

The Detection of RCS Worm Epidemics*

Kurt Rohloff
BBN Technologies
10 Moulton St.
Cambridge, MA, 02138, USA
krohloff@bbn.com

Tamer Başar
Coordinated Science Laboratory
University of Illinois at Urbana-Champaign
1308 West Main St.
Urbana, IL, 61801, USA
tbasar@control.csl.uiuc.edu

ABSTRACT

This paper discusses the problem of automatically detecting the existence of Random Constant Scanning (RCS) worm epidemics on the Internet by observing packet traffic in a local network. The propagation of the RCS worm is modelled as a simple epidemic. An optimal hypothesis-testing approach is presented to detect simple epidemics under idealized conditions based on the cumulative sums of log-likelihood ratios. It is shown that there are limitations on the ability of this optimal method to detect several important subclasses of RCS worm epidemics even under idealized conditions.

Categories and Subject Descriptors

G.3 [Probability and Statistics]: [stochastic processes, time series analysis]

General Terms

Experimentation, Security, Theory

1. INTRODUCTION

A potentially virulent type of worm, known as a Random Constant Scanning (RCS) worm, usually targets hosts running a particular operating system with a specific vulnerability that make the host susceptible to infection. A computer host infected by an RCS worm continually and randomly scans Internet addresses, and for every address selected for scanning, the infected host attempts to make a new connection with the other host at the selected address. Once a connection is established, the infectious host attempts to send infectious packets to the other host. If the other host is susceptible to the infection, then that host becomes infected and the scanning and infection process continues at both hosts.

*This research was performed while the first author was at the Coordinated Science Laboratory in Urbana, Illinois. This research was supported by the NSF grant CCR 00-85917 ITR.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

WORM'05, November 11, 2005, Fairfax, Virginia, USA.
Copyright 2005 ACM 1-59593-229-1/05/0011 ...\$5.00.

Due to the branching nature of scanning worm epidemics, these infections have the potential to propagate very fast over the Internet. For instance, during the CodeRed1v2 Internet worm attack of 2001 over 359,000 computers were infected in under 14 hours [8]. During the more aggressive Slammer Internet worm attack of 2003 more than 90 percent of 75,000 vulnerable computers were infected in less than 10 minutes [7]. Although it might be feasible for a human to be able to respond to an attack on the scale of relatively slow worms in order to protect a local network, human driven responses are infeasible for local network protections on global scale, especially when vigilance is required twenty-four hours a day, seven days a week, fifty-two weeks a year. This realization motivates the need for automated worm attack detection methods with response times and reliability significantly better than what could be provided by a human network administrator. There have been several approaches to the automated anomaly-based detection of RCS Internet worms, most notably in [16] where Kalman filtering methods are used in conjunction with observations based on ingress and egress filters on routers for local networks.

This paper discusses the abilities of anomaly-based RCS worm epidemic detection methods based on the observation of undesired scanning behavior due to the propagation of an epidemic. In particular, this paper discusses in a mathematical rigorous fashion why detection techniques that are only based on passively monitoring local IP addresses cannot quickly detect worm epidemics with a low false alarm rate even under idealized conditions by observing scanning behavior on unused local network addresses. There is an inherent level of “background radiation” of undesired scanning that occurs even in unused addresses in local networks connected to the Internet [9] which can hinder the ability of anomaly-based detection methods to diagnose the existence of a worm epidemic in its early stages of growth.

The worm epidemic detection methods discussed in this paper are developed from the field of detection and estimation theory which has been classically applied to the signal processing of radar systems and fault detection in manufacturing systems. A brief introduction to this field is given in this paper, but a more in-depth review of this material can be found in [10]. Of particular concern to the material in this paper is the notion of sequential change detection which is discussed more deeply in [2, 14]. In order to simplify the worm propagation models used in this paper it is assumed that when a host attempts to scan and propagate a worm epidemic there is a uniform distribution of any

one Internet address of being scanned, but this restriction can be easily removed. A relatively weak defender model is assumed in this paper in order to focus on the optimal detection policies outlined below. Note that scan detectors which consider other factors such as failed connection attempts [15] or other anomalies may lead to a much better and faster detection. Other relevant work on sequential hypothesis testing for scanning detection includes [6, 12].

Section 2 discusses models for the traffic observed on a local network during an RCS worm epidemic that incorporates both scanning due to the worm and background noise. Section 3 provides an introduction to the field of hypothesis testing and presents an optimal sequential hypothesis-testing method for the detection of scanning worm epidemics under idealized conditions. Section 4 discusses situations where even this optimal epidemic detection method under idealized conditions may not be able to achieve desired performance levels. The paper concludes with a brief discussion of the results and possible areas for future research in Section 5.

2. THE MODELING OF INTERNET WORM PROPAGATION

Of critical importance to the RCS worm epidemic detection problem discussed in this paper is the need for an effective model of the scanning behavior due to such an epidemic observed on a local network. Furthermore, in order to model the scanning behavior due to an RCS worm epidemic it is necessary to have an effective model of the worm's global propagation behavior. From data collected from previous RCS worm attacks such as CodeRed and Slammer, it has been found that simple epidemic models can effectively capture the behavior of worm propagation in a population susceptible to infection [7, 8, 13]. The simple epidemic model is one of the most basic epidemiological models from the fields of public health and epidemiology and has a vast literature dedicated to it [1, 3, 4, 5]. The simple epidemic model for the propagation of an RCS worm is now presented in Subsection 2.1. The simple epidemic model is then adapted to model the scanning behavior due to an RCS worm observed on a local network in Subsection 2.2.

2.1 The Simple Epidemic Model

For the simple epidemic model of a worm's propagation, it is assumed that in a given network (such as the Internet) there are n unique host addresses. Of these addresses, $n_s \leq n$ hosts could potentially become infected by the worm. The set of potential hosts is split into the *infected* and *susceptible* subpopulations, where $i(t)$ is the number of hosts which are infected by the worm at time t , and $s(t)$ is the number of hosts which could become infected, but are not at time t . Note that due to the random scanning propagation behavior of RCS worms, $s(t)$ and $i(t)$ are inherently random variables that are more accurately modelled as density-dependent Markov jump processes [11]. As discussed in [1], it is difficult if not impossible to find a closed form expression for the covariance of a density-dependent Markov jump process. However, as discussed in [11], for most situations when modeling the scanning behavior of an RCS worm epidemic observed on a local network, it is generally reasonable to model $s(t)$ and $i(t)$ as deterministic variables. Therefore, a deterministic model for $s(t)$ and $i(t)$ is used herein.

At time 0, let $(s(0), i(0)) = (s_0, i_0)$ where $i_0 \geq 1$ is the initial infected population. It is assumed that an infected host scans for susceptible hosts at a constant rate $\beta > 0$ which is called the *infection parameter*. Due to the fast dynamics of the epidemics modeled in this paper and for general simplicity in the model, it is also assumed that infectious hosts are not removed from the general population so that for all $t \geq 0$, $i(t) + s(t) = n_s$. As discussed in [4], $\frac{\beta}{n}s(t)i(t)$ is the rate at which the $i(t)$ infected hosts propagate the epidemic and

$$\frac{di}{dt} = \frac{\beta}{n}s(t)i(t) = \frac{\beta}{n}(n_s - i(t))i(t).$$

With $i(0) = i_0$,

$$i(t) = \frac{i_0 n_s}{i_0 + (n_s - i_0)e^{-\beta \frac{n_s}{n} t}}.$$

2.2 Observed Local Scanning Model

The simple epidemic model of Subsection 2.1 is now used to develop a model for the traffic observed on a local network due to an RCS worm epidemic. When an infected host attempts to spread the worm, the addresses selected by the host for scanning is assumed to be selected with a uniform distribution as discussed above. Consequently, if $\{A_1, A_2, A_3, \dots\}$ represents the sequence of host addresses globally selected for scanning by all hosts infected by a worm during an epidemic propagation where A_i is the address selected for the i th scanning, then $\{A_1, A_2, A_3, \dots\}$ are independently, uniformly distributed random variables over the set of n addresses.

Suppose there are n_t unique unused addresses in the local network. Define $\{b_1, b_2, b_3, \dots\}$ to be another set of binary random variables such that $b_i = 1$ if A_i is in the local address space and $b_i = 0$ otherwise. Note that because $\{A_1, A_2, A_3, \dots\}$ are independently and uniformly distributed, $\{b_1, b_2, b_3, \dots\}$ are similarly independently and uniformly distributed such that for all i , $b_i = 1$ with probability $\frac{n_t}{n}$.

Define $\{d_1, d_2, d_3, \dots\}$ to be another set of random variables where d_1 is the index of the first probe in $\{b_1, b_2, b_3, \dots\}$ such that $b_{d_1} = 1$, d_2 is the index of the second probe in $\{b_1, b_2, b_3, \dots\}$ such that $b_{d_2} = 1$ and so on such that d_j is the index of the j th probe in $\{b_1, b_2, b_3, \dots\}$ such that $b_{d_j} = 1$. For example, if $\{b_1, b_2, b_3, \dots\} = \{0, 0, 1, 1, 1, 0, 1, 0, \dots\}$, then for this example, $\{d_1, d_2, d_3, \dots\} = \{3, 4, 5, 7, \dots\}$. By definition $\{d_1, d_2 - d_1, d_3 - d_2, \dots\}$ are random variables with an independent gamma distribution with parameter $\frac{n_t}{n}$. Because $\frac{n_t}{n}$ can generally be assumed to be very small, then $\{d_1, d_2 - d_1, d_3 - d_2, \dots\}$ can effectively be approximated as independent exponentially distributed random variables with parameter $\frac{n_t}{n}$. (Recall however that the gamma distribution is defined over non-negative integers while the exponential distribution is defined over the non-negative reals.) Therefore, values for $\{d_1, d_2 - d_1, d_3 - d_2, \dots\}$ can be found in a straightforward manner and can be used to compute $\{d_1, d_2, d_3, \dots\}$.

Now define $\{t_1^w, t_2^w, \dots\}$ to be the times at which infected hosts scan addresses in the local network. Because $i(t) = \frac{i_0 n_s}{i_0 + (n_s - i_0)e^{-\beta \frac{n_s}{n} t}}$, the infection density changes with time, so the distribution of the random variables $\{t_1^w, t_2^w, \dots\}$ most likely does not follow immediately from a well-known distribution such as the uniform or exponential distributions.

However, with the assumption of the simple epidemic propagation model discussed above, if given a scan index $d_i \in \{d_1, d_2, d_3, \dots\}$, then an expression for t_i^w , the time of the i th scan corresponding to d_i can be computed. To demonstrate this, define $p^w(t)$ to be the number of scans which have occurred globally due to a worm epidemic at time t and define $i(p^w)$ to be the number of globally infected hosts when p^w scans have occurred globally. Note that $i(p^w(t)) = i(t)$. Recall that $\frac{dp^w}{dt} = \beta i$ and $\frac{di}{dt} = \frac{\beta}{n}(n_s - i)i$. Therefore

$$\frac{di}{dp^w} = \frac{(n_s - i)}{n},$$

and

$$i(p^w) = n_s - (n_s - i_0)e^{-\frac{p^w}{n}}.$$

Hence,

$$\frac{dt}{dp^w} = \frac{1}{\beta \left(n_s - (n_s - i_0)e^{-\frac{p^w}{n}} \right)}.$$

This can therefore be used to show that

$$t(p^w) = \frac{n}{\beta n_s} \left[\frac{p^w}{n} + \ln \left(\frac{\beta \left(n_s - (n_s - i_0)e^{-\frac{p^w}{n}} \right)}{\beta i_0} \right) \right]$$

$$p^w(t) = \beta n_s t + n \ln \left[\frac{i_0 + (n_s - i_0)e^{-\beta \frac{n_s}{n} t}}{n_s} \right]. \quad (1)$$

Due to above, $\{t_1^w, t_2^w, t_3^w, \dots\} = \{t(d_1), t(d_2), t(d_3), \dots\}$. An example of a plot of scanning times $\{t_1^w, t_2^w, t_3^w, \dots\}$ for a Slammer outbreak can be seen in Figure 1 when $n_t = 15$.




Figure 1: Plot of Scanning Times on Local Network During a Simulated Slammer Outbreak.

As was discussed in [9], not all of the scanning behavior observed on a set of locally unassigned addresses in a local network is solely due to the active propagation of a worm epidemic. In [9] it is shown that the traffic on a local network is inherently noisy with a substantial volume of “background radiation”: undesired scanning behavior on a local network not due to the propagation of a worm. Therefore, if one is attempting to detect the existence of an RCS worm epidemic by solely observing the scanning traffic on a set of unused local addresses, the observation of scanning traffic

on an unused IP address is not an immediate indicator of the existence of an Internet worm.

For the sake of mathematical simplicity, in this paper the interarrival times between background radiation packets observed on a local network are assumed to be independently and exponentially distributed with parameter η . By visual inspection of the graphs in [9], approximately 5 TCP background radiation packets hit a given IP address per hour on the Internet on average, and about 0.2 radiation packets per IP address per hour for UDP packets. (These averages most likely vary with time, but they are most likely good “ballpark” figures.) A simulation of the background radiation scan arrival times $\{t_1^n, t_2^n, t_3^n, \dots\}$ observed on a local address space with $n_t = 15$ can be seen in Figure 2.

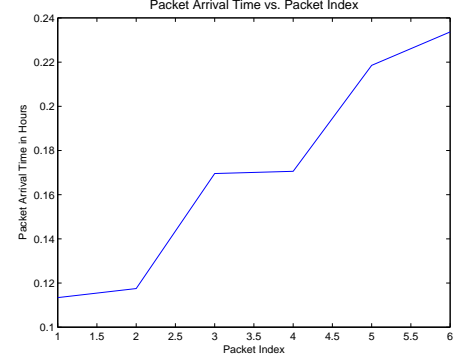


Figure 2: Plot of Scanning Times on Local Network Due to Background Radiation.

Now, with a mathematical model for $\{t_1^w, t_2^w, t_3^w, \dots\}$ for a given worm and $\{t_1^n, t_2^n, t_3^n, \dots\}$ for a given level of background radiation, a set $\{t_1, t_2, t_3, \dots\} = \{t_1^w, t_2^w, t_3^w, \dots\} \cup \{t_1^n, t_2^n, t_3^n, \dots\}$ can be defined where $\{t_1, t_2, t_3, \dots\}$ are the ordered times of all scans of the local address space such that t_i is the time of the arrival of the i th packet in unused address space of the local network. A simulation of this set of composed scan arrival times $\{t_1, t_2, t_3, \dots\}$ observed on a local address space with $n_t = 15$ can be seen in Figure 3.

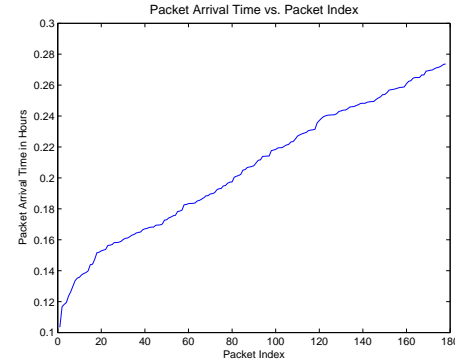


Figure 3: Plot of Scanning Times on Local Network Due to a Simulated Slammer Outbreak and Background Simulation.

3. HYPOTHESIS TESTING FOR WORM EPIDEMICS

An optimal Sequential Probability Ratio Test (SPRT) is now presented for the detection of a worm epidemic propagating over the Internet under idealized conditions. A brief introduction to SPRT is given here, but a more extensive overview is provided in Chapter 4 of [2].

In the framework of simple hypothesis testing it is assumed that there are two possible hypotheses to describe a system, H_0 and H_1 , corresponding to two (possibly non-stationary) probability distributions (P_0 and P_1 respectively) on observations that could be made of the system. Based on a set of observations $Y^n = \{o_1, o_2, o_3, \dots\}$ of system behavior, it should be decided whether hypothesis H_0 and H_1 better describes the system. In the context of the worm detection problem in this paper, given a set of observations of local scanning times $\{t_1, t_2, t_3, \dots\}$, the hypothesis testing problem is to decide whether or not a worm epidemic caused those observations. That is, hypothesis H_0 is said to hold if solely background radiation scanning caused the observations $\{t_1, t_2, t_3, \dots\}$, while hypothesis H_1 is said to hold if the scanning behavior due to a worm combined with background radiation scanning caused the observations $\{t_1, t_2, t_3, \dots\}$.

Suppose a set of n observations $Y^n = \{o_1, o_2, \dots, o_n\}$ of a system with two possible hypotheses H_0 and H_1 are given. Let $g(Y^n)$ be a hypothesis decision function for the observations Y^n such that if $g(Y^n) = 1$, then hypothesis H_1 is chosen to be the hypothesis of the current state of the system, and if $g(Y^n) = 0$, then hypothesis H_0 is chosen to be the hypothesis of the current state of the system. (In general it is possible that $g(Y^n)$ may at times be undefined for some input in order to indicate that neither hypothesis is chosen.) Define $p_0(Y^n)$ to be the probability that the observations Y^n are generated by the system when hypothesis H_0 holds. Similarly, define $p_1(Y^n)$ to be the probability that the observations Y^n are generated by the system when hypothesis H_1 holds. If

$$S_n = \ln \left(\frac{p_1(Y^n)}{p_0(Y^n)} \right),$$

a decision function $g(Y^n)$ can be defined such that

$$g(Y^n) = \begin{cases} 1 & \text{if } S_n \geq h \\ 0 & \text{if } S_n \leq -a \end{cases} \quad (2)$$

Because S_n is a logarithm of a probability ratio, this decision function $g(\cdot)$ is called a probability ratio test where h and a are boundaries (thresholds) on the hypothesis decisions such that $-a \leq h$. Note that it may be possible based on an observation Y^n to have a false alarm such that $g(Y^n) = 0$ when H_1 holds or $g(Y^n) = 1$ when H_0 holds. Note that if $-a \leq S_n \leq h$, then $g(Y^n)$ is undefined.

Sometimes it may not always be possible or desirable to have a fixed number of observations Y_n in order to decide between H_0 and H_1 . For instance, for the RCS worm detection problem, once a worm commences scanning addresses in a local network, the worm epidemic should be detected as soon as possible without regard to the number of observations made. To this end, *sequential analysis* has been developed as a theory for solving hypothesis testing problems when the sample size of the observations is not fixed a priori [2]. That is, observations of the behavior of a system might be made in an online manner and it should be

decided which hypothesis of the system state holds as soon as possible while maintaining a desirable performance level such as a sufficiently small false alarm rate.

Let for the above probability ratio test $g(\cdot)$, let there be a set of observations $Y = \{o_1, o_2, \dots\}$, where T is a random variable such that

$$T = \min\{n > 1 : g(Y^n) \in \{0, 1\}\}.$$

The random variable T represents the minimum number of observations of system behavior necessary in order to choose a hypothesis using the decision function g . For $Y = \{o_1, o_2, \dots\}$, define the *Sequential Probability Ratio Test* as follows:

$$g'(Y) = \begin{cases} 1 & \text{if } g(Y^T) = 1 \\ 0 & \text{if } g(Y^T) = 0 \end{cases} \quad (3)$$

The decision function $g'(\cdot)$ is a sequential decision function in that it selects a hypothesis using the minimum number of observations Y^T with respect to $g(\cdot)$.

The random variable T is called the *stopping time* because it is the time such that once T observations have been made, it is no longer necessary to collect more data to assign a hypothesis after T . The Average Run Length (ARL) is the mean number of observations $E_\theta(T)$ necessary for testing the hypotheses to obtain a given error rate if hypothesis H_θ holds. The ARL is an important performance measure of sequential analysis methods. Also related to the ARL is the notion of two different classes of false alarms. The variables α_0 and α_1 are defined to be the error rates for a sequential analysis methods where α_0 is the rate at which H_1 is thought to be true when H_0 holds, and α_1 is the rate at which H_1 is thought to be true when H_0 holds. Generally, as the threshold levels $-a$ and h are adjusted for the probability ratio test, the ARL decreases as the false alarm rates increase.

The SPRT is known to be an optimal sequential decision method [2]. Suppose due to the thresholds $-a, h$ used in $g(\cdot)$ to define $g'(\cdot)$ the ARL is $E_\theta(T)$ and false alarm rates are α_0 and α_1 . Let there be another sequential decision function $\tilde{g}'(\cdot)$ with ARL $E_\theta(\tilde{T})$ and false alarm rates $\tilde{\alpha}_0$ and $\tilde{\alpha}_1$ such that $\tilde{\alpha}_0 \leq \alpha_0$ and $\tilde{\alpha}_1 \leq \alpha_1$. Then, $E_0(\tilde{T}) \geq E_0(T)$ and $E_1(\tilde{T}) \geq E_1(T)$.

It is now shown how the above SPRT can be used for the detection of RCS worm epidemics on the Internet. Although it is not realistic, it is assumed for the sake of discussion in this section and the next that in formulating the SPRT test the parameters of the background radiation are known, along with the exact parameters of a possible worm if H_1 holds (namely β, i_0 and n_s .) This knowledge is used in the following section to demonstrate some fundamental performance limitations of sequential analysis methods for the detection RCS worms.

Consider the observed scanning interarrival time data seen in Figure 4. The top set of data is a graph of simulated TCP interarrival times due to both background radiation and the propagation of an RCS worm with parameters similar to CodeRed1v2. The radiation simulation uses parameters taken from the observations in [9]. The CodeRed1v2 propagation begins propagation at $t = 0$, and for the first several hours, there is little or no observed scanning due to the worm. However, during the last several hours, the worm is fully propagated and scanning the local host at its peak rate. As can be seen from the simulation, it may not be

immediately obvious that a worm is propagating by simply “looking” at scanning data.

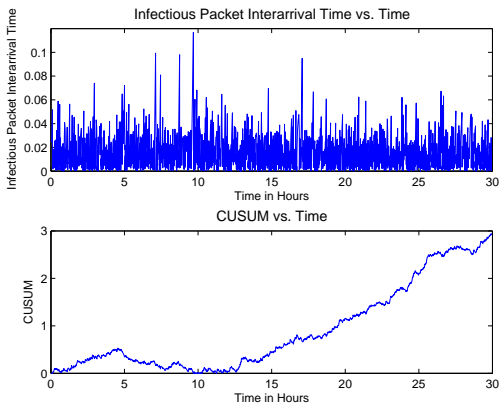


Figure 4: Plots of Observed Scanning Data and Computed CUSUM Data for a CodeRed Epidemic Using SPRT.

Suppose a set of scanning packet arrival times $\{t_1, t_2, \dots, t_n\}$ are observed on the local network. Let $p(t)$ be the function that maps a packet scanning time to its scanning index. If there is no worm on the Internet, let $p^0(t) = \eta t$ be the function which maps a packet arrival time to its scanning index, but if there is a worm on the Internet, with Equation 1 above,

$$p^1(t) = (\beta n_s + \eta)t + n \ln \left[\frac{i_0 + (n_s - i_0)e^{-\frac{\beta n_s t}{n}}}{n_s} \right]$$

is the function which maps a packet arrival time to its scanning index. Hence, for a given set of local scanning time observations $\{t_1, t_2, \dots, t_n\}$, two sets of packet indices can be computed for the two worm existence hypotheses H_0 and H_1 . Define $\{p_1^0, p_2^0, \dots, p_n^0\}$ to be the packet indices under the assumption of hypothesis H_0 using $p^0(\cdot)$ and let $\{p_1^1, p_2^1, \dots, p_n^1\}$ be the packet indices under the assumption of hypothesis H_1 using $p^1(\cdot)$. Recall that $\{p_1^0, p_2^0 - p_1^0, \dots, p_n^0 - p_{n-1}^0\}$ and $\{p_1^1, p_2^1 - p_1^1, \dots, p_n^1 - p_{n-1}^1\}$ should both be independently and exponentially distributed with parameter n_t/n . Given an exponential distribution with parameter γ , define $pdf(\kappa, \gamma)$ to be the probability density function of this exponential distribution at κ . Therefore, given $\{t_1, t_2, \dots, t_n\}$ and due to properties of the natural logarithm, it can be shown that

$$\begin{aligned} S_n &= S_{n-1} + \ln \frac{pdf(p_n^1 - p_{n-1}^1, n_t/n)}{pdf(p_n^0 - p_{n-1}^0, n_t/n)} \\ &= S_{n-1} + \ln \frac{pdf(p^1(t_n) - p^1(t_{n-1}), n_t/n)}{pdf(p^0(t_n) - p^0(t_{n-1}), n_t/n)} \end{aligned}$$

if $S_0 = 0$. This method of computing S_n is known as the *Cumulative Sum* (CUSUM) method and gives an efficient online method to compute $g^r(Y^n)$ in the worm detection scenario. The bottom graph in Figure 4 shows how the CUSUM for S_n changes with time as scanning data is observed.

4. LIMITATIONS TO DETECTION

Now that the SPRT method for sequential hypothesis testing has been introduced, we discuss in how this optimal detection method can be used to demonstrate fundamental limitations to the detection of RCS worm epidemics.

Let us consider a class of SPRT tests for the detection of RCS worms where h should be chosen so that the ARL is relatively small but not so small that the false positive rate α_1 is too high. Similarly, let us assume a lower threshold $-a$ that is very small so when the SPRT test runs, α_0 is very small and it is decided that no worm exists very slowly when H_0 holds. The intuition behind this assumption is that it is desirable to know if there exists a worm as soon as possible, but it is not necessary to know that there is no worm with great urgency.

As might be intuitive, there are strong connections between the threshold levels $-a, h$, the ARL and the false alarm rates of the SPRT. For instance, the ARL is monotonically increasing with respect to h . This can be seen in first plot of Figure 5 which shows the number of susceptible hosts which are infected when a worm epidemic is detected versus the threshold level h for a worm with parameters similar to CodeRed1v2. Similarly, the bottom plot of Figure 5 shows how the false positive rate α_1 depends on the threshold level h for a worm with parameters similar to CodeRed1v2.

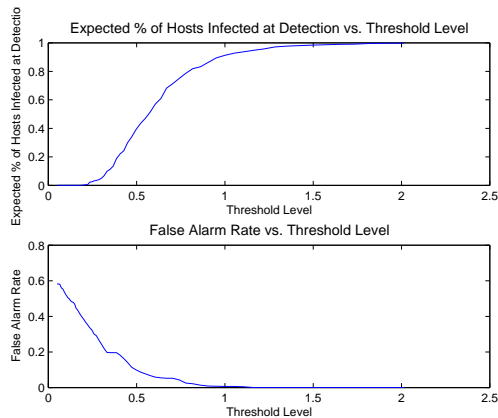


Figure 5: Plots of the Expected Percentage of Susceptible Hosts which are Infected when a Worm Epidemic is Detected vs. the Log-Likelihood Threshold and the Expected False Alarm Rate vs. the Log-Likelihood Threshold for the CodeRed SPRT.

Suppose a threshold level h should be chosen so that a desired ARL is achieved when a worm exists on the Internet. As indicated above, because α_1 is monotonically decreasing with respect to h , so by choosing a desired ARL, this places a fundamental limitation on the false alarm rate α_1 that can be achieved for the given ARL. With the data in Figure 5, it is plotted in Figure 6 how the SPRT false alarm rate α_1 depends on the desired ARL when H_1 holds. Figure 6 indicates that even with full knowledge of the worm and noise parameters, if an RCS worm epidemic should be detected during the early stages of propagation, most likely a reasonable false alarm rate (much less than 10%) cannot be achieved for a worm similar to CodeRed1v2, especially when there is no knowledge of the worm parameters.

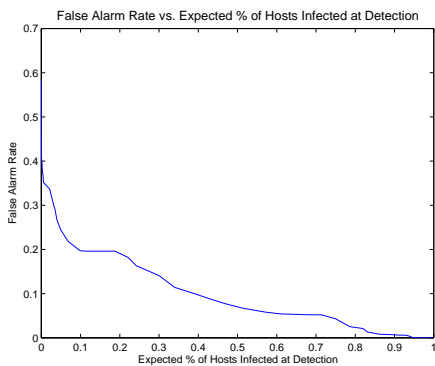


Figure 6: Plot of the Expected False Alarm Rate vs. the Expected Percentage of Susceptible Hosts which are Infected when a Worm Epidemic is Detected for a CodeRed1v2 worm SPRT.

However, as seen in Figure 7, the trade-off between expected detection time and false alarm rate is much more promising for an RCS worm with parameters similar to the Slammer worm. From Figure 7 it can be seen that under the ideal conditions of full knowledge of the worm’s parameters, a very small false alarm rate α_1 can be obtained as long as the percentage of hosts infected at detection is at least 14%.

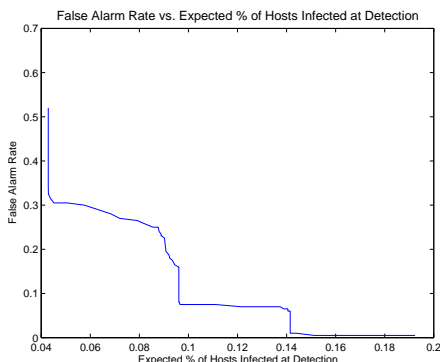


Figure 7: Plot of the Expected False Alarm Rate vs. the Expected Percentage of Susceptible Hosts which are Infected when a Worm Epidemic is Detected for a Slammer worm SPRT.

This result at first seems counter intuitive - that a fast worm like is easier to detect than a slower worm. However, when one considers that it is generally easier to detect a faster change in mean packet interarrival time, this should help to explain why Slammer worms can be detected with better false alarm rates.

5. DISCUSSION

This paper has discussed the anomaly-based RCS worm detection problem in the context of detection and estimation theory. A point-process model for the observations of scanning behavior due to a worm and background radiation has been proposed. An optimal SPRT worm detection method

has been proposed under the idealized condition of knowledge of a worm’s parameters. Fundamental limitations to the detection of RCS worms have been discussed based on simulations of RCS worms with the SPRT detection method. It has been shown that in some sense aggressive RCS worms like Slammer are generally easier to detect than slower RCS worms such as CodeRed1v2.

6. REFERENCES

- [1] H. Andersson and T. Britton. *Stochastic Epidemic Models and Their Statistical Analysis*. Number 151 in Lecture Notes in Statistics. Springer-Verlag, 2000.
- [2] M. Basseville and I. Nikiforov. *Detection of Abrupt Changes: Theory and Applications*. Prentice-Hall, New York, 1993.
- [3] F. Brauer and C. Castillo-Chávez. *Mathematical Models in Population Biology and Epidemiology*. Number 40 in Texts in Applied Mathematics. Springer-Verlag, New York, 2001.
- [4] D. Daley and J. Gani. *Epidemic Modelling: An Introduction*. Cambridge University Press, 1999.
- [5] H. Hethcote. The mathematics of infectious diseases. *SIAM Review*, 42(4):599–653, 2000.
- [6] J. Jung, V. Paxson, A. W. Berger, and H. Balakrishnan. Fast portscan detection using sequential hypothesis testing. In *Proc. of the IEEE Symposium on Security and Privacy*, 2004.
- [7] D. Moore, V. Paxson, S. Savage, C. Shannon, S. Staniford, and N. Weaver. Inside the slammer worm. *IEEE Security and Privacy*, 1(4):33–39, 2003.
- [8] D. Moore, C. Shannon, and J. Brown. Code-red: A case study on the spread and victims of an Internet worm. In *Proc. of the Internet Measurement Workshop (IMW)*, 2002.
- [9] R. Pang, V. Yegneswaran, P. Barford, V. Paxson, and L. Peterson. Characteristics of Internet background radiation. In *Proc. of the 4th ACM SIGCOMM Conference on Internet Measurement*, 2004.
- [10] H. Poor. *An Introduction to Signal Detection and Estimation*. Springer Texts in Electrical Engineering. Springer-Verlag, New York, 1994.
- [11] K. Rohloff and T. Başar. Stochastic behavior of random constant scanning worms. In *Proc. of 14th ICCCN*, 2005.
- [12] S. E. Schechter, J. Jung, and A. W. Berger. Fast detection of scanning worm infections. In *Proc. of The Seventh International Symposium on Recent Advances in Intrusion Detection (RAID)*, 2004.
- [13] S. Staniford, V. Paxson, and N. Weaver. How to Own the Internet in your spare time. In *Proc. of the 11th USENIX Security Symposium (Security '02)*, 2002.
- [14] A. Wald. *Sequential Analysis*. Dover, New York, 1947.
- [15] N. Weaver, S. Staniford, and V. Paxson. Very fast containment of scanning worms. In *Proc. of the 13th USENIX Security Symposium (Security '04)*, 2004.
- [16] C. Zou, L. Gao, W. Gong, and D. Towsley. Monitoring and early warning for Internet worms. In *Proc. of the 10th ACM conference on Computer and communications security*, pages 190–199. ACM Press, 2003.