2017 HOTSOS AGENDA

http://hot-sos.com

Tuesday, April 4

0800 - 0730 **Registration**

0830 - 0900 **Welcome and Introduction**

0900 - 1000 **Keynote: Security of Cyber-Physical Systems: Challenges and Approaches**
Insup Lee, UPenn

1000 - 1030 **Posters and Break**

1030 - 1200 **SESSION 2**

Paper: Leveraging Unique CPS Properties to Design Better Privacy-Enhancing Algorithms
Jairo Giraldo, Alvaro A. Cárdenas, Murat Kantarcioglu
University of Texas at Dallas

Paper: Characterizing Scientific Reporting in Security Literature: An analysis of ACM CCS and IEEE S&P Papers
**Morgan Burcham, *Mahran Al-Zyoud, *Jeffrey C. Carver, †Mohammed Alsaleh, **Hongying Du, †Fida Gilani, ‡Jun Jiang, **Akond Rahman; **Özgür Kafalı, †Ehab Al-Shaer, **Laurie Williams*
*University of Alabama, †UNC Charlotte,
**NC State University, ‡UNC Chapel Hill

Paper: Optimal Security Investments in a Prevention and Detection Game
Carlos Barreto, Alvaro A. Cárdenas, Alain Bensoussan
University of Texas at Dallas

1200 - 1330 **Lunch *(on your own)***

1330 - 1600 **SESSION 3**

1330 - 1430 **Keynote Presentation**
Jules Polonetsky, Future of Privacy Forum

1430 - 1530 **TBA**
TBA

1530 - 1600 **Posters and Break**

1600 - 1730 **SESSION 4 - Parallel Tutorial Session**

The Bugs Framework (BF) "Hands-On"
Irena Bojanova and Paul E. Black
National Intitute of Standards and Technology

System Monitoring for Security
Xusheng Xiao, Case Western Reserve University

1730 - 1800 **BREAK**

1800 **Shuttle to American Visionary Art Museum**
(Meet in the Hotel Lobby - Don't be late!)

1830 - 2100 **Dinner - American Visionary Art Museum**

Wednesday, April 5

0800 - 0530 **Registration**

0900 - 1200 **SESSION 5**

0900 - 1000 **Keynote: Differential Privacy and Data Analysis**
Aaron Roth, UPenn

1000 - 1030 **Posters and Break**

1030 - 1330 **SESSION 6**

1030 - 1200 **Paper: Learning a Privacy Incidents Database**
**Pradeep K. Murukannaiah, †Chinmaya Dabral, †Karthik Sheshadri, †Esha Sharma, †Jessica Staddon*
*Rochester Institute of Technology, †NC State University

Paper: On the Trade-off Between Privacy and Utility in Collaborative Intrusion Detection Systems-A Game Theoretical Approach
**Richeng Jin, †Xiaofan He, *Huaiyu Dai*
*NC State University, †Lamar University

Paper: Use of Phishing Training to Improve Security Warning Complience: Evidence from a Field Experiment
**Weining Yang, *Aiping Xiong, †Jing Chen, *Robert W. Proctor, *Ninghui Li*
*Purdue University, †New Mexico State University

1200 - 1330 **Lunch *(on your own)***

1330 - 1500 **SESSION 7**

Paper: Global Variation in Attack Encounters and Hosting
**Ghita Mezzour, †Kathleen M. Carley, †L. Richard Carley*
*International University of Rabat, †Carnegie Mellon University

Paper: An Approach to Incorporating Uncertainty in Network Security Analysis
Hoang Hai Nguyen, Kartik Palani, David M. Nicol
University of Illinois at Urbana-Champaign

Paper: Surveying Security Practice Adherence in Software Development
Patrick Morrison, Laurie Williams
North Carolina State University

1500 - 1515 **Closing Remarks**

15:15 **SYMPOSIUM ADJOURNED**



ABSTRACTS

SECURITY OF CYBER-PHYSICAL SYSTEMS: CHALLENGES AND APPROACHES

Insup Lee

Cyber-physical systems (CPS) are used to monitor and control real-world systems, ranging from implantable life-critical medical devices to large safety-critical infrastructures. Securing CPS introduces additional challenges since the attack surface is significantly increased, compared to conventional systems. In addition to cyber intrusions that apply to all computer systems, attacks on CPS can be through interference to the physical environment of CPS. A number of such attacks have emerged recently. These incidents suggest that conventional cyber-only security approaches will not be effective in dealing with such attacks. This talk describes approaches to making CPS resilient to cyber-physical attacks by exploiting spatial and temporal redundancy as well as dynamics of the underlying physical system.

LEVERAGING UNIQUE CPS PROPERTIES TO DESIGN BETTER PRIVACY-ENHANCING ALGORITHMS

Jairo Giraldo, Alvaro A. Cárdenas, Murat Kantarcioglu

Cyber-Physical Systems (CPS) have unique properties that can be exploited to design new privacy-enhancing technologies that minimize the negative impact to the utility of CPS. In this paper we show two examples of these properties.

The first example looks at how differential privacy degrades CPS performance due to the large noise addition, and we then show how the inherent noise of CPS can be leveraged to reduce the additional noise added by differential privacy algorithms, and therefore, minimize the negative impact on the system utility and safety. In the second example we look at the ability to sample at sensor readings on demand, and how this flexibility can be used to design adaptive sensor sampling algorithms that hide sensitive information without the need to add noise.

CHARACTERIZING SCIENTIFIC REPORTING IN SECURITY LITERATURE: AN ANALYSIS OF ACM CCS AND IEEE S&P PAPERS

Morgan Burcham, Mahran Al-Zyoud, Jeffrey C. Carver, Mohammed Alsaleh, Hongying Du, Fida Gilani, Jun Jiang, Akond Rahman; Özgür Kafalı, Ehab Al-Shaer, Laurie Williams

Scientific advancement is fueled by solid fundamental research, followed by replication, meta-analysis, and theory building. To support such advancement, researchers and government agencies have been working towards a "science of security". As in other sciences, security science requires high-quality fundamental research addressing important problems and reporting approaches that capture the information necessary for replication, meta-analysis, and theory building. The goal of this paper is to aid security researchers in establishing a baseline of the state of scientific reporting in security through an analysis of indicators of scientific research as reported in top security conferences, specifically the 2015 ACM CCS and 2016 IEEE S&P proceedings. To conduct this analysis, we employed a series of rubrics to analyze the completeness of information reported in papers relative to the type of evaluation used (e.g. empirical study, proof, discussion). Our findings indicated some important information is often missing from papers, including explicit documentation of research objectives and the threats to validity. Our findings show a relatively small number of replications reported in the literature. We hope that this initial analysis will serve as a baseline against which we can measure the advancement of the science of security.

OPTIMAL SECURITY INVESTMENTS IN A PREVENTION AND DETECTION GAME

Carlos Barreto, Alvaro A. Cárdenas, Alain Bensoussan

Most security defenses can be breached by motivated adversaries, therefore in addition to attack-prevention technologies, firms investing in cyber-security for their information technology infrastructure need to consider attack-detection and restoration tools to detect intruders, and restore their system to a safe condition. Attackers face similar investment alternatives; they need to invest resources to finding vulnerabilities in a protected system, and once the protection has been broken, they need to invest in the infrastructure necessary to exploit these attacks and maintain stealthy persistence in the compromised infrastructure. We model this dual considerations as a dynamic programming problem between attackers and defenders and then study the Nash equilibrium of this game. Our goal is to find models and alternatives that can help us understand optimal security investments in prevention and detection against advanced rational adversaries.

KEYNOTE: NAVIGATING PRIVACY ISSUES IN A DATA DRIVEN WORLD

Jules Polonetsky

As 2017 begins, are we entering a time of crisis or a time of opportunity for individual privacy? Law enforcement demands for consumer data continue to grow and surveillance by intelligence agencies continues to drive civil liberties debates. Online tracking for analytics and advertising has been extended to mobile devices, to interactive television and to smart home devices. Social media sharing has achieved near ubiquity, with services integrating location, facial recognition, and live video sharing. With connected cars, our motor vehicles become data collectors and with drones our public spaces can be more easily monitored. Big data strains against fair information practices of consent, limited purpose and data minimization. Algorithmic decision making and machine learning wreak havoc with efforts to provide transparency. Artificial Intelligence may leave us unsure who will even be accountable for data driven determinations.

But yet, the very stress on privacy created by these advances could create opportunities for progress. Academic-industry cooperation can be a competitive advantage in working on issues that cross borders and span continents, even as new governmental roles are shaped and defined. Based on his experience working with more than 130 senior privacy officers at companies, as well as with leading academics and privacy advocates, Jules will discuss some expectations of privacy and data management in the coming months and years, including a review of current areas of privacy developments around a variety of industries and technologies. He will also examine opportunities to ensure data is available in a trusted manner for evidence based policymaking.

TUTORIAL: THE BUGS FRAMEWORK (BF) "HANDS-ON"

Irena Bojanova, Paul E. Black

Advancements of scientific foundation in cybersecurity rely on the availability of accurate, precise, and non-ambiguous definitions of software weaknesses (bugs) and descriptions of software vulnerabilities. The Bugs Framework (BF) organizes software weaknesses into distinct classes, such as buffer overflow (BOF), injection (INJ), faulty operation (FOP), and control of interaction frequency (CIF). Each BF class has an accurate and precise definition and comprises: * Attributes that identify the software fault; * Causes that bring about the fault; * Consequences the fault could lead to; * Sites in code where the fault might occur.

Through a "hands-on" approach the attendees will be able to analyze definitions and (static) attributes of bugs' classes, along with their related dynamic properties, such as proximate, secondary and tertiary causes, consequences and sites. The focus will be on at least three of the developed BF classes, as well as on examples of applying the BF taxonomy to describe vulnerabilities such as Heartbleed and Ghost. The audience will be involved in describing particular software vulnerabilities, and in discussions about the benefits of BF. The organizers are the BF Principal Investigators and are proposing this tutorial as a way to help researchers and practitioners more accurately and quickly diagnose, describe, and measure security vulnerabilities.

TUTORIAL: SYSTEM MONITORING FOR SECURITY

Xusheng Xiao

Intrusive multi-step attacks, such as Advanced Persistent Threat (APT) attacks, have plagued many well-protected businesses with significant financial losses. These advanced attacks are sophisticated and stealthy, and can remain undetected for years as individual attack steps may not be suspicious enough. To counter these advanced attacks, a recent trend is to leverage ubiquitous system monitoring for collecting the attack provenance for a long period of time and perform attack investigation for identifying risky system behaviors. In this tutorial, I first provide an introduction to popular techniques and tools of collecting system monitoring data, such as auditd, Sysinternals/ETW, and sysdig. I then describe how dependency analysis can be applied on the system monitoring data for attack investigation, and how to perform data reduction techniques to scale up the analysis to monitor more hosts. I will also present another technique that mines the patterns of the low-level system behaviors collected by system monitoring (e.g., a process reads a file) and uses these patterns to identify high-level software behaviors that security analysts are interested in (e.g., file compressions and ssh login). Finally, we discuss the current challenges of analyzing system monitoring data for security problems, concluding with future research directions.

KEYNOTE: DIFFERENTIAL PRIVACY AND DATA ANALYSIS

Aaron Roth

In this talk, we will give a friendly introduction to Differential Privacy, a rigorous methodology for analyzing data subject to provable privacy guarantees, that has recently been widely deployed in several settings. The talk will specifically focus on the relationship between differential privacy and machine learning, which is surprisingly rich. This includes both the ability to do machine learning subject to differential privacy, and tools arising from differential privacy that can be used to make learning more reliable and robust (even when privacy is not a concern).

LEARNING A PRIVACY INCIDENTS DATABASE

Pradeep K. Murukannaiah, Chinmaya Dabral, Karthik Sheshadri, Esha Sharma, Jessica Staddon

A repository of privacy incidents is essential for understanding the attributes of products and policies that lead to privacy incidents. We describe our vision for a novel privacy incidents database and our progress toward building a prototype. Key challenges in gathering such a database include bootstrapping and sustainability. We propose a semi-automated framework that can recognize privacy incidents and related information from various online sources such as news, blogs, and social media. The crux of our framework is an incident classifier that identifies whether a piece of text in natural language is related to a privacy incident or not. We curate a dataset consisting of 1324 news articles of which 543 articles are about one or more privacy incidents. We train the incident classifier on this dataset, considering a variety of feature engineering, feature selection, and classification techniques. We find that our incident classifier yields an F1 measure of 93.1%, which is about 12% higher than the keyword search-based baselines we adopt.

ON THE TRADE-OFF BETWEEN PRIVACY AND UTILITY IN COLLABORATIVE INTRUSION DETECTION SYSTEMS-A GAME THEORETICAL APPROACH

Richeng Jin, Xiaofan He, Huaiyu Dai

Intrusion Detection Systems (IDSs) are crucial security mechanisms widely deployed for critical network protection. However, conventional IDSs become incompetent due to the rapid growth in network size and the sophistication of large scale attacks. To mitigate this problem, Collaborative IDSs (CIDSs) have been proposed in literature. In CIDSs, a number of IDSs exchange their intrusion alerts and other relevant data so as to achieve better intrusion detection performance. Nevertheless, the required information exchange may result in privacy leakage, especially when these IDSs belong to different self-interested organizations. In order to obtain a quantitative understanding of the fundamental tradeoff between the intrusion detection accuracy and the organizations' privacy, a repeated two-layer single-leader multi-follower game is proposed in this work. Based on our game-theoretic analysis, we are able to derive the expected behaviors of both the attacker and the IDSs and obtain the utility-privacy tradeoff curve. In addition, the existence of Nash equilibrium (NE) is proved and an asynchronous dynamic update algorithm is proposed to compute the optimal collaboration strategies of IDSs. Finally, simulation results are shown to validate the analysis.

USE OF PHISHING TRAINING TO IMPROVE SECURITY WARNING COMPLIANCE: EVIDENCE FROM A FIELD EXPERIMENT

Weining Yang, Aiping Xiong, Jing Chen, Robert W. Proctor, Ninghui Li

The current approach to protect users from phishing attacks is to display a warning when the webpage is considered suspicious. We hypothesize that users are capable of making correct informed decisions when the warning also conveys the reasons why it is displayed. We chose to use traffic rankings of domains, which can be easily described to users, as a warning trigger and evaluated the effect of the phishing warning message and phishing training. The evaluation was conducted in a field experiment. We found that knowledge gained from the training enhances the effectiveness of phishing warnings, as the number of participants being phished was reduced. However, the knowledge by itself was not sufficient to provide phishing protection. We suggest that integrating training in the warning interface, involving traffic ranking in phishing detection, and explaining why warnings are generated will improve current phishing defense.

GLOBAL VARIATION IN ATTACK ENCOUNTERS AND HOSTING

Ghita Mezzour

Countries vary greatly in the extent to which their computers encounter and host attacks. Empirically identifying factors behind such variation can provide a sound basis for policies to reduce attacks worldwide. However, the main current approach to identify these factors consists of expert opinions with limited empirical validation. In this work, we empirically test hypotheses regarding social and technological factors behind such international variation. We use Symantec's Intrusion Prevention System (IPS) telemetry data collected from around 10 million Symantec customers worldwide.

We find that web attacks and fake applications are most prominent in Western Europe and North America. Our results indicate a relationship between countries' wealth and technological sophistication and attack exposure, indicating that attackers probably target developed countries to maximize their profits. Moreover, Eastern Europe hosts disproportionate quantities of attacks. Our statistical analysis reveals a relationship between attack hosting and the combined effect of widespread corruption and computing resources. Surprisingly, China is not among the top 10 attack hosting countries and Africa hosts the smallest quantities of attacks. Our work has important policy implications.

AN APPROACH TO INCORPORATING UNCERTAINTY IN NETWORK SECURITY ANALYSIS

Hoang Hai Nguyen, Kartik Palani, David M. Nicol

Attack graphs used in network security analysis are analyzed to determine sequences of exploits that lead to successful acquisition of privileges or data at critical assets. An attack graph edge corresponds to a vulnerability, tacitly assuming a connection exists and tacitly assuming the vulnerability is known to exist. In this paper we explore use of uncertain graphs to extend the paradigm to include lack of certainty in connection and/or existence of a vulnerability. We extend the standard notion of uncertain graph (where the existence of each edge is probabilistically independent) however, as significant correlations on edge existence probabilities exist in practice, owing to common underlying causes for dis-connectivity and/or presence of vulnerabilities. Our extension describes each edge probability as a Boolean expression of independent indicator random variables. This paper (i) shows that this formalism is maximally descriptive in the sense that it can describe any joint probability distribution function of edge existence, (ii) shows that when these Boolean expressions are monotone then we can easily perform uncertainty analysis of edge probabilities, and (iii) uses these results to model a partial attack graph of the Stuxnet worm and a small enterprise network and to answer important security-related questions in a probabilistic manner.

SURVEYING SECURITY PRACTICE ADHERENCE IN SOFTWARE DEVELOPMENT

Patrick Morrison, Laurie Williams

Software development teams are increasingly incorporating security practices in to their software development processes. However, little empirical evidence exists on the costs and benefits associated with the application of security practices. Balancing the trade off between the costs in time, effort, and complexity of applying security practices and the benefit of an appropriate level of security in delivered software requires measuring security practice benefits and costs. The goal of this research is to support researcher investigations of software development security practice adherence by building and validating a set of security practices and adherence measures through literature review and survey data analysis. We extracted 16 software development security practices from a review of the literature, and established a set of adherence measures based on technology acceptance theory. We built a survey around the 13 most common practices and our adherence measures. We surveyed 11 security-focused open source projects to collect empirical data as a test of our theorizing about practice adherence. In our collected survey data, each of the 13 security practices we identified was used daily by at least one survey participant. Tracking vulnerabilities and applying secure coding standards are the practices most often applied daily. In our data, Ease of use, Effectiveness, and Training, measured via Likert items, did not always show the expected theoretical relationship with practice use. In our data, Training is positively correlated with practice use, while Effectiveness and Ease of use vary in their correlations with practice use on a practice by practice basis.