Cooperative Response Strategies for Large Scale Attack Mitigation

D. Nojiri, J. Rowe, K. Levitt
University of California, Davis
dnojiri@ucdavis.edu, rowe@cs.ucdavis.edu, levitt@cs.ucdavis.edu

Abstract

We introduce models for cooperative mitigating response strategies to suppress large scale Internet worm attack. In our models, cooperating members communicate with others using a "friend protocol" that spreads attack reports to potentially vulnerable uninfected sites. We use mathematical models for the simplest strategies and a simulation for more complex models of mitigation. We investigate the performance of different strategies both in the presence of large scale worms and of false alarms.

1 Introduction

Responding to Internet-scale attacks poses many new problems. First, and most obvious, is the collective nature of the victims. Single organizations attempting to manage attacks using only local knowledge will have very little effect on Internet-scale worm incidents. Response mechanisms will necessarily extend beyond the borders of any single organization's network. The second problem is closely related to the first; how do different inter-organization control structures perform when responding to Internet-scale attacks? That is, what type of control works best against specific models of distributed coordinated attack. Hierarchical control, for example, may produce a faster mitigating response to a detected spreading attack but only in a localized region of control. Peer-to-peer control mechanisms give more widespread coverage, but propagation of the necessary information is slower. There are several issues that complicate automated mitigation control strategy:

- Trust between organizations is never complete. Most organizations would refuse to implement an automated response strategy on the advice of a competitor alone, for example. Close cooperative arrangements might be acceptable, however, between strategic business partners. Models of mitigating automated response will have to take this into account.
- Responses will typically be costly. Control mechanisms, then, need to be able to distribute the expense of response equitably among cooperating partners. A mitigating response may completely thwart the attacker, but in the process hinder critical business transactions of a single partner. Such a strategy would never be acceptable to that partner and is therefore infeasible when using a peer-to-peer type control structure. An authoritative node in a control hierarchy, however, might direct the single partner to act for the good of all sub-nodes collectively.
- There will be an ambient level of false alarms. These false reports will lead to a continuous triggering of unnecessary responses. The models for automated mitigation must not lead to massive propagation of false alarm messages and the accompanying potential for mass automated denial-of-service. We continue to model different methods of cooperative mitigation and study how variants of our model perform in the face of Internet-wide coordinated attack.

1.1 Cooperative Peer-to-peer Strategies

In this paper we focus on peer-to-peer control structures and investigate their efficacy in stopping large scale Internet worms. Under this type of control, all policies are decided upon locally within a single organization; no directives from external central authorities are considered. In our model, direct cooperation occurs only between a limited number of *friend* organizations. When a site detects suspicious worm-like behavior, its initial cooperation strategy is to share the information with its friend organizations. Those organizations, upon receiving the report, act according to their local policy as to their preferred course of action. In our model, these actions are limited to blocking the suspect behavior and/or sharing the information with its own set of friend sites. This combination of blocking and sharing produces a propagating mitigating response whose rate of spread is similar to that of the worm itself. In this paper we present two separate approaches to modeling worm mitigation strategies. First are purely mathematical models of worm propagation when simple defense strategies are

0-7695-1897-4/03/$17.00 (C) 2003 IEEE

in place. Secondly, to address more complex strategies, we present simulated cooperative models. We discuss the desirable features of a peer-to-peer cooperative response strategy when faced with an Internet-wide malicious worm attack.

2 Mathematical Models

For simple worm mitigation strategies, mathematical models with closed form solutions are possible. Considerable work has been done describing models for worm propagation in the absence of any coherent mitigation measures; some of these will be described below. These are similar to mathematical models of population dynamics and disease propagation seen in the biological sciences. Based upon these, we develop models showing how worm propagation is modified in the presence of simple mitigation mechanisms.

2.1 Mathematical Models of Virus/Worm Propagation

Staniford's Virus Propagation Model S. Staniford, et al. in their paper[5] analyze last year's Code Red worm by developing a quantitative model of its propagation. Their model is as follows:

N	:	Initial total # of vulnerable hosts
a	:	proportion of infected hosts
K	:	# of hosts each infected host can find and compromise

Since the infection ability of a worm is proportional to the density of the target hosts, the successful infection in dt is $K(1-a)dt$. There are Na worms in total, thus, the rate at which infected hosts increase during the time period of dt is

$$Nda = (Na)K(1-a)dt.$$

Dividing by N,

$$\frac{da}{dt} = aK(1-a) \tag{1}$$

with solution:

$$a = \frac{e^{K(t-T)}}{1+e^{K(t-T)}} \tag{2}$$

where T is a constant that fixes the time position of the incident. This equation produces the S-curve growth behavior typically seen in population growth models with limited environmental carrying capacity.

Kephart's virus infection model[3] J. Kephart and S. White created another mathematical model by representing an individual system as a node in a graph. Directed edges from a given node j to other nodes represent the set of individuals that can be infected by j. We introduce it briefly here, and refer to the original paper [3] for details.

i	:	the proportion of the infected hosts
$\bar{b}$	:	The expected number of nodes around a node
β	:	The infection rate of virus
δ	:	The cure rate of each node

They showed a deterministic differential equation describing the time evolution of $i(t)$:

$$\frac{di}{dt} = \beta\bar{b}i(1-i) - \delta i \tag{3}$$

Note that if the second term, which describes the cure rate of a host, is taken out this becomes the same as Staniford's model.

2.2 Dynamic response strategy

Based on Kephart's model, we develop a mathematical model of a simple worm mitigation strategy. We consider a graph where each node represents a mitigation-enabled member, and is connected by directed edges to its cooperating friends. In this model the variables are defined as,

M	:	the total # of response members
a	:	the number of infected members
c	:	the proportion of alerted members
F	:	the # of friends of each response member
α	:	the # of alerts a response member needs before it changes its state

For a given member, the expected number of cooperating friends who remain in the normal, unalerted state is

$$F \cdot (1-c).$$

In our simple strategy, a cooperating member will increase the severity value of the alert messages that it shares as more infection attempts are seen. Thus the number of alerts a particular responding member sends in a certain period of time dt is

$$F(1-c) \cdot \sigma a \cdot dt$$

where σ is the severity value assigned to the alert. The system-wide total number of alerts in dt, then, is

$$F(1-c)\sigma a \cdot Mc \cdot dt.$$

From this we obtain the following differential equation describing the time evolution of the proportion of response members that have been alerted to the presence of a malicious worm.

$$\frac{dcM}{dt} = \frac{F(1-c)\sigma aMc}{\alpha}$$

Dividing by M,

$$\frac{dc}{dt} = \frac{F(1-c)\sigma ac}{\alpha} \tag{4}$$

To obtain the proportion of members who have already been compromised by the malicious worm, we change (1) to include the fact that cooperatively alerted members will be able to block worm activity. We consider two types of infection attempts: local infection and global infection. When a particular infected host tries to propagate the worm to a site controlled by a different response member, it must pass two response devices: the local response device blocking outgoing infection attempts and the remote response device that protects the potential target by blocking incoming infections. The probability that both of the response devices are not alerted is $(1-c)^2$. Thus, the rate for the global infection is

$$aK(1-a)(1-c)^2.$$

The local infection rate is the same as (1) because there is no response device between the infection source and the target. Since the average # of hosts connected to a response device is N/M, the probability that a worm chooses a remote host as a target is $1-1/M$. The probability that a worm chooses a local host as a target is $1/M$. Combining the global and local infection rates along with these probabilities, we have

$$\frac{da}{dt} = aK(1-a)(1-c)^2 \cdot \left(1 - \frac{1}{M}\right) + aK(1-a) \cdot \frac{1}{M}. \tag{5}$$

We have a pair of simultaneous differential equations: (4) and (5). To solve them numerically, we use the following equations:

$$\begin{cases} t_{k+1} &= t_k + h \\ a(t_{k+1}) &= a(t_k) + h \cdot \frac{da}{dt}\Big|_{a(t_k),c(t_k)} \\ c(t_{k+1}) &= c(t_k) + h \cdot \frac{dc}{dt}\Big|_{a(t_k),c(t_k)} \\ a(t_0) &= a_0 \\ c(t_0) &= c_0 \end{cases} \tag{6}$$

where h is a step interval, and a_0 and c_0 are an initial proportion of infected hosts and alerted response devices (detectors), respectively. The result of the numerical solution is shown in Figure 1, where the solid line indicates the proportion of members infected by the worm over time, and the dashed line shows the proportion of members that have been alerted to the presence of the worm and are responding. The horizontal time axis is arbitrary.

Note that both curves follow the typical S-curve growth pattern. This is due to the similarity of the cooperation strategy to the spread of the worm itself. Single infected members propagate alerts to cooperating partners, effectively producing a “white worm” message propagation behavior. To effectively counter a malicious worm incident, the number of cooperating partners must be large enough to overtake the spread of the worm itself and protect uninfected sites. It is worth noting that in spite of limiting the worm spread, this model also exacts the maximum cost of response; not only does each site protect against worm infection, each site blocks potentially misclassified good behavior as well.

Dynamic response strategy with back-off mechanism We introduce a new model that reduces the penalty paid when good behavior is inadvertently blocked due to cooperative mitigation measures. Alerted members will now back off and remove any blocking filters after a particular time interval has elapsed. The back-off rate depends on the proportion of infected hosts. The more worms a response member sees, the more slowly it will back off. Thus, the rate can be

$$\epsilon \cdot (1-a) \cdot c,$$

where ϵ is a constant which indicates how fast a response member does back off [1]. The differential equation (4) becomes

$$\frac{dc}{dt} = \frac{F(1-c)\sigma ac}{\alpha} - \epsilon(1-a)c. \tag{4'}$$

3 Simulated Models of Mitigation Strategies

Effective mitigation strategies that suppress large scale worm propagation in a cost-effective manner might involve more complex decision making by cooperating members. The effect of these strategies taken by an individual member upon the global dynamics of the worm, and the global cost of the aggregate response can no longer be modeled by simple mathematical formulae. We have developed a simulation to investigate the global properties of these more complex strategies. Given that the mathematical models are very close to population dynamics in the biological sciences, we base our simulations on the Swarm simulation package[7]. Swarm simulates interactive agents, and is widely used to model population dynamics and simple

[1] Note this is very simple case. More precisely, a member will back off only when it has no infected host in its domain, thus the back-off rate can be written as $\epsilon(1-a)^{N/M} \cdot c$.

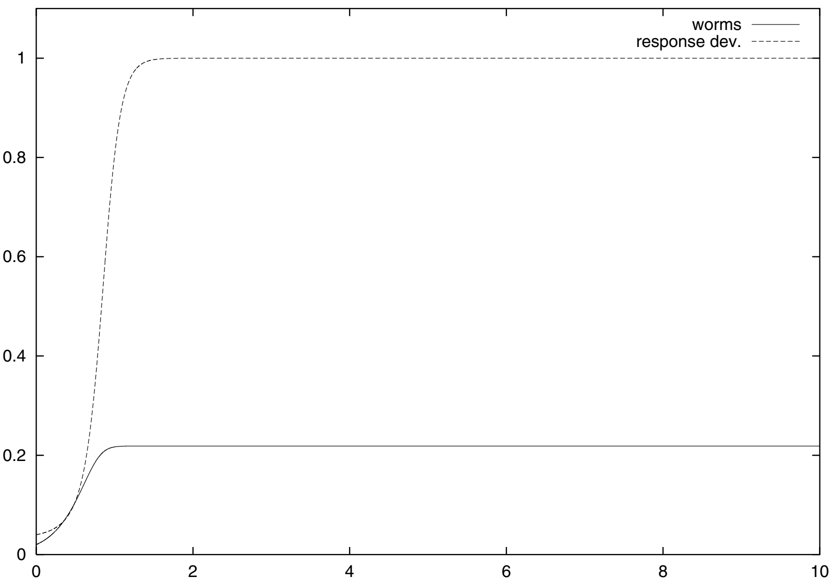

Figure 1. Numerical solution

social behaviors in organisms. In this section, we describe settings and parameters involved in the simulation and show results in some different cases.

Topology We simplify the Internet topology by considering it as a flat network. In our simulation we set up 5832 vulnerable hosts and 729 cooperating members connected to a shared network by idealized response devices. Members cooperate by sharing worm reports. When the number of worm reports received exceeds a certain threshold, a member's response device protects its collection of vulnerable hosts from infection, potentiallly at the cost of inadvertantly blocking misidentified desirable activity as well. As seen in figure 2, the network is separated inside and outside by the response devices. Each host is connected to the outside network through a response device, and each response device has a direct connection to all other response devices.

Response devices Each response device has eight hosts in its local network, and can watch both traffic coming in from the outside network and traffic going out from the local network. A response device has two states: normal and alerted. It sends alerts to its cooperating friends to inform them of the ongoing worm attack. When in the normal state, it receives alerts and raises its alert severity level, but does not send alerts to other response devices. Once its alert severity threshold is exceeded, it begins blocking worm infection attempts in both incoming and outgoing traffic, and it cooperatively shares alert messages with its friends. When a response device sees no worm activities for a certain period of time, it backs off (changes its state to "normal" and stops blocking worm activity). In this mitigation strategy model, response devices are controlled by the following parameters:

- The average number of vulnerable hosts protected by each member
- The number of cooperating friends
- Threshold for state change
- Back-off speed
- Alert severity

When a cooperating member receives an alert from its friend, it increments its alert level by the value of the alert severity. The state of a member's response device changes to "alerted" when its alert level exceeds the threshold for state change as set by the model parameter. It transitions back to "normal" when the alert level falls below the threshold.

Hosts We assume that all the vulnerable hosts in our model have a certain security hole which allows a worm to gain control of the system. Hosts have two states: normal and infected. A host recovers after a certain period of time, and it becomes immune to the worm. This models the effect of efforts to patch the vulnerability, clean up the damage and continue normal operation. Parameters controlling host behavior in our model are:

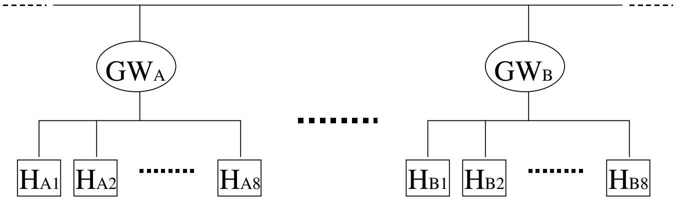

Figure 2. Topology

- Total host population
- Initial infected host density
- Infection probability - the probability that an infection attempt succeeds
- Scanning method - random, permutation, or hitlist scanning
- Infection frequency - The # of time steps between infection attempts

Figure 3 shows the result of the simulation where a worm infects a single new host each propagation time interval, the average number of cooperating friends is 16, the alert severity sent by a member detecting the worm is three, and the severity threshold is seven (i.e., three message from friends are needed to trigger a response). The solid line shows the fraction of vulnerable hosts that have been infected by the worm while the dashed line shows the proportion of a member's response devices that are in the alerted state, actively blocking suspect behavior. At the beginning of the simulation, worms spread quickly because there are only a small number of response devices in the alerted state. As an important calibration point, notice that the early portion of the time evolution graph is the same as the mathematical model of worm spread without back-off and patch mechanisms. As the worms spreads, some of the members detect worm activity, and subsequently send alerts to their friends. The proportion of alerted response devices grows very quickly to catch up with and overtake the proportion of worms. As the percentage of alerted response devices approaches 100%, the response devices protecting uninfected hosts in their local domain back off, while those which have infected hosts stay alerted. At this point the worm has been captured by the alerted members and confined within their local networks. With these parameters, the infected host percentage is suppressed to less than 20%. As time progresses the model assumes that infected hosts are patched and, when all of its locally protected hosts have been patched, the alerted members back off independently. The worm reaches its maximum infection extent in 15 time steps. Alerted members, however, remain in their blocking state for nearly 90 time steps, paying the cost of response nearly 6 times longer than the period that the worm actively infects new hosts. By manipulating the parameters, one can test the effectiveness of different mitigating defense strategies in the face of a large-scale worm attack.

3.1 The number of friends

Here, we show how the average number of friends in our mitigating strategy model affects the global features of a large scale worm attack. As in figure 3, figure 4 shows the proportion of alerted response devices and the proportion of infected hosts. The solid lines, dashed lines, and dotted lines correspond to the case where the number of cooperating friends are 8, 12 and 16 respectively. Also as before, infected hosts propagate the worm to one additional vulnerable host every propagation time interval. As expected, the greater the number of cooperating friends in the strategy, the greater the suppression of the worm, and the shorter the time to recovery. When only 8 friends cooperate, nearly 50% of all vulnerable hosts are compromised before the alert is go]received by all sites. As we shall see later, however, a greater number of friends performs worse in the face of false alarms.

3.2 Slow worms

Generally, worms spread very quickly, and the faster they move, the more dangerous they are. Cooperative peer-to-peer mitigating response strategies will not be as effective against these very fast “flash” worms. Our preliminary studies show that hierarchically controlled models will likely be more effective in this case. There are, however, stealthy worms which spread very slowly to hide their activities. Figure 5 shows the way this type of worm spreads, and how our defense strategy reacts.

Even after a majority of the response devices become alerted and infected hosts are encapsulated by them, the worm continues to spread. This is for the following reason: when response devices do not see worm activities often, they back off, assuming hosts are patched. But in reality hosts are not patched, they simply have yet to be infected by the slow moving worm. Slow worms, then, keep

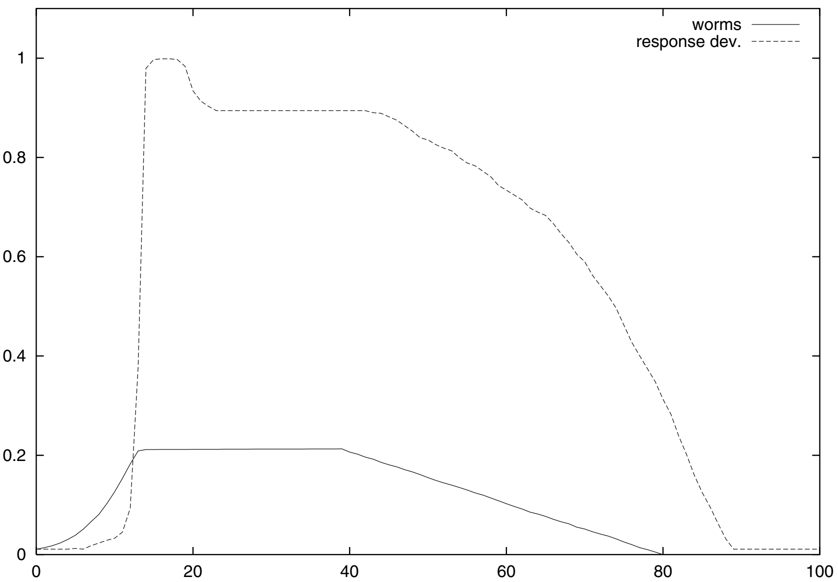

Figure 3. Typical case of the simulation

spreading after the local response devices back off. To keep response devices alerted one can simply make the back-off speed slower, but this is not desirable in terms of cost when faced with faster spreading worms. What is needed is more complex analysis that classifies worms according to the features of their spread so that the best strategy can be selected by members and shared with cooperating friends. These analysis techniques are beyond the scope of our current work.

3.3 False alarms

The ability of a particular mitigation strategy in supressing a large scale worm attack is only part of the measure of its effectiveness. Internet scale worm attacks are relatively rare but a cooperative mitigation system will need to operate at all times. It is safe to assume that under normal conditions there will be a constant ambient level of false worm reports issued by members. Friends receiving the reports and accepting them as true worm activity will pay the needless cost of reacting to a non-attack. The greater number of cooperating friends, the greater the cost associated with false alarms. Raising the threshold of reports needed in order to respond will reduce the cost associated with false alarms, but will also decrease the proportion of sites protected from attack. Figure 6 shows the reactions of response devices with different values of the threshold for state change when there is a false alarm instead of worms. Here we assume that, at a certain time, 5% of all members falsely report a worm attack to their cooperating friends. The dotted, dashed and solid lines show the case where the alert threshold is 5, 7 and 10 respectively. With an alert threshold of 5, nearly 75% of all members block when only 5% originally reported the false alarm. At threshold of 10, a negligible number of members has responded. The more sensitive response devices are, the more of them react to a false alarm, and consequently, the cost rises. One can simply reduce the sensitivity (threshold) of response devices, but doing so increases the risk in the case of real worm attacks.

3.4 Optimal friend lists

Finally, we show the difference of the response abilities between randomly constructed friend lists and optimized friend lists. To have optimized friend lists, we use the following algorithm developed by Imase and Itoh [8], which constructs a nearly optimal directed graph for given nodes n and degree d with diameter $\lceil \log_d n \rceil$.

1. Let vertices of a directed graph G be labeled as $0, 1, 2, \ldots, n-1$.
2. For every $i, j \in 0, 1, \ldots, n-1$, let there be an arc from vertex i to vertex j if $j = id + \alpha \pmod{n}$, $\alpha = 0, 1, 2, \ldots, d-1$.

We map the vertices of G to the response devices, and if response device i has an arc to response device j, i has j in its friend list. [2] Figure 7 show the differences between

[2] The diameter, $\lceil \log_d n \rceil$, obtained by this algorithm, is at most one larger than the lower bound of the minimum possible diameter of a directed

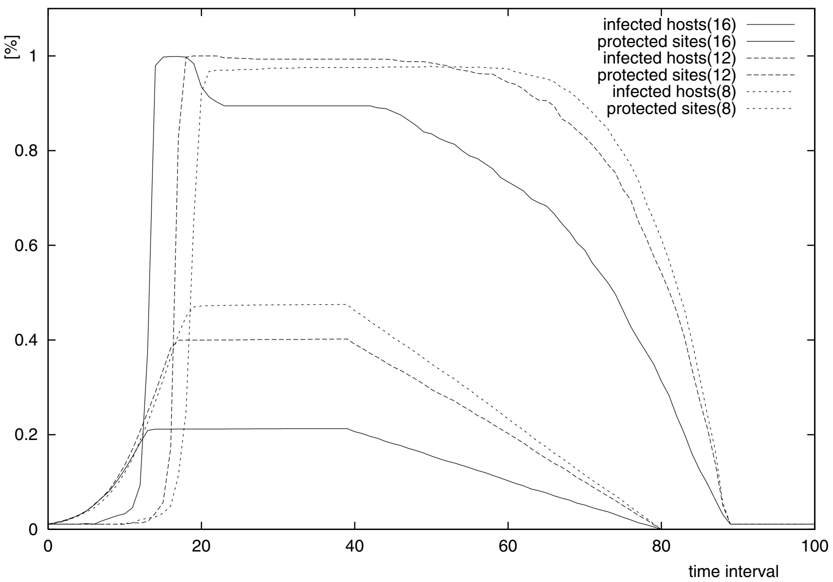

Figure 4. Number of friends: 8, 12, 16

randomly constructed friend lists and nearly optimal friend list. As seen in the graph, using optimal friend lists, one can have quicker reaction of response devices, thus fewer hosts are infected. One can also have better back-off execution, which results in reduction of the cost for the response devices.

4 Conclusions and Future Work

We have described models for peer-to-peer mitigation strategies to suppress large scale worm attacks on the Internet. Our mathematical models show that an effective counter to large scale worms is a controlled "white worm" propagating faster to potentially vulnerable sites. These simple models fail to take the costs of mitigating response into account. For this we have used a simulation to model more complex state-dependent mitigation strategies. Our simulations show that, in general, a larger number of cooperating friends does better in suppressing worms, but is much worse when faced with ambient false alarms that will inevitably be a part of a normal operating environment. Strategies that are effective against fast worms will be largely ineffective against slow, stealthy worms. Automated analysis to classify the propagation behavior is needed for efficient strategy selection. Improvement in the performance of peer-to-peer cooperative strategy is seen when one uses an optimized friend list rather than one constructed randomly.

We will continue to develop mathematical models that can be used to calibrate our simulation results. In our simulation we plan to extend our mitigation models to include more complex strategies to include friends with varying trust levels, member populations containing more than one strategy, more complex worm scenarios and the effect of aggregate worm classification analysis. We continue to pursue in parallel studies of hierarchically controlled mitigation strategies and hybrids.

graph of order n and degree d. So, we have quite efficient connections among response devices.

References

[1] Carey Nachenberg, 'Computer Virus - Coevolution,' Communication of ACM, vol.40, 1997

[2] Carey Nachenberg, 'Computer Parasitology'

[3] Jeffrey O. Kephart and Steve R. White, 'Directed-Graph Epidemiological Models of Computer Viruses,' Procedings of the 1991 IEEE Computer Society Symposium on Research in Security and Privacy; Oakland, California

[4] N. Weaver, V. Paxon, S. Staniford, 'Large Scale Malicious Code: A Research Agenda'

[5] Stuart Staniford, V. Paxon, N. Weaver, 'How to 0wn the Internet in Your Spare Time',

[6] Changchun Zou, 'Code Red worm propagation modeling and analysis'

[7] Swarm Development Group ,'http:/www.swarm.org'

[8] M. Imase, M. Itoh ,'Design to Minimize Diameter on Building-Block Network' ,IEEE Trans. on Computers, Vol. C-30, pp. 439-442, June 1981

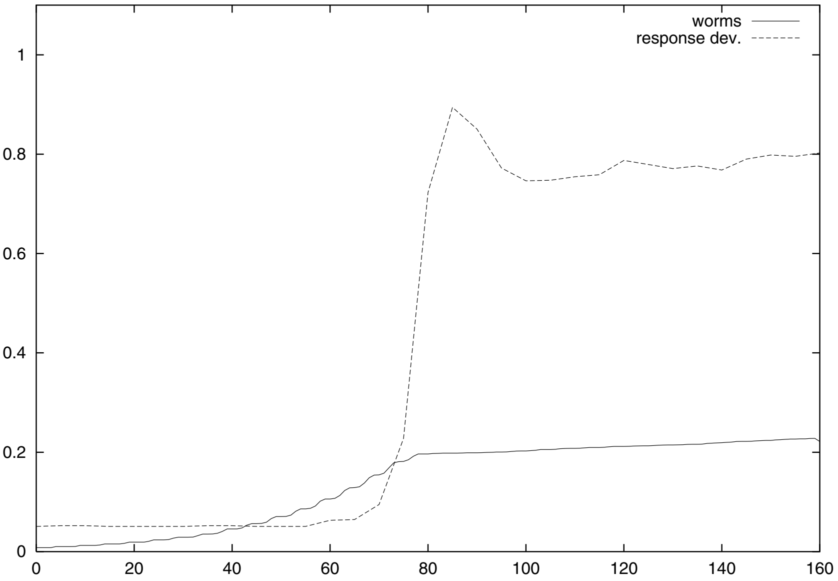

Figure 5. Slow worms

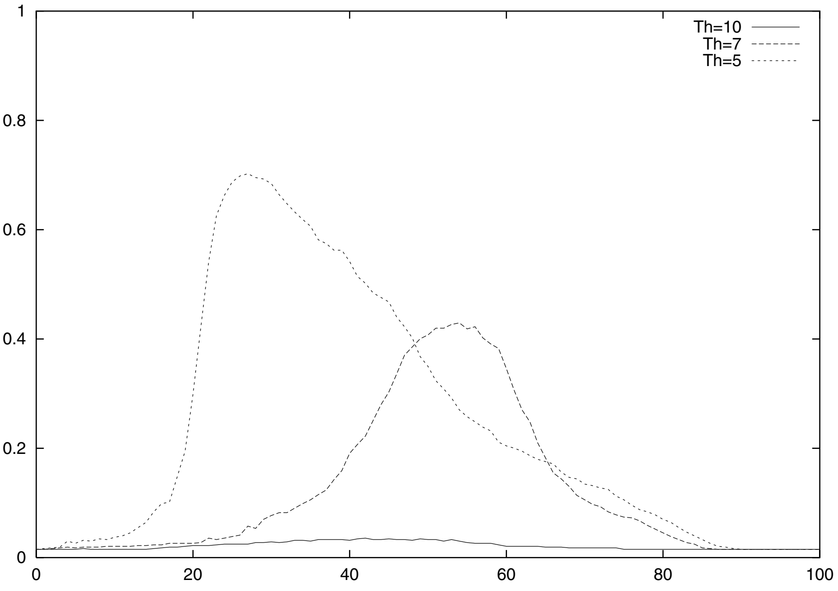

Figure 6. False alarm (Threshold $= 10, 7, 5$)

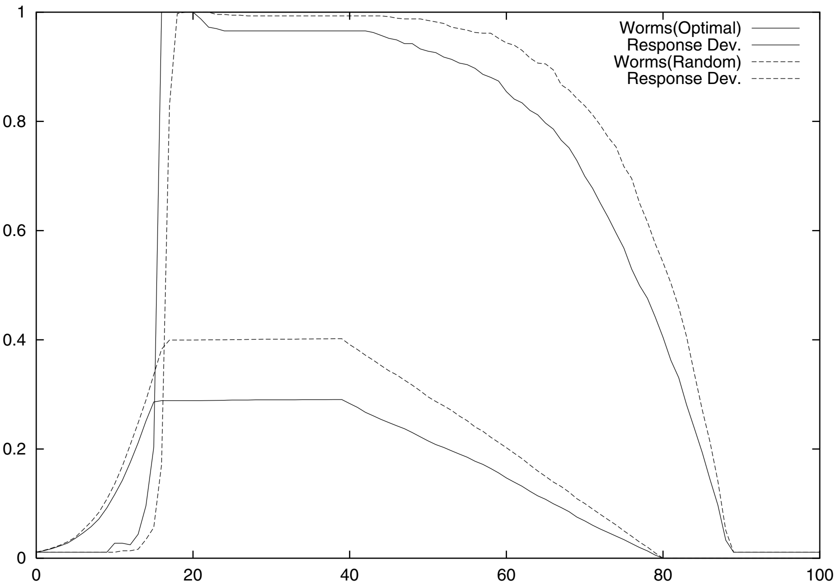

Figure 7. Optimal friend lists vs. Random friend lists