

Epidemic Spread in Mobile Ad Hoc Networks: Determining the Tipping Point

Nicholas Valler, B. Prakash, Hanghang Tong, Michalis Faloutsos, Christos Faloutsos

► **To cite this version:**

Nicholas Valler, B. Prakash, Hanghang Tong, Michalis Faloutsos, Christos Faloutsos. Epidemic Spread in Mobile Ad Hoc Networks: Determining the Tipping Point. 10th IFIP Networking Conference (NETWORKING), May 2011, Valencia, Spain. pp.266-280, 10.1007/978-3-642-20757-0_21 . hal-01583416

HAL Id: hal-01583416

<https://hal.inria.fr/hal-01583416>

Submitted on 7 Sep 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Epidemic Spread in Mobile Ad Hoc Networks: Determining the Tipping Point

Nicholas C. Valler^{1*}, B. Aditya Prakash², Hanghang Tong³, Michalis Faloutsos¹, and Christos Faloutsos²

¹ Dept. of Computer Science and Engineering, University of California, Riverside
{nvalley, michalis}@cs.ucr.edu

² Computer Science Dept., Carnegie Mellon University
{badityap, christos}@cs.cmu.edu

³ IBM T.J. Watson Research Center
htong@us.ibm.com

Abstract. Short-range, point-to-point communications for mobile users enjoy increasing popularity, particularly with the rise in Bluetooth-equipped mobile devices. Unfortunately, virus writers have begun exploiting lax security in many mobile devices and subsequently developed malware exploiting proximity-based propagation mechanisms (e.g. Cabir or CommWarrior). So, if given an ad-hoc network of such mobile users, will a proximity-spreading virus survive or die out; that is, can we determine the “tipping point” between survival and die out? What effect does the average user velocity have on such spread? We answer the initial questions and more. Our contributions in this paper are: (a) we present a framework for analyzing epidemic spreading processes on mobile ad hoc networks, (b) using our framework, we are the first to derive the epidemic threshold for any mobility model under the SIS model, and (c) we show that the node velocity in mobility models does not affect the epidemic threshold. Additionally, we introduce a periodic mobility model and provide evaluation via our framework. We validate our theoretical predictions using a combination of simulated and synthetic mobility data, showing ultimately, our predictions accurately estimate the epidemic threshold of such systems.

1 Introduction

The prevalence and increased functionality of mobile phones present an unique opportunity for malicious software writers. According to popular reports, 45.4 million U.S.

* This material is based upon work supported by the Army Research Laboratory under Cooperative Agreement No. W911NF-09-2-0053, the National Science Foundation under Grants No. IIS-1017415, CNS-0721736 and CNS-0721889, NETS-0721889 and NECO-0832069 and a Sprint gift. Any opinions, findings, and conclusions or recommendations in this material are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory, the U.S. Government, the National Science Foundation, or other funding parties. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on.

citizens own smartphones in 2010 [12], roughly 15% of the total population. These devices are becoming more and more essential for the daily lives of end-users, especially as smartphones offer more and more capabilities. At the same time, mobile devices are equipped with device-to-device (a.k.a. point-to-point) communication technologies (such as Bluetooth), where the communication does not use the phone’s service provider’s infrastructure. There are estimated over 920 million Bluetooth-equipped devices worldwide shipped in 2008, making Bluetooth the most common point-to-point communication protocol in today’s smartphones. To date, at least two smartphone worms have been found in the wild, Cabir and CommWarrior, which spread using Bluetooth.

Given a system of mobile agents, such as smartphones, what can we say about the propagation of a virus⁴ within the system? A key question is to identify the tipping point, known as the system’s *epidemic threshold*, or *take-off point*, below which a virus is guaranteed to “die out.” For the epidemic models, we focus on one of the most popular models, the flu-like one susceptible-infected-susceptible - SIS (see Section 5 where we handle other models). There, agents maintain no immunity, and become susceptible, immediately after they heal. Our key contributions are as follows:

1. *Framework and Formula*: We present a framework for estimating the epidemic threshold on *any, arbitrary* mobile ad hoc network model. To the best of our knowledge, ours is the first theoretical study for the general mobility case. The idea is to derive a sequence of adjacency matrices, and then compute the first eigenvalue of the so-called *system matrix* (see Theorem 1). There, we show that the epidemic threshold depends only on this first eigenvalue, and nothing else.
2. *Closed Formulas*: We show how to use our framework to derive simple, approximate (but accurate) formulas for several, popular special cases (Random Walk model, Levy flight model).
3. *Insensitivity to Velocity*: To our surprise, our results show that the epidemic threshold does *not* depend on the node velocity ($v > 0$). Our experiments confirm the accuracy of our approximations, as well as our ‘insensitivity’ observation.

Jumping ahead, Figure 1 showcases the accuracy of our results (Lemma 3) of the epidemic threshold for the SIS (=flu-like) model on the so-called ‘Levy Flight’ mobility model. See section 4 for more details - but the point is that our prediction for the take-off point (= *epidemic threshold*, indicated with a black arrow) is exactly where all curves take off.

We have two additional contributions: through extensive simulation experiments, we show that similar insensitivity results hold for other popular mobility models like Levy flight, Random Waypoint etc.; and, moreover, we introduce the *periodic mobility* model, which is very popular in biological virus epidemiology [1, 4], and show how to use our framework to estimate its epidemic threshold.

The rest of this paper has the typical organization: background (Section 2, proposed framework and theorems (Section 3), experiments (Section 4), additional observations (Section 5), and finally, related work and conclusions. (Sections 6 and 7).

⁴ We focus on proximity propagation, which in mobile networks is affected by mobility. However, virus for smartphones can also propagate through email, mms, or direct access to the web. We do not consider these cases, sine they are not directly affected by mobility




Fig. 1. Accuracy of our results (Lemma 3), for the ‘Levy-Flight’ model. Take-off plot, plotting the max number of infections vs. strength of the virus. Notice that our predicted take-off point (black arrow) agrees with the simulations, for several node velocities.

2 Preliminaries

In this section, we present a general background on proximity-based epidemic spreading models and formulate our problem statement.

2.1 Epidemic Model: SIS (flu-like)

The SIS epidemic model resembles a flu-like virus, where nodes have no immunity. Healthy (‘S’ = susceptible) nodes become sick (‘I’ = infected) stochastically from their infected neighbors with a probability β . Alternatively, a sick node becomes healthy (and open to re-infection), with a probability δ . These two parameters are also referred to as the *birth* rate (β) and *death* rate (δ) of the virus.

The tipping point τ , or *epidemic threshold*, of an SIS epidemic model is the condition under which an infection will die out exponentially quickly irrespective of initial infection, as opposed to spreading out, causing an epidemic (technically, a pandemic). For a survey on SIS and numerous other epidemic models, see Hethcote [13], or [7, 11].

2.2 Problem Formulation

Using the background discussed above, we now formulate our problem statement. See Table 1 for definitions of various symbols. In this paper, we consider an epidemic on a mobile network, which provides an underlying contact structure for the virus to use as it propagates. By doing so, at any point in time, the system is non-homogenous, as nodes may only transmit the virus to its neighbors.

Given:

1. Mobile ad hoc network mobility models, described below.
2. The SIS model parameters, i.e., the virus birth and death probabilities β and δ .

Find:

The epidemic threshold τ or *tipping point* for the system such that for $\tau < 1$ an infection will die out quickly, irrespective of initial conditions.

Table 1. Terminology

Symbol	Definition and Description
General Terms:	
$\mathbf{A}, \mathbf{B}, \dots$	matrices (bold upper case)
$\mathbf{A}(i, j)$	element at the i^{th} row and j^{th} column of \mathbf{A}
$\mathbf{A}(i, :)$	i^{th} row of matrix \mathbf{A}
$\mathbf{A}(:, j)$	j^{th} column of matrix \mathbf{A}
\mathbf{I}	standard $n \times n$ identity matrix
$\mathbf{a}, \mathbf{b}, \dots$	column vectors
$\mathcal{I}, \mathcal{J}, \dots$	sets (calligraphic)
$\lambda_{\mathbf{B}}$	first eigenvalue (in absolute value) of a matrix \mathbf{B}
Mobility Terms:	
M	mobility model
$P_{i,t}$	position of node i at time t
N	number of nodes
A	simulation area
ρ	node density (N/A)
ΔT	Time step
T	number of different alternating behaviors
$\mathbf{A}_1, \dots, \mathbf{A}_T$	T corresponding size ($n \times n$) symmetric alternating adjacency matrices
Epidemic Terms:	
β	virus transmission probability in the SIS model
δ	virus death probability in the SIS model
τ	epidemic threshold
Acronyms and Terms:	
USS	Uniform Steady-State Approximation
EAAM	Eigenvalue of Average Adjacency Matrix Approximation
Take-off Plot	Max number of infected agents vs. Epidemic Threshold τ

Our problem naturally leads to other issues like the effect of node velocity in models on the threshold, giving approximations in specific cases etc. We elaborate on them in the upcoming sections.

3 Framework

In this section, we detail our framework for analyzing mobility models and then move on to specific approximations and questions arising out of the framework. We will present extensive simulations demonstrating the results later in Section 4. Also please see Section 4.2 for a description of the mobility models.

Note that node-to-node contacts at a particular time can be represented by an *adjacency matrix* \mathbf{A} . We next provide a general theorem expressing the epidemic threshold for mobility models.

3.1 Epidemic Thresholds on Mobility Models

Theorem 1 (Mobility model threshold). *If a mobility model can be represented as a sequence of connectivity graphs $L = \{\mathbf{A}_1, \mathbf{A}_2, \dots, \mathbf{A}_T\}$, one adjacency matrix \mathbf{A}_t for each time step $t \in \{1..T\}$, then the epidemic threshold is:*

$$\tau = \lambda_S \tag{1}$$

where $\lambda_{\mathbf{S}}$ is the first eigenvalue of matrix \mathbf{S} and $\mathbf{S} = \prod_i \mathbf{S}_i$ and $\forall_i \in \{1..T\}$ $\mathbf{S}_i = (1 - \delta)\mathbf{I} + \beta\mathbf{A}_i$ (\mathbf{I} is the standard $N \times N$ identity matrix).

Proof. If the mobility model can be represented as a sequence of graphs, then the epidemic threshold depends on the first eigenvalue of the system matrix [25]. Hence, $\tau = \lambda_{\prod_i((1-\delta)\mathbf{I}+\beta\mathbf{A}_i)}$. \square

We can now give a simpler closed-form approximation for the threshold in Equation 1 in the following lemma:

Lemma 1 (EAAM Approximation for Threshold). *Under the same conditions as in Theorem 1, the following is an approximation for the epidemic threshold:*

$$\tau \approx \frac{\beta}{\delta} \times \lambda_{\mathbf{A}_{avg}} \quad (2)$$

where $\mathbf{A}_{avg} = \sum_i \mathbf{A}_i/T$ is the average adjacency matrix.

Proof. Note that,

$$\begin{aligned} \mathbf{S} &= \prod_i ((1 - \delta)\mathbf{I} + \beta\mathbf{A}_i) \\ &= (1 - \delta)^T \mathbf{I} + \beta \sum_i \mathbf{A}_i + O(\beta^2) + O(\beta * \delta) + O(\delta^2) \\ &\approx (1 - T\delta)\mathbf{I} + T\beta\mathbf{A}_{avg} \end{aligned} \quad (3)$$

where we neglected second or lower order terms involving β and δ . Hence, we find that $\mathbf{B} = (1 - T\delta)\mathbf{I} + T\beta\mathbf{A}_{avg}$ is a first order approximation for the $\mathbf{S} = \prod_i \mathbf{S}_i$ matrix. Hence from Theorem 1 we want $\lambda_B < 1$ which implies Equation 2. \square

We will refer to the above approximation as the ‘Eigenvalue of the Average Adjacency Matrix’ (EAAM) approximation.

3.2 Specific Approximations

Lemma 2 (Random-Walk Threshold). *In the random-walk mobility model and under the SIS model, the following is an approximate epidemic threshold:*

$$\tau \approx \beta/\delta \times \pi R^2 \times N/A \quad (4)$$

where R is the radius of influence of each node.

Proof. Under the random-walk model, at every point of time, each node is at a random (x, y) position, uniformly distributed on the field of interest. Each node has a radius of possible connections (like the BlueTooth radius) R . Consequently each node has $d = \pi R^2 \times N/A$ neighbors on average (ignoring boundary effects). The connectivity graph at each time step is roughly a random graph with average degree d . Hence it has first eigenvalue $\lambda_1 = d$ on average. Hence this is approximately equivalent to having a static graph under the SIS model where the epidemic threshold [7] is $\tau = \beta/\delta \times \lambda_1$. We now obtain the lemma after obvious substitutions. \square

In fact, we can go further and generalize this to *any* mobility model where the geographic steady state distribution is uniform.

Lemma 3 (Uniformly-Distributed Steady State (USS) Threshold). *For any mobility model where the geographic distribution of nodes at the steady state is uniform over the area of interest and under the SIS model, the following is an approximate epidemic threshold:*

$$\tau \approx \beta/\delta \times \pi R^2 \times N/A \quad (5)$$

where R is the radius of influence of each node.

Proof. The proof for Lemma 2 goes through even here precisely because of the geographically uniformly distributed nature of the steady state. Each node has roughly the same number of connections and hence the adjacency graph is approximately a homogenous graph with constant first eigenvalue. The result follows as before. \square

We will refer to the above approximation as the ‘Uniform Steady-State’ (USS) approximation. Mobility models like Levy Flight and Random-Walk are examples of models with a geographically uniform distribution of the nodes at the steady state. Lemma 3 allows us to quickly estimate the threshold for these and many other models.

3.3 Insensitivity to Node Velocity

As there is no factor depending on the node velocity in Lemma 3, we conclude the following surprising implication:

Corollary 1 (Node velocity and threshold). *The node velocity ($v > 0$) does not affect the epidemic threshold in mobility models where the steady state has a geographically uniform steady state distribution like Random Walk, Levy Flight etc.*

We conjecture that the velocity does not affect the threshold even for non-geographically uniformly distributed steady state mobility models like Random Waypoint. We provide empirical results supporting this claim later in Section 4.

Conjecture 1 (Effect of velocity). The node velocity ($v > 0$) does not affect the epidemic threshold in the Random Waypoint mobility model.

Does Velocity have an impact at all? The above discussion raises the point whether the node velocity has any effect at all on the dynamics of the epidemic spreading. We expect that the velocity of motion does have an effect, when we are above threshold. Furthermore, simulations resulted in a non-obvious observation. The velocity had an impact on the steady-state number of infected agents in the system. We elaborate more on these issues in Section 5.

4 Simulation Methodology and Results

4.1 Experimental Setup

To facilitate the simulation, we wrote a custom Python2.6 simulation program using the NumPy/SciPy python libraries. All simulations were conducted on a 4 core Intel(R) Xeon(R) CPU operating at 2.53 GHz and 72 GB of memory running CentOS-5.5 (Linux kernel 2.6). We varied the number of agents (nodes) N between 250 and 1500 within a simulation field of area $A = 40,000m^2$ (200m by 200m). Thus, node density ρ , commonly defined a N/A , was between 0.125 and .125 nodes per m^2 . All nodes had a

transmission range of 5.0 meters. We did not account for signal attenuation, reflection nor other wireless phenomena. Prior to the beginning of the simulation, nodes were distributed on the simulation field in a uniform fashion. Simulations were generally run for a period of 100s with time intervals of $\Delta T = 0.1$ seconds.

We studied three mobility models common to mobile ad hoc networking: Random Walk, Levy Flight and Random Waypoint. In the following sections, we provide detail on each model as well as simulation results. The position $P_{i,t}$ of each node in the system at time t is a function of mobility model and previous position and time step ΔT , such that $P_{i,t+1} = M(P_{i,t}, \Delta T)$, where M is the mobility model.

The purpose of our simulations was to determine the role of the mobility model in the spread of malware in a point-to-point contact network loosely describing Bluetooth communication technology.

4.2 Mobility Models

Random Walk. The Random Walk (RW) mobility model (also referred to as *Brownian Motion*) was originally formulated to describe the seemingly random motion of particles. Numerous variations exist, here we describe our implementation.

Each node i in the system is parameterized by *speed* (V_i) and *angle* (θ_i). Both V_i and θ_i are drawn uniformly from systemwide predefined ranges, $[v_{min}, v_{max}]$ and $[0, 2\pi)$, respectively. Clearly, such a system is memoryless. The model we employ varies from the simple RW model by introducing a *flight time* for each node, T_i . Flight time is drawn uniformly from a range $[\tau_{min}, \tau_{max}]$. The spatial distribution of the RW mobility model is uniform over the simulation field. According to our framework, we predict the RW mobility model will follow Lemma 2.

Levy Flight. Levy Flight mobility models have recently attracted attention due to their statistical similarities with human mobility [27]. At the beginning of each flight, each node selects an angle uniformly from within $(0, 2 * \pi]$, a flight time drawn from some distribution, a flight length and a pause time. Flight length and pause time are drawn from Levy distributions $p(t) \propto |t|^{-(1+\alpha)}$ and $\psi(t) \propto t^{-(1+\beta)}$, where time $t > 0$, respectively. When $\alpha = 2$ and $\beta = 2$ the result is a special case of the Levy distribution resulting in a Gaussian distribution. As with the Random Walk, the spatial distribution of the Levy Flight mobility model is uniform over the simulation field.

Random Waypoint. The Random Waypoint (RWP) mobility model is often cited as the *de facto* mobility model in ad hoc networks. As originally proposed by Johnson et. al [14], the RWP mobility model each node i is described by three parameters: *current location* (P_i), *speed* (V_i), *waypoint* (W_i), and *pause time* (ρ). In general, the RWP mobility model operates as follows: Initially, a node is stationary. After a pause time ρ , the node selects a waypoint uniformly from the simulation field, then, travels along the shortest path to its waypoint P at a velocity V_i drawn uniformly from $(v_{min}, v_{max}]$. Upon arrival at their waypoint, each node will pause for a time $t = \rho$. After the pause period is done, each node will repeat the process. The spatial distribution of the RWP mobility model is bell-shaped [5].

4.3 Summary of Results

Accuracy of Approximations, Random Walk. We present a series of simulation studies of the Random Walk mobility model in Figure 2. In these studies, we varied the birth (β) and death (α) parameters of the SIS infection model. We refer to the resulting




Fig. 2. Accuracy of Framework Approximations, Random Walk Take-off plots for EAAM and USS. Notice that our predictions (black arrows) are accurate.

plots as “take-off plots,” which show the maximum number of infected agents seen in our simulation against our approximated epidemic threshold. For each plot, we labeled the estimated take-off point according to the specific threshold approximation.

Figure 2(A), we approximated the epidemic threshold using the first eigenvalue of the average adjacency matrix, \mathbf{A}_{avg} (the EAAM threshold approximation of Lemma 1). We indicate the predicted threshold value at $\tau = 1$. As expected, no epidemic was present at values of $\tau < 1$. At values of $\tau > 1$, we see explosive growth in the max number of infected agents.

In Figure 2(B), we plot the the USS approximation of the epidemic threshold (Lemma 3), i.e. $\beta/\delta \times \pi R^2 \times N/A$. As in the EAAM threshold, USS behaves as expected (i.e. no epidemic below $\tau = 1$). In fact, in each of these figures, we note that no infected agents (aside from patient zero) are present for epidemic threshold values below 1. Compared to EAAM, this threshold value takes off at values closer to $\tau = 1$.

Accuracy of Approximations, Levy Flight. As with RW, the spatial distribution of nodes following Levy Flight mobility model is uniformly distributed on the simulation field. Thus, we expect Levy flight to perform similar to RW. Figure 3 presents a take-off plots for Levy Flight simulations. Again, we note that no infected agents exist below either threshold approximations. Furthermore, we find that USS performs better than EAAM.

Accuracy of Approximations, Random Waypoint. The next series of simulations were conducted using the popular Random Waypoint mobility model. We also selected this model specifically because it does not result in a uniform spatial distribution of nodes, therefore does not fall under USS. The results of these simulations are presented in Figure 4. Surprisingly, both threshold approximations perform well against




Fig. 3. Accuracy of Framework Approximations, Levy Flight Take-off plots for EAAM and USS. Again our predictions (black arrows) prove accurate.

the RWP mobility model. As the earlier mobility models exemplified, USS performs better than EAAM.

Insensitivity to Velocity. In order to illustrate Corollary 1 and validate Conjecture 1, we performed a series of simulation in which we varied the node velocity, the results of which are illustrated in Figures 2, 3, 4.

As expected, for both the RW and Levy Flight mobility models, the take-off points were not greatly affected by increasing the nodes velocity. Furthermore, Figure 4 shows that the take-off point for the RWP mobility model was not affected by velocity, affirming Conjecture 1.

5 Discussion

We elaborate here on the effect of node velocity on the dynamics of epidemic spreading. We also introduce the periodic mobility model and present an analysis of it via our framework. In addition, we touch upon other epidemic models as well.

5.1 More on impact of node velocity

As discussed previously in Section 3, node velocity does not seem to effect the threshold in many models. We now ask whether the velocity affects the epidemic at all?

For an “above threshold” system, two more parameters are of interest: (a) *steady-state maximum*, and (b) *warm-up period*. The steady-state maximum is the maximum number of infected agents in the system till steady state, whereas the warm-up period is the time necessary to reach steady state. We expect that the velocity of motion does have an effect, when we are above threshold. Clearly, speed will effect the speed of propagation of the virus and thus the warm-up period. Higher velocity means better mixing of agents, and thus faster convergence to the steady state. This observation is also demonstrated through simulations.

Figure 5 (best viewed in color) shows the number of infected agents per unit time (in seconds) for both the Random Walk and Random Waypoint mobility models. The




Fig. 4. Accuracy of Framework Approximations. Random Waypoint Take-off plots for EAAM and USS.

velocity varied between a fixed 1 m/s fixed and an uniform selected 15 – 20 m/s, as indicated in the legends. We performed a longer simulation in order for the systems to settle in a steady state.

Less intuitive is that velocity appears to affect the number of infected agents at steady-state. For example, in Figure 5(A), the line corresponding to 1 m/s appears to reach a steady-state of approximately 65 – 70% infected agents, whereas, at 15 – 20 m/s, the steady-state is roughly 90%. The steady-state for velocities between these two extremes lay in-between. We suspect the degree of mixing, influenced by node velocity, is the root cause of the above observation.

5.2 The Periodic Mobility Model

As we indicated in Section 3, our framework predicts the epidemic threshold of mobility models that can be represented as a series of adjacency matrices. The *Periodic Mobility Model* is a special case of such a series, where a set of k adjacency matrices $\{\mathbf{A}_1, \mathbf{A}_2, \dots, \mathbf{A}_k\}$ are repeated periodically.

This is a typical model used in biological virus studies [4] to model general movements of a population. As an example, let \mathbf{A}_1 be an adjacency matrix of people during the day (say, at the office). Let \mathbf{A}_2 be an adjacency matrix representing contacts/interactions during the evening (say, at home). So the series formed by repeating $\{\mathbf{A}_1, \mathbf{A}_2\}$ represents the daily, repeated interactions of our population. The periodic model offers a realistic, yet general model of mobility, capturing general patterns rather than specific movements of the system.

Lemma 4 (Periodic Model Threshold). *Under the periodic mobility model with k alternating behaviors repeating periodically, the epidemic threshold is given by $\tau = \lambda_S$ where $\mathbf{S} = \prod_{i=1}^k \mathbf{S}_i$ and, as before, $\mathbf{S}_i = (1 - \delta)\mathbf{I} + \beta\mathbf{A}_i$.*

Proof. Omitted for brevity, similar to Theorem 1.

As an example, Figure 6 shows the take-off plot of a periodic mobility models where $N = 500$ nodes, $k = 2$ adjacency matrices. As predicted by Lemma 4, the max number of infected agents over the simulation period takes off at $\lambda_S = 1$.




Fig. 5. Number of Infected Agents vs. Time on the Random Walk and Random Waypoint. Node velocities were varied as indicated in the legends. Steady-State number of infected agents increases with node velocity, while warm-up period shrinks (best viewed in color).

5.3 Other Epidemic Models

Given recent results on epidemic thresholds on static networks [24], we believe that our results will carry through for many other epidemic models as well e.g. SIR (mumps-like), SIRS, SEIR, MSEIR etc [13] which capture differences between the way various diseases spread.

Conjecture 2 (Other Epidemic models). Our results for all the mobility models discussed in this paper hold for all the epidemic models covered in [24] as well.

6 Related Work

Here, we review the related work. It is worth pointing out that while most of existing studies about epidemic spread on mobile networks focus on (1) some particular types of network structures, and/or (2) one specific mobility model; our framework is very general and it applies to *arbitrary* network structure, and all the three popular mobility model.

General Epidemic Modeling. Bailey provides the canonical text on epidemic modeling [2]. A more recent survey is provided by Hethcote in [13]. Kephart and White [15, 16] were among the first to propose epidemiology-based models to analyze the spread of computer viruses. The model they suggest provides a good approximation of virus propagation in networks where contact among individuals is essentially homogeneous. Recent discoveries suggest real networks (including social networks [9], router and AS networks [10], and Gnutella overlay graphs [28]) follow a power-law structure instead, prompting a re-evaluation of the homogeneity assumption common in the works above.

Epidemics on Static Networks. Observation suggests that real networks are not homogeneous, rather, overwhelming evidence suggests real networks follow a power law structure instead. By introducing an underlying structure for a disease to spread, such




Fig. 6. Take-off Plot for System Matrix Eigenvalue of a Periodic Mobility Model. A_1 and A_2 had $N = 500$ nodes spread uniformly across a 200×200 simulation field. Notice the accurate our prediction (black arrow).

as a static network, removes the original homogeneous assumption pioneered by those reference above. Newman [21] studied the epidemic thresholds for multiple competing viruses on special, *random* graphs. Pastor-Satorras and Vespignani studied viral propagation for such power-law networks [20, 22, 23]. They developed an analytic model for the Barabási-Albert (BA) power-law topology [3]. However, their derivation depends on some assumptions which does not hold for many real networks [17, 10]. Pastor-Satorras et al. [23] also proposed an epidemic threshold condition, but this uses the “mean-field” approach, where all graphs with a given degree distribution are considered equal. There is no particular reason why all such graphs should behave similarly in terms of viral propagation. Chakrabarti et. al. [7] observe that epidemic threshold of an arbitrary graph can be captured in a single parameter, the first eigenvalue of the adjacency matrix $\lambda_{1,A}$. Their observation was rigorously confirmed in [8] and independently by [30]. We again leverage the above observations to formulate our solution in Sect. 3

Epidemics on Mobile Networks. Prompted by the emergence of mobile devices, such as Bluetooth-equipped smartphones, researchers introduced mobility to epidemic spread. Mickens et. al. were among the first to examine device-to-device spreading of malicious software in mobile ad hoc networks [19, 18]. In their work, they present a *queue*-based technique for the RWP model to overcome the limitations of the earlier homogeneous models of Kephart and White. In a similar work, Yan et. al. extend the observations of Mickens et. al. by examining additional mobility models and their effect on epidemic spreading of a SIS virus [31]. Their work is unique in that it models virus propagation, in detail, a Cabir-like bluetooth worm, including the Bluetooth stack and unique worm properties.

Mobility Models. The mobility models used in Sect. 4 are fairly common, with significant literature devoted to the subject. For an overview on mobility models, we refer our readers to the following surveys [6]. The RWP mobility model has been extensively used, despite well know flaws. For a discussion of the merits of RWP, refer to [32]. The Levy mobility model was first described in [29], yet has been used extensively to model human and animal movements [27, 26].

7 Conclusions

To conclude, recent malware in the wild, using device-to-device virus propagation schemes, prompted our study of the epidemic threshold in mobile ad hoc networks. Our contributions in this paper are:

1. *Framework*: We present a framework for the determining the epidemic threshold (for the SIS model) on *any* mobility model which can be converted into a series of adjacency matrices and give a formula for it (Theorem 1).
2. *Closed Formulas*: We also give a closed-form approximation for the SIS epidemic threshold on general mobility models (Lemma 1).
3. *Insensitivity to Velocity*: We analyze the impact of velocity in popular mobility models like Random walk, Levy Flight, Random waypoint etc. and find that it unexpectedly does not affect the threshold (Lemmas 2, 3 and Conjecture 1).

In addition, we introduced the “periodic mobility model,” popular in other fields like epidemiology [1, 4], to the networking community and solved it using our framework (Lemma 4). Finally we presented extensive simulations to demonstrate our analysis and results. Future work may concentrate on providing theoretical analysis on the effect of velocities on the steady state behavior of an epidemic on various mobility models.

References

1. Anderson, R.M., May, R.M.: Infectious Diseases of Humans. Oxford University Press (1991)
2. Bailey, N.T.J.: The Mathematical Theory of Infectious Diseases and its Applications. Charles Griffin and Company Ltd, London (1975)
3. Barabasi, A.L., Albert, R.: Emergence of Scaling in Random Networks. Science 286(5439), 509–512 (1999)
4. Barrett, C.L., Bisset, K.R., Eubank, S.G., Feng, X., Marathe, M.V.: EpiSimdemics: an Efficient Algorithm for Simulating the Spread of Infectious Disease over Large Realistic Social Networks. ACM/IEEE Conf. on Supercomputing (2008)
5. Bettstetter, C., Resta, G., Santi, P.: The Node Distribution of the Random Waypoint Mobility Model for Wireless Ad Hoc Networks. IEEE Trans. on Mobile Computing 2(3) (2003)
6. Camp, T., Boleng, J., Davies, V.: A Survey of Mobility Models for Ad Hoc Network Research. Wireless Communication and Mobile Computing pp. 483–502 (2002)
7. Chakrabarti, D., Wang, Y., Wang, C., Leskovec, J., Faloutsos, C.: Epidemic thresholds in real networks. ACM TISSEC 10(4) (2008)
8. Chakrabarti, D., Leskovec, J., Faloutsos, C., Madden, S., Guestrin, C., Faloutsos, M.: Information Survival Threshold in Sensor and P2P Networks. In: IEEE ICC (2007)
9. Domingos, P., Richardson, M.: Mining the Network Value of Customers. In: ACM SIGKDD '01. pp. 57–66. ACM Press, New York, NY, USA (2001)
10. Faloutsos, M., Faloutsos, P., Faloutsos, C.: On power-law relationships of the Internet topology. SIGCOMM pp. 251–262 (Aug-Sept 1999)
11. Ganesh, A., Massoulié, L., Towsley, D.: The effect of network topology in spread of epidemics. IEEE INFOCOM (2005)

12. Gonsalves, A.: Android Phones Steal Market Share. Information Week (04 2010), http://www.informationweek.com/news/mobility/smart_phones/showArticle.jhtml?articleID=224201881
13. Hethcote, H.W.: The mathematics of infectious diseases. *SIAM Rev.* 42(4), 599–653 (2000)
14. Johnson, D.B., Maltz, D.A.: Dynamic Source Routing in Ad Hoc Wireless Networks. In: *Mobile Computing* (1996)
15. Kephart, J.O., White, S.R.: Directed-Graph Epidemiological Models of Computer Viruses. In: *IEEE Security and Privacy*. pp. 343–359 (May 1991)
16. Kephart, J.O., White, S.R.: Measuring and Modeling Computer Virus Prevalence. In: *IEEE Security and Privacy 1993*. pp. 2–15 (May 1993)
17. Kumar, S., Raghavan, P., S.Rajagopalan, Tomkins, A.: Trawling the Web for Emerging Cyber-Communities. *Computer Networks* 31(11–16), 1481–1493 (1999)
18. Mickens, J.W., Noble, B.D.: Modeling Epidemic Spreading in Mobile Environments. In: *ACM WiSe '05*. pp. 77–86. ACM, New York, NY, USA (2005)
19. Mickens, J.W., Noble, B.D.: Analytical models for epidemics in mobile networks. *IEEE WiMob 2007* 0, 77 (2007)
20. Moreno, Y., Pastor-Satorras, R., Vespignani, A.: Epidemic Outbreaks in Complex Heterogeneous Networks. *The European Physical Journal B* 26, 521–529 (2002)
21. Newman, M.: Threshold Effects for Two Pathogens Spreading on a Network. *Phys. Rev. E* (2005)
22. Pastor-Satorras, R., Vespignani, A.: Epidemic Dynamics and Endemic States in Complex Networks. *Physical Review E* 63, 066117 (2001)
23. Pastor-Satorras, R., Vespignani, A.: Epidemic Dynamics in Finite Size Scale-Free Networks. *Physical Review E* 65, 035108 (2002)
24. Prakash, B.A., Chakrabarti, D., Faloutsos, M., Valler, N., Faloutsos, C.: Got the Flu (or Mumps)? Check the Eigenvalue! arXiv:1004.0060v1 [physics.soc-ph] (2010)
25. Prakash, B.A., Tong, H., Valler, N., Faloutsos, M., Faloutsos, C.: Virus Propagation on Time-Varying Networks: Theory and Immunization Algorithms. *ECML/PKDD* (2010)
26. Rhee, I., Shin, M., Hong, S., Lee, K., Chong, S.: Human Mobility Patterns and their Impact on Delay Tolerant Networks. In: *ACM HotNets IV* (2007)
27. Rhee, I., Shin, M., Hong, S., Lee, K., Chong, S.: On the Levy-Walk Nature of Human Mobility. In: *IEEE INFOCOM 2008*. pp. 924–932. IEEE (2008)
28. Ripeanu, M., Foster, I., Iamnitchi, A.: Mapping The Gnutella Network: Properties of Large-Scale Peer-to-Peer Systems And Implications For System Design. *IEEE Internet Computing Journal* 6(1) (2002)
29. Shlesinger, M.F., Klafter, J., Wong, Y.M.: Random Walks with Infinite Spatial and Temporal Moments. In: *J. Stat. Phys.* vol. 27, pp. 499–512 (1982)
30. Van Mieghem, P., Omic, J., Kooij, R.: Virus Spread in Networks. *IEEE/ACM Trans. Netw.* 17(1), 1–14 (2009)
31. Yan, G., Flores, H.D., Cuellar, L., Hengartner, N., Eidenbenz, S., Vu, V.: Bluetooth Worm Propagation: Mobility Pattern Matters! In: *ACM ASIACCS '07*. pp. 32–44. New York, NY, USA (2007)
32. Yoon, J., Liu, M., Noble, B.: Random Waypoint Considered Harmful. In: *INFOCOM 2003*. vol. 2, pp. 1312 – 1321 (2003)