

On Using Mobility to Propagate Malware

Sandeep Sarat Andreas Terzis
Computer Science Department
Johns Hopkins University
{sarat,terzis}@cs.jhu.edu

Abstract—Mobility can be exploited to spread malware among wireless nodes. In this paper, we present an analytical model for estimating the evolution of infections spanning multiple network domains that host mobile nodes. We validate the accuracy of the proposed model by comparing its predictions to simulations driven by realistic mobility patterns. Our results show that such a mobile infection requires less than a day to infect the majority of a mobile population with thousands of wireless nodes spanning hundreds of network domains. Moreover, if mobile nodes are allowed to infect nodes within the same domain that are connected to the wired network, then an even smaller number of mobile nodes can inflict comparable damage in similar time frames. Unfortunately, these infections generate negligible activity at global malware monitoring stations (*e.g.*, network telescopes and honeypots), which contributes to their stealthiness. By observing the infection’s spatial evolution we show that popular domains are infected during the early stages of the infection. This observation is likely to be useful in designing countermeasures against mobile infections. By placing monitors in approximately 10% of the most visited domains, we can detect the mobile worm before it reaches a majority of the population. Finally, we elucidate why simply placing telescopes in just the popular domains is not sufficient for early detection.

I. INTRODUCTION

Mobility pervades networked devices today. For example, millions of users access the Internet through laptops and PDAs equipped with WiFi cards connected to thousands of Access Points (APs) located on campuses, coffee shops, airports, etc. This increase in connectivity however comes at a high price – failure to secure these communication channels provides a new propagation vector for spreading malware (*i.e.* self-replicating malicious code such as worms and viruses). As a matter of fact, the exploitation of these channels is not just our speculation: variants of the Zotob/Mytob worm are suspected to have used physical movement of computers across network domain boundaries as a propagation strategy [23]. More recently, a series of malware that attempt to exploit Bluetooth connections as an infection mechanism were reported in the media [2]. The accepted practice to protect from such worms today is to place mobile nodes in a de-militarized zone (DMZ), separate from the rest of the network. In such a scenario, all communication between the mobile nodes and the wired nodes passes through a firewall. However, mobile nodes can still infect each other through contacts within these de-militarized zones.

Unfortunately, modelling efforts have not followed the pace of malware evolution as most previous work describes how infections spread over wired networks. To better understand this impending threat, we develop a concise analytical model that predicts the speed of infections over populations of

nomadic users traversing a collection of network access points. The accuracy of the model is validated through simulations driven by realistic mobility models, drawn from university-wide traces at Dartmouth College [8]. We found that in networks with thousands of users and hundreds of APs the infection can reach 65% of the total population within only one day, a relatively short time considering that infections follow the slow pace of node movements across network domains. Furthermore, if mobile nodes are allowed to infect co-located nodes connected to the wired network, a scenario modelling imperfect DMZs, we observed that even a small proportion of vulnerable mobile nodes can propagate the infection to the majority of the network domains within a single day.

Due to the high propagation speed of these worms, human defense mechanisms are rendered implausible. Moreover, the threat from this class of infections stems from the fact that mobile nodes trivially bypass existing perimeter defenses, such as firewalls. Since cross-domain transfer of the infection is accomplished by the physical migration of infected nodes, it is difficult to contain them, when no controls exist to police the movement of nodes across domains. Such gaps in network defenses can lead to global worm outbreaks. Finally, the detection of these worms is challenging due to their stealthiness. This characteristic is a consequence of the fact that the majority of current detection techniques relies on traffic anomalies measured at network monitors (*network telescopes* [14]). Unfortunately since mobile infections scan within the domains of infected nodes, suspicious probes on telescopes deployed at remote domains would be absent. This observation motivates the need for developing novel malware containment technologies. One promising direction towards this goal involves exploiting the spatial characteristics of the infection. Specifically, we observed that by placing monitors in approximately 10% of the most visited domains, we can detect the mobile worm before it reaches the majority of the population. While this seems a straightforward solution to the early detection problem, we argue that monitor placement is still a challenging problem with many intricacies.

The structure of rest of the paper is as follows: In the next section we present previous models for malware and mobility patterns. Section III introduces the model for predicting the spread of infections among populations of mobile users. We compare the model’s predictions to simulation results driven by realistic mobility traces in Section IV where we also investigate a number of variants of this worm. In Section V we compare the mobile worm to a ‘traditional’ (*i.e.* globally

scanning) worm and provide intuition about the temporal evolution of the infection by connecting it to the structure of the mobility graph in Section VI. Finally, we discuss the issues involved with telescope placement in Section VII and close in Section VIII with future research directions.

II. RELATED WORK

A large volume of research has focused on modelling Internet worms. Among these, the classic homogeneous worm model assumed all-to-all node connectivity and that every susceptible node was a target of equal probability [10]. More recent models accounted for non-uniform scanning strategies [7], as well as for the fact that node population is not uniformly distributed over the IP address space [15]. However, much of the prior work ([6], [18], [20] among many others), primarily considers how malware propagates in wired networks. Instead, we explore how mobility can facilitate the spread of infections among groups of nomadic users traversing different network attachment points such as WiFi Access Points. In this case unlike previous scenarios, each infected node has a time-varying infection transmission probability depending on its local scope.

In the context of mobile networks, Anderson *et al.* derived the speed of mobile worms through simulations [3]. While our results seem to be in broad agreement, we focus our attention on the actual infection evolution, so as to infer the worm characteristics. Similar trace-driven studies covering infections over Bluetooth networks were performed by Su *et al.* [19]. Unlike those previous studies, which are limited to simulations performed using a particular trace, we propose a general analytical model that predicts the evolution of infections over a wide range of mobility patterns. Epidemic spreading in ad-hoc networks has been studied by Mickens and Noble in [13]. That work explained why traditional epidemic models fail in the case of mobile networks and proposed a new framework for such networks. While that study focused on worms spreading within a single ad-hoc wireless network, our model explains how infections are carried across a variety of networks by the physical movement of mobile users.

The mobility model we use is similar to the semi-Markov model presented in [12]. Lee *et al.* developed a cumulative model for different user groups to obtain the AP-user mobility patterns. Instead, we model the mobility patterns of individual users. We choose to do so, because the derived mobility model is then used to calculate the contact rates between mobile node pairs. As we will later show, this is the key factor that determines the rate at which the infection travels among individual nodes.

Today, it is generally considered good practice to place mobile nodes in a DMZ separated from wired nodes. Various enterprise solutions exist for doing so, *e.g.* Cisco's network admission control [1]. We believe that these perimeter defenses by themselves are insufficient and a more fine-grained approach is needed to detect and contain mobile worms. We present an outline of such defenses in Sections VI and VII.

III. WORM MODEL

We model infections spreading over collections of mobile users who connect to the Internet through a revolving set of network access points. This model consists of two types of entities: **(a)** network domains through which users connect to the Internet and **(b)** mobile nodes, *e.g.* laptops and PDAs, that are susceptible to infections and move across these domains. In this context, domains act as mixing regions in which mobile nodes can reach each other. We assume that an infected mobile node can infect another susceptible mobile node if they reside in the same domain, even for a short period of time. This is a realistic assumption because an infected mobile node can eavesdrop on communications from all the other wireless nodes in the same domain and attempt to infect them directly.

The evolution of an infection can be modelled as a discrete time, replication process over the set \mathcal{V} of vulnerable nodes. We denote the probability that node i is infected at time step t by $p_{i,t}$. Furthermore, let β_{ij} be the probability that node i contacts node j . Given these conditions, node i is not infected at time step t iff it was not infected by time step $t-1$ and no infected nodes in the domain it resides contacted node i during the last time step. Because these events are independent, this probability can be expressed as:

$$1 - p_{i,t} = (1 - p_{i,t-1}) \prod_{j \neq i} (1 - \beta_{ji} p_{j,t-1})$$

$$1 - p_{i,t} = 1 - p_{i,t-1} - \sum_j \beta_{ji} p_{j,t-1}$$

Here, we use the approximation $(1-a)(1-b) \approx 1-a-b$ when $a \ll 1$, $b \ll 1$. Thus we have,

$$p_{i,t} \approx p_{i,t-1} + \sum_j \beta_{ji} p_{j,t-1} \quad (1)$$

By representing $(p_{1,t}, p_{2,t}, \dots)$ as a row vector P_t and assigning $\beta_{ii} = 1$ (*i.e.*, the probability that a node i contacts itself is trivially one), we can rewrite Equation (1) in a matrix form as:

$$P_t = P_{t-1} M \quad (2)$$

where $M = [\beta_{ij}]$ is the system matrix, containing the pairwise contact probabilities. From the definition of P_t , $p_{i,t}$ is the probability that node i is infected at time t . Therefore, the expected number of infected nodes after time t is given by

$$E[|I|] = \sum_{i=1}^{|\mathcal{V}|} p_{i,t} = \|P_t\|_1 \quad (3)$$

where I is the set of all infected nodes. This type of matrix multiplication view of an infection is common in epidemic modelling (*e.g.* [20]).

We initiate the infection by infecting a single node, say k . The initial conditions are then as follows:

$$p_{i,0} = \begin{cases} 1 & \text{if } i = k, \\ 0 & \text{Otherwise} \end{cases}$$

If multiple nodes are initially infected (also known as patient zeros), the corresponding indices in P_0 are set to unity.

A. Mobility Model

It is evident that in order to estimate the expected number of infected nodes in Equation (3) we need to calculate the contact probabilities β_{ij} . In turn, these probabilities depend on the number of domains a node visits and the duration of time that the node resides in each domain. We therefore need a mobility model that describes the movement of mobile nodes across network domains.

We model the mobility pattern of individual nodes using semi-Markov chains. We chose the more general semi-Markov model because it was shown that node residence times do not follow the exponential distribution [5], [11], but are better modelled by heavy-tailed distributions. The state space $S = \{1, \dots, m\}$ of the homogeneous semi-Markov chain is the set of all network domains. The transition matrix P describing the chain is then an $m \times m$ matrix, while $\bar{D} = [\bar{d}_i]$ is an $m \times 1$ vector, which gives the mean residence time of the node in each domain.

We can then derive the steady-state transition probability distribution $\tilde{\pi}$ by solving the following set of equations:

$$\begin{aligned} \tilde{\pi} &= \tilde{\pi}P \\ \sum_{i=1}^m \tilde{\pi}_i &= 1 \end{aligned}$$

Given the fraction of time $\tilde{\pi}$ that the user stays in each state and the mean residence times \bar{D} for each state, it is easy to calculate the steady-state probability π_i of the user staying in domain i :

$$\pi_i = \frac{\bar{d}_i \tilde{\pi}_i}{\sum_{j=1}^m \bar{d}_j \tilde{\pi}_j} \quad (4)$$

From Equation (4) we can subsequently compute the contact rate β_{xy} between nodes x and y . This value is equal to the probability that both x and y are in the same domain at some point in time. Without loss of generality, we say that when a node is in the ‘‘OFF’’ state (*i.e.* it is not operational) then it resides in the domain with index 1. Since, the infection does not propagate when nodes are not connected, we do not include the percentage of time in the ‘‘OFF’’ state in the calculation of the contact rates. The contact rates are then given by:

$$\beta_{xy} = \sum_{i=2}^m \pi_i^x \pi_i^y \quad (5)$$

where π_i^x is the percentage of the time spent by x in domain i . We substitute Equation (5) into Equation (2) to obtain the number of infectees as a function of time.

The last complication is that Equation (2) proceeds on discrete time steps of uniform duration, while nodes actually have variable domain staying times. We address this discrepancy by using the mean residence time across all domains as the discrete time step in Equation (2). While doing so compromises the accuracy of the analytical model, as the simulation results from Section IV demonstrate, even with this compromise the model is able to accurately track the infection’s evolution.

IV. EVALUATION

We derive the parameters of the mobility model described in the previous section from traces of actual mobile user behaviors, available from Dartmouth college [8]. Each trace is a time sequence of the access points the mobile users visit (identified by their MAC addresses). Traces also contain the special ‘‘OFF’’ location, signifying a user’s departure from the network. The trace we use contains 626 different access points and tracks the movement of mobile users from 9/23/2003 to 12/10/2003. Approximately 6% of the users in our trace visited just a single domain before entering the ‘‘OFF’’ state. We removed such users, since states in their semi-Markov chain are not recurrent and their steady state probabilities in states other than the ‘‘OFF’’ state are trivially zero. In all, we had 6101 users. We assume that all the mobile users in the system are vulnerable. We observed similar infection curves when only a fraction of the mobile users were vulnerable. Furthermore, the infection model can easily incorporate scenarios in which only a subset of the mobile nodes are vulnerable by appropriately defining the set of vulnerable nodes, \mathcal{V} . The mean domain residence time of the users is approximately 67 minutes. We use this value as the discrete time step in Equation (2).

A. Mobile node infection

We compare the model’s predictions with results provided by detailed simulations. The custom simulator we developed emulates the movements of mobile users over the same collection of APs and tracks the evolution of the infection after an initial node (Patient zero) is infected. As before, we assume that the infection passes from an infected node to any other node that resides in the same network at the same time. We ran 100 simulations, each time randomly choosing a different initial node to infect.

Figure 1 graphs the evolution of the infection as a function of time. In addition to the infection curve predicted by the analytical model, we present three representative simulation runs. These curves represent the 5th, 50th, and 95th percentiles across all simulations, where rank is calculated based on the time when 70% of vulnerable hosts are infected. Intuitively, these curves represent a slow, average, and fast infection instance depending on which node was infected first.

First, we note that the model provides a decent approximation of the average infection evolution, faithfully tracking the curves that represent the simulations. Furthermore, the infection spreads to approximately 60% of the users within a single day. Given that the worm requires under a day to infect




Fig. 1. Percentage of infected users as a function of time as predicted by the analytical model and as demonstrated by simulation.

the majority of the population, we experimented by starting the infection at different days during the period covered by the network trace. In all cases we observed patterns very similar to those in Figure 1. We also found that the evolution speed varied depending on the time of day when the first node was infected. Worms that started during the daytime spread faster than those started at night. This is due to the decreased movement of nodes during the night hours.

B. Mixing mobile and static nodes

So far we have assumed that mobile users cannot infect nodes connected to the static (wired) network, This model corresponds to current security practices according to which WiFi APs are separated from the rest of the network (*e.g.* a company’s intranet) by firewalls. However, firewalls are complex devices that are notoriously difficult to configure. Therefore, it is possible that a misconfigured firewall would allow infected wireless devices to contact hosts residing in the static part of the network. More commonly, laptops can connect directly to the static portion of the network after they have roamed across several wireless domains (*e.g.* during a business trip) effectively bypassing the barrier between the static and mobile compartments of a network domain.

In this scenario, static hosts can be infected by mobile nodes and subsequently carry the infection to other vulnerable nodes. Therefore, it is no longer necessary for mobile nodes to simultaneously reside in the same domain for the infection to spread; a mobile node entering a network domain can contract the infection by infected static nodes in that domain. In order to understand how these infections spread, we modified the original simulator to assume the worst case scenario, wherein an infected mobile node instantly infects any domain that it enters. The “instant infection” assumption is valid even for a uniform scanning worm (*i.e.* which follows a naive strategy of random scanning and therefore one of the slower spreading

worms). Even with a scan rate of 10 scans/sec and domains with as few as 10% vulnerable nodes, one static node on the average is infected within the first second from the entry of an infected user to the domain.





Fig. 2. (a)Rate of domain infections as a function of time with the total mobile population (b) Rate of infection with only 25% of the mobile nodes

Figure 2.(a) presents the number of *network domains* infected as a function of time when mobile nodes can infect the domains they visit. The infection spreads to about 65% of the domains within a day. It then slows down considerably and takes a long time to infect the remaining domains. This result might seem straightforward, given that 65% of mobile nodes contract the infection within one day. In order to investigate the relationship between the number of mobile nodes carrying the infection and its spread over the set of network domains, we repeated the previous experiment, with a randomly selected subset of 1500 wireless nodes (25% of the original population). The surprising result, as Figure 2.(b) indicates, is that infection rates in this case are comparable to the previous case, *i.e.* the infection reaches $\sim 60\%$ of the domains within a day. This result indicates that the worm speed is not significantly hampered by the significantly smaller set of

cross-domain carriers. This phenomenon can be explained by the association graph usually observed in social networks [9]. In that context, as well as in the context of network domains visited by mobile hosts, domain popularity has been shown to follow a heavy tailed distribution, whereas a small number of domains are extremely popular followed by a large number of less popular domains. As a result, the smaller subset of nodes is still likely to frequent at the very popular domains thus fuelling the growth of the infection.

V. DETECTION

Thus far we have shown that a mobile infection can take up to a day to affect a significant portion of the vulnerable population. Although this is fast enough to make human defense mechanisms implausible, it is considerably slower compared even to the naïve uniform scanning strategy, or more sophisticated variants such as flash worms that can spread over the entire Internet in a few minutes [17].

The fact that such worms spread more slowly might lead to the conclusion that they are *easier* to contain. This, however, is false. On the contrary, mobile infections are more difficult to detect using conventional approaches, such as distributed network monitors [4], [15]. In the paragraphs that follow, we explain the underlying reason for this negative result.

A. Detection Speed

We compare the expected time to detect a mobile infection to the average detection time of a uniform scanning worm. Here we assume that a single network telescopes is used to detect the infection. We define detection time as the time elapsed from the first infection until the first probe arrives to the address space monitored by the telescope(s). Suppose, that the telescope covers a large fraction, $\alpha = 0.5$, of the IP space used in the network domain where it is deployed. Then, the expected time T to detect the first instance of the infection for a uniform scanning worm is given by:

$$\begin{aligned} H(T) &= \int_0^T I(t) \cdot s \, dt \\ &\approx \int_0^T e^{sft} \cdot s \, dt \\ &= \frac{N}{\alpha} \\ \Rightarrow T &= \frac{1}{s \cdot f} \ln\left(\frac{N \cdot f}{\alpha} + 1\right) \end{aligned} \quad (6)$$

where $H(t)$ is the number of IP addresses scanned by all the infected nodes in $[0, t]$, s is the scan rate, N is the total number of domains, and f is the average density of vulnerable nodes.

Substituting conservative values for $s = 20$ scans/min (the Witty worm had a scan rate of roughly 1200 scans/min [22]), $N = 1000$, and $f = 0.01$ in Equation (6) we find that a uniform worm will be detected within 15 minutes on the average. By this time the worm has spread to less than 2% of the vulnerable population (calculated from the equation for the uniform scanning worm). Furthermore, the placement of the telescope is immaterial to the detection time. Thus, we conclude that such a telescope can be an effective early warning device for typical worms.

On the other hand, since mobile worms scan only their local network, detection time is governed by the speed with which infected mobile nodes enter the domain where the telescope is located. Considering the same (randomly placed) single telescope, detection will occur when the worm has spread to half of the domains on the average. Figure 2 provides the time for the worm to spread to 50% of the vulnerable domains as ~ 15 hours. Within this time, the worm infection has already taken off, infecting a large number of hosts. Once the worm enters the domain which contains the network monitor, detection is much faster. On the other hand, since detection time is dominated by the time necessary for the worm to enter the domain, using larger telescopes within a domain does not significantly reduce detection speed.

In short, unlike traditional uniform-scanning worms, telescope size is not important and random placement is of little use. On the other hand, given that the worm first infects popular domains first, it is prudent to place worm monitors in those domains.

VI. SPATIAL EVOLUTION

Until now we have investigated the temporal behavior of the infection. However, an equally interesting aspect is the infection’s spatial evolution, that is how the infection spreads over the collection of network domains the mobile nodes visit. We note that Figures 1 and 2 flatten out considerably after an almost vertical growth during the middle phase of the evolution graph. This behavioral change can be explained by dividing the spatial evolution of the infection into a number of distinct phases. The infection initially “moves” in the direction of domains which are extremely popular, since many nodes visit them. This is the slow take-off phase. These popular domains (we call them *hubs*) are closely connected by the group of mobile nodes which frequent them, thus forming a *dense core* of the network graph. When the infection reaches this core, an exponential increase in the number of infected hosts occurs, as the majority of vulnerable nodes frequently visit the core. Finally, the infection gradually slows down after it has consumed the core and extends towards domains with low contact rates (*i.e.* unpopular domains). Figure 3 illustrates this phenomenon where it is clear that popular domains are infected within the first few hours of the infection.

A. Popularity

We define the popularity of a domain as the cumulative number of node-hours that nodes spend in that domain. This definition accounts for both the distinct number of nodes visiting the domain as well as the length of time a node resides in the domain.

Intuitively, placing network monitors in the most popular domains yields the earliest detection times. To quantitatively measure the effect of placing multiple monitors, we placed monitors in the top $x\%$ of the domains and measured the detection times. As Figure 4 shows installing monitors in 10% of the domains reduced the detection time to about 10 hours. During this time the worm has spread to less than




Fig. 3. The first time an infected node is seen at a network domain as a function of the domain's popularity, defined as the number of cumulative node-hours occupancy of a domain.




Fig. 4. Detection time when monitors are deployed in the top x% of the domains.

10% of the hosts (as seen from Figure 1). Installing additional monitors provides only marginal benefits, reducing in the limit the detection time to a little over 9 hours.

VII. DISCUSSION

Deploying wireless network monitors may involve modifying APs to scan through packets they forward looking for traces of malware or deploying honeypots acting as decoys. As we showed, placing such monitors in the top 10% of the domains can help detect the worm early enough. However, this strategy in itself is not sufficient to guarantee early detection. We present two arguments to support this claim.

A. Popularity is dynamic

First, we investigate how domain popularities change over time and the effect these changes have on detection time. For this purpose we use the access points from the previous dataset [8] to calculate the popularity of each domain on a weekly basis. We then choose an initial set of the 50 most popular APs ($\sim 10\%$ of the total AP population) during the first week of the network trace and measure how this set compares with the set of top 50 APs for every other week. The similarity between the first and every other weekly set is estimated by calculating the dot product between the two sets and dividing the result by 50. In this case a product of one indicates that the sets are identical, while zero indicates that no common members exist between the two sets.

Figure 5.(a) plots how the similarity between the top 50 APs evolved during year 2004. It is evident that there are wide variations with two prominent dips around weeks 30 and 50. Closer inspection of the CRAWDAD dataset revealed that during the Fall and Spring sessions, the APs in the residential buildings were the most popular. On the other hand, APs in the academic buildings and athletic centers were highly ranked during inter sessions, explaining the aforementioned changes. Figure 5.(b) shows the corresponding median worm detection time over time, when monitors are statically placed in the top 50 domains according to the popularity results of the first week. While it may seem that the difference in the detection time is only a matter of two hours, varying between 10.5 and 12.5 hours, the effects of this difference are dramatic. As Figure 1 indicates, this disparity results in an infection spread of $<5\%$ in the case of 10.5 hours, as opposed to $\sim 30\%$ when the detection time is 12.5 hours. Thus, reducing the detection time window is crucial to providing sufficient time if the worm defenses are to be effective.

B. Evasive worms

The second reason why static placement of monitors is insufficient, is that worms can potentially detect their presence and avoid the networks in which these monitors are deployed. Rajab *et al.* have presented an efficient *probe-response attack* that can be used to discover the locations of network monitors deployed on the (wired) Internet [16]. A similar technique could potentially be applied in the context of mobile infections. In this case, worm instances probe the domain they currently reside, using standard network tools such as ping and ARPs, or even passively eavesdrop all ongoing communications to the AP. If a domain is believed to host a monitor the worm will not attempt to infect any mobile nodes in that domain, thus avoiding detection.

On the other hand, if avoiding popular domains, in which monitors are deployed, slows down the infection to the point where human intervention is practical, then the threat posed from these *evasive* infections is minimal. To verify whether this is true, we simulated such an evasive worm that does not try to infect the 50 most popular domains, and measured its infection speed. Unfortunately, as Figure 6 indicates, the

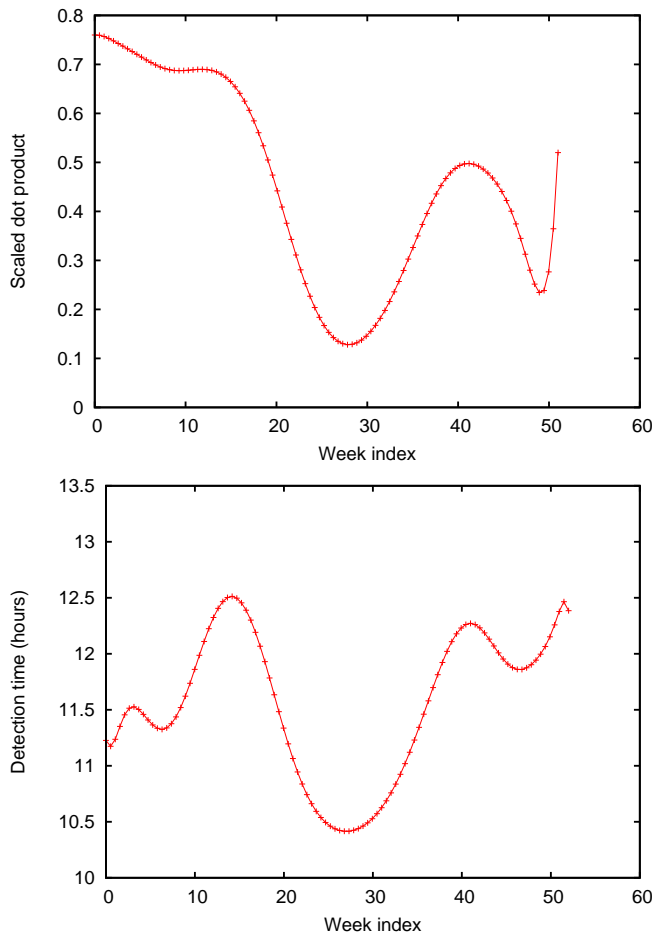


Fig. 5. (a) Similarity between the popularity of the top 50 domains on a weekly basis for 2004 (b) Median detection time if the monitors are deployed statically.

infection rate is still significant, with 60% of the hosts infected within 3 days on the average.

From the two arguments presented above it is clear that placing monitors in the most popular domains is not a complete solution to the problem of early detection. The actual strategy for their placement is part of our future work.

VIII. SUMMARY AND FUTURE DIRECTIONS

We presented and validated an analytical model that describes the evolution of worms that exploit node mobility to propagate. We evaluated infection speeds in different scenarios: first, when mobile users can only infect each other as they move across a collection of network domains and second when infections can spread from mobile users to static nodes. Our ultimate goal is to use this model to design effective detection and containment mechanisms for this novel category of worms. While we touched upon the detection mechanisms for this type of infections, an important topic for future work is the in-depth study of the worm mitigation mechanisms.

Even with effective detection mechanisms, the feasibility of policing nodes as they enter popular domains is not




Fig. 6. Worm evolution when the worm is inactive in the top 50 domains.

straightforward. Numerous practical concerns for containment mechanisms designed for mobile infections must be addressed, including how to exploit topological information to limit the damage from potentially infected nodes, how to appropriately apply the notion of hard-LANs [21] in this setting, and how to track (in a tamper-resistant manner) the movement of nodes across network domains.

ACKNOWLEDGMENTS

This work is supported in part by National Science Foundation grant CNS-0627611. We gratefully acknowledge the use of trace data from the CRAWDAD archive at Dartmouth College.

REFERENCES

- [1] Cisco network admission control. CISCONAC: Available at http://www.cisco.com/en/US/netsol/ns466/networking_solutions_package.html.
- [2] Macosx malware latches onto bluetooth vulnerability. Available at http://www.theregister.co.uk/2006/02/17/macosex_bluetooth_worm, 2006.
- [3] E. Anderson, K. Eustice, S. Markstrum, M. Hansen, and P. Reiher. Mobile contagion: Simulation of infection and defense. In *PADS '05: Proceedings of the 19th Workshop on Principles of Advanced and Distributed Simulation*, pages 80–87, Washington, DC, USA, 2005. IEEE Computer Society.
- [4] M. Bailey, E. Cooke, F. Jahanian, J. Nazario, and D. Watson. Internet motion sensor: A distributed blackhole monitoring system. In *Proceedings of the ISOC Network and Distributed System Security Symposium (NDSS)*, 2005.
- [5] M. Balazinska and P. Castro. Characterizing Mobility and Network Usage in a Corporate Wireless Local-Area Network. In *1st International Conference on Mobile Systems, Applications, and Services (MobiSys)*, San Francisco, CA, May 2003.
- [6] G. S. Canright and K. Engo-Monsen. Epidemic Spreading over Networks - A View from Neighbourhoods. *Teletronikk*, 2005(1), 2005. Available at: http://www.telenor.com/teletronikk/volumes/pdf/1.2005/Page_065-085.pdf.
- [7] Z. Chen, L. Gao, and K. Kwiat. Modeling the Spread of Active Worms. In *Proceedings of IEEE INFOCOMM*, volume 3, pages 1890 – 1900, 2003.

- [8] Crawdad: A community resource for archiving wireless data at dartmouth. Available at: <http://www.crawdad.cs.dartmouth.edu/data.php>.
- [9] S. Eubank, V. S. A. Kumar, M. V. Marathe, A. Srinivasan, and N. Wang. Structural and algorithmic aspects of massive social networks. In *Proceedings of the fifteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 718–727, 2004.
- [10] H. Hethcote. The Mathematics of Infectious Diseases. In *SIAM Reviews*, Vol. 42 No. 4, 2000.
- [11] R. Jain, A. Shivaprasad, D. Lelescu, and X. He. Towards a model of user mobility and registration patterns. *SIGMOBILE Mob. Comput. Commun. Rev.*, 8(4):59–62, 2004.
- [12] J.-K. Lee and J. C. Hou. Modeling steady-state and transient behaviors of user mobility:: formulation, analysis, and application. In *MobiHoc '06: Proceedings of the seventh ACM international symposium on Mobile ad hoc networking and computing*, pages 85–96, New York, NY, USA, 2006. ACM Press.
- [13] J. W. Mickens and B. D. Noble. Modeling epidemic spreading in mobile environments. In *WiSe '05: Proceedings of the 4th ACM workshop on Wireless security*, pages 77–86, New York, NY, USA, 2005. ACM Press.
- [14] D. Moore. Network Telescopes: Observing Small or Distant Security Events. In *11th USENIX Security Symposium, Invited Talk*, Aug. 2002.
- [15] M. A. Rajab, F. Monrose, and A. Terzis. On the effectiveness of Distributed Worm Monitoring. In *Proceedings of Usenix Security*, 2005.
- [16] M. A. Rajab, F. Monrose, and A. Terzis. Fast and Evasive Attacks: Highlighting the challenges ahead. In *Proceedings of the 9th International Symposium on Recent Advances in Intrusion Detection (RAID)*, Sept. 2006.
- [17] S. Staniford, D. Moore, V. Paxson, and N. Weaver. The Top Speed of Flash Worms. In *Proceedings of the ACM Workshop on Rapid Malcode (WORM)*, pages 33–42, Oct. 2004.
- [18] S. Staniford, V. Paxson, and N. Weaver. How to Own the internet in your spare time. In *Proceedings of the 11th USENIX Security Symposium*, Aug. 2002.
- [19] J. Su, K. W. Chan, A. G. Miklas, K. Po, A. Akhavan, S. Saroiu, E. de Lara, and A. Goel. A Preliminary Investigation of Worm Infections in a Bluetooth Environment. In *4th Workshop on Rapid Malcode*, 2006.
- [20] Y. Wang, D. Chakrabarti, C. Wang, and C. Faloutsos. Epidemic spreading in real networks: An eigenvalue viewpoint. In *22nd Symposium on Reliable Distributed Computing, Florence, Italy, Oct. 6-8, 2003.*, 2003.
- [21] C. Weaver, D. Ellis, S. Staniford, and V. Paxson. Worms vs Perimeters: The Case for Hard-LANs. In *Proceedings of the 12th Annual IEEE Symposium on High Performance Interconnects*, 2004.
- [22] The CAIDA Dataset on the Witty Worm - March 19-24, 2004, Colleen Shannon and David Moore, <http://www.caida.org/passive/witty/>. Support for the Witty Worm dataset and the UCSD Network Telescope are provided by Cisco Systems, Limelight Networks, DHS, NSF, CAIDA, DARPA, Digital Envoy, and CAIDA Members.
- [23] Zotob causes carnage in corporate networks. Available at: http://www.netfastusa.com/xq/asp/id.1338/p.5-6-1/qx/PressRelease_view.htm.