HoneyStat: Local Worm Detection Using Honeypots

David Dagon, Xinzhou Qin, Guofei Gu, Wenke Lee
Julian Grizzard, John Levine, and Henry Owen
{dagon, xinzhou, guofei, wenke}@cc.gatech.edu,
{grizzard, levine, henry.owen}@ece.gatech.edu

Georgia Institute of Technology

Abstract. Worm detection systems have traditionally used global strategies and focused on scan rates. The noise associated with this approach requires statistical techniques and large data sets (e.g., 2^{20} monitored machines) to yield timely alerts and avoid false positives. Worm detection techniques for smaller local networks have not been fully explored.

We consider how local networks can provide early detection and compliment global monitoring strategies. We describe HoneyStat, which uses modified honeypots to generate a highly accurate alert stream with low false positive rates. Unlike traditional highly-interactive honeypots, HoneyStat nodes are script-driven, automated, and cover a large IP space.

The HoneyStat nodes generate three classes of alerts: memory alerts (based on buffer overflow detection and process management), disk write alerts (such as writes to registry keys and critical files) and network alerts. Data collection is automated, and once an alert is issued, a time segment of previous traffic to the node is analyzed. A logit analysis determines what previous network activity explains the current honeypot alert. The result can indicate whether an automated or worm attack is present.

We demonstrate HoneyStat's improvements over previous worm detection techniques. First, using trace files from worm attacks on small networks, we demonstrate how it detects zero day worms. Second, we show how it detects multi vector worms that use combinations of ports to attack. Third, the alerts from HoneyStat provide more information than traditional IDS alerts, such as binary signatures, attack vectors, and attack rates. We also use extensive (year long) trace files to show how the logit analysis produces very low false positive rates.

Keywords: Honeypots, Intrusion Detection, Alert Correlation, Worm Detection

1   Introduction

Worm detection strategies have traditionally relied on artifacts incidental to the worm infection. For example, many researchers measure incoming scan rates (often using darknets) to indirectly detect worm outbreaks, e.g., [ZGGT03]. But since these techniques measure noise as well as attacks, they often use costly algorithms to identify worms. For example, [ZGGT03] suggests using a Kalman filter [Kal60] to detect worm attacks. In [QDG$^+$], this approach was found to work with a large data set but proved inappropriate for smaller networks.

To improve detection time and decrease errors caused by noise, the solution so far has been to increase monitoring efforts, and gather more data. The intuition is that

with more data, statistical models perform better. Thus, researchers have suggested the creation of global monitoring centers [MSVS03], and collecting information from distributed sensors. These efforts are already yielding interesting results [YBJ04,Par04].

Although the need for global monitoring is obvious, the value this has for local networks is not entirely clear. For example, some local networks might have enough information to conclude a worm is active, based on additional information they are unwilling to share with other monitoring sites. Likewise, since global detection strategies require large amounts of sensor data, worm outbreaks may be detected only after local networks fall victim. Also, we see significant problems in gaining consensus among different networks, which frequently have competing and inconsistent policies regarding privacy, notification, and information sharing. Without doubt, aggregating information from distributed sensors makes good sense. However, our emphasis is on local networks and requires a complimentary approach. In addition to improving the *quantity* of monitoring data, researchers should work to improve the *quality* of the alert stream.

In this paper, we propose the use of honeypots to improve the accuracy of alerts generated for local intrusion detection systems. To motivate the discussion, we describe in Section 3 the worm infection cycle we observed in honeypots that led to the creation of HoneyStat. Since honeypots usually require labor-intensive management and review, we describe in Section 4 a deployment mechanism used to automate data collection. HoneyStat nodes collect three types of events: memory, disk write and network events. Section 4 describes these in detail, and discusses a way to compare and correlate intrusion events. Using logistic regression, we analyze previous network traffic to the honeypot to see what network traffic most explains the intrusion events. Intuitively, the logit analysis asks if there is a common set of network inputs that precede honeypot intrusions. Finding a pattern suggests the presence of an automated attack or worm.

To demonstrate HoneyStat's effectiveness, in Section 6, we describe our experience deploying HoneyStat nodes, and a retrospective analysis of network captures. We also use lengthy (year long) network trace files to analyze the false positive rate associated with the algorithm. The false positive rate is low, due to two influences: (a) the use of honeypots, which only produce alerts when there are successful attacks, and (b) the use of user-selected confidence intervals, which let one define a threshold for alerts. Finally, in Section 7, we analyze whether a local detection strategy with a low false positive rate (like HoneyStat) can make an effective worm detection tool. We consider the advantages this approach has for local networks.

2   Related Work

Honeypots A honeypot is a vulnerable network decoy used for several purposes: (a) distracting attackers, (b) gathering early warnings about new attack techniques, (c) facilitating in-depth analysis of an adversary's strategies [Spi03,Sko02]. By design, a honeypot should not receive any network traffic, nor will it run any legitimate production services. This greatly reduces the problem of false positives and false negatives often found in other types of IDS systems.

Traditionally, honeypots have been used to gather intelligence about how human attackers operate [Spi03]. The labor-intensive log review required of traditional honey-

pots makes them unsuitable for a real-time IDS. In our experience, data capture and log analysis time can require a 1:40 ratio, meaning that a single hour of activity can require a week to fully decipher [LLO$^+$03].

The closest work to our own is [LLO$^+$03], which uses honeypots in an intrusion detection system. We have had great success at the Georgia Institute of Technology utilizing a Honeynet as an IDS tool, and have identified a large number of compromised systems on campus, mostly the result of worm-type exploits. Other researchers have started to look at honeypot alert aggregation techniques [JX04], but presume a centralized honeypot farming facility. Our simplified alert model allows for distributed honeypots, but is more focused on defending local networks.

Researchers have also considered using virtual honeypots, particularly with honeyd [Pro03]. Initially used to help prevent OS fingerprinting, honeyd is a network daemon that exhibits the TCP/IP stack behavior of different operating systems. It has since been extended to emulate some services (e.g., NetBIOS). Conceptually, honeyd is a daemon written using libpcap and libdnet. To emulate a service, honeyd requires researchers to write programs that completely copy the service's network behavior. Assuming one can write enough modules to emulate all aspects of an OS, a single machine can delay worms by inducing needless connections.

Recently, honeyd was offered as a way to detect and disable worms [Pro03]. We believe this approach has promise, but must overcome a significant hurdle before it is used as an early warning IDS. It is not clear how a daemon emulating a network service can catch zero day worms. If one knows a worm's attack pattern, it is possible to write modules that will behave like a vulnerable service. But before this is known, catching zero day worms requires emulating even the *presumably unknown* bugs in a network service. We were unable to find solutions to these limitations, and so do not consider virtual networks as a means of providing an improved alert stream. Instead, we used full honeypots.

More closely related to our work, [Kre03] suggested automatic binary signature extraction using honeypots. This work used honeyd, flow reconstruction, and pattern detection to generate IDS rules. The honeycomb approach also has promise, but uses a very simple algorithm (longest common substring) to correlate payloads. This makes it difficult to identify polymorphic worms and worms that use multiple attack vectors. Honeycomb was well suited for its stated purpose, however: extracting string signatures for automated updates to a firewall.

Worm Detection Worm propagation and early detection have been active research topics in the security community. In worm propagation, researchers have proposed an epidemic model to study worm spreading, e.g., [Sta01,ZGT02,CGK03,WL03]. For early detection, researchers have proposed statistical models, e.g., Kalman Filter [ZGGT03], analyzing repeated outgoing connections [Wil02], ICMP messages collected at border routers to infer worm activity [BGB03] and victim counter-based detection algorithms [WVGK04]. All these approaches require a large deployment of sensors or a large monitoring IP space (e.g., 2^{20} IP addresses). Others suggest a "cyber Center for Disease Control" to coordinate data collection and analysis [SPN02]. Researchers have also proposed various data collection and monitoring architectures, e.g., "network telescopes" [Moo02b] and an "Internet Storm Center" [Ins].

Our objective is also to conduct early worm detection. However, considering the current diffi culty and challenges in large space monitoring system (e.g., conflicts in data sharing, privacy, and coordinated responses), our detection mechanism is based on local networks, in particular, local honeypots for worm detection. In our prior work [QDG$^+$] we analyzed the current worm early detection algorithms, i.e., [ZGGT03] and [WVGK04], and found instability and high false positives when applying these techniques to local monitoring networks.

Event Correlation Several techniques have been proposed for the alert/event correlation, e.g., pre-/post-condition-based pattern matching [NCR02,CM02,CLF03], chronicles formalism [MMDD02], clustering technique [DW01] and probabilistic-based correlation technique [VS01,GHH01]. All these techniques count on some prior knowledge of attack step relationships. Our approach is different in its need to detect *zero-day* worm attacks, and does not depend on prior knowledge of attack steps. Statistical alert correlation was presented in [QL03]. Our work is different in that our correlation analysis is based on variables collected over *short* observations. Time series-based analysis proposed in [QL03] is good for relatively long observation variables and requires a series of statistical tests in order to accurately correlate observations.

3   Worm Infection Cycles

If local networks do not have access to the volume of data used by global monitoring systems, what local resources can they use instead? Studying worm infections gives some insights, and identifi es what data can be collected for use in a local IDS.

Model of Infection A key assumption in our monitoring system is that the worm infection can be described in a systematic way. We fi rst note that worms may take three types of actions during an infection phase. The Blaster worm is instructive, but we do not limit our model to this example. Blaster consists of a series of modules designed to infect a host [LUR03].

Memory Events The infection process, illustrated in Figure 1(a), begins with a probe for a victim providing port 135 RPC services. The service is overflowed, and the victim spawns a shell listening on a port, usually 4444. (Later generations of the worm use different or even random ports.) This portion of the infection phase is characterized by memory events. No disk writes have taken place, and network activity cannot (yet) be characterized as abnormal, since the honeypot merely ACKs incoming packets. Still, a buffer overflow has taken place, and the infection has begun by corrupting a process.

Network Events The Blaster shell remains open for only one connection and closes after the infection is completed. The shell is used to instruct the victim to download (often via tftp) an "egg" program. The egg can be obtained from the attacker, or a third party (such as a free website, or other compromised hosts.) The time delay between the initial exploit and the download of the egg is usually small in Blaster, but this may not always be the case. Exploits that wait a long period to download the egg risk having the service restarted, canceled, or infected by competing worms (e.g., Nachi). Nonetheless, some delay may occur between the overflow and the "egg" transfer. All during this time, other harmless network traffi c may arrive. This portion of the infection phase is often characterized by network traffi c. Downloading the egg for example requires the

honeypot to initiate TCP (SYN) or UDP traffic. In some cases, however, the entire worm payload can be included in the initial attack packet [Sta01,Moo02a,ZGT02]. In such a case, network events may not be seen until much later.

Disk Events Once the Blaster egg is downloaded, it is written to a directory so it may be activated upon reboot. Some worms (e.g., Witty [LUR04]) do not store the payload to disk, but do have other destructive disk operations. Not every worm creates disk operations.

These general categories of events, although present in Blaster, do not limit our analysis to just the August 2003 DCOM worm. A local detection strategy must anticipate future worms lacking some of these events.

Improved Data Capture Traditional worm detection models deal with worm infection at either the start or end of the cycle shown in Figure 1(a). For example, models based on darknets consider only the rate and sometimes the origin of incoming scans, the traffic at the top of the diagram. The Destination Source Correlation (DSC) model [GSQ+04] also considers scans, but also tracks outgoing probes from the victim (traffic from the bottom of the diagram). The activity in the middle of the cycle (including memory and disk events) can be tracked.

Even if no buffer overflow is involved, as in the case of mail-based worms and LANMAN weak password guessing worms (e.g., pubstro worms), the infection still follows a general pattern: a small set of attack packets obtain initial results, and further network traffic follows, either from the egg deployment, or from subsequent scans.

Intrusion detection based only on incoming scan rates must address the potentially high rate of noise associated with darknets. As noted in Figure 1(a), every phase of the infection cycle may experience non-malicious network traffic. Statistical models that filter the noise (e.g., Kalman) require large data sets for input. It is no wonder, then, that scan-based worm detection algorithms have recently focused on distributed data collection.

4   HoneyStat Configuration and Deployment

The foregoing analysis of the worm infection cycle generally identified three classes of events that one might track in an IDS: memory, disk and network events. As noted above in Section 2, it is difficult to track all of these events in virtual honeypots or even in stateful firewalls. Networks focused on darknets, of course, have little chance of getting even complete network events, since they generally blackhole SYN packets, and never see the full TCP payload.

A complete system is needed to gather the worm cycle events and improve the data stream for an IDS. We therefore use a HoneyStat node, a minimal honeypot created in an emulator, and multihomed to cover a large address space. The deployment typically would not be interesting to attackers, because of its minimal resources (limited memory, limited drive size, etc.) Worms, however, are indiscriminating and use this configuration.

In practice, we can use VMware GSX Server as our honeypot platform. Currently, VMware GSX Server V3 can support up to 64 isolated virtual machines on a single hardware system [VMW04]. Mainstream operating systems (e.g. Windows, Linux,

FreeBSD) all support multihoming. For example, Windows NT allows up to 32 IP addresses per interface. So if we use a GSX server with 64 virtual machines running windows and each windows having 32 IP addresses, then a single GSX machine can have $64 * 32 = 2^{11}$ IP addresses.

In practice, we found nodes with as little as 32MB RAM and 770MB virtual drives were more than adequate for capturing worms. Since the emulators were idle for the vast majority of time, many instances could be started on a single machine. Although slow and unusable from a user perspective, these virtual honeypots were able to respond to worms before any timeouts occur.

The honeypots remain idle until a *HoneyStat event* occurs. We define three types of events, corresponding to the worm infection cycle discussed in Section 3.

1. `MemoryEvent`. A honeypot can be configured to run buffer overflow protection software, such as a StackGuard [Inc03], or similar process-based monitoring tools. Likewise, Windows logs can be monitored for process failures and crashes. Any alert from these tools constitutes a HoneyStat event. Because there are no users, we found that one can use very simple anomaly detection techniques that would otherwise trigger enormous false positive rates on live systems. Since there are no users, even simple techinques work well.
2. `NetworkEvents`. The honeypots are configured to generate no outgoing traffic. If a honeypot generates SYN or UDP traffic, we consider it an event.
3. `DiskEvents`. Within the limits of the host system, we can also monitor honeypot disk activities and trap writes to key file areas. For example, writes to systems logs are expected, while writes to `C:\WINNT\SYSTEM32` are clearly events. In practice, we found that kqueue [Lem01] monitoring of flat virtual disks was reasonably efficient. One has to enumerate all directories and files of interest, however.

Data recorded during a HoneyStat event includes: (a) The OS/patch level of the host. (b) The type of event (memory, net, disk), and relevant capture data. For memory events, this includes stack state or any core, for network events this is the outgoing packet, and for disk events this includes a delta of the file changes, up to a size limit. (c) A trace file of all prior network activity, within a bound t_p, discussed in Section 5.

Based on our analysis in Section 3, we believe this to be a complete set of features necessary to observe worm behavior. However, new worms and evasive technologies will require us to revisit this heuristic list of features. Additionally, if HoneyStat is given a larger mission (e.g., e-mail virus detection or trojan analysis instead of just worm detection), then more detailed features must be extracted from the honeypots.

Once events are recorded, they are forwarded to an analysis node. This may be on the same machine hosting the honeypots, or (more likely) a central server that performs logging and propagates the events to other interested nodes. Figure 1(b) shows a conceptual view of one possible HoneyStat deployment. In general, the analysis node has a secure channel connecting it with the HoneyStat servers. Its primary job is to correlate alert events, perform statistical analysis, and issue alerts.

HoneyStat events are placed in a work queue and analyzed by worker threads. We can prioritize certain types of events, such as when outgoing connections attempt to reach many different IPs (i.e., it looks like a worm scanning.) These events are obviously
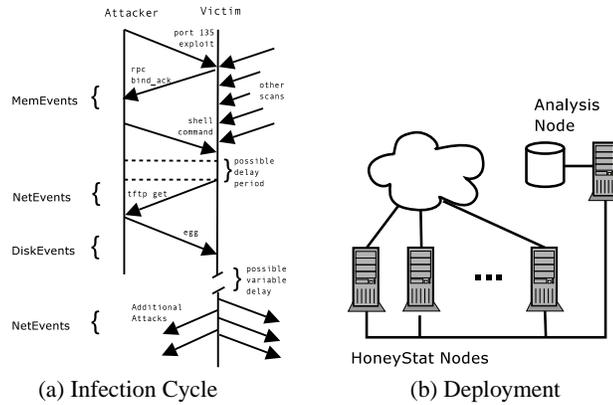
(a) Infection Cycle   (b) Deployment

Fig. 1. a) A time line of a Blaster worm attack. Because of modular worm architectures, victims are first overflowed with a simple RPC exploit, and instructed to obtain a separate worm "egg", which contains the full worm. The network activity between the initial overflow and download of the "egg" constitutes a single observation. Multiple observations allow one to filter out other scans arriving at the same time. b) HoneyStat nodes interact with malware on the Internet. Alerts are forwarded through a secure channel to an analysis node for correlation.
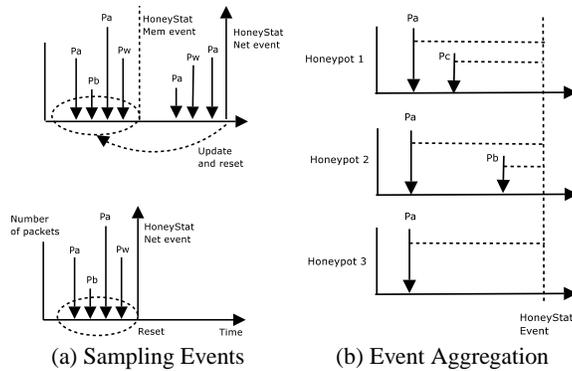

(a) Sampling Events   (b) Event Aggregation

Fig. 2. a) In the top diagram, a HoneyStat `MemEvent` occurs, and the honeypot is allowed to continue, in hopes of capturing an egg or payload. (The event is immediately analyzed, however, without delay.) If a subsequent `NetEvent` occurs, we update the previous `MemoryEvent` event and reset. In the bottom diagram, we see a `NetEvent` without any prior `MemoryEvent`, indicating our host-based IDS did not spot any anomaly. We immediately reset, and analyze the circled traffic segment. b) Aggregating multiple honeypot events can help spot a pattern. In all three diagrams, activity P_a is observed, even though others (P_b, P_c) appear closer in time to the event.

more important than others, since they suggest an automated spread mechanism instead of a single connection back to the attacker's IP. This idea borrows from the work of researchers who observed the importance of tracking distinct destination IPs in worm detection systems [JPBB04,WVGK04]. Future work will explore queue processing and other prioritizations.

Several actions are taken when a HoneyStat event is analyzed.

1. First, we check if the event corresponds to a honeypot that has already been recorded as "awake" or active. If the event is a continuation of an ongoing infection, we simply annotate the previous event with the current event type. For example, if we first witness a `MemoryEvent`, and then see a `DiskEvent` for the same honeypot, we update the `MemoryEvent` to include additional information, such as the `DiskEvent` and all subsequent network work activity. The intuition here is that `MemoryEvents` are usually followed by something interesting, and it is worth keeping the honeypot active to track this.

2. Second, if the event involved `NetworkEvents` (e.g., either downloading an egg or initiating outgoing scans), the honeypot reporting the event is reset. The idea here is two-fold. In keeping with the principle of honeypot Data Control [LLO$^+$03], we need to prevent the node from attacking other machines. Also, once network activity is initiated, we have enough attack behavior recorded to infer that the worm is now infective. If only `DiskEvents` or `MemoryEvents` are observed, the node is not reset.

   Since the honeypot is deployed using an emulator, resets are fast in practice. One merely has to kill and restart the emulator, using in round-robin style a fresh copy of the virtual disk. The disk image is kept in a suspended state, and no reboot of the guest OS is required. The reset delay is slight, often seconds or a minute, and always in practice completes before TCP timeouts occur. The effect of startup time on detection is considered in the analysis in Section 6.

3. Third, the analysis node examines basic properties of the event, and determines whether it needs to redeploy other honeypots to match the affected OS. The intuition here is that HoneyStat nodes are often deployed to cover a variety of operating systems: Linux, Windows, and with different patch levels. If one of the systems falls victim to a worm, it makes sense to redeploy most of the other nodes to run the vulnerable OS. This improves the probability that an array of HoneyStat nodes will capture similar events. Again, the delay this causes for detection is discussed in Section 7.

4. Finally, the HoneyStat event is correlated with other observed events. If a pattern emerges, this can indicate the presence of a worm or other automated attacks. Any reasonable correlation of events can be done. In the next section, we present a candidate analysis based on logistic regression.

As an example, in Figure 2(b), we see three different honeypots generating events. Prior input to the honeypots includes a variety of sources. For simplicity, the example in Figure 2(b) merely has three different active ports, P_a, P_b, P_c. Intuitively, we can use the time difference between the honeypot event and the individual port activity to infer what caused the honeypot to become active. But if all these events are from the

same worm, then one would expect to see the same inputs to all three honeypots. In this case, only P_a is common to all three. A logistic regression presents a more flexible way of discovering the intersection of all inputs and provides a better explanation why a honeypot has become active.

4.1 Honeypot Evasion

As more honeypots are used in IDS settings, attackers may attempt to evade by having worms detect and avoid honeypot traps. Honeypot researchers have observed that a few assembly instructions behave slightly differently in various (often incomplete) emulators, and that emulated hardware may have predictable signatures (e.g., BIOS Strings, MAC address ranges for network cards) [Cor04,Sei02]. One can prevent trivial honeypot detection by patching emulated VMs [Kor04], and removing any obvious indicators like registry keys.

Hand crafted assembly instructions designed to detect VMWare present a different problem. Since Intel chips don't support multiple zero ring contexts, some instructions will elicit a VM monitor error, allowing attackers to evade the honeypot trap. This can be countered by filtering incoming traffic to identify the limited instruction set designed to detect VMWare. Failing this (e.g., if the emulator detection code is polymorphic), one can always just treat the crashed emulator as a HoneyStat memory event. This yields a more limited alert (e.g., you miss disk events), but allows HoneyStat to correlate what caused the error.

Attackers might also attempt to make machine observations, e.g., the time needed to perform lengthy calculations. This potentially benefits defenders, since worms may have a slower propagation rate, allowing for human intervention and earlier detection. Ultimately, we believe the honeypot evasion problem may devolve into a classic cat-and-mouse game, not unlike virus detection. In this case, however, the tables are turned, and it is the attacker who must perform reliable detection in a changing environment.

5 Logistic Analysis of HoneyStat Events

Our key objective is the detection of *zero-day* worms, or those without a known signature. Without the ability to perform pattern matching, our task is analogous to *anomaly detection*. We therefore use a statistical analysis of the events to identify worm behavior. Statistical techniques, e.g., [MHL94,AFV95,PN97,QVWW98], have been widely applied in anomaly detection, . In our prior work, we applied time series-based statistical analysis to alert correlation [QL03].

Our preference was for a technique that can effectively correlate variables collected in a *short* observation window with a short computation time. Time series-based analysis is good for a relatively long observation and requires a series of statistical tests in order to accurately correlate variables. It is also often not suitable for real-time analysis because of its computationally intensive nature. Therefore, in this work, we instead apply logistic analysis [HL00] to analyze port correlation.

Logistic regression is a non-linear transformation of the traditional linear regression model. Instead of correlating two continuous variables, logistic regression considers

(in the simplest case) a dichotomous variable and continuous variables. That is, the dependent variable is a boolean "dummy" variable coded as 0 or 1, which corresponds to a state or category we wish to explain. In our case, we treat the honeypot event as a dichotomous variable, i.e., the honeypot is either awake (1) or quiescent (0). Logit analysis then seeks to explain what continuous variables explain the changes in the honeypot state, from asleep to awake.

We settled on using a logit analysis only after considering other, more restrictive analysis techniques. A simple linear regression, for example, would compare continuous-to-continuous variables. In the case of honeypots, this would require either measuring rates of outgoing packets, or identifying some other continuous measurement in the memory, network and disk events. Since it only takes one packet to be infected or cause an infection to spread, a simple linear regression approach would not clearly identify "sleeper worms" (a false negative scenario) and worms on busy networks (a false positive potential). Additionally, measuring outgoing packet rates would also include a significant amount of noise, since honeypots routinely complete TCP handshakes for the services they offer (e.g., normal, non-harmful webservice, mail service, ftp connections without successful login, etc.). Using continuous variables based on outgoing rates may only be slightly better than using incoming scan rates.

The basic form of the model expresses a binary expectation of the honeypot state, $E(Y)$ (asleep or awake) for k events, as seen in Eq. (1).

$$E(Y) = \frac{1}{1 + e^{-Z}}, \quad \text{where } Z = \beta_0 + \epsilon + \sum_{j=1}^{k} \sum_{i=1}^{n_j} (\beta_{i,j} X_{i,j}) \tag{1}$$

In Eq. (1), j is a counter for each individual honeypot event, and i is a counter for each individual port traffic observation for a specific honeypot. Each $\beta_{i,j}$ is the regression coefficient corresponding to the $X_{i,j}$ variable, a continuous variable representing each individual port observation. We have one error term ϵ and one constant β_0 for the equation. To set values of $X_{i,j}$, we use the inverse of time between an event and the port activity. Thus, if a `MemoryEvent` (or honeypot event j) occurs at time t, and just prior to this, port i on that same honeypot experienced traffic at time $t - \delta_t$, the variable $X_{i,j}$ would represent the port in the equation, and would have the value of $\frac{1}{\delta_t}$. This biases towards network traffic closer in time to the event, consistent with our infection model discussed in Section 3.

An example shows how honeypot events are aggregated. Suppose one honeypot event is observed, with activity to ports $\{P_1, P_2, \ldots, P_n\}$. We calculate the inverse time difference between the port activity and the honeypot event, and store the values for $X_{1,1}, X_{2,1}, \ldots X_{n,1}$ in a table that solves for Y. Suppose then a second event is recorded, in the same class as the first. We add the second event's values of $X_{1,2}, X_{2,2}, \ldots, X_{n,2}$ to the equation. This process continues. After each new event is added, we resolve for Y, and calculate new values of β. After sufficient observations, the logit analysis can identify candidate ports that explain why the honeypots are becoming active.

The inverse time relation between event and prior traffic allows one to record arbitrary periods of traffic. Traffic that occurred too long ago will, in practice, have such a low value for $X_{i,j}$ that it cannot affect the outcome. As a convenience, we cut off

prior traffic t_p at 5 minutes, but even this arbitrary limit is generous. Future work will explore use of other time treatments, such as $\frac{1}{\delta_t^2}$, and $\frac{1}{\sqrt{\delta_t}}$, as a means of further biasing toward more recent network events. Note that this assumption prevents HoneyStat from tracking worms that sleep for a lengthy period of time before spreading. These worms are presumably self-crippling, and have a slow enough spread rate to allow for human intervention.

A key variable in this analysis includes the *Wald statistic*, which lets us test whether a variable's coefficient is zero. The Wald statistic is merely the ratio of the coefficient to its standard error, with a single degree of freedom [HL00]. The Wald statistic can be used to reject certain variables, and exclude them from a model. For example, if ports $P_0, P_1, \ldots P_n$ were observed prior to a honeypot event, we might exclude some of these ports based on the ratio of their coefficient $\beta_{i,j}$, and their standard error. Thus, the Wald statistic essentially poses a null hypothesis for each variable, and lets us exclude variables with zero coefficients. (After all, a variable with a zero β value does not contribute to solving Eq. 1). This analysis is helpful since it reduces noise in our model. However, since it uses a simple ratio, when the standard error is large, it can lead one to not reject certain variables. Thus, the Wald statistic can be used to remove unlikely variables, but might not always remove variables that have no affect.

Applying logistic analysis involves the following steps. First, for a particular honeypot event j, we estimate the coefficients, i.e., $\beta_{0,j}, \beta_{1,j} \ldots \beta_{n,j}$, using *maximum likelihood evaluation* [HL00] (MLE). In this step, we try to find a set of coefficients that minimize the prediction error. Stated another way, MLE assigns values that will maximize the probability of obtaining the observed set of data. (This is similar to the least squares method under simple regression analysis.) Second, we use the Wald statistic to evaluate each variable, and remove those below a user-selected threshold of significance level, say, 5%. The intuition of this step is that we try to evaluate whether the "causal" variable in the model is *significantly* related to the outcome. In other words we essentially ask the question: Is activity on port x significantly related to the honeypot activity or was it merely random?

If the analysis results in a single variable explaining changes in the honeypot, then we report the result as an alert. If the results are not conclusive, the event data is stored until additional events are observed, triggering a renewed analysis. Of course, since the events involve breakins to honeypots, users may also wish to receive informational alerts about these events.

6   HoneyStat in Practice

To evaluate HoneyStat's potential as a local worm detection system, we tested two key aspects of the algorithm: (a) does it properly identify worm outbreaks, and (b) what false positive rate does it produce? Testing showed that HoneyStat could identify worm outbreaks, with a low false positive rate. Our testing with available data showed the false positive rate of zero. This result is encouraging, given the enormous data set used. Nonetheless, a zero false positive rate may be due to properties of the data set and we will continue to run more experiments.

6.1   Worm Detection

In [QDG$^+$], we used data from six honeypots that became active during the Blaster worm outbreak in August 2003. The trace data used for the analysis also included network traffic from some 100 /24 darknet IPs. Figure 3 shows an aggregate view of traffic to all the honeypots on August 9 and 11, as well as background traffic to the darknets.
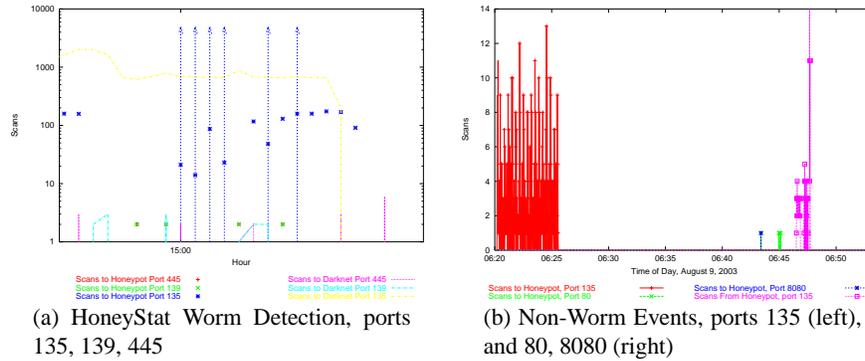

(a) HoneyStat Worm Detection, ports 135, 139, 445

(b) Non-Worm Events, ports 135 (left), and 80, 8080 (right)

Fig. 3. a) *HoneyStat worm detection for Blaster.* The Blaster attack on August 11, 2003, is detected by the honeypots. Upward arrows, not drawn to scale, indicate the presence of outgoing traffic from the HoneyStat nodes. Traffic prior to the honeypot activity is analyzed, using the inverse of time difference, so that more recent activities more likely explain the change in the honeypot. A logit analysis shows that prior scans to port 135 explains these episodes–effectively identifying the blaster worm. b) *Avoiding false positives.* Here, we see a trojaned honeypot node becoming active days prior to the Blaster worm outbreak. However, since this event is seen only in isolation (one honeypot), it does not trigger a worm alert. Traffic to ports 80 and 8080 does not bias the later analysis.

If we mark the honeypot activity as `NetEvents`, we can examine the previous network activity to find whether a worm is present. As shown in Table 1, a logit analysis of the honeypot data shows that of all the variables, port 135 explains the tendency of honeypots to become active. (In our particular example, one can even visually confirm in Figure 3(a) that honeypot activity took place right after port 135 traffic arrived.) The standard error reports the error for the estimated β, and the significance column reports the chance that the variable's influence was merely chance. The Wald statistic indicates whether the β statistic significantly differs from zero. The significance column is the most critical for our analysis, since it indicates whether the variable's estimated β is zero. The lower the score, the less chance the variable had no influence on the value of Y. Thus, we eliminate any variable with a significance above a threshold (5%). From this, the observations for ports 80, 8080, and 3128 can be discounted as not a significant explanation for changes in Y.

Table 1. Logit Analysis of Multiple HoneyStat Events

Variable	β	Standard Error	Wald	Significance
port_80	-17463.185	2696276.445	.000	.995
port_135	3.114	.967	10.377	.001
port_139	1869.151	303.517	37.925	.000
port_445	-1495.040	281.165	28.274	.000
port_3128	-18727.568	9859594.820	.000	.998
port_8080	10907.922	10907.922	6919861.448	.999
constant	.068	1.568	.210	1.089

In this case, the logit analysis performs two useful tasks. First, we use the significance column to rule out variables above a certain threshold, leaving only ports 135, 139 and 445. Second, the analysis lets us rank the remaining variables by significance.

The logit analysis did not pick one individual port as explaining the value of Y. The alert that issues therefore identifies three possible causes of the honeypot activity. As it turns out, this was a very accurate diagnosis of the Blaster outbreak. Recall that just prior to Blaster's outbreak on port 135, there were numerous scans being directed at ports 139 and 445. The port 135 exploit eventually became more popular, since only a few machines were vulnerable on 445 and 139. We are aware of no statistical test that could focus on port 135 alone, given the high rate of probing being conducted on ports 139 and 445. This required human insight and domain knowledge to sort out.

One complicating factor occurs when two zero-day worms attack at the same time, using different ports. For example, consider in the example what would happen if traffic to port 80 were a worm after all. In such a case, more data observations will be required to separate out which events support a leading theory of causation. The logit analysis will eventually select a pattern for one of the worms, and with the removal of those observations, the second worm can be identified. (Recall that once identified, the "explained" data is removed from the analysis queue.) Future work will explore the use of Best Subsets logistic regression models[HL00], to avoid the linear identification of multiple worms.

The number of observations required for logistic regression appears to be a matter of some recent investigation. In [HL00], the authors (eminent in the field) note "there has been surprisingly little work on sample size for logistic regression". Some rough estimates have been supplied. They note that at least one study show a minimum of 10 events per parameter are needed to avoid over/under estimations of variables [HL00]. Since each honeypot activity observation is paired with a corresponding inactivity observation, HoneyStat would need to generate five HoneyStat events to meet this requirement. Section 7 notes how waiting for this many observations potentially affects worm detection time.

Since each event involves an actual compromise of a system, one could also report alerts with a lower confidence level. While we might want more samples and certainty, we can at the very least rank likely ports in an alert.

6.2   Benefits of HoneyStat

HoneyStat provides the following benefits to local networks: (a) It provides a very accurate data stream for analysis. Every event is the result of a successful attack. This significantly reduces the amount of data that must be processed, compared to Kalman filter, and other traditional scan-based algorithms. (b) Since HoneyStat uses complete operating systems, it detects zero day worms, for which there is no known signature. (c) HoneyStat is agnostic about the incoming and outgoing ports for attack packets, as well as their origin. In this way, it can detect worms that enter on port P_a, and exit on port P_b.

Thus, HoneyStat reports *an explanation* of worm activation, and not merely the *presence* of a worm. Other information, such as rate of scans, can be obtained from the traffic logs captured for the logit analysis. [Kre03] has already suggested a simple method of quickly extracting a binary signature, in a manner compatible with Honey-Stat.

6.3   False Positive Analysis

Analyzing the false positive rate for HoneyStat is subtle. Since honeypot events always involve breakins and successful exploits, it might seem that honeypot-based alert systems would produce no false positives. This is not the case. Although the underlying data stream consists of serious alerts (successful attacks on honeypots), we still need to analyze the potential for the logit analysis to generate a false positive. Two types of errors could occur. First, normal network traffic could be misidentified as the source of an attack. That is, a worm could be present, but the analysis may identify other, normal traffic as the cause. Second, repeated human breakins could be identified as a worm. We do not consider this second failure scenario, since in such a case, the manual breakins are robotic in nature, and (for all practical purposes) indistinguishable from, and *potentially just as dangerous* as any worm.

Model Failure It is not feasible to test HoneyStat on the Internet. This would require waiting for the outbreak of worms, and dedicating a large IP space to a test project. We can instead perform an retrospective analysis of a tracefile to estimate the chance of a false positive.

Using a honeypot activity log, dating from July 2002 to March 2004, we used uniform random sampling to collect background traffic samples, and injected a worm attack. The intuition is this: we wish to see if a HoneyStat logit analysis were to cause a false positive. This could occur if normal non-malicious background traffic occurs in such a pattern that random sampling produces a candidate solution to the logistic regression.

The data we use for the background sampling came from the Georgia Tech Honeynet project. We have almost two years of network data captured from the Honeynet. The first year of data was captured on a Generation I Honeynet, which is distinguishable by the use of a reverse firewall serving as the gateway for all the Honeypots. The second year of data was captured from a Generation II Honeynet, which is distinguishable by the use of a packet filtering bridge between all of the Honeypots and their gateway. The data is available to other researchers in a sanitized form.

A random sampling of over 250 synthetic honeypot events did not produce a false positive. This certainly does not prove that HoneyStat is incapable of producing a false positive. Rather, this may reflect the limited range of the data. A much larger data set is required to fully explore the potential of logistic regression to misidentify variables. Even if false positives are found, it should be noted that these are not the usual false positives, or type I errors found in IDS. Instead, a false positive with a HoneyStat node is half right: there are breakins to honeypots, even if the algorithm were to misidentify the cause.

7   HoneyStat as an IDS Tool

The previous sections have shown that HoneyStat can detect worm attacks with a low false positive rate. This shows that it could be incorporated into a local intrusion detection system. A more important question is whether this strategy can detect worm outbreaks early. In this section, we present an analytical model.

A HoneyStat deployment can effectively detect worms that use random scan techniques. The work in [WPSC03] presents a complete taxonomy of worms. In [QDG$^+$], we discussed how a detection algorithm with a suitably low false positive can protect local networks. Accordingly, we evaluate HoneyStat against worms that use only random scanning strategies [ZGGT03]. Since we are interested in local network protections, the results should apply to other types of scanning worms [QDG$^+$]. Realistically we assume the vulnerable hosts are uniformly distributed in the real assigned IPv4 space (all potential victims are located in this space, denoted as $T = 10^9$), not the whole IPv4 space (denoted as $\Omega = 2^{32}$). Assume N is the total number of vulnerable machines on the Internet, n_i is the number of whole Internet victims at time tick i and s is the scan rate of worm (per time tick). So the scans entering space T at time tick $i + 1$ should be $k_{i+1} = s n_i \frac{T}{\Omega}$. Within this space, the chance of one host being hit is $1 - (1 - \frac{1}{T})^{k_{i+1}}$. Then we have worm propagation equation Eq. (2).

$$n_{i+1} = n_i + [N - n_i] \left(1 - (1 - \frac{1}{T})^{s n_i \frac{T}{\Omega}} \right) \qquad (2)$$

In fact because T and Ω are very big, $(1 - (1 - \frac{1}{T})^{s n_i \frac{T}{\Omega}} \doteq \frac{1}{T} s n_i \frac{T}{\Omega} = \frac{s n_i}{\Omega}$. So the spread rate is almost the same as seen in previous models (e.g., Analytical Active Worm Propagation (AAWP) model [CGK03], epidemic model [KW93,KCW93,SPN02] etc.) Now suppose we have a honeypot network with size D ($D \subseteq T$). The initial number of vulnerable hosts is u_0. Generally a network with size D has DN/T vulnerable hosts on average. But with HoneyStat, each network has its own mix of vulnerable OS distributions. Since most worms target Windows, we can intentionally let most of our honeypots run Windows so that we present a higher number of initially vulnerable hosts to the worm. Without loss of generality we suppose $u_0 = D\alpha$. We let α be the minimum ratio for vulnerable hosts. The number of local victims at time tick i is v_i and $v_0 = 0$ which means initially there is no victim in our honeypot network. The time for the first HoneyStat node to become active is t_1 (clearly t_1 is the first time tick i when

$v_i \geq 1$). We have

$$v_{i+1} = u_0 \left(1 - (1 - \frac{1}{T})^{\sum_{j=0}^{i} sn_j \frac{T}{\Omega}} \right) \text{ when } i + 1 < t_1 + t_r \qquad (3)$$

Here $u_0 = D\alpha$. We let t_r represent the time required to reconfigure most of the non-vulnerable honeypots to run the same OS and patch level as the first victim. (In other words, this is the time required, say, to convert most of the Linux honeypots to Windows, if Windows is first attacked.) Since we need more observations for the logit analysis to work, as noted in Section 5, we shift some of our honeypots to match the vulnerable OS. As noted above, restarting the suspended OS happens quickly, before any traffic times out. We nonetheless express this delay as response time t_r. After the delay, the number of new vulnerable hosts becomes $u_1 = D\gamma$. We let γ represent the maximum ratio for vulnerable hosts. Normally $\gamma < 1$ because we may not convert *all* of our HoneyStat nodes to the operating system that is attacked. We might keep a few to detect other worms. Now we have

$$v_{i+1} = 1 + u_1 \left(1 - (1 - \frac{1}{T})^{\sum_{j=t_1+t_r}^{i} sn_j \frac{T}{\Omega}} \right) \text{ when } i + 1 >= t_1 + t_r \qquad (4)$$

Here $u_1 = D\gamma$. We can calculate the time (in terms of the whole Internet infection percentage) when our first HoneyStat node is infected. Table 2 and Figure 5 use the effect of different α and D. For example, we can see that using $D = 2^{10}$ and $\alpha = 10\%$, the first victim is found when only 0.9786% Internet vulnerable hosts are infected.

Table 2. Time (infection percentage) when HoneyStat network has a first victim

α	$D = 2^8$	$D = 2^9$	$D = 2^{10}$	$D = 2^{11}$	$D = 2^{12}$	$D = 2^{13}$	$D = 2^{14}$	$D = 2^{15}$	$D = 2^{16}$
10%	3.9141%	1.9558%	0.9786%	0.4895%	0.2448%	0.1223%	0.0613%	0.0307%	0.0155%
25%	1.5634%	0.7825%	0.3910%	0.1959%	0.0981%	0.0491%	0.0247%	0.0124%	0.0063%
50%	0.7825%	0.3910%	0.1959%	0.0981%	0.0491%	0.0247%	0.0124%	0.0063%	0.0033%
75%	0.5210%	0.2606%	0.1305%	0.0655%	0.0328%	0.0165%	0.0083%	0.0043%	0.0022%
100%	0.3910%	0.1959%	0.0981%	0.0491%	0.0247%	0.0124%	0.0063%	0.0033%	0.0017%

When the first HoneyStat node becomes infected, the other nodes switch to the vulnerable OS. This takes time t_r, after which there will be $u_1 = D\gamma$ vulnerable hosts. After redeployment, the chance of getting the next victim improves. This is shown in Eq. (4). The effect of D, γ and t_r is shown in Figure 4. Here, we can see that after redeployment we will quickly get enough victims when the whole Internet has a low infection percentage. This occurs because we obtain more vulnerable honeypots after the HoneyStat array is switched to match the OS of the first victim. Therefore, we get higher chances of being hit by the worm. For example, if $\alpha = 0.25, D = 2^{16}, \gamma = 0.7, t_r = 10$, it is still very early to have 4 victims in the HoneyStat network, when only 0.013% Internet vulnerable hosts are infected. To have 10 victims, still only 0.0261%
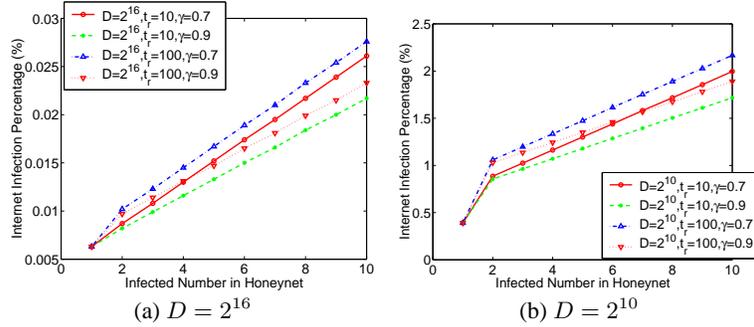
Fig. 4. Effect of HoneyStat network size, D, maximum percentage of vulnerable hosts, γ, and time to redeploy after first victim, t_r, on the victim count. These graphs, drawn with $\alpha = 0.25$; N=500,000; scanrate=20 per time tick; Hitlist=1, show that with a larger IP space monitored by HoneyStat, D, the detection time (as a precent of the infected Internet) improves greatly. Even with only 2^{10} IPs monitored, detection time is quick, requiring only a little more than 1% of the Internet to be infected.

Internet vulnerable hosts are infected. And we can see that $t_r = 10$ or $t_r = 100$, $\gamma = 0.7$ or $\gamma = 0.9$ do not affect the outcome very much. Instead, the size of honeynet D is the most important factor. Thus, the delay in switching HoneyStat nodes does not play a critical role in overall worm detection time.
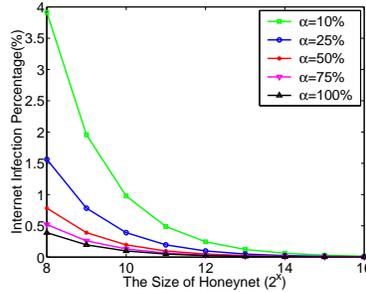

Fig. 5. Effect of α and D on Time (Infection Percentage) when HoneyStat network has a first victim. N=500,000, Scan rate=20 per time tick

In section 4, we noted that machines can be massively multihomed, so that one piece of hardware can effectively handle hundreds of IP addresses in multiple virtual machines. From the discussion of D above, 2^{11} is already a reasonable number of IP addresses that can be used in our local early worm detection. Assuming we had a few

computers sufficient to allow $D = 2^{11}$ and $\alpha = 0.25$, we can see from Table 2 that the first victim appears when on average 0.1959% of Internet vulnerable hosts are infected. Suppose $\gamma = 0.75, t_r = 10$, then to have 5 victims in our honeynet (or enough to have a minimal number of data points suggested in Section 6), it is still very early when only 0.4794% of the Internet's vulnerable hosts are infected. When one IP is infected, we reset the OS so that it can be infected again. This kind of "replacement" policy makes the whole honeynet work as we have discussed above although there are only, say, 64 virtual machines running on every GSX server.

8   Conclusion

Local detection systems deserve further exploration. We have suggested that in addition to increasing the *quantity* of data used by alert systems, the *quality* can be improved as well. It has been said that if intrusion detection is like finding a needle in a haystack, then a honeypot is like a stack of needles. Literally every event from a honeypot is noteworthy. Honeypots are therefore used to create a highly accurate alert stream.

Using logistic regression, we have shown how a honeypot alert stream can detect worm outbreaks. We define three classes of events to capture memory, disk and network activities of worms. The logit analysis can eliminate noise sampled during these events, and identify a likely list of causes. Using extensive data traces of previous worm events, we have demonstrated that HoneyStat can identify worm activity. An analytical model suggests that, with enough multihomed honeypots, this provides an effective way to detect worms early.

While participation in global monitoring efforts has value, we believe local network strategies require exploration as well. Further work could include identification of additional logistic models to sort through large sets of data, coordination of shared honeypot events, integration with other intrusion detection techniques, and response.

Acknowledgments

References

[AFV95]   D. Anderson, T. Frivold, and A. Valdes. Next-generation intrusion detection expert system (NIDES): A summary. Technical Report SRI-CSL-95-07, Computer Science Laboratory, SRI International, Menlo Park, California, May 1995.

[BGB03]   V.H. Berk, R.S. Gray, and G. Bakos. Using sensor networks and data fusion for early detection of active worms. In *Proceedings of the SPIE AeroSense*, 2003.

[CGK03]   Z. Chen, L. Gao, and K. Kwiat. Modeling the spread of active worms. In *Proceedings of the IEEE INFOCOM 2003*, March 2003.

[CLF03]   S. Cheung, U. Lindqvist, and M. W. Fong. Modeling multistep cyber attacks for scenario recognition. In *Proceedings of the Third DARPA Information Survivability Conference and Exposition (DISCEX III)*, Washington, D.C., April 2003.

[CM02]   F. Cuppens and A. Miège. Alert correlation in a cooperative intrusion detection framework. In *Proceedings of the 2002 IEEE Symposium on Security and Privacy*, pages 202–215, Oakland, CA, May 2002.

[Cor04]   Joseph Corey. Advanced honey pot identification and exploitation. (fake) Phrack, No. 63, 2004.

[DW01]   H. Debar and A. Wespi. The intrusion-detection console correlation mechanism. In *4th International Symposium on Recent Advances in Intrusion Detection (RAID)*, October 2001.

[GHH01]   R.P. Goldman, W. Heimerdinger, and S. A. Harp. Information modleing for intrusion report aggregation. In *DARPA Information Survivability Conference and Exposition (DISCEX II)*, June 2001.

[GSQ$^+$04] Guofei Gu, Monirul Sharif, Xinzhou Qin, David Dagon, Wenke Lee, and George Riley. Worm detection, early warning and response based on local victim information. Submitted for review, 2004.

[HL00]   D.W. Hosmer and S. Lemeshow. *Applied Logistic Regression*. Wiley-Interscience, 2000.

[Inc03]   Immunix Inc. Stackguard. http://www.immunix.org/stackguard.html, 2003.

[Ins]   SANS Institute. http://www.sans.org.

[JPBB04]   Jaeyeon Jung, Vern Paxson, Arthur W. Berger, and Hari Balakrishnan. Fast portscan detection using sequential hypothesis testing. In *2004 IEEE Symposium on Security and Privacy*, 2004.

[JX04]   Xuxian Jiang and Dongyan Xu. Collapsar: A vm-based architecture for network attack detention center. http://www.cs.purdue.edu/homes/jiangx/collapsar/, 2004.

[Kal60]   R.E. Kalman. A new approach to linear filtering and prediction problems. *Transaction of the ASME–Journal of Basic Engineering*, March, 1960.

[KCW93]   J.O. Kephart, D.M. Chess, and S.R. White. Computers and epidemiology. 1993.

[Kor04]   Kostya Kortchinsky. Vmware fingerprinting counter measures. The French Honeynet Project, 2004.

[Kre03]   Christian Kreibich. Honeycomb automated ids signature creation using honeypots. http://www.cl.cam.ac.uk/ cpk25/honeycomb/, 2003.

[KW93]   J.O. Kephart and S.R. White. Measuring and modeling computer virus prevalence. In *Proceedings of IEEE Symposium on Security and Privacy*, 1993.

[Lem01]   Jonathan Lemon. Kqueue: A generic and scalable event notification facility. pages 141–154, 2001.

[LLO$^+$03] John Levine, Richard LaBella, Henry Owen, Didier Contis, and Brian Culver. The use of honeynets to detect exploited systems across large enterprise networks. In *Proceedings of the 2003 IEEE Workshop on Information Assurance*, 2003.

[LUR03]   LURHQ. Msblast case study. http://www.lurhq.com/blaster.html, 2003.

[LUR04]   LURHQ. Witty worm analysis. http://www.lurhq.com/witty.html, 2004.

[MHL94]   B. Mukherjee, L. T. Heberlein, and K. N. Levitt. Network intrusion detection. *IEEE Network*, May/June 1994.

[MMDD02] B. Morin, L. Mé, H. Debar, and M. Ducassé. M2d2: A formal data model for ids alert correlation. In *Proceedings of the 5th International Symposium on Recent Advances in Intrusion Detection (RAID)*, October 2002.

[Moo02a]   D. Moore. Code-red: A case study on the spread and victims of an internet worm. http://www.icir.org/vern/imw-2002/imw2002-papers/209.ps.gz, 2002.

[Moo02b]   D. Moore. Network telescopes: Observing small or distant security events. http://www.caida.org/outreach/presentations/2002/usenix_sec/, 2002.

[MSVS03] D. Moore, C. Shannon, G. M. Voelker, and S. Savage. Internet quarantine: Requirements for containing self-propagating code. In *Proceedings of the IEEE INFOCOM 2003*, March 2003.

[NCR02] P. Ning, Y. Cui, and D.S. Reeves. Constructing attack scenarios through correlation of intrusion alerts. In *9th ACM Conference on Computer and Communications Security*, November 2002.

[Par04] Janak J Parekh. Columbia ids worminator project. http://worminator.cs.columbia.edu/, 2004.

[PN97] P. A. Porras and P. G. Neumann. EMERALD: Event monitoring enabling responses to anomalous live disturbances. In *National Information Systems Security Conference*, Baltimore MD, October 1997.

[Pro03] Niels Provos. A virtual honeypot framework. http://www.citi.umich.edu/techreports/reports/citi-tr-03-1.pdf, 2003.

[QDG$^+$] X. Qin, D. Dagon, G. Gu, W. Lee, M. Warfield, and P. Allor. Technical report.

[QL03] X. Qin and W. Lee. Statistical causality analysis of infosec alert data. In *Proceedings of the 6th International Symposium on Recent Advances in Intrusion Detection (RAID 2003)*, Pittsburgh, PA, September 2003.

[QVWW98] D. Qu, B. Vetter, F. Wang, and S.F. Wu. Statistical-based intrusion detection for OSPF routing protocol. In *Proceedings of the 6th IEEE International Conference on Network Protocols*, Austin, TX, October 1998.

[Sei02] Kurt Seifried. Honeypotting with vmware - basics. 2002.

[Sko02] E. Skoudis. *Counter Hack*. Upper Saddle River, NJ: Prentice Hall PTR, 2002.

[Spi03] Lance Spitzner. *Honeypots: Tracking Hackers*. Addison Wesley, 2003.

[SPN02] S. Staniford, V. Paxson, and N.Weaver. How to Own the Internet in Your Spare Time. In *Proceedings of 2002 Usenix Security Symposium*, 2002.

[Sta01] S. Staniford. Code red analysis pages: July infestation analysis. http://www.silicondefense.com/cr/july.html, 2001.

[VMW04] Inc. VMWare. Gsx server 3. http://www.vmware.com/products/server, 2004.

[VS01] A. Valdes and K. Skinner. Probabilistic alert correlation. In *Proceedings of the 4th International Symposium on Recent Advances in Intrusion Detection (RAID)*, October 2001.

[Wil02] Matthew M. Williamson. Throttling viruses: Restricting propagation to defeat malicious mobile code. Technical report, 2002. HPL-2002-172.

[WL03] Matthew M. Williamson and Jasmin Léveillé. An epidemiological model of virus spread and cleanup. Technical report, 2003. HPL-2003-30.

[WPSC03] N. Weaver, V. Paxson, S. Staniford, and R. Cunningham. A taxonomy of computer worms. In *2003 ACM Workshop on Rapid Malcode (WORM'03)*. ACM SIGSAC, October 2003.

[WVGK04] J. Wu, S. Vangala, L. Gao, and K. Kwiat. An efficient architecture and algorithm for detecting worms with various scan techniques. In *Proceedings of the 11th Annual Network and Distributed System Security Symposium (NDSS'04)*, February 2004. to appear.

[YBJ04] Vinod Yegneswaran, Paul Barford, and Somesh Jha. Global intrusion detection in the domino overlay system. In *Proceedings of NDSS*, 2004.

[ZGGT03] C. C. Zou, L. Gao, W. Gong, and D. Towsley. Monitoring and early warning for internet worms. In *Proceedings of 10th ACM Conference on Computer and Communications Security (CCS'03)*, October 2003.

[ZGT02] C. C. Zou, W. Gong, and D. Towsley. Code red worm propagation modeling and analysis. In *Proceedings of 9th ACM Conference on Computer and Communications Security (CCS'02)*, October 2002.