Worm Epidemiology

Thomas M. Chen, Nasir Jamil
Department of Electrical Engineering, Southern Methodist University, Texas, USA
Email: tchen@engr.smu.edu

ABSTRACT

Worms spread by replicating themselves to vulnerable hosts through the Internet. Mathematical epidemiology studies the dynamics of outbreaks spreading through a network population. We describe a community-of-households epidemic model for worms and show how it can be useful to analyze defenses such as dynamic quarantine and rate throttling.

Key words: worm, virus, epidemiology, quarantine, rate throttling

I. INTRODUCTION

Worms and viruses are self-replicating malicious software that spread through the Internet much like infectious diseases spread in human populations [1]. Both worms and viruses have the distinguishing capability to transfer copies of themselves from infected hosts to vulnerable hosts. Viruses are program fragments attached to normal programs or files. They take over control when the normal program is executed to make copies of the virus code. In contrast, worms are automated stand-alone programs that take advantage of network connectivity to seek out and copy themselves to vulnerable new targets [2].

One of the best known examples was the SQL Slammer/Sapphire worm released on January 25, 2003 [3]. It exploited a buffer overflow vulnerability in Microsoft SQL servers. The entire worm including the exploit code was carried in a single 404-byte UDP packet (28-byte IP/UDP header and 376-byte payload). When a vulnerable SQL server was infected by Slammer, it was put into a simple execution loop to send out UDP packets containing a copy of the worm as quickly as possible to randomly generated IP addresses (32-bit numbers).

Because infected computers were sending out copies of the worm as fast as they could, the worm caused heavy congestion throughout parts of the Internet. It shut down thousands of Bank of America ATM machines and disrupted Continental Airline's ticketing system. At the peak of the outbreak, infections were doubling every 8.5 seconds. SQL Slammer was able to hit 90 percent of the vulnerable population (about 90,000 SQL servers) within 10 minutes.

The initial dynamics of the SQL Slammer outbreak can be understood from basic epidemiology models. Mathematical epidemiology has a long history that can be traced back to at least 1760 when Daniel Bernoulli presented a mathematical argument for the effectiveness of smallpox immunization in France [4]. Epidemiology applies deterministic or stochastic models to disease outbreaks with two major goals. The first goal is to predict the future

outcome of an outbreak (number of infections as a function of time). The second goal is to evaluate possible defensive strategies such as immunization or quarantine. Historically, epidemiology was valuable in devising the World Health Organization's smallpox vaccination program which effectively eradicated smallpox globally [5].

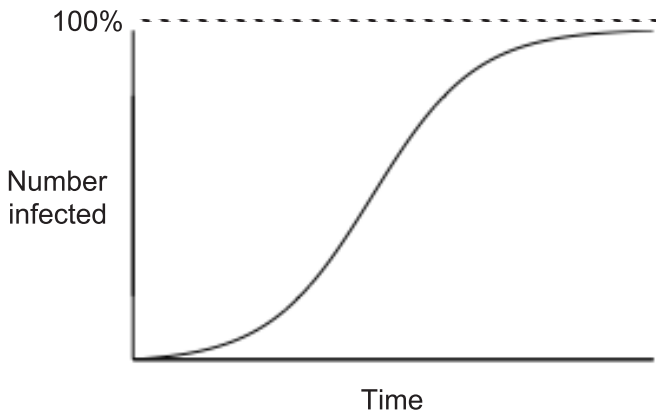

Fig.1 Epidemic rate in simple SI model

The simple epidemic model is also known as the SI (susceptible → infective) model. A population has a fixed number N hosts. Each host begins in the "susceptible" (vulnerable) state and changes to "infective" (or infected) state after contact with an infective host. After a host becomes infected, it will stay in the infective state permanently. More advanced epidemic models consider additional states and more complicated state transitions. Let $S(t)$ and $I(t)$ denote the number of susceptibles and infectives at time t, where $S(t) + I(t) = N$. By totally random mixing, each susceptible is assumed to make an average βN contacts per unit time but the probability of meeting an infective each time is I/N. The parameter β is the pairwise infection rate or infectious contact rate. Hence, the number of infectives increases at a rate of

$$\frac{d}{dt}I = (\beta N)(S/N)I = \beta IS = \beta I(N-I) \quad (1)$$

Given the initial condition $I(0)=I_0$, the epidemic curve is the logistic function

$$I(t) = \frac{I_0 N}{I_0 + (N - I_0)e^{-\beta N t}} \quad (2)$$

The rate of the epidemic is exponential in the early phase, as seen in Fig. 1. When an infective comes into contact with other hosts, the other hosts are very likely to be susceptibles. Thus, the infection spreads easily at an exponential rate. In the later phase, most of the population is already infected. When an infective comes into contact with other hosts, the other hosts are very likely to be already infected. The rate of the outbreak slows down asymptotically because it becomes much harder to find the remaining few susceptibles.

While the simple SI model fits the initial spread of the SQL Slammer worm, it is not a good fit for the later phase of the Slammer outbreak [3]. The main reason is that Slammer worked against itself by spreading so quickly that network links became seriously congested. Network congestion slows down an outbreak because infected hosts can not reach new targets. Also, in the later phase of the Slammer outbreak, human countermeasures (patching, filtering) helped to slow down the spreading.

III. A COMMUNITY-OF-HOUSEHOLDS MODEL

The simple SI model does not consider any existing network structure. It is more realistic to view the Internet as a heterogeneous population, often described as a "network of networks." The Internet is known to consist of separately administered but interconnected autonomous systems or routing domains. This Internet structure can be captured by the community-of-households epidemic model, where each household represents a subnetwork attached to the Internet through an access router as shown in Fig.2. The model has the important feature that the infectious contact rates between households can be different from infectious contact rates between individuals within the same household [4,6]. The community-of-households model has been used

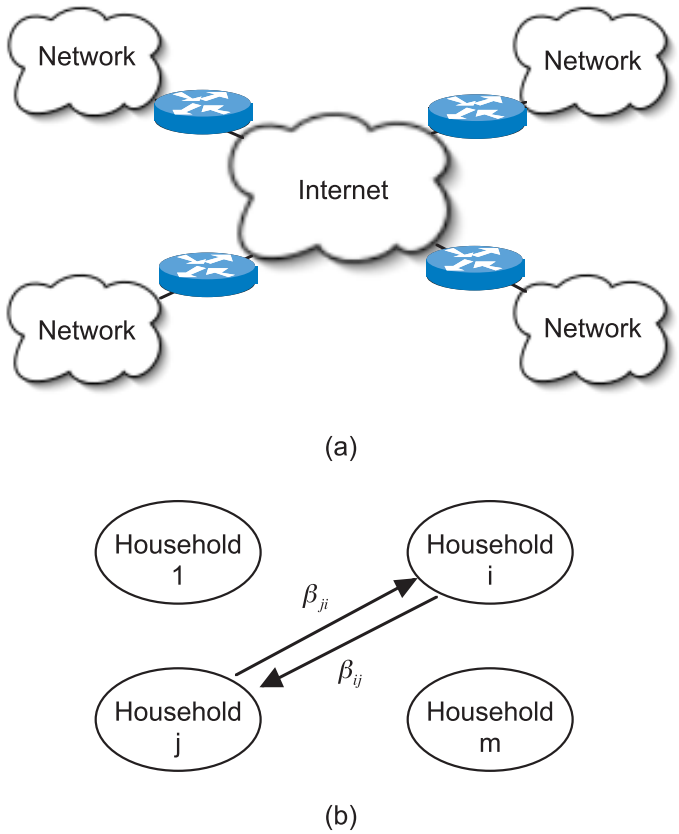

Fig.2 Community-of-households model

since 1955 in biological epidemiology for populations divided into groups with a higher infection rate within groups than between groups [7].

Suppose there are m households, and N_j is the fixed size of household j. Let $I_j(t)$ and $S_j(t)=N_j-I_j(t)$ be the number of infectives and susceptibles in household j, respectively. According to the community-of-households model, the epidemic is governed by a system of differential equations:

$$\frac{d}{dt}I_j = S_j\sum_{i=1}^{m}\beta_{ij}I_i = (N_j - I_j)\sum_{i=1}^{m}\beta_{ij}I_i \quad (3)$$

The parameter β_{ij} is the pairwise infectious contact rate of infectives in household i to susceptibles in household j. It is clear that the number of infectives in household j will increase due to intra-household contacts with rate β_{jj} and contacts with other households with rates β_{ij} $(i \neq j)$. Unfortunately, the system of equations (3) must generally be solved numerically except for the simplest special cases.

IV. WORM DEFENSES

4.1 Preventive

We describe how the community-of-households model can be used to evaluate different possible worm defenses. Today, a variety of strategies are used to protect networks against new worm attacks. The best strategy is prevention of attack by keeping operating systems and application programs up to date on patches and antivirus software updated with signatures. Patches are frequently released by software developers to fix vulnerabilities after they are discovered. Hosts with fewer vulnerabilities will be less likely to be compromised by a new worm.

The effect of preventive measures is to reduce the initial vulnerable population. Since infectives will not be able to interact with as many susceptibles, the epidemic rate will be slowed down as shown in Fig. 3. The community-of-households model can account for preventive measures by reducing the initial vulnerable household populations $\{N_j\}$. The numbers of hosts protected by up-to-date patching and antivirus signatures can be subtracted from the household populations.

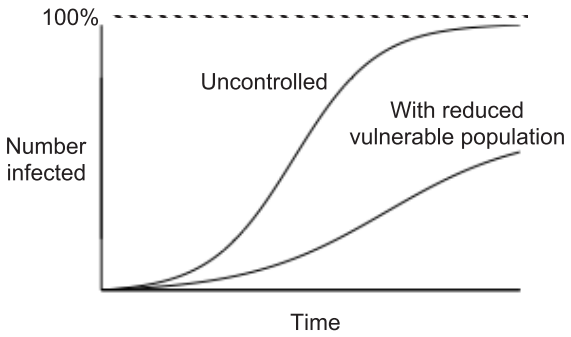

Fig.3 Effect of preventive measures

4.2 Quarantine

In addition to prevention, reactive blocking mechanisms include firewalls, intrusion prevention systems,

and routers with access control lists [2]. These are effective in blocking (quarantining) worm traffic if a worm signature is known [8]. Commercial systems usually use a combination of signature-based and heuristic behavior-based (anomaly) detection. Signature-based detection is preferred due to its accuracy in detecting known worms. However, a new worm may not have a matching signature, and new signatures usually take hours to days to develop, test, and distribute after an unknown worm is discovered. Behavior-based detection is promising for catching unknown new worms without a matching signature, but can result in a high rate of false positives (false alarms). False positives are problematic because legitimate traffic may be blocked and lost.

Ideally, worm blocking will stop an epidemic after the time needed to detect the worm and develop a signature (if a new signature is needed), as shown in Fig. 4. The community-of-households model can account for worm blocking by setting all β_{ij}=0 and β_{jj}=0 in (3), after the worm signature becomes available. However, in practice, it may be difficult to block all worm traffic between every pair of hosts. It is more feasible to block worm traffic at the routers in Fig. 2. This would prevent worms from spreading between subnetworks, but worms may likely continue to spread within subnetworks that are already infected when quarantine begins. This means β_{ij}=0 for $i \neq j$, but $\{\beta_{jj}\}$ are unchanged in the model. In this case, blocking worm traffic between subnetworks will not realize the ideal dampening in Fig. 4. The epidemic rate will be slowed down, but infected subnetworks will eventually become saturated.

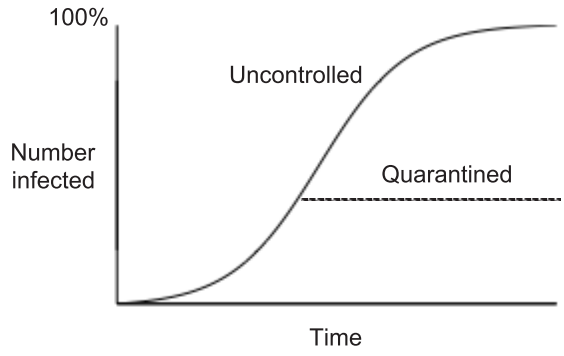

Fig. 4 Effect of quarantining

4.3 Rate throttling

As an alternative, rate throttling has been proposed as a non-destructive approach [9]. The idea is to limit the number of new outbound connections for each host. It has been found that normal hosts show a low rate of outbound connections to different hosts (lower than 2 new connections/sec). On the other hand, hosts infected with worms will exhibit much higher rates of outbound connections because they are searching for new targets. The idea of rate throttling is to limit the rate of new outbound connections for every host such that normal hosts should not be effected, but infected hosts will be significantly slowed down. Even in the case of false positives where a normal host is mistaken for an infected host, legitimate traffic may be delayed at worst but not blocked (lost). Hence, rate throttling has the major advantage that detection accuracy is not critically important. Another advantage of rate throttling is that it can work from the beginning of an epidemic; in contrast, blocking or quarantining can not be exercised until a worm signature is developed.

In the community-of-households model, rate throttling is reflected by reducing all $\{\beta_{ij}\}$ rate parameters. The effect of rate throttling is to slow down the epidemic rate, as shown in Fig. 5. The effectiveness of rate throttling depends on minimizing the β_{ij} rate parameters, but if they are too low, normal users will object to long delays to connect to other hosts. The problem in rate throttling is reducing the β_{ij} rate parameters as much as possible without inconveniencing normal users.

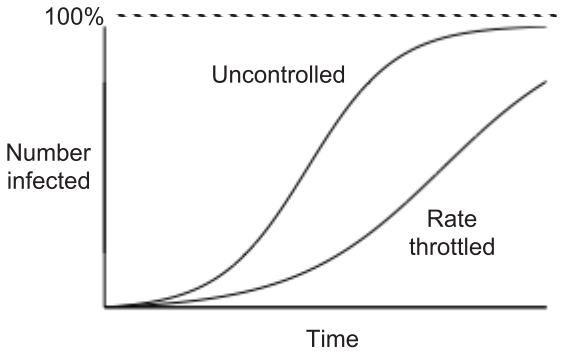

Fig. 5 Effect of rate throttling

V. CONCLUSIONS

In this paper, we have presented the community-of-households epidemic model as a means for studying and evaluating the dynamics of worm outbreaks. The model is simple but accounts for the "network of networks" organization of the Internet. Different defense strategies can be represented and evaluated by the model by appropriate settings of the pairwise infectious contact rates.

For more accurate reflection of reality, the community-of-households epidemic model can be made more complicated in various ways. The most obvious improvement is addition of an additional "recovered" state for infected hosts that are disinfected and then protected against future re-infection. In practice, worm infections can be removed by antivirus software or a clean re-installation of the operating system. Another possible improvement is to make the pairwise infectious contact rates depend on the network load instead of stay constant. When the network becomes congested, the infectious rates should decrease because it will be harder for infected hosts to reach other hosts.

VI. REFERENCES

[1] P. Szor, *The Art of Computer Virus Research and Defense*, Upper Saddle River, New Jersey: Addison-Wesley, 2005.

[2] J. Nazario, *Defense and Detection Strategies against Internet Worms*, Boston, Massachusetts: Artech House, 2004.

[3] D. Moore, et al., "Inside the Slammer worm," *IEEE Security & Privacy*, vol. 1, July 2003, pp. 33-39.

[4] D. Daley J. Gani, *Epidemic Modeling: An Introduction,* Cambridge, UK: Cambridge U. Press, 1999.

[5] N. Bailey, *The Mathematical Theory of Infectious Diseases and its Applications*, 2nd ed., New York: Oxford U. Press, 1975.

[6] N. Becker, J. Hopper, "The infectiousness of a disease in a community of households," *Biometrika*, vol. 70, 1983, pp. 29-39.

[7] S. Rushton, A. Mautner, "The deterministic model of a simple epidemic for more than one community," *Biometrika*, vol. 42, 1955, pp. 126-132.

[8] D. Moore, et al., "Internet quarantine: requirements for containing self-propagating code," *IEEE Infocom 2003*, San Francisco, California, 2003, pp. 1901-1910.

[9] M. Williamson, "Throttling viruses: restricting propagation to defeat malicious mobile code," *18th Annual Comp. Sec. Appl. Conf. (ACSAC 2002)*, Dec. 9-13, 2002, Las Vegas, Nevada, pp. 61-68.

BIOGRAPHIES

Thomas M. Chen *is an Associate Professor in the Department of Electrical Engineering at Southern Methodist University in Dallas, Texas. He received his PhD in electrical engineering from the University of California at Berkeley, and MS and BS degrees in electrical engineering from MIT. Prior to joining SMU, he was a senior member of technical staff at GTE Laboratories (now Verizon). He is the Editor-in-chief of IEEE Communications Magazine, a senior technical editor for IEEE Network, and a past associate editor for ACM Transactions on Internet Technology. He was the recipient of the IEEE Communications Society's Fred W. Ellersick best paper award in 1996. He co-authored ATM Switching Systems (Artech House, 1995). His research is in network security and traffic control.*

Nasir Jamil *is a PhD student in the Department of Electrical Engineering at Southern Methodist University in Dallas, Texas. He is currently working at Nortel Networks in Richardson, Texas.*