Modeling Epidemic Spreading in Mobile Environments

James W. Mickens and Brian D. Noble
EECS Department, University of Michigan
Ann Arbor, MI, 48103
jmickens,bnoble@eecs.umich.edu

ABSTRACT

The growing popularity of mobile networks makes them increasingly attractive to virus writers, and malicious code targeting mobile devices has already begun to appear. Unfortunately, standard techniques for modeling computer virus propagation cannot be applied to mobile settings. We describe why these models fail and introduce a new framework called *probabilistic queuing* which treats node mobility as a first-order concern. A network is modeled by multiple queues which emulate the skewed connectivity levels common in mobile environments. Each queue represents a separate epidemiological population, and as nodes shuttle between queues, they bring their infections with them. Simulations show that for realistic mobility parameters, our model is more accurate than the standard Kephart-White framework.

Categories and Subject Descriptors: C.2.0 [Computer Systems Organization]: Computer-Communication Networks — *Security and Protection*, C.4 [Performance of Systems]: Modeling Techniques, G.3 [Mathematics of Computing]: Probability and Statistics

General Terms: Security, Theory, Algorithms

Keywords: Mobile networks, computer viruses, proximity attacks

1. INTRODUCTION

With the continuing proliferation of portable wireless devices such as laptops and PDAs, mobile networks are becoming an important part of our everyday networking infrastructure. However, the growth of mobile networking is leading to new security challenges. As the wired Internet became more popular, there was a corresponding surge in the amount of malicious code which used the Internet as its transmission mechanism. Similarly, as mobile networks become more common, they too will become attractive targets for virus writers. Just as boot sector viruses were supplanted by viruses that spread through email attachments and other Internet vectors [4], the emergence of widespread mobile networking will lead to new types of malicious code. Indeed, IBM's 2004 Business Security Report forecast that malware propagation amongst mobile devices would be an increasingly dangerous problem [10]. Devising epidemiological models for mobile environments is therefore an important research area.

Malicious code targeting mobile devices has already begun to emerge. For example, the Brador virus [21] infects Pocket PCs running Windows CE, installing a backdoor which allows a remote attacker unlimited access to the device. The Cabir worm [9] infects cell phones running the Symbian operating system. Taking control of the phone's Bluetooth interface, Cabir continually scans for other Bluetooth-enabled devices and tries to infect any such device which enters the scanning range. The Mabir [7] and Symbos_Comwar [15] worms use similar scanning techniques and also propagate via MMS messages.

Brador only spreads through traditional application-level vectors like email attachments and downloadable web objects. But via Bluetooth scanning, Cabir, Mabir, and Symbos_Comwar can launch *proximity attacks* upon vulnerable machines that are physically nearby. This means that epidemiological models for mobile networks must treat the movement of devices as a first-class concern. These models must also consider that, as shown in Figure 1, some geographic locations may be more heavily visited than others. These "hot spots" generate skewed connectivity distributions, since a node in a hot spot will have many more neighbors in communication range than a node in an unpopular location. Hot spots therefore represent more fertile breeding grounds for malicious code that uses proximity attacks. Viral propagation is also influenced by border effects. Since walls and physical obstacles can exclude nodes from large geographical areas, devices near these objects often have low connectivity and thus are poor vectors for viral infection. Epidemiological models for mobile networks must capture such wall phenomenon as well.

In this work, we investigate the behavior of malicious code like Cabir which spreads via proximity-based, point-to-point wireless links; we focus on this method of infection because it is unique to mobile environments and has received little research attention from the security community. This paper makes three primary contributions. First, it shows that naive application of standard epidemiological models to mobile environments leads to erroneous predictions. These mispredictions are often as severe as forecasting an endemic network-wide infection when the virus will actually die out quickly. Second, this paper explains why the standard models fail, namely, because they ignore node velocity and the non-homogeneous connectivity distributions that often arise in mobile environments. Third, this work proposes a new framework for understanding epidemics in mobile environments. This new model, called *probabilistic queuing*, explicitly incorporates notions of node mobility and connectivity skew. It provides an accurate threshold condition which relates the virulence of malicious code to the likelihood that it will cause an endemic network-wide infection. It also provides accurate estimates of these persistent infection levels.

This work was sponsored in part by NSF Grant CNS-0509089.

WiSE'05, September 2, 2005, Cologne, Germany.
Copyright 2005 ACM 1-59593-142-2/05/0009 ...$5.00.

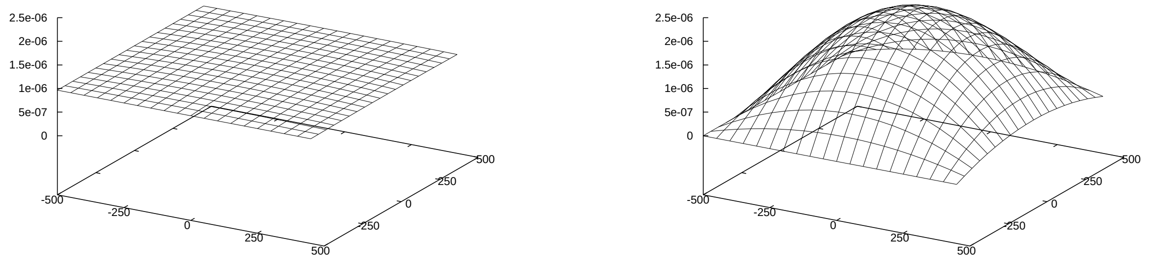

(a) In the random waypoint model, large pause times result in an effectively flat spatial distribution of nodes.

(b) As pause times shrink, the spatial distribution develops a pronounced peak; such peaks result in non-homogeneous connectivity distributions. In environments that are not governed by the random waypoint model, spatial peaks can arise because of obstacles or "popular content" regions that are frequently visited.

Figure 1: Spatial Distributions for a Square Arena Using the Random Waypoint Model

2. WHY THE KEPHART-WHITE MODEL FAILS

Before introducing our new framework, we describe the Kephart-White epidemiological model [13]. We then provide several examples that demonstrate the failure of the Kephart-White model in mobile environments. Our analysis of these failures will guide the design of probabilistic queuing.

2.1 The Kephart-White Model

The classic Kephart-White model [13] uses a differential equation to describe computer virus propagation. The model assumes a susceptible-infected-susceptible environment—a machine enters the system in a healthy state, and it can catch and subsequently be cured of the infection an infinite number of times. The Kephart-White (KW) model also assumes a homogeneous network topology in which all nodes have similar levels of connectivity or "out-degree." Thus, the network can be succinctly described by a single parameter $\langle k \rangle$ which represents the average connectivity of a node.

Defining I as the fraction of nodes infected at a particular moment, the KW model uses the following equation to describe viral propagation:

$$\frac{dI}{dt} = \beta \langle k \rangle I(1-I) - \delta I \quad (1)$$

where t is time, β is the viral birth rate along every edge from an infected node, and δ is the cure rate at each infected node. β, δ, and $\langle k \rangle$ are assumed to be constant. The KW equation has a steady state solution of:

$$I = 1 - \frac{\delta}{\beta \langle k \rangle} \quad (2)$$

However, such an endemic infection occurs only if:

$$\beta \langle k \rangle > \delta \quad (3)$$

In other words, an epidemic arises only when the expected viral output of an infected node is greater than the probability that an infected node will be cured.

In the mobile setting, $\langle k \rangle$ represents the average number of devices within wireless communication range of an arbitrary node. β represents the probability that a diseased node transmits the infection to a healthy neighbor during some small time period Δt. δ represents the probability that an infected node is cured during Δt. When we graph the global percentage of infected nodes versus time, the time axis will be in units of the viral time scale Δt. In this paper, we always use a Δt of 100 milliseconds.

2.2 The Kephart-White Model in Mobile Environments

Suppose that each mobile device has a communication range of 100 meters and travels in a square arena with 1000 meter sides. Further suppose that node movement is guided by the random waypoint model [5], pause time is 0, and that through simulation or mathematical analysis, we can derive $\langle k \rangle$ for the network. Figures 2 and 3 demonstrate several ways in which the KW model will be inaccurate for this mobile environment. Note that each simulation result represents the average of five runs, and all simulations began with node speeds and locations drawn from the appropriate steady-state distributions [16].

First consider Figure 2, which shows results for a 30 node and 60 node mobile network. In both networks, node speeds were drawn from the range $[5m/s, 20m/s]$. After 200,000 simulated seconds, we find that $\langle k \rangle_{N=30}$ is 1.22 and $\langle k \rangle_{N=60}$ is 2.37. Given these connectivities, a virus with β=0.7 and δ=0.25, and 10% of nodes initially infected, the KW model predicts high endemic infection rates in both networks. However, the simulation results disagree. In the 30 node network the KW model predicts an endemic infection of 70.7%, but the infection actually dies out completely. In the 60 node network a persistent infection emerges, but it has a level of roughly 64%, not 84.9% as predicted by the KW model.

In this example, the failure of the KW model is largely due to its strict reliance on the average connectivity statistic. Only considering mean connectivity discards useful information when the underlying distribution has significant variance. Figure 4 shows the connectivity distributions for the 30 node and 60 node sce-

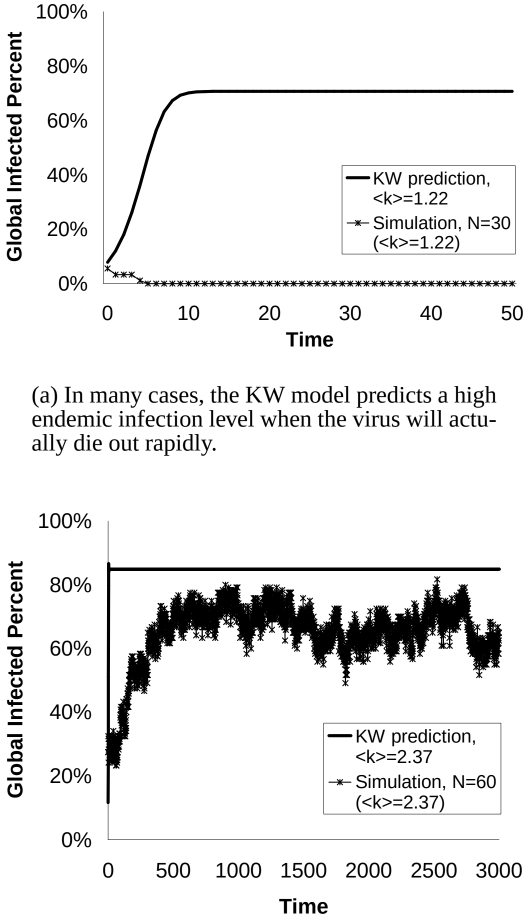

(a) In many cases, the KW model predicts a high endemic infection level when the virus will actually die out rapidly.

(b) In other situations, the KW model correctly forecasts a persistent infection but overpredicts its magnitude.

Figure 2: The failure of KW predictions — over-reliance on connectivity averages

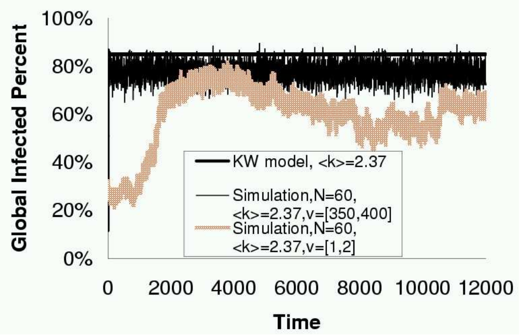

The KW model does not have a parameter for node speed. Thus, it cannot distinguish between two networks with the same connectivity distribution but different node velocities. Higher node velocities lead to greater node mixing and more virulent epidemics.

Figure 3: The failure of KW predictions — no conception of node speed

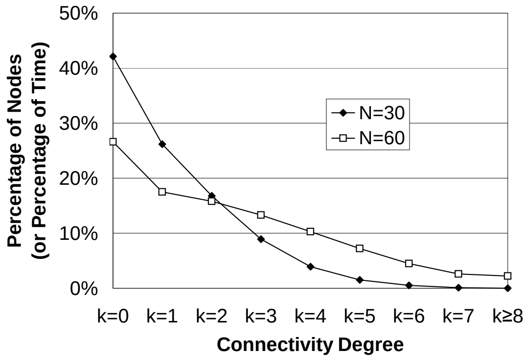

These are the connectivity probabilities for random waypoint mobile networks with 30 and 60 nodes; each node has a 100 meter communication range, and the arena is 1000 meters by 1000 meters. Each probability represents the likelihood that a node has the given connectivity at an arbitrary moment; alternatively, it represents the amount of time that a node spends with the given connectivity.

Figure 4: Connectivity Distributions for N=30 and N=60

narios; these distributions were generated analytically using techniques from Bettstetter [1] that we summarize in Section 3.1. We see that in the 30 node network, a device spends 42.1% of its time with no neighbors. In contrast, a device in the 60 node scenario is neighborless for only 26.6% of the time. These differences in disconnected time lead to differences in epidemic behavior. As the disconnected fraction grows, sick nodes have more opportunities to be cured without threat of concurrent reinfection. Similarly, there are fewer opportunities for sick nodes to infect healthy ones.

The KW model cannot detect this difference in disconnected fraction—it only sees a reduced $\langle k \rangle$ in the 30 node network relative to the 60 node network. However, the homogeneity assumption of the KW model is broken by the distributions shown in Figure 4. Sometimes a node will have more neighbors than $\langle k \rangle$, but other times it will have fewer or, importantly, none at all. Due to this latter occurrence, the predicted endemic infection rates from the KW model are depressed or even reduced to zero in real life.

In Figure 3, we show another problem with applying the KW model to mobile environments. We study two networks net_{slow} and net_{fast} in which the arena size and virus profile are the same as in the previous examples. Each network has 60 nodes, but in net_{slow} node speeds are drawn from $[1, 2]$, while in net_{fast} speeds are taken from the range $[350, 400]$. The KW model has no parameter for node velocity and predicts a steady state infection level of 84.9% for both networks. However, net_{slow} actually has an endemic infection level of about 56%, whereas net_{fast} has an endemic infection level of about 74%. This is despite the fact that the two networks have the same connectivity distribution, and despite the fact that they are being attacked by malicious code with the same β and δ.

How can we explain such velocity-dependent differences in effective virulence? Intuitively speaking, nodes in net_{fast} are “better mixed” than nodes in net_{slow} — during a given stretch of time, they communicate with a wider variety of neighbors than the devices in net_{slow}. Higher mixing rates boost viral propagation, since diseased nodes will have more opportunities to communicate with healthy ones. Also note that while nodes in the two networks spend the same percentage of time with zero neighbors, devices in

net_{slow} have longer *uninterrupted* periods of disconnection. When velocities are low, a node that wanders into an empty part of the arena will stay there for a while. This gives the node many consecutive opportunities to be cured without the threat of concurrent infection. In net_{fast}, such windows are much smaller in terms of raw temporal duration.

In summary, the examples from Figures 2 and 3 demonstrate the two problems with applying the KW model to mobile networks. First, since the KW model relies on mean connectivity as its sole topological metric, it cannot capture the non-trivial connectivity variances found in mobile environments. Second, the KW model is insensitive to node speed, which is an essential parameter in mobile networks.

3. A NEW MODEL FOR EPIDEMICS IN MOBILE NETWORKS

To remedy the problems with applying the KW model to mobile environments, we propose a new analytic framework. Our *probabilistic queuing* model explicitly accounts for both node velocities and the non-homogeneous connectivity patterns induced by this mobility.

3.1 Mathematical Background

To create a probabilistic queuing system, we first must characterize the mobility parameters of the underlying network. As a concrete example, we summarize Bettstetter's framework for describing mobility in random waypoint networks [1, 2, 5]. However, we emphasize that probabilistic queuing is agnostic to the choice of mobility model, and we show in Section 4.3 that it still outperforms the KW model for networks that are not governed by random waypoint movement.

In the random waypoint model, nodes travel within a large arena, typically either a rectangle or a circle. Each node iteratively picks a random destination, travels there, pauses for a constant time t_{pause}, and then picks another random destination. Each waypoint is independently chosen, and before leaving for a new waypoint, a node chooses a random speed from the uniform distribution $[v_{min}, v_{max}]$.

Given a square arena having sides of length a, the average trip length is:

$$E\{L\} = 0.5214a \tag{4}$$

and the average time needed to complete such a trip is:

$$E\{T\} = \frac{ln(v_{max}/v_{min})}{v_{max} - v_{min}} E\{L\} + t_{pause} \tag{5}$$

Allowing $p_p = \frac{t_{pause}}{E\{T\}}$, the spatial distribution function over $-a/2 \leq x, y \leq a/2$ is:

$$sdf(x,y) \approx \frac{p_p}{a^2} + (1-p_p)\frac{36}{a^6}(x^2 - \frac{a^2}{4})(y^2 - \frac{a^2}{4}) \tag{6}$$

Examples of this function are depicted in Figure 1. Given that a node is at some location (x_i, y_i), the probability that it is within communication range of another random node is:

$$c(x_i, y_i) = \int_{y_i - r}^{y_i + r} \int_{x_i - \sqrt{r^2 - (y-y_i)^2}}^{x_i + \sqrt{r^2 - (y-y_i)^2}} sdf(x,y)\; dx\; dy \tag{7}$$

where r is the communication radius of a wireless radio and is the same for all nodes. The average probability over the entire arena that two randomly placed nodes will be within communication range is:

$$\bar{c} = \frac{\int_{-a/2}^{a/2} \int_{-a/2}^{a/2} c(x,y)\; dx\; dy}{a^2} \tag{8}$$

The probability that a node at location (x, y) has k_i neighbors is given by:

$$Pr(x, y, k = k_i) = \left.\begin{matrix} N-1 \\ k \end{matrix}\right) c(x,y)^k (1 - c(x,y))^{N-k-1} \tag{9}$$

where N is the total number of mobile nodes. The average probability over the entire arena that a node has k neighbors is:

$$Pr(k = k_i) = \frac{\int_{-a/2}^{a/2} \int_{-a/2}^{a/2} Pr(x, y, k = k_i)\; dx\; dy}{a^2} \tag{10}$$

We can interpret $Pr(k = k_i)$ for $k_i \in [0, N-1]$ as the percentage of time that a node has k_i neighbors.

3.2 A New Epidemic Threshold

Given a set of mobility parameters which describe an ad hoc communication topology, our most basic question involves the epidemic threshold: how virulent must malicious code be to create an endemic infection amongst the mobile devices? More specifically, given values for a, N, v_{min}, v_{max}, and r, what values of β and δ lead to persistent global infections?

The standard Kephart-White model predicts endemic infection when $\beta\langle k \rangle > \delta$; in other words, epidemics occur when the infection pressure overwhelms the cure pressure. As shown in Figure 2, this epidemic threshold is not always accurate in mobile networks. The key problem is that the KW threshold ignores mobility and thus misses the impact of node speed on infection dynamics.

To remedy this problem, we must explicitly consider the connectivity fluctuations induced by mobility. Consider a particular node traveling in an arena containing many other nodes. $E\{T\}$ represents the expected time that it takes the node to travel between two waypoints. The line segment between consecutive waypoints can be conceptualized as a queue or pipe which takes $E\{T\}$ seconds to traverse. Using Equation 10, we can determine the expected percentage of queue time that is spent with a given number of neighbors. More specifically, for $E\{T\} * Pr(k = 0)$ time units, the node has no neighbors and is only subject to curing attempts [1]. For $E\{T\} * Pr(k = (k_i > 0))$ time units, the node is subjected to infection pressure proportional to βk_i and cure pressure proportional to δ. If the cumulative infection pressure in a queue is greater than the cumulative cure pressure, the node will likely be sick for the majority of its queue time, and it will be capable of infecting its neighbors for the majority of its queue time. Conversely, if the cumulative infection pressure is less than the cumulative cure pressure, we expect the node to spend the majority of its queue time in a healthy state.

These observations suggest that for an epidemic to occur in a mobile network, the following condition must be true:

$$\sum_{k_i=0}^{N-1} \beta k_i Pr(k = k_i) E\{T\} > c\delta E\{T\} \tag{11}$$

[1] Note that $E\{T\}$ must be expressed in terms of the viral time scale. For example, if the infection and cure probabilities are defined over 100 millisecond intervals, then $E\{T\}$ must be expressed in units of 100 milliseconds.

where the left-hand side represents the infection pressure over a travel segment and the right-hand side represents the cure pressure. The small constant c accounts for stochastic fluctuations in global connectivity. Since node movement is random and uncoordinated, there will be punctuated periods of time during which most nodes have very few neighbors or none at all. For an epidemic to persist, the virus must be strong enough to ensure that a critical mass of diseased nodes will emerge from such periods with their infections intact. In Section 4.1, we empirically observe that $c \approx 3.5$ for random waypoint networks.

Reconsider our examples of net_{slow} and net_{fast}. Although both networks have the same connectivity distribution, net_{slow} has a larger $E\{T\}$ than net_{fast} and thus longer travel queues. Now we can capture the fact that, for a given connectivity level, nodes in net_{slow} have this level for a longer absolute time period per segment traversal.

3.3 Making Steady-State Predictions

Having established an epidemic threshold condition for mobile networks, our next task is to predict what the endemic infection level will be. More specifically, given a virus profile (β, δ) and mobility parameters a, N, v_{min}, v_{max}, and r, what is the global steady-state infection level?

To answer this question, we extend the concept of travel queues. In the previous section, we investigated the behavior of an individual node. To travel between two waypoints, the node entered a queue at some time t and exited the queue at time $t + E\{T\}$. Using the connectivity distribution generated by Equation 10, we considered the fraction of $E\{T\}$ that was devoted to a particular connectivity level. However, we can interpret the connectivity distribution in an alternate way, using it to tell us the global percentage of nodes having a given connectivity level at an arbitrary moment. From this perspective, there are $N * Pr(k = k_i)$ nodes with connectivity k_i at any given time. If we make the simplifying assumption that the uninterrupted stretches of time that a node has a particular connectivity level are large relative to the viral Δt, we can model the mobile network as a set of N queues. Each queue Q_{k_i} contains $N * Pr(k = k_i)$ nodes. Upon entering Q_{k_i}, a node spends $E\{T\} * Pr(k = k_i)$ time units in it before exiting.

Epidemiologically, we treat each queue as a separate Kephart-White population described by the global (β, δ) and a local $\langle k \rangle$ equal to k_i. Such assumptions of local connectivity homogeneity are intuitively justifiable—if a mobile node has k_i neighbors, many of these neighbors are likely to be in communication range with each other and have roughly k_i neighbors as well. The connectivity distribution from Equation 10 tells us both the size of these "neighborhoods" and the length of time that a node stays in each neighborhood.

To initialize the queuing model, we set the integer system time to 0 and insert $N * Pr(k = k_i)$ nodes into each Q_{k_i}. We stamp each queue's initial node set with an exit time of $E\{T\} * Pr(k = k_i)$. Each queue's infection level is then set to some $I_{initial} \in [0.0, 1.0]$. After this initialization, we iteratively update the system clock in increments of the viral Δt, performing two tasks after each update. First, we simulate Kephart-White dynamics in each queue, such that $dI_{Q_{k_i}}/dt = \beta k_i I_{Q_{k_i}}(1 - I_{Q_{k_i}}) - \delta I_{Q_{k_i}}$. Second, we check whether any node sets have exit times less than or equal to the current time. If so, we remove the node set from its queue, divide it into N equally sized pieces, and enqueue one of these pieces into each Q_{k_i}. Finally, each queue updates its infected percentage $I_{Q_{k_i}}$ to reflect its newly enlarged population and the infection percentage of the just-enqueued nodes. At any moment, the global number of infected nodes is equal to $\sum_{k_i=0}^{N-1}[[I_{Q_{k_i}}] * [Q_{k_i}.length]]$. To predict the steady state infection percentage, we simply iterate the queue maintenance algorithm until the global number of infected nodes has stabilized. Figure 5 provides a visual depiction of a probabilistic queuing system.

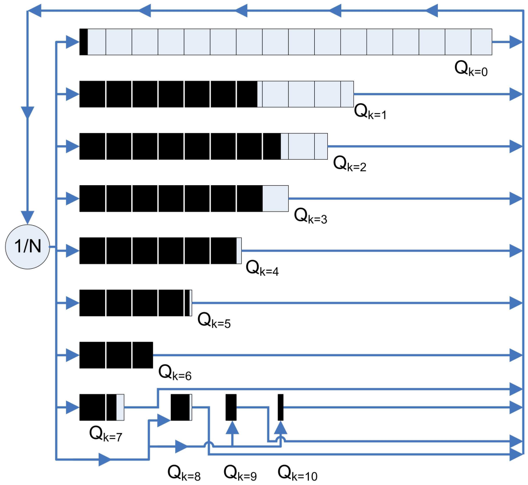

This is the steady state of a probabilistic queuing system for N=60, a=1000, r=100, $v \in [5, 20]$, β=0.7, and δ=0.25. The number of squares in a queue corresponds to the number of nodes it contains. The proportion of dark squares to light squares represents the ratio of infected nodes to healthy nodes in the queue. For example, $Q_{k=0}$ stores 15.86 nodes, of which 0.39 are infected; $Q_{k=7}$ only stores 1.59 nodes, but 1.41 of them are infected. The global number of infected nodes is the sum of the infected count in each queue. In this example, 37.57 out of 60 nodes are infected.

Figure 5: Modeling epidemics using probabilistic queues

Although we partition dequeued node sets into equally sized pieces, each piece will spend a different time waiting in its next queue. Since the time spent in Q_{k_i} is proportional to $Pr(k = k_i)$, our partitioning scheme simulates the effects of mobility and non-homogeneous connectivity distributions. For very large k_i, $E\{T\} * Pr(k = k_i)$ may be shorter than the viral timescale, i.e., nodes pushed into such queues should stay in the queues for less than the viral Δt. Since the maintenance algorithm steps in increments of Δt, it cannot accurately model such queues. Thus, we collapse the queues for such large $k_{high1}, k_{high2}..., k_{N-1}$ into a single queue whose connectivity is the average of these values and whose traversal time is $[E\{T\} * Pr(k \geq k_{high1})] > \Delta t$.

4. EVALUATION

To evaluate probabilistic queuing and the KW model in mobile environments, we wrote a custom simulator which gave us fine-grained control over mobility parameters and viral profiles. Each simulation took place in a square arena with 1000 meter sides. Unless otherwise noted, node movement was guided by the random waypoint model. To emphasize the impact of mobility on viral propagation, we typically used pause times of zero. Each mobile device had a 100 meter communication radius, which corresponds to the range of a Class 1 Bluetooth radio. The simulator did not model path effects or transmission interference. For each (β, δ) pair, the viral Δt was 100 milliseconds. In the figures presented in this section, each simulation result represents the average of five trials. Each trial ran for a maximum of 200,000 virtual seconds, terminating early if the virus was completely extinguished before this time period elapsed. Simulations were initialized with asymptotic node positions and velocities [16]. For evaluation metrics, standard deviations are often given in parentheses.

4.1 The Epidemic Threshold

Given a virus profile, a particular set of mobility parameters, and an epidemiological model, we define two evaluation metrics for the threshold condition: the *raw accuracy* and the *threshold misprediction penalty*. Raw accuracy refers to the percentage of predictions that were correct, i.e., an epidemic was predicted and one emerged, or an epidemic was not predicted and one did not emerge. We define the threshold misprediction penalty as follows. If the threshold condition correctly predicts whether an epidemic arises, the prediction has a penalty of 0.0%. Otherwise, the penalty is the actual endemic infection level if no epidemic was predicted, or the predicted endemic infection level if an epidemic was predicted. When comparing two epidemiological models, the one with the higher raw accuracy is best at predicting whether a virus will die out. Large misprediction penalties indicate that when a model mispredicts the emergence of an epidemic, it forecasts big epidemics that never materialize, or no epidemics when big ones actually arise. An epidemiological model could have both high raw accuracy and large misprediction penalties.

In Figure 6, we compare the threshold predictions of the KW model and the probabilistic queuing model for several viral profiles. Node speeds were drawn from $v \in [5, 20]$, and we display results for N values of 30, 40, 50, and 60. These N values are the most useful ones for evaluating epidemic thresholds because with a 1000 meter by 1000 meter arena and communication ranges of 100 meters, the critical mass of nodes needed to sustain an infection is typically between 30 and 60. In Figure 6, the probabilistic queuing threshold is more accurate by (13/16)=81% versus (4/16)=25%.

We evaluated the two threshold conditions with all permutations of $N \in [30, 40, 50, 60, 80, 100]$, $v \in [[1, 2], [5, 20], [350, 400]]$, and $\beta, \delta \in$[(0.7,0.25), (0.5,0.25), (0.25,0.25), (0.04,0.01), (0.02, 0.01), (0.01,0.01)]. The KW threshold had an accuracy of 58.9% and an average misprediction penalty of 59.7%(±19.5%), whereas the probabilistic queue threshold had an accuracy of 80.9% and an average misprediction penalty of 55.0%(±6.6%). Relative to its performance in Figure 6, the KW threshold improved due to two reasons. First, the KW threshold was usually correct in the $N = 80$ and $N = 100$ cases, which were much easier to predict than the others. Second, KW threshold performance sometimes approaches or surpasses that of probabilistic queuing for $v \in [350, 400]$. This is because the time scale assumption from Section 3.3 is no longer true. At such fast velocities, there is very high node mixing, so the uninterrupted stretches of time that a node spends at a particular connectivity level are no longer large compared to the viral Δt. This degrades the fidelity of the queuing abstraction as a realistic approximation of connectivity fluctuation. We discuss this phenomenon in more detail in Section 4.2 and describe how to restore the time scale property.

In general, the probabilistic queuing threshold does much better than the KW threshold for realistic node speeds, i.e., $v \in ([1, 2], [5, 20])$. It particularly excels when the KW model predicts a weak endemic infection. Such scenarios arise when the viral birth force across a single link is on par with the cure force. For example, with β=0.01, δ=0.01, and $v \in [5, 20]$, the queue threshold is 88.9% accurate whereas the KW threshold is only 27.8% accurate.

To investigate the impact of pause time, we simulated a mobile network with $v \in [5, 20]$, β=0.5, δ=0.25, and pause times of 100, 1000, or 5000 seconds; we used the same N values given above. In networks with pause time, the queuing threshold had a raw accuracy of 70.0% whereas the KW threshold had an accuracy of 25.9%. The queueing threshold had an average misprediction penalty of 58.1%(±4.0%), and the penalty for the KW model was 67.7%(±10.1%). Thus, non-zero pause times slightly reduce the accuracy of the queuing threshold but greatly reduce the accuracy of the KW threshold. The reason is that as pause time increases, node mixing diminishes, meaning that cliques of nodes spend longer amounts of time together. The KW model cannot capture this effect but our queuing model can, since pause times increase $E\{T\}$ and thus the raw time that nodes spend in each queue.

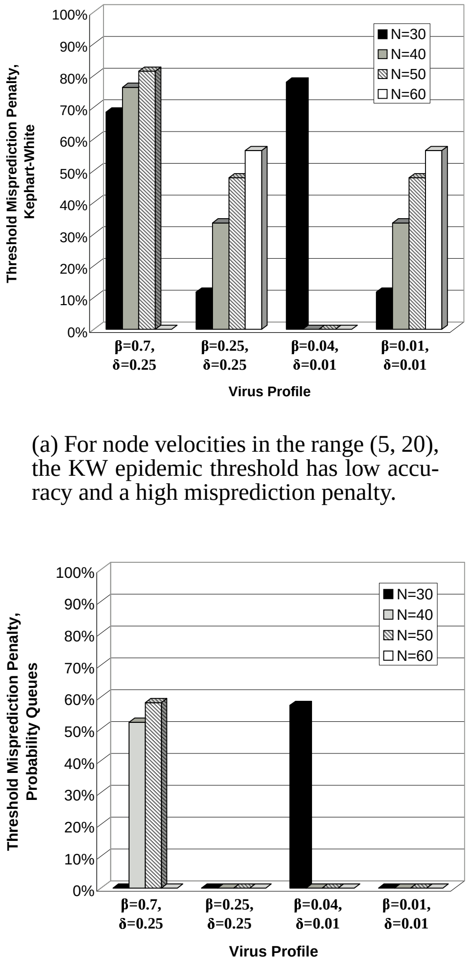

(a) For node velocities in the range (5, 20), the KW epidemic threshold has low accuracy and a high misprediction penalty.

(b) For the same velocity range, the probabilistic queuing threshold is much more accurate. The misprediction penalty is moderately smaller.

Figure 6: Comparative Accuracies of Epidemic Threshold Conditions

Recall that our threshold condition uses the constant c to represent worst-case stochastic fluctuation in network-wide connectivity. In the results reported above, we use a c value of 3.5. This value was empirically derived. However, we would like for c to be analytically derived from the variance of the connectivity distribution and the other mobility parameters. Generating such a formula is an important area for future work.

Note that our threshold condition can predict "no epidemic" even though our probabilistic queuing system stabilizes to a non-zero infected percentage. The reason is that the threshold condition explicitly (if clumsily) accounts for punctuated drops in global connectivity via the parameter c. The maintenance algorithm for the probabilistic queues does not mimic these rare yet important deviations. Interestingly, this means that the queues settle to the steady state infection level that "would have resulted" if no drastic connectivity fluctuations had occurred. In these scenarios, manual inspection of the simulation traces often reveals the infection level stochastically fluctuating around the queuing steady state before rapidly falling to zero at a random moment. If the mobile network has few nodes, the drop-off usually happens quickly and the steady state infection level in the queues is of little utility. However, as the number of mobile devices grows, bursts of extremely low global connectivity become less frequent. Even though the infection may still eventually die off, the queuing steady state often provides a good estimate of the global infection percentage before this extinction occurs. Incorporating such punctuated fluctuations in the queue maintenance algorithm is another important area for future research.

4.2 Steady State Predictions

Using a threshold condition, we can predict whether an endemic infection will occur. If we forecast that an epidemic will arise, we would like to estimate its magnitude as a function of the mobility parameters and the virus profile. In Figure 7, we give several examples of such predictions from the KW model and probabilistic queuing. For reasonable network parameters, probabilistic queuing is much more accurate than the KW model.

Suppose that an endemic infection actually occurs for a set of mobility parameters and a viral profile. We define an epidemiological model's *steady state prediction error* as the absolute difference between the predicted and observed infection percentage. For the KW model, this statistic is only defined when the KW threshold condition is satisfied, since only then is a non-zero steady state prediction made. For a probabilistic queuing system, we define the error whenever an infection actually arises, regardless of whether the queuing threshold condition is satisfied. This is because the queuing system can stabilize to a non-zero infection level despite the threshold condition failing, and the steady state error should be independent of threshold condition accuracy.

For pause times of zero and the mobility parameters investigated in the previous section, the KW model had an average steady state error of 12.5% (±7.4%). The queuing model only had an error of 4.0% (±3.6%). When considering the positive pause time scenarios described in Section 4.1, the KW model had an average steady state error of 27.0% (±7.6%), as compared to 9.6% (±7.3%) for the queueing model.

The specific example of N=60, β=0.5, and δ=0.25 in Figure 7(a) offers an instructive insight into viral dynamics in mobile networks. For these particular parameter values, *unstable epidemics* ensue. In five simulation runs of 200,000 virtual seconds, one epidemic died immediately, one lasted the entire 200,000 seconds, and the others lasted between one-fourth and three-fourths of the maximum possible time. In the graph, the simulated epidemic level is the average of the infection levels while the virus was still alive in each trial. Probabilistic queuing accurately predicts this average, but such an average hides an interesting temporal instability. Ideally, we would like to describe endemic infections using Markov model probability distribution functions, which have previously been applied to epidemics atop homogeneous topologies [3]. Given the mobility parameters, the virus profile, and an initial set of infected nodes, such a model would describe the probability of each possible epidemic level at a particular point in the future. However, it is unclear how to incorporate notions of mobility into these kinds of frameworks.

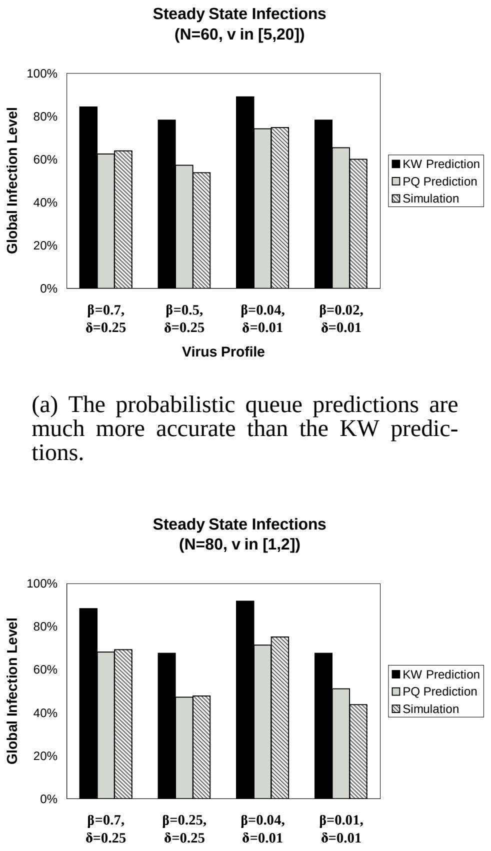

(a) The probabilistic queue predictions are much more accurate than the KW predictions.

(b) The KW model continues to predict steady state infection levels which are too high.

Figure 7: Endemic infection for common network parameters

As described in Section 3.3, the probabilistic queuing model assumes that the uninterrupted length of time that a node spends at a particular connectivity level is large compared to the viral Δt. This time scale property ensures that a node's uninterrupted journey through a connectivity queue accurately reflects its real-life fluctuations in connectivity. Unfortunately, the time scale assumption is violated when nodes move extremely quickly or the network contains a very large number of nodes. For example, nodes moving at 400 m/s might have the same connectivity distribution as nodes in a different network moving at 5 m/s. However, nodes in the slow network keep the same neighbors (and thus the same connectivity levels) for longer stretches of time; thus, the fast nodes switch between different connectivity levels much quicker. Similarly, a network with very many nodes offers more opportunities for connectivity flux than a sparsely populated network.

In common mobile environments, the time scale assumption will hold. For example, if the network of interest represents people using PDAs to query an interactive museum, or laptop users com-

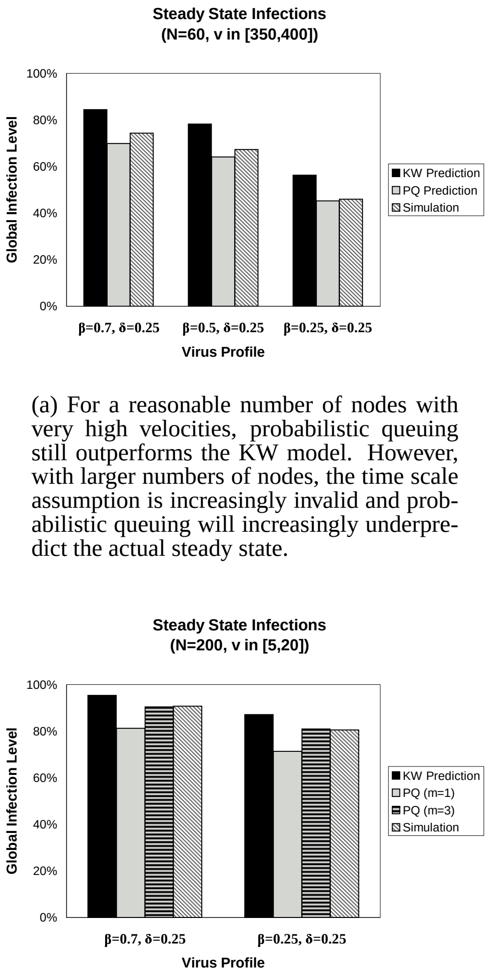

(a) For a reasonable number of nodes with very high velocities, probabilistic queuing still outperforms the KW model. However, with larger numbers of nodes, the time scale assumption is increasingly invalid and probabilistic queuing will increasingly underpredict the actual steady state.

(b) The time scale property is also violated in this example, which features reasonable node velocities but a very large number of nodes. To restore the time scale property, we must simulate the increased node mixing.

Figure 8: Epidemics when the time scale property is violated

municating with a wireless access point, node velocities and spatial densities will satisfy the time scale assumption. However, one could imagine scenarios which violate the property, such as networks deployed across automobiles. The key insight is that even though the time scale assumption is violated, the fundamental intuition behind probabilistic queuing remains valid. In these scenarios, queues still simulate skewed connectivity distributions and node mobility in these scenarios, although they do not capture the proper level of node mixing. In Figure 8, we see that for a reasonable number of nodes moving at a very high speed, the queuing model begins to consistently underpredict the steady state infection level. The problem worsens as the number of nodes increases. It also worsens as β grows larger relative to δ, since the infection rate (but not the cure rate) is sensitive to the number of neighbors encountered per unit time.

How can we modify probabilistic queuing to simulate increased mixing rates? The simplest solution is to divide the time spent in each queue by some *mixing constant*, denoted m. This division does not alter the underlying connectivity distribution of the network, since the proportion of queue sizes to each other is unchanged. The division merely increases the rate at which queues exchange nodes, which is the precise analogue of the increased mixing found in very dense or very fast networks. Figure 8(b) provides an example. In a network with 200 nodes and $v \in [5, 20]$, the KW model overpredicts the endemic infection and probabilistic queuing underpredicts. However, using an m of 3, the probabilistic queue predictions are very close to the actual results.

As with the threshold constant c, we currently lack an analytic formula which derives m from a given set of mobility parameters. We derived $m = 3$ for Figure 8(b) empirically. However, the introduction of m is not merely a “hack” to get the model to work for this particular example. Using an m of 3 for N=200 and $v \in [5, 20]$, the accuracy of the probabilistic queuing model increased for all permutations of β and δ. This suggests that the notion of node mixing is a fundamental and important property of mobile networks. Furthermore, mixing is a useful concept beyond the study of viral propagation. For example, consider a set of mobile sensor nodes. One might want to guarantee that each zone of the arena is always covered by at least one node, but one might also want to guarantee that each node directly communicates with each other node within some bounded amount of time. The former characteristic is governed by the spatial distribution function, but the latter is governed by the degree of node mixing. Generating an analytic form for m is an important area for future research.

4.3 Strongly Non-homogeneous Spatial Distributions

Up to this point, we have studied viral dynamics in random waypoint networks. As shown in Figure 1, when pause time is large, the spatial distribution is relatively flat and thus the connectivity distribution is fairly homogeneous. As pause times approach zero, connectivity becomes skewed, but the skew is still smooth and symmetric about the center of the arena. In real-life mobile networks, some locations may be much more popular than others. A natural question is “how do viruses spread when spatial distributions are very strongly skewed and asymmetric?” In Figure 9, we show an example of such a strongly non-homogeneous topology. In this network, nodes are highly attracted to one of three hotspots located at (-3a/8,-3a/8), (-3a/8, 3a/8), and (a/4,a/4). When a node picks a new waypoint, it chooses one of these locations with probability 15% per hotspot and a random location with probability 55%. As depicted in Figure 9, there are three spatial density spikes at the hotspot locations. The paths between these hotspots are also well-traveled.

Using simulations, we determined the connectivity distribution for this network for various values of N; examples of these distributions are given in Figure 10. Via simulation, we also determined that $E\{L\}$ was 544 meters and $E\{T\}$ was 50.2 seconds or 502 time units with respect to the viral Δt. Armed with these parameters, we constructed the corresponding probabilistic queuing system and compared its performance to that of the KW model.

Figure 11 depicts epidemic behavior for β=0.7, δ=0.25, and $v \in [5, 20]$. For each value of N, the KW threshold and the queuing threshold predict an endemic infection; however, no epidemic actually arises for N=40, and unstable epidemics arise for N=60. In all

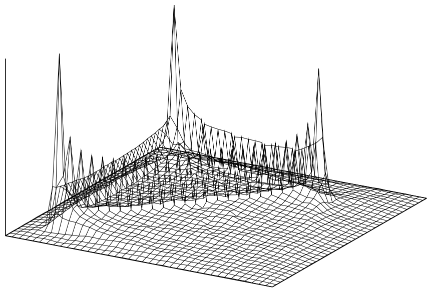

Figure 9: Spatial distribution function of a strongly non-homogeneous topology with three attractors

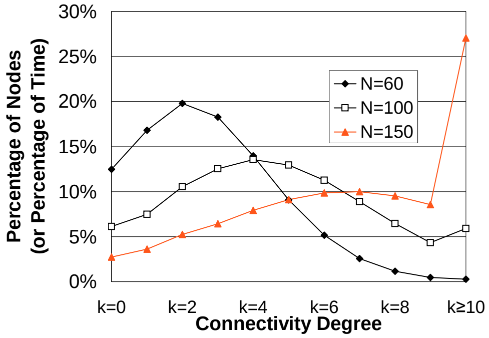

Figure 10: Sample connectivity distributions for spatial attractor example

cases, our queuing model produces more accurate steady state predictions than the KW model. However, the relative improvement is not as large as it is for random waypoint networks with similar mobility parameters. The reason is that the connectivity distribution alone can only hint at the "spikiness" in the underlying spatial distribution function. One potential method to capture these peaks is to divide the arena into zones and associate a separate connectivity queue with each zone. A queue could only exchange nodes with queues in adjacent zones, and these exchanges would be in proportion to each zone's spatial density. This would mimic the preferred movement pathways induced by strong attractors. We are currently investigating the properties of this grid model.

4.4 Epidemics in Class 2 Bluetooth Networks

Our study of mobile network epidemics has focused on devices with 100 meter communication ranges. This range corresponds to the transmission capability of a Class 1 Bluetooth radio. Class 2 radios with ranges of 10 meters are also popular. However, networks composed of Class 2 devices will have extremely low connectivities unless node density is very high. For example, in a Class 2 random waypoint network containing 100 nodes in a 1000 meter square arena, 96.2% of the devices will have zero neighbors at an arbitrary moment. Even for a 1000 device network, 69.4% of a node's time will be spent with no neighbors. These Class 2 networks will be impervious to all but the most virulent malicious code. Such code would have to be extremely aggressive in scanning for vulnerable neighbors and exceptionally difficult to purge. For these reasons, effective viruses in Class 2 networks will likely eschew point-to-point contact as the primary infection vector. For example, they might rely on spreading via email attachments, leaping onto a user's mobile device when he synchronizes it with his PC.

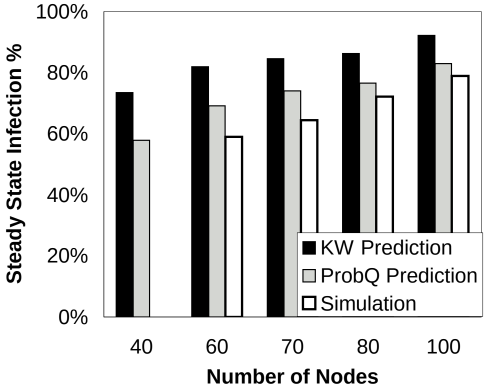

Figure 11: Steady state infection levels, spatial attractor network

5. RELATED WORK

The Kephart-White model [13] is the standard framework for studying computer viruses in homogeneous topologies. As we have shown in Section 2, it is inappropriate for modeling epidemics in non-homogeneous mobile networks. Researchers have found that some non-homogeneous computer networks have connectivities guided by power laws [8], where $Pr(k = k_i) = k^{-\gamma}$ for some $\gamma \in [2, 3]$. There are several epidemiological models for such power-law topologies [17, 18]. However, such research is not applicable to mobile networks for several reasons. First, mobile network connectivity does not typically follow a power law. For example, in a 1000 meter square arena containing 80 nodes doing random waypoint travel, $Pr(k = 1) = 0.174$ and $Pr(k = 4) = 0.108$; connectivity is non-homogeneous, but not to an exponential degree. If the arena has spatial attractors which pack a few nodes into a small region, connectivities may resemble those of a power-law network. However, epidemic frameworks for power-law networks do not incorporate notions of node mobility, and we have shown that ignoring mobility results in erroneous predictions.

Given the adjacency matrix **A** of an arbitrary communication topology, the epidemic threshold can be expressed as $1/\lambda_1$, where λ_1 is the largest eigenvalue of **A** [20]. However, in mobile environments, the adjacency matrix and its associated eigenvalues change over time. It is unclear how to construct a "probabilistic" adjacency matrix that could capture this flux and still have eigenvalues with useful properties.

Equation 8 gives the probability that two random waypoint nodes are within communication range at an arbitrary moment. Using this probability, we could treat the mobile topology as a random graph [11] and try to use component analysis [6] to reason about its topological properties. This approach fails for two reasons. First, results from random graph theory are only useful for dense networks, and mobile topologies often lack the requisite number of nodes [19]. Second, epidemics depend not just on the variety of cluster types, but on their membership churn; once again, the temporal dimension is not captured by random graph theories.

Spatially coupled epidemiological models are the closest mathematical analogues to probabilistic queuing. For example, Keeling and Rohani consider disease spreading amongst two distinct populations that can exchange members [12]. Each population has a separate differential equation representing its infection rate, but each equation has a coupling term which describes the infection spillover due to cross-population mixing. The Keeling-Rohani model makes several key assumptions which are invalid in the mobile setting, e.g., it assumes perfectly homogeneous mixing, and a distinction is made between a node's "home" and "foreign" domain. Nevertheless, the notion of coupled populations resembles in spirit our notion of coupled queues. The major difference between our model and their model is the representation of time. Keeling and Rohani's coupling constant is a dimensionless ratio relating the time a node spends in a foreign domain to the time spent in its home domain. In a mobile environment, the relative time spent at each connectivity level is important, but the *raw* amount of time is also important. This is the lesson of Figure 3 — two networks with the same connectivity distribution but different node speeds will have different epidemics.

Using a modified KW model, Khelil *et al* investigated flooding-based information dissemination in mobile networks [14]. They assumed a perfectly homogeneous mixing rate based on the ratio of the number of nodes to the size of the arena. As we have shown, assumptions of homogeneity are typically unwarranted and often lead to severe mispredictions in epidemic simulation.

6. CONCLUSION

Traditional epidemiological models fail to capture the unique topological properties of mobile networks. Node mobility introduces non-homogeneous connectivity distributions that cannot be represented using a simple average. Mobility also creates continual churn in each node's neighbor set. Our new epidemiological framework uses a queue abstraction to model these important phenomenon. By representing different connectivity levels as distinct queues, we capture the skewed connectivity distributions inherent to mobile networks. As nodes travel between different queues, they simulate the neighbor churn experienced by real nodes. By tying the travel time through a queue to both the connectivity level it represents and to node velocity, we capture a temporal factor that other viral models cannot. Simulations show that for realistic mobility parameters, probabilistic queuing offers more accurate predictions than the Kephart-White model. There are several important areas for future work, such as finding analytical derivations for the mixing rate and the connectivity fluctuation parameter. However, we believe that probabilistic queuing already offers useful and interesting insights into mobile epidemiology, a topic whose importance will grow as mobile networks become more popular and thus more alluring to attackers.

7. REFERENCES

[1] C. Bettstetter. On the Connectivity of Ad-hoc Networks. *The Computer Journal, Special Issue on Mobile and Pervasive Computing*, 47(4):432–437, July 2004.

[2] C. Bettstetter, H. Hartenstein, and X. Perez-Costa. Stochastic Properties of the Random Waypoint Mobility Model. *ACM/Kluwer Wireless Networks*, 10(5):555–567, September 2004.

[3] L. Billings, W. Spears, and I. Schwartz. A Unified Prediction of Computer Virus Spread in Connected Networks. *Physics Review Letters*, 297(3):261–266, May 2002.

[4] Larry Bridwell. *Ninth Annual Computer Virus Prevalence Survey*, 2004. Published by ICSA Labs.

[5] J. Broch, D. Maltz, D. Johnson, Y. Hu, and J. Jetcheva. A performance comparison of multi-hop wireless ad hoc network routing protocols. In *Proceedings of MobiCom*, pages 85–97, Dallas, TX, October 1998.

[6] D. Callaway, M. Newman, S. Strogatz, and D. Watts. Network Robustness and Fragility: Percolation on Random Graphs. *Physics Review Letters*, 85(25):5468–5471, December 2000.

[7] E. Chien. *Security Response: SymbOS.Mabir*, 2005. Symantec Corporation.

[8] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On Power-Law Relationships of the Internet Topology. In *Proceedings of SIGCOMM*, pages 251–262, Cambridge, MA, September 1999.

[9] P. Ferrie, P. Szor, R. Stanev, and R. Mouritzen. *Security Response: SymbOS.Cabir*, 2004. Symantec Corporation.

[10] International Business Machines Corporation. *Global Business Security Index Report*, 2004.

[11] S. Janson, T. Luczak, and A. Rucinski. *Random Graphs*. Wiley-Interscience, New York, NY, 2000.

[12] M. Keeling and P. Rohani. Estimating Spatial Coupling in Epidemiological Systems: a Mechanistic Approach. *Ecology Letters*, 5(1):20–29, 2002.

[13] J. Kephart and S. White. Directed-graph epidemiological models of computer viruses. In *Proceedings of the IEEE Computer Symposium on Research in Security and Privacy*, pages 343–359, May 1991.

[14] A. Khelil, C. Becker, J. Tian, and K. Rothermel. An Epidemic Model for Information Diffusion in MANETs. In *Proceedings of the 5th ACM International Workshop on Modeling, Analysis, and Simulation of Wireless and Mobile Systems*, pages 54–60, Atlanta, GA, September 2002.

[15] M. Lactaotao. *Security Information: Virus Encyclopedia: SYMBOS_COMWAR.A: Technical Details*, 2005. Trend Micro Incorporated.

[16] W. Navidi, T. Camp, and N. Bauer. Improving the Accuracy of Random Waypoint Simulations Through Steady-State Initialization. In *Proceedings of the 15th International Conference on Modeling and Simulation*, pages 319–326, March 2004.

[17] R. Pastor-Satorras and A. Vespignani. Epidemic Spreading in Scale-Free Networks. *Physics Review Letters*, 86(14):3200–3203, April 2001.

[18] R. Pastor-Satorras and A. Vespignani. Epidemic Dynamics in Finite Size Scale-Free Networks. *Physics Review Letters*, 65:035108, March 2002.

[19] P. Santi and D. Blough. The Critical Transmitting Range for Connectivity in Sparse Wireless Ad Hoc Networks. *IEEE Transactions on Mobile Computing*, 1(2):25–39, January-March 2003.

[20] Y. Wang, D. Chakrabarti, C. Wang, and C. Faloutsos. Epidemic Spreading in Real Networks: An Eigenvalue Viewpoint. In *Proceedings of the Symposium on Reliable Distributed Computing*, pages 25–34, Florence, Italy, October 2003.

[21] R. Wong and I. Yap. *Security Information: Virus Encyclopedia: WINCE_BRADOR.A: Technical Details*, 2004. Trend Micro Incorporated.