Random Mobility and the Spread of Infection

Aditya Gopalan, Siddhartha Banerjee, Abhik K. Das, Sanjay Shakkottai
Department of Electrical and Computer Engineering
The University of Texas at Austin, USA
Email: {gadit, sbanerjee, akdas, shakkott}@mail.utexas.edu

Abstract—We study infection spreading on large static networks when the spread is assisted by a small number of additional *virtually mobile* agents. For networks which are "spatially constrained", we show that the spread of infection can be significantly sped up even by a few virtually mobile agents acting *randomly*. More specifically, for general networks with bounded virulence (e.g., a single or finite number of random virtually mobile agents), we derive upper bounds on the order of the time taken (as a function of network size) for infection to spread. Conversely, for certain common classes of networks such as linear graphs, grids and random geometric graphs, we also derive lower bounds on the order of the spreading time over all (potentially network-state aware and adversarial) virtual mobility strategies. We show that up to a logarithmic factor, these lower bounds for adversarial virtual mobility match the upper bounds on spreading via an agent with random virtual mobility. This demonstrates that random, state-oblivious virtual mobility is in fact order-wise optimal for dissemination in such spatially constrained networks.

I. INTRODUCTION

Various natural and engineered phenomena around us involve the spread of information or infection through different kinds of networks. Rumours and news stories propagate among people linked by various means of communication, diseases diffuse as epidemics through populations by various modes, plants disperse pollen/seeds and thus genetic traits geographically, riots spread across pockets of communities, advertisers aim to disseminate information about goods through networks of consumers, and computer viruses, email worms and software patches piggy-back across computer networks. Understanding how infection/information/innovation can travel across networks has been a subject of extensive study in disciplines ranging from epidemiology [1], [2], sociology [3], [4] and computer science [5], [6], [7] to physics [8], information theory/networking [9], [10], [11], [12], [13] and applied mathematics [14], [15], [16], yielding valuable insights into qualitative and quantitative aspects of network spread behaviour.

In this work, we model and study network-wide spread with two distinct components – a basic *static spread* component in which infection spreads naturally and *locally* through neighbouring nodes in the network, and an additional *virtually mobile spread* component in which the infection is carried to nodes far from its origin by suitable "virtually mobile agents", helping it spread *globally*. Specifically, we develop a rigorous framework with which we quantify the effect that a small number of *external* (*i.e.* not constrained by the underlying graph), *omniscient* (*i.e.* network-state aware) and *adversarial*

(*i.e.* free to infect any portion(s) of the network) virtually mobile agents can have on the time it takes for infection to spread throughout the network.

We stress that the terms "static spread" and "virtually mobile spread" (or virtual mobility) are used merely as surrogates for any situation involving the spread of infection/information via (spatially and/or timescale-wise) *heterogeneous* modes. In the context of wireless communication, for instance, consider the increasingly studied propagation [17], [18], [19], [20] of viruses and worms that exploit the connectivity afforded by both *(a)* modern short-range personal communication technologies like Bluetooth, and *(b)* long-range media such as SMS/MMS and the Internet. To paraphrase Kleinberg [21], outbreaks due to such worms are well-modelled by local spreading on a fixed network of nodes in space (*i.e.* short-range Bluetooth wireless transmissions between neighbouring quasi-static users) aided by relatively unrestricted paths through the network (*i.e.* long-range, faster-timescale emails and messages through SMS/MMS/Internet). Thus, "static spread" and "virtually mobile spread" here mean short-range Bluetooth transmissions between users and long-range network-wide emails/messages respectively.

Other, more classically-studied, examples of local spread assisted by forms of virtual mobility include those of natural disease epidemics [1] and bioterror attacks [22], where infection can spread *(a)* locally through spatial pathways (*i.e.* interpersonal contact) and *(b)* globally through faster, large-scale geographic means (*e.g.* human movement through airline routes) [23].

In all these and allied situations, it can be seen that, in addition to the local or static spreading behaviour of an infection over an underlying network, a form of external "virtual mobility" unconstrained by the structure of the underlying network causes long-range proliferation of the infection. Also, in most cases, the virtual mobility is such that *any* uninfected node in the network is susceptible to infection by it. Thus, virtual mobility (resp. a virtually mobile agent) is entirely distinct from, and more general than actual or physical mobility (resp. a physical mobile agent) whose infection spreading abilities are inherently constrained by its geographic nature. We wish to investigate, in this work, the effect that the virtual mobility has on the time taken for the entire network to get infected.

Given the applicability of our virtually mobile spread model, a fundamental characterization of the impact of all possible adversarial "patterns" of virtual mobility on the spreading time across the whole network is useful for two chief rea-

sons. Firstly, whenever malicious forms of infection such as epidemics or bioterror attacks originate and threaten to spread via both local and virtually mobile means, it becomes important to understand the *worst-case* virtually mobile spread behaviour (this is the component that can potentially accelerate the spread) in order to deploy appropriate countermeasures. Secondly, in cases where propagation is in fact desirable and the virtual mobility can actually be *controlled* (*e.g.* viral advertising [7], network protocol design [6] and diffusion of innovations [3]), an adversarial study of virtual-mobility-assisted spreading can constructively help in designing fast spreading strategies.

A. Main Contributions

We consider large graphs $G = (V, E)$ in which infection starts spreading in continuous time at a designated node according to the standard Susceptible-Infected (SI) dynamics [9] (also termed the *contact process* [14], [15]) with *i.i.d.* exponentially distributed propagation times through the edges. To model the spread of infection via additional long-range virtual mobility or virtually mobile agents as discussed earlier, we allow *every* node in the graph to get infected at a potentially different (including zero) exponential rate at each instant, with the restriction that the sum of all these *mobile* infection rates is bounded and does not scale with the network size. This model is quite general in terms of the strategies that the virtually mobile agent(s) can employ – it permits the virtually mobile agent(s) to focus all their infection efforts on a single node, or on a sub-collection of nodes based on geography/topology and/or infection state, or indeed try to infect all nodes in the network but at very small infection rates (the only constraint we have is that the sum of the infection rates is bounded). Thus, our model fully incorporates *adversarial* and *omniscient* aspects of virtually mobile/long-range spreading. Throughout the paper, the main metric we study is the (random) time taken for an infection to spread to all the nodes of a network.

In this setting, our results are somewhat surprising – in spite of the adversarial "power" the virtually mobile agents have for choosing the infection sites, we show that for commonly studied topologies such as grids and spatial random graphs, it turns out that a simple random strategy where the virulence from the virtual mobility is concentrated on a single node at a time (i.e., infect only one randomly chosen node at a time) is order-optimal. More formally, our main contributions in this paper are as follows:

(a) We develop general *upper bounds* on the order of the infection time for large graphs – both in expectation and with high probability – when the virtual mobility infection pattern is simply *random*, i.e. when every node is susceptible to virtually mobile infection at the same constant rate irrespective of other factors such as the present infection-state of the graph. The bounds are based on the extent to which the graph can be partitioned into pieces with appropriate diameter/conductance.

(b) For certain common classes of structured and random graphs like rings/line graphs, d-dimensional grid graphs

and random geometric graphs in the connected regime, we use the theory of first-passage percolation [15] to derive *lower bounds* on the order of infection times – again, both in expectation and with high probability – over *all* (possibly state-aware) virtual mobility patterns. These lower bounds are shown to match the upper bounds on infection time for random virtual mobility up to a logarithmic factor, showing that *random virtual mobility suffices to spread infection/information at the fastest possible (order-wise) rates* through such graphs that are sparse, geographically constrained and locally similar.

B. Related Work

Prior work concerning network spread, though diverse in scope and treatment, does not address the impact of adversarial, virtual mobility-assisted spreading in networks. Moreover, to the best of our knowledge, it lacks a consistent analytical framework in which the effects of different forms of virtual mobility on spreading time can be compared. There has been much work in studying the static spread of infection/innovation using various notions of influence and susceptibility, both numerically using field data/extensive simulations [8], [3], [4], [5] and analytically [9], [14], [15], [16]. For the case of spreading with virtual mobility, many numerical studies have investigated the spread of infectious diseases with specific mobility patterns, *e.g.* via airline networks [23], heterogeneous geographic means [1], [2], and recently, electronic pathways [21], [17], [18], [19], [20]. Several notable works in communication engineering include studies in which all network nodes are simultaneously physically mobile – for designing gossip algorithms [13], [6] and improving the capacity of wireless networks [12] – and analyses of rumour spreading on fully-connected graphs [10], [11]. Other design-oriented studies include investigations of optimal seeding in networks for maximum spread from a computational perspective [7], and efficient routing over spatial networks with fixed long-range links [24]. We refer the reader to [25] for more references.

II. MODEL FOR SPREADING WITH VIRTUAL MOBILITY

Consider a sequence of graphs $G_n = (V_n, E_n)$ indexed by n, with the n-th graph having n nodes. For instance, G_n could be the ring graph with n nodes, or a (2-dimensional) $\sqrt{n} \times \sqrt{n}$ grid. For convenience, we will often drop the subscript n for all quantities pertaining to the graph G_n when the context is clear.

We model the spread of an infection on the graph G_n (or G) using a continuous-time *spreading process* $(S(t))_{t \geq 0}$. At each time t, $S(t) = (S_1(t), \ldots, S_n(t)) \in \{0, 1\}^V$ denotes the "infection state" of the nodes in V: $S_i(t) = 0$ (resp. $S_i(t) = 1$) indicates that node i is "healthy" (resp. "infected") at time t. Let us denote by $\mathcal{N}(S(t))$ the number of infected nodes at time t, *i.e.* $\mathcal{N}(S(t)) \triangleq |\{i \in V : S_i(t) = 1\}|$. The evolution of $S(t)$ is assumed to be driven by the following modes of infection spread (for ease of notation, we label the nodes in V from 1 to n):

- *Static spread:* Initially, at $t = 0$, all nodes are healthy, except for a single node (node 1) which is infected. Once any node is infected, it attempts to infect each of its neighbouring healthy nodes at an exponential rate of β, *i.e.* the time taken for the infection to spread from that node to a neighbour is an independent Exponential(β) random variable. We call this form of infection spread as *basic* or *static* spread.
- *Mobile spread:* We assume that a "virtually mobile agent" external to the network G is capable of infecting (healthy) nodes at all times. More precisely, at each time t, each of the healthy nodes is susceptible to infection at an exponential rate which can depend on the state of the network $S(t)$. Moreover, as a reasonable limit on the power of the virtual mobility to spread infection (*i.e.* the infection virulence), we stipulate that the sum of the (exponential) rates of infection of the healthy nodes via this virtually mobile agent does not exceed a constant μ, say $\mu = 1$. We call this form of infection spread as *virtually mobile* spread.

Remark 1: These definitions of basic and virtually mobile spread mathematically model the two modes of static/local spread and long-range virtual mobility-based spread, as discussed in the introduction, respectively.

Remark 2: In this work, we set both β and the virulence μ to be fixed constants independent of network size n, time t or network state; however, in general, these quantities could potentially vary with n, t or network state to capture various disparities and dependencies.

Our model of virtually mobile spread naturally allows for a variety of *spreading policies* for virtual mobility. A spreading policy π specifies the exponential rates of mobile infection due to virtual mobility (or a virtually mobile agent) for all nodes at all times, which may in general depend on the current network state. Thus, π can be thought of as a map from the set of all network states $\{0,1\}^V$ to the set of rates for nodes: $\{(r_1, \ldots, r_n) : r_i \geq 0, \sum_i r_i \leq \mu\}$. Thus, under a spreading policy π, node i is susceptible to infection at a *total* exponential rate of $\beta \times |\{j \in V : S_j(t) = 1, (i,j) \in E\}| + \pi(S(t))_i$. As examples, consider the following spreading policies:

- The *purely static* or *no-virtually-mobile* spreading policy π_0 which always maps into the all-zeros vector (*i.e.* as if there is no virtual/long-range mobility in spreading), and
- The *random spreading* policy π_r which always maps into the all-ones vector scaled by $1/n$ (*i.e.* as if a virtually mobile agent picks a node in the network uniformly at random to infect).

Remark 3: As a natural extension, our framework of "virtually mobile" spreading policies can capture static (local) spreading on the popular *randomly-rewired networks* introduced by Watts-Strogatz [26] and Kleinberg [24] as long-range connectivity models for routing on social networks. Broadly, such graphs are constructed by taking (structured) ring or grid graphs and randomly adding long-range links between nodes. If we define a spreading policy on the original ring or grid

graph that lets infection spread through exactly these long-range links (the virulence μ now depends on the network state), then this form of virtual mobility on the ring/grid graph is the same as static spreading on the randomly rewired version of the graph.

Let Π denote the set of all possible spreading policies. For each policy $\pi \in \Pi$, define the random variable $T_\pi \overset{\triangle}{=} \inf\{t \geq 0 : S(t) = \mathbb{1}_n\}$ to be the *finish time* (or *spreading time*) for the policy π, *i.e.* the time at which all nodes in V get infected. Our primary focus throughout the paper is to characterize the order of the finish time with random virtual mobility (*i.e.* T_{π_r}), and to compare it with the best possible finish time over all spreading policies ($\inf_{\pi \in \Pi} \mathbb{E}[T_\pi]$) on certain classes of graphs.

III. UPPER BOUNDS ON FINISH TIME FOR SPREADING WITH RANDOM VIRTUAL MOBILITY

We present two main results here – upper bounds on the finish time of the random spread policy for a general graph G both in expectation and with high probability.

A. Diameter-based Bound

The result states that if G can be broken into a (large) number of uniformly-sized pieces, then the time taken by random spreading to finish is of the order of the number of pieces or the piece size, whichever dominates. This result will be our chief tool in the subsequent sections for estimating finish times for various types of graphs.

Theorem 1 (Time taken by random spread: Diameter version)**.** *Suppose that for each n, the graph G_n admits a partition $G_n = \bigcup_{i=1}^{g(n)} G_{n,i}$ by $g(n)$ connected subgraphs $G_{n,i}$, each with size $\Theta(s(n))$ and diameter $O(d(n))$. Then,*
  - (a) *(Mean finish time)* $\mathbb{E}[T_{\pi_r}] = O(h(n) \log n)$, *where $h(n) \equiv \max(g(n), d(n))$.*
  - (b) *(Finish time concentration) If $g(n) = \Omega(n^\delta)$ for some $\delta > 0$, then for any $\gamma > 0$ there exists $\alpha = \alpha(\gamma) > 0$ such that $\mathbb{P}[T_{\pi_r} \geq \alpha h(n) \log n] = O(n^{-\gamma})$.*

In other words, we can imagine spreading assisted by random virtual mobility on large graphs to be dominated by both (a) the time taken for spread to start in each piece or area of the graph and (b) the worst possible time taken within each piece for infection to spread statically.

Proof: The proof is using stochastic dominance and graph partitioning into suitable *shortest-path spanning trees*.

Let $(S(t))_{t \geq 0}$ (on the state space $\{0,1\}^V$) denote the spreading process driven by the random spread strategy π_r. Note that in $S(\cdot)$, the net exponential rate (say r_i) of each subgraph $G_{n,i} \equiv G_i$ being infected by virtual mobile spread is $r_i = \Theta(s(n)/n) > 0$. We define an associated 'slowed-down' virtual-mobile-assisted spreading process $(\tilde{S}(t))_{t \geq 0}$ as follows:

- *Phase 1:* The virtual mobile proceeds as usual, attempting to infect each node with an exponential rate of μ/n, *until at least one node in each subgraph G_i has been infected.* Let T_1 be the first time at which all the G_i have been

infected thus. There is *no virtual mobile spread* after time T_1.

- *Phase 2:* At time T_1, for each subgraph G_i, only the *first node* in G_i, say N_i, infected in Phase 1 is assumed to be infected, and all other nodes in G_i are considered healthy, even if some of them were infected on Phase 1 after N_i. The process $\tilde{S}(\cdot)$ proceeds from time T_1 by the usual static spread dynamics *within* each G_i, *i.e.* with the caveat that infection *does not spread* across edges linking different subgraphs. Denote by T_2 the additional time taken (since T_1) for all nodes to get infected.

A standard coupling argument establishes that $\mathcal{N}(S(t))$ stochastically dominates $\mathcal{N}(\tilde{S}(t))$ at all times t. Thus, the finish time for $\tilde{S}(\cdot)$ stochastically dominates that of $S(\cdot)$, *i.e.*

$$T_{\pi_r} \leq_{st} T_1 + T_2. \tag{1}$$

We next estimate the means of T_1 and T_2 and their tail probabilities to finish the proof. The analysis for T_1 follows a coupon-collecting argument: memorylessness of the exponential distribution implies that T_1 is distributed as the maximum of $g(n)$ *i.i.d.* exponential random variables with parameter $\Theta(s(n)/n) = \Theta(1/g(n))$. Hence, using a well-known result about the expectation of the maximum of *i.i.d.* exponentials, we obtain

$$\mathbb{E}[T_1] = O\left(\frac{H_{g(n)}}{1/g(n)}\right) = O(g(n) \log g(n)), \tag{2}$$

where $H_k \triangleq \sum_{i=1}^{k} i^{-1} = O(\log k)$ is the kth harmonic number. Also, by a union bound over the tails of $g(n)$ *i.i.d.* exponential random variables, for any $\kappa > 0$ we can estimate the tail of T_1:

$$\mathbb{P}[T_1 \geq \kappa g(n) \log g(n)] \leq g(n) e^{-(\Theta(1/g(n))\kappa g(n) \log g(n))}$$
$$= g(n)^{-\Theta(\kappa)+1}. \tag{3}$$

To estimate the statistics of T_2, we further consider the following 'slower'' mode of (static) spreading in phase 2: for each subgraph G_i (with diameter $O(d(n))$), let W_i be a *shortest-path spanning tree* of G_i rooted at the node N_i which is infected in phase 1. Such a tree has diameter $O(d(n)))$ and can in principle be obtained by performing a Breadth-First Search (BFS) on G_i starting at N_i. If we now insist that the phase-2 static infection process in G_i spreads *only via the edges* of W_i, then again, a standard coupling can be used to show that the time $\tilde{T}_2$ when all nodes in G get infected thus stochastically dominates T_2.

We will need the following simple lemma for the remainder of the proof:

Lemma 1. *For real numbers* a_{ij}, $1 \leq i \leq m$, $1 \leq j \leq n$, $\max_{i=1}^{m} \sum_{j=1}^{n} a_{ij} \leq \sum_{j=1}^{n} \max_{i=1}^{m} a_{ij}$.

For each tree W_i, let its leaves be labelled $N_{i1}, \ldots, N_{il(i)}$. Each leaf N_{ij} has a unique path p_{ij} starting from N_i to itself, of length $O(d(n))$. Let $\hat{T}_{jk}$ be the time taken for the infection to spread across the kth edge on this path p_{ij}, *i.e.* the (exponentially distributed) interval between the time when the

$(k-1)$th node on the path is infected up to the time when the kth node is infected. Then, the time $\hat{T}_{2,i}$ taken for all nodes in W_i (or G_i) to get infected can be upper-bounded by using Lemma 1:

$$\hat{T}_{2,i} = \max_{j=1}^{l(i)} \sum_{k=1}^{|p_{ij}|} \hat{T}_{jk} \leq \sum_{k=1}^{O(d(n))} \left(\max_{j=1}^{l(i)} \hat{T}_{jk} \right),$$

and a further application of the lemma bounds the phase-2 finish time $\hat{T}_2 = \max_{i=1}^{g(n)} \hat{T}_{2,i}$ as

$$\hat{T}_2 \leq \max_{i=1}^{g(n)} \sum_{k=1}^{O(d(n))} \left(\max_{j=1}^{l(i)} \hat{T}_{jk} \right) \leq \sum_{k=1}^{O(d(n))} \left(\max_{i=1}^{g(n)} \max_{j=1}^{l(i)} \hat{T}_{jk} \right).$$

The term in brackets is simply the maximum of the infection spread times across all stage-k edges of all the trees W_i within G. Hence, it is stochastically bounded above by the maximum of n i.i.d Exponential(β) random variables (say $Z_1, \ldots, Z_n$), using which we can write

$$\mathbb{E}[T_2] \leq \mathbb{E}[\hat{T}_2] \leq \sum_{k=1}^{O(d(n))} O(H_n/\beta) = O(d(n) \log n). \tag{4}$$

Again, using the union bound to estimate the tail probability of T_2, we have, for any $\kappa > 0$,

$$\mathbb{P}[T_2 \geq \kappa d(n) \log n] \leq \mathbb{P}[\hat{T}_2 \geq \kappa d(n) \log n]$$
$$\leq O(d(n))\mathbb{P}[Z_1 \geq \kappa \log n] \leq n \cdot n e^{-\beta \kappa \log n} = n^{-\beta \kappa + 2}. \tag{5}$$

We now have all the pieces required for the proof. Combining (1), (2) and (4), along with the fact that $g(n) = O(n)$, proves the first part of the theorem. For the second part, the hypothesis that $g(n) = \Omega(n^\delta)$, together with (3), gives

$$\mathbb{P}[T_1 \geq \kappa h(n) \log n] \leq \mathbb{P}[T_1 \geq \kappa g(n) \log g(n)] \leq n^{-\delta \Theta(\kappa)+\delta},$$

which, together with (1) and (5), gives

$$\mathbb{P}[T_{\pi_r} \geq 2\kappa h(n) \log n] \leq \mathbb{P}[T_1 + T_2 \geq 2\kappa h(n) \log n]$$
$$\leq n^{-\delta \Theta(\kappa)+\delta} + n^{-\beta \kappa + 2} \leq 2n^{-\min\{\delta(\Theta(\kappa)-1), \beta \kappa - 2\}}.$$

Choosing κ such that $\min\{\delta(\Theta(\kappa) - 1), \mu\kappa - 2\} \geq \gamma$ now yields the promised bound in the second part of the theorem. ∎

Remark: The factor of $\log n$ stated in the theorem actually appears only due to T_1; a more refined analysis of the phase-2 time T_2 shows that T_2 is order-wise $d(n)$ (in expectation and *w.h.p.*), the analysis is omitted for lack of space but can be found in [25]. An important implication thus is that if a spreading policy infects the subgraphs G_i *sequentially* (instead of randomly as with π_r), then the finish time is $O(h(n))$ in expectation and *w.h.p.*

B. Conductance-based Bound

As with the diameter, we can also bound the finish time with random virtual mobility in terms of a different structural property intimately related to spreading ability in graphs – the *conductance* (also called the *isoperimetric constant*). The

conductance $\Psi(G)$ of a graph $G = (V, E)$ is defined as $\Psi(G) \triangleq \inf_{S \subset V: 1 \leq S \leq \frac{|V|}{2}} \frac{E(S, V \setminus S)}{|S|}$, where for $A, B \subseteq V$, $E(A, B)$ denotes the number of edges that have exactly one endpoint each in A and B. The conductance of a graph is a widely studied measure of how fast a random walk on the graph converges to stationarity [27]; the higher the conductance, the lesser 'bottlenecks' it offers for spreading.

Analogous to Theorem 1, the next result formalizes the idea that spreading on a graph is dominated by the larger of (a) the number of pieces it can be broken into, and (b) the reciprocal of the piece conductance. We refer the reader to [25] for proof.

Theorem 2 (Time taken by random spread: Conductance version). *Suppose that for each n, the graph G_n admits a partition $G_n = \bigcup_{i=1}^{g(n)} G_{n,i}$ by $g(n)$ connected subgraphs $G_{n,i}$, each with size $\Theta(s(n))$ and conductance $\Theta(\Psi(n))$. Then,*

(a) *(Mean finish time)* $\mathbb{E}[T_{\pi_r}] = O(k(n) \log g(n))$, *where* $k(n) \equiv \max\left(g(n), \frac{\log s(n)}{\Psi(n)}\right)$.

(b) *(Finish time concentration) There exists $\kappa > 0$ independent of n such that*

$$\mathbb{P}[T_{\pi_r} \geq \kappa k(n) \log g(n)] = O\left((\log g(n))^{-2}\right).$$

IV. LOWER BOUNDS ON FINISH TIME OVER ALL ADVERSARIAL VIRTUAL MOBILITY STRATEGIES

In the previous section, we have estimated the time that random virtual mobility takes to infect all nodes in a network. A natural question at this point is: How does the time taken by random virtual mobility compare with the best (*i.e.* lowest) possible spreading time among all other spreading strategies? In this section, we show that for certain commonly studied *spatially limited* networks such as line/ring networks, d-dimensional grids and random geometric graphs, *random spreading yields the best order-wise time* (up to a logarithmic factor) to spread infection, even among virtual mobile strategies that could use the state of the network to decide their virtual mobility patterns. In particular, for each of these classes of graphs, we establish lower bounds on the finish time of *any* spreading strategy that match the upper bounds established in the previous section, thus demonstrating the finish-time optimality of random spreading.

A. Ring/Linear Graphs

For each n, let $G_n = (V_n, E_n)$ be the ring graph with n contiguous vertices $V_n \triangleq \{v_1, \ldots, v_n\}$, $E_n \triangleq \{(v_i, v_j) : j - i \equiv 1 \pmod{n}\}$. In the context of Theorem 1, let us partition G_n into $\sqrt{n}$ successive $\sqrt{n}$-sized segments, *i.e.* $G_{n,i}$ is the subgraph induced by $v_{(i-1)\sqrt{n}+1}, \ldots, v_{i\sqrt{n}}$, where i ranges from $1, \ldots, \sqrt{n}$. The diameter of each segment is $\sqrt{n}$, and a straightforward application of the theorem gives

Corollary 1 (Time for random spread on ring graphs). *For the random spread policy π_r on the ring/line graph G_n,*

(a) $\mathbb{E}[T_{\pi_r}] = O(\sqrt{n} \log n)$,

(b) *For any $\gamma > 0$ $\exists \alpha = \alpha(\gamma) > 0$ such that $\mathbb{P}[T_{\pi_r} \geq \alpha \sqrt{n} \log n] = O(n^{-\gamma})$.*
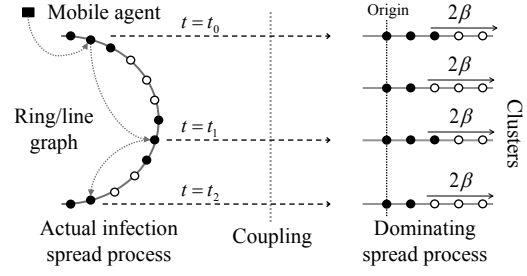

Fig. 1. The line graph: Coupling infection spreading with virtual mobility to a dominating 'cluster-growth' process

i.e., the finish time on an n-ring, with random virtual mobility, is $O(\sqrt{n} \log n)$ in expectation and with high probability.

Our next main result is to demonstrate that the finish time on a grid or line graph with *any* (possibly infection-state aware) virtual mobility spread strategy must be $\Omega(\sqrt{n})$, both in expectation and with high probability. This establishes that for ring graphs (or 1-dimensional grids), random virtual mobility is as good as any other form of controlled virtual mobility in an order-wise (up to a logarithm) sense.

Theorem 3 (Lower Bound on Finish Time for Ring Graphs). *For the ring graph G_n with n nodes, there exists $c > 0$ independent of n such that for any spreading policy π,*

$$\mathbb{P}\left[T_\pi < c\sqrt{n}\right] = O\left(e^{-\Theta(1)\sqrt{n}}\right).$$

Moreover, $\inf_{\pi \in \Pi} \mathbb{E}[T_\pi] = \Omega(\sqrt{n})$.

Proof: Along with the spreading process $(S^\pi(t))_{t \geq 0}$ induced by the policy π, consider a random process $(\tilde{S}(t))_{t \geq 0}$ described as follows:

(a) At all times t, $\tilde{S}(t)$ consists of an integer number $(\tilde{C}_t)$ of sets of points called *clusters*, where $(\tilde{C}_t)_{t \geq 0}$ is a Poisson process with intensity $\mu = 1$, and $\tilde{C}_0 = 1$ (the 1 denotes an 'initial' cluster in which static infection starts spreading).

(b) Once a new cluster is formed at some time s, it *grows*, *i.e.* adds points, following a Poisson process of intensity 2β.

Via a standard coupling argument, it can be shown that for *all* spreading strategies $\pi \in \Pi$, at all times $t \geq 0$, the total number of points in $\tilde{S}(t)$ (denoted by $\tilde{N}_t$) stochastically dominates that in $S^\pi(t)$. Essentially, this is due to two reasons: first, that the rate of 'seeding' of new clusters by π is at most as fast as that in $\tilde{S}(\cdot)$; secondly, each cluster in $\tilde{S}(\cdot)$ grows independently and without interference from other existing clusters, as opposed to clusters that could 'merge' in the process $S^\pi(\cdot)$. Figure 1 graphically depicts the structure of the dominating process $\tilde{S}(\cdot)$. Let $\tilde{T} \triangleq \inf\{t \geq 0 : \tilde{N}_t = n\}$ be the time when the number of points in $\tilde{S}(\cdot)$ first hits n. Owing to the stochastic dominance $\mathcal{N}(S^\pi(t)) \leq_{st} \tilde{N}_t$, we have that

$$\tilde{T} \leq_{st} T_\pi \ \forall \pi \in \Pi. \tag{6}$$

Knowing the way $\tilde{S}(\cdot)$ evolves, we can calculate $E[\tilde{N}_t]$:

$$\mathbb{E}[\tilde{N}_t] = \mathbb{E}[\mathbb{E}[\tilde{N}_t|\tilde{C}_t]] = \sum_{k=0}^{\infty} \frac{e^{-t}t^k}{k!} \mathbb{E}[\tilde{N}_t|\tilde{C}_t = k].$$

Since $\tilde{C}_t$ is a Poisson process, conditioned on $\{\tilde{C}_t = k\}$, the k cluster-creation instants are distributed uniformly on $[0, t]$. Let the times of these arrivals be $\tilde{T}_1, \ldots, \tilde{T}_k$; then $[\tilde{T}_i, t]$ is the time for which the ith cluster has been growing. Since every cluster grows at a rate of 2β, conditioned on $\{\tilde{C}_t = k\}$, the expected size of the ith cluster is $2\beta(t - \tilde{T}_i)$, $1 \le i \le k$. Also, the expected size of the '0-th' cluster at time t is $2\beta t$. Using $\mathbb{E}[\tilde{T}_i|\tilde{C}_t = k] = t/2$, we obtain $\mathbb{E}[\tilde{N}_t|\tilde{C}_t = k] = 2\beta t + \sum_{i=1}^{k} \mathbb{E}[2\beta(t - \tilde{T}_i)|\tilde{C}_t = k] = \beta(k + 2)t$, thus

$$\mathbb{E}[\tilde{N}_t] = \sum_{k=0}^{\infty} \frac{e^{-t}t^k}{k!} \beta(k+2)t = \beta t^2 + 2\beta t.$$

$$\Rightarrow P(\tilde{T} > t) = P(\tilde{N}_t < n) = 1 - P(\tilde{N}_t \ge n)$$
$$\ge 1 - \frac{\mathbb{E}[\tilde{N}_t]}{n} \ge 1 - \frac{\beta(t+1)^2}{n},$$

by an application of Markov's inequality. This means that

$$\mathbb{E}[\tilde{T}] = \int_0^{\infty} \mathbb{P}(\tilde{T} > x)dx \ge \int_0^{\sqrt{\frac{n}{\beta}}-1} \mathbb{P}(\tilde{T} > x)dx$$
$$\ge \int_0^{\sqrt{\frac{n}{\beta}}-1} \left(1 - \frac{\beta(x+1)^2}{n}\right) dx = \Theta(\sqrt{n}).$$

Together with (6), this forces $\inf_{\pi \in \Pi} \mathbb{E}[T_\pi] = \Omega(\sqrt{n})$, and the first part of the theorem is proved. For the second part, if we denote the size of the ith created cluster at time $s \ge T_i$ by $\tilde{X}_i(s)$, then for any time t we can write

$$\left(\bigcap_{i=0}^{2et} \{\tilde{X}_i(t + T_i) < 4e\beta t\}\right) \bigcap \{\tilde{C}_t < 2et\}$$
$$\subseteq \left(\bigcap_{i=0}^{\tilde{C}(t)} \{\tilde{X}_i(t + T_i) < 4e\beta t\}\right) \bigcap \{\tilde{C}_t < 2et\}$$
$$\subseteq \left(\bigcap_{i=0}^{\tilde{C}(t)} \{\tilde{X}_i(t) < 4e\beta t\}\right) \bigcap \{\tilde{C}_t < 2et\} \subseteq \{\tilde{N}_t < 8\beta e^2 t^2\}.$$

Applying a standard Chernoff bound ($\mathbb{P}[Y \ge 2e\lambda] \le (2e)^{-\lambda}$ for $Y \sim \text{Poisson}(\lambda)$) to $\tilde{C}_t \sim \text{Poisson}(t)$ and $\tilde{X}_i(t + T_i) \sim \text{Poisson}(2\beta t)$ above, we can write

$$\mathbb{P}[\tilde{N}_t \ge 8\beta e^2 t^2] \le \mathbb{P}[\tilde{C}_t \ge 2et] + \sum_{i=1}^{2et} \mathbb{P}[\tilde{X}_i(t + T_i) \ge 4e\beta t]$$
$$\le (2e)^{-t} + 2et \cdot (2e)^{-2\beta t} = O(e^{-t(1 \wedge 2\beta)}).$$

The proof is concluded using the stochastic dominance (6):

$$\mathbb{P}\left[T_\pi < \sqrt{\frac{n}{8\beta e^2}}\right] \le \mathbb{P}\left[\tilde{T} < \sqrt{\frac{n}{8\beta e^2}}\right]$$
$$= \mathbb{P}\left[\tilde{N}_{\sqrt{\frac{n}{8\beta e^2}}} > n\right] = O\left(e^{-\Theta(1)\sqrt{n}}\right).$$
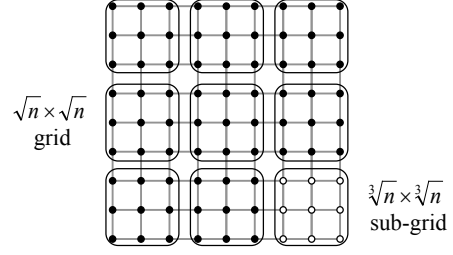
∎

Fig. 2. A planar grid: Tiling into sub-grids

B. d-Dimensional Grid Graphs

This section shows that the simple, state-oblivious random virtual mobility spreading strategy achieves the optimal order-wise finish time even on d-dimensional grid networks where $d \ge 2$. For such a dimension d, the d-dimensional grid graph $G_n = (V_n, E_n)$ on n nodes is given by $V_n \triangleq \{1, 2, \ldots, n^{1/d}\}^d$, and $E_n \triangleq \{(x, y) \in V_n \times V_n : ||x - y||_1 = 1\}$ (throughout, for any l we assume $n^{1/l}$ to be integer to avoid cumbersome notation).

Consider a partition of G_n into $n^{1/(d+1)}$ identical and contiguous 'sub-grids' G_{ni}, $i = 1, \ldots, n^{1/(d+1)}$. By this, we mean that each G_{ni} is induced by a copy of $\{1, 2, \ldots, n^{1/(d+1)}\}^d$ (and thus has $n^{d/(d+1)}$ nodes). For instance, in the case of a planar $\sqrt{n} \times \sqrt{n}$ grid, imagine tiling it horizontally and vertically with $\sqrt[3]{n}$ identical $\sqrt[3]{n} \times \sqrt[3]{n}$ sub-grids (Figure 2). With such a partition, an application of Theorem 1 shows:

Corollary 2 (Time for random spread on d-grids). *For the random spread policy π_r on an n-node d-dimensional grid G_n,*

   (a) $\mathbb{E}[T_{\pi_r}] = O\left(n^{1/(d+1)} \log n\right)$,
   (b) *For any $\gamma > 0$ there exists $\alpha = \alpha(\gamma) > 0$ such that $\mathbb{P}[T_{\pi_r} \ge \alpha n^{1/(d+1)} \log n] = O(n^{-\gamma})$.*

i.e., the finish time with random virtual mobility on a d-dimensional n-node grid is $O\left(n^{1/(d+1)} \log n\right)$ in expectation and with high probability.

In what follows, we show that *any* virtual mobile spreading policy on a grid must take time $\Omega(n^{1/(d+1)})$ to finish infecting all nodes with high probability, and consequently also in expectation. Barring a logarithmic factor, this shows that random virtual mobility is as good as any other (possibly state-aware) spreading policy on the class of grids.

Theorem 4 (Lower bound on Finish Time for d-grids). *Let G_n be the symmetric d-dimensional grid graph with n nodes. Then, there exists $c = c(d) > 0$ not depending on n such that for any spreading policy π,*

$$\mathbb{P}\left[T_\pi \le cn^{1/(d+1)}\right] = O\left(e^{-\Theta(1)\cdot n^{1/(2d+2)}}\right).$$

Moreover, $\inf_{\pi \in \Pi} \mathbb{E}[T_\pi] = \Omega\left(n^{1/(d+1)}\right)$.

Proof: For the spreading policy π, we introduce a (dominating) spreading process $(\tilde{S}(t))_{t \ge 0}$ in which
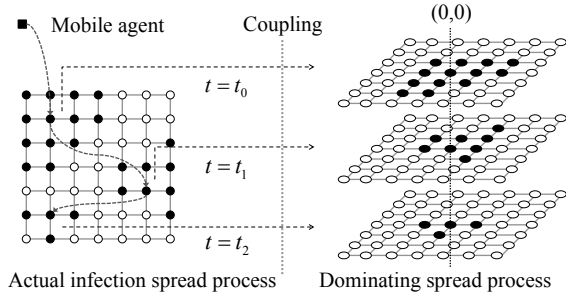
Fig. 3. The grid graph: Coupling infection spreading with virtual mobility to a dominating 'cluster-growth' process

- At all times t, $\tilde{S}(t)$ consists of an integer number ($\tilde{C}_t$) of sets of points called *clusters*, where $(\tilde{C}_t)_{t \geq 0}$ is a Poisson process with intensity $\mu = 1$, and $\tilde{C}_0 = 1$ (the 1 denotes an 'initial' cluster in which static infection starts spreading).
- Each cluster grows as an independent copy of a static infection process on an exclusive infinite d-dimensional grid $\mathbb{Z}^d$ starting at $(0, 0, \ldots, 0)$.

We note that this process is similar to the one considered in the proof of Theorem 3, but with the difference that the growth of each cluster follows the natural infection dynamics in a grid-structured graph. Again, a standard coupling argument shows that at all times $t \geq 0$, the total number of points in $\tilde{S}(t)$ (denoted by $\tilde{N}_t$) stochastically dominates that in $S^\pi(t)$; indeed, this is is due to "virtually mobile" cluster seeding at the highest possible exponential rate and the absence of colliding infections (Figure 3). As before, letting $\tilde{T} \overset{\triangle}{=} \inf\{t \geq 0 : \tilde{N}_t = n\}$ be the time when the number of points in $\tilde{S}(\cdot)$ first hits n, we record the stochastic dominance

$$\mathcal{N}(S^\pi(t)) \leq_{st} \tilde{N}_t \; \forall \pi \in \Pi \; \Rightarrow \; \tilde{T} \leq_{st} T_\pi \; \forall \pi \in \Pi. \quad (7)$$

We will need the following key lemma, due the theory of first-passage percolation [15], which essentially lets us control the extent to which infection on an infinite grid has spread at time t:

Lemma 2. Let $(\tilde{Z}(t))_{t \geq 0} \in \{0, 1\}^{\mathbb{Z}^d}$ represent a static/basic infection spread process on the infinite d-dimensional lattice $\mathbb{Z}^d$ starting at node $(0, 0, \ldots, 0)$ at time 0. Then, there exist positive constants l, c_1, c_2 such that for $t \geq 1$,

$$\mathbb{P}[\mathcal{N}(\tilde{Z}(t)) > t^d l^d] \leq c_1 t^{2d} e^{-c_2 \sqrt{t}}.$$

As in the proof of Theorem 3, denoting by $\tilde{X}_i(s)$ the size of the ith created cluster of $\tilde{S}(\cdot)$ at time $s \geq T_i$, we can write, for $t \geq 0$,

$$\left(\bigcap_{i=0}^{2et} \{\tilde{X}_i(t + T_i) < t^d l^d\} \right) \bigcap \{\tilde{C}_t < 2et\} \subseteq \{\tilde{N}_t < 2el^d t^{d+1}\}.$$

Each of the random variables $\tilde{X}_i(t + T_i)$ is distributed as the number of infected nodes in a static infection process on an infinite grid at time t. Thus, using Lemma 2 and the

aforementioned Chernoff bound for $\tilde{C}_t \sim \text{Poisson}(t)$, we can write

$$\mathbb{P}\left[\tilde{N}_t \geq (2el^d)t^{d+1} \right]$$
$$\leq \mathbb{P}\left[\tilde{C}_t \geq 2et \right] + \sum_{i=1}^{2et} \mathbb{P}\left[\tilde{X}_i(t + T_i) \geq t^d l^d \right]$$
$$\leq (2e)^{-t} + 2et \cdot c_1 t^{2d} e^{-c_2 \sqrt{t}} = O(e^{-c_2 \sqrt{t}}).$$

With the stochastic dominance (7), this forces

$$\mathbb{P}\left[T_\pi \leq \left(\frac{n}{2el^d} \right)^{1/(d+1)} \right] \leq \mathbb{P}\left[\tilde{T} \leq \left(\frac{n}{2el^d} \right)^{1/(d+1)} \right]$$
$$= \mathbb{P}\left[\tilde{N}_{\left(\frac{n}{2el^d} \right)^{1/(d+1)}} \geq n \right] = O\left(e^{-c_3 \cdot n^{1/(2d+2)}} \right), \quad (8)$$

establishing the first part of the theorem. To see how this implies the second part, note that by the probability estimate (8) and the Borel-Cantelli lemma,

$$\mathbb{P}\left[\tilde{T} \leq \left(\frac{n}{2el^d} \right)^{1/(d+1)} \text{ for finitely many } n \right] = 1,$$
$$\Rightarrow \liminf_{n \to \infty} \frac{\tilde{T}}{n^{1/(d+1)}} \geq c_4 \overset{\triangle}{=} \frac{1}{(2el^d)^{1/(d+1)}} > 0 \quad a.s.$$

By Fatou's lemma,

$$\liminf_{n \to \infty} \mathbb{E}\left[\frac{\tilde{T}\pi}{n^{1/(d+1)}} \right] \geq \mathbb{E}\left[\liminf_{n \to \infty} \frac{\tilde{T}}{n^{1/(d+1)}} \right] \geq c_4 > 0.$$

This shows that $\mathbb{E}[T_\pi] \geq \mathbb{E}[\tilde{T}] = \Omega\left(n^{1/(d+1)} \right)$ for any $\pi \in \Pi$, concluding the proof of the theorem. ∎

Proof of Lemma 2: Let

$$\tilde{B}(t) \overset{\triangle}{=} \{v \in \mathbb{Z}^d : \tilde{Z}_v(t) = 1\} \subset \mathbb{Z}^d \; (\subset \mathbb{R}^d)$$

be the set of infected nodes at time t in $\tilde{Z}$. We will use the following version of a result, from percolation on lattices with exponentially distributed edge passage times, about the 'typical shape' of $\tilde{B}(t)$ [15]:

(Theorem 2 in [15]) There exists a fixed (i.e. *not depending on* t) cube $B_0 = \left[-\frac{l}{2}, \frac{l}{2} \right]^d \subset \mathbb{R}^d$, and constants $c_1, c_2 > 0$, such that for $t \geq 1$,

$$\mathbb{P}\left[\tilde{B}(t) \subset tB_0 \right] \geq 1 - c_1 t^{2d} e^{-c_2 \sqrt{t}}. \quad (9)$$

It follows from (9) that for $t \geq 1$,

$$\mathbb{P}[\mathcal{N}(\tilde{Z}(t)) > t^d l^d] = \mathbb{P}[|\tilde{B}(t)| > t^d l^d]$$
$$\leq \mathbb{P}[\tilde{B}(t) \not\subseteq tB_0] \leq c_1 t^{2d} e^{-c_2 \sqrt{t}}.$$
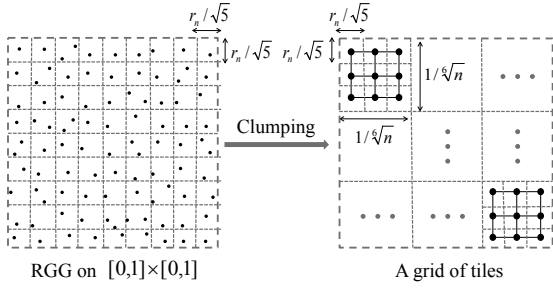
∎

Fig. 4. The random geometric graph: Clumping into chunks of small 'tiles'

C. Random Geometric Graphs

We turn to a popular random graph model for representing various physical networks – the *Random Geometric Graph (RGG)*. For simplicity we consider the planar version of the RGG, in which n points (*i.e.* nodes) are picked *i.i.d.* uniformly in the unit square $[0,1] \times [0,1]$. Two nodes x, y are connected by an edge if and only if $||x - y|| \leq r_n$, where r_n is often called the *coverage radius*. The RGG $G_n = G_n(r_n)$ consists of the n nodes and edges as above.

It is well-known that when the coverage radius r_n is above a critical threshold of $\sqrt{\frac{\log n}{\pi}}$, the RGG is connected with high probability [28]. In this section, we state and prove two results that show that random spreading on RGGs in this critical connectivity regime is as good as any other form of omniscient virtual mobility. First, we show with high probability that random spreading finishes in time $O(\sqrt[3]{n} \log n)$, and follow it up with a converse result that says that no other policy can better this order (up to the logarithmic factor) with significant probability. This directly parallels the earlier results about finish times on 2-dimensional grids, where random virtual-mobility-based spread exhibits the same optimal order of growth.

Theorem 5 (Upper bound on spreading time with random virtual mobility on RGGs)**.** *For the planar random geometric graph $G_n(r_n)$, if $r_n \geq \sqrt{\frac{5 \log n}{n}}$, then there exists $\alpha > 0$ such that $\lim_{n \to \infty} \mathbb{P}[T_{\pi_r} \geq \alpha \sqrt[3]{n} \log n] = 0$.*

Proof: Divide the unit square $[0,1] \times [0,1]$ into (row and column-wise) square *tiles* of side length $r_n / \sqrt{5}$ each; there are thus $5/r_n^2$ such tiles, say $k_1, \ldots, k_{5/r_n^2}$. If n points are thrown uniformly randomly into $[0,1] \times [0,1]$, then, with $\mathcal{E}$ denoting the event that some tile is empty, it can be shown that

$$\mathbb{P}[\mathcal{E}] \leq \frac{5}{r_n^2} \mathbb{P}[\text{tile 1 empty}] = \frac{5}{r_n^2}\left(1 - \frac{r_n^2}{5}\right)^n \leq \frac{1}{\log n}. \quad (10)$$

By construction, note that the maximum distance between points in two horizontally or vertically adjacent tiles is exactly r_n. Hence, two nodes in horizontally or vertically adjacent tiles are always connected by an edge. Also, a node in a tile is not connected to any node in a tile at least three tiles away in either dimension. If we now divide $[0,1] \times [0,1]$ into (bigger) square *chunks* of side length $1/\sqrt[6]{n}$ each, there are

$\sqrt[3]{n}$ such square chunks, each containing a $\frac{\sqrt{5}}{r_n \sqrt[6]{n}} \times \frac{\sqrt{5}}{r_n \sqrt[6]{n}}$ grid of square tiles (Figure 4). In the case where no tile is empty, it follows from the arguments in the preceding paragraph that the diameter of the subgraph induced within each chunk is $O\left(\frac{1/\sqrt[6]{n}}{r_n}\right) = O\left(\frac{1}{r_n \sqrt[6]{n}}\right) = O\left(\frac{\sqrt[3]{n}}{\sqrt{\log n}}\right)$ since $r_n \geq \sqrt{\frac{5 \log n}{n}}$. Also, since we are interested in upper-bounding T_{π_r}, we disallow the spread of infection between (possibly connected) nodes in all pairs of diagonally adjacent tiles by appealing to stochastic dominance as before. An application of Theorem 1 now shows that $\mathbb{E}[T_{\pi_r} | \mathcal{E}] = O(\sqrt[3]{n} \log n)$, and for some $\alpha, \gamma > 0$, $\mathbb{P}[T_{\pi_r} \geq \alpha \sqrt[3]{n} \log n \mid \mathcal{E}] = O(n^{-\gamma})$. Using (10), we conclude that

$$\mathbb{P}[T_{\pi_r} \geq \alpha \sqrt[3]{n} \log n] = O\left(1/\log n\right) \overset{n \to \infty}{\longrightarrow} 0.$$

$\blacksquare$

Towards a lower bound on the spreading time on a RGG over *all* spreading policies, consider an infinite planar grid with *additional one-hop diagonal edges*, i.e. $G = (V, E)$ where $V = \mathbb{Z}^2$, $E = \{(x, y) \in \mathbb{Z}^2 : ||x - y||_\infty \leq 1\}$. Let an infection process $(S(t))_{t \geq 0}$ start from $0 \in \mathbb{Z}^2$ at time 0 according to the standard static spread dynamics, i.e. with each edge propagating infection at an exponential rate μ, and let $I(t)$ denote the set of infected nodes at time t. The following key lemma helps control the size of $I(t)$, *i.e.* the extent of infection at time t:

Lemma 3 (Extent of infection at time t on grid with diagonal edges)**.** *There exists $c_1 > 0$ such that for any $c_2 > 0$ and t large enough, $\mathbb{P}[\exists x \in I(t) : ||x||_\infty \geq (c_1 \mu + c_2)t] = O\left((c_1 \mu + c_2)t \cdot e^{-c_2 t}\right)$.*

The reader is referred to [25] for the proof. Using Lemma 3 in a manner analogous to that used in proving Theorem 4 for the d-dimensional grid, we finally have

Theorem 6 (Lower Bound on Finish Time for RGGs)**.** *For the planar random geometric graph G_n with $r_n = O(\sqrt{\log n / n})$ and any spreading policy π, $\exists \beta > 0$ such that $\lim_{n \to \infty} \mathbb{P}\left[T_\pi \geq \beta \frac{\sqrt[3]{n}}{\log^{4/3} n}\right] = 1$.*

We omit the detailed proof here due to space limitations. The interested reader can refer to [25] for details.

Proof idea: Divide the unit square $[0,1] \times [0,1]$ row and column-wise into $r_n \times r_n$ *tiles*; there are thus $1/r_n^2$ such tiles, say $k_1, \ldots, k_{1/r_n^2}$. By standard balls-and-bins arguments, with the n nodes thrown randomly into these $n/\log n$ tiles, each tile receives a maximum of $O(\log n)$ nodes with high probability. Under this event, take each tile to be the vertex of a square grid where adjacent diagonals are connected. Also, set the rate of infection spread on every edge of this grid to be $\text{Exp}(\mu \log^2 n)$. This effectively upper-bounds the best rate of spread among neighbouring tiles, and using Lemma 3 to control the maximum rate of spread we can lower bound T_π for any π.

$\blacksquare$

V. Conclusion

We have modelled the spread of infection on networks with static and virtual-mobility-assisted spreading components. For

general graphs, we have bounded above the time taken to spread infection with random long-range virtual mobility across the whole graph. In the cases of common "spatially constrained" graphs like lines, grids and random geometric graphs, we have also established lower bounds on the time taken to spread infection across all adversarial forms of long-range spreading. These bounds match the upper bounds for random virtual mobility up to a logarithmic factor, showing that random long-range spreading is order-wise optimal for such graphs.

Future work involves – *(i)* extending the analysis to other classes of graphs; we conjecture that all degree-wise sparse, spatially-induced graphs exhibit optimal infection spreading time with merely random virtual mobility, whereas expanders and other low-diameter graphs do not require any long-range virtual mobility to improve the spreading time; *(ii)* investigating the design of reasonably simple *non-random* virtual mobile spreading policies that are spreading time-optimal without the logarithmic factor – we hope this will help provide insights into optimal seeding for quick dissemination across networks.

REFERENCES

[1] F. G. Ball, "Stochastic multitype epidemics in a community of households: Estimation of threshold parameter R* and secure vaccination coverage," *Biometrika*, vol. 91, no. 2, pp. 345–362, 2004.

[2] B. Wang, K. Aihara, and B. J. Kim, "Immunization of geographical networks," in *Complex (2)*, pp. 2388–2395, 2009.

[3] E. M. Rogers, *Diffusion of Innovations, 5th Edition*. Free Press, original ed., August 2003.

[4] M. Granovetter, "The Strength of Weak Ties," *The American Journal of Sociology*, vol. 78, no. 6, pp. 1360–1380, 1973.

[5] J. O. Kephart and S. R. White, "Directed-graph epidemiological models of computer viruses," in *IEEE Symposium on Security and Privacy*, pp. 343–361, 1991.

[6] D. Kempe, J. Kleinberg, and A. Demers, "Spatial gossip and resource location protocols," *J. ACM*, vol. 51, no. 6, pp. 943–967, 2004.

[7] D. Kempe, J. Kleinberg, and E. Tardos, "Maximizing the spread of influence through a social network," in *KDD '03: Proc. 9th ACM SIGKDD International Conference on Knowledge Discovery and Data mining*, pp. 137–146, ACM, 2003.

[8] R. Pastor-Satorras and A. Vespignani, "Epidemic dynamics in finite size scale-free networks," *Phys. Rev. E*, vol. 65, p. 035108, Mar 2002.

[9] A. J. Ganesh, L. Massoulié, and D. F. Towsley, "The effect of network topology on the spread of epidemics," in *INFOCOM*, pp. 1455–1466, 2005.

[10] B. Pittel, "On spreading a rumor," *SIAM J. Appl. Math.*, vol. 47, no. 1, pp. 213–223, 1987.

[11] S. Sanghavi, B. Hajek, and L. Massoulié, "Gossiping with multiple messages," in *INFOCOM*, pp. 2135–2143, 2007.

[12] M. Grossglauser and D. Tse, "Mobility Increases the Capacity of Ad Hoc Wireless Networks," *IEEE/ACM Transactions on Networking*, vol. 10, no. 4, p. 477, 2002.

[13] A. D. Sarwate and A. G. Dimakis, "The impact of mobility on gossip algorithms," in *INFOCOM*, pp. 2088–2096, 2009.

[14] R. Durrett and X.-F. Liu, "The contact process on a finite set," *Annals of Probability*, vol. 16, pp. 1158–1173, 1988.

[15] H. Kesten, "On the speed of convergence in first-passage percolation," *The Annals of Applied Probability*, vol. 3, no. 2, pp. 296–338, 1993.

[16] F. Ball, D. Mollison, and G. Scalia-Tomba, "Epidemics with two levels of mixing.," *Ann. Appl. Probab.*, vol. 7, no. 1, pp. 46–89, 1997.

[17] P. Wang, M. C. González, C. A. Hidalgo, and A.-L. Barabasi, "Understanding the spreading patterns of mobile phone viruses," *Science*, vol. 324, pp. 1071–1076, May 2009.

[18] J. Cheng, S. H. Wong, H. Yang, and S. Lu, "Smartsiren: virus detection and alert for smartphones," in *MobiSys '07: Proceedings of the 5th international conference on Mobile systems, applications and services*, (New York, NY, USA), pp. 258–271, ACM, 2007.

[19] J. W. Mickens and B. D. Noble, "Modeling epidemic spreading in mobile environments," in *WiSe '05: Proceedings of the 4th ACM Workshop on Wireless Security*, (New York, NY, USA), pp. 77–86, ACM, 2005.

[20] J. Su, K. K. W. Chan, A. G. Miklas, K. Po, A. Akhavan, S. Saroiu, E. de Lara, and A. Goel, "A preliminary investigation of worm infections in a Bluetooth environment," in *WORM '06: Proceedings of the 4th ACM workshop on Recurring malcode*, (New York, NY, USA), pp. 9–16, ACM, 2006.

[21] J. Kleinberg, "The Wireless Epidemic," *Nature*, vol. 449, pp. 287–288, September 2007.

[22] E. Kaplan, D. L. Craft, and L. M. Wein, "Analyzing bioterror response logistics: the case of smallpox," *Math Biosci*, vol. 185, p. 33072, September 2003.

[23] V. Colizza, A. Barrat, M. Barthélemy, and A. Vespignani, "The role of the airline transportation network in the prediction and predictability of global epidemics," *Proceedings of the National Academy of Sciences*, vol. 103, no. 7, pp. 2015–2020, 2006.

[24] C. Martel and V. Nguyen, "Analyzing Kleinberg's (and other) small-world models," in *PODC '04: Proc. 23rd annual ACM symposium on Principles of Distributed Computing*, pp. 179–188, ACM, 2004.

[25] A. Gopalan, S. Banerjee, A. Das, and S. Shakkottai, "Random mobility and the spread of infection," tech. rep., The University of Texas at Austin, July 2010.

[26] D. J. Watts and S. H. Strogatz, "Collective dynamics of 'small-world' networks," *Nature*, vol. 393, pp. 440–442, June 1998.

[27] D. Shah, "Gossip algorithms," *Found. Trends Netw.*, vol. 3, no. 1, pp. 1–125, 2009.

[28] P. Gupta and P. R. Kumar, "Critical power for asymptotic connectivity in wireless networks," in *Stochastic Analysis, Control, Optimization and Applications: A Volume in Honor of W.H. Fleming* (W. M. McEneany, G. Yin, and Q. Zhang, eds.), pp. 547–566, Boston: Birkhauser, 1998.