Effectiveness of Rate Control in Slowing Down Worm Epidemics

Nasir Jamil and Thomas M. Chen
Department of Electrical Engineering
Southern Methodist University
Dallas, Texas 75275
Email: nasir@mail.smu.edu, tchen@engr.smu.edu

***Abstract*— Rate control is an automated defense to slow down a worm outbreak to buy time for conventional defenses to take effect. In this study, we apply the *community of households* model from biological epidemiology to evaluate rate control strategies. We find that rate throttling of outbound worm traffic, implemented in the network or hosts, can be effective in slowing down a new worm outbreak given sufficient coverage of hosts. An outbreak is slowed down exponentially as the fraction of protected hosts is increased. We also find that throttling both inbound and outbound traffic can be much more effective than rate throttling only outbound traffic.**

I. Introduction

Automated worm programs take advantage of network connectivity to spread from infected hosts to vulnerable hosts. A worm such as SQL Slammer has been called "bandwidth limited" because compromised hosts were put into a simple execution loop to send out 404-byte UDP packets containing a copy of the worm to random IP addresses as fast as they could transmit packets [1]. Slammer was observed to saturate the bandwidth on many links and would probably have taken advantage of more bandwidth had it been available.

Faster worm outbreaks will compel organizations to depend more on automated blocking defenses compared to manual responses such as software patching and reconfiguring router access lists and firewalls. Current automated worm defenses consist of antivirus software, firewalls, and intrusion detection or prevention systems (IDS/IPS) [2], [3]. These defenses attempt to detect and then block worms by a combination of misuse detection and anomaly detection. Misuse detection based on signature matching is the preferred approach of commercial antivirus and IDS/IPS due to its accuracy in detecting known (and similar) worms. However, new signatures may take hours to days to develop, test, and distribute after an unknown worm is discovered.

Anomaly detection looking for deviations from "normal" baseline behavior has the potential to detect new worms without a known signature. Commercial products often include heuristic behavior-blocking rules in addition to signature matching. For example, SMTP worm blocking checks if the process initiating SMTP outgoing email is an attachment in the current email; if it is, this self-mailing software is blocked [3]. Another heuristic rule may look for any signs of a buffer overflow exploit which is a type of attack often used by worms. While anomaly detection is considered useful to catch broad classes of unknown worms, it has been prone to false positives because normal behavior is difficult to characterize precisely.

A high rate of false positives is problematic because legitimate traffic may be blocked. Since blocking is a destructive action, detection accuracy is critical but detection with no false positives is not likely. As an alternative, Williamson and others have advocated the use of rate throttling as a "benign" action [4]. In false positive cases, legitimate traffic will be mostly delayed or at worst dropped if the throttle queue overflows. Hence detection accuracy is not as critical with rate throttling as with blocking (filtering). It is still desirable to minimize the impact on legitimate traffic, so detection accuracy continues to be an issue. Detection of worm-like epidemic behavior is discussed later.

Rate throttling of worm traffic is particularly appealing for high-speed networks because it effectively counteracts the abundant bandwidth offered by the network. Moreover, rate control is entirely complementary to existing worm defenses. In epidemiological terms, rate control does not provide for any "removals" (returning an infected host to normal state protected from subsequent infection). Removals are enabled by signature-based antivirus software or software patching. Rate control serves to slow down and minimize the spread of a new worm outbreak to buy time for conventional signature-based worm defenses.

Host-based rate control such as Williamson's virus throttling is attractive because hosts can best monitor themselves for signs of worm intrusions. However, host-based controls are vulnerable to compromise if the host is infected and may be difficult to deploy universally because most hosts are not managed centrally. It is clear that the same approach of limiting the rate of new outbound connections can be implemented in access routers. Network-based rate control has the advantage that a single access router could cover many hosts. Also, routers with rate control could be deployed easily by network providers. Implementation of rate control in access routers is discussed in Section 2. The existing epidemiological models for worm flow are discussed in Section 3. Little theoretical work has been done to evaluate the network-wide effectiveness of rate control. We apply the *community of households* epidemic model to compare different rate control

strategies in Section 4. Numerical results are presented in Section 5.

II. Router-based Rate Control

The problem in implementing rate control in access routers is detection of unknown worm traffic, although detection accuracy is not as critical as for conventional defenses that filter traffic. Essentially, any infectious "epidemic-like" traffic, which might include spam and certain legitimate applications, is suspect and would be rate throttled. Several traffic indicators of epidemic-like behavior have been studied.

A. Traffic Indicators of Epidemics

Rate of new connections: It is obvious that an infected host should attempt to contact as many vulnerable hosts as quickly as possible, in order to spread successfully before it is stopped by network-based and host-based defenses. Williamson's throttle limits the number of outgoing connections to less than 5 unique hosts per second [4]. Porras et al. takes essentially the same approach further to limit the number of new outbound connections that any host can make within each time interval [5]. Ganger et al. proposed to correlate the rate of new outbound connections with DNS (domain name system) lookups [6].

Rate of connection failures: A random scanning worm will hit a number of addresses that are unused or unreachable on a particular port. For a TCP-based worm, these cases would result in a significant number of TCP resets returned in response to TCP SYN packets. Chen et al. proposed to record the sources of these failed TCP connections, and when the failure rate of an offending source crosses a threshold, the source is deemed to be possibly infected [7], [8]. Schechter et al. proposed a monitoring system to observe the first-contact connection requests made by hosts on a LAN and the rate of connection failures [9]. Berk et al. suggested to monitor ICMP Destination Unreachable messages [10].

Fan-out of identical packets: Singh et al. noted that an indicator of worm behavior is a high volume of identical traffic [11]. Their EarlyBird system does packet matching by generating a hashed signature of every packet's payload and destination port and then comparing it against previously seen signatures. Martin et al. proposed a similar signature-based scheme which enables a router to count how many times it has forwarded a packet of the same signature, and to trigger an alarm when a threshold is exceeded [12]. This approach by itself is limited to non-polymorphic worms which do not change their forms signficantly.

Matching inbound and outbound packets: Chen noted that an incoming worm will infect a host that will subsequently try to exploit the same vulnerability in the outbound direction [13]. They proposed an algorithm for access routers to match inbound and outbound packets with the same destination port numbers. Gu et al. proposed a similar idea [14]. Toth proposed a detection algorithm that involves looking for a similar correlation between inbound and outbound packets, with additional information about connection failures [15].

B. Inbound versus Outbound Rate Control

Access routers could look for the traffic characteristics above to identify likely infected hosts and rate throttle their suspicious outbound traffic. An access router could also use some of these traffic characteristics to identify incoming worm probes, and impose a rate throttle on traffic from these offending sources. However, the detection of inbound epidemic-like traffic is more complex than detection in the outbound direction. Inbound traffic is coming from many different sources, whereas outbound epidemic traffic is generated from a single source. Due to the complexity, inbound detection requires a significant detection time compared to outbound throttling which can be always on.

The capability for combined inbound and outbound rate throttling is an advantage of network-based rate control that is not as easily done with host-based rate control. The case of extreme rate throttling in both directions acts as quarantining. Moore et al. examined several general issues related to quarantining [16]. Also, Zou et al. investigated a temporary quarantining scheme where hosts suspected of being infected are quarantined only for some length of time [17].

III. Epidemic Modeling

Rate control has been proposed as a worm defense, but little is currently known about the network-wide effectiveness or optimal strategies for deployment. Wong et al. compared rate limiting done at hosts, access routers, and backbone routers [18]. They concluded that host-based and access router-based rate control result in a slowdown that is linear to the number of hosts or routers with the rate limiting filter, and that rate control should be deployed at the backbone routers to be most effective. However, the analysis was approximate and relied on a star network topology (leaf nodes interconnected through a centralized hub) that was not representative of the Internet.

As a practical matter, access routers with rate control can not be expected to be deployed universally. If a fraction of access routers are capable of rate control, it is necessary to evaluate the slowdown of a worm outbreak as a function of the percent of coverage. In order to gain insight into these problems, we investigate a heterogeneous *community of households* epidemic model that has been widely used in biological epidemiology [19].

The main purpose of rate control is to slow down the spread of a newly released worm to buy more time for traditional defense mechanisms, which depend on developing a worm signature and re-configuring routers and firewalls. A measure of the effectiveness of a rate control mechanism is how much the outbreak is dampened over an interval of time. Here, the interval of interest is the time between the initial release of the worm and the deployment of conventional defense mechanisms at some time T. In epidemiological terms, the number of infectives can only increase during this time and there are no removals in the time window of interest (before time T).

A. Homogeneous Epidemic Model

Before introducing the community of households epidemic model, it is illustrative to describe the homogeneous epidemic model. A disease outbreak in a homogeneous biological population is usually characterized by the "simple epidemic" or SI (susceptible $\rightarrow$ infective) model [19], [20]. Let $S(t)$ and $I(t)$ denote the number of susceptibles and infectives at time t, where $S(t) + I(t) = N$. By homogeneous mixing, each infective is assumed to make an average βN contacts per unit time but the probability of meeting a susceptible each time is S/N. The parameter β is the pairwise infection rate or infectious contact rate. Hence, the number of infectives increases at a rate of

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

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

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

According to (2), an outbreak will reach an infection level pN at time

$$T_p = \frac{\ln p(N - I_0) - \ln I_0(1 - p)}{\beta N} \tag{3}$$

The SI model appears to be a good candidate for early stages of random scanning worm epidemics. These worms target IP addresses pseudo-randomly which seems to conform to the assumption of homogeneous mixing. Moore et al. showed that the logistic curve predicted by the SI model could fit the observed data for the growth of the Code Red worm well [21]. Liljenstam et al. fit the SI model to the initial spread of the SQL Slammer worm [22]. Zou et al. agreed with the close fit for the early stages of the Code Red outbreak but pointed out a greater than predicted slowdown in the later stages [23]. The discrepancy in the later stages was attributed to the fact that the SI model did not account for network congestion and human countermeasures (such as patching, filtering and isolation).

B. Heterogeneous Epidemic Model

It has long been recognized in biological epidemiology that the assumption of homogeneous and uniform mixing in the simple epidemic model is unrealistic and unnecessarily restrictive. A more realistic assumption is a heterogeneous population consisting of different subpopulations. The *community of households* model views subpopulations as households, where the infectious contact rates between separate households can be different from infectious contact rates between individuals within the same household. Let m be the number of households, and N_j is the size of household j ($1 \leq j \leq m$). The number of infectives and susceptibles in household j are $I_j(t)$ and $S_j(t) = N_j - I_j(t)$, respectively. The dynamics of the epidemic change according to a system of differential equations for $j = 1, \ldots, m$:

$$\frac{d}{dt}I_j = S_j \sum_{i=1}^{m} \beta_{ij} I_i = (N_j - I_j) \sum_{i=1}^{m} \beta_{ij} I_i \tag{4}$$

The parameter β_{ij} is the pairwise infectious contact rate of infectives in household i to susceptibles in household j. According to (4), 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, a general solution for the community of households epidemic is not known, and the system of equations must be solved numerically except for the simplest special cases.

The community of households model has been studied for biological epidemics since 1955 when Rushton and Mautner proposed a deterministic epidemic in a population divided into groups with a higher infection rate within groups than between groups [24]. Subsequent researchers have focused on the problems of parameter estimation [25]; final epidemic size [26]; and mostly immunization strategies [27]–[34]. Unfortunately, these analyses have typically assumed a removal process (where infectives can recover or die) and the results are not directly applicable to the rate control problem considered here.

In the context of worms, Liljenstam et al. proposed worm simulations following the community of households model where the households represent autonomous systems [22], [35], [36]. The Internet is known to consist of separately administered but interconnected autonomous systems or routing domains. The infection parameters were estimated to fit historical worm traffic data. Wagner et al. chose to model worm propagation through an Internet structured as an interconnection of multiple subnetworks [37] characterized by different bandwidth and latency.

IV. Epidemic Models for Different Rate Throttling Scenarios

We consider the network structure shown in Fig. 1 as a community of households. Each household represents a subnetwork attached to an access router.

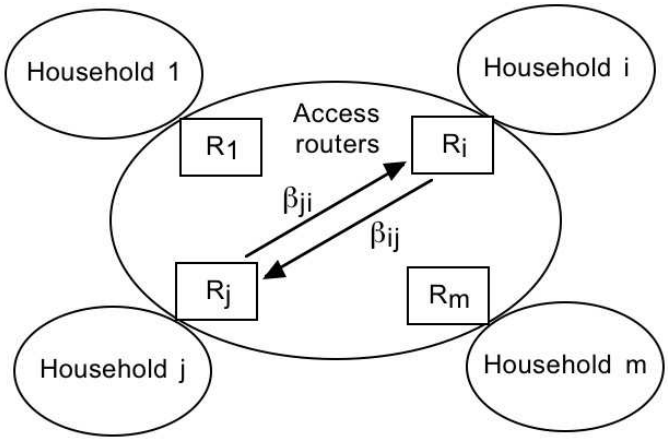

Fig. 1. Interconnected subnetworks represented as a community of households

1) Host-based Rate Throttling: We assume that a fraction q of hosts are capable of throttling their outbound worm traffic, while the other hosts are uncontrolled. For simplicity, we also assume that infection rate parameters may take only two possible values, one much lower than the other. Without rate control, the infection rates are $\beta_{ij} = \beta$ for all i, j. When rate control is exercised, the rate throttling will cause certain infection rate parameters to be attenuated to $\beta_{ij} = \alpha\beta$. In this case, the population can be viewed as two households as shown in Fig. 2. The first household of size qN has intra-household and outbound infection rates $\alpha\beta$. The second household of size $(1-q)N$ has intra-household and outbound infection rates β. The epidemic is governed by the system of two differential equations:

$$\begin{aligned} \frac{d}{dt}I_1 &= (qN - I_1)(\alpha\beta I_1 + \beta I_2) \\ \frac{d}{dt}I_2 &= ((1-q)N - I_2)(\alpha\beta I_1 + \beta I_2) \end{aligned} \tag{5}$$

If the attenuation factor is very close to 0, then (5) becomes

$$\begin{aligned} \frac{d}{dt}I_1 &\approx (qN - I_1)\beta I_2 \\ \frac{d}{dt}I_2 &\approx ((1-q)N - I_2)\beta I_2 \end{aligned} \tag{6}$$

2) Router-based Rate Throttling: If access router k is capable of rate throttling outbound traffic, then the infection rate parameters β_{kj} for all $j \neq k$ would be attenuated by a factor of α. That is, the dynamics of the epidemic would be changed to

$$\begin{aligned} \frac{d}{dt}I_j &= (N_j - I_j)(\alpha\beta_{kj}I_k + \sum_{i\neq k}\beta_{ij}I_i) \\ &\quad , \ j = 1, \ldots, m; \ j \neq k \\ \frac{d}{dt}I_k &= (N_k - I_k)\sum_{i=1}^{m}\beta_{ik}I_i \end{aligned} \tag{7}$$

Additional routers with rate control could be incorporated into the system of equations similarly by attenuating the appropriate infection rate parameters.

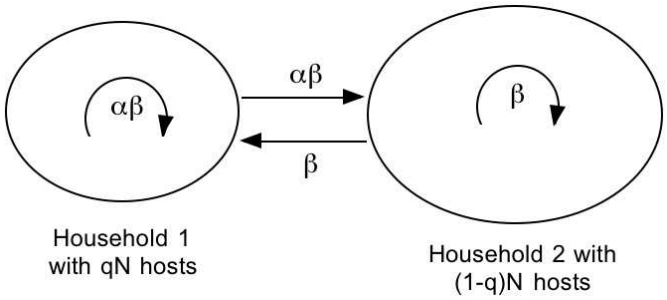

Fig. 2. A community of two households, one representing hosts with and without outbound rate throttling

If access router k is capable of throttling inbound and outbound directions simultaneously, then the infection rate parameter β_{ik} for all $i \neq k$ would also be attenuated by a factor of α. However, inbound rate throttling needs a certain detection time before it can be initiated. Assuming that t_1 is the time for inbound worm detection, equations (9)-(10) represent the outbreak for time 0 to t_1, and the equations below represent the outbreak from time t_1 to T:

$$\begin{aligned} \frac{d}{dt}I_j &= (N_j - I_j)(\alpha\beta_{kj}I_k + \sum_{i\neq k}\beta_{ij}I_i) \\ &\quad , \ j = 1, \ldots, m; \ j \neq k \\ \frac{d}{dt}I_k &= (N_k - I_k)(\beta_{kk}I_k + \sum_{i\neq k}\alpha\beta_{ik}I_i) \end{aligned} \tag{8}$$

Again, additional routers with inbound rate throttling could be factored into the equations similarly.

V. Numerical Results

A. Router-based Outbound Rate Throttling

For the sake of simulation, we assume that households are all equal size $N_1 = \cdots = N_m = N/m$. Suppose that a fraction q of access routers are capable of throttling outbound worm traffic, while the remaining routers are not controlled. Thus, the number of routers with rate control are qm. For simplicity, we assume qm is an integer and the capable routers are numbered 1 to qm (which routers are capable does not matter because all households are equal size). Thus, households 1 to qm have an intra-household infection rate β but outbound inter-household infection rate $\alpha\beta$. The remaining households $qm+1$ to m have intra-household and outbound inter-household infection rates β. These households interact with each other homogeneously and may be viewed as a single large household of size $(1-q)N$. Hence the epidemic in this case changes according to the system of differential equations:

$$\begin{aligned} \frac{d}{dt}I_j &= (\frac{N}{m} - I_j)(\beta I_j + \sum_{i\leq qm, i\neq j}\alpha\beta I_i + \beta I_{qm+1}) \\ &\quad , \ j = 1, \ldots, qm \\ \frac{d}{dt}I_{qm+1} &= ((1-q)N - I_{qm+1})(\beta I_{qm+1} \\ &\quad + \sum_{i=1}^{qm}\alpha\beta I_i) \end{aligned} \tag{9}$$

If the attenuation factor is very close to 0, then (9) becomes

$$\begin{aligned} \frac{d}{dt}I_j &\approx (\frac{N}{m} - I_j)\beta(I_j + I_{qm+1}) \\ &\quad , \ j = 1, \ldots, qm \\ \frac{d}{dt}I_{qm+1} &\approx ((1-q)N - I_{qm+1})\beta I_{qm+1} \end{aligned} \tag{10}$$

The solutions were found numerically and plotted in Fig. 3 for increasing values of q. For these results, we used parameter

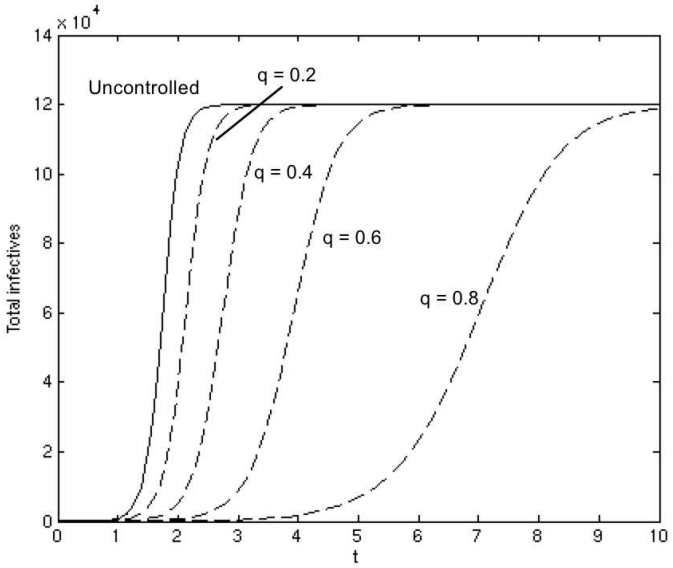

Fig. 3. Total epidemic size dependent on the fraction q of the population protected by router-based outbound rate throttling

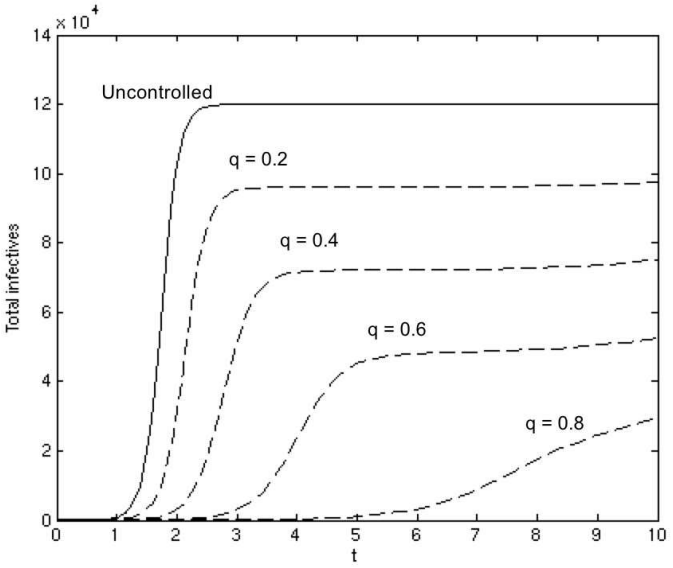

Fig. 4. Total epidemic size dependent on the fraction q of the population protected by router-based inbound and outbound rate throttling

values for the SQL Slammer worm: $\beta = 5.6 \times 10^{-5}$ and $N = 120,000$. In addition, we set $m = 10$ households and the initial conditions to $I_1(0) = \cdots = I_{qm}(0) = 0$, $I_{qm+1}(0) = 1$.

B. Router-based Inbound and Outbound Rate Throttling

Suppose that a fraction q of access routers are capable of throttling inbound and outbound worm traffic simultaneously, while the remaining routers are not controlled. Thus, protected households will be effectively quarantined. For simplicity, we again assume qm is an integer and the capable routers are numbered 1 to qm. All households are equal size with N/m hosts. Households 1 to qm can be quarantined with an intra-household infection rate β but inter-household infection rate $\alpha\beta$ in both inbound and outbound directions. The remaining households $qm + 1$ to m have intra-household and outbound inter-household infection rates β. These households may be viewed as a single large household of size $(1 - q)N$.

Recall that inbound rate throttling requires a detection time t_1 before it can be initiated. From time 0 to t_1, the outbreak is governed by (9), the same as outbound-only throttling. Here we assume t_1 is the time for the uncontrolled outbreak to reach 25 percent saturation, which is $t_1 = 1.5$ according to (3). From time t_1 onward, the epidemic with simultaneous inbound and outbound throttling is governed by the system of differential equations:

$$\begin{aligned} \frac{d}{dt}I_j &= (\frac{N}{m} - I_j)(\beta I_j + \sum_{i \neq j} \alpha\beta I_i) \quad , \; j = 1, \ldots, qm \\ \frac{d}{dt}I_{qm+1} &= ((1-q)N - I_{qm+1})(\beta I_{qm+1} + \sum_{i=1}^{qm} \alpha\beta I_i) \end{aligned} \tag{11}$$

If the attenuation factor is very close to 0, then (11) becomes

$$\begin{aligned} \frac{d}{dt}I_j &\approx (\frac{N}{m} - I_j)\beta I_j \qquad , j = 1, \ldots, qm \\ \frac{d}{dt}I_{qm+1} &\approx ((1-q)N - I_{qm+1})\beta I_{qm+1} \end{aligned} \tag{12}$$

It is evident that the outbreaks in each household will progress independently of each other, which is the goal of this strategy. After $t = t_1$, the number of infectives in each household will be

$$\begin{aligned} I_j(t) &\approx \frac{I_j(t_1)N}{I_j(t_1)m + (N - I_j(t_1)m)e^{-\beta N(t=t_1)/m}} \\ &\qquad , \; j = 1, \ldots, qm \\ I_{qm+1}(t) &\approx \frac{I_{qm+1}(t_1)(1-q)N}{I_{qm+1}(t_1) + ((1-q)N - I_{qm+1}(t_1))e^{-\beta(1-q)N(t-t_1)}} \end{aligned} \tag{13}$$

and the total epidemic size will be

$$I(t) = I_1(t) + \cdots + I_{qm+1}(t) \tag{14}$$

We wish to compare the effectiveness of throttling both inbound and outbound directions with throttling only outbound traffic. Numerical solutions of (14) are plotted in Fig. 4 using the same parameter values as before. The results show that the epidemic spreads quickly in the external Internet but the outbreaks within the rate throttled subnetworks are slowed down substantially due to their isolation from each other. As expected, the rate throttling prevents "mass action" mixing of many infectives and susceptibles. Clearly, in comparing Fig. 4 with Fig. 3, throttling inbound and outbound traffic appears to be much more effective than outbound-only rate throttling, and the difference is more dramatic with the amount of coverage q.

VI. Conclusions

In this paper, we have attempted to demonstrate the usefulness of the community of households epidemic model. The model is flexible enough to account for different rate control strategies without being overly complicated. Unfortunately, only simple special cases of the model can be analyzed in closed form, but the general model can be solved numerically.

We found that rate throttling, implemented in the network or hosts, can be effective in slowing down a worm epidemic given sufficient coverage of hosts. Moreover, throttling both outbound and inbound traffic can be much more effective than outbound-only rate throttling. Throttling both directions is effective because it severely limits the infectious contacts between infected hosts and potential targets. More generally, we can conclude that complete quarantining will be the most effective strategy because mixing between households is totally stopped.

We believe the community of households model can account for the effects of network congestion that helps to slow down the late stages of an epidemic. This is a problem for future investigation.

References

[1] D. Moore, V. Paxson, S. Savage, C. Shannon, S. Staniford, and N. Weaver, "Inside the Slammer worm," *IEEE Sec. & Privacy*, vol. 1, pp. 33–39, July 2003.

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

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

[4] M. Williamson, "Throttling viruses: restricting propagation to defeat malicious mobile code," in *18th Annual Comp. Sec. Appl. Conf.*, (Las Vegas, NV), Dec. 9–13, 2002.

[5] P. Porras, L. Briesemeister, K. Skinner, K. Levitt, J. Rowe, and Y.-C. A. Ting, "A hybrid quarantine defense," in *ACM Workshop on Rapid Malcode (WORM 2004)*, (Wash. DC), pp. 73–82, 2004.

[6] G. Ganger, G. Economou, and S. Bielski, "Self-securing network interfaces," Tech. Rep. CMU-CS-02-144, Carnegie Mellon U., Aug. 2002.

[7] S. Chen and Y. Tang, "Slowing Down Internet Worms," in *Proceedings of the 24th International Conference on Distributed Computing Systems (ICDCS'04)*, (Tokyo, Japan), pp. 312–319, Mar. 2004.

[8] S. Chen and S. Ranka, "An Internet-worm early warning system," in *IEEE Globecom 2004*, (Dallas, TX), pp. 2261–2265, Nov. 29 – Dec. 3, 2004.

[9] S. E. Schechter, J. Jung, and A. W. Berger, "Fast Detection of Scanning Worm Infections," in *Seventh International Symposium on Recent Advances in Intrusion Detection (RAID)*, (Sophia Antipolis, French Riviera, France), pp. 59–81, Sept. 15–17, 2004.

[10] V. Berk, G. Bakos, and R. Morris, "Designing a framework for active worm detection on global networks," in *First IEEE Int. Workshop on Info. Assurance (IWIAS 2003)*, pp. 13–23, Mar. 24, 2003.

[11] S. Singh, C. Estan, G. Varghese, and S. Savage, "The EarlyBird System for Real-time Detection of Unknown Worms," Tech. Rep. CS2003-0761, University of California at San Diego, Aug. 2003.

[12] M. V. Martin, J.-M. Robert, and P. van Oorschot, "A Monitoring System for Detecting Repeated Packets with Applications to Computer Worms," Tech. Rep. TR-04-02, School of Computer Science, Carleton University, Ottawa, Canada, June 2004.

[13] X. Chen and J. Heidemann, "Detecting Early Worm Propagation through Packet Matching," Tech. Rep. ISI-TR-2004-585, USC Information Sciences Institute, Marina del Rey, CA, Feb. 2004.

[14] G. Gu, M. Sharif, X. Qin, D. Dagon, W. Lee, and G. Riley, "Worm detection, early warning and response based on local victim information," in *20th Annual Comp. Sec. Applic. Conf.*, pp. 136–145, Dec. 6–10, 2004.

[15] T. Toth and C. Kruegel, "Connection-history based anomaly detection," in *2002 IEEE Workshop on Info. Assur. and Sec.*, (West Point, NY), June 17–19, 2002.

[16] D. Moore, C. Shannon, G. Voelker, and S. Savage, "Internet quarantine: requirements for containing self-propagating code," in *IEEE Infocom 2003*, (San Francisco, CA), pp. 1901–1910, 2003.

[17] C. Zou, W. Gong, and D. Towsley, "Worm propagation modeling and analysis under dynamic quarantine defense," in *ACM Workshop on Rapid Malcode (WORM 2003)*, (Wash. DC), pp. 51–60, Oct. 27, 2003.

[18] C. Wong, C. Wang, D. Song, S. Bielski, and G. R. Ganger, "Dynamic Quarantine of Internet Worms," in *Proceedings of the International Conference on Dependable Systems and Networks (DSN-2004)*, (Florence, Italy), pp. 73–82, June 28 – July 1, 2004.

[19] D. Daley and J. Gani, *Epidemic Modeling: An Introduction*. Cambridge, UK: Cambridge University Press, 1999.

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

[21] D. Moore, C. Shannon, and J. Brown, "Code-Red: a case study on the spread and victims of an Internet worm," in *Proceedings of the 2nd ACM SIGCOMM Internet Measurement Workshop (IMW)*, (Marseille, France), pp. 273–284, Nov. 6–8, 2002.

[22] M. Liljenstam, D. M. Nicol, V. H. Berk, and R. S. Gray, "Simulating realistic network worm traffic for worm warning system design and testing," in *2003 ACM Workshop on Rapid Malcode (WORM)*, (Wash. DC), pp. 24–33, Oct. 2003.

[23] C. C. Zou, W. Gong, and D. Towsley, "Code red worm propagation modeling and analysis," in *Proceedings of the 9th ACM conference on Computer and Communications Security (CCS '02)*, (Washington, DC), pp. 138–147, Nov. 18–22, 2002.

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

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

[26] D. Daley and J. Gani, "A deterministic general epidemic model in a stratified population," in *Probability, Statistics and Optimization - a Tribute to Peter Whittle* (F. P. Kelly, ed.), Chichester: Wiley, 1994.

[27] F. Ball, D. Mollison, and G. Scalia-Tomba, "Epidemics with two levels of mixing," *Annals of Applied Prob.*, vol. 7, pp. 46–89, 1997.

[28] F. Ball and O. Lyne, "Stochastic multitype SIR epidemics among a population partitioned into households," *Adv. Appl. Prob.*, vol. 33, pp. 99–123, 2001.

[29] F. Ball, "Stochastic multitype epidemics in a community of households: estimation of threshold parameter R_* and secure vaccination coverage," *Biometrika*, vol. 91, pp. 345–362, 2004.

[30] N. Becker and K. Dietz, "The effect of household distribution on transmission and control of highly infectious diseases," *Math. Biosci.*, vol. 127, pp. 207–219, 1995.

[31] N. Becker and R. Hall, "Immunization levels for preventing epidemics in a community of households made up of individuals of various types," *Math. Biosci.*, vol. 132, pp. 205–216, 1996.

[32] N. Becker and D. Starczak, "Optimal vaccination strategies for a community of households," *Math. Biosci.*, vol. 139, pp. 117–132, 1997.

[33] N. Becker and S. Utev, "The effect of community structure on the immunity coverage required to prevent epidemics," *Math. Biosci.*, vol. 147, pp. 23–39, 1998.

[34] T. Britton, "Epidemics in heterogeneous communities: estimation of R_0 and secure vaccination coverage," *J. Royal Statist. Soc. B*, vol. 63, pp. 705–715, 2001.

[35] M. Liljenstam, Y. Yuan, B. J. Premore, and D. Nicol, "A mixed abstraction level simulation model of large-scale Internet worm infestations," in *10th IEEE/ACM Symp. on Modeling, Analysis, and Simulation of Comp. Telecom. Sys. (MASCOTS 2002)*, (Fort Worth, TX), pp. 109–116, Oct. 11–16, 2002.

[36] M. Liljenstam and D. Nicol, "Comparing passive and active worm defenses," in *1st Int. Conf. on Quantitative Eval. of Sys. (QEST 2004)*, (Enschede, Netherlands), pp. 18–27, Sept. 27–30, 2004.

[37] A. Wagner, T. Dübendorfer, B. Plattner, and R. Hiestand, "Experiences with worm propagation simulations," in *ACM Workshop on Rapid Malcode (WORM 2003)*, (Wash. DC), pp. 34–41, Oct. 27, 2003.