Quantifying the Effectiveness of Mobile Phone Virus Response Mechanisms

Elizabeth Van Ruitenbeek,
Tod Courtney, and William H. Sanders
Coordinated Science Laboratory
University of Illinois at Urbana-Champaign
Urbana, IL, USA
evanrui2, tcourtne, whs@uiuc.edu

Fabrice Stevens
France Telecom Research and Development
Network and Services Security
92794 Issy les Moulineaux Cedex 9, France
fabrice1.stevens@orange-ftgroup.com

Abstract

Viruses that infect smartphones are emerging as a new front in the fight against computer viruses. In this paper, we model the propagation of mobile phone viruses in order to study their impact on the dependability of mobile phones. We propose response mechanisms and use the models to obtain insight on the effectiveness of these virus mitigation techniques. In particular, we consider the effects of multimedia messaging system (MMS) viruses that spread by sending infected messages to other phones. The virus model is implemented using the Möbius software tool and is highly parameterized, enabling representation of a wide range of potential MMS virus behavior. Using the model, we present the results of four illustrative MMS virus scenarios simulated with and without response mechanisms. By measuring the propagation rate and the extent of virus penetration in the simulation phone population, we quantitatively compare the effectiveness of mobile phone virus response mechanisms.

1. Introduction

The enhanced computational and communication capabilities of smartphones are beginning to attract viruses targeted at these increasingly sophisticated mobile phones [12]. The problem is expected to worsen as smartphones become more prevalent and as virus writers become more proficient in working with mobile phones [11].

Attacks from mobile phone viruses can compromise personal information, delete data, drain the battery [9], and steal phone services by using expensive features [3]. The impact of mobile phone viruses on phone service providers includes increased customer complaints concerning infected phones and extra network congestion due to the virus-related traffic [12]. It is imperative that the mobile phone industry anticipate and act now against these looming threats to dependable and secure mobile phone services.

Because mobile phones are communications devices with many connectivity options, there exist many possible infection vectors [11]. Mobile phones can become infected by downloading infected files using the phone Internet browser, by transferring files between phones using the Bluetooth interface, by synchronizing with an infected computer, by accessing an infected physical memory card, or by opening infected files attached to multimedia message service (MMS) messages. MMS messages are similar to text messages between mobile phones, but MMS messages are capable of including attached files, much like email with attached files.

The most threatening propagation vectors permit rapid and widespread virus penetration throughout a network of phones. Based on this criterion, one of the most significant threats is propagation by MMS message attachments [11], [12], [7]. Thus, we choose to focus on mobile phone viruses spreading via MMS messages.

Mobile phone security measures can leverage existing antivirus efforts against traditional computer viruses, but the effectiveness of these measures must be evaluated in the context of the mobile phone network environment. Mobile phone viruses are expected to follow an evolution similar to that of computer viruses, only at an accelerated pace [12].

Our model of mobile phone virus propagation leverages related work in computer virus modeling. Kephart and White introduced epidemiological models to the study of computer viruses [6]. More recent work utilizes Markov models to incorporate the probability distributions of model behavior [1]. Some other related work on models of email viruses has influenced our work on models of mobile phone viruses. In much the same way that models of email virus propagation incorporate user behavior [14], our model of mobile phone viruses considers factors such as how quickly a phone user reads a new MMS message and how likely a phone user is to open an infected attachment. In addition,

37th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN'07)
0-7695-2855-4/07 $20.00 © 2007 IEEE

our development of mobile phone response mechanisms is related to research on defenses against computer network worms [8]. Some researchers have proposed mobile phone defense measures [5], but they perform no quantitative evaluation on their proposed measures.

In this paper we present research quantifying the impact of virus propagation on the dependability of mobile phone systems. We also quantify the effectiveness of a range of potential virus mitigation techniques. Section 2 describes the general attack process of a mobile phone virus, and Section 3 describes the six response mechanisms to be evaluated. In Section 4, we discuss the implementation of the model used to generate the results presented in Section 5. Four test case virus scenarios are defined. We analyze the simulation results by comparing the virus propagation with and without response mechanisms.

2. Mobile Phone Virus Attack Process

For the mobile phone virus propagation that we model, the infection starts with a single infected phone. The virus on this phone sends MMS messages with an infected attachment file to other phones. These targeted phones are either selected from the contact list of the infected phone or selected by dialing a random phone number.

Each infected MMS message is delivered to its target phone. After the user of the target phone notices this new MMS message, the user must choose whether to accept the accompanying attachment. If the unsuspecting user accepts the infected attachment file using a phone susceptible to the virus, then the virus is installed, the target phone becomes infected, and the target phone begins to function as an attacker phone.

3. Mobile Phone Virus Response Mechanisms

In response to the mobile phone viruses spreading via MMS messages, we present mechanisms intended to slow or stop the infection dissemination. In contrast to the situation with email viruses, where the antivirus vendor is typically separate from the Internet service provider, mobile phone service providers have expressed an interest in developing and deploying antivirus measures. As a consequence, these response mechanisms can incorporate the network infrastructure hardware owned by the mobile phone service provider, as well as the information already collected by the phone service provider for billing purposes.

In this section, we propose six response mechanisms for mobile phone viruses. We categorize the response mechanisms as actions taken at one of three response points during the virus propagation process: the point of reception by target phones, the point of infection on target phones, and the point of dissemination from infected phones.

3.1. Virus Response Mechanisms at the Point of Reception

The first two response mechanisms focus on preventing infected MMS messages from reaching their intended targets. These response mechanisms use the infrastructure owned by the mobile phone service provider.

Virus scan of all MMS attachments in an MMS gateway. During the normal delivery process for an MMS message, the mobile phone service provider routes the MMS message through its MMS gateway hardware. As each MMS message passes through a gateway, this virus scan response mechanism examines the MMS attachment for known virus signatures. Attachments identified as infected are prevented from reaching their intended recipients. Admittedly, when a new virus appears, there is lag time between the initial appearance and when the new virus signature can be added to the list of known viruses. Our experimental results illustrate how the length of that delay affects the relative effectiveness of this response mechanism.

Virus detection algorithm in an MMS gateway. While the virus scan response mechanism identifies specific known virus signatures, the virus detection algorithm approach is more universal and can detect previously unidentified viruses. The algorithm identifies infected MMS messages by looking for suspicious traits characteristic of a virus. When a virus is first detected, the virus detection algorithm in the MMS gateway analyzes the infected messages to determine the best way to recognize the presence of this virus in subsequent MMS messages. After the analysis period is complete, the MMS gateway detection algorithm successfully recognizes and stops each subsequent virus-infected MMS with some probability. We study how high this probability must be in order for the detection algorithm to be effective.

Both the virus scan and the virus detection algorithm operate within the MMS gateway infrastructure of the phone service provider. These response mechanisms at the point of reception stop the infected message in transit before the message reaches the target phone. The next line of defense involves stopping the virus at the point of infection.

3.2. Virus Response Mechanisms at the Point of Infection

The next two response mechanisms focus on the infected MMS messages that have already passed through the MMS gateways and have arrived in the inboxes of target phones. The goal here is to stop the virus from actually infecting the target phone. This can be accomplished by stopping the user from accepting the infected MMS attachment or by immunizing the phone against the virus attack.

Phone user education. Educating phone users about the risks associated with accepting and installing unsolicited MMS message attachments can help reduce the probability that users will choose to accept infected messages [3]. Since the user acceptance of the virus is a vital link in the virus propagation, reducing the probability of acceptance has a direct impact on the ability of a virus to spread. Many people are still unaware of the existence of mobile phone viruses, and educating those phone users would encourage them to be more cautious concerning suspicious MMS messages. Phone user education can also include warning messages when the user attempts to perform actions that would potentially compromise the security of the phone. For example, the installation of digitally unsigned executable files could trigger a warning message to the user.

Our experimental results illustrate how decreasing the probability that a phone user will accept a virus to only one-half or one-fourth of the baseline acceptance rate can limit the virus spread.

Immunization using software patches. Although the phone user education response mechanism strives to dissuade the user from accepting infected messages, other response mechanisms, such as immunization, can prevent infection even if the user accepts the MMS message attachment. The immunization response mechanism operates using software placed directly on each mobile phone.

After the service provider detects a virus that exploits a vulnerability, the service provider begins developing a patch to fix that vulnerability. Once the patch is developed, the immunization software resident on each mobile phone automatically installs any immunization patches available. Due to bandwidth constraints, all the phones cannot receive the patch simultaneously, so the patch is rolled out to the entire phone population uniformly over a period of time. The more servers that are dedicated to distributing these patches, the faster the deployment to all susceptible phones in the network. After the deployed patch arrives at a particular phone, that phone becomes immunized from the virus if not already infected, or the patch stops further propagation attempts from the phone if the phone is already infected.

Our experimental results show how both the time to develop the patch and the time to distribute the patch to the entire population of susceptible phones can influence the effectiveness of this response mechanism. Varying the patch distribution time is equivalent to varying the number of servers dedicated to deploying the patch.

Immunization and phone user education are both defensive response mechanisms to protect uninfected phones from becoming infected. However, after a phone has already been compromised, the response mechanism must act offensively to stop further dissemination of the virus.

3.3. Virus Response Mechanisms at the Point of Dissemination

The final two response mechanisms focus on containing the virus spread by preventing infected phones from disseminating more infected messages. Virus spread can be contained if propagation efforts by infected phones are detected and suppressed.

Monitoring for anomalous behavior. Some anomaly detection algorithms for mobile phones already exist [10], [2]. Before the monitoring response mechanism can detect anomalous virus behavior, the monitoring mechanism must first be trained to recognize normal user behavior. Our monitoring mechanism is a count of the number of MMS messages sent from a particular phone during a period of time. When the monitor detects an excessive number of outgoing MMS messages (above a threshold based on normal expected usage), the behavior is flagged as suspicious.

When a phone is suspected of being infected, there are several possible responses, including simply alerting the phone user, completely blocking subsequent outgoing messages from the phone, or adding a forced waiting time between outgoing messages. For the monitoring response mechanism in our experiments, the forced delay between outgoing messages is imposed on phones that exceed the specified threshold. Our studies compare the effectiveness of the monitoring response mechanism while varying the length of the enforced minimum time between outgoing messages.

Blacklist phones suspected of infection. In contrast to the monitoring response mechanism that counts all outgoing MMS messages (infected or not), the blacklisting response mechanism counts only messages suspected of being infected. Then, when the number of suspected infected messages for a phone reaches some threshold value, the service provider places that phone on a blacklist and completely stops MMS service for that phone (until the phone is proven to be uninfected). Our experiments determine how low the threshold must be for blacklisting to be effective against different types of viruses.

In summary, the six proposed response mechanisms are categorized based on the three response points in the propagation process: the point of reception by target phones, the point of infection on target phones, and the point of dissemination from infected phones. The effectiveness of these six response mechanisms is evaluated using a model of virus propagation in a mobile phone system.

4. Model Implementation

To quantify mobile phone virus spread and evaluate the effectiveness of the response mechanisms, we perform sim-

ulations using a parameterized stochastic model of a network of mobile phones. Some parameters control virus behavior and are varied to produce different virus scenarios. Other parameters control specific characteristics of the response mechanisms.

The scope of the model includes only mobile phone viruses that spread between phones via infected MMS message attachments. The model only simulates the MMS traffic due to the virus and does not track the delivery of legitimate messages between the phones. The mobile phone viruses that are simulated here infect only the phones themselves, not the phone network infrastructure. It is also assumed that the phone network infrastructure can support the extra volume of MMS messages generated by the viruses.

In this section, we describe how the model construction facilitates evaluating the effectiveness of response mechanisms. Using the general parameterized model, we show how four specific virus scenarios are defined. These four viruses are the test cases for evaluating the response mechanisms. The section concludes with brief discussions of the topology of contact list connections and the role of phone user consent in virus propagation.

4.1. Modeling Phones and Phone Networks

To simulate virus propagation, we first develop a model representing the mobile phone system in which the virus operates. The entire phone system model is developed in the Möbius software tool [4] and is composed of 1000 individual phone submodels, of which 800 are randomly designated as susceptible to infection. We assume that there is enough homogeneity in the population of mobile phones—the same operating system platform or the same application software—that 80% of the mobile phone population could be vulnerable to the same virus.

Each phone submodel represents a single phone and is initialized and assigned a unique identification number. Then the phone is given a contact list containing the identification numbers of other phones. The contact lists are reciprocal; if phone 22 is in the contact list of phone 83, then phone 83 is in the contact list of phone 22. The contact lists connect phones so that MMS messages can be sent between them.

The submodel for each phone contains two functionalities: receiving and sending infected messages. The portion of the model that receives messages is the only active part of the model for phones that are still uninfected. The incoming infected MMS messages wait in the inbox until the phone user makes a decision whether to accept (open) the MMS message attachment. The decision to accept the MMS message occurs with some defined probability. If the user rejects (deletes) the MMS message attachment, then the infection attempt was unsuccessful. However, if the user chooses to accept the MMS message attachment, then that phone becomes newly infected.

After a phone becomes infected, the portion of the phone submodel that sends out infected messages becomes enabled. Several parameters control the frequency at which outgoing infected messages are dispersed. The virus may restrict the total number of infected messages sent from a particular phone within a certain time period (e.g., 30 messages per day).

Because the model is implemented in a parameterized fashion, many different virus behaviors can be simulated. For example, the propagation process can identify new target phones either by using the contact lists of infected phones or by randomly selecting mobile phone numbers. Another example of the parameterized options is that each infected message can be addressed to single or multiple recipients.

4.2. Four Illustrative Virus Scenarios

The flexibility of our parameterized phone virus propagation model enables the study of a large variety of possible viruses. However, to perform any meaningful analysis, we must choose feasible sets of input parameters that characterize potential viruses. We define four example virus scenarios that demonstrate a range of attack approaches based on real mobile phone viruses such as CommWarrior.

Virus 1. When a mobile phone is infected with Virus 1, the phone immediately begins to send infected MMS messages to the phones in its contact list. To avoid alerting the phone user that something is amiss, the virus waits at least 30 minutes between consecutive infected messages, and each message is sent to a single recipient.

Virus 1 also limits itself to sending 30 messages between reboots of the phone. This limit is based on behavior seen in the mobile phone virus CommWarrior. The time between phone reboots is on average approximately 24 hours.

Virus 2. Compared with Virus 1, Virus 2 attempts to spread much more aggressively, engaging in behavior that a phone user might more readily recognize as suspicious. As with Virus 1, a phone infected with Virus 2 immediately begins to send infected MMS messages to the phones in its contact list; however, Virus 2 waits a minimum of only one minute between consecutive infected messages instead of the minimum 30-minute wait for Virus 1. In addition, Virus 2 addresses each infected MMS message to multiple recipients (up to 100 recipients per message). These factors dramatically increase the speed at which Virus 2 can reach all the contacts in the contact list of an infected phone.

The main throttle on the number of messages that Virus 2 spawns is that only 30 infected MMS messages can be sent from each infected phone per 24-hour period. Because

the minimum wait between infected messages is so short for Virus 2, those 30 messages are all sent very near the start of each 24-hour period. This non-uniform nature of the active infection spread of Virus 2 will be evident in the simulation results.

Virus 3. Virus 3 propagates by dialing random mobile phone numbers. In France, all mobile phone numbers start with the same prefix, and approximately one third of the possible phone numbers with the mobile phone prefix are valid mobile phone numbers. This parameter—the fraction of valid random mobile phone numbers—can be adjusted to reflect other circumstances.

When a phone becomes infected with Virus 3, the phone immediately begins to send MMS messages to random mobile phone numbers. One-third of the attempted phone numbers are valid. The minimum wait between these infected messages is one minute, and each message is sent to only one phone number. This virus imposes no daily limits on the number of infected messages sent, so the spread of Virus 3 is very rapid.

Virus 4. The final example virus is the most stealthy virus of the four. When a phone is infected with Virus 4, the phone does not immediately begin sending out infected messages as the other viruses do. After an initial one-hour dormancy period, this stealthy virus waits until the phone user sends or receives a legitimate MMS message and then automatically either appends the infection to outgoing MMS messages or sends infected reply messages in response to incoming MMS messages. Although the model implementation does not include legitimate message traffic, the model still simulates sending infected messages in conjunction with legitimate incoming and outgoing traffic. The model does so by sending out infected messages at the same rate that a phone might expect to send and receive legitimate messages. The virus is less likely to be noticed by the phone user because the user already expects some data transmission to occur while sending or receiving legitimate messages.

To perform a quantitative analysis of virus spread, we choose combinations of parameter values to simulate the four specific virus scenarios described above. Some response mechanisms are more effective against some types of viruses than others, so this suite of virus test cases can demonstrate the strengths and weaknesses of each response mechanism.

4.3. Phone Contact List Network Topology

Since the contact lists define the connections over which three of the four example viruses spread, the MMS contact lists should appropriately reflect the structure of connectivity within a real phone network.

Although the structure of connectivity through mobile phone contact lists is unknown, email address books can be represented by a power law network [14]. Since a contact list is populated based on the same general social network principles as an email address book, it is not unreasonable to use a power-law random graph to represent the contact list connections within a phone population.

To generate a random graph to represent realistic contact list connections between phones, we utilize the software package Network Graphs for Computer Epidemiologists (NGCE) [13], which is an open source software package for generating network graphs. We modify this graph generation software to produce a contact list output file to be read as input by our Möbius model. Since we expect the sizes of the contact lists for a population of 1000 to conform to a certain distribution, we are able to manipulate the graph package input parameters to produce contact lists with an average contact list size of 80.

4.4. Probability of User Consent

Although the topology of the contact list network can influence the penetration and speed of a mobile phone virus, the virus propagation is also affected by the probability that a phone user will consent to the installation of an infected attachment file. Since users are likely to become more suspicious (and less likely to accept the attachment) as they receive more and more infected MMS messages, the model uses a dynamic probability of acceptance that is dependent on the total number of infected messages that the phone user has previously received.

The decreasing probability of acceptance curve is defined as some initial quantity called the Acceptance Factor divided by the quantity two to the power of the number of infected messages received by that phone. Thus, when the Acceptance Factor is 0.468, as it is in our simulations, the probability of acceptance for the nth received message is $0.468 \div 2^n$. Thus, given that the user receives a large number of infected messages, the probability that a user will eventually give consent to accept an infected file is 0.40.

5. Experimental Results

To evaluate the relative effectiveness of the six proposed response mechanisms, we have defined four test-case virus scenarios using feasible combinations of input parameters. Baseline experiments simulate virus propagation unconstrained by any response mechanism. Experimental results from simulations of the model then demonstrate how effective each response mechanism is against each of the four virus scenarios.

One measure to gauge the effectiveness of a response mechanism is a count of the total number of infected phones

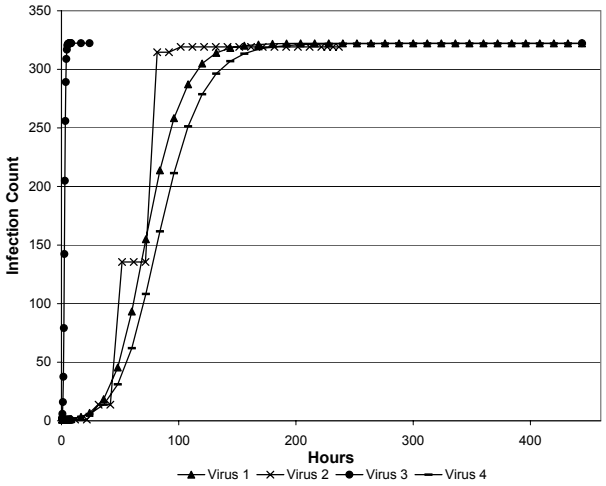

Figure 1. Baseline Infection Curves without Response Mechanisms

in the simulation population. The simulation population contains 800 phones that are susceptible to infection, and the total probability that any given phone user will eventually accept an infected message and become infected is 0.40. Therefore, given enough time, all the unrestrained viruses (without any active response mechanisms) can be expected to infect approximately $800 \times 0.40 = 320$ phones, assuming each susceptible phone receives enough infected messages.

Before we can evaluate the effectiveness of any response mechanism, we must first examine the baseline virus spread without any response mechanisms.

5.1. Baseline Studies

All four of the virus scenarios produce classic virus infection curves, although Virus 2 displays a more jagged curve. As shown in all four curves in Figure 1, the infected population grows at a rate that is first increasing and then decreasing as the number of infected phones reaches a plateau. The virus propagation occurs on different time scales for different viruses. The progression of Virus 2 is tracked over 10 days, and Viruses 1 and 4 are examined over an 18-day period. In contrast, Virus 3 travels so quickly that the simulations only record the infection spread over a 24-hour period. (For that reason, the baseline infection spread for Virus 3 is better observed in Figure 6, which also includes the monitoring response mechanism results for Virus 3.)

The baseline infection curve for Virus 2 resembles a step function more than a smooth curve due to the definition of the virus. The minimum waiting time between infected messages being sent from an infected phone is only one minute (contrasted with a 30-minute wait for Viruses 1 and 4), so the virus sends its whole allotment of 30 messages allowed per day within the first hour of each 24-hour period. This results in a step-like infection curve.

Because of the model parameters held constant, the peak number of infected phones is 320 for all four virus scenarios without response mechanisms. In all four scenarios, 800 phones are susceptible, and each phone user has a 0.40 probability of eventually accepting the virus and becoming infected (provided the phone receives enough opportunities to accept the virus). However, the reaction mechanisms affect the propagation of different viruses in different ways.

5.2. Response Mechanism Studies

Some response mechanisms completely stop further virus propagation, but others simply slow the propagation rate of the virus. Both types of response can be useful. Ideally, the response mechanism would always quickly and completely stop the propagation of a mobile phone virus. However, some viruses spread so quickly that a first response mechanism that slows the spread could buy time to enable activation of a secondary response mechanism that completely halts the propagation process. In the following studies, each response mechanism is evaluated independently.

Virus scan of all MMS attachments in MMS gateways. A virus scan of all MMS attachments as they pass through an MMS gateway is completely effective against viruses with known virus signatures. For that reason, after the new virus signature is added to the list of known viruses, the gateway virus scan is able to completely halt virus propagation.

The gateway virus scan response mechanism is evaluated for three cases. The time required to identify and add the new virus signature to the list (after the virus reaches a detectable level) is varied from 6 hours to 12 hours to 24 hours. As Figure 2 illustrates for Virus 1, a prompt response is most effective because the infection is contained before the virus spread reaches the rapid propagation portion of the curve. When the activation delay is only six hours, the infection only reaches 5% of the infection level in the baseline. Even for an activation delay as large as 24 hours, the virus spread is still contained to 25% of the baseline infection level.

For Viruses 1, 2, and 4, the results with the gateway virus scan look similar because the response mechanism is able to respond while the virus spread is still in its early stages. In contrast, the gateway virus scan is completely ineffectual against rapid viruses like Virus 3 because the virus has already completely penetrated the entire susceptible population before the new virus signature is added to the watch list.

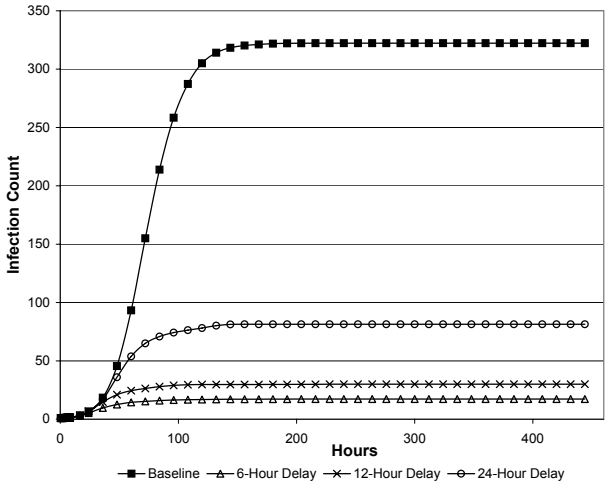

Figure 2. Virus Scan: Varying the Activation Time Delay (Virus 1 shown)

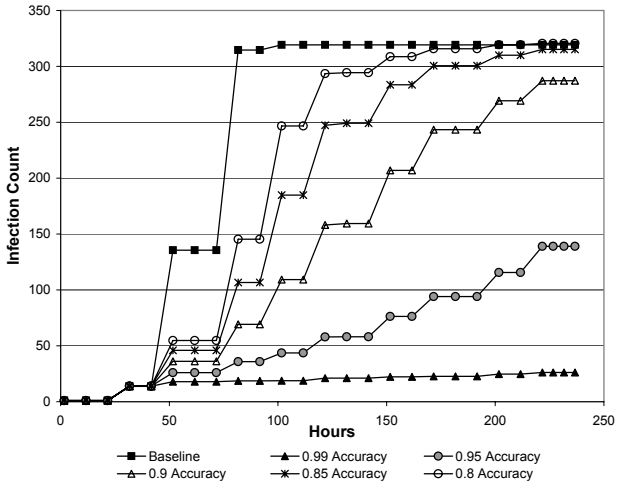

Figure 3. Virus Detection Algorithm: Varying Detection Accuracy (Virus 2 shown)

Thus, the relationship between the speed of the virus propagation and the response mechanism deployment is critical to the success of this response mechanism.

Virus detection algorithm in MMS gateways. In contrast to the gateway virus scan response mechanism, the gateway virus detection algorithm is able only to slow the virus spread, not stop it. Because the detection algorithm attempts to identify infected MMS messages by looking for suspicious traits characteristic of a virus, the algorithm does not catch 100% of the infected MMS messages sent through the MMS gateways. Thus, a small percentage of infected messages still reach target recipients, so the potential for some virus spread remains.

The accuracy of the detection algorithm is a critical factor in the effectiveness of this response mechanism. Therefore, the detection algorithm is evaluated at different levels of accuracy based on the percentage of infected messages that are successfully detected and stopped: 80%, 85%, 90%, 95%, and 99% accuracy. Figure 3 displays how the infection spread of Virus 2 is slowed by the detection algorithm. When the detection algorithm accurately stops 95% of the infected messages, the number of infected phones reaches 135 after nine days of propagation. However, without this reaction mechanism, Virus 2 has infected 135 phones after only two days of propagation. The difference between two and nine days is significant because the extra time could enable the phone service provider to find a more permanent fix to the problem that could completely halt the virus spread.

Like the gateway virus scan, the gateway detection algorithm produces similar results for Viruses 1, 2, and 4. The gateway detection algorithm is ineffective against rapid-spreading Virus 3 for the same reason that the gateway virus scan is ineffective: the response mechanism cannot react fast enough.

Phone user education. Since all four of the illustrative virus scenarios require the consent of the phone user to accept the message and infect the phone, changing the probability that a user will accept an infected message has a direct effect on virus propagation.

In the baseline virus scenarios, the total probability that a user will accept an infected message is 0.40. When the phone user education response mechanism is evaluated, the virus spread is examined for the cases in which the total probability of acceptance has been reduced to 0.20 or 0.10. In each case, the 0.20 total probability of acceptance produced a final infection level at one-half the baseline level. Similarly, the 0.10 total probability of acceptance produced a final infection level at one-quarter the baseline level. Figure 4 shows the baseline spread for each virus scenario (with total probability of acceptance equal to 0.40), as well as the phone user education response mechanism (with the total probability of acceptance reduced to 0.20). When the user education response mechanism is enabled, the total number of infected phones plateaus at approximately 80, which is 25% of the number of infected phones in the baseline case. This reduced plateau at 80 infected phones is observed in the infection curve for all four viruses with user education enabled.

Because reducing the probability that a phone user will accept infected MMS messages is the most consistent defense against any type of mobile phone virus requiring user consent, phone user education should be part of any long-term virus response effort. Decreasing the probability of acceptance both slows and eventually stops the virus spread.

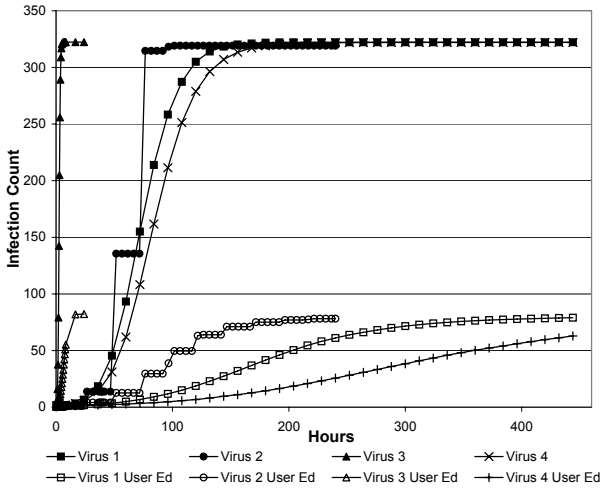

Figure 4. Phone User Education: Effective for All Viruses

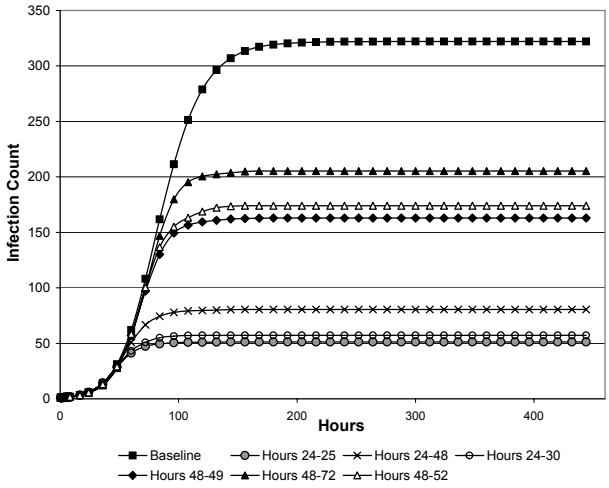

Figure 5. Immunization Using Patches: Varying the Deployment Times (Virus 4 shown)

The caveat is that education is an ongoing effort due to the constant influx of new users.

Immunization using software patches. Mobile phone immunization involves the installation of patches to fix vulnerabilities that a virus could otherwise exploit. Both the time to develop the patch and the time to install the patch on every susceptible phone contribute to the delay in fully activating this response mechanism. As the analysis of other response mechanisms has demonstrated, the time required to fully deploy a response mechanism can have a large bearing on its effectiveness, especially concerning rapidly spreading viruses such as Virus 3.

Of the six tested variations of the immunization response mechanism, three require 24 hours to develop the patch after the virus becomes detectable, and the other three require 48 hours. Within each set of three, the length of time to deploy the patch to all susceptible phones varies from 1 hour to 6 hours to 24 hours. As shown by the results for Virus 4 in Figure 5, the patch development time determines how long the virus is permitted to spread unrestrained. Each curve is identified by the hours during which the deployment is in progress. For example, the "Hours 24-30" curve displays the results when patch development requires 24 hours and distribution requires an additional 6 hours. The three most effective cases, in which the patch is developed in only 24 hours, start limiting the virus spread earlier in the propagation curve than do those cases that require 48 hours to develop the patch.

Regardless of the patch development time, the length of time to fully distribute the immunization patch (1, 6, or 24 hours) influences how much more the virus can spread during the patch distribution process. When patch deployment begins 24 hours after the initial virus detection and occurs uniformly over a 24-hour period, approximately 60% more phones become infected than if the patch deployment had occurred over only one hour. However, the trade-off to a fast deployment is that many servers are necessary to handle the large amount of bandwidth, which can be expensive.

Viruses 1 and 2 once again show results comparable with Virus 4, and Virus 3 once again resists the efforts of a response mechanism. Virus 3 moves too fast for a patch to be developed and deployed in time to be effective.

Monitoring for anomalous behavior. The final two response mechanisms are responsible for limiting the attempts of infected phones to send outgoing infected MMS messages. Since monitoring detects sharp peaks in activity, monitoring for anomalous behavior is most effective against aggressive viruses that attempt to send an extremely large number of messages within a short time period. Once activated, the monitoring response mechanism introduces a forced waiting period between any two consecutive messages, which greatly slows the pace of virus propagation.

The monitoring response mechanism is evaluated while the length of the enforced waiting period is varied from 15 to 30 to 60 minutes. Figure 6 displays the effect of the monitoring response mechanism on fast-moving Virus 3. The speed of Virus 3 makes it resistant to response mechanisms with long activation times, but that same aggressive nature is what enables the monitoring response mechanism to identify its suspicious behavior. Even when the imposed waiting time between all outgoing messages from a suspected infected phone is only 15 minutes, this response mechanism can still constrain the infection level to under 150 phones for up to 20 hours. In contrast, the baseline Virus 3 can in-

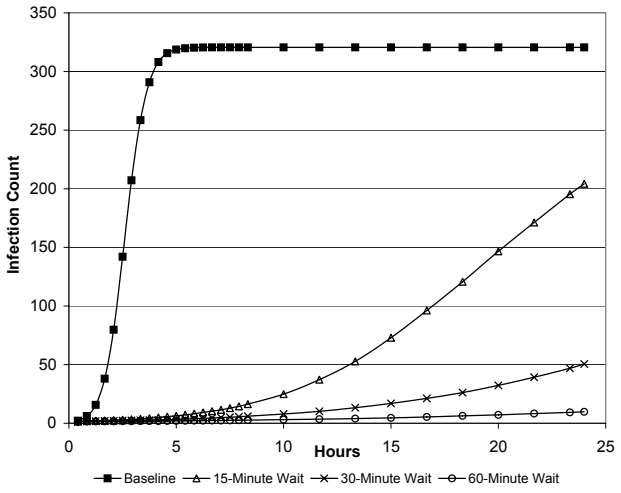

Figure 6. Monitoring: Varying the Wait Time for Suspicious Phones (Virus 3 shown)

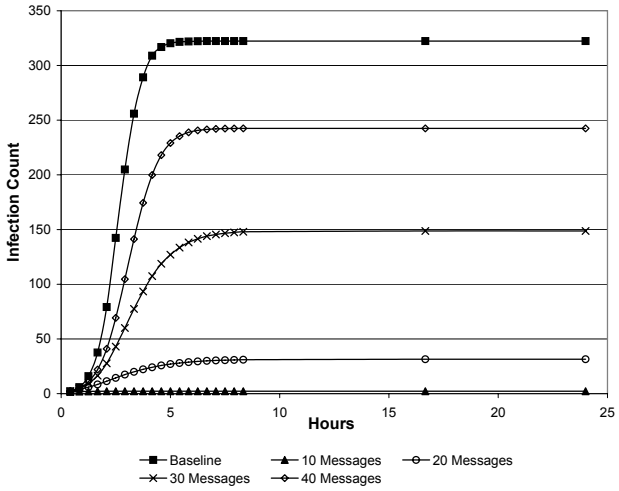

Figure 7. Blacklisting: Varying the Activation Threshold (Virus 3 shown)

fect 150 phones in only about two and one-half hours. The monitoring response mechanism buys time for a secondary response mechanism to be implemented and stop a rapidly-spreading virus.

Although very effective against Virus 3, the monitoring response mechanism is ineffectual against Viruses 1, 2, and 4 because the self-imposed constraints of those viruses limit the total number of messages sent from each phone per unit time. As a result, the volume of infected messages sent from any one infected phone within a monitoring observation period is not radically different from the volume of normal message traffic. Thus, the monitoring response mechanism does not effectively detect Viruses 1, 2, or 4.

Blacklist phones suspected of infection. The blacklist response mechanism blocks all outgoing messages from a phone after the number of suspected infected messages exceeds some threshold. The threshold should ideally be as high as possible to avoid false positive activation of the blacklist response, but the threshold must be low enough to effectively restrict the dissemination of infected messages. To study the effectiveness of the blacklist response mechanism, infected phones are blacklisted after 10, 20, 30, or 40 infected messages.

Blacklisting at a threshold level of 10 infected messages is somewhat effective for Viruses 1 and 4. The infection penetration is restricted to approximately 60% of the baseline infection penetration. However, blacklisting at higher thresholds is ineffective for these viruses.

Blacklisting is completely ineffective for Virus 2 at any threshold level because Virus 2 sends each infected message to many recipients, so the number of infected messages sent from a phone does not accurately capture the amount of virus propagation activity.

The blacklist response mechanism is most effective against Virus 3 (Figure 7) because Virus 3 propagates to random phones without using contact lists. Only one-third of the randomly-addressed infected messages are sent to valid recipients, but all of those infected messages count toward the threshold limit of suspected infected messages. Therefore, blacklisting with a threshold level of 30 infected messages implemented against a virus with random propagation is equivalent, in terms of effectiveness, to blacklisting with a threshold level of 10 against a virus with contact list propagation (where all contact phone numbers are assumed to be valid).

5.3. Optimal Response Strategy

Each response mechanism is designed to slow or stop the propagation of mobile phone viruses, but different response mechanisms are needed to fight different types of viruses.

For rapidly propagating viruses like Virus 3, the most effective response mechanisms are based on monitoring for anomalous behavior, such as the excessive volume of outgoing messages generated by an infected phone. The specific response strategy implemented after the virus is detected determines whether the response mechanism merely slows the virus spread (as monitoring did) or completely stops infected messages from being sent from the infected phone (as blacklisting did).

For more slowly propagating viruses, a more discriminating response mechanism is necessary to identify the presence of a mobile phone virus. This response can occur in the MMS gateway infrastructure owned by the phone service provider or in individual mobile phones. The ad-

vantage to stopping infected messages in MMS gateways is that the mechanism is controlled by the phone service provider and is largely hidden from phone users. Also, response mechanisms in the MMS gateways could mitigate some traffic congestion due to infected messages.

Though possibly less straightforward to implement, educating phone users about the risks of mobile phone viruses should not be neglected. Since user education has universal effectiveness, this response mechanism could enhance the overall effectiveness of any virus mitigation strategy.

Because mobile phone viruses have the potential to attack in many different ways, an optimal response strategy must incorporate mechanisms to counteract a wide variety of virus behaviors. Although the results presented here use a population size of 1000 phones, additional experiments with a 2000-phone population demonstrate that our results scale nicely to larger population sizes.

Our results would also be valuable in conjunction with implementation cost data for each response mechanism. Since the implementation costs could vary greatly depending on the implementation details and the existing infrastructure of an individual service provider, broad cost-based comparisons between response mechanisms without company-specific cost data would be difficult to justify. However, we can still assume that there are increasing costs associated with implementing a stronger version of the same response mechanism. Given this, the results of our experiments are useful for locating the point of diminishing returns for each individual response mechanism, the point where implementing a faster or more accurate response mechanism does not much improve the success rate of the response mechanism.

6. Conclusions

Mobile phone viruses present an emerging problem that threatens the dependability and security of mobile phone communications. We proposed six response mechanisms that respond to this threat at three response points in the propagation process: the point of reception by target phones, the point of infection on target phones, and the point of dissemination from infected phones. To quantify the effectiveness of these response mechanisms, we developed a model to simulate virus propagation with and without response mechanisms. Within the model, four specific virus test cases were used to evaluate the effectiveness of the response mechanisms. The experimental results revealed that response mechanisms must be agile enough to respond quickly to rapidly propagating viruses and discriminating enough to detect more stealthy, slowly propagating viruses. An optimal virus response strategy must be able to address many different types of virus behavior.

This work can be extended with an evaluation of combinations of reaction mechanisms, particularly when a response mechanism that only slows virus propagation requires a secondary mechanism to completely halt virus spread. This same virus propagation modeling approach can also be used to evaluate response mechanisms for mobile phone viruses that spread through means other than MMS messages, such as viruses that spread using the Bluetooth interface on a phone.

Acknowledgments The authors would like to thank Olivier Billet of France Telecom and Sankalp Singh of the University of Illinois at Urbana-Champaign for useful technical discussions and Jenny Applequist for her editorial comments. The authors thank France Telecom for funding support for this research.

References

[1] L. Billings, W. Spears, and I. Schwartz. A unified prediction of computer virus spread in connected networks. *Physics Letters Review*, pages 261–266, May 2002.

[2] A. Boukerche and M. Notare. Behavior-based intrusion detection in mobile phone systems. *Parallel and Distributed Computing*, (9):1476–1490, 2002.

[3] D. Dagon, T. Martin, and T. Starner. Mobile phones as computing devices: The viruses are coming! *Pervasive Computing, IEEE*, 3(4):11–15, Oct.-Dec. 2004.

[4] D. Deavours, G. Clark, T. Courtney, D. Daly, S. Derisavi, J. M. Doyle, W. H. Sanders, and P. G. Webster. The Möbius framework and its implementation. *IEEE Trans. on Software Engineering*, 28(10):956–969, Oct. 2002.

[5] C. Guo, H. Wang, and W. Zhu. Smart-phone attacks and defenses. *HotNets III*, Nov. 2004.

[6] J. Kephart and S. White. Directed-graph epidemiological models of computer viruses. In *IEEE Comp. Soc. Symp. on Research in Security and Privacy*, pages 343–359, 1991.

[7] N. Leavitt. Mobile phones: The next frontier for hackers? *Computer, IEEE Computer Society*, 38(4):20–23, Apr. 2005.

[8] M. Liljenstam and D. Nicol. Comparing passive and active worm defenses. *Conf. on Quant. Eval. of Sys.*, Sept. 2004.

[9] R. Racicm, D. Ma, and H. Chen. Exploiting MMS vulnerabilities to stealthily exhaust mobile phone's battery. *SECURECOMM*, 2006.

[10] B. Sun, F. Yu, K. Wu, and V. Leung. Mobility-based anomaly detection in cellular mobile networks. *ACM Workshop on Wireless Security*, 2004.

[11] Trend Micro. *Security for Mobile Devices: Protecting and Preserving Productivity*, Dec. 2005.

[12] S. Viveros. The economic impact of malicious code in wireless mobile networks. *4th Intl. Conf. on 3G Mobile Communication Technologies*, pages 1–6, Jun. 2003.

[13] V. Vlachos, V. Vouzi, D. Chatziantoniou, and D. Spinellis. NGCE — network graphs for computer epidemiologists. In *Advances in Informatics: 10th Panhellenic Conf. on Informatics*, pages 672–683, Berlin, Nov 2005. Springer-Verlag.

[14] C. C. Zou, D. Towsley, and W. Gong. Email worm modeling and defense. *Computer Communications and Networks*, pages 409–414, Oct. 2004.