

NSF CyberTrust Center Proposal

Center for Internet Epidemiology and Defenses

Stefan Savage, Geoffrey M. Voelker, George Varghese
University of California, San Diego

Vern Paxson, Nicholas Weaver
International Computer Sciences Institute

October 2004–September 2009

1 Overview

Late last year the Computing Research Association convened a panel of security experts to formulate the next decade's *Grand Challenge* problems in Information Security and Assurance. The panel, drawn from leading academic and industrial members of the CRA's 200 research organizations [17], identified four key problems deserving vigorous attention. First among these was: *Eliminate epidemic-style attacks within 10 years.*

Since 2001, large-scale Internet worm and virus outbreaks have *profoundly* demonstrated the threat posed by self-propagating programs. The combination of widespread software homogeneity and the Internet's unrestricted communication model creates an ideal climate for infectious pathogens. Worse, each new generation of outbreaks demonstrates increasing speed, virulence, and sophistication. While the "Code Red" worm of July, 2001 is widely cited as an example of a potent epidemic attack, it is simple compared to today's contagions. Code Red exploited only a single Web server vulnerability, took over 14 hours to spread across the 360,000 vulnerable hosts it infected, and had a simple "payload" — a flooding attack against the White House web server — that was easily thwarted [50, 72]. Subsequent pathogens have made "advances" in all of these dimensions — they have spread through multiple vectors [38, 39, 9]; shoehorned the entire worm into a single packet for lightning-quick infection [44, 48]; probed all 4 billion possible Internet addresses for potential victims in under 10 minutes, with spread limited solely by the Internet's *contagion carry capacity* [44]; left behind "backdoors" for subsequent access [10, 39] and installed engines for relaying spam [38] and mounting controlled denial-of-service attacks [39]; focused attacks against the software update sites needed to cleanse the infection [73]; patched victims to prevent infection by rival worms [75]; maliciously overwrote random disk blocks [48]; parasitically exploited the backdoors installed by previous worms [9, 74]; and have even spread through populations of network intrusion detection devices, exploiting flaws in the very servers people have deployed to protect themselves [48].

This entire "progression" in pathogen sophistication and virulence has taken *less than three years*. Even more disturbing is the growing degree to which Internet contagion is becoming a vehicle for commercial gain [42, 61], because as attackers become highly motivated (i.e., financially as opposed to vandalism), they increasingly bring greater resources to bear on the "arms race" between attackers and defenders. Thus, we in fact expect the rate of innovation to *accelerate*.

This evolution leads to the sobering conclusion that we may not even have 10 years with which to address this Grand Challenge. It is a truly pressing problem, and one which will require sustained energetic efforts to address. As a major element of such an effort, we envision a dedicated *Center for Internet Epidemiology and Defenses*, and we have assembled a team with the deep experience, technical excellence, zeal, and commitment necessary to turn this vision into a reality.

The present state-of-the-art in understanding and defending against Internet epidemics has not in fact advanced very far since the Code Red episode of July, 2001. In part, our lack of deeper understanding and viable defense mechanisms is due to the absence of a stable scientific foundation for reasoning about the unique character of Internet-borne infections. While it is tempting to repurpose the epidemiological models of infectious disease in humans [29], Internet pathogens are in fact quite different — they are authored by intelligent adversaries. Consequently, traditional stochastic analyses are highly fragile tools for predicting the dynamics or limitations of future outbreaks. Similarly, while traditional computer security research has provided many techniques for reasoning about the security of *particular* resources, few of these approaches address the qualitatively different *scale* posed by network epidemics. By their nature, Internet pathogens do not target any particular host; rather, they cast a wide net and automatically spread to where defenses are weakest. We believe that addressing these characteristics will require the development of a new science of *Internet epidemiology*.

However, it is not sufficient to simply understand Internet pathogens: we must also be able to defend against them. The current approach of detection, characterization, and containment was developed to address the spread of file-based viruses and has not changed significantly over the last five years. Typically, new pathogens are detected in an ad hoc fashion by a combination of intrusion detection systems and administrator legwork. Then, after isolating an instance of the worm or virus, skilled security professionals manually characterize a *signature* for it, and finally this signature is used to contain subsequent infections via updates to anti-virus software and network filtering products. While this approach is qualitatively sound, it is quantitatively insufficient. The response time between the onset of a new outbreak and any meaningful defense is limited to human time scales — typically taking hours, if not days. However, in our recent simulations of Internet worm dynamics, we demonstrated that effective worm containment can require a reaction time of well under sixty seconds [51]. More concretely, consider that in the time it took to read this section, the Slammer worm had contacted well over a billion distinct Internet hosts [45].

Fully automated defenses are therefore a *requirement* for addressing Internet epidemics. However, these same defenses must also be highly precise and accurate, rarely impairing legitimate network communications. This problem is fundamentally difficult due to the adversarial nature of these attacks. Worms and viruses can be designed to exhibit a wide variety of behavior — they can spread very quickly or in a slow stealthy manner, they can mutate, mimic normal traffic, or piggyback upon existing services. Finally, it is not simply enough to address these challenges, but it we must engineer these solutions in a manner that is scalable, operationally viable, and can be deployed in a cost-effective manner. Developing these defenses will not only require advances in the understanding of Internet epidemics, but also in the engineering of active large-scale defenses.

Our goal in building the Center for Internet Epidemiology and Defenses is to address both of these fundamental Grand Challenge needs: to better understand the behavior, dynamics and limitations of Internet epidemics, and to develop systems that can automatically detect and defend against new outbreaks in real-time. The remainder of this proposal describes our approaches to each goal, the broader impact through outreach activities and technology transfer, our management plan, and an in-depth proposal for technical evaluation.

2 Research Plan

As indicated in our overview, our research plan is composed of two interrelated efforts. The first focuses on developing a new science of empirical data and analyses for understanding Internet epidemics, while the second focuses on applying this understanding to the engineering of strong defense systems.

2.1 Internet Epidemiology

A fundamental element of defending against worms is understanding their behavior in its full complexity. Internet-scale worms seen in the wild have exhibited surprising dynamics, including forms of propagation unanticipated by researchers (e.g., Slammer’s “bandwidth-limited” spreading [45]) and emergent behavior due to the worm “ecosystems” that the Internet’s endemic malware can form (e.g., the competitive interplay between Code Red 1, Code Red 2, and Nimda [72, 49]).

Indeed, modern network-borne epidemics are entities of such scale that their analysis is in fact an empirical science. However, for similar reasons of scale, they are extremely difficult to monitor with high accuracy and precision. Consequently, most research, development and policy efforts related to these threats have been based more on guesswork than data. To this end, this section focuses on the techniques for both detecting and tracking worms and viruses in-the-large. Such measurements include both purely passive observation (“network telescopes”) and active experimentation to elicit responses (“honeyfarms”). This work is motivated by three key goals that together form a new science and practice of Internet Epidemiology: *(i)* understanding how in reality worms propagate on an Internet scale, including how these dynamics interact with the distribution of vulnerabilities, network resources and network topology; *(ii)* investigating the core elements necessary to construct a global-scale *early warning system* to detect incipient network epidemics — both to aid in future defense systems and to dynamically activate heavy-weight monitoring infrastructures for analyzing fine-grained interactions; and *(iii)* large-scale forensic and economic analysis such as evidentiary attribution and actuarial modelling that translate scientific results into legal, policy and behavioral changes. Each of these goals requires major research efforts, as discussed in the remainder of this section.

2.1.1 Network Telescopes

A basic challenge for analyzing worldwide network phenomena such as worms is acquiring sufficiently broad visibility into their workings. Monitoring at a single location may miss the early stages of a worm’s spread or, more generally, may lack the diverse perspectives necessary for capturing the worm’s large-scale behavior (how many nodes infected, speed of propagation, changes in activity, etc.).

A powerful tool for acquiring such broader visibility is a “network telescope.” Network telescopes monitor traffic sent to communication dead-ends such as unallocated portions of the IP address space. Since there is no legitimate reason for a host to send packets to those destinations, such traffic provides strong evidence of malicious activity. Such activity includes DDoS “backscatter”, port scanning and, critically for our purposes, probe activity from *scanning* worms. These latter reflect a large class of worms that search for victims by attempting to connect to random Internet addresses. (We discuss approaches for the very difficult problem of detecting *non-scanning* worms in Section 2.2.4.)

Thus, network telescopes give us a powerful tool for monitoring a large class of global worm outbreaks. In an operational context, a rise in the observed scanning rate can be used to detect the onset of new scanning worms. This early warning can then be used to alert defense elements to increase their sensitivity to particular services (for example, by introducing delay to packets associated with those services). Telescopes can also yield excellent forensic information, enabling a more detailed understanding of the spread of a worm, and potentially capturing information about “Patient 0,” the first infected machine. The ability to track the initial point of infection offers significant opportunities to understand how the attack penetrated an internal network, or to discover the initial point of infection for legal or situational-awareness purposes (*attribution*).

The most powerful network telescope in use today is that operated by UCSD/CAIDA. It monitors a set of unused network addresses that, in aggregate, constitute 1/256 of all *possible* Internet addresses and therefore give unprecedented insight into Internet-scale malicious activity. This monitor has proven invaluable in analyzing both the prevalence of denial-of-service attacks in the Internet [51] and the dynamics of a number of large-scale Internet worms [49, 45]. However, its capabilities for worm analysis are diminishing because, increasingly, scanning worms no longer randomly scan the entire address space uniformly but instead use biased address scanning such as preferring nearby addresses over distant ones [10, 9]. In addition, due to the widely publicized successes of the UCSD/CAIDA telescope, attackers are becoming aware of its presence and are able to modify their spreading algorithms to specifically avoid the addresses it monitors.

Consequently, our goal is to build a larger-scale telescope with significantly more sampling diversity to address these problems. This telescope will be structured as two layers. Its front-end sensors will be spread across a large number of address blocks and several monitoring points to achieve sampling diversity. We will leverage not just unallocated address blocks (which attackers can learn about fairly easily) but also unused subblocks within allocated blocks. This latter “dark address space” is much more difficult for an attacker to learn about and also enables highly diverse distribution of the sensors. In principle, any large site or ISP could provide a slice of their unused address space for use in this regard. We have an experimental setup along these lines in operation at the Lawrence Berkeley National Laboratory using 10 of their internally unallocated /24’s (so a total of 2,560 addresses) and a similar setup at UCSD using a /16 network (65,535 addresses) “on loan” from a friendly third-party. We are involved in advanced discussions with several other institutions and service providers about contributing additional resources for this project.

A critical issue with building this distributed network telescope is how to collect the traffic from all the disparate sensor locations to a single (or a few) centralized analysis location. We envision pursuing three approaches. The first is to deploy end hosts to which multiple addresses are assigned. A monitor running on the host then tunnels traffic it receives on those addresses over to the analysis point. (We believe, for example, that we will be able to deploy such an approach across a significant portion of the distributed PlanetLab overlay.) A second approach entails configuring routers to directly tunnel (e.g., via GRE) any traffic destined for particular dark address space. A third approach consists of configuring specific routes for the unallocated subblocks so that the traffic comes directly to the analysis point. This is attractive because it is the least burdensome on the network. However, it is only applicable for subblocks allocated to the access provider connected to the analysis point, or for globally routable address prefixes (generally, a /21 or wider). Fortunately, we believe that in some cases (e.g., ESNET’s connectivity to LBNL) we will be able to employ this approach.

Another crucial issue is *filtering*. For a very large telescope, the volume of data collected is simply enormous. For example, the UCSD/CAIDA telescope creates more than 15 GB of data per day. However, for many forms of analysis we can often distill a great majority of the traffic down to highly aggregated summaries. Therefore, one area of our research concerns forms of filtering implemented either at the front-end sensors, or at intermediary points used to hierarchically aggregate traffic. For worm detection, another important form of filtering is dropping traffic seen by the sensor if the traffic does not correspond to possible worm activity. For example, the responses that constitute DoS backscatter are not of interest for worm detection, since they reflect activity quite distinct from the mechanisms by which worms spread [52]. To investigate the design of efficient filters for this task, we are currently beginning a large measurement study to classify the different types of traffic seen on dark address space and its relevance to different types of attacks.

2.1.2 Honeyfarms

It is important to note that network telescopes are entirely passive entities. While they can observe increased activity from random scanning worms, they do not respond to these scans, cannot be infected themselves, and cannot detect non-random spreading vectors such as e-mail or instant messenger viruses.

However, by combining the broad reach of the network telescope with the active capabilities of a *honeypot* it is possible to accomplish far more. A *honeypot* is a single sacrificial machine whose sole purpose is to be compromised for the purposes of detecting the presence, techniques, and possibly motivations of an attacker. The critical value offered by honeypots is that they behave exactly as a real host would and therefore allow accurate, *in situ* examinations of a worm's behavior. Moreover, since these machines are not used for any legitimate purpose, activity on a honeypot is almost certain to represent real malicious intent.

However, being individual resources, honeypots are poor at providing "early warning" of new worm outbreaks. While any infection of a honeypot will be detected immediately, the probability of any particular honeypot being infected early in a worm's growth is low. Conceptually, one might scatter a large number of well-monitored honeypots across the Internet to address this limitation, but this is impractical because each honeypot requires administration, represents the cost of an individual machine, and if compromised can be used to create malicious false positives.

Instead, we envision creating a more sophisticated and lower-cost analysis system by combining clusters of honeypots with the large address spaces managed by network telescopes — a system we collectively term a *honeyfarm*. A honeyfarm is associated with a network telescope node and receives a subset of the traffic collected by the telescope. By monitoring the behavior of the honeyfarm's servers to detect any attempts to initiate anomalous *outbound* connections, we can detect that a worm is attempting to propagate and simultaneously store a copy of the worm payload for future analysis. From this traffic we envision automatically creating *susceptibility signatures* (by observing which configurations are compromised), *behavioral signatures* (by recording actions the worm takes), and *attack signatures* (by identifying unique features in the infection traffic). We can furthermore *automatically test* the sensitivity of these signatures by applying them to instances of the pathogen captured in additional honeypots.

We propose building the honeyfarm using virtual machines (based on VMWare's GSX server) to host large numbers of system, software and service configurations. Thus, we can create the illusion of a wide spectrum of potential vulnerabilities using a much smaller number of actual systems. Not only can each CPU support a number of honeypots, each running in different virtual machines, but each such honeypot can in turn virtualize a large number of "potential" server IP addresses. For example, suppose a honeypot Web server can sustain 1,000 hits/sec. If a worm is targeting generally idle Web servers (as did Code Red and Nimda), then we can configure the honeypot server to emulate 1,000 different (lightly loaded) Web servers seemingly scattered across the dark address space observed by the telescope. Consequently, the scaling of the honeyfarm approach is mainly limited only by the "background" activity rate across the range of monitored address prefixes. We believe that industrial-strength honeyfarms will in fact be able to emulate 10 million hosts under normal conditions.

Finally, a key point regarding honeyfarms is that we can apply them more broadly than just for detecting random-scanning worms. Since we can fully decouple the routing of traffic to the honeyfarm with the operation of the honeypot servers themselves, it is possible to interdict a far wider range of worms and viruses than the pure network telescope approach. For example, we could use Web links around the Internet to direct spiders towards pages hosted by honeyfarm servers; or create large numbers of email addresses whose domains resolve to honeyfarm mail servers; or subscribe honeyfarm nodes to peer-to-peer networks; or enter honeyfarm nodes into Instant Messaging "buddy lists." Each of these would serve to link the honeyfarm into the application-layer network used by a particular type of (non-address-scanning) worm.

2.1.3 Forensic and Economic Analysis

While developing the science of Internet epidemiology is a critical goal in and of itself, it is equally important that we relate this emerging discipline to the surrounding legal and economic context.

One vital challenge to establishing "cyber trust" relates to the fact that some of the tools and techniques employed to safeguard networks and information systems — such as passively tapping traffic, surreptitiously interdicting and altering communication flows, or spoofing purported replies from honeypots — are also used by attackers to perpetrate unlawful activities and exact damage on our information infrastructure. In this environment, it is increasingly important to construct clear technical, social and legal assurances that define "acceptable" versus "unacceptable" behavior and distinguish between lawful and unlawful acts. Therefore, the development of any technologies for automatically detecting, characterizing, filtering or mitigating network attacks should be informed by legal and forensic risk assessments at the early stages. A key element of our proposal includes research on understanding and developing guidelines for characterizing and protecting relevant property and privacy interests in the context of advances in automating defense mechanisms. We will specifically assess the design of technologies that ensure appropriate intervention points for legal process controls. This effort will build heavily on project member Kenneally's current work in value-sensitive

technology solutions for automated data integration and analysis of large-scale, sensitive datasets.

Similarly, it is equally important that forensic evidence derived from epidemic measurements and defense technologies offer robust value for law enforcement activities. Much of the data gathered using techniques such as network telescopes, honeyfarms and even many automated defense mechanisms have probabilistic assumptions and inferred, rather than observed, conclusions. As part of this proposal we will evaluate the evidentiary impact of these systems, including challenges to establishing the authenticity, reliability, relevance and reproducibility of evidence related to proving unlawful network activity. An important advance in this area would entail establishing confidence intervals in the algorithms developed for the purpose of distinguishing and responding to malicious activity in order to assign probabilistic value to evidence supporting legal positions.

The work will also include determinations of the legal (civil, common law, criminal, regulatory) implications of automated network actions, particularly as it relates to inference of intent, presumptions and burdens of proof, assignment of responsibility, proving source identity, and understanding and defining the “reasonableness” of deploying automated defense mechanisms in light of legal cost-benefit determinations of liability.

Finally, the quantitative data generated by the network telescope and honeyfarm systems provides a global view into which organizations are affected by each epidemic. Combined with active measurements and third-party demographic and economic data, we will investigate the network security risk variables and concentration controls present today. In partnership with CNA, a global insurance carrier, we will use these results to inform the development of an economic actuarial model for exposure to aggregate risk and liability. Such a quantitative model will be critical for supporting the emerging Cyber-insurance industry and consequently for providing economic incentives for corporate investment in deploying new network defense technologies.

2.2 Developing Defenses

While one major thrust of our research effort is focused on “understanding” the behavior of Internet epidemics, the other half of the effort is to use this information to drive the development of defensive technologies. Even in the presence of widely distributed vulnerabilities we must prevent worms and viruses from causing massive damage. As part of the Center’s efforts, we will explore several complementary approaches to this problem, described below.

2.2.1 Vulnerability Signatures

Worm outbreaks to date have all leveraged vulnerabilities that were publicly known and for which a patch had been made available previously. Unfortunately, for a variety of reasons, including drastic variability in administrative procedures and impact on quality assurance, new patches are frequently not applied when they are released [59]. However, since the vulnerability is known in advance, it is potentially possible to recognize attempts to exploit it. Microsoft Research’s Shield system is developing this idea in the context of end systems (i.e., individual hosts) as a way to provide short-term “emulation” of the protection provided by patches [83]. We envision extending this approach to a pure network mechanism to promote effective deployment.

We propose to monitor network flows for content that could potentially exploit a known vulnerability. For example, any exploitation of the MS SQL vulnerability targeted by the Slammer worm could have been detected by searching for any packet with a payload greater than 16 bytes directed to UDP port 1434. However, this example is simplistic and there are many challenges in making this approach practical. In many cases a vulnerability signature may span multiple packets and may require the network to approximate changes in the state of the end host. This problem is notoriously difficult, but one for which our team has expertise from extensive experience with network intrusion detection research and operation [54]. Moreover, accommodating the high speeds demanded by network packet processing will require tradeoffs in accuracy and precision that are not currently well understood.

2.2.2 Signature Extraction

An ideal network defense would instantly detect a new worm, extract a precise signature that uniquely describes the malicious payload, and then broadcast this information to allow network devices to block any network transactions carrying the worm. In fact, this approach roughly mirrors what is currently done today in the anti-virus industry. However, virus signatures are constructed manually based on human inspection. As a result, the time to respond to a new outbreak can be hours or days. In contrast, we know from experience that a quickly-spreading network worm can blanket the Internet in 10 minutes [44]. Consequently, an effective signature-based worm defense must remove humans from the loop and derive signatures in seconds.

While this requirement seems daunting on its face, there is good reason to believe that it may be achievable for a large class of worms that have modest byte runs of invariant payload (e.g., initial exploit and/or decryption routine). In fact, in our preliminary work we have found that all of the Internet worms identified to date have this property; one can isolate a signature for the worm by comparing multiple instances of suspicious network flows and looking for common substrings in the data portion [67].

The early results of our initial investigation of classes of automatic signature extraction algorithms are very promising. By combining a histogram measuring all popular substrings on the network with an associated table measuring the number of source and destination addresses associated with each substring, we find that known Internet worms all quickly reveal themselves [67]. In addition, running on live traffic our algorithms have found signatures for several **new** worms (MyDoom, SoBig.F, Witty) long before they were publicly known, with very few false positives.

By using approximate data structures, such as our multi-resolution bitmaps and multistage filters [22, 24], we can maintain the necessary data structures at very high speeds and with low memory overhead. We currently operate our prototype, Earlybird, at UCSD. Although purely a software implementation, it is able to run in excess of 100 Mbps and identify worm signatures in seconds. We believe we can scale a hardware implementation to 10 Gbps.

However, there is significant work yet to be done. We are still just beginning to understand the dynamics of this approach, what additional inputs may be used to improve its response time or fidelity, how its various parameters should be tuned, and the fundamental tradeoffs in detecting a worm early in its spread with the increased possibility of false positives. Moreover, as attackers adapt to this approach there is additional research needed in extracting signatures for worms with ever-varying content — a potential class of malicious code frequently termed “metamorphic”. Finally, there are many systems and networking problems in protecting such systems from explicit evasion attempts and denial-of-service attacks on their algorithms.

2.2.3 Scan Detection

Another fundamental technique in the anti-worm repertoire is detecting instances of infected hosts by their communication patterns. There are several dimensions that can be examined. Most simply, a host distinguishes itself by sending packets to a large number of different destinations, particularly when exhibiting a large proportion of failed connections relative to successful connections [34, 87]. More complex scan detectors will also look at the particular destinations that are being accessed. Does the distribution exhibit undue uniformity over some range? How much does it vary from the destination address distribution of other hosts or from the same host in a previous time period? Are “dark” address ranges being contacted?

Each of these approaches may suffer from false positives, both due to normal behavior (e.g., DNS servers contact many hosts) and due to “benign” scanning (e.g., P2P applications without rendezvous points). However, taken together we aim to construct a framework that can combine these signals to create highly accurate and tunable detectors. Additional complications arise from the scaling complexities inherent in the Internet context. For example, keeping track of pending connections and per-address state can be difficult when monitoring a large access link or within low-cost “Hard-LAN” devices [85], and analyzers deployed within backbones must grapple with observational ambiguities that arise due to frequent asymmetric routing.

2.2.4 Application Analysis

Worms using more sophisticated spreading mechanisms — “topological,” “meta-server,” and “contagion” worms [86] — are *much* more challenging to detect than scanning worms because their network traffic lacks the tell-tale anomalous fanout (and repeated connection attempt failures) characteristic of address scanning. These worms generate much less of a “sore thumb” network traffic signature than do simpler random-scanning worms. Without sophisticated analysis of the network traffic corresponding to the specific application being exploited, it is essentially impossible to detect the propagation of such worms.

Consequently, a critical requirement for robust worm detection is developing detailed application-level analysis. We propose performing such application analysis in three overlapping phases. First, we will extend the open-source Bro intrusion detection system developed by one of the PIs [54, 96, 97] in accordance with University and Institute policies. Bro already includes a framework for developing analyzers for different network applications, spanning analysis of a broad set of transport protocols: TCP, UDP, ICMP, RPC. Our enhancements will aim to augment Bro with analyzers for five different classes of applications: content distribution (e.g., HTTP, SMTP, FTP, NFS), peer-to-peer (e.g., KaZaA, Gnutella, NNTP), remote user access (e.g., SSH, Telnet, Rlogin), interactive messaging (e.g., IM,

IRC), and network infrastructure (e.g., DNS, BGP/OSPF/RIP, SNMP). Note that Bro already has analyzers for several of these, but not with the full functionality we require.

In the second phase, we will develop misuse-detection algorithms that analyze the semantics of a given application's network communications in order to detect requests or replies that are likely malicious. In the third phase, we will devise application-specific models of the communication patterns associated with different applications. These models will be parameterized by monitoring routine network traffic both within enterprises and across the Internet to relate particular communications with different application elements (clients, servers, peers, proxies, relay agents). We will then devise anomaly-detection algorithms that assess the significance of deviations from the normal communication patterns with respect to the models, in order to spot the presence of worms that exploit the structure of the application's communications.

2.2.5 Network Response

Having detected a worm or infected host, we then must act. In addition to blocking traffic wholesale, we can significantly retard the spread of the worm using connection rate-limiting [77], which has low negative impact even in the presence of false positives. Finally, with sufficient protocol understanding a network defense can sometimes transform the traffic (as in "protocol scrubbing" [40]) to remove its potential to infect the end host. This approach is ideal when the accuracy of detection is low and the side-effects of transformation are minimal (e.g., transforming Unicode strings to UTF-8 format).

However, in all of these actions it is critical to modulate the effects of a response with the accuracy and precision of the detection mechanism. For example, while it quickly became known that the Code Red worm spread on TCP port 80, only the most paranoid security posture will suggest blocking all Web traffic in response. An important part of our work is to understand the interactions between countermeasures and detectors and the tradeoffs between different options. Real systems will require hybrid responses that escalate from less stringent countermeasures to stronger ones over the lifetime of a worm, as detection becomes more accurate and precise. We have begun preliminary exploration of incorporating such "worm escalation" feedback in a containment system in [87].

Finally, all of these countermeasures again suffer from the scaling complexities of high-speed implementation in network hardware. While there are demonstrations of individual technologies, such as our deterministic high-speed string matching [76], making these general and flexible enough to protect against a dedicated adversary is a vital challenge we must surmount. Along these lines, many of the network detection and response devices will need to be placed in the local area network, creating "Hard-LAN" devices [85], as both conventional coarse-grained perimeters (exemplified by firewalls and intrusion detection systems) and end-host antivirus systems have severe limitations when used for worm defense. Developing such LAN devices requires substantial research, as the LAN vantage point and scale requires two orders-of-magnitude lower cost/performance when compared with access-link locations. In this regard, one promising platform for such development that we will leverage is the highly programmable, open-source XORP project [1, 28] headed by one of our team members (Ghosh).

2.2.6 Global Correlation and Coordination

Internet-scale worm defense is fundamentally different from existing network security systems due to its deeply distributed nature. In no other domain do we need to synthesize a single global determination ("a worm is spreading with characteristics X and Y ") from such distant inter-domain signals. The problem is even harder, though, because then in addition we must turn that global determination into local action at many thousands of sites throughout the network.

Clearly, a simplistic centralized approach is doomed. We will explore alternatives such as hierarchically structured designs like those employed for aggregating alerts from intrusion detection systems [70], fully decentralized designs that leverage recent advances in constructing very large-scale peer-to-peer systems [58, 71, 60], and rapidly disseminating information through them via *gossip*-based protocols [18, 80].

In addition, in our investigations we must bear in mind that any global coordination mechanism *must* be capable of withstanding targeted, and potentially massive, attack, as the worm may commit some of the resources it acquires during its propagation to attempt to either overwhelm the coordination mechanism or poison its analyses with misleading sensor data. Thus, we need to identify systems that avoid bottlenecks and single points of failure, and can leverage their own scale via large-number statistics to achieve robustness even in the presence of false inputs.

3 Evaluation Plan

A critical aspect of designing architectures and algorithms for successful worm defense is evaluating the systems under a wide variety of network configurations and plausible attacks. For instance, a particular defense may be sensitive to the topology of the network, available bandwidth, the distribution of the defense infrastructure, the character of the background traffic, and the virulence of the worm. The potentially immense scales and widely-varying network characteristics that must be explored make this extremely challenging. For example, how do we experiment with a 10,000-router topology using a physical testbed?

We have a four-fold strategy for addressing these challenges. First, we will enhance ModelNet [79], a system developed by one of our team members (Vahdat) that provides a scalable network emulation environment. ModelNet enables unmodified applications (e.g., susceptible Web servers or attacking worms) to run on unmodified operating systems while their packet communication is transparently subject to the characteristics of a user-specified, large-scale network topology. Currently, ModelNet can successfully model (on a commodity UNIX workstation) the hop-by-hop characteristics, including queuing and congestion, of a 10,000-node router topology with an aggregate bisection bandwidth of 1 Gbps. We can include additional bandwidth by distributing the task of emulation across multiple ModelNet nodes. The applications generating the traffic run across a cluster of machines, with 10–100 logical “edge” nodes, each with their own unique IP address, attached to user-specified locations in the router topology.

While the existing ModelNet infrastructure provides a powerful emulation environment, a number of technical challenges must be overcome before the infrastructure is suitable for large-scale experimentation with worm defenses. The core ModelNet emulation engine must be modified to enable “extensible” routers that can assist in various worm defense architectures, as well as embedding live defenses. Currently, emulated routers employ simple packet abstractions and do not consider packet contents or maintain auxiliary state. Challenges include performing user-specified per-packet functionality in ModelNet while maintaining emulation accuracy and speed and defining an effective extensibility API. We will also investigate scaling ModelNet significantly beyond current performance limits, ultimately supporting emerging 10 Gbps Ethernet hardware in ModelNet.

Second, we will use the DETER testbed currently being developed in a joint NSF/DHS project, for which Dr. Paxson is one of the PIs. We can both deploy ModelNet functionality on top of DETER nodes as well as using other emulation abstractions. Regarding this latter, we have begun developing worm simulation “scale-down” mechanisms as part of the DETER/EMIST project, and will further develop this resource to test our defensive developments.

Third, by deploying network telescopes, honeyfarms, and prototype defenses within our institutions, we will both directly test their efficacy against existing endemic worms and evaluate whether our defenses are effective whenever a new outbreak occurs. Likewise, the endemic (now generally non-threatening) and easily blocked nature of existing pathogens allows these worms to be used for *in vivo* testbed experiments.

Finally, in addition to exploiting actual live attacks to strike, we will develop a low-footprint daemon for use in both testbed and live network environments. Although this daemon does not self-propagate (and therefore is not a worm), its communication patterns will be programmed to emulate propagating worms. Thus it will allow us to test network-level defenses, both in ModelNet, on the DETER testbed, and on any operational subnets that can tolerate a potentially disruptive experiment, without our needing to actually create or employ live pathogens. Critically, this enables us to test defenses against sophisticated worms, such as “metaserver” or “contagion” attacks [86], which have not yet been seen in the wild.

4 Education and Outreach

The primary goal of our educational and outreach efforts is to develop an extensive curriculum for educating a new generation of network security students and workforce professionals in Internet Epidemiology and Defenses (IED). The new challenges posed by IED threats necessitate new skills, analysis techniques, and software systems. The Center will address its education and outreach goals through three broad efforts: integrating our research results into our undergraduate and graduate curricula, distributing this curricula and key project materials through existing outreach efforts in networking, and offering focused education to both the research and professional network security workforce via an annual IED outreach workshop. We discuss each briefly in turn.

In the context of higher education, we plan to directly incorporate our research results into networking and security curricula designed for upper division undergraduate and graduate courses. All the PIs have a history of commitment to education and curricula development. Their courses are frequently ranked by students among the top courses in

their department. Further, Voelker received the 2001 School of Engineering Teaching Award at UCSD, Vahdat was awarded the 2003 Duke University David and Janet Vaughn Distinguished Teaching Award, and Varghese received a “Best Teacher in Computer Science” award in 2001 as well as a “Mentor of the Year” award in 1997. The effectiveness of these courses rests partly on incorporating research material into the coursework. The upper-division undergraduate courses offered by the PIs are project-oriented and incorporate current research, and undergraduates working with the PIs have co-authored a number of conference publications to great success as a result. For example, one of our undergraduate researchers, Stefan Schoenmachers, was a finalist for the CRA outstanding undergraduate award and received a highly-competitive NSF fellowship in 2004. Moreover, all of our graduate courses involve an original research effort culminating in a project report and presentation during an end of term “mini-conference”. A number of these course efforts have subsequently led to conference publications.

We plan to disseminate IED course curriculum through two means. First, we will collect IED education materials developed across all of the undergraduate and graduate courses taught by Center faculty, including lectures, projects, exams, data, and tutorials, and make them available through the Center’s *IED Web Portal*. Second, we will incorporate our educational materials into two existing CAIDA educational projects, the Internet Engineering Curriculum (IEC) [4] and the Internet Teaching Labs (ITL) [5]. The IEC serves as a central repository for teaching materials on Internet topics, which we will extend with a package on IED. The ITL bundles hands-on network laboratory materials and equipment for use at universities around the country. We will create a lab component for experimentation with detecting and analyzing epidemics and deploying defenses against them.

Finally, we will develop, organize, and lead a new IED outreach workshop. This workshop will train network security professionals, researchers and graduate students in basic and leading-edge techniques for understanding and defending against Internet epidemics. It will be a day long workshop including lectures, demonstrations, and tutorials. Initially, we plan to hold the outreach workshop in conjunction with the existing ACM Workshop on Rapid Malcode (WORM) that is already led by PIs Savage and Paxson. Also, using the materials developed for the outreach workshop, Center PIs will further disseminate research results and techniques by giving tutorials at conferences such as SIGCOMM and USENIX. Our group has a long history of providing this kind of professional research and workforce education. PI Paxson has taught tutorials on network measurement and security at ACM SIGCOMM, ACM CCS, ITC, and IEEE/ACM Supercomputing, and developed courses for the University of California Extension and the Interop Graduate Institute; PI Varghese has taught tutorials on high-speed protocol processing at ACM SIGCOMM, IPAM, PODC and SIGMETRICS (including Best Tutorial, SIGMETRICS ’97), Kenneally has developed curricula for judges and other legal professionals focused on the forensic use of digital evidence, and PI Weaver has (in collaboration with Dan Ellis of MITRE) already developed a day-long IED technical tutorial, “The Worm and Virus Threat” [20], successfully presented at ACSAC 2003 and several US government agencies. We will further develop this material for presentation at universities, conferences, corporations, and US Government institutions.

5 Partnership and Technology Transfer

A project of this scale requires the active support and participation of a wide variety of partners. We describe our current partners, our history of involvement, and our plan for collaboration below. The Center will undoubtedly attract other partners as well, but we consider our initial set already sufficient for our goals.

AT&T. Our group has a long-standing peer relationship with members of AT&T’s Research Labs. ICIR was originally funded by AT&T and developed strong personal relationships with their researchers; for example, PI Paxson has co-authored 11 papers with AT&T researchers. UCSD also has close ties with AT&T, who is the founding sponsor for UCSD’s new Center for Networked Systems. As a part of this relationship UCSD has hosted annual meetings of AT&T research staff and they in turn have supported large numbers of UCSD graduate internships (10 this year). As a direct result of one of these collaborations, Cisco will be changing their routing protocol implementation across all of their high-end routers. Many of our groups members have also been involved in joint service activities for the networking and security communities. For example, AT&T’s Steve Bellovin recently co-organized with PI’s Paxson and Savage, and Stuart Staniford, the DIMACS Workshop on Large-Scale Internet Attacks, and PI Paxson works closely with AT&T’s Jennifer Rexford as vice-chair and chair of ACM SIGCOMM, respectively.

AT&T is a key partner in this effort for many reasons. First, they have a unique set of network vantage points for monitoring and acting on network epidemics. Second, managing these threats is considered core to their business. In a recent press interview, Hossein Eslambolchi (AT&T’s CTO) stated that security was one of the six strategic areas of research for AT&T Labs. “We are looking at innovations related to network-based security. We need a lot better ways

to do forensic analysis of viruses and worms” [41]. Finally, AT&T is unique among the large network providers in having a world-class research group who are natural collaborators and serve as a gateway between the operational and business requirements of Internet security and the research developments provided by our Center.

Microsoft is at the forefront of the worm and virus problem because, so often, it is their software which is targeted. They are actively involved in research relating to epidemic defense, notably the Shield project described earlier, and are close colleagues. Our collaboration with Microsoft has taken many flavors, including graduate student internships, site talks (PI Savage was a member of Microsoft’s recent Software Security Summer Institute and Microsoft Research’s Helen Wang recently traveled to UCSD to discuss the progress of Shield), service activities (PI’s Savage and Paxson serve with Wang on the WORM Program Committee), and providing access to code, network trace data and financial support. We expect to be direct collaborators in advancing vulnerability-filtering defenses, and also to influence Microsoft’s overall strategy for dealing with epidemic attacks.

Hewlett-Packard. HP Research’s Matt Williamson recently authored a groundbreaking paper on automated virus throttling that pioneered a large class of scan-detection defenses [88]. Hewlett-Packard will be donating equipment to this proposal and will be active collaborators with developing more powerful scan detection. We also have a history of internships at HP Labs (3 this summer) and expect collaboration anchored by internships to continue to prove fruitful.

Intel. There are two areas of collaboration planned with Intel. First, PI’s Paxson and Savage are founding members of a new distributed “network health” project at Intel’s Berkeley Labs. Led by Joe Hellerstein, this project will devote Intel resources to developing a very large scale distributed end-host system (10s–100s of thousands of users) for monitoring global network activities. We intend to leverage this infrastructure to correlate the results of our Honeyfarm infrastructure and to seed the Honeyfarm with input from stealthy worms (e.g., E-mail viruses, instant-messaging attacks) otherwise invisible. The second area of collaboration is with Raj Yavatkar’s Network Processor Group, using Intel’s IXP platform for experimenting with high-speed algorithms for scan detection and signature extraction.

CNA Insurance. CNA is a leading global insurance carrier that is aggressively developing Cyber-insurance policies. We have been active collaborators with CNA’s Mark Silvestri on the needs of the insurance industry and in developing demographic and actuarial models of worms and viruses. Mark will serve on our technical advisory board and as a liaison between our efforts and the broader insurance-industry policy arena. This collaboration is critical to achieving our vision. It is not enough to simply develop successful anti-worm defenses — they must also be broadly deployed. Unfortunately, the history of research is littered with good ideas that foundered on the rocks of “demonstrable return on investment.” A strong Cyber-insurance market provides a lever for directly affecting these corporate IT investments. If a new technology can be shown to reduce risk, and policy premiums are discounted as a result, corporate investments in leading security technology can be rationalized much more readily.

Internet2. Internet2, through its Abilene network, provides one of the largest non-commercial data networks in the world. Its national backbone is provisioned for 10 Gbps per link and it carries all data traffic between its 200-plus member institutions. However, Internet2 was conceived not only as a vehicle for high-bandwidth connectivity, but also to provide a unique research platform as well. Through the Abilene Observatory, Internet2 provides fine-grained access to flow data, routing data, status reports, etc. They also serve as a testing ground for the deployment of leading edge research prototypes and as an early tester of beta products from startups. PI Savage is a member of Internet2’s Network Research Liaison Council and consulted on the creation of the Abilene Observatory. Together with PI’s Voelker and Paxson, he has a strong relationship with members of the Internet2 research team as well as Steve Corbato, who serves as Director of Backbone Network Infrastructure.

We expect to use the Abilene Observatory (and also ESnet — see accompanying letter of support) as part of our epidemic analysis and to feed our network honeyfarm infrastructure. We will also work with Internet2 members to deploy our worm and virus defense prototypes as we have at UCSD and the Lawrence Berkeley National Laboratory (where PI Paxson holds a joint appointment).

CAIDA. Senior members Moore and Shannon are both members of UCSD’s Cooperative Association for Internet Data Analysis (CAIDA). CAIDA is a unique entity in network research and is perhaps the largest non-profit organization focused on network data measurement, collection and analysis. They have come to serve as one of the primary sources and clearinghouses for Internet-related data as well as a liaison between the research community and the operational portions of the network service provider and network equipment communities. Through CAIDA, the Center has strong ties to Cisco, AOL, Verio, MCI/Worldcom and relevant parts of the Department of Homeland Security such as the National Communications Services (NCS). As well, CAIDA hosts regular graduate internships and these students will have opportunities to collaborate with the Center. Finally, through CAIDA we have a long history of making leading datasets directly available to the research community.

DETER. PI Paxson is also an investigator in the DETER project funded by NSF and DHS. As part of this Center’s

activity, he will serve as liaison to DETER and we will provide them with software, experiments and data for driving realistic worm simulations and defense evaluations. In addition, we will provide an updated and modified ModelNet system that will allow DETER to safely emulate worms in networks with tens of thousands of nodes.

IRTF. The Internet Research Task Force (www.irtf.org) is the research counterpart to the Internet Engineering Task Force (IETF) open standards body. The IRTF encompasses a number of research groups focusing on areas that have relevance for the evolution and transfer of Internet technology in general and Internet standards in particular. From his heavy involvement with the IRTF, PI Paxson is ideally positioned to foster the efforts of the Center serving as the nucleus of a broad worm-defense research community.

6 Management Plan

The Center management team consists of the PIs Savage, Paxson, Voelker, Weaver, and Varghese. PI Savage will be the Project Director for the Center, serve as the central point of contact with NSF, and lead Center efforts at UCSD; Co-PIs Voelker and Varghese will closely interact with Savage to assist in Center management. PI Paxson will lead Center efforts at ICSI, supervise ICSI personnel, and with co-PI Weaver provide overall technical leadership at ICSI.

Some of the Center staff is focussed on particular elements of the overall effort. Kenneally's expertise will be dedicated to legal, economic/insurance, and evidentiary issues. Ghosh will supervise the development of XORP-based worm defense devices. Vahdat will lead the expansion of Modelnet's capabilities.

Other Center staff members already have extensive experience across many different facets of the problem space. Along with the management team above, this includes Moore, Shannon, Snoeren, Marzullo, Calder and Allman. The fluidity of these personnel thus give us a vital degree of management flexibility in adapting focus as the research efforts unfold.

PI Paxson has long-established close working relationships with each ICSI co-PI and will ensure the integration of their activities with the greater Center efforts. Faculty PIs and Senior Personnel will closely supervise graduate student researchers and staff supported by the Center. PIs Savage and Voelker will supervise the efforts of CAIDA Senior Personnel Moore and Shannon. In addition, the Center will draw upon the very extensive ties the PIs have with the worm research, operational response, and commercial worlds to create an Advisory Board composed of representatives from the academic, commercial, and operational network and security stakeholders to guide the development of the Center.

The PI teams will interact and communicate at various time scales to manage Center efforts. First, on a continuous basis, PIs, Senior Personnel and students will interact via a shared mailing list, shared CVS repository and shared secure file space (already in daily production use between the two sites). Second, PIs and Senior Personnel will meet with graduate students and staff on a weekly basis to ensure short-term progress on all projects. Third, the PI teams for both locations will conduct monthly teleconferences to communicate status reports and to coordinate longer-term efforts, building upon the energetic style the institutions have developed collaborating already for more than a year on worm research. Finally, we will conduct annual Center meetings alternating between the UCSD and ICSI locations. These meetings will bring together all Center personnel, including PIs, staff, and students, together with the Advisory Board and representatives of Center industrial partners (see letters of support). These meetings will serve a number of objectives: by gathering all personnel, they will ensure comprehensive and coordinated progress and direction of Center activities; with Advisory Board attendance, they provide an ongoing opportunity for practical feedback from external perspectives; and, by including industrial partners, they spark the technical transfer process. We will allocate a portion of our travel budget to support travel between team locations for these annual meetings as well as for ad hoc travel as necessary to respond to unforeseen incidents and opportunities.

The Center PIs will coordinate technology transfer with Center partners. PIs will initiate technology transfer opportunities by visiting partner locations to communicate Center efforts and contributions. Actual transfer of specific technology will happen in a number of ways, including the annual Center meetings, student internships at Center partner locations, and the dissemination of data sets, analyses, software artifacts, and scientific papers via the Center Web portal. PI Paxson will serve as liaison to the Internet standards bodies through his extensive involvement with the IETF's Internet Research Task Force. Similarly, Kenneally will serve as liaison to the law enforcement and legal communities as well as leading our relationship with the insurance industry. Finally, the system programmer/analyst proposed for the Center will serve a vital role in transferring Center technology by ensuring the production level quality of data sets and software artifacts.

The Center will fulfill its mission of education and workforce development using various means. The Center team

will create and disseminate Internet Epidemiology and Defense (IED) course curricula and projects via the Center Web portal. In addition, the Center PIs will organize and lead the IED outreach workshop (Section 4) to educate national academic and industrial researchers and personnel. The proposed administrative assistant for the UCSD team is critical to the success of education and workforce development. In addition to supporting the research activities of the PIs and the Center graduate students, this assistant will manage the Center Web portal, disseminate course curricula and project materials, and organize the annual Center meetings and IED outreach workshops.

The schedule and milestones for the Center are as follows. We note that in addition the Center must always remain poised to respond nimbly and in a supple fashion to new Internet epidemics as they break and to unforeseen turns in the on-going arms race with attackers.

Year 1:

- Construct and deploy large-scale distributed network telescope.
- Investigate distributed data collection, filtering, and analysis methods.
- Develop initial epidemic defenses: vulnerability signature creation, payload signature extraction, and scan detection.
- Implement prototype LAN-based scan-suppression devices on top of XORP.
- Conduct legal and forensic risk assessment associated with initial methods for distributed data collection, filtering, analysis, and prototype defense mechanisms. Outline current state of the law relevant to privacy and property rights, acceptable behavior, and evidentiary issues that affect the legal risk. This will be a continuous activity throughout the course of the project, in response to the dynamics of real-world legal developments.
- Present detailed analyses on the behavior, severity, and scope of critical network epidemics as they occur. Effort triggered by real events and spans the lifetime of the Center.

Year 2:

- Experiment with and evaluate distributed telescope methods using a combination of synthetic epidemic events as well as experience with real events gathered over the first year.
- Construct and deploy honeyfarms associated with network telescopes. Implement application routing techniques to concentrate worm and virus epidemic events on honeyfarms. Optimize virtual machine technology to scale honeyfarms to support vast address ranges.
- Extend Bro intrusion detection system to implement analyzers for extensive network application protocols.
- Evaluate hardware-assisted Bro as a candidate for Hard-LAN and XORP-based detection and response.
- Migrate XORP-based suppression device to possible FPGA implementation.
- Develop guidelines in context of legal risk assessment to inform design, implementation and application of detection and response mechanisms.
- Frame and inform the legal debate by publishing use case scenarios that apply the developed value-sensitive technologies.

Year 3:

- Using data from real events gathered over the previous year, evaluate honeyfarms as a function of operating system and application distribution, address space distribution, and scale.
- Evaluate tradeoffs of passive detection methods using network telescopes with active detection methods of honeyfarms.
- Develop misuse-detection algorithms for the Bro IDS to analyze the semantics of application network communication.

- Assess candidate mechanisms for global coordination, begin prototype implementation of the most promising for XORP containment platform.
- Conduct a forensic impact evaluation of data gathered via automated tools. Evaluate effect of tool implementation and abstraction errors on data reliability.

Year 4:

- Deploy initial Internet epidemic defenses, including drop filters in routers, connection rate limiting, traffic transformation, and limited coordination.
- Experiment with and evaluate using synthetic worm events and attacks against the system.
- Devise application-specific models of application communication patterns for the Bro IDS to detect deviations from normal behavior.
- Develop metrics to quantify confidence intervals in the prototype algorithms. Extrapolate to practical information assurance applications.
- Validate proactive defenses based upon real epidemics; evaluate FPGA implementation of scan suppression system.

Year 5:

- Develop high-speed implementations of epidemic defenses and countermeasures. Scale to performance requirements of core Internet routers.
- Continuous deployment and measurement of epidemic defenses.
- Evaluation of system based upon external, real epidemics over time.
- Validation of proposed techniques, additional research to deal with unexpected behaviors and conditions.
- Architect fusion of worm detection coordination framework with Intel's Public Health system.
- Conduct probabilistic modeling of aggregate network security threats and develop corresponding risk concentration controls to inform development of economic actuarial model for exposure to aggregate risk and liability.
- Multi-year longitudinal analyses of data gathered by both passive epidemiology tools (e.g., network telescopes) and active tools (e.g., honeyfarms).

7 Results from Prior NSF Support

Vern Paxson is PI of NSF Grant STI-0334088, "Viable Network Defense for Scientific Research Institutions." This new project explores a number of research issues that arise in operating high-speed network intrusion detection at sites that rely upon such monitoring as a primary means to maintain their "open" nature. Preliminary results include [68] and elements of [87, 34]. Paxson is co-PI of NSF Grant NRT-0335290, "NRT: Collaborative Research: Testing and Benchmarking Methodologies for Future Network Security Mechanisms." This new project aims to develop sound methodologies for evaluating Internet security mechanisms in the context of the NSF/DHS-sponsored DETER testbed. Initial efforts have focused on formulating an Internet-scale phenomenological assessment of the *Slammer* worm and studying how to faithfully "scale down" such a large event. Paxson is co-PI of NSF Grant ANI-0205519, "Addressing Fundamental Issues for Robust Internet Performance." His work under this project has concerned analyses of *distributed stresses*, in particular Internet worms. Among the results to date are [45, 46, 86, 53], and elements of [87]. Paxson was co-PI of NSF Grant 9711091, "Creating a National Internet Measurement Infrastructure." This one-year pilot project led to the initial creation of *NIMI* (National Internet Measurement Infrastructure). *NIMI* development subsequently proceeded funded by DARPA, and is now supported under NSF Award 0222846, "An Open Infrastructure for Network Performance and Security Monitoring." *NIMI* currently comprises 50 measurement platforms around the Internet, and has been used to conduct a number of published measurement studies. Among the results directly related to measurement infrastructure are [57, 55, 56, 27].

Atanu Ghosh is PI of NSF Grant ANI-0129541, “An eXtensible Open Router Platform (XORP).” This project will provide a stable platform for routing research, both for new routing protocols and to build testbeds. Some early results include [28].

Stefan Savage, Geoffrey Voelker and David Moore are PIs on recent NSF Grant CCR-031160, “Quantitative Network Security Analysis,” analyzing the prevalence of denial-of-service backscatter, port scanning, and supporting early worm analysis such as [48]. Savage is also co-PI on the late-2003 award EIA-0303622, “FWGrid: A Research Infrastructure for Next Generation Systems and Applications,” providing infrastructure support at UCSD for building a large computation and storage infrastructure. Moore has also been a PI on the recently concluded NCR-9711092, “CAIDA: Cooperative Association for Internet Data Analysis,” which funded his involvement in a wide variety of network measurement and analysis efforts [33, 31, 30, 32, 25, 19, 62, 63, 43, 35, 47] and has created datasets widely used throughout the community. Finally, Moore is a co-PI (with George Varghese) on ANI-0137102 “New Directions in Accounting and Traffic Measurement” which has produced several key techniques for accurately estimating traffic workloads and events [24, 22, 21, 23].

Varghese has been a PI on four other NSF grants over the last decade, most recently ANI 0074004, “Terabit Lookups”, which has produced a wide range of innovations in high-speed packet processing [81, 2, 84, 66, 69].

Amin Vahdat has been supported by a 2000 NSF CAREER award, CCR-9984328, “Balancing Performance, Security, and Resource Utilization in Wide-Area Distributed Systems,” which led to a body of work on *informed transcoding* [11, 12, 13, 14, 15], overlay networks [3, 36, 37, 78], and distributed resource management SHARP [16, 26]. He was also co-PI on Grant ITR-0082912, “System Support for Automatic and Consistent Replication of Internet Services,” which focused on consistency in replicated systems, and the inherent tradeoffs between availability [90, 91, 92, 93, 94, 95]. Finally, Vahdat is PI on ongoing NSF grant, CCR-0306490, “Evaluating Global-scale Distributed Systems using Scalable Network Emulation,” which has supported development and deployment of ModelNet [79, 89].

Brad Calder has been involved in 8 NSF grants (four completed and four current). The four completed grants include an NSF CAREER Award CCR-9733278, “Value and Memory Access Profiling for Compiler Optimization”, CCR-9808697, “Hardware Generation of Threads in a Multithreading Processor”, CCR-0073551, “Predicate-Sensitive Software and Hardware Analysis to Enable Optimization and Speculation”, and CCR-0105743, “Critical Path Computing”. These projects have produced a range of leading and widely cited results in computer architecture, including value profiling [7], cache conscious data placement [6], threaded multiple path execution for SMT [82], selective value prediction [8], automated phase classification [64], wide word pipelined memory network processors [65].

Alex Snoeren just received an NSF CAREER Award, CNS-0347949, “CAREER: Decoupling Policy from Mechanism in Internet Routing” focused on separating network forwarding policies from route discovery mechanisms.

Mark Allman, Erin Kenneally, Colleen Shannon, Stuart Staniford, and Nick Weaver have not previously served as PIs on NSF grants.

8 References Cited

- [1] XORP: eXtensible Open Router Platform. <http://www.xorp.org/>.
- [2] F. Baboescu and G. Varghese. Scalable packet classification. In *Proceedings of the ACM SIGCOMM Conference*, San Diego, California, Aug. 2001.
- [3] R. Braynard, D. Kostić, A. Rodriguez, J. Chase, and A. Vahdat. Opus: an Overlay Peer Utility Service. In *Proceedings of the 5th International Conference on Open Architectures and Network Programming (OPENARCH)*, June 2002.
- [4] CAIDA. CAIDA Internet Engineering Curriculum Repository. <http://www.caida.org/outreach/iec>.
- [5] CAIDA. CAIDA Internet Teaching Laboratory. <http://www.caida.org/outreach/itl>.
- [6] B. Calder, K. Chandra, S. John, and T. Austin. Cache-conscious data placement. In *Proceedings of the Eighth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS-VIII)*, San Jose, 1998.
- [7] B. Calder, P. Feller, and A. Eustace. Value profiling. In *International Symposium on Microarchitecture*, pages 259–269, 1997.
- [8] B. Calder, G. Reinman, and D. M. Tullsen. Selective value prediction. In *ISCA*, pages 64–74, 1999.
- [9] CERT. CERT Advisory CA-2001-26 Nimda Worm, <http://www.cert.org/advisories/ca-2001-26.html>.
- [10] CERT. Code Red II: Another Worm Exploting Buffer Overflow in IIS Indexing Service DLL, http://www.cert.org/incident_notes/in-2001-09.html.
- [11] S. Chandra, C. Ellis, and A. Vahdat. Multimedia Web Services for Mobile Clients Using Quality Aware Transcoding. In *Proceedings of the Second ACM/IEEE International Conference on Wireless and Mobile Multimedia*, August 1999.
- [12] S. Chandra, C. Ellis, and A. Vahdat. Application-Level Differentiated Multimedia Web Services Using Quality Aware Transcoding. *IEEE Journal on Selected Areas of Communications (JSAC)*, 18(12), December 2000.
- [13] S. Chandra, C. S. Ellis, and A. Vahdat. Differentiated Multimedia Web Services Using Quality Aware Transcoding. In *INFOCOM 2000 - Nineteenth Annual Joint Conference of the IEEE Computer And Communications Societies*, March 2000.
- [14] S. Chandra, C. S. Ellis, and A. Vahdat. Managing the Storage and Battery Resources in an Image Capture Device (Digital Camera) using Dynamic Transcoding. In *Third ACM International Workshop on Wireless and Mobile Multimedia (WoW-MoM'00)*, August 2000.
- [15] S. Chandra, A. Gehani, C. S. Ellis, and A. Vahdat. Transcoding Characteristics of Web Images. In *Multimedia Computing and Networking 2001 (MMCN'01)*, January 2001.
- [16] B. N. Chun, Y. Fu, and A. Vahdat. Bootstrapping a distributed computational economy with peer-to-peer bartering. In *Proceedings of the Workshop on Economics of Peer-to-Peer Systems*, Berkeley, California, June 2003.
- [17] Computing Research Association. Grand Research Challenges in Information Security & Assurance, <http://www.cra.org/Activities/grand.challenges/security/home.html>.
- [18] A. Demers, D. Greene, C. Hauser, W. Irish, J. Larson, S. Shenker, H. Sturgis, D. Swinehart, , and D. Terry. Epidemic algorithms for replicated database maintenance. In *In Proceedings of the Sixth ACM Symp. on Principles of Distributed Computing (PODC)*, pages 1–12, August 1987.
- [19] C. Dovrolis, P. Ramanathan, and D. Moore. What do packet dispersion techniques measure? In *INFOCOM 2001*, Alaska, Apr 2001. <http://www.caida.org/outreach/papers/2001/consti/>.
- [20] D. Ellis and N. Weaver. The worm and virus threat, 2003. Presented at ACSAC 2003. Slides avialble by request.
- [21] C. Estan, S. Savage, and G. Varghese. Automated measurement of high-volume traffic clusters. In *Proceedings of the ACM/USENIX Internet Measurement Workshop (IMW)*, Marseille, France, Nov. 2002.
- [22] C. Estan, S. Savage, and G. Varghese. Automatically inferring patterns of resource consumption in network traffic. In *Proceedings of the ACM SIGCOMM Conference*, Karlsruhe, Germany, Aug. 2003.
- [23] C. Estan and G. Varghese. New directions in traffic measurement and accounting. In *Proceedings of the ACM SIGCOMM Conference*, Pittsburgh, Pennsylvania, Aug. 2002.

- [24] C. Estan, G. Varghese, and M. Fisk. Bitmap algorithms for counting active flows on high speed links. In *Proceedings of the USENIX/ACM Internet Measurement Conference*, Miami, Florida, Oct. 2003.
- [25] M. Fomenkov, k. claffy, B. Huffaker, and D. Moore. Macroscopic Internet Topology and Performance Measurements From the DNS Root Name Servers. In *Usenix LISA*, San Diego, CA, 4-7 Dec 2001. Usenix. <http://www.caida.org/outreach/papers/2001/Rssac2001a/>.
- [26] Y. Fu, J. S. Chase, B. Chun, S. Schwab, and A. Vahdat. Sharp: An architecture for secure resource peering. In *Proceedings of the 19th ACM Symposium on Operating System Principles (SOSP)*, Bolton Landing, NY, Oct. 2003.
- [27] J. M. González and V. Paxson. pktD: A packet capture and injection daemon. In *Proceedings of Passive/Active Measurement (PAM)*, 2003.
- [28] M. Handley, O. Hodson, and E. Kohler. XORP: Open platforms for network research. In *Proceedings of the 1st ACM Workshop on Hot Topics in Networks (HotNets)*, Princeton, NJ, Nov. 2002.
- [29] H. W. Hethcote. The mathematics of infectious diseases. *SIAM Review*, 42(4):599–653, 2003.
- [30] B. Huffaker, M. Fomenkov, D. Moore, and k. claffy. Macroscopic analyses of the infrastructure: measurement and visualization of Internet connectivity and performance. In *PAM 2001*, Amsterdam, Netherlands, Apr 2001. RIPE NCC. <http://www.caida.org/outreach/papers/2001/SkitViz/>.
- [31] B. Huffaker, M. Fomenkov, D. Moore, E. Nemeth, and k. claffy. Measurements of the Internet topology in the Asia-Pacific Region. In *INET '00*, Yokohama, Japan, 18-21 July 2000. The Internet Society. http://www.caida.org/outreach/papers/2000/asia_paper/.
- [32] B. Huffaker, M. Fomenkov, D. Plummer, D. Moore, and k. claffy. Distance Metrics in the Internet. In *IEEE International Telecommunications Symposium (ITS)*, Brazil, Sept 2002. IEEE. <http://www.caida.org/outreach/papers/2002/Distance/>.
- [33] B. Huffaker, D. Plummer, D. Moore, and k. claffy. Topology discovery by active probing. In *Symposium on Applications and the Internet (SAINT)*, Nara, Japan, Jan 2002. SAINT. <http://www.caida.org/outreach/papers/2002/SkitterOverview/>.
- [34] J. Jung, V. Paxson, A. W. Berger, and H. Balakrishnan. Fast portscan detection using sequential hypothesis testing. In *IEEE Symposium on Security and Privacy*, 2004.
- [35] K. Keys, D. Moore, R. Koga, E. Lagache, M. Tesch, and k. claffy. The architecture of CoralReef: an Internet traffic monitoring software suite. In *PAM 2001*, Amsterdam, Netherlands, Apr 2001. RIPE NCC. <http://www.caida.org/outreach/papers/2001/CoralArch/>.
- [36] D. Kostić, A. Rodriguez, J. Albrecht, A. Bhirud, and A. Vahdat. Using random subsets to build scalable network services. In *Proceedings of the 4th USENIX Symposium on Internet Technologies and Systems (USITS)*, Seattle, WA, Mar. 2003.
- [37] D. Kostić, A. Rodriguez, J. Albrecht, and A. Vahdat. Bullet: High bandwidth data dissemination using an overlay mesh. In *Proceedings of the 19th ACM Symposium on Operating System Principles (SOSP)*, Bolton Landing, NY, Oct. 2003.
- [38] Lurhq. Mydoom word advisory, <http://www.lurhq.com/mydoomadvisory.html>.
- [39] Lurhq. Sobig.a and the Spam You Received Today, <http://www.lurhq.com/sobig.html>.
- [40] R. Malan, D. Watson, and F. Jahanian. Transport and application protocol scrubbing. In *Proceedings of IEEE Infocom Conference*, Tel-Aviv, Isreal, Mar. 2000.
- [41] C. D. Marsan. AT&T developing early warning tool. Network World article 9/29/03, Nov. 2003.
- [42] R. R. Miller. A peek at script kiddie culture, <http://software.newsforge.com/software/04/02/28/0130209.shtml>.
- [43] D. Moore, K. Keys, R. Koga, E. Lagache, and k. claffy. CoralReef software suite as a tool for system and network administrators. In *Usenix LISA*, San Diego, CA, 4-7 Dec 2001. Usenix. <http://www.caida.org/outreach/papers/2001/CoralApps/>.
- [44] D. Moore, V. Paxson, S. Savage, C. Shannon, S. Staniford, and N. Weaver. Inside the Slammer Worm. *IEEE Magazine of Security and Privacy*, pages 33–39, July/August 2003 2003.

- [45] 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, July 2003.
- [46] D. Moore, V. Paxson, S. Savage, C. Shannon, S. Staniford, and N. Weaver. The spread of the Sapphire/Slammer worm, Jan. 2003.
- [47] D. Moore, R. Periakaruppan, J. Donohoe, and k. claffy. Where in the world is netgeo.caida.org? In *INET '00*, Yokohama, Japan, 18-21 Jul 2000. The Internet Society. http://www.caida.org/outreach/papers/2000/inet_netgeo/.
- [48] D. Moore and C. Shannon. The spread of the witty worm. <http://www.caida.org/analysis/security/witty/>.
- [49] D. Moore, C. Shannon, and J. Brown. Code-red: a case study on the spread and victims of an Internet worm. In *Proceedings of the ACM/USENIX Internet Measurement Workshop (IMW)*, Marseille, France, Nov. 2002.
- [50] D. Moore, C. Shannon, and k. claffy. Code-Red: a Case Study on the Spread and Victims of an Internet Worm. In *Proceedings of the Second Internet Measurement Workshop*, pages 273–284, November 2002.
- [51] D. Moore, C. Shannon, G. M. Voelker, and S. Savage. Internet quarantine: Requirements for containing self-propagating code. In *Proceedings of the IEEE Infocom Conference*, San Francisco, California, Apr. 2003.
- [52] D. Moore, G. Voelker, and S. Savage. Interring Internet Denial-of-Service Activity. In *Proceedings of the 10th USENIX Security Symposium*. USENIX, August 2001.
- [53] R. Pang and V. Paxson. A high-level programming environment for packet trace anonymization and transformation. In *Proceedings of ACM SIGCOMM*, 2003.
- [54] V. Paxson. Bro: a system for detecting network intruders in real-time. *Computer Networks*, 31(23–24):2435–2463, 1999.
- [55] V. Paxson, A. Adams, and M. Mathis. Experiences with NIMI. In *Proceedings of Passive/Active Measurement (PAM)*, 2000.
- [56] V. Paxson, A. Adams, and M. Mathis. Experiences with NIMI (revised). In *Proceedings of IEEE Symposium on Applications and the Internet (SAINT)*, 2002.
- [57] V. Paxson, J. Mahdavi, A. Adams, and M. Mathis. An architecture for large-scale Internet measurement. *IEEE Communications*, 36(8), August 1998.
- [58] S. Ratnasamy, P. Francis, M. Handley, R. Karp, and S. Shenker. A scalable content addressable network. In *Proceedings of ACM SIGCOMM 2001*, 2001.
- [59] E. Rescorla. Security Holes.. Who Cares? In *Proceedings of the 12th USENIX Security Symposium*, 2003.
- [60] A. Rowstron and P. Druschel. Pastry: Scalable, decentralized object location, and routing for large-scale peer-to-peer systems. In *IFIP/ACM International Conference on Distributed Systems Platforms (Middleware)*, pages 329–350, 2001.
- [61] S. Schechter and M. Smith. Access For Sale: A New Class of Worm. In *The First ACM Workshop on Rapid Malcode (WORM)*, 2003.
- [62] C. Shannon, D. Moore, and k. claffy. Characteristics of fragmented IP traffic on Internet links. In *ACM Internet Measurement Workshop 2001*, San Francisco, CA, Nov 2001. <http://www.caida.org/outreach/papers/2001/Frag/>.
- [63] C. Shannon, D. Moore, and k. claffy. Beyond Folklore: Observations on Fragmented Traffic. *To appear in IEEE/ACM Transactions on Networking*, Dec 2002. <http://www.caida.org/outreach/papers/2002/Frag/>.
- [64] T. Sherwood, E. Perelman, G. Hamerly, and B. Calder. Automatically characterizing large scale program behavior. In *Proceedings of the Tenth International Conference on Architectural Support for Programming Languages and Operating Systems*, 2002.
- [65] T. Sherwood, G. Varghese, and B. Calder. A pipelined memory architecture for high throughput network processors. In *Proceedings of the ACM International Symposium on Computer Architecture (ISCA)*, San Diego, California, June 2003.
- [66] S. Sikka and G. Varghese. Memory efficient state lookups with fast updates. In *Proceedings of the ACM SIGCOMM Conference*, Stockholm, Sweden, Aug. 2000.
- [67] S. Singh, C. Estan, G. Varghese, and S. Savage. The earlybird system for real-time detection of unknown worms. Technical Report CS2003-0761, Univerisity of California, San Diego, Aug. 2003.

- [68] R. Sommer and V. Paxson. An architecture for independent state in network intrusion detection, 2004. In submission.
- [69] V. Srinivasan, S. Suri, and G. Varghese. Packet classification using tuple space search. In *Proceedings of the ACM SIGCOMM Conference*, Cambridge, Massachusetts, Sept. 1999.
- [70] S. Staniford-Chen, S. Cheung, R. Crawford, M. Dilger, J. Frank, J. Hoagland, K. Levitt, C. Wee, R. Yip, and D. Zerkle. GrIDS – A graph-based intrusion detection system for large networks. In *Proceedings of the 19th National Information Systems Security Conference*, 1996.
- [71] I. Stoica, R. Morris, D. Karger, M. F. Kaashoek, and H. Balakrishnan. Chord: A scalable peer-to-peer lookup service for internet applications. In *Proceedings of ACM SIGCOMM 2001*, pages 149–160, 2001.
- [72] Stuart Staniford and Vern Paxson and Nicholas Weaver. How to Own the Internet in Your Spare Time. In *Proceedings of the 11th USENIX Security Symposium*. USENIX, August 2002.
- [73] Symantec. W32.blaster.worm, <http://securityresponse.symantec.com/avcenter/venc/data/w32.blaster.worm.html>.
- [74] Symantec. W32.HLLW.Doomjuice, <http://securityresponse.symantec.com/avcenter/venc/data/w32.hllw.doomjuice.html>.
- [75] Symantec. W32.Welchia.Worm, <http://securityresponse.symantec.com/avcenter/venc/data/w32.welchia.worm.html>.
- [76] N. Tuck, S. Sherwood, B. Calder, and G. Varghese. Deterministic Memory-Efficient String Matching Algorithms for Intrusion Detection. In *IEEE INFOCOM*, 2004.
- [77] J. Twycross and M. M. Williamson. Implementing and Testing a Virus Throttle. In *Proceedings of the 12th USENIX Security Symposium*. USENIX, August 2003.
- [78] A. Vahdat, J. S. Chase, R. Braynard, D. Kostić, P. Reynolds, and A. Rodriguez. Self-organizing subsets: From each according to his abilities, to each according to his needs. In *Proceedings of the International Workshop on Peer To Peer Systems (IPTPS)*, Cambridge, Massachusetts, Mar. 2002.
- [79] A. Vahdat, K. Yocum, K. Walsh, P. Mahadevan, D. Kostić, J. Chase, and D. Becker. Scalability and accuracy in a large-scale network emulator. In *Proceedings of the 5th ACM/USENIX Symposium on Operating System Design and Implementation (OSDI)*, Boston, MA, Dec. 2002.
- [80] R. van Renesse and K. P. Birman. Astrolabe: A robust and scalable technology for distributed system monitoring, management, and data mining. *ACM Transactions on Computer Systems*, 21(3), May 2003.
- [81] M. Waldvogel, G. Varghese, J. Turner, and B. Plattner. Scalable high-speed prefix matching. *ACM Transactions on Computer Systems*, 19(4):440–482, Nov. 2001.
- [82] S. Wallace, B. Calder, and D. M. Tullsen. Threaded multiple path execution. In *ISCA*, pages 238–249, 1998.
- [83] H. J. Wang, C. Guo, D. R. Simon, and A. Zugenmaier. Shield: Vulnerability-driven network filters for preventing known vulnerability exploits. ms research technical report msr-tr-03-81.
- [84] P. Warkhede, S. Suriand, and G. Varghese. Fast packet classification for two-dimensional conflict-free filters. In *Proceedings of IEEE Infocom Conference*, Anchorage, Alaska, Apr. 2001.
- [85] N. Weaver, D. Ellis, S. Staniford, and V. Paxson. Worms vs. perimeters, the case for Hard-LANs, 2004. In submission.
- [86] N. Weaver, V. Paxson, S. Staniford, and R. Cunningham. A taxonomy of computer worms. In *Proceedings of the First ACM CCS Workshop on Rapid Malcode (WORM)*, Washington, D.C., Oct. 2003.
- [87] N. Weaver, S. Staniford, and V. Paxson. Very fast containment of scanning worms. In *Proceedings of the USENIX Security Symposium*, San Diego, CA, Aug. 2004.
- [88] M. M. Williamson. Throttling Viruses: Restricting Propagation to Defeat Mobile Malicious Code. In *ACSAC*, 2002.
- [89] K. Yocum, E. Eade, J. Degeys, D. Becker, J. Chase, and A. Vahdat. Toward Scaling Network Emulation using Topology Partitioning. In *Proceedings of the 11th International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS)*, October 2003.
- [90] H. Yu and A. Vahdat. Building Replicated Internet Services Using TACT: A Toolkit for Tunable Availability and Consistency Tradeoffs. In *Proceedings of the Second International Workshop on Advanced Issues of E-Commerce and Web-based Information Systems*, June 2000.

- [91] H. Yu and A. Vahdat. Efficient Numerical Error Bounding for Replicated Network Services. In *Proceedings of the 26th International Conference on Very Large Databases (VLDB)*, September 2000.
- [92] H. Yu and A. Vahdat. Combining Generality and Practicality in a Conit-Based Continuous Consistency Model for Wide-Area Replication. In *The 21st IEEE International Conference on Distributed Computing Systems (ICDCS)*, April 2001.
- [93] H. Yu and A. Vahdat. The Costs and Limits of Availability for Replicated Services. In *Proceedings of the 18th ACM Symposium on Operating Systems Principles (SOSP)*, October 2001.
- [94] H. Yu and A. Vahdat. Minimal Replication Cost for Availability. In *Proceedings of the ACM Principles of Distributed Computing*, July 2002.
- [95] H. Yu and A. Vahdat. Consistent and automatic service regeneration. In *Proceedings of the 1st ACM/USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, San Francisco, CA, Mar. 2004.
- [96] Y. Zhang and V. Paxson. Detecting Backdoors. In *Proceedings of the 9th USENIX Security Symposium*. USENIX, August 2000.
- [97] Y. Zhang and V. Paxson. Detecting Stepping Stones. In *Proceedings of the 9th USENIX Security Symposium*. USENIX, August 2000.