose because it is exactly one of the *bottlenecks* for the evolution.

The following lemma makes a formal comparison of \mathcal{G} and $\tilde{\mathcal{G}}$.

Lemma 7. $\tilde{\mathcal{G}}$ evolves stochastically slower than \mathcal{G} , i.e.

$$\mathbb{P}(T_2 > x) < \mathbb{P}(\tilde{T}_2 > x), \quad \forall x > 0 \quad (16)$$

where $T_2 = \min \{t : N_i(t) > n/8, \forall i\}$ and $\tilde{T}_2 = \min \{t : \tilde{N}_i(t) > n/8, \forall i\}$ are the stopping time of Phase 2 for \mathcal{G} and $\tilde{\mathcal{G}}$, respectively.

Proof: Whenever a node i sends a message M_k to j in \mathcal{G} : (a) if $M_k \in \tilde{\mathcal{S}}_i$, then i selects M_k with probability S_i/k , and a random useless message otherwise; (b) if $M_k \notin \tilde{\mathcal{S}}_i$, i always sends a random noise message. The initial condition $\tilde{\mathcal{S}}_i = \mathcal{S}_i$ guarantees that $\tilde{\mathcal{S}}_i \subseteq \mathcal{S}_i$ always holds with this coupling method. Hence, the claimed stochastic order holds. ■

Lemma 8. Denote by $\tilde{T}_2^B := \min \{t_B : \tilde{N}_i^B(t_B) > n/8, \forall i\}$ the stopping time block of Phase 2 in $\tilde{\mathcal{G}}$. Then there exists a constant c_{14} independent of n such that

$$\mathbb{P}(\tilde{T}_2^B \leq 4 \log_{c_{14}} n) \leq 1 - n^{-2}.$$

Sketch of Proof of Lemma 8: We first look at a particular message M_i , and use union bound later after we derive the concentration results on the stopping time associated with this message. We observe the following facts: after a mixing time duration, the number of users $N_{i,A_k}(t)$ containing M_i at each subsquare A_k is approximately *uniform*. Since $N_i^B(t_B)$ is the lower bound on the number of copies of M_i across this time block, concentration results suggest that $N_{i,A_k}(t) = \Omega(N_i^B(t_B)/m)$. Observing from the mobility model that the position of any node inside a subsquare is *i.i.d.* chosen, we can derive

$$\mathbb{E}(\tilde{N}_i^B(t_B + 1) - \tilde{N}_i^B(t_B) \mid \tilde{N}_i^B(t_B)) \geq \frac{\tilde{c}_9}{2} \tilde{N}_i^B(t_B) \quad (17)$$

for some constant \tilde{c}_9 . A standard *martingale* argument then yields an upper bound on the stopping time. ■

This lemma implies that after at most $4 \log_{c_{14}} n$ time blocks, the number of nodes containing all messages will exceed $n/8$ with high probability. Therefore, the duration \tilde{T}_2 of Phase 2 of $\tilde{\mathcal{G}}$ satisfies $\tilde{T}_2 = O(k \log n)$ with high probability. This gives us an upper bound on T_2 of the original evolution \mathcal{G} .

3) *Phase 3:* This phase ends when $N_i(t) = n$ for all i with $t = 0$ denoting the end of Phase 2. Assume that $N_{i,A_j}(0) > \frac{n}{16m}$ for all i and all j , otherwise we can let the process further evolve for another mixing time duration $\Theta(\log n/v^2(n))$.

Lemma 9. Denote by T_4 the duration of Phase 3, i.e. $T_4 = \min \{t : N_i(t) = n \mid N_i(0) \geq n/8, \forall i\}$. Then there exists a constant c_{18} such that

$$\mathbb{P}(T_4 \leq \frac{64}{c_{18}} k \log n) \geq 1 - \frac{15}{16n^2} \quad (18)$$

Sketch of Proof of Lemma 9: The random push strategies are efficient near the start (exponential growth), but the evolution will begin to slow down after Phase 2. The concentration effect allows us to obtain a different evolution bound as

$$\begin{aligned} & \mathbb{E}(N_i(t+1) - N_i(t) \mid N_i(t)) \\ &= \mathbb{E}(n - N_i(t) - (n - N_i(t+1)) \mid N_i(t)) \\ &\geq \frac{c_{18}}{16k} (n - N_i(t)) \end{aligned}$$

Constructing a different submartingale based on $n - N_i(t)$ yields the above results. ■

4) *Discussion:* Combining the stopping time in all three phases, we can see that: the spreading time $T_{\text{mp}}^d = \min \{t : \forall i, N_i(t) = n\}$ satisfies

$$T_{\text{mp}}^d \leq O\left(\frac{\log^3 n}{v^2(n)}\right) + O(k \log n) + O(k \log n) = O(k \log^2 n).$$

It can be observed that, the mixing time bottleneck will not be critical in multi-message dissemination. The reason is that

the mixing time is typically much smaller than the optimal spreading time. Hence, the nodes have sufficient time to spread out to everywhere. The key step is how to see the network with a sufficiently large number of copies at the initial stage of the spreading process.

ACKNOWLEDGEMENT

This work has been supported in part by DARPA IT-MANET program.

REFERENCES

- [1] D. Qiu and R. Srikant, "Modeling and performance analysis of bittorrent-like peer-to-peer networks," *ACM SIGCOMM*, pp. 367–378, 2004.
- [2] X. Yang and G. de Veciana, "Service capacity of peer to peer networks," *IEEE INFOCOM*, vol. 4, pp. 2242 – 2252 vol.4, March 2004.
- [3] B. Pittel, "On spreading a rumor," *SIAM Journal of Applied Mathematics*, vol. 47, no. 1, pp. 213–223, 1987.
- [4] J. Tsitsiklis, *Problems in decentralized decision making and computation*, PhD dissertation, LIDS, MIT, Cambridge, 1984.
- [5] A. Kashyap and R. Srikant, "Quantized consensus," *Automatica*, vol. 43, pp. 1192–1203, 2007.
- [6] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah, "Randomized gossip algorithms," *IEEE Transactions on Information Theory*, vol. 52, no. 6, pp. 2508–2530, June 2006.
- [7] E. Modiano, D. Shah, and G. Zussman, "Maximizing throughput in wireless networks via gossiping," *ACM SIGMETRICS*, pp. 27–38, 2006.
- [8] A. Eryilmaz, A. Ozdaglar, D. Shah, and E. Modiano, "Distributed cross-layer algorithms for the optimal control of multihop wireless networks," *IEEE/ACM Transactions on Networking*, vol. 18, no. 2, pp. 638–651, Apr. 2010.
- [9] A. Demers, D. Greene, C. Houser, W. Irish, J. Larson, S. Shenker, H. Sturgis, D. Swinehart, and D. Terry, "Epidemic algorithms for replicated database maintenance," *Proceeding of the sixth annual ACM symposium on principles of distributed computing*, pp. 1–12, 1987.
- [10] D. Shah, "Gossip algorithms," *Foundations and Trends in Networking*, vol. 3, no. 1, pp. 1–125, April 2009.
- [11] D. Mosk-Aoyama and D. Shah, "Fast distributed algorithms for computing separable functions," *IEEE Transactions on Information Theory*, vol. 54, no. 7, pp. 2997–3007, July 2008.
- [12] S. Deb, M. Medard, and C. Choute, "Algebraic gossip: a network coding approach to optimal multiple rumor mongering," *IEEE Transactions on Information Theory*, vol. 52, no. 6, pp. 2486–2507, June 2006.
- [13] S. Sanghavi, B. Hajek, and L. Massoulie, "Gossiping with multiple messages," *IEEE Transactions on Information Theory*, vol. 53, no. 12, pp. 4640–4654, Dec. 2007.
- [14] M. Borokhovich, C. Avin, and Z. Lotker, "Tight bounds for algebraic gossip on graphs," *IEEE International Symposium on Information Theory*, June 2010.
- [15] A. Sarwate and A. Dimakis, "The impact of mobility on gossip algorithms," *IEEE INFOCOM*, pp. 2088–2096, April 2009.
- [16] R. Zheng, "Information dissemination in power-constrained wireless networks," *IEEE INFOCOM*, pp. 1–10, April 2006.
- [17] S. Subramanian, S. Shakkottai, and A. Arapostathis, "Broadcasting in sensor networks: The role of local information," *IEEE/ACM Transactions on Networking*, vol. 16, no. 5, pp. 1133–1146, Oct. 2008.
- [18] G. Resta and P. Santi, "On the fundamental limits of broadcasting in wireless mobile networks," *IEEE INFOCOM*, pp. 1–5, 14–19 2010.
- [19] M. Grossglauser and D. Tse, "Mobility increases the capacity of ad hoc wireless networks," *IEEE/ACM Transactions on Networking*, vol. 10, no. 4, pp. 477–486, Aug 2002.
- [20] A. El Gamal, J. Mammen, B. Prabhakar, and D. Shah, "Optimal throughput-delay scaling in wireless networks – Part I: the fluid model," *IEEE Trans. on Information Theory*, vol. 52, no. 6, pp. 2568–2592, June 2006.
- [21] X. Lin, G. Sharma, R. Mazumdar, and N. Shroff, "Degenerate delay-capacity tradeoffs in ad-hoc networks with Brownian mobility," *IEEE Trans. on Information Theory*, vol. 52, no. 6, pp. 2777–2784, June 2006.
- [22] Y. Chen, S. Shakkottai, and J. G. Andrews, "Sharing multiple messages over mobile networks," Tech. Rep. [Online]. Available: <http://www.stanford.edu/~yxchen/publications/InfocomGossipMobile.pdf>