NSF Proposal

# TWC: Frontier: Collaborative:
# **Beyond Technical Security: Developing an Empirical Basis for Socio-Economic Perspectives**

Stefan Savage, Geoffrey M. Voelker, James Fowler, Alex Snoeren,
Lawrence Saul, Kirill Levchenko, Erin Kenneally
*University of California, San Diego*

Vern Paxson, Mark Allman, Chris Grier, Chris Hoofnagle, Dan Klein,
Christian Kreibich, Deirdre Mulligan, Nicholas Weaver
*International Computer Sciences Institute*

Damon McCoy
*George Mason University*

September 2012 – August 2017

# 1 Project Description

Security is at once a technical property of a system and a socio-economic property of the environment in which it operates. The former property—focused on real or potential abuse resulting from unintended design or implementation vulnerabilities—encompasses the vast majority of security research and practice. Indeed, the very term "trustworthy system" implicitly encodes the notion that security is a distinct and comprehensive property of the computing artifacts we use: secure against all adversaries and user failings. But history has shown that a perspective solely focused on technical considerations becomes mired in a relentless arms race. More importantly, this perspective misses an entire half of the problem space: the human element. Fundamentally, we care about security only because adversaries are motivated to attack us and because users prove susceptible to being victimized.

Our proposal squarely focuses on addressing this other half of the problem space. We hold that while security is a phenomenon *mediated* by the technical workings of computers and networks, it is ultimately a conflict driven by economic and social issues which merit a commensurate level of scrutiny. Security is a game among actors: adversaries, defenders, and users. Together their behavior defines the shape of the threats we face, how they evolve over time, and, we argue, how they can best be addressed.

**A Socio-Economic Motivation.** The need to address the social and economic elements of security has become ever more pressing due to the *scale* of Internet activities. The last decade has seen transformative changes to the nature of Internet threats, not only quantitatively but qualitatively, due to the platform nature of the ecosystem. The low-stakes, individualistic attacks of ten years ago have become goal-driven and organized with a well-oiled efficiency. Behind this change lies a combination of factors, but key among them is how the growth of e-commerce has implicitly imparted small amounts of value in large numbers of machines. Each person or host that exposes a ready means for extracting value—almost regardless of how modest—becomes a potential victim for economically-motivated attacks. But whereas simply robbing individuals on the street for $50 each will not amount to much, the Internet removes the overhead of physicality and enables savvy, well-organized adversaries to extract millions in precisely this manner. This seemingly minor shift has seeded a bloom of economically-motivated attacks that is still expanding today, fueling and energetically monetizing the vast majority of malicious activity that ultimately touches upon all Internet users, from botnets to keyloggers, from fast-flux networks to exploit kits.

Scale does not merely drive an economic engine, but has changed the social structure for attackers as well. At the turn of the 21st century, most attackers were limited by what they could do themselves: what they knew, what they could access, and how much they could innovate. The open nature of Internet communications, however, soon allowed like-minded individuals to effectively gather and coordinate, meeting to exchange ideas and eventually to create markets [58]. Today, the markets for abusive goods and services overflow with options. Miscreants can purchase the installation of arbitrary malware on compromised hosts on demand for $8 per thousand [31]. New malware is bought and sold, underground consultants ply their clients with dedicated support, third-world labor is harnessed to address tasks either menial or too difficult to solve by computer [122, 123], exploit "kits" compete on efficacy and the quality of their user interface, specialized "bulletproof" services offer to host abusive content protected against takedown, and so on.

These activities have become possible because the social network of cyber-attackers has grown sufficiently large and organized to support specialization and stratification. Freed of needing to do everything "soup-to-nuts", each actor can instead focus on perfecting one element. Infrastructure actors can make the best peer-to-peer bulletproof hosting service or the easiest-to-use info-stealing malware, while financial credential thieves can focus on developing the most robust "cash out" networks. Moreover, while much of this activity is ultimately funded by economically-motivated attacks (e.g., such as click fraud [116], fake AV [152], spam [99] or outright financial theft [153]), the underlying infrastructure ecosystem is available to the full spectrum of attackers regardless of motivation. Commercial criminals, hacktivists, corporate spies,

and nation-state actors can all leverage the underground ecosystem, reaping the fruits of its COTS-like efficiency. These actors can furthermore obtain cover for their activities when necessary by appearing to simply be "yet another crook" out to steal a profit.

More broadly, the socio-economic lens is not merely useful for understanding attackers, but victims as well. Today's consumers, corporations, and governments alike make large investments in security technology with little understanding of their ultimate return-on-investment. By rote, we tell consumers to patch their systems, use a firewall and run anti-virus software, but we lack any incisive ability to quantify how well this multi-billion dollar cost actually helps protect them. How should we intelligently guide our investments to ensure effective, affordable outcomes? Can we place computer security on the same *empirical evidence basis* as is happening in health care?

We likewise understand little about the role played by the social nature of victims, although clearly it has growing significance. Society is based on certain degrees of mutual trust, and people have developed cues and protocols that guide how much faith we put in others in our daily life. Here again, however, scale upends the assumptions underlying these precepts. Today, the ease with which we adopt online personas and relationships has created a collective blind spot that attackers find all-too-easy to exploit. Is your online friend really a friend? When can you count on them? Are they even a person? The first generation of "social attacks" is predictably oriented towards advertising or malware distribution [154, 155, 159], but a variety of recent work suggests that far more complex deceptions loom [69, 96].

Even putting aside attacks, online social networks represent a powerful driver for shaping behavior. As in the physical world, one's online social group can influence how we make decisions regarding what to do and where to go online, how to behave and what to allow. While there is little research even quantifying which online behaviors lead to bad security outcomes, it seems reasonable to ask the extent to which unsafe behaviors and practices are in fact learned via social networks—and, in turn, the extent to which we might employ the power of social networks to shape behavior as a means for staunching this spread.

In summary, we hold as the central tenet of our proposed effort that grappling with these socio-economic dimensions is of fundamental importance for achieving a secure future information infrastructure. Rather than security research remaining focused on the latest technical consideration-of-the-day and the accompanying "arms race" dynamics, we need new perspectives of security as viewed through the lens of the large-scale drivers, motivators, needs and activities. Many of the ideas we frame have a speculative quality, since much is left to be understood. But this opacity is precisely why such research is *crucially needed*. As such, the scope of our proposal is broad and focuses on the plethora of opportunities for meaningful impact, rather than an in-depth examination of any one particular approach.

**An Empirical Approach.** We view it as fundamental that developing a sound understanding of these socio-economic dimensions requires research grounded in *empiricism*. Experience has repeatedly highlighted the vital role that measurement plays when investigating phenomena-of-scale, which almost by definition do not conform with abstract expectations. In addition, we highlight the potential for empirical exploration to illuminate major new defensive possibilities that can pay outsize dividends in the struggle to keep adversaries at bay. Case in point, in recent work using this approach we identified a *payment processing* bottleneck in the abusive advertising ecosystem (e.g., spam email, "blackhat" search engine optimization, blog abuse). This finding, previously undocumented and unrecognized, has swiftly become a highly successful commercial intervention (under the auspices of the International Anti-Counterfeiting Coalition and a range of brand-holders) [70]. We seek to generalize this sort of approach to consideration of other economic and social factors. While we necessarily cannot at this point specify precisely how, our extensive past experiences make us strongly confident that further work in this regard holds similar potential for impact.

To this end, our proposal is one of empirical assessment and evaluation. We bring to this problem a team with decades of experience in computer security research grounded in methodical measurement, machine learning as applied to very large datasets, natural language processing for inferring roles of and

relationships among actors, analysis of social structures as manifest in online social networks, and legal, policy, and privacy perspectives. We seek to understand, through measurement and analysis, the shape of key economic and social forces—as seen at scale—in terms of relevance for both attackers and defenders. In particular, our proposed effort has four key components:

- *Economics of E-crime*. Most of today's Internet attacks are profit-driven, which has come to require complex enterprises with many moving pieces. We propose to pursue in-depth empirical analyses of a range of online criminal activities, including both activities that are themselves profit-making (e.g., financial credentials theft, advertising) as well as the broad range of enabling infrastructure activities that involve "sale to the trade" (e.g., underground proxies, search engine abuse, account theft, pay-per-install markets).

- *The Role of Online Social Networks*. OSNs such as Facebook and Twitter provide a broad new "social attack surface" for exploiting users' trust to a variety of ends. We will empirically map out both the attacker ecosystem that preys on this environment, and the extent to which unsafe online behavior is itself learned and transmitted via this same vehicle. An especially important new threat in this sphere concerns the use of social networks by adversaries, not for monetary gain, but to manipulate public opinion or implement *de facto* censorship.

- *Underground Social Networks*. We have scant understanding of the implicit social networks that link today's attackers, yet they play a critical role in fostering the pace of innovation and the efficiency of the cybercrime market economy. Drawing upon disparate data sources, we will study the nature of "trust amongst thieves" and map out how relationships among these criminals are established, maintained, and evolve over time within this medium. We will also analyze in depth how particular actors advanced within the echelons, with a goal of illuminating what enabled their success and identifying common patterns in the development of trust and cooperation. For example, underground data we have acquired through a collaboration partner (Krebs) illuminates how major botmasters such as "Cosma2k" (*Rustock*) and "Docent" (*MegaD*) established their reputations over time on spammer forums, parlaying that notoriety into higher commissions at competing affiliate programs.

- *Efficacy of Intervention*. Where the insights developed in the other elements of our research agenda come to fruition is the degree to which they allow us to develop more effective security interventions. To this end we emphasize outcome-based approaches rooted in empiricism. First, we will embark on large-scale studies to measure the efficacy of today's security interventions, both in the large (e.g., botnet takedown, domain blacklisting) and at the level of individual users (e.g., anti-virus, firewall, patching, "safe" online behaviors). Second, drawing upon our understanding of the economic and social factors that underlie online attacks, we will design analyses to identify bottleneck elements: where an intervention might most effectively undermine an entire group of abusive activities. Finally, we will explore how user choices (security hygiene, risky behaviors) affect security outcomes, and the possibilities for shaping these in novel positive ways.

As we have formulated and initially pursued this research agenda, our team has developed a set of key capabilities that we will draw upon for our proposed effort. These include: an industrial strength "bot-farm" for executing malware that interacts with remote Internet systems in a controlled, safe fashion [88]; the ability to reverse-engineer command-and-control protocols to support *infiltration* of large-scale malware systems [31, 39, 76]; high-fidelity (full browser execution) large-scale Web crawling and screen scraping [19, 99, 154]; and a framework for legal and ethical purchasing of underground services and products, with accompanying tracking of payment processing [99]. In concert with these capabilities, we have also over the course of our previous efforts acquired extensive, and in some cases arguably unique, large-scale sources of data that we will employ in this new effort. We frame these in our *Data Management Plan*.

## 2 The Economics of E-Crime

Beginning around the middle of the last decade, arguably the single most significant change in the cybersecurity landscape has been the explosive rise of profit-fueled *cybercrime*, i.e., the onset of Internet attacks ultimately motivated by making money. In sheer scale, such attacks now dwarf those of any other nature. This change has grave implications for defenders because it has brought with it both (1) an acceleration of the unending arms race between attackers and defenders, and (2) an efficient, specialized underground economy that greatly enhances the ease with which miscreants—of any flavor—can enter into and effectively exploit the cybercrime ecosystem.

This shift in the nature of cyberattacks opens up a new front in terms of possible countermeasures that defenders can employ. Obtaining a detailed, structural understanding of how profit-driven cybercrime works—i.e., its underlying economics and associated business practices—has the potential to illuminate new defensive measures that might prove far more efficient for countering attacks than purely technical measures. Put another way, we might find it much more effective to disrupt the ability of attackers to profit from Internet attacks, than it is to pointwise defend against the huge number of ways that our systems and users repeatedly prove themselves vulnerable. If we can undermine the profitability of conducting cyberattacks, we can remove the primary driver that underlies them today.

However, developing accurate models of the cybercrime ecosystem, and then from these devising effective defensive interventions, requires both a wealth of technical capability and nimble, often opportunistic, execution of measurement studies. Our team brings unique experience and skills in this regard. In past work we have undertaken the first measurement studies of: goods and services bartered on underground markets [58], the response and conversion rates associated with email spam [76], freelance labor and CAPTCHA-solving in support of service abuse [122, 123], the revenue of pharmaceuticals advertised via email spam [78], underground services that install arbitrary malware on newly compromised systems [31], and a "soup-to-nuts" analysis of the infrastructure required to monetize the sale of fake products via spam [99].

While these studies represent significant undertakings in terms of time and resources, in truth we have only just come to understand the full scope of the underground economy, and there remain far more questions than answers. Informed by this experience, we have identified three primary classes of underground activities that guide our measurement activities going forward: *advertising*, *theft*, and *infrastructure*. The first two, advertising and theft, describe the two dominant value strategies for bringing new capital to the underground; advertising-based schemes require a certain degree of consumer participation to monetize (e.g., buying goods via credit card), while theft involves direct monetization of user data (e.g., bank account credentials). However, much of the underground economy is in the third leg of this triad: infrastructure. Infrastructure refers to the plethora of goods and services (e.g., malware distribution, proxies, phishing kits, botnets, marketplaces) that carry no external value outside the underground economy, but are essential to its efficient functioning; these are the "cost centers" focused entirely on sale to underground participants.

Among these three domains, we have by far the most experience on the advertising front, with a very heavy skew towards email spam advertising [19, 85, 86, 99]. However, email spam is only one activity among many. We have also recently started exploring the characteristics of "social spam" on Twitter, and identified qualitative differences between it and email spam, including indications that it is conducted by a different set of actors [63] using different business models (e.g., exploiting affiliate marketing programs of well-known sites such as Amazon [155]). Based on this experience we seek to revisit our past analyses but across the full range of advertising vectors, both established (e.g., "blackhat SEO" [102], forum/blog spam, abuse of legitimate advertising channels) and emerging (e.g., Facebook, Google+, Twitter). We have experience that the nature of these vectors can play a large role in the class of activities and business models that they engender. For example, we have identified a large underground eco-system that uses "black hat" SEO techniques to capture clicks from search engines and resell this organic search traffic to miscreants for

their scams. However, this "stolen" traffic is itself differentiated by keyword (much like Google's Adsense), which in turn impacts the kinds of business models that make sense. Undifferentiated "popular" traffic has minimal value for advertising particular goods (e.g., counterfeit software) but, being cheap, is very attractive for those engaged in "surprise-based" scams such as fake anti-virus or pure drive-by malware distribution. Our work will generalize these kinds of findings and construct a full taxonomy of advertising abuse vectors, highlighting which of their characteristics dictate their underlying business value to cybercriminals.

As stated earlier, the other major mechanism by which value flows *into* the underground ecosystem is via theft. Here the nature of the abuse has two components, inflow and outflow; acquiring user financial credentials (e.g., via info-stealing software, phishing, e-commerce site compromise) and then extracting that value and laundering it. To study inflow, we will systematically leak "honeytokens"—well-known financial credentials—at scale into a variety of theft vectors including phishing sites, info stealers (e.g., Zeus/Spyeye), and underground forums. Through partnerships with several financial services companies, we now have the ability to track the use of these honeytoken credentials after they have been leaked. Thus, we can measure the latency of different channels from time-of-theft to time-of-first-use, and identify signatures of approaches by which distinct miscreants test account validity and attempt to cash out value. We will additionally study this outflow by measuring the recruitment of *money mules* (patsies used to launder funds) by pseudonymously responding to solicitations that "mulers" send out via spam and on job-posting sites. In this manner we can measure the scope of such laundering, the number of actors specializing in it, and their segmentation across different areas.

Finally, the infrastructure domain includes a vast number of components that we plan to explore, including the anonymization layer of proxy networks, the competitive market of exploit kits (which encode standard Web browser exploits into slick, flexible packages), anti-virus testing services (for validating the non-detectability of newly developed malware), and so on. Of these, one that holds great interest is the Pay-Per-Install (PPI) ecosystem [31]. PPI describes a business model whereby retail customers pay a third-party service to install their malware on compromised Internet hosts (for prices as low as $5 per 1k in Asia and $100 per 1k in the US), and the service in turn pays independent contractors to compromise hosts to fulfill their retail demand. In a short period of time the PPI model appears to have transformed and commoditized malware distribution (indeed, we have evidence that most, if not all, of the Rustock botnet was deployed in just such a manner). Since malware distribution in fact provides the foundation upon which the entire underground platform economy operates, this development is deeply intriguing.

Moreover, our previously-developed ability to infiltrate PPI programs provides an unprecedented opportunity to measure these distribution systems from the *inside*. By operating emulators that mimic the protocol used to download malware onto newly compromised systems, we in essence obtain a real-time feed of new e-crime activity: all major fresh malware in support of the full spectrum of new cybercrime enterprises. We will develop analysis capabilities for automatically classifying the kinds of malware being distributed, to what locales, and for what purposes. Further, we will investigate the factors that drive the pricing of PPI installations, as seen by both affiliate members (those who compromise systems) and customers (those who pay to disseminate malware). We are particularly interested in exploring the following key questions: (1) how does the interplay between supply and demand affect pricing, (2) to what extent does lack of transparency fuel both oversubscription and arbitrage between PPI programs, (3) why do we observe huge (20×) geographic diversity in pricing (perhaps reflecting ease of compromise? monetization potential?), and (4) what is the potential for pricing to serve as a *metric* for measuring both the market impact of significant events (e.g., botnet takedowns), and, more broadly, future progress (or setbacks) in securing Internet systems as seen in aggregate. We discuss this last point further in Section 5.

## 3 The Role of Online Social Networks

In recent years, online social networks (OSNs) have fundamentally transformed how people use the Internet. Today, OSNs such as Facebook and Twitter attract nearly one billion people from around the world, enabling users to share their lives through photos, stories and links to Web content. Implicit to the interactions within a social network is the notion of trust; users form relationships with their friends and valued media outlets, in turn receiving access to content generated by each relationship. While the capacity to connect people has created numerous positive opportunities for collaboration and social change, it simultaneously exposes people to an enormous "social attack surface" that continues to grow.

At the root of this problem are poor defense mechanisms that fail to protect people while they engage in online social interactions. People are all too willing to publicly expose sensitive personal details, form relationships with untrusted users, and act on potentially harmful messages they receive. The naivete of users towards the safety of OSNs makes them an idea channel for affecting social control and abuse. The first generation of such abuse targets transient trust. This *social spam* uses OSNs as a vehicle for direct advertising through the formation of social relationships. A more recent variant, which we call *social shilling*, seeks to indirectly influence the market by manipulating reviews and opinions for specific products or services rather than attracting sales through clicks. Finally, a third type of OSN abuse, *social control*, is not economically motivated at all and instead focuses on promoting an ideological or political agenda while suppressing opposing views.

Our group has extensive experience investigating these issues as they have emerged. We have identified a broad ecosystem enabling the purchase and sale of OSN accounts and "friends" [123] and shown that these marketplaces enable very large-scale campaigns at low-cost and with great resilience to intervention [155]. We have shown that OSNs have far better success rates than other sources of advertising (in part due to the implicit trust in social relationships) [63] and that social spammers engage in a business ecosystem that overlaps, but is distinct from, that of spam-email and search engine optimization [154]. We have also developed large-scale machine learning approaches for dynamically identifying and filtering such abuse, which has not only guided the selection of features that identify spam, but has quantitatively demonstrated that social spam is distinct in nature and scope from other advertising vectors [107, 154].

Our work has only begun to highlight the potential of OSN-based threats. As part of this proposal, we plan several major investigations in this space: account abuse, social spam, social shilling and, looking forward, how to identify and manage the risks of OSN-based social control.

The first of these efforts is to fully map out the enabling components that underlie OSN abuse. To commit abuse *at scale* requires large numbers of accounts and fraudulent personas. There are two methods of acquiring such accounts: wholesale new account creation and mass account theft. With a retail market that offers a thousand Web mail accounts for $8 and as many Facebook accounts for $20, the low price belies a complex ecosystem and specialized services.

In prior work we documented how ad-driven online services (such as OSNs and Web mail) created a market for retail CAPTCHA-solving services driven by cheap outsourced human labor [122]. The act of registration is simply the beginning of the journey a new account faces. Accounts are created via broad Internet proxy networks (so their IP addresses appear from different locations and from within the appropriate country), third-party services are employed to provide phone validation (in turn fueling a secondary market for toll fraud abuse), the accounts are carefully crafted to create the semblance of reality, and social relationships are formed or purchased, only then are accounts are ready for sale (and sometimes sold multiple times). By contrast, stolen accounts are far simpler to use because their background is already assured; however, they in turn require matching those attackers with access to credentials from compromised online Web site "dumps" or those with keylogger data from Zeus/Spyeye malware with buyers of particular properties.

We propose to fully map this ecosystem, the actors within each part, the underlying economics and where the key dependencies lie. Our active collaborations aid us in this work through key industry partnerships

that provide data and insight around particular areas of their concern. Combined with the legally-cleared capabilities we have developed to directly engage with online criminals and purchase accounts and support services, we have a unique opportunity to measure account abuse from the standpoint of the victim, the provider and the criminal—sometimes all at once. For example, the cost structure of a purchased account is a complicated matter that balances the price paid and return on investment. For instance, the income an account can generate relates to the maximal number of messages sent before termination (itself frequently a function of message rate and targeting), the conversion rate per message and the marginal revenue per sale. To complicate matters further, externalities such as the number of times the account has been previously sold and the opportunity cost versus other vectors also impact the cost of an account. We believe we are in a solid position to dissect these processes and provide each component with a precise empirical footing. This in turn defines the space of activities that will attract compromised accounts and provides a framework to reason about the efficacy of different interventions (e.g., does it matter if fraudulent account termination happens twice as fast, or not?). With our industry partners, we can observe how the market and corresponding technologies react to interventions and assess the mobility of different actors in this space.

The use of OSN accounts in a multiplicity of ways is implicit in the above discussion. From our own prior work it is clear that they intersect with a disparate set of resources, advertising networks, malware sites and monetization strategies than are present in email spam. Not all social network abuse focuses on getting a "click". We have also started to witness online social networks used to *indirectly* influence markets via shilling. Online reviews, product touts in tweets or on Facebook, anti-feedback for competing products (so-called astroturfing) are currently used at scale to promote particular market goals, typically without any disclosure of their financial interest in the outcome (or their fundamental non-personhood). Ultimately, understanding how these business models work is essential for identifying the appropriate intervention.

For both social spam and social shilling, we benefit from the large-scale crawling infrastructure we have developed over several years of past work [19], allowing us to pervasively capture activity on large portions of major OSNs [63, 154], visit the embedded links and referrals in real-time and then classify the site promoted [99, 104, 106, 107, 154, 155]. However, while social spam can benefit directly from the sponsorship classification methodology we developed [99], shilling is a different problem. To disambiguate independent third-party commentary from coordinated fraud will require new tools focused on behavioral features in the data. For example, we see evidence that many of today's astroturfers have little cover traffic outside their message, are highly bursty and have minimal message diversity. Similarly, "fake" users frequently exhibit very different social network structures from real users [155, 168]. We plan to explore the best way to incorporate such features into online machine learning tools as we have in the past [107, 154] and validate against the third-party data we have access to.

Finally, members of our team have also explored how OSNs can effect larger social influence. In what is undoubtedly the largest randomized control trial of conditioned political engagement to date, we worked with Facebook to present voting encouragement messages to over 60 million users during the last congressional election. While prompts indeed impacted turnout, we were able to demonstrate that the effect of a user's social network, e.g., by letting people know their friends had voted, was far larger and accounted for over 2 million voters mobilized who would not have otherwise participated in the 2010 election [28].

However, while improving enfranchisement is an admirable goal, it is clear that this same approach also applies in far more negative ways. Even in the same domain, one might try to suppress turnout in one region and increase it in another to sway an election. Indeed, in the recent December elections in Russia, it was widely documented that online protesters communicating via Twitter were targeted by bots that overwhelmed key channels (i.e., hashtags) with misinformation. Due to our measurement infrastructure, we were able to capture the activity of the bots (roughly 25,000 accounts by our measure) and have discovered that they were not newly created, but in fact belong to a far larger botnet (roughly 1M Twitter accounts) that includes accounts used for commercially-motivated social spam and a pool of accounts that remain dormant in preparation for attack. This activity is yet more evidence that the underground ecosystem is available to

the full spectrum of actors.

We propose to build on these recent experiences, identify online social conflict zones and map how they obtain and develop their infrastructure, how it differs from commercially-motivated actors and try to infer any linkage between commercial and socio-political groups. This aspect of our work is the most long term, but for the same reason we believe it has the potential for the longest-term impact. If online social networks continue to thrive, it will be important that we are able to distinguish between true online democracy and mass action, and its Potemkin shell.

# 4 Underground Social Networks

Today, a few dozen underground forums provide the social hub for criminal activity, attracting new members and enforcing an evolving social hierarchy. These environments—steeped in anonymity and mutual distrust—create new, complex social dynamics around reputation, interactions, and rivalry [124]. Invisible to most of the cybersecurity research community, this milieu is where much of adversarial innovation occurs.

Underground economies exist, however, in a fundamental tension between being open and closed. On the one hand, open markets are a basic tenet of economic efficiency—larger markets mean greater reach and profits. On the other hand, the illicit nature of the goods and services being traded (e.g., stolen credit card numbers, hijacked accounts, compromised hosts) artificially limits the market because each advertisement or transaction exposes both buyer and seller to legal risk. A risk that is both real and palpable: after the well-publicized FBI sting of the DarkMarket forum (a market for stolen credit card numbers) [143], many forums went underground, fearing infiltration by law enforcement and informants.

Even in the absence of such external pressure, the most important factor limiting the underground economy remains *trust*. Trust is a necessary component of any business transaction, giving both parties confidence that the other will uphold its end of the bargain. Legitimate businesses derive trust from established institutions—banks, regulatory agencies, courts, and the like—allowing them to operate efficiently. In the underground economy, no courts exists to adjudicate on the terms of a contract or prevent outright theft.

Absent such institutions, buyers and sellers must establish trust explicitly. In such contexts, prior interactions carry much weight in establishing trust. The desire to grow and expand, however, requires forming new relationships, the formation of which is mediated by the *social network* of the participants. Joining key forums and programs already requires referrals from existing members, and in many cases also an instant messaging "interview."

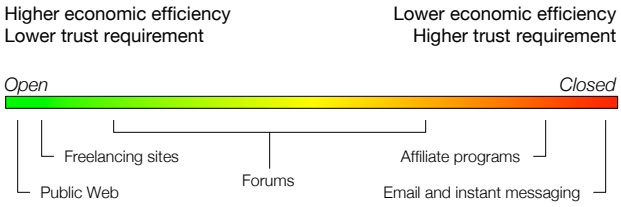


Figure 1: Openness allows participants to reach larger markets, but also exposes them to greater risk.

Broadly speaking, then, underground markets exist on a spectrum ranging from fully public settings like open forums or freelancing sites, to closed interactions via email and instant messages (Figure 1). The nature of the markets that arises in each context thus depends on the goods and services being traded (which in turn largely determine the available trust-enabling institutions), as well as the social network of the participants. How these factors—social networks and institutions of trust—shape the resulting markets is the central question we aim to answer.

The markets and interactions that arise at different ends of the open-closed spectrum do vary considerably. Some of the earliest underground markets took place over Internet Relay Chat, and our study of those markets helped pioneer the analysis of such activities [58]. More recently, we studied one of the largest freelancing sites (Freelancer.com), which provides not only a forum for employers to find workers for Web-outsourceable work (e.g., site design and programming), but also a reputation mechanism and escrow service. Because of its large and public nature (putting it at the "open" end of the spectrum), most of the jobs posted are benign; however, we found almost a third were direct precursors to Internet service abuse [95, 123]. Such large, open markets allow the building blocks of complex cybercrime to be produced and assembled in the open, leveraging the availability of cheap skilled labor.

Underground Internet forums lie in the middle of the open-closed spectrum. While some are public and easy to join, closed forums usually require an invitation from a current member or the forum operator. Our recent study of six underground forums (L33tCrew, HackSector, FreeHack, Cards, BlackHatWorld, and HackeL1te) [124] sheds light on their social dynamics. In addition to identifying the goods being traded, we also studied the effect of the underlying social network on trade. For example, are users be more inclined to do business with someone well-connected in the social network? The answer is a strong Yes. Having a social degree of 90 or greater generated $3\times$ as many responses to trade offers than a degree of 10 or less.

At the "closed" end of the spectrum, we have also recently undertaken an analysis of the operation of several large counterfeit pharmaceutical affiliate programs[1] using data leaked to the operational security community. This data includes internal database dumps of the affiliate programs (including all customer and affiliate transactions), as well as years of chat logs of the program operators. Our initial analysis focused on just these databases, from which we were able to reconstruct the internal workings of the organizations, isolate the activities of their independent advertisers (many of whom we have been able to tie directly to major botnets), and infer a detailed balance sheet for the whole business grossing over $50 million per year.

Mining individual data sources in this manner holds great promise for enabling us to identify the emergence of new classes of goods and services, and the evolution of business models. The real gold, however, comes from the capability to link services and individuals across *multiple* data sources. To this end, our overarching goal is to ultimately understand: the role of trust in underground economies, the role of social networks in developing trust and facilitating business transactions, and the points in underground social networks that provide the most potent opportunities for disruption?

Pursuing these questions requires automating the analysis of tremendous amounts of unstructured and semi-structured information: What is being discussed? By whom? How does it relate to other information about that actor or that actor's aliases? Indeed, the very anonymity of the underground creates unique challenges, as actors frequently operate using multiple identities. This same challenge, however, exemplifies one of the advantages afforded by linking data across data sources. Figure 2 shows a real example, constructed by analyzing disparate data sources in our possession, illustrating how we can draw upon a variety of data to identify the authors of major botnets. (All references in text are to parts of the figure.)

Over the course of the past two years, we have collected (❶) and visited (❷) the sites advertised in billions of spam messages we received via data feeds from industry partners [99]. In addition, we perform the same processing on spam generated using malware in our own "botfarm", a constrained environment in which we execute instances of key major botnets [88]. This capability allowed us to identify which botnets (including Rustock, in this instance) advertise particular domain names in their messages. Using a leaked affiliate program database provided by one of our collaborators (❸), we confirmed that the domain names being spammed by Rustock were associated with the program, and in particular commissions for visitors to those sites were awarded to an affiliate using the handle "Cosma2k". We determined that over the

[1]Today, a spammer wishing to sell counterfeit pharmaceuticals joins an *affiliate program*. The spammer takes care of advertising links to the store site, while the program handles store operation, payment processing, and product delivery. Affiliate programs have revolutionized modern spamming and search engine manipulation (so-called "blackhat SEO"). Affiliate programs thus act as an enabling mechanism for much of today's cyber crime and Internet service abuse.

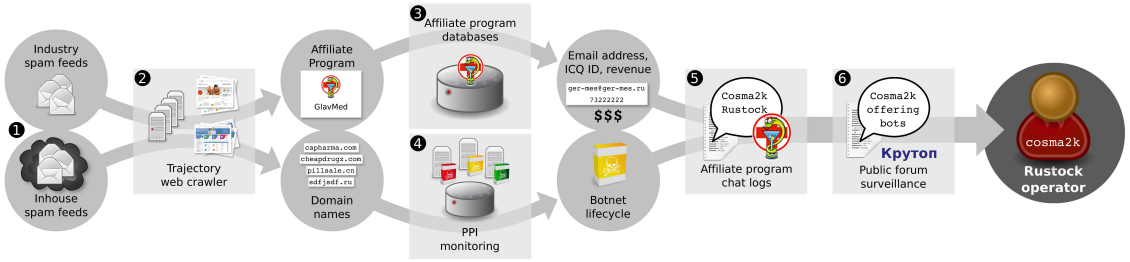


Figure 2: Linking information from multiple sources allows us to identify the operators of several major botnets.

course of 4 years, Cosma2k and his aliases earned over $2M in such commissions from this single program. (Through our PPI malware feed (❹), we amassed evidence that the Rustock botmaster was the main client of the LoaderAdv/GangstaBucks pay-per-install program.) In addition, a range of corroborating information about the affiliate appeared in leaked chat logs of operators of the same affiliate program (❺), including Cosma2k's competition with other botmasters for better commissions. Finally, mapping the set of aliases used by this affiliate (which share an ICQ instant messaging identifier), we can link the same actor with another individual offering to buy compromised hosts on a popular underground forum (❻).

While this sort of data mining can yield invaluable information about an individual's role in the cybercrime underground, as described above, it is a highly labor-intensive manual process. The central challenge is to automate this process so we can pursue such analysis at much larger scale. Our experience has shown that language-independent machine learning approaches can provide good results for identifying coarse features such as topic modeling [95], but identifying finer-grained structure, such as to whom an ICQ instant messaging service identifier refers to in a given conversation, requires much more powerful techniques. In particular, this problem of "coreference resolution" (a difficult problem in its own right) is here compounded by the use of domain-specific slang as well as the highly context-sensitive nature of informal communication such as instant messaging.

Given these considerations, for our goal of automating both the *extraction* of such data as well as constructing inferences like those in our earlier example, we plan to adapt to this domain techniques our team has recently developed for automatic acquisition and modeling of knowledge for coreference [67]. The initial elements of this work (during Years 1 and 2) will proceed separately from this project, funded under the auspices of our MURI project *Infiltration of Botnet Command-&-Control and Support Ecosystems*. Further out (starting in Year 3), however, we will then need to push beyond coreference analysis. While coreference is mainly about entities, information extraction in contexts such as chat logs is often instead concerned with *events*. In particular, we will need to enrich our models to capture the structure of events, with their own attributes and properties, in which entities participate according to their domain roles. Our team's recent work in this domain suggests that such a unified approach can prove viable [66].

The resulting database will combine a vast array of raw underground data sources—both unstructured conversational data and structured data from network sources, active experiments, and third-party databases—along with NLP-based algorithms to extract key identifiers, prices and topics, enabling us to link actors based on how much and in what ways they interface. This repository will provide a unique resource for understanding cybercriminal social networks, including how relations evolve between actors, how new cyberattack strategies propagate, and how vulnerable they are to disruption of reputation or trust.

## 5 Efficacy of Intervention

Those tasked with securing operational systems have long recognized the aptness of the adage "*security is economics*." That is, security in practice focuses not on achieving airtight defenses regardless of cost, but rather determining efficient ways to allocate limited resources. This perspective places a premium on under-

standing which defensive measures reap the most benefit. Unfortunately, today we rarely have an informed, evidence-based foundation for making such decisions, and security investments often have unknown value. For example, anti-malware products and services constitute a multi-billion dollar industry, driven by standard IT security practices that require hosts to install $30–$50 AV software—yet access to a compromised host sells on the underground for mere pennies [31]. Such disparities raise the question of whether we are spending our money and resources wisely. We argue that to achieve effective *and* cost-effective defenses, approaches to intervention and the evaluation of their efficacy must both fundamentally evolve. We will pursue three main avenues of research in this regard, as follows.

### 5.1 Attacker Metrics

Security has proven notoriously challenging to quantify compared to traditional performance evaluation [169]. Approaches that focus on defense-oriented metrics, such as proportion of spam filtered or number of malware executables found, leave unanswered the larger question: is *enough* spam being filtered or malware discovered to undermine the overall goals of attackers? Answering this strategic question requires a holistic view of how specific countermeasures affect the overall underground enterprise.

For example, in a previous study infiltrating spamming botnets, we measured the impact of various spam defenses from the perspective of the spammer [76]. Impressively, we found that free Web mail providers filtered 99.99% of spam sent to their users. More impressively, however, we found that the sheer volume at which the botnet spammed still enabled the spammer to receive well over $1 million in commissions each year. Thus, while spam filtering clearly has a crucial role for keeping email usable, spam filtering itself will not suffice to stop spam. (Indeed, it spurs spammers to increase volume and hone their evasive techniques.)

We advocate for instead considering more directly relevant metrics by framing them in terms of how *attackers* perceive how well their enterprises are working. As a touchstone, a natural attacker metric is the prices for underground goods and services. For instance, if one of our interventions successfully disrupts the supply of compromised hosts into the market, then the price-per-host should rise to reflect the increased scarcity. If we make spam less effective, then we should see (1) prices for third-party spamming services increase, (2) spammers dropping out of the market when costs exceed their profit-making potential, and/or (3) increased prices of goods at spam-advertised sites to compensate for fewer customers. Similarly, if we succeed in disrupting trust among specialists in the underground social networks, then we should observe increased prices offered for specialized goods and services due to increased "friction" in the market.

Since the prices of underground goods and services directly reflect the difficulties imposed by security defenses, we can obtain system-wide visibility into the effects of interventions by building infrastructure to monitor changes in price and supply. As a concrete example, we recently directly measured the impact of Google's improved account validation, which now requires the use of SMS challenges in certain circumstances. Over a very short period, the cost of fresh Google accounts on the underground rose from $8 per thousand to over $200 per thousand. While attackers subsequently adapted by partnering with elements of the toll-fraud ecosystem, the price has remained five times higher than similar accounts for other services that lack this intervention. Thus, we can directly assess the benefit—in terms of economic friction—that this approach has on reducing account creation abuse.

Given our existing data access and crawling capabilities, we are in position to weave together data from a broad variety of sources, including underground social forums [124], IRC marketplaces [58], Web trading sites, and freelance labor markets [123] to create an encyclopedic scoreboard of the underground economy. Such a database can form the yardstick by which we evaluate interventions, not only our own but also opportunistically due to actions from other security researchers, industry (e.g., botnet takedowns by Microsoft, privacy policy changes by Facebook), law enforcement (e.g., FBI's seizure of Megaupload's hosting service), and government (e.g., if legislation like SOPA is enacted).

## 5.2 Bottleneck Intervention

In addition to measuring the efficacy of interventions in a more meaningful way, we need to reconsider the basic nature of interventions. Traditional interventions target specific activities, such as propagating malware or sending spam, that at a higher level represent only individual components of a larger ecosystem that both blends and depends upon many interrelated activities and actors. We argue for instead pursuing much more potent interventions of a fundamentally broader scope, where we aim to undermine entire sectors of the ecosystem. To do so, we first need to understand how the ecosystem functions by developing comprehensive models that identify infrastructure components, how they depend upon each other, the actors behind the components, and the relationships among them. From such models we can then identify structural bottlenecks within the infrastructure and among the relationships, and target precisely those weaknesses.

Our recent work on the spam ecosystem exemplifies this approach to intervention [99]. We pursued a holistic analysis to quantify the full set of resources employed to monetize spam email by selling fake products—including naming, hosting, payment, and fulfillment—by capturing extensive measurements of three months of diverse spam data, broad crawling of naming and hosting infrastructures, and hundreds of purchases from spam-advertised sites. We then related these resources to the organizations that administer them to characterize the relative prospects for defensive interventions at each link in the spam value chain.

This effort provided the first strong evidence of *payment bottlenecks* in the spam value chain: 95% of spam-advertised pharmaceutical, replica and software products are monetized using merchant services from just a handful of banks. In part on the basis of this work, the Internet Anti-Counterfeiting Coalition (IACC), in conjunction brand holders and major payment card networks (and the US Office of Intellectual Property Enforcement), is implementing precisely this bottleneck intervention: shutting down merchant accounts used to process payment for online counterfeit sites to *demonetize* the underlying enterprise [70].

We collaborated closely with many of the key parties in this activity and are now poised to conduct a major empirical study of this payment-level intervention: how well it works, how adversaries react, how demonetization flows down through the market, and the challenges the intervention poses for use as a general policy instrument. Further, we can take advantage of a unique vantage point for conducting such a study: we are one of the few organizations able to reliably identify affiliate program structure (per [99]); we can view banking relationships in near-real-time due to our partnership with multiple credit card issuers; and for many of these organizations, we have the ability to witness the impact of these interventions from *within* due to our prior field work [74].

More generally, we aim to apply this type of analysis, suitably tailored, to a broad range of security domains, including fake anti-virus [152], click fraud [116], and pay-per-install services [31]. For example, we have anecdotal experience that online financial fraud gangs (e.g., credit card theft or ACH fraud) are ultimately limited not by access to money, but by the availability of "cash out" bandwidth, typically in the form of organized *money mules*. We will undertake a large-scale tracking study of stolen financial credentials in collaboration with a major financial services organization to identify the depth and breadth of such cash-out networks. In another dimension, the "anonymity layer" provided by underground proxy services appears to likely represent a key bottleneck for attribution, and thus deterrence. We will pursue a major effort to identify the number, scope, and popularity of such services to determine the most effective interventions for improving miscreant attribution.

## 5.3 Data-Driven Security Behavior

Looking more broadly, we observe that many security outcomes depend fundamentally on imperfect or unsafe actions taken by users. Indeed, a large literature documents individual ways in which user behavior plays a key role in victimization [45, 92, 93, 94]. A cost-effective way to modify these behavioral patterns could fundamentally shift the threat landscape. But while security education efforts have been extensive,

it is striking how little meaningful empirical data we have regarding the impact of user behavior on broad security outcomes. What are the important correlates of victimization? Are they mainly related to standard security hygiene admonishments that users receive (i.e., use a firewall, run anti-virus software, patch your system, don't open attachments from unknown people)? Or do the dynamic usage patterns of individual users play a greater role (i.e., do they use file-sharing networks, do they have a broad or narrow online social network, are they gamers, do they visit "dangerous" sites)?

Our previous work in the context of studying home users in Europe and India (and dorm residents in the US) suggests that hygiene does not play nearly as significant a role as user behavior [111]. Those conclusions, however, have limited strength due to the dearth of ground-truth data that was available for validating the inferences. We thus propose conducting a more detailed and grounded assessment of how security outcomes relate to a range of features, drawing upon network flow and payload capture data available under our relationships with the network operations groups UC San Diego, UC Berkeley, and the Lawrence Berkeley National Laboratory. This data will allow us to assess hygiene and behavior manifestations such as the presence of firewalling, use of antivirus software, patching, visiting blacklisted sites, and accessing file-sharing services. By then combining this data with operational security trouble tickets to label outcomes, we can explore the relationship between these factors and subsequent security incidents. Finally, if we identify key features as important correlates, we have the possibility to embed select security measurements into our *Netalyzr* measurement platform, which provides a broad, global reach into the end user vantage point [89].

We aim to use this data to form evidence-based security policies, such as prioritizing which technical measures or user behaviors to target to best improve security outcomes. In addition, these large, labeled datasets will enable us to then train classifiers and make predictions about which machines or users are most likely to be compromised, similar in spirit to our previous work on applying machine learning to automatically classify the security risk of URLs [106, 107, 154] and to predicting whether and when particular vulnerabilities will be exploited [29].

The biggest potential impact, however, is in terms of education, which has the potential advantage of being both broad and long-lived. Unfortunately, high quality, "teachable moment" education can be expensive, and even when automated requires active participation [92, 147]. However, we hypothesize that the same conditioning properties of online-social networks that we discussed as a *vulnerability* in Section 3 may provide an important lever for improving user behavior. Our team's previous work has highlighted the powerful role that social networks can play in other behavior-driven outcomes such as obesity [40]. Similarly, a user's social networks may already serve as an implicit carrier for bad security practices; numerous anecdotes indicate that security behavior appears shared among friends. However, the extent to which this interaction represents true information transmission rather than mere common mindsets is unclear. To this end, we propose to experiment with employing online social networks as a *positive* vehicle for improving security behaviors, piggybacking on the existing UCSD Social Wellness Project. This effort, led by one of our team members, explores the efficacy of online social networks for influencing health practice (e.g., exercise, vaccination) and affective states (e.g., happiness). As with our prior Facebook voting study [28], we will provide an opportunity for users to easily report (e.g., via a Facebook button) their positive security behaviors and negative security experiences. The data this feedback provides will enable us to assess the degree to which key features of a social network and a user's place within it influence outcomes.

# 6 Broader Impacts

If successful, this research has the potential to dramatically impact society by undermining entire cybercrime ecosystems: disrupting underground activities, infrastructure, and social networks through strategic intervention. Inhibiting the flow of money reduces the profitability of these activities, thereby subverting the key incentive underlying modern cybercrime. The results have the potential to be widely transformative,

impacting virtually all Internet users given the indiscriminate nature of cybercrime attacks.

The PIs are deeply committed to education and curricula development. Student evaluations routinely rank their courses among the top in their departments, and the PIs have received numerous teaching awards, with PI Voelker most recently receiving the UCSD Academic Senate Distinguished Teaching Award, and PI Paxson the Jim and Donna Gray Faculty Award for Excellence in Undergraduate Teaching, both in 2011. The effectiveness of the courses rests partly on incorporating research materials into the course work. All the graduate and advanced undergraduate courses offered by the PIs are project-oriented, and course projects regularly lead to subsequent research publications.

The resources from this grant and the context of this work will create numerous educational opportunities for students and practitioners at a variety of levels. At the graduate level, we will integrate the theme of socio-economic perspectives on Internet security to train advanced students in the graduate security and networking courses regularly taught at all three universities. The material will include in-depth studies of the technical infrastructure and economic and social behavior underlying pervasive cybercrime activities such as spamming, host compromise, click fraud, and identity theft. We will also develop course modules on the same topic for inclusion in our undergraduate security courses, with a specific goal to engage more undergraduates in research. Undergraduates working with the PIs have co-authored research publications and many have gone on to pursue graduate study. We will enthusiastically continue these efforts.

We will also seek new opportunities to develop courses that span disciplines. PI Savage, for example, has been co-developing an interdisciplinary course on cybersecurity and public policy with Prof. Peter Cowhey in the UCSD International Relations and Pacific Studies department. The course includes the lessons learned from our ongoing research in this space, yet structured for student audiences beyond computer science.

The PIs also have a long history of outreach to high school students. The UCSD PIs participate in the California State Summer School for Mathematics and Science (COSMOS) residential summer program for talented high school students. High school students in COSMOS select clusters in sub-disciplines of science and engineering and complete hands-on projects in partnership with faculty and graduate students, and both PIs Savage and Voelker participate. In addition, PI Weaver annually gives a series of talks to students in the Berkeley Foundation for Opportunities in Information Technology (BFOIT), a program that provides outreach to historically underrepresented minorities and women high school students who have a desire to become leaders in the field of computer science, engineering, and information technology.

Our team also frequently engages in public policy outreach activities to government agencies. PI Savage has given focused briefings to members of the US Intellectual Property Enforcement Coordinator, the Federal Trade Commission, the Food and Drug Administration, the Treasury's Office of Illicit Finance, staffers for the Senate Commerce committee, DARPA's ISAT, and a range of US law enforcement and intelligence representatives. Similarly, PI Paxson's briefings of the UK House of Lords Science and Technology Committee, the British Consulate, the United States Air Force Strategic Advisory Board, the US Department of State, and the acting director of the DHS National Cyber Security Division, as well as his service with the President's Council of Advisors on Science and Technology Technical Advisory Group for Networking and Information Technology, have provided venues for communicating our findings to major policy-makers.

Due to the broad nature of our research efforts, we will also promote workforce education by developing training materials appropriate for security professionals, law enforcement and civil regulatory agencies, and legal scholars and professionals. In particular, we expect the data and models that we generate to be unique in the community and of value to a wide range of researchers (including in computer science, economics, sociology and criminology), practitioners, and policy makers in addressing Internet security issues. The PIs have an established track record of creating tutorials for major conferences, including ACM SIGCOMM, ACM CCS, and USENIX Security, that provided detailed technical overviews of large-scale host compromise circa 2005. As part of this effort, we will develop similar tutorials for security practitioners focused on both the technical and socio-economic aspects of contemporary cybercrime. We will also perform broad outreach by supporting a research workshop in the area. PIs Paxson and Savage originally formed the

ACM Workshop on Rapid Malcode (WORM), which then merged with the HotBots workshop to form the USENIX Workshop on Large-scale Exploits and Emergent Threats (LEET). Paxson and Savage are stewards of the workshop, and will evolve its charter to match the evolution of socio-economic threats.

Finally, we will widely disseminate our research products, course and tutorial materials, software tools, data sets, presentation videos, and other education and outreach resources through a project Web portal.

# 7 Results from Prior NSF Support

**UCSD.** The UCSD investigators have been PI and/or Co-PI of numerous NSF awards in the areas of network security, modeling, and reliability. These grants and their results include: CCR-0311690, "Quantitative Network Security Analysis" [73, 119, 120, 121]; CCR-0411307, "Generating Realistic Network Traffic and Topologies" [37, 38, 64, 65]; CNS-0433668/NSF-0433702 (joint with ICSI), "Collaborative Center for Internet Epidemiology and Defenses" [19, 20, 30, 35, 50, 51, 52, 53, 68, 72, 74, 75, 76, 77, 85, 86, 97, 98, 103, 105, 106, 109, 117, 118, 142, 145, 149, 157, 158, 171, 173]; CNS-0722031, "Privacy Preserving Attribution & Provenance" [1]; CNS-0831138, "Understanding and Exploiting Economic Incentives in Internet-based Scams" [29, 78, 95, 99, 101, 107, 108, 122, 123, 124, 141, 159, 172]; TC-0963702, "Foundations, Architectures, and Methodologies for Secure and Private Cyberphysical Vehicles" [36, 82]; CNS-1018910, "Symbiosis in Byzantine Fault Tolerance and Intrusion Detection"; CNS-1116904, "Understanding Network Failures" [156]. In part with this support, the UCSD PIs have collectively supervised dozens of both Ph.D. and M.S. students.

**ICSI.** The ICSI investigators on our team have been PI and/or Co-PI of numerous NSF awards across a wide range of networking topics. In the area of network security, these include: STI-0334088, "Viable Network Defense for Scientific Research Institutions"; NRT-0335290, "Testing and Benchmarking Methodologies for Future Network Security Mechanisms"; NSF-0433668/NSF-0433702 (joint with UCSD) "Cybertrust Center for Internet Epidemiology and Defenses"; CNS-0627320, "Approaches to Network Defense Proven in Open Scientific Environments"; CNS-0716636, "Exploiting Multi-Core CPUs for Parallelizing Network Intrusion Prevention"; and CNS-0716640, "Establishing a Cross-Institutional Platform for Cooperative Security Monitoring and Forensics"; CNS-0749648, "SGER: Architecting Effective Computer Security Grand Challenges"; and CNS-1015835, "Understanding and Taming the Web's Privacy Footprint". Our results to date on these projects include [3, 5, 7, 9, 10, 12, 13, 15, 17, 33, 35, 41, 42, 43, 44, 46, 47, 48, 60, 71, 79, 80, 81, 83, 84, 87, 91, 127, 128, 131, 133, 134, 137, 139, 148, 150, 151, 160, 161, 162, 167, 166, 170]. In the area of network architecture, the work of ICSI investigators spans: ANI-0205519, "Addressing Fundamental Issues for Robust Internet Performance"; CNS-0636539, "Network Fabric for Personal, Social, and Urban Sensing Applications"; CNS-0721933, "Architectural Support for Selectively-Connected End Systems: Enabling an Energy-Efficient Future Internet"; CNS-0722035, "Architectural Support for Network Trouble-Shooting"; and CNS-0831780, "Relationship-Oriented Networking". Results from these efforts include [2, 3, 4, 6, 8, 11, 12, 14, 16, 24, 25, 26, 27, 34, 41, 49, 54, 55, 56, 83, 90, 100, 114, 115, 119, 120, 127, 131, 132, 135, 144, 146, 148, 160, 164] and elements of [9, 10, 35, 44, 46, 80, 91, 133, 134, 137, 151, 161, 162, 163, 165, 167]. Finally, in the area of network measurement their awards include: NSF-9711091, "Creating a National Internet Measurement Infrastructure"; NSF-0222846, "An Open Infrastructure for Network Performance and Security Monitoring"; CNS-0905631, "Invigorating Empirical Network Research via Mediated Trace Analysis" and CNS-0831535, "Comprehensive Application Analysis and Control". Results from these efforts include: [32, 59, 61, 83, 110, 126, 125, 138, 140].

**GMU.** Damon McCoy was previously awarded an NSF/CRA CI Fellowship and has produced results across a wide range of networking topics. In the area of network security and privacy, these include [18, 21, 22, 23, 57, 62, 112, 113, 129, 130].

# 8 Collaboration Plan

The project PIs have an established track record of long-term success at close project collaboration. Spanning eight years, we have worked intensely together on two large projects, starting in 2004 with an NSF CyberTrust Center, "Collaborative Internet Epidemiology and Defenses," and then beginning in 2009 with an ONR MURI project, "Infiltration of Botnet Command-&-Control and Support Ecosystems". We have co-authored dozens of papers on multiple research efforts, with authors from both sites contributing specialized and complementary expertise. We regularly meet on a weekly basis, with both overall group as well as project-specific conference calls. We share project infrastructure, including operational capabilities such as a large-scale botfarm, high-fidelity Web crawling, source code repositories, multi-TB databases, and collaborative resources on a group Wiki. We share extensive data sets, both gathered by our team as well as contributed from external industrial sources. We periodically have cross-site meetings and perform cross-site personnel exchanges, including UCSD students visiting ICSI as interns for multiple months, and ICSI research staff taking visiting scholar positions at UCSD for up to a year. Finally, we extend our close collaborations to industry partners as well, including ongoing shared collaborations with Google, Microsoft, Yahoo, and Twitter.

In sum, we are a group that closely and enthusiastically works together, with a demonstrated record of research success at the highest level. We emphasize the importance of such team cohesion for pursuing research in this problem space: our past experience has shown several times that the ability to nimbly shift focus and resources to an emergent opportunity can prove vital for enabling a particular study.

## 8.1 Participant Roles

PI Savage will be Project Director and serve as the central point of contact with NSF. He will be supported administratively by Jennifer Folkestad, who is partially funded by the project for this purpose. Together with Co-PIs Voelker and Levchenko, Savage will lead the overall project's technical activities at UCSD. Individual UCSD Co-PIs will also lead technical activities according to their strengths and specializations. Lawrence Saul, for instance, will provide machine learning expertise for analyzing the large data sets gathered throughout the project. James Fowler will provide the political, social science, and behavioral economic expertise for understanding the social structure of both attackers and victims, and reasoning about effective intervention in social structures.

Erin Kenneally, Senior Personnel from UCSD, will provide operational legal and ethics guidance for the project. She is a licensed attorney and cyber forensics analyst at the San Diego Supercomputer Center who specializes in evidentiary, procedural, and policy implications related to digital forensics, information security, and privacy technology. As a co-author of the recent Department of Homeland Security (DHS) report, "The Menlo Report: Ethical Principles Guiding Information and Communication Technology Research", she is one of the foremost experts on ethics in ICT research. As a member of our joint projects over the last eight years, Kenneally has provided essential advice on both the legal and ethical issues of the cybersecurity research our group has performed, and she will continue to provide the same critical role in this project. More broadly, Kenneally will also serve as liaison to the law enforcement and legal communities.

PI Paxson will lead project efforts at ICSI and supervise ICSI personnel. He will be supported by ICSI's administrative staff, the cost of which is funded by the institute's overhead structure. Working with Co-PIs Allman, Grier, Kreibich and Weaver, Paxson will provide overall technical leadership at ICSI. Senior Personnel Dan Klein will lead efforts on automating natural-language based analysis of online economic and social interactions, such as markets and forums. Senior Personnel Deirdre Mulligan will provide expert advice in exploring the potentials and implications of security interventions, both with regard to evaluating interventions developed by the team's research efforts, and by framing new research pursuits the team might undertake in light of particular intervention angles. Senior Personnel Chris Hoofnagle will provide expert

advice in assessing privacy considerations for the team's research thrusts and for potential interventions the team devises, particularly in the arena of social networking.

Co-PI Damon McCoy will lead project efforts at GMU. McCoy has critical expertise and experience interfacing with the black hat and underground communities, providing invaluable insight into the means, methods, and motivations of attackers, miscreants, and nefarious actors and organizations.

Finally, successfully undertaking our past efforts has required the development of a variety of relationships with operators at large Web services, network and system administrators at multiple organizations, multiple domain registrars, policy and legal consultants, and law enforcement representatives. We will continue to leverage these relationships as well as developing new ones in this project.

## 8.2 Project Management

Broadly speaking, PI Savage will manage project efforts at UCSD and PI Paxson will manage project efforts at Berkeley (ICSI and UCB). Savage will serve as Project Director, but in practice they will coordinate the overall project together through frequent contact over the phone, email, and instant messaging, as well as in person at regular project meetings and opportunistically at conferences and workshops.

In addition to coordinating the project's activities, PIs Savage and Paxson share responsibility for managing the research evaluation process and making critical decisions about prioritizing resources. We will conduct formal project assessments during our biannual inter-site visits with input from other PIs and senior personnel. For each subproject we will assess its place in the priority list, make go/no-go milestones to manage risk around particular technical approaches, and evolve the overall milestone schedule accordingly. We will also evaluate external feedback from the research community (from presented papers, talks, etc.).

In making technical decisions we will adhere to the following guidelines:

- *Prioritize joint-use technologies*. In particular, we will favor those projects that effectively leverage cross-site collaboration.
- *Act on serendipitous opportunities*. In this quickly-evolving field it is common that the window of opportunity to evaluate a new threat or enterprise, monitor its evolution, and analyze its actors is quite short. For those events representing a significant research opportunity, we will be flexible in temporarily redirecting the project's attention to ensure that we can gather vital data.
- *Consensus of active PIs*. For significant resource reallocations we will seek the consensus of the active PIs affected. Historically, our group has worked extremely well together so we believe that any reallocations will readily prove tenable. However, in the case of significant disagreement, lead PIs Paxson and Savage will make the final decisions on such matters.

## 8.3 Coordination Mechanisms

The proposed effort will comprise many individuals spread across three geographic areas, four campuses (UCSD, ICSI, UCB, and GMU) and distinct administrative units (including CSE, SDSC, and CalIT2 at UCSD). Thus, ensuring that this activity constitutes a single unified project requires consistent and frequent coordination. Through our joint efforts spanning eight years in the CCIED and MURI Botnet Infiltration projects, we have successfully addressed this requirement through a number of mechanisms. Based on these successes, we plan to continue with these approaches in the proposed project:

- *Inter-site weekly phone conferences*. The UCSD, ICSI, and GMU participants currently engage in regular weekly group meetings to provide status updates, report problems, ideas and challenges, and to prepare for inter-site exchanges. Such weekly conference call meetings have proven extremely effective at coordinating activities spanning sites, and we will continue this practice in this project.

- *Inter-site PI phone conferences*. PIs at each site will coordinate regularly via phone every two weeks. On the agenda are recent technical progress, prioritization of effort around upcoming deadlines (paper submissions, agency reporting and deployment activities) and coordination of inter-site plans.
- *Biannual cross-site visits*. Twice a year a significant fraction of all project participants will meet physically to describe the efforts of the past half year, brainstorm on new approaches, and provide a venue for evaluating progress.
- *Cross-site personnel exchanges*. The group has a long history of cross-site personnel exchanges that further facilitate close collaboration and interaction across sites, generate additional familiarity among project members, and expose student and staff members to a broader range of environments. These exchanges have been both long-term and short-term, including UCSD students visiting ICSI as interns for multiple months, ICSI research staff taking visiting scholar positions at UCSD for up to a year, and ICSI research staff and students visiting UCSD for periods of intense collaboration. In this project, we plan to continue such cross-site personnel exchanges, extending our past interactions to also include GMU personnel as well (e.g., we anticipate having GMU students spending terms at UCSD and interning at ICSI).
- *Shared source code repository*. UCSD, ICSI, and GMU share a subversion revision control repository maintained at ICSI for holding shared source code. Building prototypes from a shared source base imposes a design discipline that ensures a fundamental degree of coordination and collaboration to maintain system integrity and consistency.
- *Shared infrastructure, tools, and data*. We also share critical infrastructure such as the GQ botnet and malware large-scale botfarm hosted at ICSI, tools such as the Bro network security monitor [136] and the Netlyzr network connectivity analyzer [89], and data such as spam, URL, DNS, blacklist, and network telescope feeds.
- *Project Wiki*. We will establish a Wiki site that both serves as a public face for the project as well as houses extensive, private collaborative resources such as notes, datasets, draft documents, logs, and presentation materials.
- *Project-wide mailing lists*. As we have historically, we will maintain a project-wide mailing list that includes all technical participants, used both for providing updates on project activities, discussing broad technical questions, and focusing attention on timely events and opportunities.
- *Project-specific mailing lists*. Individual sub-projects will also maintain dedicated technical mailing lists devoted to coordinating detailed technical issues within each project.
- *Annual meeting*. Once a year we will bring together all project personnel, including PIs, staff, and students. These meetings will serve a number of objectives and, by gathering all personnel, ensure comprehensive and coordinated progress and direction of all activities. In the past, we have typically held annual meetings across all groups at the end of the academic year.

## 8.4 Budgetary Support

As detailed in our budget justification, the project budget includes a moderate amount of support for managing the project and supporting collaboration. This support includes administrative support for lead PI Savage in the form of partial salary funding for Jennifer Folkestad at UCSD. The budget line items for graduate student support cover expenses for cross-site internships, and the budget line items for travel cover expenses for cross-site meetings. Similarly, the proposed budget for supplies and equipment cover items that will also be used as shared infrastructure.

# References


[1] M. Afanasyev, T. Kohno, J. Ma, N. Murphy, S. Savage, A. C. Snoeren, and G. M. Voelker. Privacy-preserving Network Forensics. *Communications of the Association for Computing Machinery*, 54(5):78–87, May 2011.

[2] Z. Al-Qudah, H. Alzoubi, M. Allman, M. Rabinovich, and V. Liberatore. Efficient Application Placement in a Dynamic Hosting Platform. In *International World Wide Web Conference*, Apr. 2009.

[3] Z. Al-Qudah, M. Rabinovich, and M. Allman. Web Timeouts and Their Implications. In *Passive and Active Measurement Conference (PAM)*, Apr. 2010.

[4] M. Allman. Personal Namespaces. In *Proc. HotNets*, Nov. 2007.

[5] M. Allman. Comments On Selecting Ephemeral Ports. *ACM Computer Communication Review*, 39(2), Apr. 2009.

[6] M. Allman. On Building Special-Purpose Social Networks for Emergency Communication. *ACM Computer Communication Review*, 40(5), Oct. 2010.

[7] M. Allman, P. Barford, B. Krishnamurthy, and J. Wang. Tracking the Role of Adversaries in Measuring Unwanted Traffic. In *Workshop on Steps to Reduce Unwanted Traffic on the Internet (SRUTI)*, July 2006.

[8] M. Allman and E. Blanton. Notes on Burst Mitigation for Transport Protocols. *ACM Computer Communication Review*, 35(2), Apr. 2005.

[9] M. Allman, E. Blanton, and V. Paxson. An Architecture for Developing Behavioral History. In *Proceedings of USENIX Workshop on Steps to Reducing Unwanted Traffic on the Internet*, July 2005.

[10] M. Allman, E. Blanton, V. Paxson, and S. Shenker. Fighting Coordinated Attackers with Cross-Organizational Information Sharing. In *Proc. HotNets*, 2006.

[11] M. Allman, K. Christensen, B. Nordman, and V. Paxson. Enabling an Energy-Efficient Future Internet Through Selectively Connected End Systems. In *Proc. HotNets*, Nov. 2007.

[12] M. Allman, C. Kreibich, V. Paxson, R. Sommer, and N. Weaver. The Strengths of Weaker Identities: Opportunistic Personas. In *Proceedings of the USENIX Hot Security Workshop*, August 2007.

[13] M. Allman, C. Kreibich, V. Paxson, R. Sommer, and N. Weaver. Principles for Developing Comprehensive Network Visibility. In *USENIX Workshop on Hot Topics in Security*, July 2008.

[14] M. Allman, L. Martin, M. Rabinovich, and K. Atchinson. On Community-Oriented Internet Measurement. In *Proc. Passive and Active Measurement Conference*, Apr. 2008.

[15] M. Allman and V. Paxson. Issues and Etiquette Concerning Use of Shared Measurement Data. In *ACM SIGCOMM/USENIX Internet Measurement Conference*, Oct. 2007.

[16] M. Allman and V. Paxson. A Reactive Measurement Framework. In *Proc. Passive and Active Measurement Conference*, Apr. 2008.

[17] M. Allman, V. Paxson, and J. Terrell. A Brief History of Scanning. In *ACM SIGCOMM/USENIX Internet Measurement Conference*, Oct. 2007.

[18] M. AlSabah, K. Bauer, I. Goldberg, D. Grunwald, D. McCoy, S. Savage, and G. M. Voelker. DefenestraTor: throwing out windows in Tor. In *Proceedings of the 11th international conference on Privacy enhancing technologies*, PETS'11, pages 134–154, Berlin, Heidelberg, 2011. Springer-Verlag.

[19] D. S. Anderson, C. Fleizach, S. Savage, and G. M. Voelker. Spamscatter: Characterizing Internet Scam Hosting Infrastructure. In *Proceedings of the USENIX Security Symposium*, Boston, MA, Aug. 2007.

[20] S. N. Bannur, L. K. Saul, and S. Savage. Judging a Site by its Content: Learning the Textual, Structural, and Visual Features of Malicious Web Pages. In *Proceedings of the ACM Workshop on Artificial Intelligence and Security (AISEC)*, Chicago, IL, Oct. 2011.

[21] K. Bauer, D. Mccoy, B. Greenstein, D. Grunwald, and D. Sicker. Physical Layer Attacks on Unlinkability in Wireless LANs. In *Proceedings of the 9th International Symposium on Privacy Enhancing Technologies*, PETS '09, pages 108–127, Berlin, Heidelberg, 2009. Springer-Verlag.

[22] K. Bauer, D. McCoy, D. Grunwald, T. Kohno, and D. Sicker. Low-resource Routing Attacks Against TOR. In *Proceedings of the 2007 ACM workshop on Privacy in electronic society*, WPES '07, pages 11–20, New York, NY, USA, 2007. ACM.

[23] K. Bauer, M. Sherr, D. McCoy, and D. Grunwald. ExperimenTor: a Testbed for Safe and Realistic TOR Experimentation. In *Proceedings of the 4th Conference on Cyber Security Experimentation and Test*, CSET'11, pages 7–7, Berkeley, CA, USA, 2011. USENIX Association.

[24] S. Bhandarkar, A. L. N. Reddy, M. Allman, and E. Blanton. Improving the Robustness of TCP to Non-Congestion Events, Aug. 2006. RFC 4653.

[25] E. Blanton and M. Allman. Using TCP Duplicate Selective Acknowledgement (DSACKs) and Stream Control Transmission Protocol (SCTP) Duplicate Transmission Sequence Numbers (TSNs) to Detect Spurious Retransmissions, Feb. 2004. RFC 3708.

[26] E. Blanton and M. Allman. On the Impact of Bursting on TCP Performance. In *Passive and Active Measurement Workshop*, Mar. 2005.

[27] J. Blanton, E. Blanton, and M. Allman. Using Spurious Retransmissions to Adapt the Retransmission Timeout. Technical Report 08-005, International Computer Science Institute, Aug. 2008.

[28] R. M. Bond, C. J. Fariss, J. J. Jones, A. D. I. Kramer, C. Marlow, J. E. Settle, and J. H. Fowler. A Massive Scale Experiment in Social Influence and Political Mobilization. Unpublished manuscript, 2011.

[29] M. Bozorgi, L. K. Saul, S. Savage, and G. M. Voelker. Beyond Heuristics: Learning to Classify Vulnerabilities and Predict Exploits. In *Proceedings of the ACM SIGKDD Conference*, pages 105–114, Washington D.C., July 2010.

[30] E. Buchanan, R. Roemer, H. Shacham, and S. Savage. When Good Instructions Go Bad: Generalizing Return-oriented Programming to the SPARC. In *Proceedings of the ACM Conference on Computer and Communications Security*, Alexandria, VA, Oct. 2008.

[31] J. Caballero, C. Grier, C. Kreibich, and V. Paxson. Measuring Pay-per-Install: The Commoditization of Malware Distribution. In *Proceedings of the 20th USENIX Security Symposium*, 2011.

[32] T. Callahan, M. Allman, and V. Paxson. A Longitudinal View of HTTP Traffic. In *Passive and Active Measurement Conference (PAM)*, Apr. 2010.

[33] T. Callahan, M. Allman, and M. Rabinovich. Pssst, Over Here: Communicating Without Fixed Infrastructure. In *IEEE Infocom Mini-Conference*, Mar. 2012.

[34] T. Callahan, M. Allman, M. Rabinovich, and O. Bell. On Grappling with Meta-Information in the Internet. *ACM Computer Communication Review*, 41(5), Oct. 2011.

[35] M. Casado, T. Garfinkel, W. Cui, V. Paxson, and S. Savage. Opportunistic Measurement: Extracting Insight from Spurious Traffic. In *Proceedings of the 4th ACM Workshop on Hot Topics in Networks (HotNets-IV)*, College Park, MD, Nov. 2005.

[36] S. Checkoway, D. McCoy, D. Anderson, B. Kantor, H. Shacham, S. Savage, K. Koscher, A. Czeskis, F. Roesner, and T. Kohno. Comprehensive Experimental Analyses of Automotive Attack Surfaces. In *Proceedings of the USENIX Security Symposium*, San Francisco, CA, Aug. 2011.

[37] Y.-C. Cheng, M. Afanasyev, P. Verkaik, P. Benkö, J. Chiang, A. C. Snoeren, S. Savage, and G. M. Voelker. Automating Cross-Layer Diagnosis of Enterprise Wireless Networks. In *Proceedings of the ACM SIGCOMM Conference*, Kyoto, Japan, Aug. 2007.

[38] Y.-C. Cheng, J. Bellardo, P. Benkö, A. C. Snoeren, G. M. Voelker, and S. Savage. Jigsaw: Solving the Puzzle of Enterprise 802.11 Analysis. In *Proceedings of the ACM SIGCOMM Conference*, pages 39–50, Pisa, Italy, Sept. 2006.

[39] C. Y. Cho, J. Caballero, C. Grier, V. Paxson, and D. Song. Insights from the Inside: A View of Botnet Management from Infiltration. In *Proceedings of the 3rd USENIX Workshop on Large-Scale Exploits and Emergent Threats (LEET)*, April 2010.

[40] N. Christakis and J. Fowler. The Spread of Obesity in a Large Social Network over 32 Years. *New England Journal of Medicine*, 357(4):370–379, 2007.

[41] D. D. Clark, J. Wroclawski, K. R. Sollins, and R. Braden. Tussle in Cyberspace: Defining Tomorrow's Internet. In *ACM SIGCOMM*, Aug. 2002.

[42] W. Cui, V. Paxson, and N. Weaver. GQ: Realizing a System to Catch Worms in a Quarter Million Places. Technical report, International Computer Science Institute, Berkeley, CA, 2006.

[43] W. Cui, V. Paxson, N. Weaver, and R. Katz. Protocol-Independent Adaptive Replay of Application Dialog. In *Network and Distributed Security Symposium (NDSS)*, 2006.

[44] S. Dharmapurikar and V. Paxson. Robust TCP Stream Reassembly in the Presence of Adversaries. In *USENIX Security Symposium*, Aug. 2005.

[45] J. Downs, M. Holbrook, and L. Cranor. Decision Strategies and Susceptibility to Phishing. In *Proceedings of the Symposium On Usable Privacy and Security*, July 2006.

[46] H. Dreger, A. Feldmann, M. Mai, V. Paxson, and R. Sommer. Dynamic Application-Layer Protocol Analysis for Network Intrusion Detection. In *USENIX Security Symposium*, 2006.

[47] H. Dreger, A. Feldmann, V. Paxson, and R. Sommer. Operational Experiences with High-Volume Network Intrusion Detection. In *ACM CCS*, Oct. 2004.

[48] H. Dreger, C. Kreibich, V. Paxson, and R. Sommer. Enhancing the Accuracy of Network-based Intrusion Detection with Host-based Context. In *Conference on Detection of Intrusions and Malware and Vulnerability Assessment (DIMVA)*, 2005.

[49] W. Eddy, S. Ostermann, and M. Allman. New Techniques for Making Transport Protocols Robust to Corruption-Based Loss. *ACM Computer Communication Review*, 34(5), Oct. 2004.

[50] B. Enright, G. Voelker, S. Savage, C. Kanich, and K. Levchenko. Storm: When Researchers Collide. *USENIX ;login:*, 33(4):6–13, Aug. 2008.

[51] R. Farrow and S. Savage. Interview with Stefan Savage: On the Spam Payment Trail. *USENIX ;login:*, 36(4):7–20, Aug. 2011.

[52] C. Fleizach, M. Liljenstam, P. Johansson, G. M. Voelker, and A. Méhes. Can You Infect Me Now? Malware Propagation in Mobile Phone Networks. In *Proceedings of the ACM Workshop on Recurring Malcode (WORM)*, Washington D.C., Nov. 2007.

[53] C. Fleizach, G. M. Voelker, and S. Savage. Slicing Spam with Occam's Razor. In *Proceedings of Conference on Email and Anti-Spam (CEAS)*, Mountain View, CA, Aug. 2007.

[54] S. Floyd and M. Allman. Specifying New Congestion Control Algorithms, Aug. 2007. RFC 5033, BCP 133.

[55] S. Floyd and M. Allman. Comments on the Usefulness of Simple Best-Effort Traffic, July 2008. RFC 5290.

[56] S. Floyd, M. Allman, A. Jain, and P. Sarolahti. Quick-Start for TCP and IP, Jan. 2007. RFC 4782.

[57] J. Franklin, D. McCoy, P. Tabriz, V. Neagoe, J. Van Randwyk, and D. Sicker. Passive data link layer 802.11 wireless device driver fingerprinting. In *Proceedings of the 15th conference on USENIX Security Symposium - Volume 15*, Berkeley, CA, USA, 2006. USENIX Association.

[58] J. Franklin, V. Paxson, A. Perrig, and S. Savage. An Inquiry into the Nature and Causes of the Wealth of Internet Miscreants. In *Proceedings of the ACM Conference on Computer and Communications Security*, Alexandria, VA, Oct. 2007.

[59] J. Gonzalez and V. Paxson. *pktd*: A Packet Capture and Injection Daemon. In *Passive and Active Measurement Workshop*, 2003.

[60] J. M. Gonzalez and V. Paxson. Enhancing Network Intrusion Detection With Integrated Sampling and Filtering. In *Recent Advances in Intrusion Detection (RAID)*, 2006.

[61] R. Govindan and V. Paxson. Estimating Router ICMP Generation Delays. In *Proceedings of Passive and Active Measurement*, Mar. 2002.

[62] B. Greenstein, D. McCoy, J. Pang, T. Kohno, S. Seshan, and D. Wetherall. Improving Wireless Privacy with an Identifier-free Link Layer Protocol. In *Proceedings of the 6th international conference on Mobile systems, applications, and services*, MobiSys '08, pages 40–53, New York, NY, USA, 2008. ACM.

[63] C. Grier, K. Thomas, V. Paxson, and M. Zhang. @spam: The Underground on 140 Characters or Less. In *Proceedings of the ACM Conference on Computer and Communications Security (CCS)*, October 2010.

[64] D. Gupta, S. Lee, M. Vrable, S. Savage, A. C. Snoeren, G. Varghese, G. M. Voelker, and A. Vahdat. Difference Engine: Harnessing Memory Redundancy in Virtual Machines. In *Proceedings of the 8th ACM/USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, pages 309–322, San Diego, CA, Dec. 2008. Award paper.

[65] D. Gupta, K. Yocum, M. McNett, A. C. Snoeren, A. Vahdat, and G. M. Voelker. To Infinity and Beyond: Time-Warped Network Emulation. In *Proceedings of the 3rd ACM/USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, pages 87–100, San Jose, CA, May 2006.

[66] A. Haghighi and D. Klein. An Entity-Level Approach to Information Extraction. In *Proceedings of ACL*, pages 291–295, Uppsala, Sweden, July 2010. Association for Computational Linguistics.

[67] A. Haghighi and D. Klein. Coreference Resolution in a Modular, Entity-Centered Model. In *Proceedings of NAACL*, pages 385–393, Los Angeles, California, June 2010. Association for Computational Linguistics.

[68] T. Halvorson, J. Szurdi, G. Maier, M. Felegyhazi, C. Kreibich, N. Weaver, K. Levchenko, and V. Paxson. The BIZ Top-Level Domain: Ten Years Later. In *Proceedings of the Passive and Active Measurement Workshop*, Vienna, Austria, Mar. 2012.

[69] M. Huber, S. Kowalski, M. Nohlberg, and S. Tjoa. Towards Automating Social Engineering Using Social Networking Sites. In *International Conference on Computational Science and Engineering (CSE'09)*, August 2009.

[70] IACC Has New Tools To Cut Off Money to Bad Sites. `https://iacc.org/news-media-resources/press-releases/iacc-has-new-tools-to-cut-off-money-to-bad-sites.php`, September 2011.

[71] 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.

[72] F. Junqueira, R. Bhagwan, A. Hevia, K. Marzullo, and G. M. Voelker. Surviving Internet Catastrophes. In *Proceedings of the USENIX Annual Technical Conference*, Anaheim, CA, Apr. 2005.

[73] F. Junqueira, R. Bhagwan, K. Marzullo, S. Savage, and G. M. Voelker. The Phoenix Recovery System: Rebuilding from the Ashes of an Internet Catastrophe. In *Proceedings of the 9th USENIX Workshop on Hot Topics in Operating Systems (HotOS-IX)*, pages 73–78, Lihue, HI, May 2003.

[74] C. Kanich, N. Chachra, D. McCoy, C. Grier, D. Wang, M. Motoyama, K. Levchenko, S. Savage, and G. M. Voelker. No Plan Survives Contact: Experience with Cybercrime Measurement. In *Proceedings of Workshop on Cyber Security Experimentation and Test (CSET)*, Aug. 2011.

[75] C. Kanich, C. Kreibich, K. Levchenko, B. Enright, V. Paxson, G. M. Voelker, and S. Savage. Spamalytics: an Empirical Analysis of Spam Marketing Conversion. In *Proceedings of the ACM Conference on Computer and Communications Security*, pages 3–14, Alexandria, VA, Oct. 2008.

[76] C. Kanich, C. Kreibich, K. Levchenko, B. Enright, G. M. Voelker, V. Paxson, and S. Savage. Spamalytics: An Empirical Analysis of Spam Marketing Conversion. *Communications of the Association for Computing Machinery*, 52(9):99–107, Sept. 2009.

[77] C. Kanich, K. Levchenko, B. Enright, G. M. Voelker, and S. Savage. The Heisenbot Uncertainty Problem: Challenges in Separating Bots from Chaff. In *Proceedings of the USENIX Workshop on Large-scale Exploits and Emergent Threats (LEET)*, San Franciso, CA, Apr. 2008.

[78] C. Kanich, N. Weaver, D. McCoy, T. Halvorson, C. Kreibich, K. Levchenko, V. Paxson, G. M. Voelker, and S. Savage. Show Me the Money: Characterizing Spam-advertised Revenue. In *Proceedings of the USENIX Security Symposium*, San Francisco, CA, Aug. 2011.

[79] J. Kannan, J. Jung, V. Paxson, and C. E. Koksal. Detecting Hidden Causality in Network Connections, 2005. Technical Report, University of California, Berkeley.

[80] J. Kannan, J. Jung, V. Paxson, and C. E. Koksal. Semi-Automated Discovery of Application Session Structure. In *Proceedings of ACM Internet Measurement Conference*, 2006.

[81] S. Kornexl, V. Paxson, H. Dreger, A. Feldmann, and R. Sommer. Building a Time Machine for Efficient Recording and Retrieval of High-Volume Network Traffic. In *ACM Internet Measurement Conference*, 2005.

[82] K. Koscher, A. Czeskis, F. Roesner, S. Patel, T. Kohno, S. Checkoway, D. McCoy, B. Kantor, D. Anderson, H. Shacham, and S. Savage. Experimental Security Analysis of a Modern Automobile. In *Proceedings of the IEEE Symposium and Security and Privacy*, Oakland, CA, May 2010.

[83] J. Koskela, N. Weaver, A. Gurtov, and M. Allman. Securing Web Content. In *ACM CoNext Workshop on ReArchitecting the Internet (ReArch)*, Dec. 2009.

[84] C. Kreibich. Brooery: A Graphical Environment for Analysis of Security-Relevant Network Activity. In *FREENIX*, 2005.

[85] C. Kreibich, C. Kanich, K. Levchenko, B. Enright, G. M. Voelker, V. Paxson, and S. Savage. On the Spam Campaign Trail. In *Proceedings of the USENIX Workshop on Large-scale Exploits and Emergent Threats (LEET)*, San Franciso, CA, Apr. 2008.

[86] C. Kreibich, C. Kanich, K. Levchenko, B. Enright, G. M. Voelker, V. Paxson, and S. Savage. Spamcraft: An Inside Look at Spam Campaign Orchestration. In *Proceedings of the USENIX Workshop on Large-scale Exploits and Emergent Threats (LEET)*, pages 4:1–4:9, Boston, MA, Apr. 2009.

[87] C. Kreibich and R. Sommer. Policy-controlled Event Management for Distributed Intrusion Detection. In *4th International Workshop on Distributed Event-Based Systems*, 2005.

[88] C. Kreibich, N. Weaver, C. Kanich, W. Cui, and V. Paxson. Practical Containment for Measuring Modern Malware Systems. In *Proceedings of the ACM Internet Measurement Conference*, Berlin, CA, Nov. 2011.

[89] C. Kreibich, N. Weaver, B. Nechaev, and V. Paxson. Netalyzr: Illuminating the edge network. In *Proceedings of the ACM Internet Measurement Conference (IMC)*, pages 246–259, Melbourne, Australia, November 2010.

[90] R. Krishnan, J. Sterbenz, W. Eddy, C. Partridge, and M. Allman. Explicit Transport Error Notification (ETEN) for Error-Prone Wireless and Satellite Networks. *Computer Networks*, 46, Oct. 2004.

[91] A. Kumar, V. Paxson, and N. Weaver. Exploiting Underlying Structure for Detailed Reconstruction of an Internet-scale Event. In *ACM Internet Measurement Conference*, 2005.

[92] P. Kumaraguru, J. Cranshaw, A. Acquisti, L. Cranor, J. Hong, M. Blair, and T. Pham. School of Phish: A Real-Word Evaluation of Anti-Phishing Training. In *Proceedings of the Symposium On Usable Privacy and Security*, 2009.

[93] P. Kumaraguru, Y. Rhee, S. Sheng, S. Hasan, A. Acquisti, L. Cranor, and J. Hong. Getting Users to Pay Attention to Anti-Phishing Education: Evaluation of Retention and Transfer. In *Proceedings of the 2nd Annual eCrime Researchers Summit*, October 2007.

[94] P. Kumaraguru, S. Sheng, A. Acquisti, and L. C. an J. Hong. Lessons from a Real World Evaluation of Anti-phishing Training. In *Proceedings of the third eCrime Researchers Summit*, October 2008.

[95] D. kyum Kim, M. Motoyama, G. M. Voelker, and L. K. Saul. Topic Modeling of Freelance Job Postings to Monitor Web Service Abuse. In *Proceedings of the ACM Workshop on Artificial Intelligence and Security (AISEC)*, Chicago, IL, Oct. 2011.

[96] T. Lauinger, V. Pankakoski, D. Balzarotti, and E. Kirda. Honeybot, Your Man in the Middle for Automated Social Engineering. In *Proceedings of the 3rd USENIX conference on Large-scale exploits and emergent threats: botnets, spyware, worms, and more*, pages 11–11. USENIX Association, 2010.

[97] B. Laxton, K. Wang, and S. Savage. Reconsidering Physical Key Secrecy: Teleduplication via Optical Decoding. In *Proceedings of the ACM Conference on Computer and Communications Security*, Alexandria, VA, Oct. 2008.

[98] K. Levchenko, R. Paturi, and G. Varghese. On the Difficulty of Scalably Detecting Network Attacks. In *Proceedings of the ACM Conference on Computer and Communications Security*, Washington, D.C., Oct. 2004.

[99] K. Levchenko, A. Pitsillidis, N. Chachra, B. Enright, M. Félegyházi, C. Grier, T. Halvorson, C. Kanich, C. Kreibich, H. Liu, D. McCoy, N. Weaver, V. Paxson, G. M. Voelker, and S. Savage. Click Trajectories: End-to-End Analysis of the Spam Value Chain. In *Proceedings of the IEEE Symposium and Security and Privacy*, pages 431–446, Oakland, CA, May 2011.

[100] D. Liu, M. Allman, S. Jin, and L. Wang. Congestion Control Without a Startup Phase. In *Protocols for Fast, Long Distance Networks (PFLDnet) Workshop*, Feb. 2007.

[101] H. Liu, K. Levchenko, M. Félegyházi, C. Kreibich, G. Maier, G. M. Voelker, and S. Savage. On the Effects of Registrar-level Intervention. In *Proceedings of the USENIX Workshop on Large-scale Exploits and Emergent Threats (LEET)*, pages 1–8, Boston, MA, Mar. 2011.

[102] L. Lu, R. Perdisci, and W. Lee. SURF: Detecting and Measuring Search Poisoning. In *Proceedings of the ACM Conference on Computer and Communications Security (CCS)*, October 2011.

[103] J. Ma, J. Dunagan, H. J. Wang, S. Savage, and G. M. Voelker. Finding Diversity in Remote Code Injection Exploits. In *Proceedings of the ACM Internet Measurement Conference*, Rio de Janeiro, Brazil, Oct. 2006.

[104] J. Ma, A. Kulesza, M. Dredze, K. Crammer, L. K. Saul, and F. Pereira. Exploiting Feature Covariance in High-Dimensional Online Learning. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, Sardinia, Italy, May 2010.

[105] J. Ma, K. Levchenko, C. Kriebich, S. Savage, and G. M. Voelker. Automated Protocol Inference: Unexpected Means of Identifying Protocols. In *Proceedings of the ACM Internet Measurement Conference*, Rio de Janeiro, Brazil, Oct. 2006.

[106] J. Ma, L. K. Saul, S. Savage, and G. M. Voelker. Beyond Blacklists: Learning to Detect Malicious Web Sites from Suspicious URLs. In *Proceedings of the ACM SIGKDD Conference*, pages 1245–1254, Paris, France, June 2009.

[107] J. Ma, L. K. Saul, S. Savage, and G. M. Voelker. Identifying Suspicious URLs: An Application of Large-Scale Online Learning. In *Proceedings of the International Conference on Machine Learning*, pages 681–688, Montreal, Quebec, June 2009.

[108] J. Ma, L. K. Saul, S. Savage, and G. M. Voelker. Learning to Detect Malicious URLs. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 2(3):30:1–30:24, Apr. 2011.

[109] J. Ma, G. M. Voelker, and S. Savage. Self-stopping Worms. In *Proceedings of the ACM Workshop on Rapid Malcode (WORM)*, pages 12–21, Washington D.C., Nov. 2005.

[110] G. Maier, A. Feldmann, V. Paxson, and M. Allman. On Dominant Characteristics of Residential Broadband Internet Traffic. In *ACM/USENIX Internet Measurement Conference (IMC)*, Nov. 2009.

[111] G. Maier, A. Feldmann, V. Paxson, R. Sommer, and M. Vallentin. An Assessment of Overt Malicious Activity Manifest in Residential Networks. *Detection of Intrusions and Malware, and Vulnerability Assessment*, 2011.

[112] D. Mccoy, K. Bauer, D. Grunwald, T. Kohno, and D. Sicker. Shining Light in Dark Places: Understanding the Tor Network. In *Proceedings of the 8th international symposium on Privacy Enhancing Technologies*, PETS '08, pages 63–76, Berlin, Heidelberg, 2008. Springer-Verlag.

[113] D. McCoy, J. A. Morales, and K. Levchenko. Proximax: Fighting Censorship with an Adaptive System for Distribution of Open Proxies. In *Proceedings of the International Conference on Financial Cryptography and Data Security*, St Lucia, Feb. 2011.

[114] A. Medina, M. Allman, and S. Floyd. Measuring Interactions Between Transport Protocols and Middleboxes. In *ACM SIGCOMM/USENIX Internet Measurement Conference*, Oct. 2004.

[115] A. Medina, M. Allman, and S. Floyd. Measuring the Evolution of Transport Protocols in the Internet. *ACM Computer Communication Review*, 35(2), Apr. 2005.

[116] B. Miller, P. Pearce, C. Grier, C. Kreibich, and V. Paxson. What's Clicking What? Techniques and Innovations of Today's Clickbots. In *Proceedings of the Conference on Detection of Intrusions and Malware and Vulnerability Assessment*, July 2011.

[117] A. Mizrak, S. Savage, and K. Marzullo. Detecting Compromised Routers via Packet Forwarding Behavior. *IEEE Network*, 22(2), Mar. 2008.

[118] A. Mizrak, S. Savage, and K. Marzullo. Detecting Malicious Packet Losses. *IEEE Transactions on Parallel and Distributed Systems*, 20(2), Feb. 2009.

[119] 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.

[120] D. Moore, V. Paxson, S. Savage, C. Shannon, S. Staniford, and N. Weaver. The Spread of the Sapphire/Slammer Worm. CAIDA Report, Jan. 2003.

[121] 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*, pages 1901–1910, San Francisco, CA, Apr. 2003.

[122] M. Motoyama, K. Levchenko, C. Kanich, D. McCoy, G. M. Voelker, and S. Savage. Re: CAPTCHAs – understanding CAPTCHA-solving from an economic context. In *Proceedings of the USENIX Security Symposium*, pages 435–452, Washington, D.C., Aug. 2010.

[123] M. Motoyama, D. McCoy, K. Levchenko, S. Savage, and G. M. Voelker. Dirty Jobs: The Role of Freelance Labor in Web Service Abuse. In *Proceedings of the USENIX Security Symposium*, San Francisco, CA, Aug. 2011.

[124] M. Motoyama, D. McCoy, S. Savage, and G. M. Voelker. An Analysis of Underground Forums. In *Proceedings of the ACM Internet Measurement Conference*, Berlin, CA, Nov. 2011.

[125] B. Nechaev, M. Allman, V. Paxson, and A. Gurtov. A Preliminary Analysis of TCP Performance in an Enterprise Network. In *USENIX Internet Network Management Workshop/Workshop on Research on Enterprise Networking (INM/WREN)*, Apr. 2010.

[126] B. Nechaev, V. Paxson, M. Allman, and A. Gurtov. On Calibrating Enterprise Switch Measurements. In *ACM/USENIX Internet Measurement Conference (IMC)*, Nov. 2009.

[127] T. Ouyang, S. Ray, M. Allman, and M. Rabinovich. A Large-Scale Empirical Analysis of Email Spam Detection Through Transport-Level Characteristics. Technical Report 10-001, International Computer Science Institute, Jan. 2010.

[128] T. Ouyang, S. Ray, M. Allman, and M. Rabinovich. Can Network Characteristics Detect Spam Effectively in a Stand-Alone Enterprise? In *Passive and Active Measurement Conference*, Mar. 2011.

[129] J. Pang, B. Greenstein, M. Kaminsky, D. McCoy, and S. Seshan. Wifi-Reports: Improving Wireless Network Selection with Collaboration. *IEEE Transactions on Mobile Computing*, 9:1713–1731, December 2010.

[130] J. Pang, B. Greenstein, D. McCoy, S. Seshan, and D. Wetherall. Tryst: The Case for Confidential Service Discovery. In *HotNets VI: The Sixth Workshop on Hot Topics in Networks*, Nov. 2007.

[131] R. Pang, M. Allman, V. Paxson, and J. Lee. The Devil and Packet Trace Anonymization. In *Computer Communication Review*, 2006.

[132] R. Pang and V. Paxson. A High-Level Programming Environment for Packet Trace Anonymization and Transformation. In *ACM SIGCOMM*, Aug. 2003.

[133] R. Pang, V. Paxson, R. Sommer, and L. Peterson. binpac: A yacc for Writing Application Protocol Parsers. In *ACM Internet Measurement Conference*, 2006.

[134] R. Pang, V. Yegneswaran, P. Barford, V. Paxson, and L. Peterson. Characteristics of Internet Background Radiation. In *ACM Internet Measurement Conference*, Oct. 2004.

[135] A. Parker, S. Reddy, T. Schmidt, K. Chang, G. Saurabh, M. Srivastava, M. Hansen, J. Burke, D. Estrin, M. Allman, and V. Paxson. Network System Challenges in Selective Sharing and Verification for Personal, Social, and Urban-Scale Sensing Applications. In *Proc. HotNets*, Nov. 2006.

[136] V. Paxson. Bro: a System for Detecting Network Intruders in Real-Time. *Computer Networks*, 31(23-24):2435–2463, 1999.

[137] V. Paxson. Strategies for Sound Internet Measurement. In *ACM SIGCOMM Internet Measurement Conference*, Oct. 2004.

[138] V. Paxson, A. Adams, and M. Mathis. Experiences with NIMI. In *Proceedings of Passive and Active Measurements (PAM)*, 2000.

[139] V. Paxson, K. Asanovic, S. Dharmapurikar, J. Lockwood, R. Pang, R. Sommer, and N. Weaver. Rethinking Hardware Support for Network Analysis and Intrusion Prevention. In *Proceedings of the USENIX Hot Security Workshop*, August 2006.

[140] V. Paxson, J. Mahdavi, A. Adams, and M. Mathis. An Architecture for Large-Scale Internet Measurement. *IEEE Communications*, 1998.

[141] A. Pitsillidis, K. Levchenko, C. Kreibich, C. Kanich, G. M. Voelker, V. Paxson, N. Weaver, and S. Savage. Botnet Judo: Fighting Spam with Itself. In *Proceedings of the Network and Diestributed System Security Symposium (NDSS)*, San Diego, CA, Feb. 2010.

[142] A. Pitsillidis, Y. Xie, F. Yu, M. Abadi, G. M. Voelker, and S. Savage. How to Tell an Airport from a Home: Techniques and Applications. In *Proceedings of the 9th ACM Workshop on Hot Topics in Networks (HotNets-IX)*, pages 13:1–13:6, Monterey, CA, Oct. 2010.

[143] K. Poulsen. 56 Arrested in DarkMarket Sting, Says FBI.

[144] C. Reis, S. D. Gribble, T. Kohno, and N. Weaver. Detecting In Flight Page Changes with Web Tripwires. In *Proc. of the 5th USENIX Symposium on Networked Systems Design & Implementation*, 2008.

[145] T. Ristenpart, E. Tromer, H. Shacham, and S. Savage. Hey, You, Get Off of My Cloud: Exploring Information Leakage in Third-Party Compute Clouds. In *Proceedings of the ACM Conference on Computer and Communications Security*, Chicago, IL, Nov. 2009.

[146] P. Sarolahti, M. Allman, and S. Floyd. Determining an Appropriate Sending Rate Over an Underutilized Network Path. *Computer Networks Special Issue on Protocols for Fast, Long-Distance Networks*, 51(7), May 2007.

[147] S. Sheng, B. Magnien, P. Kumaraguru, A. Acquisti, L. Cranor, J. Hong, and E. Nunge. Anti-Phishing Phil: The Design and Evaluation of a Game That Teaches People Not to Fall for Phish. In *Proceedings of the Symposium On Usable Privacy and Security*, Pittsburgh, PA, July 2007.

[148] C. Shue, A. Kalafut, C. Taylor, and M. Allman. On Building Inexpensive Network Capabilities. *Under submission*, 2011.

[149] S. Singh, C. Estan, G. Varghese, and S. Savage. Automated Worm Fingerprinting. In *Proceedings of the 6th ACM/USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, pages 45–60, San Francisco, CA, Dec. 2004.

[150] R. Sommer and V. Paxson. Exploiting Independent State For Network Intrusion Detection. In *ACSAC*, Dec. 2005.

[151] S. Staniford, D. Moore, V. Paxson, and N. Weaver. The Top Speed of Flash Worms. In *ACM CCS WORM*, Oct. 2004.

[152] B. Stone-Gross, R. Abman, R. A. Kemmerer, C. Kruegel, and D. G. Steigerwald. The Underground Economy of Fake Antivirus Software. In *Tenth Workshop on Economics of Information Security (WEIS)*, June 2011.

[153] B. Stone-Gross, M. Cova, L. Cavallaro, B. Gilbert, M. Szydlowski, R. Kemmerer, C. Kruegel, and G. Vigna. Your Botnet is My Botnet: Analysis of a Botnet Takeover. In *Proceedings of the 16th ACM conference on Computer and communications security*, October 2009.

[154] K. Thomas, C. Grier, J. Ma, V. Paxson, and D. Song. Design and Evaluation of a Real-Time URL Spam Filtering Service. In *Security and Privacy (SP), 2011 IEEE Symposium on*, pages 447–462. IEEE, May 2011.

[155] K. Thomas, C. Grier, V. Paxson, and D. Song. Suspended Accounts In Retrospect: An Analysis of Twitter Spam. In *Proceedings of the Internet Measurement Conference*, November 2011.

[156] D. Turner, K. Levchenko, A. C. Snoeren, and S. Savage. California Fault Lines: Understanding the Causes and Impact of Network Failures. In *Proceedings of the ACM SIGCOMM Conference*, New Delhi, India, Aug. 2010.

[157] M. Varvello and G. M. Voelker. SecondLife: a Social Network of Humans and Bots. In *Proceedings of the ACM International Workshop on Network and Operating Systems Support for Digital Audio and Video (NOSSDAV)*, pages 9–14, Amsterdam, the Netherlands, June 2010.

[158] M. Vrable, J. Ma, J. Chen, D. Moore, E. VandeKieft, A. C. Snoeren, G. M. Voelker, and S. Savage. Scalability, Fidelity and Containment in the Potemkin Virtual Honeyfarm. In *Proceedings of the 20th ACM Symposium on Operating Systems Principles (SOSP)*, pages 148–162, Brighton, UK, Oct. 2005.

[159] D. Wang, S. Savage, and G. M. Voelker. Cloak and Dagger: Dynamics of Web Search Cloaking. In *Proceedings of the ACM Conference on Computer and Communications Security*, Chicago, IL, Oct. 2011.

[160] N. Weaver and M. Allman. On Constructing a Trusted Path to the User. Technical Report 09-009, International Computer Science Institute, Dec. 2009.

[161] N. Weaver, D. Ellis, S. Staniford, and V. Paxson. Worms vs. Perimeters: The Case for Hard-LANs. In *Hot Interconnects 12*, Aug. 2004.

[162] N. Weaver, I. Hamadeh, G. Kesidis, and V. Paxson. Preliminary Results Using ScaleDown to Explore Worm Dynamics. In *ACM CCS WORM*, Oct. 2004.

[163] N. Weaver and V. Paxson. A Worst-Case Worm. In *Third Annual Workshop on Economics and Information Security (WEIS04)*, May 2004.

[164] N. Weaver, V. Paxson, S. Staniford, and R. Cunningham. A Taxonomy of Computer Worms. In *First ACM CCS Workshop on Rapid Malcode (WORM)*, Oct. 2003.

[165] N. Weaver, V. Paxson, S. Staniford, and R. Cunningham. Large Scale Malicious Code: A Research Agenda, 2003. DARPA-sponsored report.

[166] N. Weaver and R. Sommer. Stress Testing Cluster Bro. In *DETER Community Workshop*, 2007.

[167] N. Weaver, S. Staniford, and V. Paxson. Very Fast Containment of Scanning Worms. In *USENIX Security Symposium*, Aug. 2004.

[168] C. Yang, R. Harkreader, and G. Gu. Die Free or Live Hard? Empirical Evaluation and New Design for Fighting Evolving Twitter Spammers. In *Proceedings of the International Symposium on Recent Advances in Intrusion Detection*, 2011.

[169] K. Yee. Aligning Security and Usability. *IEEE Security & Privacy, IEEE*, 2(5):48–55, 2004.

[170] V. Yegneswaran, P. Barford, and V. Paxson. Using Honeynets for Internet Situational Awareness. In *ACM SIGCOMM HOTNETS*, 2005.

[171] S. Yilek, E. Rescorla, H. Shacham, B. Enright, and S. Savage. When Private Keys are Public: Results from the 2008 Debian OpenSSL Debacle. In *Proceedings of the ACM Internet Measurement Conference*, Chicago, IL, Nov. 2009.

[172] Q. Zhang, T. Ristenpart, S. Savage, and G. M. Voelker. Got Traffic? An Evaluation of Click Traffic Providers. In *Proceedings of the WICOM/AIRWeb Workshop on Web Quality (WebQuality)*, pages 19–26, Hyderabad, India, Mar. 2011.

[173] G. Zyba, G. M. Voelker, M. Liljenstam, A. Méhes, and P. Johansson. Defending Mobile Phones from Proximity Malware. In *Proceedings of the IEEE Infocom Conference*, pages 1503–1511, Rio de Janeiro, Brazil, Apr. 2009.