Infection Spread in Wireless Networks with Random and Adversarial Node Mobilities

Yi Wang
Viterbi School of Engineering
University of Southern California
Los Angeles, CA 90089, USA
wangyi@usc.edu

Shyam Kapadia
Viterbi School of Engineering
University of Southern California
Los Angeles, CA 90089, USA
kapadia@usc.edu

Bhaskar Krishnamachari
Viterbi School of Engineering
University of Southern California
Los Angeles, CA 90089, USA
bkrishna@usc.edu

ABSTRACT

We study the process of the spread of an infection among mobile nodes moving on a finite, grid based map. A random walk and a novel adversarial model are considered as two extreme cases of node mobility. With N nodes, we present analytical and simulation results for both mobility models for a square grid map with size $\sqrt{G} \times \sqrt{G}$. A key finding is that with random mobility the total time to infect all nodes decreases with N while with an adversarial model we observe a reverse trend. Specifically, the random case results in a total infection time of $\Theta(\frac{G \log G \log N}{N})$ as opposed to the adversarial case where the total infection time is found to be $\Theta(\sqrt{G} \log N)$. We also explore the possibility of emulating such an infection process as a mobile interaction game with wireless sensor motes, and the above results are complimented by traces obtained from an empirical study with humans as players in an outdoor field.

Categories and Subject Descriptors

C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design—*Wireless communication*

General Terms

Performance

Keywords

infection spread, wireless network, node mobility, random, adversarial

1. INTRODUCTION

Data dissemination in mobile wireless ad hoc networks with a sparse density of nodes poses a variety of challenges. A popular and sometimes the only viable option for disseminating and gathering useful information in these kind of networks is epidemic routing [21, 5, 12]. As the name implies, epidemic routing operates much like the spread of an epidemic or an infection. In the absence of any central entities, it represents the fastest way in which information can be spread throughout a network. This mechanism is widely employed in the design of protocols for wireless networks especially delay tolerant networks [10] also sometimes known as intermittently connected networks. Some contexts in which epidemic routing like mechanisms have been effectively employed include node/neighbor discovery, data dissemination, routing, data gathering etc.

Traditionally, in social and environmental sciences, various studies have explored the spread of epidemics to gain insights into ways in which diseases spread. A rational school of thought suggests that if some information could be learnt about the way an infection spreads, it might make it possible to either halt or reduce the speed at which the infection proceeds. It is ironic that a process mimicking the spread of an infection is in fact an extremely useful way of exchanging information in sparse intermittently connected mobile networks (ICMNs). However, the information being disseminated among the nodes in ICMNs may contain malicious information. Consequently, this information can rapidly spread to all the nodes causing critical damage to the functioning of the network.

Some examples in which such networks may be deployed include (a) a conference setting for exchanging information between participants, (b) a network set up on-the-go for a disaster relief operation, (c) various wireless sensor network related applications (example environment monitoring, habitat monitoring etc.) where some/all nodes may be mobile (d) sophisticated artificial intelligence and robotics applications where teams of robots are deployed for exploration. In all these settings, even the presence of a single rogue participant node that may release a 'virus' into the setting can quickly spread across to all other nodes leading to a system malfunction. In this study, we formulate our problem based on this simple model: initially one node is infected and the infection is spread to another node via packet transmissions if both are in radio range of each other. Once a non-infected node is infected, it has the ability to infect other innocent nodes. Interesting metrics of interest include total time to infect a certain percentage of the nodes in the network. The mean total infection time refers to the time required to infect the entire network.

In general, in mobile wireless network studies, mobility patterns of nodes have been found to have a significant impact on the observed performance metrics [4, 6, 7]. A widely

MobilityModels'08, May 26, 2008, Hong Kong SAR, China.

employed mobility model of interest is a random mobility model where node movements are 'blind' and memoryless. In other words, a node moves independently of others and is unaware of the environment. This represents one extreme in the mobility space that we explore where the spread of an infection is studied. At the other extreme, we study a context-aware mobility model for the nodes. In this case, infector nodes move toward non-infected nodes who in turn try to evade them for as long as possible. We refer to this model as an *adversarial* mobility model. As the name implies, nodes view each other as adversaries or friends depending on their state and their behavior changes accordingly. While the adversarial mobility model may not be widely applicable under all scenarios with ICMNs, it presents a reasonable choice for a context-aware mobility model.

Our primary contributions are as follows. We investigate the spread of an infection in ICMNs with two extreme node mobility behaviors, namely, a random and an adversarial model. For nodes moving in a grid-based map, we quantify the mean total infection time as a function of key parameters such as number of nodes and the size of the field via mathematical analysis for both mobility models. Extensive simulations are conducted to validate our analysis. A key finding is that the total infection time with random node mobility *decreases* with the total number of nodes, while with adversarial node mobility this metric *increases* with the number of nodes. To gain further insights into the infection process, we encode this infection process into a game played by humans with the aid of wireless sensor motes. The results obtained from real game play show a closer match with the trends seen with the adversarial model as compared to the random model. However, the results do not strictly follow the simulation results. We analyze the collected trace data and identify reasons that explain this deviation. These empirical results serve in closing the loop of our study by complimenting our earlier findings with the adversarial mobility model.

The rest of the paper is organized as follows. In Section 2, we provide an overview of some related studies that have appeared in the literature. We provide mathematical analysis of the infection process using both the random and adversarial mobility models in Section 3. Section 4 presents simulation results that validate our analysis. In Section 5, we present empirical results of the infection based flooding game obtained from real play using mobility traces gathered from human movements. Finally, Section 6 provides brief conclusions and future research directions.

2. RELATED WORK

Several variants of epidemic routing have appeared in recent literature [5, 12] for the purpose of data delivery in wireless mobile ad-hoc networks with sparse density. Common metrics of interest in these studies include message delivery latency, throughput, message delivery overhead, and buffer usage. A closely related study to ours is the one conducted by Dimitriou *et al.* [9] who model the infection spread process by k concurrent random walks, each corresponding to a node that may either be infected or non-infected. However, the main focus of this study is comparison of worst case upper bounds for total infection time over different underlying graphs such as lollipop graph, clique, and expander graphs.

Our proposal is certainly complementary to these works but has some key differences. We investigate the infection process for a random and an adversarial mobility model on a grid-based topology. The random walk based model employed in our study is well known and has been widely employed by researchers [2, 3, 11, 13] chiefly because it provides a baseline against which performance of other mobility models can be compared. Our adversarial mobility model comes from the idea of pursuer-evader games [14, 8, 18]. This kind of mobility pattern is not only ubiquitous in nature, but can also provide a powerful framework of investigation.

The empirical part of our study implements the infection process as a mobile interaction game [22] with wireless sensor motes. Packet transmissions among motes emulates the infection spreading process and game ends when all the participants have been infected. Although our idea of using wireless sensor motes for such a game is novel, some other studies have deployed motes for games in more static contexts. One example is a game called Trove [17], where the participants negotiate with each other in a closed environment such as a room to reach a hidden treasure. Another example is the use of static sensor motes as resources such as virtual objects or characters in a physical environment in a game called 'Save the Princess' [16].

3. ANALYSIS OF THE INFECTION PROCESS

We provide a discrete space-time analytical model that captures the performance of the infection process for the two different node mobility behaviors: (a) random walk mobility model (b) adversarial mobility model. Let N be the number of nodes involved in the infection process, including the initial infector. For the ease of analysis, the field in which the nodes are moving with random walk mobility model is assumed to be a 2D torus with size $\sqrt{G} \times \sqrt{G}$ (similar settings have been employed in other studies [19, 20]).[1] In the adversarial mobility model, the field is assumed to be a grid with boundaries rather than a torus. This is because with an open field such as a torus, if nodes have similar skills, an infected node may never be able to infect any non-infected ones. The size of the grid is also $\sqrt{G} \times \sqrt{G}$, hence, there are a total of G grid cells in both cases. Two nodes are assumed to be in communication range as long as they are co-located in the same grid cell at the same time.

3.1 Random walk mobility model

This mobility model is 'blind' in that, at each step, each node moves from a cell to a neighboring cell oblivious to the positions and roles of the other nodes. As a precursor to developing an expression for the total infection time, next, we present some preliminary assumptions for this mobility model.

- At time $t = 0$ all nodes are placed in the torus uniformly at random. At any given time slot, a node in a cell (i, j) selects one of the four adjacent cells, i.e., cells $(i+1, j)$, $(i-1, j)$, $(i, j+1)$, and $(i, j-1)$ with equal probability, and moves to the selected cell. For a torus all the operations should be modulo $\sqrt{G}$.
- Initially only one node is infected. All nodes perform independent random walks, and a node is infected if it

[1] in Section 4.1, we verify this assumption by simulations, showing that the results for random walks on torus and grid with boundaries have similar trend.

meets an infected one in the same cell. Note that the independence assumption is valid when $G >> N$ which is frequently the case where nodes are moving around in a large field. The infection process is complete when all nodes are infected.

Aldous *et al.* [1] show that the the meeting time of two random walks in such a setting can be modeled as an exponential distribution with mean $C = \Theta(G \log G)$. Note that initially only one node carrying the infection will leave $N-1$ non-infected nodes. A transition occurs in the infection process when the infected node meets the first non-infected node. The node movements can be modeled as independent exponentials and since initially there are $N-1$ non-infected and one infected node such pairs, the average time taken to infect the first non-infected node is given by,

$$\overline{\delta_1} = \frac{C}{N-1} \tag{1}$$

Equation 1 shows that the average time to infect the first node reduces as the total number of nodes increases. The result is intuitive because for the same field size, as the number of nodes goes up, the probability of meeting the first non-infected node increases thereby reducing the time to first infection.

We generalize the above result as follows. At any given time, there are i nodes that are infected, consequently, $(N-i)$ nodes remain non-infected. Using the same idea as above, the average time until any of the i infected nodes meet any of the $(N-i)$ non-infected nodes is $\frac{C}{i\cdot(N-i)}$. Hence, the expected time until all the nodes have been infected, i.e. the expected total infection time, is given by:

$$\overline{\delta_N} = \sum_{i=1}^{N-1} \frac{C}{i \cdot (N-i)} \tag{2}$$

The above equation can also be written as:

$$\begin{aligned} \overline{\delta_N} &= \frac{C}{N} \cdot \sum_{i=1}^{N-1} \frac{1}{i} + \frac{1}{N-i} \\ &= 2 \cdot \frac{C}{N} \cdot \sum_{i=1}^{N-1} \frac{1}{i} \\ &= 2 \cdot \frac{C}{N} \cdot H_{N-1} \end{aligned} \tag{3}$$

Where H_{N-1} is the $(N-1)^{th}$ Harmonic Number.

When N is large, $H_{N-1} = \ln(N-1) + \gamma$ in which γ is the Euler-Mascheroni constant 0.5772. Hence equation 3 becomes:

$$\begin{aligned} \overline{\delta_N} &= 2 \cdot \frac{C}{N} \cdot (0.5772 + \log(N-1)) \\ &= \Theta(\frac{G \log G \log N}{N}) \end{aligned} \tag{4}$$

Equation 4 gives a tight bound for the average time to infect all N nodes on a 2D torus with size $\sqrt{G} \times \sqrt{G}$. It also indicates that the average total infection time reduces with N. The result is somewhat counter-intuitive: it suggests that as more and more nodes are injected, the total time to infect all of them goes down. This is because the node movements are random and as N increases, the time to infect the next non-infected node progressively decreases. The above result is verified by simulations in Section 4.1.

3.2 Adversarial mobility model

In some cases, the infected nodes may be more pro-active in spreading the infection by following and trapping non-infected ones. An example scenario is a team of self reliant robots being infected by some rogue robot(s). In this case, the infected and non-infected nodes view each other as 'enemies'. Before the analysis of such an infection process, we first list the assumptions for the adversarial mobility model.

- Each node has location information of all the other nodes. This can be done via a GPS device or employing other location based services.
- At each time step, assume that the uninfected nodes move first, followed by the infected ones. A node can move to any one of its eight (five if in an edge cell, three if in a corner cell) neighboring cells. It may also decide to stay put in its current cell.
- We assume that initially there's only one infected node. Without loss of generality, let the starting position of this initial infector be the bottom-left corner of the grid. The initial positions of the remaining $N-1$ nodes may be anywhere except the bottom-left corner, chosen uniformly at random.
- At any given time, a node will adopt one of the following three roles: *pursuer* (nodes who are infected), *evader* (non-infected nodes being chased by pursuers) and *rest-node* (non-infected nodes not being chased by pursuers). We assume that a node is able to identify the active roles of the other nodes.
- An evader e becomes infected if a pursuer p manages to move into the current cell of e at any time step. Once infected, a node cannot be disinfected. Let $(x_p, y_p), (x_e, y_e)$ denote the position of pursuer p and 'chosen' evader e, respectively.
- We assume a homogeneous environment in which all nodes are assumed to have the same skill set (for e.g., speed, infection range etc.).

Based on the above preliminaries, we provide efficient strategies for pursuers, evaders, and rest-nodes.

Pursuer Strategy: A pursuer p identifies its nearest non-infected node (which becomes the evader e) and moves to a neighboring cell such that $|x_p - x_e|$ and $|y_p - y_e|$ are reduced, if applicable. The pursuer does not chase another node until it infects the current identified evader.[2]

LEMMA 1. *Following the pursuer strategy, if p starts from any boundary of the square grid, it will be able to catch evader e in less than or equal to* $\sqrt{G} - 1$ *steps no matter how e moves.*

PROOF. Without loss of generality, suppose p starts (at $t = 0$) at any one of the left boundary cells $(1, y_{init})$ in which $1 \leq y_{init} \leq \sqrt{G}$. Because of the grid boundary constraints,

[2]In real scenarios, switching to chasing another noninfected node may be more beneficial. However, due to the fact that noninfected nodes always move away from pursuers (see evader and rest-node strategies) in our model, chasing a particular node until caught is the most efficient strategy for pursuers. We will leave the building of more realistic models and sophisticated strategies as future works.

there exists a time step t (where $1 \leq t \leq \sqrt{G}-1$) at which p and e have common values on their y coordinates for the first time. In other words, the pursuer 'catches' the evader in the y-dimension.

For each of the following time steps after t, no matter how the evader moves, the pursuer is able to choose a neighboring cell such that y_p will still equal to y_e. While on the x-dimension, the remaining time for the 'catch' is at most $\sqrt{G}-1-t$ steps. This is because after the first t steps, if p has not caught e, the x coordinate for p will be $(1+t)$, hence, the maximum distance between p and e can be $\sqrt{G}-(1+t)$.

As a result, the total time for the pursuer p to catch the evader e is less than or equal to $\sqrt{G}-1$. □

Although each evader will be eventually infected by a pursuer in bounded time, we seek to find a strategy to let the evader 'survive' for as long as possible. It turns out that for general initial positions of p and e, if the evader simply moves to the opposite direction from the pursuer to 'enlarge' its distance from p, the time that he can avoid being infected is maximized and matches the upper bound in lemma 1. The formal definition of evader strategy is:

Evader Strategy: The evader e moves to a neighboring cell that increases $|x_p - x_e|$ and $|y_p - y_e|$, if such a cell exists. In case multiple choices exist , the evader chooses one arbitrarily.

LEMMA 2. *Given that a pursuer p starts from any boundary cell of the square grid, and an evader e starts from any cell in the map other than the neighboring cell of p on the same boundary, the evader will be able to survive in at least* $\sqrt{G}-1$ *steps, if it follows the above evader strategy.*

PROOF. Suppose p starts at any cell of the left boundary cells $(1, y_{init})$. By the above strategy the evader will keep moving towards right in order to enlarge the value of $|x_p - x_e|$. Since p cannot catch e until e hits the right boundary, it requires the pursuer p to move at least $\sqrt{G}-1$ steps to reach the right boundary to catch the evader. □

The special case which lemma 2 excludes is when the evader e starts from p's neighboring cell on the same boundary (i.e: p starts at $(1, y_{init})$ and e starts at either $(1, y_{init}+1)$ or $(1, y_{init}-1)$). In this case, the time for p to catch e will depend on the value of y_{init} and is at most $\sqrt{G}-1$. Because the probability for this special case to happen is very low, and it only affects the first infection, we will not take it into consideration when calculating the total infection time.

So far we have not considered nodes that are neither pursuers nor evaders, i.e: neither infected nor being chased by infectors. We call these nodes: rest-nodes. A simple strategy for them is simply moving to a position in order to avoid getting in contact with both current pursuers as well as evaders ('potential pursuers'). Hence, the rest nodes should follow the following strategy:

Rest-node Strategy: The rest-nodes calculate their current distance to all the pursuers and evaders. They move to a neighboring cell that maximizes the summation of all these distances, if such a cell exists. When this strategy is adopted, it may be the case that remaining stationary in their current cell is the best strategy for the rest-nodes.

Observation: In any closed field like a square grid, the 'catch' always happens at the boundary cells. We provide an intuitive argument for the same. If the evader is not yet at a boundary, it can always find a neighboring cell to move to in order to increase its distance from the current pursuer. Since the grid provides finite boundaries, eventually the evader in its quest to escape reaches a field boundary, gets trapped and hence, caught by the pursuer.

Once the initial infector infects the first non-infected node the result is two infected nodes at a boundary cell. Consequently, for optimal total infection time, they both can operate in parallel and simultaneously infect others. In other words, each of the infected nodes chooses a different non-infected node if one exists and chases them down (1-to-1 match between pursuer and evader). Hence, at the end of an additional $\sqrt{G}-1$ steps, there will be a total of 4 nodes that will be infected. This process continues until all the nodes are infected signaling the end of the game.

Hence, in the infection process with N nodes on a $\sqrt{G} \times \sqrt{G}$ grid, the total infection time $\overline{\delta_N}$ can be written as:

$$\begin{aligned} \overline{\delta_N} &= (\sqrt{G}-1) \cdot \lceil \log_2 N \rceil \\ &= \Theta(\sqrt{G} \log_2 N) \end{aligned} \quad (5)$$

The total infection time is a non-decreasing function of the number of nodes N. This result is in contrast to the random walk mobility case presented earlier in Section 3.1, where we observed that the total infection time is a decreasing function of the number of nodes. In the next section, we will present simulation results that validate our analysis for both mobility models.

4. SIMULATIONS RESULTS AND DISCUSSION

4.1 Random walk mobility model

We performed simulations of the infection process using a discrete time model. Two different maps and movement patterns are considered: (a) The one used in Section 3.1, where the map is assumed to be a 2D torus and a node can only move to one of its four neighboring cells. (b) The map is a 2D grid with boundaries and a node randomly moves to one of its eight neighboring cells in each time step. Initially, all the nodes are distributed uniformly at random in the field and they perform independent random walks. One node is randomly selected as the initial infector. When an infected node meets a non-infected one, the latter gets infected and thereafter, both nodes continue performing random walks, i.e. the node movement patterns are unchanged. As before, a simulation run ends when all nodes are infected. All presented results are averages over 500 runs. Each simulation run employs a different seed that decides the initial placement of the nodes. Also, all the time metrics are reported in units of discrete time steps. We found that the 95% confidence intervals for the result metrics are quite tight and hence we do not present error bars in the figures.

Figure 1 captures the average total infection time for the random mobility model as a function of two parameters: the number of nodes (N) and the side length of the map ($\sqrt{G}$). As seen in Figure 1(a), for a given fixed size map, the average total infection time decreases with N. In Figure 1(b), we observe that for a fixed number of nodes, the total infection time increases with G. These observations match Equation 4 if we fix G and N, respectively.

In Figure 1(a), it is also important to note that even if nodes move in a 2D grid with boundaries and each node

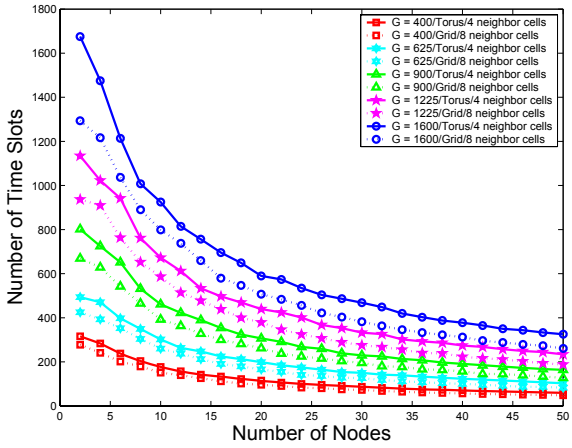

(a) Variation with Number of Mobile Nodes

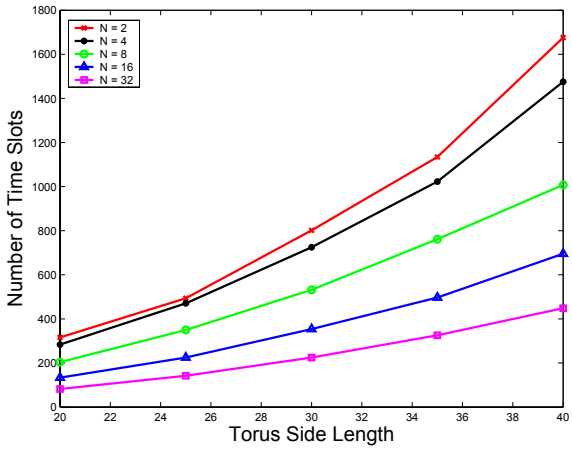

(b) Variation with Field Size

Figure 1: Shows the average total infection time for the random mobility model as a function of the number of mobile nodes and field size in figures (a) and (b) respectively.

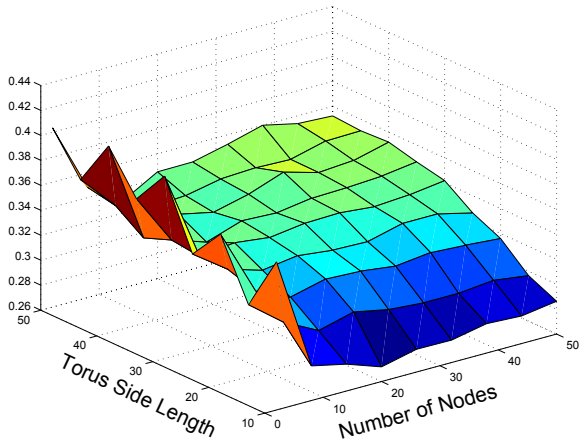

Figure 2: The ratio of average total infection time over $\frac{G \log G \log N}{N}$, which converges around 0.35.

can move into one of its eight neighboring cells in each time step, the trend of total infection time is same as the case of 2D torus. This observation serves as a validation for the assumptions made in Section 3.1 and the comparison between two different mobility models.

In order to verify the tightness of the bound suggested by Equation 4, we plot the value of average total infection time obtained from simulations on torus and corresponding value obtained from analysis for $N = 2, ..., 50$ and $\sqrt{G} = 10, ..., 50$. Figure 2 confirms that this ratio stabilizes around 0.35, as the number of nodes and torus size are increased. Notice that the peaks in Figure 2 are more pronounced for lower values of N. For low N, there is relatively larger difference between harmonic number $H_{N-1} = 1$ and $\log N$ employed in the approximation in the analysis.

4.2 Adversarial mobility model

A similar simulation study was performed for the adversarial mobility case using a similar discrete time model as was employed before. At the beginning of the simulation run, an infector is placed at the left-bottom corner and all other nodes are distributed uniformly at random in the square field.

Again, as before, we assume that the infection range of the nodes is one cell so that an evader will be infected if there is a pursuer in the same cell as itself. The size of the square grid is varied as $\{4 \times 4, 6 \times 6, 8 \times 8 , 10 \times 10 \}$. The number of nodes is varied in the range [2, 32]. A single run of the infection process ends when all the nodes are infected. Again, all results are averages over 500 runs and for each single run a randomly generated seed determined the initial location of the non-infected nodes.

Figure 3 shows the average total infection time for the adversarial mobility model as a function of G and N. In contrast to the random mobility case, for a given map size, we observe that the total infection time increases with N (Figure 3(a)). The total infection time also increases as the field size is increased (Figure 3(b)) which shows similar trend as in random mobility model. However, here, the average total infection time $\overline{\delta_N}$ varies linearly with the field length $\sqrt{G}$ whereas in the random mobility case $\overline{\delta_N}$ increases super-linearly with $\sqrt{G}$.

As before, to validate the tightness of the proposed analytical bound, we plot the the ratio of total infection time obtained with simulations with $\sqrt{G} \cdot \log_2 N$ for different N and $\sqrt{G}$ values in Figure 4. We find that this ratio stabilizes around 1, which implies the estimation of total infection time given by Equation 5 is quite accurate.

Comparing the values of total infection time for the two mobility models for the same G and N, we find that it takes longer to infect mobile nodes in the random mobility case. This is not surprising. Intuitively, there is a greater probability of rendezvous between nodes when infected nodes are actively (with 'open eyes') chasing non-infected ones as opposed to the case when all nodes are 'blind-folded' and the infectors rely on fortuitous encounters with non-infected ones for spreading the infection. More formally, for the adversarial mobility model the time for one infected node to catch another node is deterministic, and is given by the length of the map $(\sqrt{G} - 1)$. However, for random mobility this time is random due to the randomness of node mobility and is given by $\Theta(G \log G)$.

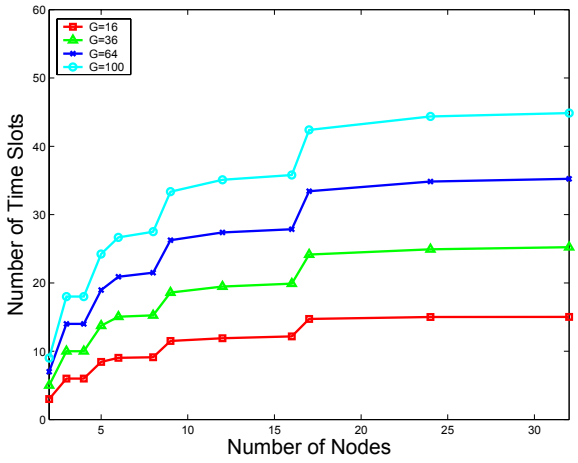

(a) Variation with Number of Mobile Nodes

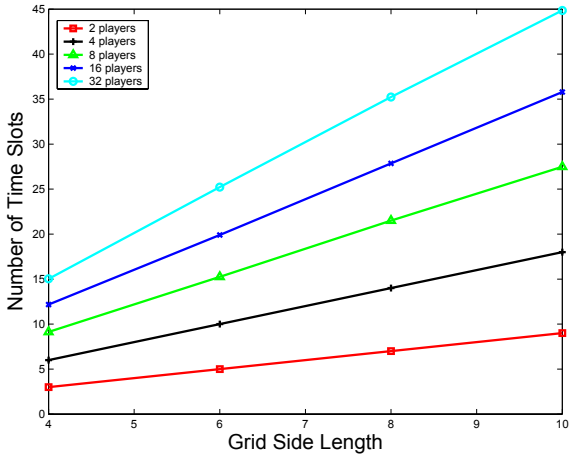

(b) Variation with Field Size

Figure 3: Shows the average total infection time for the adversarial mobility model as a function of the number of mobile nodes and field size in figures (a) and (b) respectively.

Simulation results of the infection process in continuous time and space offer similar observations and trends. For brevity and to avoid redundancy we do not include those results here.

5. EMPIRICAL RESULTS FROM THE MOBILE INTERACTION GAME

We have implemented a gaming version of the infection process on telosb motes [15] and conducted a few test runs with students from our laboratory in an outdoor field. This experiment of implementing a mobile interaction game besides providing a source of leisure and entertainment, also serves an additional purpose. The collected trace logs provide useful data for off-line analysis thereby allowing the possibility of using different types of mobile interaction games as test beds for evaluating performance of different wireless protocols. Next, we briefly describe some details of our implementation highlighting some differences in the set-up compared to the simulations and analysis presented earlier.

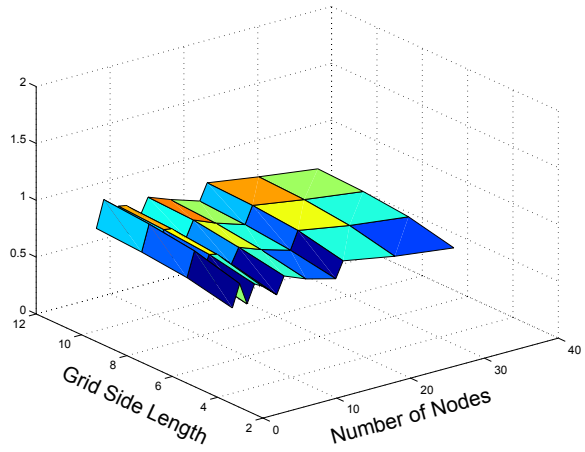

Figure 4: The ratio of average total infection time over $\sqrt{G}\log_2 N$, which converges around 1.

- A group of players each carrying a wireless sensor device (mote) gather around in a field with size $30m \times 35m$, representing the nodes in our study. Players are told to avoid being infected for as long as possible. No particular strategies are assigned to players.
- At the beginning, the clocks of all the motes are synchronized to keep track of the time when different motes get infected.
- Initially one player is infected (called initiator), his mote has its red LED switched ON. This reveals the identity of the infected player to the other players.
- Each infected mote, initially only that of the initiator, broadcasts an infection packet periodically, every BROADCAST_INTERVAL seconds. To visually aid the players, the blue and green LEDs are toggled to indicate that packets are being transmitted and received respectively.
- All other motes in the range of an infected mote may receive the packet[3]. If total number of infected packets received by a non-infected mote exceeds a certain threshold, that mote also gets infected. Let this threshold be denoted by INFECTION_THRESHOLD. This is to emulate the fact that an infection usually requires two nodes to stay together for a certain amount of time.
- The moment a mote gets infected its red LED is turned ON. This is an indication to the player that its behavior must change from that of an evader to a pursuer.
- When all the players have been infected indicated by lit red LEDs on the motes, an individual run of the game is complete.

Appropriate logs are created and stored on the flash memory of the motes during individual runs of the game to enable off-line processing of the logs to determine the exact

[3]Sometimes, we observed that players in radio range of other infected players did not get infected, because they used to shield the chip antenna of the telosb motes with their hands.

Number of Players	Total Infection Time
N = 10	91.69
N = 8	88.22
N = 6	74.11
N = 4	157.20

Table 1: Result for total infection time with respect to different number of players

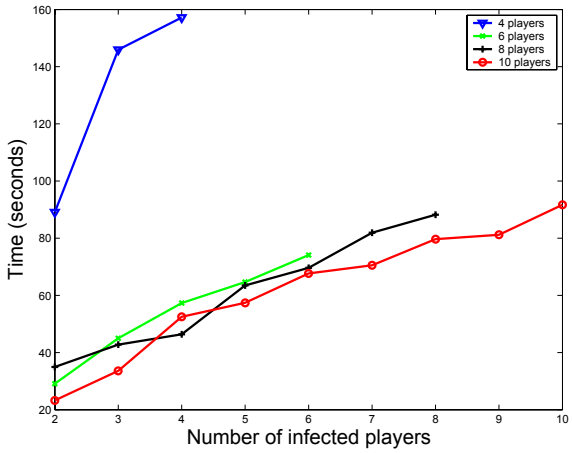

Figure 5: Shows the time-line for the flooding game indicating the individual times when the i^{th} player got infected.

time when each mote was infected. A reset function enabled multiple runs of the game to be executed one after another. We set the value of INFECTION_THRESHOLD to 10 while the BROADCAST_INTERVAL value was set to 100 milliseconds.

Before presenting the empirical results we note a few caveats. The presented results represent two game runs for each N value, hence, they may be influenced by the behavior of the players that participated in the run. Moreover, the same subset of players were juggled across the different game runs: players that were generally enthusiastic to play the game initially tended to get tired at the end. As a result when we conducted our final set of experimental runs for $N = 4$ players, we found that the total infection time was extremely high and the results were biased because the time taken by the pursuer to infect the first evader was exorbitant.

Table 1 presents the total infection time when the number of players is varied as $\{10, 8, 6, 4\}$. Here, except for the the $N = 4$ case where the first two individual infections took an extremely long time, the results are quite reasonable and we can observe the increasing trend in total infection time as number of players grows.

However, when we look into the details of the infection process, see Figure 5, even if we ignore the $N = 4$ case, we find that the time taken to infect the first non-infected player is different for different number of players. Recall that lemma 1 and 2 indicated that this value should be a constant $(\sqrt{G} - 1)$ for a given grid side length $\sqrt{G}$. In the following, we provide a brief discussion explaining this observation.

As one might expect in practical scenarios, we found that the capabilities and physical skills of different players are different whereas in the analysis and simulation studies we assume homogeneous skills across all nodes. Specifically, certain players are more athletic and nimble compared to others. For example, if an extremely fast runner is the initial infector for the case with $N = 10$, and the result of that run is compared to the case where an extremely slow runner is the initial infector with $N = 6$, it may turn out that the time taken to infect the first non-infected player with $N = 10$ is much smaller than with $N = 6$. This example can be extrapolated so that the total infection time with $N = 10$ may be lower than with $N = 6$ which is contradictory to our earlier results.

One may also argue that it is easier for one to infect the first non-infected player when the game has a larger N as compared to a smaller one. This is because in practical settings, it stands to reason that the players are more or less aware of each other's skills such as how fast a player moves, how competitive a player is, etc. Now, say with $N = 8$ players, the initial infector has more choices in terms of this selection as compared to the case with $N = 4$ players. Although human beings do not always make the globally optimal decision of minimizing the total infection time, selecting and trapping a player given fewer choices generally turns out to be more difficult.

Furthermore, we observed that human beings did not perform optimally in the mobile interaction game. For example, players getting infected did not immediately start chasing other players. In some cases, multiple players ended up chasing the same non-infected player whereas operating in parallel would have resulted in a lower total infection time. In other cases, some infectors 'goofed' up and spent time chasing players that were already infected. Recall that in the simulations and analysis presented earlier, each player always has global knowledge of the state of the game and strictly obeys the outlined strategies depending on its current state. Clearly, in real-play oriented practical settings this cannot be enforced especially when human beings are participants in the experiment. The empirical observations might show a closer match with the earlier model results when entities such as robots may play the game.

In general, a small number of gaming experimental runs may not be good for performance testing purposes simply because humans are not pre-programmed and do not perform optimally as can be controlled in simulated environments. However, as data is collected from a large number of experimental runs and players become more strict in following the strategies, it is safe to assume that the bias caused by player heterogeneity and skill dynamics can be reduced and more accurate and statistically significant observations can be made.

6. CONCLUSIONS AND FUTURE RESEARCH DIRECTIONS

We have presented a systematic study of a spread of an infection across an intermittently connected wireless network. Two different node mobility models have been considered namely random mobility where node movements are 'blind' and adversarial mobility where the nodes view each other as 'enemies' and make movement decisions depending on each others' states and locations. We have studied the infection process as a function of two key parameters namely the size

of the field and the number of nodes. A key result, validated by analysis and simulations, is that total infection time for the random mobility case decreases with the number of nodes while with adversarial mobility this time increases with the number of nodes. Real world experimental runs by humans and off-line processing collected logs provide some insights into the infection process. The empirical results show a closer match with the trends seen with the adversarial model, however, a large collection of real traces is necessary for making statistical observations on the details of our study. Otherwise, the results will be biased by human heterogeneity in terms of diverse player skills.

As part of our future research, we would like to extend our model to capture node heterogeneity in terms of dynamics such as different speeds and infection ranges. Also, the analytical model can be extended to capture the empirical experiments more accurately where the infectors require multiple encounters with non-infected players in order to infect them. More sophisticated strategies such as switching short term goal and cooperation among nodes that will have a significant effect on the infection process present another promising research direction.

7. REFERENCES

[1] D. Aldous and J. Fill. Reversible markov chains and random walks on graphs. manuscript available at http://www.stat.berkeley.edu/ aldous/RWG/book.html.

[2] A. Bar-Noy, I. Kessler, and M. Sidi. Mobile users: to update or not to update? In *Proceedings of the Joint Conference of the IEEE Computer and Communications Societies(INFOCOM)*, 1994.

[3] L. Breslau, D. Estrin, K. Fall, S. Floyd, J. Heidemann, A. Helmy, P. Huang, S. McCanne, K. Varadhan, Y. Xu, and H. Yu. Advances in network simulation. In *IEEE Computer*, volume 33, pages 59–67, May 2000.

[4] J. Broch, D. A. Maltz, D. B. Johnson, Y. C. Hu, and J. Jetcheva. A performance comparison of multi-hop wireless ad hoc network routing protocols. In *Fourth Annual ACM/IEEE International Conference on Mobile Computing and Networking*, October 1998.

[5] J. Burgess, B. Gallagher, D. Jensen, and B. Levine. MaxProp: Routing for Vehicle-Based Disruption-Tolerant Networking. In *Proc. of IEEE Infocom*, April 2006.

[6] T. Camp, J. Boleng, and V. Davies. A survey of mobility models for ad hoc network research. In *Wireless Communications & Mobile Computing(WCMC): Special Issue on Mobile Ad Hoc Networking: Reaearch, Trends, and Applications*, volume 1, pages 483–502, 2001.

[7] S. R. Das, C. E. Perkins, and E. M. Royer. Performance Comparison of Two On-demand Routing Protocols for Ad Hoc Networks. In *Proceedings of the Joint Conference of the IEEE Computer and Communications Societies(INFOCOM)*, March 2000.

[8] M. Demirbas, A. Arora, and M. Gouda. A pursuer-evader game for sensor networks. In *Sixth Symposium on SelfStabilizing Systems(SSS'03)*, pages 1–16, 2003.

[9] T. Dimitriou, S. Nikoletseas, and P. Spirakis. The infection time of graphs. *Discrete Applied Mathematics*, 154:2577–2589, 2006.

[10] Kevin Fall. A delay-tolerant network architecture for challenged internets. In *SIGCOMM '03: Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications*, pages 27–34, New York, NY, USA, 2003. ACM.

[11] J. J. Garcia-Luna-Aceves and E. L. Madrga. A Multicast Routing Protocol for Ad-Hoc Networks. In *Proceedings of the Joint Conference of the IEEE Computer and Communications Societies(INFOCOM)*, 1999.

[12] A. Jindal and K. Psounis. Performance analysis of epidemic routing under contention. In *IWCMC '06: Proceeding of the 2006 international conference on Comm unications and mobile computing*, pages 539–544, New York, NY, USA, 2006. ACM Press.

[13] D. Johnson and D. Maltz. Dynamic source routing in ad hoc wireless networks. In *Mobile Computing, T. Imelinsky and H.Korth, eds*, pages 153–181. Kluwer Academic Publishers, 1996.

[14] G. F. Miller and D. Cliff. Protean behavior in dynamic games: arguments for the co-evolution of pursuit-evasion tactics. In *Proceedings of the third international conference on Simulation of adaptive behavior : from animals to animats 3*, 1994.

[15] moteiv. Telosb motes datasheet. http://www.moteiv.com/products/docs/telos-revb-datasheet.pdf.

[16] L. Mottola, A.L. Murphy, and G.P. Picco. Demonstrating Pervasive Game Development using Tiny Devices and the TinyLIME Middleware. In *5th Workshop on Network and System Support for Games(NetGames)*, Oct 2006.

[17] S. Mount, E. Gaura, M. Newman, A. Beresford, S. Dolan, and M. Allen. Trove: a Physical Game Running on an Ad-Hoc Wireless Sensor Network. In *Joint sOc-EUSAI*, pages 235–240, October 2005.

[18] L. Schenato, S. Oh, and S. Sastry. Swarm coordination for pursuit evasion games using sensor networks. In *Proc. of the International Conference on Robotics and Automation, Barcelona, Spain*, 2005.

[19] R. Shah, S. Roy, S. Jain, and W. Brunette. Data mules: Modeling and analysis of a three-tier architecture for sparse sensor networks. *Elsevier Ad Hoc Networks Journal*, 1, September 2003.

[20] T. Spyropoulos, K. Psounis, and C. Raghavendra. Single-Copy Routing in Intermittently Connected Mobile Networks. In *Proc. of IEEE SECON*, April 2004.

[21] A. Vahdat and D. Becker. Epidemic routing for partially-connected ad hoc networks. Technical report, Department of Computer Science, Duke University, 2000.

[22] Y. Wang, S. Kapadia, and B. Krishnamachari. MIGM: Mobile Interaction Games with Motes. In *Proc. of IEEE CCNC*, Janurary 2008.