IEEE Papers – Systematization of Knowledge

The goal of this call is to encourage work that evaluates, systematizes, and contextualizes existing knowledge. These papers can provide a high value to our community but may not be accepted because of a lack of novel research contributions. Suitable papers include survey papers that provide useful perspectives on major research areas, papers that support or challenge long-held beliefs with compelling evidence, or papers that provide an extensive and realistic evaluation of competing approaches to solving specific problems. Submissions are encouraged to analyze the current research landscape: identify areas that have enjoyed much research attention, point out open areas with unsolved challenges, and present a prioritization that can guide researchers to make progress on solving important challenges.

2012 Papers

Peek-a-Boo, I Still See You: Why Efficient Traffic Analysis Countermeasures Fail

http://web.cecs.pdx.edu/~teshrim/tmAnotherLook.pdf

Abstract - We consider the setting of HTTP traffic over encrypted tunnels, as used to conceal the identity of websites visited by a user. It is well known that traffic analysis (TA) attacks can accurately identify the website a user visits despite the use of encryption, and previous work has looked at specific attack/countermeasure pairings. We provide the first comprehensive analysis of general-purpose TA countermeasures. We show that nine known countermeasures are vulnerable to simple attacks that exploit coarse features of traffic (e.g., total time and bandwidth). The considered countermeasures include ones like those standardized by TLS, SSH, and IPsec, and even more complex ones like the traffic morphing scheme of Wright et al. As just one of our results, we show that despite the use of traffic morphing, one can use only total upstream and downstream bandwidth to identify —with 98% accuracy— which of two websites was visited. One implication of what we find is that, in the context of website identification, it is unlikely that bandwidth-efficient, general purpose TA countermeasures can ever provide the type of security targeted in prior work.

Third-Party Web Tracking: Policy and Technology

https://www.stanford.edu/~jmayer/papers/trackingsurvey12.pdf

Abstract—In the early days of the web, content was designed and hosted by a single person, group, or organization. No longer. Webpages are increasingly composed of content from myriad unrelated “third-party” websites in the business of advertising,

analytics, social networking, and more. Thirdparty services have tremendous value: they support free content and facilitate web innovation. But third-party services come at a privacy cost: researchers, civil society organizations, and policymakers have increasingly called attention to how third parties can track a user's browsing activities across websites. This paper surveys the current policy debate surrounding third-party web tracking and explains the relevant technology. It also presents the FourthParty web measurement platform and studies we have conducted with it. Our aim is to inform researchers with essential background and tools for contributing to public understanding and policy debates about web tracking.

Dissecting Android Malware: Characterization and Evolution

http://www.csc.ncsu.edu/faculty/jiang/pubs/OAKLAND12.pdf

Abstract—The popularity and adoption of smartphones has greatly stimulated the spread of mobile malware, especially on the popular platforms such as Android. In light of their rapid growth, there is a pressing need to develop effective solutions. However, our defense capability is largely constrained by the limited understanding of these emerging mobile malware and the lack of timely access to related samples. In this paper, we focus on the Android platform and aim to systematize or characterize existing Android malware. Particularly, with more than one year effort, we have managed to collect more than 1,200 malware samples that cover the majority of existing Android malware families, ranging from their debut in August 2010 to recent ones in October 2011. In addition, we systematically characterize them from various aspects, including their installation methods, activation mechanisms as well as the nature of carried malicious payloads. The characterization and a subsequent evolution-based study of representative families reveal that they are evolving rapidly to circumvent the detection from existing mobile anti-virus software. Based on the evaluation with four representative mobile security software, our experiments show that the best case detects 79.6% of them while the worst case detects *only* 20.2% in our dataset. These results clearly call for the need to better develop next-generation anti-mobile-malware solutions.

Prudent Practices for Designing Malware Experiments: Status Quo and Outlook

http://www.icsi.berkeley.edu/icsi/publication_details?n=3298

Abstract—Malware researchers rely on the observation of malicious code in execution to collect datasets for a wide array of experiments, including generation of detection models, study of longitudinal behavior, and validation of prior research. For such research to reflect prudent science, the work needs to address a number of concerns relating to the correct and representative use of the datasets, presentation of methodology in a fashion sufficiently transparent to enable reproducibility, and due consideration of the need not to harm others. In this paper we study the methodological rigor and prudence in 36 academic publications from 2006–2011 that rely on malware execution. 40% of these papers appeared in the 6 highest-ranked academic security conferences. We

find frequent shortcomings, including problematic assumptions regarding the use of execution-driven datasets (25% of the papers), absence of description of security precautions taken during experiments (71% of the articles), and oftentimes insufficient description of the experimental setup. Deficiencies occur in top-tier venues and elsewhere alike, highlighting a need for the community to improve its handling of malware datasets. In the hope of aiding authors, reviewers, and readers, we frame guidelines regarding transparency, realism, correctness, and safety for collecting and using malware datasets.

The Psychology of Security for the Home Computer User
http://www.cs.colostate.edu/psysec/papers/SSP.pdf

Abstract—The home computer user is often said to be the weakest link in computer security. They do not always follow security advice, and they take actions, as in phishing, that compromise themselves. In general, we do not understand why users do not always behave safely, which would seem to be in their best interest. This paper reviews the literature of surveys and studies of factors that influence security decisions for home computer users. We organize the review in four sections: understanding of threats, perceptions of risky behavior, efforts to avoid security breaches and attitudes to security interventions. We find that these studies reveal a lot of reasons why current security measures may not match the needs or abilities of home computer users and suggest future work needed to inform how security is delivered to this user group.

2011 Papers

A Formal Foundation for the Security Features of Physical Functions
https://www.cosic.esat.kuleuven.be/publications/article-2012.pdf

Abstract—Physical attacks against cryptographic devices typically take advantage of information leakage (e.g., sidechannels attacks) or erroneous computations (e.g., fault injection attacks). Preventing or detecting these attacks has become a challenging task in modern cryptographic research. In this context intrinsic physical properties of integrated circuits, such as Physical(ly) Unclonable Functions (PUFs), can be used to complement classical cryptographic constructions, and to enhance the security of cryptographic devices. PUFs have recently been proposed for various applications, including anticounterfeiting schemes, key generation algorithms, and in the design of block ciphers. However, currently only rudimentary security models for PUFs exist, limiting the confidence in the security claims of PUF-based security primitives. A useful model should at the same time (i) define the security properties of PUFs abstractly and naturally, allowing to design and formally analyze PUF-based security solutions, and (ii) provide practical quantification tools allowing engineers to evaluate PUF instantiations.

In this paper, we present a formal foundation for security primitives based on PUFs. Our approach requires as little as possible from the physics and focuses more on the main

properties at the heart of most published works on PUFs: robustness (generation of stable answers), unclonability (not provided by algorithmic solutions), and unpredictability. We first formally define these properties and then show that they can be achieved by previously introduced PUF instantiations. We stress that such a consolidating work allows for a meaningful security analysis of security primitives taking advantage of physical properties, becoming increasingly important in the development of the next generation secure information systems.

Timing- and Termination-Sensitive Secure Information Flow: Exploring a New Approach

http://www.cs.ucsb.edu/~benh/papers/kashyap11timing.pdf

Abstract—Secure information flow guarantees the secrecy and integrity of data, preventing an attacker from learning secret information (secrecy) or injecting untrusted information (integrity). Covert channels can be used to subvert these security guarantees; for example, timing and termination channels can, either intentionally or inadvertently, violate these guarantees by modifying the timing or termination behavior of a program based on secret or untrusted data. Attacks using these covert channels have been published and are known to work in practice—as techniques to prevent non-covert channels are becoming increasingly practical, covert channels are likely to become even more attractive for attackers to exploit. The goal of this paper is to understand the subtleties of timing- and termination-sensitive noninterference, explore the space of possible strategies for enforcing noninterference guarantees, and formalize the exact guarantees that these strategies can enforce. As a result of this effort we create a novel strategy that provides stronger security guarantees than existing work, and we clarify claims in existing work about what guarantees can be made.

Formalizing Anonymous Blacklisting Systems

http://cacr.uwaterloo.ca/techreports/2010/cacr2010-24.pdf

Abstract—Anonymous communications networks, such as Tor, help to solve the real and important problem of enabling users to communicate privately over the Internet. However, in doing so, anonymous communications networks introduce an entirely new problem for the service providers—such as websites, IRC networks or mail servers—with which these users interact; in particular, since all anonymous users look alike, there is no way for the service providers to hold individual misbehaving anonymous users accountable for their actions. Recent research efforts have focused on using anonymous blacklisting systems (which are sometimes called anonymous revocation systems) to empower service providers with the ability to revoke access from abusive anonymous users. In contrast to revocable anonymity systems, which enable some trusted third party to deanonymize users, anonymous blacklisting systems provide users with a way to authenticate anonymously with a service provider, while enabling the service provider to revoke access from any users that misbehave, without revealing their identities. In this paper, we introduce the anonymous blacklisting problem and survey the literature on

anonymous blacklisting systems, comparing and contrasting the architecture of various existing schemes, and discussing the tradeoffs inherent with each design. The literature on anonymous blacklisting systems lacks a unified set of definitions; each scheme operates under different trust assumptions and provides different security and privacy guarantees. Therefore, before we discuss the existing approaches in detail, we first propose a formal definition for anonymous blacklisting systems, and a set of security and privacy properties that these systems should possess. We also outline a set of new performance requirements that anonymous blacklisting systems should satisfy to maximize their potential for real-world adoption, and give formal definitions for several optional features already supported by some schemes in the literature.

Mobile Security Catching Up? - Revealing the Nuts and Bolts of the Security of Mobile Devices

http://pi1.informatik.uni-mannheim.de/filepool/publications/mobile-security-oakland-2011.pdf

Abstract—We are currently moving from the Internet society to a mobile society where more and more access to information is done by previously dumb phones. For example, the number of mobile phones using a full blown OS has risen to nearly 200% from Q3/2009 to Q3/2010. As a result, mobile security is no longer immanent, but imperative. This survey paper provides a concise overview of mobile network security, attack vectors using the back end system and the web browser, but also the hardware layer and the user as attack enabler. We show differences and similarities between "normal" security and mobile security, and draw conclusions for further research opportunities in this area.

2010 Papers

State of the Art: Automated Black-BoxWeb Application Vulnerability Testing

http://crypto.stanford.edu/~jcm/papers/pci_oakland10.pdf

Abstract—Black-box web application vulnerability scanners are automated tools that probe web applications for security vulnerabilities. In order to assess the current state of the art, we obtained access to eight leading tools and carried out a study of: (i) the class of vulnerabilities tested by these scanners, (ii) their effectiveness against target vulnerabilities, and (iii) the relevance of the target vulnerabilities to vulnerabilities found in the wild. To conduct our study we used a custom web application vulnerable to known and projected vulnerabilities, and previous versions of widely used web applications containing known vulnerabilities. Our results show the promise and effectiveness of automated tools, as a group, and also some limitations. In particular, "stored" forms of

Cross Site Scripting (XSS) and SQL Injection (SQLI) vulnerabilities are not currently found by many tools. Because our goal is to assess the potential of future research, not to evaluate specific vendors, we do not report comparative data or make any recommendations about purchase of specific tools.

Outside the Closed World: On Using Machine Learning For Network Intrusion Detection
http://www.icsi.berkeley.edu/icsi/publication_details?n=2818

Abstract—In network intrusion detection research, one popular strategy for finding attacks is monitoring a network's activity for *anomalies*: deviations from profiles of normality previously learned from benign traffic, typically identified using tools borrowed from the machine learning community. However, despite extensive academic research one finds a striking gap in terms of actual deployments of such systems: compared with other intrusion detection approaches, machine learning is rarely employed in operational "real world" settings. We examine the differences between the network intrusion detection problem and other areas where machine learning regularly finds much more success. Our main claim is that the task of finding attacks is fundamentally different from these other applications, making it significantly harder for the intrusion detection community to employ machine learning effectively. We support this claim by identifying challenges particular to network intrusion detection, and provide a set of guidelines meant to strengthen future research on anomaly detection.

All You Ever Wanted to Know About Dynamic Taint Analysis and Forward Symbolic Execution (but might have been afraid to ask)
http://users.ece.cmu.edu/~ejschwar/papers/oakland10.pdf

Abstract—Dynamic taint analysis and forward symbolic execution are quickly becoming staple techniques in security analyses. Example applications of dynamic taint analysis and forward symbolic execution include malware analysis, input filter generation, test case generation, and vulnerability discovery. Despite the widespread usage of these two techniques, there has been little effort to formally define the algorithms and summarize the critical issues that arise when these techniques are used in typical security contexts.

The contributions of this paper are two-fold. First, we precisely describe the algorithms for dynamic taint analysis and forward symbolic execution as extensions to the run-time semantics of a general language. Second, we highlight important implementation choices, common pitfalls, and considerations when using these techniques in a security context.

Bootstrapping Trust in Commodity Computers
https://sparrow.ece.cmu.edu/group/pub/PaMcPe2010.pdf

Abstract - Trusting a computer for a security-sensitive task (such as checking email or banking online) requires the user to know something about the computer's state. We examine research on securely capturing a computer's state, and consider the utility of this

information both for improving security on the local computer (e.g., to convince the user that her computer is not infected with malware) and for communicating a remote computer's state (e.g., to enable the user to check that a web server will adequately protect her data). Although the recent "Trusted Computing" initiative has drawn both positive and negative attention to this area, we consider the older and broader topic of bootstrapping trust in a computer. We cover issues ranging from the wide collection of secure hardware that can serve as a foundation for trust, to the usability issues that arise when trying to convey computer state information to humans. This approach unifies disparate research efforts and highlights opportunities for additional work that can guide real-world improvements in computer security.

How Good are Humans at Solving CAPTCHAs? A Large Scale Evaluation

http://theory.stanford.edu/~jcm/papers/captcha-study-oakland10.pdf

Abstract—Captchas are designed to be easy for humans but hard for machines. However, most recent research has focused only on making them hard for machines. In this paper, we present what is to the best of our knowledge the first large-scale evaluation of captchas from the human perspective, with the goal of assessing how much friction captchas present to the average user. For the purpose of this study we have asked workers from Amazon's Mechanical Turk and an underground captchabreaking service to solve more than 318 000 captchas issued from the 21 most popular captcha schemes (13 images schemes and 8 audio scheme). Analysis of the resulting data reveals that captchas are often difficult for humans, with audio captchas being particularly problematic. We also find some demographic trends indicating, for example, that non-native speakers of English are slower in general and less accurate on English-centric captcha schemes. Evidence from a week's worth of eBay captchas (14,000,000 samples) suggests that the solving accuracies found in our study are close to real-world values, and that improving audio captchas should become a priority, as nearly 1% of all captchas are delivered as audio rather than images. Finally our study also reveals that it is more effective for an attacker to use Mechanical Turk to solve captchas than an underground service.