Articles for Oral Presentation

COMPSCI 725, S2 2016
Version 1.01 of 2016-07-22: added alternate download URLs for unregistered DOIs
Clark Thomborson

Instructions:

1. Read through this document, to discover three articles which you'd be willing to present orally in COMPSCI 725.
2. Send an email to me on cthombor@cs.auckland.ac.nz, listing your three choices by **surname** of the **first** author and the **year of publication** e.g. "Fagan 2016, Kang 2016, Kohls 2016". Hint: use the bracketed information before the article's bibliographic details. List your highest preference first.
3. Be careful to include your UPI in all course-related email! Your instructor may be unable to guess your University identity from your name and external email address.

Process notes:

1. I will assign up to three students to an article, using a **first-come-first-served** (FCFS) allocation strategy. I'll make assignments approximately daily, at some unpredictable time each day.
2. You may delay sending your email until the end of the enrolment period – on Friday of the second week of lectures. However, because I'm assigning by FCFS, such delay will lessen your chances of being assigned your first preference.
3. I am currently trying to discover some provision in Canvas, whereby each of you could self-schedule your oral presentation day, on a first-come-first-served basis. If I find a way to do this[1], I'll let you know, in an announcement by email and on https://www.cs.auckland.ac.nz/courses/compsci725s2c/lectures/.

Usable Security

1. [Fagan 2016] Michael Fagan and Mohammad Maifi Hasan Khan. "Why Do They Do What They Do?: A Study of What Motivates Users to (Not) Follow Computer Security Advice". In *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016),* USENIX Association, pp. 59-75, 2016. [Download]

   Abstract: Usable security researchers have long been interested in what users do to keep their devices and data safe and how that compares to recommendations. Additionally, experts have long debated and studied the psychological underpinnings and motivations for users to do what they do, especially when such behavior is seen as risky, at least to experts. This study investigates user motivations through a survey conducted on Mechanical Turk, which resulted in responses from 290 participants. We use a rational decision model to guide our design, as well as current thought on human motivation in general and in the

[1] The Canvas affordances for instructors are poorly documented and are unstable over time. Instability of feature and affordance is increasingly-often portrayed as a desirable attribute, rather than as a design defect, in computer systems. See e.g. https://www.brown.edu/it/canvas/: "… consistent introduction of new functionality without major upgrades so less disruption for users…" Such "consistent" (and mostly imperceptible) instability is loads of fun for the security team, as well as for the users. Never a dull moment in a fast-changing world!

realm of computer security. Through quantitative and qualitative analysis, we identify key gaps in perception between those who follow common security advice (i.e., update software, use a password manager, use 2FA, change passwords) and those who do not and help explain participants' motivations behind their decisions. Additionally, we find that social considerations are trumped by individualized rationales.

2. [Kang 2016] Ruogu Kang, Laura Dabbish et al. "'My Data Just Goes Everywhere': User Mental Models of the Internet and Implications for Privacy and Security". In *Eleventh Symposium on Usable Privacy and Security (SOUPS 2015),* USENIX Association, pp. 39-52, 2015. [Download]

   Abstract: Many people use the Internet every day yet know little about how it really works. Prior literature diverges on how people's Internet knowledge affects their privacy and security decisions. We undertook a qualitative study to understand what people do and do not know about the Internet and how that knowledge affects their responses to privacy and security risks. Lay people, as compared to those with computer science or related backgrounds, had simpler mental models that omitted Internet levels, organizations, and entities. People with more articulated technical models perceived more privacy threats, possibly driven by their more accurate understanding of where specific risks could occur in the network. Despite these differences, we did not find a direct relationship between people's technical background and the actions they took to control their privacy or increase their security online. Consistent with other work on user knowledge and experience, our study suggests a greater emphasis on policies and systems that protect privacy and security without relying too much on users' security practices.

3. [Sundaramurthy 2016]. Sathya Chandran Sundaramurthy, John McHugh, et al. Turning Contradictions into Innovations or: How We Learned to Stop Whining and Improve Security Operations. In *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016),* USENIX Association, pp. 237-251, 2016. [Download]

   Abstract: Efforts to improve the efficiency of security operation centers (SOCs) have emphasized building tools for analysts or understanding the human and organizational factors involved. The importance of viewing the viability of a solution from multiple perspectives has been largely ignored. Multiple perspectives arise because of inherent conflicts among the objectives a SOC has to meet and differences between the goals of the parties involved. During the 3.5 years that we have used anthropological fieldwork methods to study SOCs, we discovered that successful SOC innovations must resolve these conflicts to be effective in improving operational efficiency. This discovery was guided by Activity Theory (AT), which provided a framework for analyzing our fieldwork data. We use the version of AT proposed by Engestrom to model SOC operations. Template analysis, a qualitative data analysis technique, guided by AT validated the existence of contradictions in SOCs. The same technique was used to elicit from the data concrete contradictions and how they were resolved. Our analysis provide evidence of the importance of conflict resolution as a prerequisite for operations improvement. AT enabled us to understand why some of our innovations worked in the SOCs we studied (and why others failed). AT helps us see a potentially successful and repeatable mechanism for introducing new

technologies to future SOCs. Understanding and supporting all of the spoken and unspoken requirements of SOC analysts and managers appears to be the only way to get new technologies accepted and used in SOCs.

Obfuscation, Stenography, Tamperproofing

4. [Bernstein 2015] Daniel Bernstein, Andreas Hülsing, et al. "Bad Directions in Cryptographic Hash Functions." In *Information Security and Privacy: 20th Australasian Conference (ACISP 2015)*, pp. 488-508. Springer, 2015. [Download]

   Abstract: A 25-gigabyte "point obfuscation" challenge "using security parameter 60" was announced at the Crypto 2014 rump session; "point obfuscation" is another name for password hashing. This paper shows that the particular matrix-multiplication hash function used in the challenge is much less secure than previous password-hashing functions are believed to be. This paper's attack algorithm broke the challenge in just 19 minutes using a cluster of 21 PCs.

5. [Hoffman 2016] Johannes Hoffman, Teemu Rytilahti, et al. "Evaluating Analysis Tools for Android Apps: Status Quo and Robustness Against Obfuscation". In *Proceedings of the Sixth ACM Conference on Data and Application Security and Privacy (CODASPY'16)*, pp. 139—141. ACM, 2016. [Download]

   Abstract. The recent past has shown that Android smartphones became the most popular target for malware authors. Malware families offer a variety of features that allow, among the others, to steal arbitrary data and to cause significant monetary losses. This circumstances led to the development of many different analysis methods that are aimed to assess the absence of potential harm or malicious behavior in mobile apps. In return, malware authors devised more sophisticated methods to write mobile malware that attempt to thwart such analyses. In this work, we briefly describe assumptions analysis tools rely on to detect malicious content and behavior. We then present results of a new obfuscation framework that aims to break such assumptions, thus modifying Android apps to avoid them being analyzed by the targeted systems. We use our framework to evaluate the robustness of static and dynamic analysis systems for Android apps against such transformations.

6. [Kohls 2016] Katherina Kohls, Thrsten Holz, et al. "Skypeline: Robust hidden data transmission for VoIP." In *Proceedings of the 11th ACM on Asia Conference on Computer and Communications Security (CCS'16)*, pp. 887-888. ACM, 2016. [Download]

   Abstract: Internet censorship is used in many parts of the world to prohibit free access to online information. Different techniques such as IP address or URL blocking, DNS hijacking, or deep packet inspection are used to block access to specific content on the Internet. In response, several censorship circumvention systems were proposed that attempt to bypass existing filters. Especially systems that hide the communication in different types of cover protocols attracted a lot of attention. However, recent research results suggest that this kind of covert traffic can be easily detected by censors. In this paper, we present SkypeLine, a censorship circumvention system that leverages Direct-Sequence Spread

Spectrum (DSSS) based steganography to hide information in Voice-over-IP (VoIP) communication. SkypeLine introduces two novel modulation techniques that hide data by modulating information bits on the voice carrier signal using pseudo-random, orthogonal noise sequences and repeating the spreading operation several times. Our design goals focus on undetectability in presence of a strong adversary and improved data rates. As a result, the hiding is inconspicuous, does not alter the statistical characteristics of the carrier signal, and is robust against alterations of the transmitted packets. We demonstrate the performance of SkypeLine based on two simulation studies that cover the theoretical performance and robustness. Our measurements demonstrate that the data rates achieved with our techniques substantially exceed existing DSSS approaches. Furthermore, we prove the real-world applicability of the presented system with an exemplary prototype for Skype.

7. [Ming 2016] Jiang Ming, Zhi Xin, et al. "Replacement Attacks: Automatically Impeding Behavior-Based Malware Specifications." In *Applied Cryptography and Network Security (ACNS 2015)*, LNCS 9092, Springer-Verlag, pp. 497-517, 2016. [Download]

   Abstract: As the underground market of malware flourishes, there is an exponential increase in the number and diversity of malware. A crucial question in malware analysis research is how to define malware specifications or signatures that faithfully describe similar malicious intent and clearly stand out from other programs. It is evident that the classical syntactic signatures are insufficient to defeat state-of-the art malware. Behavior-based specifications which capture real malicious characteristics during runtime, have become more prevalent in anti-malware tasks, such as malware detection and malware clustering. This kind of specification is typically extracted from system call dependence graphs that a malware sample invokes. In this paper we present replacement attacks to poison behavior-based specifications by concealing similar behaviors among malware variants. The essence of the attacks is to replace a behavior specification to its semantically equivalent one, so that similar malware variants within one family turn out to be different. As a result, malware analysts have to put more efforts to re-analyze similar samples. We distill general attacking strategies by mining more than 5,000 malware samples' behavior specifications and implement a compiler-level prototype to automate replacement attacks. Experiments on 960 real malware samples demonstrate effectiveness of our approach to impede multiple malware analyses based on behavior specifications, such as similarity comparison and malware clustering. In the end, we provide possible counter-measures to strengthen behavior-based malware analysis.

8. [Rasthofer 2016] Siegfried Rasthofer, Steven Arzt, et al. "Harvesting Runtime Values in Android Applications That Feature Anti-Analysis Techniques". In *Proceedings of the Annual Network and Distributed System Security Symposium* (NDSS), The Internet Society, 2016. [Download]

   Abstract: It is generally challenging to tell apart malware from benign applications. To make this decision, human analysts are frequently interested in runtime values: targets of reflective method calls, URLs to which data is sent, target telephone numbers of SMS messages, and many more. However,

obfuscation and string encryption, used by malware as well as goodware, often not only render human inspections, but also static analyses ineffective. In addition, malware frequently tricks dynamic analyses by detecting the execution environment emulated by the analysis tool and then refraining from malicious behavior.

In this work we therefore present HARVESTER, an approach to fully automatically extract runtime values from Android applications. HARVESTER is designed to extract values even from highly obfuscated state-of-the-art malware samples that obfuscate method calls using reflection, hide sensitive values in native code, load code dynamically and apply anti-analysis techniques. The approach combines program slicing with code generation and dynamic execution.

Experiments on 16,799 current malware samples show that HARVESTER fully automatically extracts many sensitive values, with perfect precision. The process usually takes less than three minutes and does not require human interaction. In particular, it goes without simulating UI inputs. Two case studies further show that by integrating the extracted values back into the app, HARVESTER can increase the recall of existing static and dynamic analysis tools such as FlowDroid and TaintDroid.

Privacy

9. [Avancha 2012] Sasikanth Avancha, Amit Baxi, and David Kotz. "Privacy in Mobile Technology for Personal Healthcare". *ACM Computing Surveys 45(1)*, pp. 3:1-3-54. [Download]

    Abstract: Information technology can improve the quality, efficiency, and cost of healthcare. In this survey, we examine the privacy requirements of mobile computing technologies that have the potential to transform healthcare. Such mHealth technology enables physicians to remotely monitor patients' health and enables individuals to manage their own health more easily. Despite these advantages, privacy is essential for any personal monitoring technology. Through an extensive survey of the literature, we develop a conceptual privacy framework for mHealth, itemize the privacy properties needed in mHealth systems, and discuss the technologies that could support privacy-sensitive mHealth systems. We end with a list of open research questions.

10. [Dinev 2013] Tamara Dinev, Heng Xu, et al. "Information privacy and correlates: an empirical attempt to bridge and distinguish privacy-related concepts". *European Journal of Information Systems 22(3)*, pp. 295–316, 2013. [Download] [Alternate download]

    Abstract: Privacy is one of the few concepts that has been studied across many disciplines, but is still difficult to grasp. The current understanding of privacy is largely fragmented and discipline-dependent. This study develops and tests a framework of information privacy and its correlates, the latter often being confused with or built into definitions of information privacy per se. Our framework development was based on the privacy theories of Westin and Altman, the economic view of the privacy calculus, and the identity management

framework of Zwick and Dholakia. The dependent variable of the model is perceived information privacy. The particularly relevant correlates to information privacy are anonymity, secrecy, confidentiality, and control. We posit that the first three are tactics for information control; perceived information control and perceived risk are salient determinants of perceived information privacy; and perceived risk is a function of perceived benefits of information disclosure, information sensitivity, importance of information transparency, and regulatory expectations. The research model was empirically tested and validated in the Web 2.0 context, using a survey of Web 2.0 users. Our study enhances the theoretical understanding of information privacy and is useful for privacy advocates, and legal, management information systems, marketing, and social science scholars.

11. [Englehardt 2016] Steven Englehardt and Arvind Narayan. "Online tracking: A 1-million-site measurement and analysis: Draft July 11th 2016", web manuscript published by author Arvind Narayan of Princeton University. Available: http://randomwalker.info/publications/OpenWPM_1_million_site_tracking_measurement.pdf, 22 July 2016. [Download]

    Abstract: We present the largest and most detailed measurement of online tracking conducted to date, based on a crawl of the top 1 million websites. We make 15 types of measurements on each site, including stateful (cookie-based) and stateless (fingerprinting-based) tracking, the effect of browser privacy tools, and the exchange of tracking data between different sites ("cookie syncing"). Our findings include multiple sophisticated fingerprinting techniques never before measured in the wild.

    This measurement is made possible by our web privacy measurement tool, OpenWPM, which uses an automated version of a full-fledged consumer browser. It supports parallelism for speed and scale, automatic recovery from failures of the underlying browser, and comprehensive browser instrumentation. OpenWPM is open-source and has already been used as the basis of seven published studies on web privacy and security.

12. [Khattak 2016] Sheharbano Khattak, David Fifield, et al. "Do You See What I See? Differential Treatment of Anonymous Users". In *Proceedings of the Annual Network and Distributed System Security Symposium* (NDSS), The Internet Society, 2016. [Download]

    Abstract: The utility of anonymous communication is undermined by a growing number of websites treating users of such services in a degraded fashion. The second-class treatment of anonymous users ranges from outright rejection to limiting their access to a subset of the service's functionality or imposing hurdles such as CAPTCHA-solving. To date, the observation of such practices has relied upon anecdotal reports catalogued by frustrated anonymity users. We present a study to methodically enumerate and characterize, in the context of Tor, the treatment of anonymous users as second-class Web citizens.

    We focus on first-line blocking: at the transport layer, through reset or dropped connections; and at the application layer, through explicit blocks served from website home pages. Our study draws upon several data sources: comparisons of

Internet-wide port scans from Tor exit nodes versus from control hosts; scans of the home pages of top-1,000 Alexa websites through every Tor exit; and analysis of nearly a year of historic HTTP crawls from Tor network and control hosts. We develop a methodology to distinguish censorship events from incidental failures such as those caused by packet loss or network outages, and incorporate consideration of the endemic churn in web-accessible services over both time and geographic diversity. We find clear evidence of Tor blocking on the Web, including 3.67% of the top-1,000 Alexa sites. Some blocks specifically target Tor, while others result from fate-sharing when abuse-based automated blockers trigger due to misbehaving Web sessions sharing the same exit node.

13. [Mayer 2016] Jonathan Mayer, Patrick Mutchler, and John C. Mitchell. "Evaluating the privacy properties of telephone metadata". *PNAS 2016 113 (20)*, pp. 5536-5541. [Download]

    Abstract: Privacy protections against government surveillance are often scoped to communications content and exclude communications metadata. In the United States, the National Security Agency operated a particularly controversial program, collecting bulk telephone metadata nationwide. We investigate the privacy properties of telephone metadata to assess the impact of policies that distinguish between content and metadata. We find that telephone metadata is densely interconnected, can trivially be reidentified, enables automated location and relationship inferences, and can be used to determine highly sensitive traits.

Hacks and Cracks

14. [Genkin 2016] Daniel Genkin, Lev Pachmanov, et al. "Physical Key Extraction Attacks on PCs". *Commun. ACM 59(60)*, pp. 70-79, 2016. [Download]

    Abstract: Computers broadcast their secrets via inadvertent physical emanations that are easily measured and exploited.

15. [Humbert 2015] Chen Wang, Xiaonan Guo, et al. "Friend or Foe?: Your Wearable Devices Reveal Your Personal PIN". *Proc. of the 11th ACM Asia Conference on Computer and Communications Security (ASIA CCS '16)*, pp. 189-200, 2016. [Download]

    Abstract: The proliferation of wearable devices, e.g., smartwatches and activity trackers, with embedded sensors has already shown its great potential on monitoring and inferring human daily activities. This paper reveals a serious security breach of wearable devices in the context of divulging secret information (i.e., key entries) while people accessing key-based security systems. Existing methods of obtaining such secret information relies on installations of dedicated hardware (e.g., video camera or fake keypad), or training with labeled data from body sensors, which restrict use cases in practical adversary scenarios. In this work, we show that a wearable device can be exploited to discriminate mm-level distances and directions of the user's fine-grained hand movements, which enable attackers to reproduce the trajectories of the user's hand and further to recover the secret key entries. In particular, our system confirms the possibility of using embedded sensors in wearable devices, i.e., accelerometers, gyroscopes, and

magnetometers, to derive the moving distance of the user's hand between consecutive key entries regardless of the pose of the hand. Our Backward PIN-Sequence Inference algorithm exploits the inherent physical constraints between key entries to infer the complete user key entry sequence. Extensive experiments are conducted with over 5000 key entry traces collected from 20 adults for key-based security systems (i.e. ATM keypads and regular keyboards) through testing on different kinds of wearables. Results demonstrate that such a technique can achieve 80% accuracy with only one try and more than 90% accuracy with three tries, which to our knowledge, is the first technique that reveals personal PINs leveraging wearable devices without the need for labeled training data and contextual information.

16. [Lettner 2016] Julian Lettner, Benjamin Kollenda, et al. "Subversive-C: Abusing and Protecting Dynamic Message Dispatch". *2016 USENIX Annual Technical Conference (USENIX ATC 16),* pp. 209-221, 2016. [Download]

    Abstract: The lower layers in the modern computing infrastructure are written in languages threatened by exploitation of memory management errors. Recently deployed exploit mitigations such as control-flow integrity (CFI) can prevent traditional return-oriented programming (ROP) exploits but are much less effective against newer techniques such as Counterfeit Object-Oriented Programming (COOP) that execute a chain of C++ virtual methods. Since these methods are valid control-flow targets, COOP attacks are hard to distinguish from benign computations. Code randomization is likewise ineffective against COOP. Until now, however, COOP attacks have been limited to vulnerable C++ applications which makes it unclear whether COOP is as general and portable a threat as ROP.

    This paper demonstrates the first COOP-style exploit for Objective-C, the predominant programming language on Apple's OS X and iOS platforms. We also retrofit the Objective-C runtime with the first practical and efficient defense against our novel attack. Our defense is able to protect complex, real-world software such as iTunes without recompilation. Our performance experiments show that the overhead of our defense is low in practice.

17. [Starov 2016] Oleksii Starov, Johannes Dahse, et al. "No honor among thieves: A large-scale analysis of malicious web shells". Proc. of the 25th International Conference on World Wide Web, pp. 1021-1032, 2016. [Download] [Alternate download]

    Abstract: Web shells are malicious scripts that attackers upload to a compromised web server in order to remotely execute arbitrary commands, maintain their access, and elevate their privileges. Despite their high prevalence in practice and heavy involvement in security breaches, web shells have never been the direct subject of any study. In contrast, web shells have been treated as malicious blackboxes that need to be detected and removed, rather than malicious pieces of software that need to be analyzed and, in detail, understood. In this paper, we report on the first comprehensive study of web shells. By utilizing different static and dynamic analysis methods, we discover and quantify the visible and invisible features offered by popular malicious shells, and we discuss how attackers can

take advantage of these features. For visible features, we find the presence of password bruteforcers, SQL database clients, portscanners, and checks for the presence of security software installed on the compromised server. In terms of invisible features, we find that about half of the analyzed shells contain an authentication mechanism, but this mechanism can be bypassed in a third of the cases. Furthermore, we find that about a third of the analyzed shells perform homephoning, i.e., the shells, upon execution, surreptitiously communicate to various third parties with the intent of revealing the location of new shell installations. By setting up honeypots, we quantify the number of third-party attackers benefiting from shell installations and show how an attacker, by merely registering the appropriate domains, can completely take over all installations of specific vulnerable shells.

18. [Zeltmann 2016] Steven Eric Zeltmann, Nikhil Gupta, et al. "Manufacturing and Security Challenges in 3D Printing", *JOM 68(7)*, pp. 1872-1881, 2016. [Download]

Abstract: As the manufacturing time, quality, and cost associated with additive manufacturing (AM) continue to improve, more and more businesses and consumers are adopting this technology. Some of the key benefits of AM include customizing products, localizing production and reducing logistics. Due to these and numerous other benefits, AM is enabling a globally distributed manufacturing process and supply chain spanning multiple parties, and hence raises concerns about the reliability of the manufactured product. In this work, we first present a brief overview of the potential risks that exist in the cyber-physical environment of additive manufacturing. We then evaluate the risks posed by two different classes of modifications to the AM process which are representative of the challenges that are unique to AM. The risks posed are examined through mechanical testing of objects with altered printing orientation and fine internal defects. Finite element analysis and ultrasonic inspection are also used to demonstrate the potential for decreased performance and for evading detection. The results highlight several scenarios, intentional or unintentional, that can affect the product quality and pose security challenges for the additive manufacturing supply chain.