Internet security is defined by conflict. The value and power mediated by the global, interconnected systems of today's Internet in turn attract adversaries who seek to exploit these same systems for economic, political or social gain. However, the underlying complexity of the Internet infrastructure, the layering of its services, and the indirect nature of its business relationships can make it difficult to identify even the *existence* of adversaries manipulating systems for their benefit. Further, identifying that an attack is taking place is only the first step in a ongoing challenge, as adversaries have the luxury to define where and when their actions take place and responders are forced to discover the landscape of the battle after it commences.

Studying these problems is challenging. While anecdotes and serendipitous findings are common, understanding the full nature of a particular action or attack requires systematic measurement—frequently measurement of an actor who seeks to hide or camouflage their actions. Designing *effective* and *comprehensive* defenses requires sound understanding of not only specific problems, but also the fundamental limitations and costs associated with those problems. In the context of global-scale attacks such as cybercrime and censorship, acquiring this understanding is frequently challenging and requires new types of research systems and measurement methods. Remediation without systematic understanding of opponents' costs, capabilities, and objectives risks developing reactionary, incremental defenses lacking feasibility or robustness.

My research brings empirical grounding and understanding to the study of global, hidden Internet security problems. My work has focused on both politically and economically motivated attacks, spanning censorship, cybercrime, and "advanced persistent threats." In pursuit of these goals I have built Internet-scale measurement platforms and designed new empirical methods aimed at discovering complex and unseen adversarial behavior.

I have published at the top-tier computer security venues IEEE Security & Privacy (receiving the Distinguished Practical Paper award), USENIX Security, and ACM CCS. Via collaborations with Microsoft and Google, my research resulted in real-world remediations, reducing cybercrime and protecting users. My work understanding censorship has led to interdisciplinary collaboration with Harvard's Berkman Klein Center aimed at contributing to social science research.

Understanding Internet Censorship

Anecdotes and reports indicate that Internet censorship is widespread, affecting many countries around the world. Despite this prevalence, empirical Internet measurements revealing the scope and evolution of censorship remain comparatively sparse. This limitation stems from fundamental technical and ethical difficulties in obtaining large-scale, consistent, and sound data, from both numerous countries and multiple vantage points within those countries.

Understanding global manipulation practices is necessary to develop technologies and formulate policies that effectively address censorship. My work develops systems and methods that can perform widespread, ethical, longitudinal measurements of multiple types of global Internet censorship. To achieve the necessary reach and diversity, I have developed systems and methods that do not require the participation of individual users in the countries of interest. My work addresses a range of both technical and extra-technical challenges, at a scale and fidelity not previously achieved.

To enable global continuous measurement of TCP/IP Internet censorship, I developed Augur [1,2], a measurement regimen and accompanying system that soundly leverages potentially highly noisy TCP/IP side channels to measure reachability between two Internet locations without access to measurement vantage points at the locations, or even at points along the path between them. Augur includes the use of sequential hypothesis testing to develop statistical confidence in the face of network and side channel noise, and techniques for responsibly addressing censorship measurement ethics. The validation of Augur included a global censorship measurement study which examined the blocking practices of more than 180 countries and dependent territories.

Adversaries employ a variety of technical mechanisms to achieve Internet censorship. Augur enables us to understand global TCP/IP censorship, but to further build a comprehensive picture of censorship, we need techniques to examine additional commonly deployed blocking technologies. Towards this goal I then developed Iris [3,4], a scalable, accurate, and ethical method to measure global manipulation of DNS resolutions. Iris enables ongoing censorship measurement from more than 150 countries and territories. The deployment of Iris revealed new patterns in

global censorship, including the heterogeneity of censorship within countries.

A critical problem in the field of censorship measurement is ethics. The development of methodologies that are both comprehensive and ethically responsible is challenging, as these goals frequently conflict. By its nature, measuring censorship involves interacting with content that an authority has deemed objectionable; exploring the scope and scale of that content potentially creates risk for those involved in the measurement. Both Augur and Iris include approaches for reasoning about risk, as well as for reducing potential harm via the use of network infrastructure for indirect censorship measurement.

The potential of Augur and Iris have lead to ongoing collaboration with Harvard's Berkman Klein Center aimed at producing datasets that will contribute to policy decisions and social science research.

A related facet of the censorship arms race involves the blocking of censorship circumvention technologies by both state censors as well as private entities. As part of a multi-year "threat intelligence" project (see "Future Directions") I have collected longitudinal information on circumvention infrastructure across hundreds public and private sources and services. This dataset played a key role in measuring how and why circumvention technologies themselves experience blocking [5].

Combating Cybercrime

Cybercrime has evolved into a complex global ecosystem of criminal actors performing attacks ranging from ransomware to denial-of-service to advertising fraud. The motivation for these attacks is economic—criminals carry out these attacks in order to monetize users at a global scale. This economic force means criminals focus their efforts on maximizing profit rather than relying on specific technical mechanisms. Past work has shown that solutions which focus on the financial and relational aspects of cybercrime have the potential to be more effective than incremental technological defenses.

I combine a range of methods and data sources—from command-and-control infiltration, to direct measurement, to industry data—in order to construct understanding of criminal attacks which, by their nature, are designed to be difficult to distinguish. With this understanding, my work identified fundamental structural weak-points leverageable for defense, resulting in dismantling botnets, cleaning up ad networks, and protecting users.

Advertising abuse is a prevalent and lucrative cybercrime activity that exploits the wealth of the online advertising ecosystem while masking criminal identity and activities through the ecosystem's byzantine structure. The scale and structure of these attacks tilt the conflict in favor of the criminals, impacting users on a global scale. I systematically examined advertising abuse both from an external perspective and internally through collaborations with industry partners. My work uncovered the scale, structure, and nature of advertising fraud across multiple monetization strategies, attacks, and botnets.

I began exploring advertising abuse with an execution-driven study of click fraud malware and the supporting ad ecosystem [6]. By iteratively executing malware in isolation with controlled network access, I was able to build tools for automated command-and-control interaction ("milkers"). These tools allowed us to explore the ad abuse ecosystem at a scale not possible with traditional malware execution. This exploration uncovered the breath and scale of the ad fraud ecosystem, as well as the fundamentals of the business model. Despite this new understanding of the ecosystem, our view was still limited to an external ad placement perspective.

To obtain a view of the ad fraud ecosystem at both a larger scale and from an internal perspective, I next focused on an in-depth exploration of ZeroAccess (ZA). ZA was a vast and complex peer-to-peer (P2P) botnet, serving as a delivery platform for advertising abuse malware for more than four years. At its peak, ZA infected more than 1.9 million systems, resulting in millions of dollars in advertising fraud per month. In order to have a qualitatively different view of ad fraud, I collaborated with a top ad-network industry partner which provided me with back-end ad placement data. My work illuminated the rich, intertwined nature of malware-driven click fraud and the advertising ecosystem it exploited, as well as how to develop effective remediations [7]. To conduct a deep analysis, I combined an array of data sources, including P2P measurements, command-and-control telemetry from botnet infiltration, and click information from my ad-network industry partner. Using this multifaceted approach, I identified fraudulent business relationships within the advertising network stemming from complex multi-hop ad reseller chains, which were used as a focal point for remediation. I also quantified the financial impact of ZA's criminal activity.

As part of my efforts to explore the advertising fraud ecosystem I also developed an in-depth technical report on the structure and function of the ZA malware [8]. This work led to collaboration with Microsoft's Digital Crime Unit and law enforcement, with the technical report serving as Exhibit 1 in legal action against the criminal actors. During the height of ZA's spread, I worked with Microsoft and law enforcement to facilitate a technical "takedown" of the botnet, resulting in the demise of ZeroAccess.

A key result from my work exploring ZeroAccess was the complexity of multi-hop ad reseller chains, and how those chains can be used to mask and "launder" fraud. Next, I collaborated with Google to explore this issue broadly, focusing on exploring a different facet of the advertising abuse ecosystem—"ad injection"—which affected *tens of millions* of users [9]. Similar to malware-driven click fraud, ad injection generates revenue through a complex and intertwined ecosystem of intermediaries with opaque business relationships used to launder ad views and clicks. But for ad injection, software that is likely unwanted or has misrepresented its purpose injects ads into the browsing experiences of actual users. My efforts focused on exploring the structure and composition of traffic intermediaries and advertisers that served as the revenue source feeding the injection ecosystem. Our award-winning work enabled remediation for millions of users by identifying structural choke-points of three ad networks and 25 affiliate programs that were responsible for the majority of ad injections.

An important aspect of my advertising abuse work involved the execution of malware at scale. More broadly, it is frequently necessary to perform controlled executions in order to understand the behavior of malware. Throughout my work I have developed methods and systems both for running malware in a safe and controlled manner, and performing analysis of those executions. As part of this effort I have collaborated on the development of malware execution and analysis systems used to quantify the behavior and operation of remote access trojans (RATs)—manually operated malware commonly used for extortion and espionage [10].

Additional Work

Beyond understanding censorship and combating cybercrime, I have undertaken a variety of other work that also has at its heart a focus on providing illumination, empirical grounding, and effective defenses for difficult-to-ascertain global Internet security problems. These efforts have included developing open-source software that facilitates academia and industry to pursue similar work.

Advanced Persistent Threat Targets. Anecdotal industry and initial academic findings suggest state-sponsored advanced persistent threats (APTs) target a variety of organizations ranging from human rights activists to governments. These long-term, covert, highly-targeted threats use sophisticated malware, processes, and reconnaissance to manually compromise networks. Unfortunately we lack a comprehensive picture of who or what is targeted by a single actor or campaign due to the sensitive nature of the victims and the difficulty of identifying targets. Building off of industry reports and my prior global scanning work, I developed an Internet-wide scanning technique leveraging network malware signatures that identifies *all* victims of a specific government APT campaign. Working with law-enforcement, this longitudinal study has identified more than 70 public and private sector targets since 2015. This work is ongoing, with an ultimate goal of publishing a longitudinal study on target selection, recidivism, and the impact of notification.

Mobile Ad Risks. Advertising fuels the mobile application system. Developers use a wide variety of tools and libraries to embed ads and monetize their applications. Unfortunately, these ad systems frequently require additional, potentially risky permissions beyond the core functionality of the application. These additional permissions pose a risk to both user safety and privacy. I performed a measurement study identifying advertising-based over-privileging in Android applications, and built AdDroid [11], a privilege separated advertising framework for the Android platform. My work identified that 27% of measured Android applications were over-privileged due to ad libraries, and 34% of measured applications that used a user's location, did so solely for advertisements. Over-privileging could be reduced by the use of AdDroid for 46% of all advertising-supported applications, demonstrating the potential safety and privacy benefits of advertising aware API frameworks.

Open-Source Software. Throughout my research I have also contributed to a number of open-source network measurement tools. I am a co-author and co-maintainer of the ZMap [12] and ZDNS [13] open-source projects. ZMap

enables rapid scanning of the entire IPv4 address space across a number of protocols, and ZDNS enables large-scale iterative and recursive DNS lookups with normalized well-structured output. Both tools enable researchers and practitioners to perform efficient, sound large-scale Internet measurements.

Future Directions

The impact of global Internet security problems such as censorship, cybercrime, and APTs continues to increase. As systems and data also expand in complexity, the need for empirical grounding is critical. Building on my prior work, I have a number of specific research directions for both the short and long term that aim to change how we think about censorship measurement, state-sponsored attackers, and defending systems.

Longitudinal and Continuous Censorship Measurement. With the development of Augur and Iris, we now have the ability to perform *continuous* censorship measurement around the globe. This capability opens the door to new avenues of research: we can study the trends of censorship within a single country as well as across groups of countries. Using multiple vantage points within countries, we can begin to understand how censorship is deployed and the heterogeneity of deployment, at scale. Using this technology we can identify the onset of new censorship and the dynamics of blocking behavior around key political or social events. The ability to perform controlled repeated measurements also allows us to understand the efficacy of various blocking techniques and evasions, and how censors respond to these evasions. Understanding these facets of censor behavior and technology will allow computer scientists to develop more effective, well-grounded defenses and evasions, while also enabling social scientists and policy makers to study the interactions between users and censors at a scale not previously possible.

Exploring Advanced Persistent Threats. The continued growth of state-sponsored attacks highlights the need for empirical grounding and sound, rigorous understanding of these threats and how best to defend against them. Starting with my ongoing longitudinal work examining APT victims, I will leverage my expertise in malware, Internet scanning, and designing sound empirical methods to explore the landscape of APT attack techniques and victims. Initial work will focus on target selection, recidivism, and the impact of notification, with an ultimate aim of a comprehensive understanding of how to reason about and defend against nation-state adversaries.

Information Sharing. The increasing frequency of large-scale and targeted attacks has companies struggling to identify threats and remediate compromises at the speed and volume they occur. The ability to quickly and reliably share threat information (e.g., IP addresses, malware samples, URLs, tactics) between targets is emerging as a popular proactive defense, and, though of unproven utility, has become current industry best practice. This trend is supported by recent legislation (the Cybersecurity Information Sharing Act), aimed at improving information sharing within the United States.

Within industry, this information sharing is commonly known as *threat intelligence*, an emerging market valued at over three billion US dollars. But despite significant industry capital and supporting legislation, the scope, scale, quality, and efficacy of threat intelligence remains unstudied. In order to answer these questions we need sound, well-designed evaluation metrics and criteria that will allow us to develop effective deployment strategies.

To study this space I have developed a large-scale longitudinal threat intelligence collection system. This system collects well-structured data from more than 200 sources, and has been in operation for over two years. In addition to public threat intelligence data, this system also collects data from several private industry sources, negotiated via my industry contacts.

My work will draw upon this longitudinal data to develop evaluation metrics, collection and curation methodologies, and effective sharing strategies. Ultimately this work aims to change how we defend systems by enabling effective multi-vantage data-driven defenses that do not exist today.

References

[1] Paul Pearce, Roya Ensafi, Frank Li, Nick Feamster, and Vern Paxson. Augur: Internet-Wide Detection of Connectivity Disruptions. In *IEEE Symposium on Security and Privacy (S&P)*, 2017.

[2] Paul Pearce, Roya Ensafi, Frank Li, Nick Feamster, and Vern Paxson. Towards Continual Measurement of Global Network-Level Censorship. In *IEEE Security & Privacy Magazine, Special Issue*, 2018.

[3] Paul Pearce, Ben Jones, Frank Li, Roya Ensafi, Nick Feamster, Nick Weaver, and Vern Paxson. Global Measurement of DNS Manipulation. In *USENIX Security Symposium (USENIX)*, 2017.

[4] Paul Pearce, Ben Jones, Frank Li, Roya Ensafi, Nick Feamster, Nick Weaver, and Vern Paxson. Global Measurement of DNS Manipulation. In *USENIX ;login:*, Winter 2017.

[5] Rachee Singh, Rishab Nithyanand, Sadia Afroz, Paul Pearce, Michael Carl Tschantz, Phillipa Gill, and Vern Paxson. Characterizing the Nature and Dynamics of Tor Exit Blocking. In *USENIX Security Symposium (USENIX)*, 2017.

[6] Brad Miller, Paul Pearce, Chris Grier, Christian Kreibich, and Vern Paxson. What's Clicking What? Techniques and Innovations of Todays Clickbots. In *Detection of Intrusions and Malware & Vulnerability Assessment (DIMVA)*, 2011.

[7] Paul Pearce, Vacha Dave, Chris Grier, Kirill Levchenko, Saikat Guha, Damon McCoy, Vern Paxson, Stefan Savage, and Geoffrey M. Voelker. Characterizing Large-Scale Click Fraud in ZeroAccess. In *ACM Conference on Computer and Communications Security CCS*, 2014.

[8] Paul Pearce, Chris Grier, Vern Paxson, Vacha Dave, Damon McCoy, Goeffrey M. Voelker, and Stefan Savage. The ZeroAccess Auto-Clicking and Search-Hijacking Click Fraud Modules. Technical report, University of California, Berkeley, 2013.

[9] Kurt Thomas, Elie Bursztein, Chris Grier, Grant Ho, Nav Jagpal, Alexandros Kapravelos, Damon McCoy, Antonio Nappa, Vern Paxson, Paul Pearce, Niels Provos, and Moheeb Abu Rajab. Ad Injection at Scale: Assessing Deceptive Advertisement Modifications. In *IEEE Symposium on Security and Privacy (S&P)*, 2015.

[10] Brown Farinholt, Mohammad Rezaeirad, Paul Pearce, Hitesh Dharmdasani, Haikuo Yiny, Stevens Le Blond, Damon McCoy, and Kirill Levchenko. To Catch a Ratter: Monitoring the Behavior of Amateur DarkComet RAT Operators in the Wild. In *IEEE Symposium on Security and Privacy (S&P)*, 2017.

[11] Paul Pearce, Adrienne Porter Felt, Gabriel Nunez, and David Wagner. AdDroid: Privilege Separation for Applications and Advertisers in Android. In *ACM ASIA Conference on Information, Computer and Communications Security (ASIACCS)*, 2012.

[12] ZMap Internet Scanner. `https://github.com/zmap/zmap`.

[13] ZDNS: Fast CLI Utility for Large-Scale DNS Lookups. `https://github.com/zmap/zdns`.