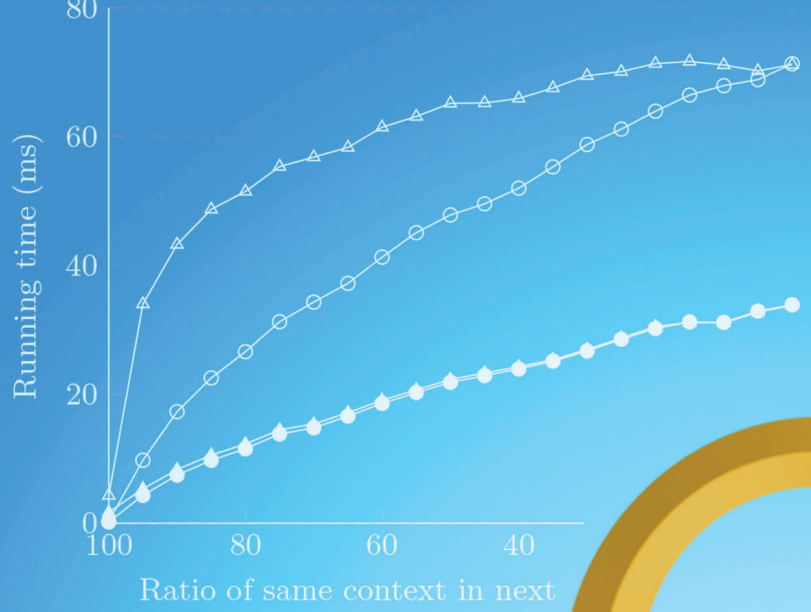

$$S_i = (\boldsymbol{loc} = s) \wedge (\boldsymbol{tab} = t) \wedge \bigwedge_{\mathcal{F} \in \mathbb{F}} (\boldsymbol{\mathcal{F}} = \mathbb{R}_i[\mathcal{F}]) \wedge \neg \boldsymbol{S}_{-i} \quad (2)$$

$$S'_{i,a} = \bigvee_{k \in \mathbb{O}} \begin{bmatrix} (\boldsymbol{loc'} = k_{tar}) \wedge (\boldsymbol{tab'} = 0) \\ \wedge \bigwedge_{\mathcal{F} \in \mathbb{F}} (\boldsymbol{\mathcal{F}'} = \mathbb{R}'_i[\mathcal{F}]) \end{bmatrix} \wedge \bigwedge_{\mathcal{A} \in \mathbb{A}} (\boldsymbol{\mathcal{A}'} = 0) \quad (3)$$

$$S'_{i,g} = (\boldsymbol{loc'} = loc) \wedge (\boldsymbol{tab'} = new\_table) \wedge \bigwedge_{\mathcal{F} \in \mathbb{F}} (\boldsymbol{\mathcal{F}'} = \mathbb{R}'_i[\mathcal{F}]) \wedge \bigwedge_{\mathcal{A} \in \mathbb{A}} (\boldsymbol{\mathcal{A}'} = \mathbb{R}'_i[\mathcal{A}]) \quad (4)$$

$$S'_{i,as} = \mathbb{R}'_i[AS\_OUT] \rightarrow [(\boldsymbol{loc'} = \mathbb{R}'_i[AS\_OUT]) \wedge (\boldsymbol{tab'} = 0) \wedge \bigwedge_{\mathcal{F} \in \mathbb{F}} Exp[\mathcal{F}] \wedge \bigwedge_{\mathcal{A} \in \mathbb{A}} (\boldsymbol{\mathcal{A}'} = 0)] \quad (5)$$


```
send(sock, msg) {
  if(!isConfidential(msg)) {
      dispatch(m
  } else
      throw Ex
}

secureSend(sock,
  assert(sock instance
  dispatch(ssl
}

connect(addr) {
  /* create an HttpUrlConnection
}

secureConnect(ad
  /* create Ht
  /* Perform re
  /* Perform app
  /* return sslSoc
}
```

SCIENCE OF SECURITY AND PRIVACY

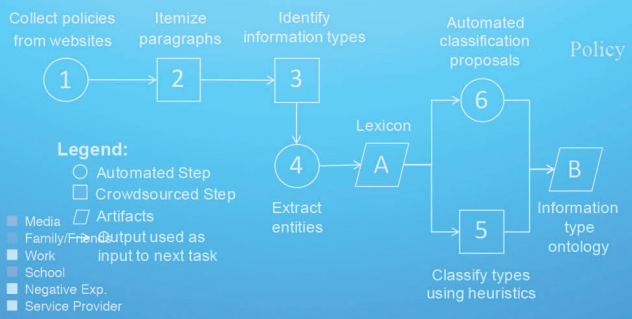

Categories

- Physical Host
- Hardware Component
- Supervisor
- Intra-application
- Guest OS
- Application Group
- Application
- Sub-application
- Automatic
- Manual
- Reconfigurable
- Non-reconfigurable
- Always-on
- On-demand

- Policy Generation
- Policy Configurability
- Policy Lifetime

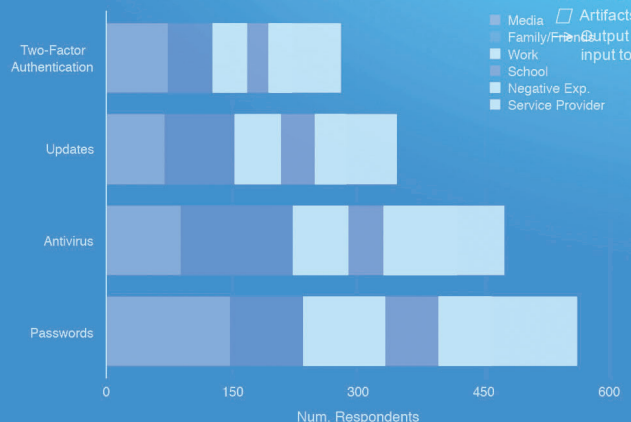

Figure 4: Advice source prevalence by behavior.

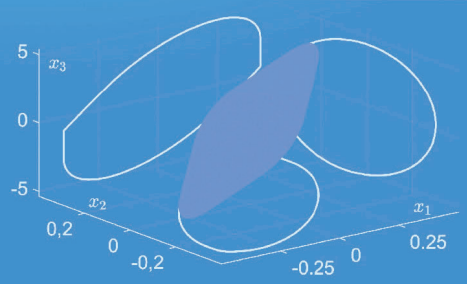

Annual Report

2016

VISION

The National Security Agency Research Directorate sponsors the Science of Security Initiative for the promotion of a foundational cybersecurity science that is needed to mature the cybersecurity discipline and to underpin advances in cyberdefense.

Science of Security (SoS) Initiative

Annual Report

SCIENCE OF SECURITY
AND PRIVACY

2016

Table of Contents

EXECUTIVE SUMMARY

In 2016, the Science of Security (SoS) initiative increased in size and scope as members of academia, industry, and government continued to research and collaborate under the sponsorship of the National Security Agency (NSA) Research Directorate. The SoS initiative focused on producing scientifically supported cybersecurity advancement in the establishment of cybersecurity as a science. By replacing ad hoc and common practice approaches to security with scientifically supported best practice methods established through rigorous research, SoS is developing strategic rather than tactical methods of approaching cybersecurity. These strategic results are needed to transform cybersecurity from a cost-disadvantaged, reactionary field to one that is efficient and proactive in both attack and defense. Established in 2012, the Science of Security fosters the establishment of security science through the pursuit of its three stated goals: engage the academic community for foundational research, promote rigorous scientific principles, and grow the SoS community. Over the past year, the Science of Security initiative has seen continued success in its pursuit of these goals. Success is measured by the progress made against each of the five Hard Problems, the major focus areas identified by NSA and the original three Lablets funded by the NSA Research Directorate. The five Hard Problems are:

- Scalability and Composability
- Policy-Governed Secure Collaboration
- Security-Metrics-Driven Evaluation, Design, Development, and Deployment
- Resilient Architectures
- Understanding and Accounting for Human Behavior

In 2016, the four Science of Security Lablets sited at Carnegie Mellon University, North Carolina State University, the University of Illinois at Urbana-Champaign, and the University of Maryland, made significant progress against the five Hard Problems. In this work, they collaborated regularly with 25 smaller subLablets throughout the country. Progress against the Hard Problems in 2016 was documented in 63 publications authored by researchers from the various Lablets. Some of the major advancements included:

- A more thorough understanding of which security warnings and behaviors users comprehend and heed, as well as how users apply security warnings and best practice security measures to their online behavior. This research allows developers to create more effective warnings for the average user.
- Increased effectiveness in problem and attack detection through the development of redundancy-based anomaly detectors and diversity-based adjudication. Additionally, problem and attack detection ability was increased through the development of automated resilience analysis and the development of two resilience metrics that were integrated to determine optimal isolation-diversity.
- Stronger Software Defined Networks (SDNs) through network optimization applications, newly developed tools for studying SDNs, studying correct virtualization behavior, and more efficient estimates of link failure impact on SDN connectivity.
- Improved quality of vulnerability prediction and identification by utilizing attack behavior simulations.
- Closer alignment of privacy policy analysis and program analysis achieved by detecting when data is being repurposed in software, resulting in new tools to warn developers about potential privacy violations.
- The development of UberSpark, a tool that applies a previously identified logic for compositionally guaranteeing safety when executing unknown code provided by an adversary in real software systems.

The SURE, SecUrity and REsilience for cyber-physical systems, program is another ongoing project sponsored by the SoS initiative with the goal of advancing foundational research. Specifically, the SURE program develops foundations and tools for designing, building, and assuring cyber-physical systems that can maintain essential system properties

in the presence of adversaries. SURE is led by Vanderbilt University and engages researchers at the Massachusetts Institute of Technology; the University of California, Berkeley; and the University of Hawaii.

In addition to engaging the academic community for foundational research, the SoS initiative has promoted rigorous research methods in academia by sponsoring academic competitions. One of the highlights of the year for SoS was the 4th Annual Best Scientific Cybersecurity Paper Competition, in which three outstanding scientific papers were recognized for their achievements in cybersecurity research. Additionally, SoS sponsored an award at the 2016 Intel International Science and Engineering Fair and was able to recognize the scientific achievements of three high school students in the field of cybersecurity.

In an effort to grow the Science of Security community, the SoS initiative continued to employ several highly effective tools. The Annual Hot Topics in the Science of Security (HoTSoS) Symposium and Bootcamp provided a forum for researchers from many different disciplines to discuss scientific approaches to cybersecurity. With over 1100 members, including many of the more than 10,000 published authors featured in the SoS-VO newsletters, the SoS Virtual Organization continues to grow and remains a centralized, online location for researchers and all interested parties to engage in discussion and access the most current research in cybersecurity.

In 2017, the SoS community will continue to grow, supported by efforts from within the NSA Research Directorate as well as industry and academia world-wide. Goals for fiscal year 2017 include: developing the foundational cybersecurity breakthroughs to underpin cybersecurity; increasing membership and engagement on the SoS-VO website; increasing collaboration among academia, industry, and government researchers; and raising awareness of the SoS community and the necessity of foundational research in the field of cybersecurity.

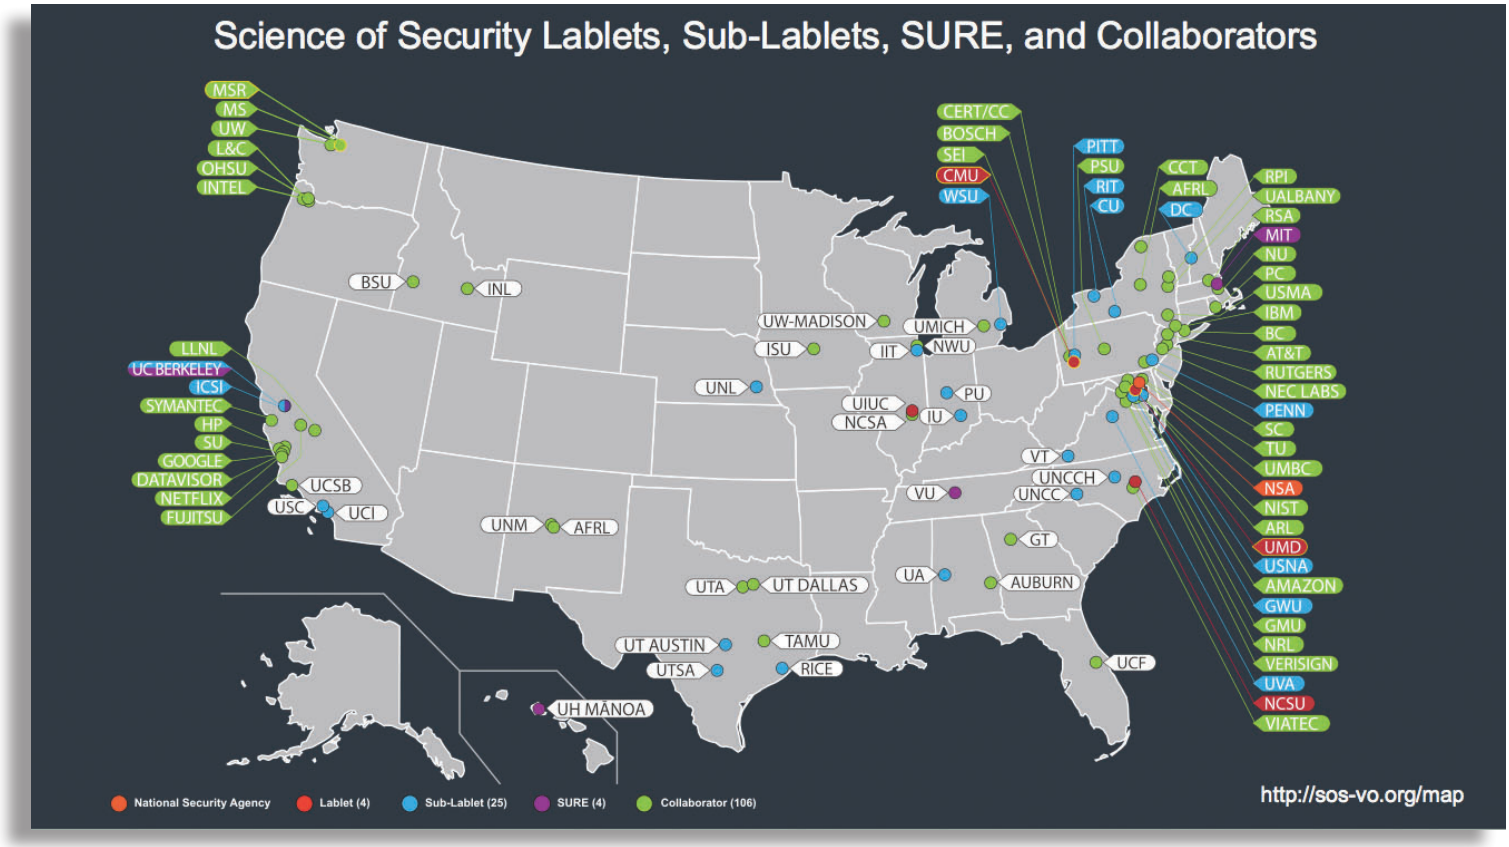

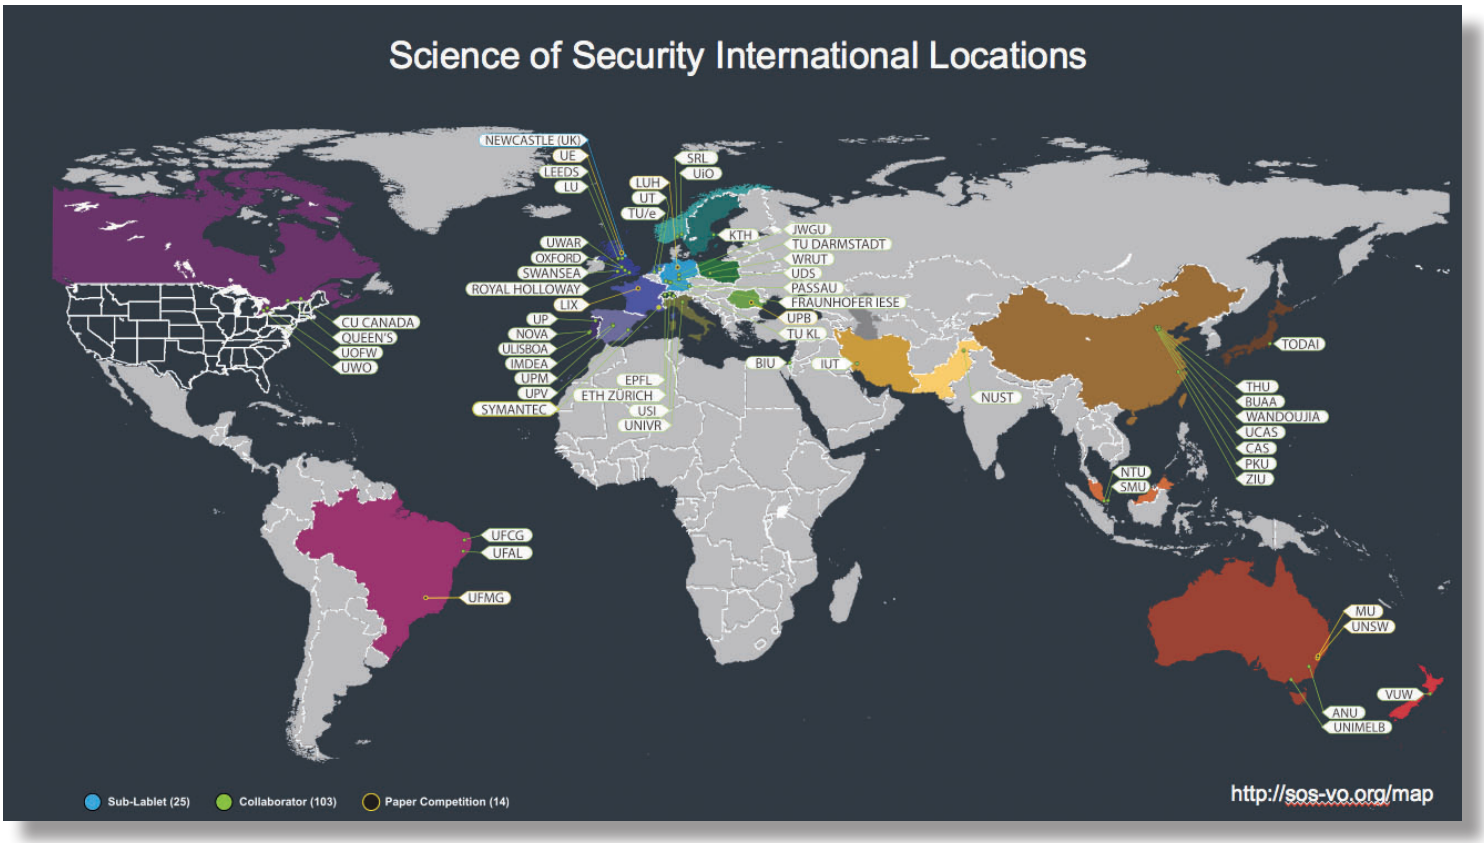

For more information about the Science of Security program browse the SoS web site at http://sos-vo.org Or scan the QR code

Science of Security

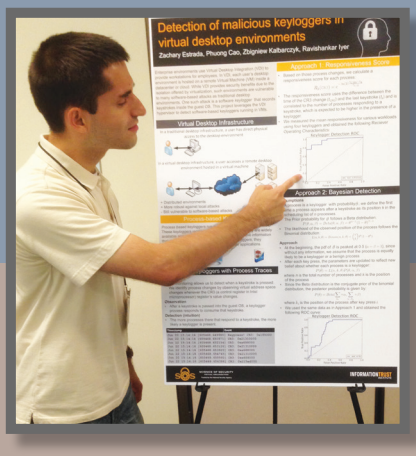

Section 1

Engaging the Academic Community for Foundational Research

The Science of Security (SoS) program focuses its academic engagement with the four Science of Security Lablets established in 2012 to further the Science of Security principles, enhance the scientific rigor of research into cybersecurity, and grow the Science of Security community. The four Lablets engage with 25 additional institutions (subLablets) and another 106 collaborators. The Principal Investigators (PIs) of the Science of Security Lablets, along with the NSA Research organization, developed five Hard Problems as a means to establish the beginnings of a common language and to gauge progress in foundational SoS research. The papers published over the past year provide tangible evidence of the impact that the Lablets' research has had on improving the Science of Security in the five Hard Problem focus areas.

The Lablets organize their research into projects, all of which relate to one or more of the five Hard Problems. Over the past year, there were 40 lablet research projects that produced 119 published papers. The Lablets also engaged regularly in community outreach, participated in international conferences and workshops, and integrated Science of Security principles into their curricula. In addition to providing quarterly and annual reports on their activities, the four Lablets, along with some of the subLablets, collaborators, and NSA researchers, met quarterly to present updates on their research projects and exchanged information about issues related to Science of Security. The Science of Security program also sponsored the Science of SecUrity and REsilience for Cyber-Physical Systems (SURE) project that involved another four universities. The status of Hard Problems, lablet activity, lablet quarterly meetings, and the SURE project are detailed in this section.

Science of Security Lablets Progress on Hard Problems 2016

The Principal Investigators (PIs) of the Science of Security Lablets, in collaboration with National Security Agency (NSA) Research, developed the Hard Problems as a measure to establish the beginnings of a common language and assess progress. These problems were selected for their level of technical challenge, their potential operational significance, and the likelihood that these problems would benefit from emphasis on scientific research methods and improved measurement capabilities. The five problems are: (1) Scalability and Composability; (2) Policy-Governed Secure Collaboration; (3) Security-Metrics-Driven Evaluation, Design, Development, and Deployment; (4) Resilient Architectures; and (5) Understanding and Accounting for Human Behavior.

The Hard Problem of Scalability and Composability deals with the development and analysis of large-scale secure systems and the study of how to improve system security through security improvement of the components. Policy-Governed Secure Collaboration aims to develop the science underlying methods to express and enforce normative requirements and policies for handling data with differing usage needs and among users in different authority domains. The hard problem of Security Metrics and Evaluation, Design, Development and Deployment addresses the measurement of properties relevant to cybersecurity and quantifying the degree to which a system satisfies those properties. Research into Resilient Architectures includes the ability of the system to statically withstand attack, the ability of a system to continue to deliver essential services in the midst of an attack, and how quickly a system can be restored to full functionality following an attack. Understanding and Accounting for Human Behavior addresses how to handle the unpredictability of human actors in cybersecurity.

The five Hard Problems are not intended to cover all cybersecurity research challenges, but rather five specific areas that need scientific progress. Fundamental research undertaken by the Lablets is tied to at least one hard problem. The following updates provide tangible evidence of the impact that the Lablets' research has had on improving the Science of Security in these five focus areas.

Scalability and Composability

This hard problem involves exploring techniques for constructing and analyzing large-scale secure systems ("scalability"), in particular focusing on the study of modular approaches that allow for proving security of the system as a whole by proving security of various components of that system ("composability"). Composable analysis approaches generally help to achieve scalability, because different parts of a system can be analyzed independently, and the analysis results from each of the parts can be efficiently combined to understand properties of the whole system. Lablet research has produced fundamental advances in the ability to perform composable security analyses:

- A number of researchers have proposed using immutability to help software developers better protect sensitive data, both from malicious software components and from unsafe access by other threads. While prior systems for enforcing immutability provided strong technical guarantees, we had no empirical data about how well they met developer needs. Lablet research in the past year used expert interviews and lab studies to better understand the needs of developers, and found that existing tools were too complex to be usable and did not match developer needs well. These results motivated current research on a new composable type system for immutability that combines a firm mathematical foundation with a design that avoids excessive complexity while still providing key features required by users, including transitive immutability semantics.

    - Michael Coblenz, Joshua Sunshine, Jonathan Aldrich, Brad Myers, Sam Weber, and Forrest Shull, "Exploring Language Support for Immutability," in *Proceedings of the International Conference on Software Engineering (ICSE)*, Austin, TX, May 14-22, 2016, pp. 736-747.

- Researchers developed a novel hybrid analysis that combines static program analysis and dynamic kernel event reporting tools. This analysis provides the observability and controllability needed to reproduce races that are historically difficult to locate and effectively reproduce at production scales. Results indicate that a tool based on this analysis can effectively reproduce races, allowing developers to correct process-level race conditions in modern software systems. The empirical study contributed to a deeper understanding of scalability by providing the foundational research needed to develop effective software analysis tools.

    - Supat Rattanasuksun, Tingting Yu, Witawas Srisa-an, and Gregg Rothermel, "RRF: A Race Reproduction Framework for Use in Debugging Process-Level Races," in *Proceedings of the International Symposium on Software Reliability Engineering (ISSRE)*, Ottawa, Canada, 2016, pp. 162-172.

- Sandboxes are a critical technology for imposing a security policy on a component within a larger system, but sometimes fail to properly enforce the desired policy because the sandbox is misused. Prior to last year, however, there were no scientific studies evaluating the extent and nature of the problem in the setting of the Java sandbox, and, as a consequence, we also lacked the necessary information to effectively mitigate the problem. Lablet researchers carried out an empirical study, finding that developers often misuse the Java sandbox, and identifying key differences between benign uses of the sandbox and exploits. We showed that these differences can be leveraged through dynamic analysis to stop a representative proportion of modern Java 7 exploits without breaking backwards compatibility with benign applications.

    - Zack Coker, Michael Maass, Tianyuan Ding, Claire Le Goues, and Joshua Sunshine, "Evaluating the Flexibility of the Java Sandbox," *Annu. Computer Security Applications Conference (ACSAC)*, Los Angeles, December 7, 2015, pp. 1-10.

- Lablet researchers developed a system called UberSpark that enforces verifiably secure-object abstractions in systems software. The system, written in C99 and Assembly, leverages abstractions to compositionally verify security properties of extensible commodity systems

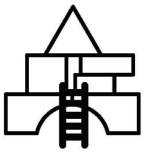

software (BIOS, OS, and hypervisor) written in low level-level languages. UberSpark enforces abstractions found in high-level languages using a combination of hardware mechanisms and lightweight static analysis. Some of these abstractions include interfaces, function-call semantics for implementations of interfaces, access control on interfaces, and forms of concurrency. UberSpark applies a previously identified logic for compositionally guaranteeing safety when executing unknown code provide by an adversary in real software systems.

  - Amit Vasudevan, Sagar Chaki, Petros Maniatis, Limin Jia, and Anupam Datta, "uberSpark: Enforcing Secure Object Abstractions for Automated Compositional Security Analysis of a Hypervisor, in *Proceedings of the USENIX Security Symposium*, Austin, TX, 2016, pp. 87-104.

- Accurately detecting security threats is challenging, because the evidence that attackers leave is often at a low level of abstraction where patterns of behavior are hard to detect. Lablet researchers developed an approach to raise the level of abstraction of threat detection by mapping sequences of observed activities to paths on the graph of the underlying software architectural model. Based on this abstraction, they applied unsupervised learning algorithms to detect insider threats and other covert attacks.

  - Eric Yuan and Sam Malek, "Mining Software Component Interactions to Detect Security Threats at the Architectural Level" in *Proceedings of the 13th Working IEEE/ IFIP Conference on Software Architecture (WICSA)*, Venice, Italy, 2016, pp. 211-220.

- Users increasingly rely on the functionality of mobile platforms such as Android, but preventing security violations on these platforms has been difficult. Although Android's permission system can be used to control access to resources, enforcing permissions is insufficient to provide desired security properties because permissions may be mismanaged, either intentionally or unintentionally–for example, malicious applications can collude to combine their permissions or trick vulnerable applications to perform sensitive actions on their behalf. Lablet research developed the COVERT system, which applies compositional static analysis of Android applications, along with model checking, to determine whether a collection of applications preserves safety properties when run together.

  - Eric Yuan and Sam Malek, "Mining Software Component Interactions to Detect Security Threats at the Architectural Level, in *Proceedings of the 13th Working IEEE/ IFIP Conference on Software Architecture* (WICSA), Venice, Italy, 2016, pp. 211-220.

- Some challenges of composability arise in the new approaches for policy-governed secure collaboration. In particular, developing systems of multiple autonomous agents (representing different stakeholders) calls for a form of composition wherein we specify the social protocol that characterizes the interactions or rules of encounter between the stakeholders. We proposed an approach called Interaction-Oriented Software Engineering (IOSE) that describes a protocol via normative relationships that capture who is accountable to whom for what, independent of implementation. Software agents representing the various parties are constructed to participate in that protocol in a manner that respects the internal preferences of the respective parties. Doing so partitions the system design problem into an accountability (protocol-level) component and a compliance (agent-level) component.

  - Amit K. Chopra and Munindar P. Singh, "From Social Machines to Social Protocols: Software Engineering Foundations for Sociotechnical Systems," in *Proceedings of the 25th International Wrld Wide Web Conference*, Montreal, Canada, April 2016, pp. 903-914.

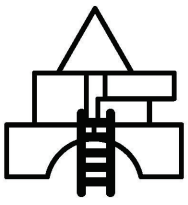

Policy-Governed Secure Collaboration

The hard problem of policy-governed secure collaboration is about developing the science underlying methods for expressing and enforcing normative requirements and policies for handling data with differing usage needs and among users in different authority domains. Over the course of the last year, we have deepened the scientific foundations that help us address key limitations of the state of the art. In particular, we previously formulated the notion of a sociotechnical system that captures its social and technical elements and the interplay between them in a formal, computational manner. Our research addresses challenges in bringing this view of sociotechnical systems for cybersecurity to fruition; understanding and reducing the complexity of policies; producing a repository for privacy incidents to promote the scientific study of privacy; and computing trust relationships.

- One of our advances was in the formal study of social architectures requisite for the adoption and enforcement of normative requirements and the creation of flexible trust relationships. We have begun to create models of social architectures in qualitative terms as a basis for further empirical validation with users. We have developed a simulation approach that accommodates both social and technical concerns in a setting where participants interact freely depending upon the environment, that is, the interacting parties are not predetermined. This approach considers integrity violations (which technical mechanisms can avoid) and conflicts (which social mechanisms in the form of norms can help reduce).

    - Medhi Mashayekhi, Hongying Du, George F. List, and Munindar P. Singh, "Silk: A Simulation Study of Regulating Open Normative Multiagent Systems," in *Proceedings of the 25th International Joint Conference on Artificial Intelligence* (IJCAI), New York, NY, July 2016, pp. 373-379.

- Previous policy approaches handled challenges such as authentication and authorization of users for performing data operations based on attribute or role-based credentials. In essence, these approaches focus on the technical architecture by providing access control on assets. However, these approaches do not adequately and explicitly characterize the correctness requirements for secure collaboration and their impact on security. That is, they do not answer the question of whether access control is appropriate given the social architecture, especially in situations that call for deviations from the stated policies. We have developed three approaches that tackle different aspects ignored in previous research. The first, Coco, provides a language and formal semantics to capture some subtleties of secure collaboration requirements, including priorities, and deals with handling conflicts that arise at runtime. The second, Custard, shows how to store normative relationships that arise at runtime in information stores and how to query those stores in high-level terms. Doing so helps both participants and other monitors to determine who is accountable for what and who has complied with or violated those accountabilities. The third approach, Revani, specifies and verifies sociotechnical systems interrelating the social tier (norms) with the technical tier (access control). It proposes design rules by which we can iteratively specify a sociotechnical system that meets stated stakeholder requirements, e.g., by relaxing or tightening the applicable norms.

    - Nirav Ajmeri, Jiaming Jiang, Rada Chirkova, Jon Doyle, and Munindar P. Singh, Coco: Runtime Reasoning about Conflicting Commitments, in *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, New York, NY, July 2016, pp. 17-23.

    - Amit K. Chopra and Munindar P. Singh, "Custard: Computing Norm States over Information Stores," in *Proceedings of the 15th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*, Singapore, May 2016, pp. 1096-1105.

    - Özgür Kafali, Nirav Ajmeri, and Munindar P. Singh, "Revani: Revision and Verification of Normative Specifications

for Privacy," *IEEE Intelligent Systems*, September 2016, vol. 31, no. 5, pp. 8-15.

- Firewall policies are notorious for configuration errors that cause security vulnerabilities. Yet, there is inadequate understanding of the complexity in human terms of firewall policies. We have been developing such understanding and are now relating it to computational models. Specifically, we introduced an approach, ModFP, to automatically modularize firewall policies, using policies that we examined to uncover subtle errors. This approach can help raise the modularity and abstraction of firewall policies and potentially facilitate authoring such policies in a variety of policy languages.

    - Haining Chen, Omar Chowdhury, Ninghui Li, Warut Khern-am-nuai, Suresh Chari, Ian Molloy, and Youngja Park, "Tri-Modularization of Firewall Policies," in *Proceedings of the 21st ACM Symposium on Access Control Models and Technologies (SACMAT)*, Shanghai, China, June 2016, pp. 37-48.

- Attempts to develop a solid understanding of privacy are stymied by the lack of empirical data on privacy incidents. In a new project launched in January 2016, we initiated and began populating a privacy incidents database to provide a foundation for empirically grounding scientific research on privacy. We conducted a human-subjects study to validate the assumption that the incidents recorded in our database are understood as privacy incidents by non-specialists, such as members of the public, whose privacy we seek to enhance.

    - Pradeep K. Murukannaiah, Jessica Staddon, Heather R. Lipford, and Bart P. Knijnenburg, "PrIncipedia: A Privacy Incidents Encyclopedia," *Privacy Law Scholars Conference* Working Paper, Washington, DC, June 2016, pp. 1-21.

- Prior to our work, trust-inference mechanisms were limited in the type of evidence they could incorporate. We propose using a semiring framework as a way to combine evidence for inferring trust, where trust relationship is modeled as a 2-D vector containing both trust and degree of certainty. The trust propagation and aggregation rules, as the building blocks of our trust inference scheme, are based upon the properties of trust relationships. In our approach, both trust and distrust (i.e., positive and negative trust) are considered, and conflict resolution is supported. We evaluate the proposed approach on real-world datasets, and show that our trust inference framework has high accuracy, and is capable of handling trust relationship in large networks.

    - Peixin Gao, Hui Miao, John S. Baras, and Jennifer Golbeck, "STAR: Semiring Trust Inference for Trust-Aware Social Recommenders," in *Proceedings of the 10th ACM Conference on Recommender Systems (RecSys)*, Boston, MA, 2016, pp. 301-308.

- Before last year, privacy risk was promoted as a driving factor for restricting information sharing, but there were no empirical methods for measuring this risk. Today, we have designed new tools to measure perceived privacy risk. Using these tools, we have found that privacy risk is not a multiplicative product of likelihood and impact, and that, when information is shared for national-security purposes, individuals are more concerned about sharing information about their identity than they are about sharing information regarding their activities.

    - Jaspreet Bhatia, Travis D. Breaux, Liora Friedberg, Hanan Hibshi, Daniel Smullen, "Privacy Risk in Cybersecurity Data Sharing,"in *Proceedings of the 3rd ACM International Workshop on Information Sharing and Collaborative Security (WISCS)*, Vienna, Austria, 2016, pp. 57-64.

- Before this research, we did not have a systematic means of identifying the implied security requirements related to specified functional requirements. We have developed a process that takes as input existing natural language requirements artifacts and produces a set of candidate security requirements for the system as output using empirically-developed security requirements patterns.

    - Maria Riaz, Jason King, Fabio Massacci, Marcelo Jenkins, Christian Quesada- López, and Laurie A. Williams, "Identifying the Implied: Findings from Three Differentiated Replications on the

Use of Security Requirements Templates," *Empirical Software Engineering*, to appear 2016.

- Maria Riaz, Jonathan Stallings, Munindar P. Singh, John Slankas, and Laurie A. Williams, "Discovering Goals for Security Requirements Engineering - DIGS Framework and Evaluation," *Int. Symposium on Empirical Software Engineering and Measurement (ESEM)*, 2016, Ciudad Real, Spain, Article 35.

- Past lablet research discovered how to express confidential data flow policies to avoid repurposing, which occurs when information is used for a new purpose beyond its original authorization. This year, lablet researchers extended that work to align the analysis of privacy policy with program analysis by detecting when data is being repurposed in software. The results include new tools to warn developers about potential privacy violations.

  - Travis D. Breaux, Daniel Smullen, and Hanan Hibshi, "Detecting Repurposing and Over-collection in Multi-Party Privacy Requirements Specifications," in *Proceedings of the IEEE 23rd International Requirements Engineering Conference*, Ottawa, Canada, Sep. 2015, pp. 166-175.

  - Rocky Slavin, Xiaoyin Wang, Mitra Bokaei Hosseini, William Hester, Ram Krishnan, Jaspreet Bhatia, Travis D. Breaux, and Jianwei Niu, "Toward a Framework for Detecting Privacy Policy Violation in Android Application Code,"in *Proceedings of the ACM/IEEE 38th International Conference on Software Engineering*, Austin, Texas, 2016, pp. 25-36.

  - B. Kwon, J. Mondal, J. Jang, L. Bilge, and T. Dumitras, "The Dropper Effect: Insights into Malware Distribution with Downloader Graph Analytics," in *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security*, Denver, CO, 2015, pp. 1118-1129.

- Before our work, it was widely believed that when a browser authenticates a connection to a particular web domain, the browser is actually connected to a server operated by that domain. We have demonstrated empirically, with a large-scale measurement study of the Web's PKI, that many websites–and a majority of the most popular ones– share their private keys with third-party hosting providers like Content Distribution Networks (CDNs). Our results reveal surprising new truths about trust relationships in the PKI, notably that a small handful of companies hold a disproportionate number of popular websites' private keys.

  - Frank Cangialosi, Taejoong Chung, Dave Choffnes, Dave Levin, Bruce M. Maggs, Alan Mislove, and Christo Wilson, "Measurement and Analysis of Private Key Sharing in the HTTPS Ecosystem," in *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security (CCS)*, Vienna, Austria, 2016, pp. 628-640.

- Before our work, studies of the Web's HTTPS-certificate ecosystem focused solely on certificates that could be validated by a standard web browser. Our work has shown that such valid certificates constitute only 12% of the HTTPS certificates advertised over the past three years, and we have demonstrated that the invalid majority can lend considerable insight into the inner workings of home-network security. We found that most invalid certificates originate from end-user devices, and we demonstrated that it is possible to use unique features of these certificates to uniquely track over 6.7M devices as they move throughout the IP address space.

  - Taejoong Chung, Yabing Liu, Dave Choffnes, Dave Levin, Bruce MacDowell Maggs, Alan Mislove, and Christo Wilson, "Measuring and Applying Invalid SSL Certificates: The Silent Majority," in Proceedings of the 2016 ACM on Internet Measurement Conference (IMC '16), Santa Monica, CA, 2016, pp. 527-541.

- Prior to our work, there was little understanding of optimal configurations for Intrusion-Detection Systems (IDSs) in the presence of unknown system dynamics and intruder-type uncertainty. We have developed a general, incomplete-information, stochastic-game framework that can

be applied to this problem, leading to higher-fidelity detection of intruders and more-effective IDS configuration.

  - Xiaofan He, Huaiyu Dai, Peng Ning, and Rudra Dutta, "Dynamic IDS Configuration in the Presence of Intruder Type Uncertainty,"in *Proceedings of the IEEE Global Conference on Communications (GLOBECOM)*, 2015.

- Before last year, there was no design rationale for computational resource-sharing in IDS networks, and little quantitative study on the benefits of IDS collaboration. We have developed such a rationale using a two-layer, game-theoretic approach, and derived conditions under which there is a guaranteed performance improvement over autonomous IDS systems.

  - Richeng Jin, Xiaofan He, and Huaiyu Dai, "Collaborative IDS Configuration: A Two-layer Game-Theoretical Approach," in *Proceedings of the IEEE Global Conference on Communications (GLOBECOM)*, Washington, DC, 2016.

- Prior to last year, we were aware of the extent of research into Intrusion Detection Systems (IDS) but lacked a systematic understanding of it. What we understand today is that empirical evaluation of IDS research is highly inconsistent. Systematization of intrusion detection knowledge, then, will have to first identify a set of standard evaluation metrics that every empirical study ought to follow, so that future systematic literature reviews may make fair comparisons.

  - Nuthan, Munaiah, Andrew Meneely, Ryan Wilson, and Benjamin Short, "Are Intrusion Detection Studies Evaluated Consistently? A Systematic Literature Review," 2016. Accessed from **http://scholarworks.rit.edu/article/1810**

Security Metrics and Models

This hard problem involves techniques for measuring security-relevant properties and quantifying the extent to which a given system satisfies a given set of security properties. Lablet research is exploring several aspects of this problem, including the analysis of real-world datasets to understand and quantify factors leading to vulnerabilities and exploits; developing metrics for predicting vulnerabilities and measuring effectiveness of countermeasures in software; and measuring stakeholder perception of various security measures. Highlights of lablet work in this area include the following:

- Prior to last year, we did not know if vulnerabilities could be predicted more effectively by incorporating attack-surface data. Today, we know that we can improve the prediction quality of models when attacker behavior is simulated via random walks on the call graph starting from the attack surface. Additionally, we have shown that the attack surface can be approximated by recording all the files that appear on stack traces from crash dumps. With this new metric in place, we are one step closer to being able to identify vulnerabilities in code as the code is written (instead of relying on patching vulnerabilities after the fact). Metrics can be used to prioritize vulnerability detection efforts to the code on the prioritized attack surface.

  - Christopher Theisen, "Risk-Based Attack Surface Approximation," ACM Student Research Competition Grand Finals (Third Place Worldwide winner), 2016.

  - Christopher Theisen, "Reusing Stack Traces: Automated Attack Surface Approximation," in *Proceedings of the 38th International Conference on Software Engineering Companion*, Doctoral Symposium, Austin, TX, 2016, pp. 859-862.

  - Christopher Theisen, "Automated Attack Surface Approximation,"in

Proceedings of the 10th Joint Meeting on Foundations of Software Engineering (ESEC/FSE 2015), Student Research Competition (First Place winner), Bergamo, Italy, pp. 1063-1065.

- Nuthan Munaiah and Andrew Meneely, "Beyond the Attack Surface: Assessing Security Risk with Random Walks on Call Graphs," in *Proceedings of the 2016 ACM Workshop on Software PROtection (SPRO '16)*, Vienna, Austria, pp. 3-14.

- Before our work, it was known that only a fraction of the vulnerabilities from the CVE database are attacked by real-world exploits. Research conducted this year helps to explain why some vulnerabilities were exploited and others were not. We designed and implemented a system for forecasting which vulnerabilities will be exploited in the wild using information mined from Twitter. Our work shows that taking into account both the vulnerability characteristics (e.g., whether the vulnerability enables remote code execution) as well as features extracted from social media (e.g., dissemination of information about offensive techniques) can be used for better-quality predictions. Our system can be used to prioritize responses to vulnerability disclosures, or to model risk for cyber-insurance applications.

    - Carl Sabottke, Octavian Suciu, and Tudor Dumitras, "Vulnerability Disclosure in the Age of Social Media: Exploiting Twitter for Predicting Real-World Exploits,"in *Proceedings of the 24th USENIX Conference on Security Symposium (Sec '15)*, Washington, DC, 2015, pp. 1041-1056.

- Today few malware families have the ability to propagate autonomously; instead, they rely on malware-delivery networks that allow malware to infect millions of hosts worldwide. Before our work, these malware-delivery ecosystems had been studied from the network side, but their client-side infrastructure was unexplored. Our work shows that a downloader-graph abstraction, which captures the client-side activity of malware-delivery networks, allows us to effectively distinguish between malicious and benign software downloads.

Resilient Architectures

The hard problem of resilience, as we emphasize, includes (1) the ability of the system to statically withstand attack, e.g., through diversity in implementation, also known as robustness; (2) the ability of a system to continue to deliver essential services (albeit potentially at a diminished level) in the midst of an attack; and (3) how quickly a system can be restored to full functionality following an attack. As applied to system architectures, the problem involves ability to monitor a system to support these objectives, methodologies for detecting the need for and reacting to system changes needed to support resiliency, understanding of how key applications may be expressed to support their resiliency, and methodologies for assessing the resiliency of the system and the workloads it supports.

- Before our work, surprisingly few systems used redundancy-based, run-time problem detection. We have used redundancy-based anomaly detectors to recognize some high-risk and difficult-to-detect attacks on web servers, a likely management interface for many Internet of Things (IoT) stand-alone elements. Additionally, we have developed diversity-based adjudication, which relies on the knowable behavior of a healthy system, to detect even zero-day attacks.

    - Roopak Venkatakrishnan and Mladen A. Vouk, "Using Redundancy to Detect Security Anomalies: Towards IoT Security Attack Detectors," ACM Ubiquity 2016, January, Article 1.

- Prior to last year, there were no techniques for automated real-time analysis of so-called "workload resilience," which speaks to the ability of a computational workload to continue in the face of system disruption, while still meeting performance and policy objectives. We have provided new means

of performing that analysis automatically, at significantly lessened computational cost.

  - John C. Mace, Charles Morisset, and Aad van Moorsel, "Impact of Policy Design on Workflow Resiliency Computation Time," in *Proceedings of the 12th International Conference on Quantitative Evaluation of Systems (QEST)*, Madrid, Spain, September 1-3, 2015, pp. 244-259.

- In previous work, any analysis of variable workload resiliency focused on the mean of the stochastic resiliency measure (i.e., probability of workload completing successfully in the presence of disruption). However, better risk analysis of resilience should consider the variance of these measures. For the first time, we've introduced the idea of considering variance, and identified some resiliency workload metrics that incorporate that.

  - John C. Mace, Charles Morisset, and Aad van Moorsel,"Resiliency Variance in Workflows with Choice," in *Proceedings of the International Workshop on Software Engineering for Resilient Systems (SERENE 2015)*, Paris, France, September 7-8, 2015, pp. 128-143.

- Prior to last year, virtual machine monitoring approaches required modifications to the guest OS, and did not support monitoring of applications. Use of VMs is a tool in the resilient architecture toolbox, and monitoring of applications within is very necessary to detect when control actions need to be taken. We have increased our ability to implement resilient architectures by providing so-called hypervisor probes, which give us the ability to collect the information needed.

  - Zachary J. Estrada, Cuong Pham, Fei Deng, Lok Yan, Zbigniew Kalbarczyk, and Ravishankar K. Iyer, "Dynamic VM Dependability Monitoring Using Hypervisor Probes," in *Proceedings of the 11th European Dependable Computing Conference-Dependability in Practice (EDCC 2015)*, Paris, France, September 7-11, 2015, pp. 61-72.

- Our work has shown how to synthesize control algorithms for the reach-avoid problem in the presence of adversaries, something that could not be done previously. The adversary model includes sensor spoofing, controller corruption, and intrusion into the actuator. We've demonstrated the effectiveness of the approach on several example problems.

  - Zhenqi Huang, Yu Wang, Sayan Mitra, and Geir Dullerud, "Controller Synthesis for Linear Dynamical Systems with Adversaries," in *Proceedings of the Symposium and Bootcamp on the Science of Security (HoTSoS,'16)*, Pittsburgh, PA, April 20-21, 2016, pp. 53-62.

- Before our research, realizing the benefits of Software-Defined Networking (SDN) for improving network resilience required deep expertise, nontrivial training, and effort. In part, these needs reflect the complexities associated with creating and applying heuristic techniques to produce acceptable configurations faster than the rate at which the inputs change. We have developed a framework for writing SDN-based network optimization applications that raises the level of abstraction and can not only tackle requirements for a broad spectrum of SDN applications (e.g., traffic engineering, policy steering, load balancing, and topology management), but also generates (near-) optimal configurations in times compatible with application needs.

  - V. Heorhiadi, M. K. Reiter, and V. Sekar, "Simplifying Software-Defined Network Optimization Using SOL," in *Proceedings of the 13th USENIX Symposium on Networked Systems Design and Implementation*, March 2016, pp. 223-238.

- Prior to last year, security isolation was viewed primarily as a static building block for resilient architectures with only one survey of security isolation techniques, and too, an unpublished draft. Our survey paper on security isolation explores how isolation is capable of providing both static and dynamic building blocks. We observed that the key challenge for using security isolation as a dynamic building block for resilient architectures is achieving adaptability and measurability without sacrificing practical constraints. A

benefit from this understanding is in identifying primitives that provide practical adaptability based on measurable events.

  - Rui Shu, Peipei Wang, Sigmund A. Gorski, III., Benjamin Andow, Adwait Nadkarni, Luke Deshotels, Jason Gionta, William Enck, and Xiaohui Gu, "A Study of Security Isolation Techniques,"ACM Computing Surveys (CSUR), vol. 49, no. 3, December 2016.

- Software-Defined Network (SDN) has emerged as an approach with great promise for providing a foundational basis for network security. This year saw significant advances in providing tools for SDN studies and an ability to reason about correct virtualization behavior.

  - Jiaqi Yan and Dong Jin, "A Virtual Time System for Linux-Container-Based Emulation of Software-Defined Networks," in *Proceedings of the 3rd ACM SIGSIM Conference on Principles of Advanced Discrete Simulation (SIGSIMS PADS '15)*, London, UK, 2015, pp. 235-246.

  - Dong Jin and David M. Nicol, "Parallel Simulation and Virtual-Machine-Based Emulation of Software-Defined Networks," *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, vol. 26, no.1, Article 8, December 2015.

- Prior to our research, no means existed of efficiently estimating the impact of link failures on SDN connectivity. It was known how to architect configurations to assure that any k number of failures could be tolerated, but the number of failures that could be tolerated beyond that (on average) was not estimated. This year we solved that problem by showing how importance sampling techniques could be used to accelerate estimation of a configurations 'mean links to disconnection' metric.

  - David Nicol and Rakesh Kumar, "Efficient Monte Carlo Evaluation of SDN Resiliency," in *Proceedings of the 2016 Annual ACM Conference on Principles of Advanced Discrete Simulation*, Banff, Canada, May 2016, pp. 143-152.

- Modern software systems are often compositions of entities that increasingly use self-adaptive capabilities to improve their behavior to achieve systemic quality goals. Self-adaptive managers for each component system attempt to provide locally optimal results, but if they cooperated and potentially coordinated their efforts it might be possible to obtain more globally optimal results and more resilient systems. The emergent properties that result from such composition and cooperation of self-adaptive systems are not well understood, difficult to reason about, and present a key challenge in the evolution of modern software systems. For example, the effects of coordination patterns and protocols on emergent properties such as the resiliency of the collectives need to be understood when designing these systems. Lablet researchers used probabilistic model checking of Stochastic Multiplayer Games (SMG) to analyze, understand, and reason about emergent properties in Collectives of Adaptive Systems (CAS). Analysis of SMGs allows us to reason about things like the worst-case scenarios, which constitutes a new contribution to understanding emergent properties in CAS. The research demonstrated how SMGs can be useful in analyzing the impact of communication topology for collections of fully cooperative systems defending against an external attack.

  - Thomas J. Glazier, Javier Cámara, Bradley Schmerl, and David Garlan, "Analyzing Resilience Properties of Different Topologies of Collective Adaptive Systems," in *Proceedings of the 3rd FoCAS Workshop on the Fundamentals of Collective Adaptive Systems*, Boston, MA, September 21, 2015.

- Prior to this research, we lacked adequate development of metrics and models for static and dynamic assessment of resilience of software. We have defined two resiliency metrics: (1) the isolation metric to quantify the counter-measure resistance on any path (source to destination) based on the network access controls including firewalls, IPSec, IDS, proxy, and (2) the diversity metric to quantify the required attack vector by adversary based on the different disjoint attack surface due

to OS and application diversity in the attack path. We then integrate both metrics to consider the optimal isolation-diversity combination for multi-stage attack scenarios.

  - Mohammad Ashiqur Rahman, Abdullah Al Faroq, Amarjit Datta, and Ehab Al-Shaer, "Automated Synthesis of Resiliency Configurations for Cyber Networks," *IEEE Conference on Communications and Network Security (CNS)*, Philadelphia, PA, USA, October 2016.

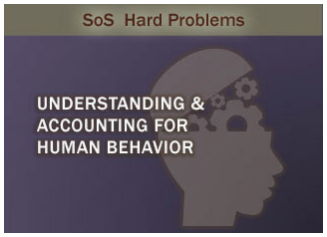

Human Behavior

The hard problem of human behavior addresses how to handle the unpredictability of human actors in cybersecurity. Computers do what we tell them to do, but humans do what they want to do, adding unpredictability and complexity to the design and implementation of computer systems. Lablet research projects were dedicated to developing models and insights of human behaviors that enable the design, modeling, and analysis of systems of people and computers with specified security properties.

- Our previous efforts revealed the cognitive and psychosocial factors that influence phishing susceptibility. We now have gained some understanding on how and why social engineering works by investigating two psychosocial factors: previous experience and personality. By studying how mental models differ between security experts and novices, we have determined that previous experience frames how individuals approach security. By classifying hundreds of phishing emails in terms of persuasion and communication, we have gained insight into how message content interacts with user traits such as personality. Such understanding could provide a basis for research on how mental models relate with phishing vulnerability and the effectiveness of phishing training.

  - Olga A. Zielinska, Allaire K.Welk, Christopher B. Mayhorn, and Emerson Murphy-Hill, "The Persuasive Phish: Examining the Social Psychological Principles Hidden in the Phishing Email Message," in *Proceedings of the Symposium and Bootcamp on the Science of Security (HoTSoS '16)*, Pittsburgh, PA, April 20-21, 2016, pp. 126-126.

  - Olga A. Zielinska, Allaire K.Welk, Christopher B. Mayhorn, and Emerson Murphy-Hill, "A Temporal Analysis of Persuasion Principles in Phishing Emails," in *Proceedings of the Human Factors and Ergonomics Society 60th Annual Meeting*, Santa Monica, CA, Human Factors and Ergonomics Society, 2016.

- Traditionally, phishing warnings focused on increasing users' awareness and understanding of phishing attacks but not the skills needed to identify phishing web pages. We incorporated a short skill training into a field phishing experiment and found that this training was essential in making a phishing warning effective. Our investigation contradicts perceived wisdom on the apparent benefits of domain highlighting in browsers. We found empirically that warnings are effective only when users are equipped with the knowledge of how to use the cued information. This study will guide the development of improved user interfaces for security warnings.

  - Aiping Xiong, Weining Yang, Ninghui Li, and Robert W. Proctor, "Improving Detection of Phishing Attacks by Directing Users' Attention to Domain Highlighted URLs," Paper presented at the *45th Annual Meeting of the Society for Computers in Psychology*, Chicago, IL, November 2015.

- Prior work did not study why users adopt some security advice and reject most of it. There was little exploration of what factors influence users' selection of different advice sources. Our work tackled this problem using both qualitative and quantitative methods. We developed a comprehensive list of advice sources and user reasoning about those sources. We found evidence of a digital security divide–wealthier, more skilled users take advice from different sources than do lower skilled users. This research can help identify improved ways to give security advice to users.

- Elissa M. Redmiles, Amelia R. Malone, and Michelle L. Mazurek, "I Think They're Trying to Tell Me Something: Advice Sources and Selection for Digital Security," in *Proceedings of the 2016 IEEE Symposium on Security and Privacy (SP 2016)*, San Jose, CA, May 2016, pp. 272-288.

- Previously, it was unclear what user-centered barriers to software patches exist and how to improve software-updating interfaces to encourage patching. Through a qualitative field study, we found that users avoid software updates because (1) updates are interruptive, (2) users are not persuaded they are essential, and (3) users cannot easily track what updates are being applied. We created a low-fidelity prototype of an improved software updating desktop interface based on our findings that (1) minimizes interruptive software update notifications, (2) provides information about updates that is easy for users to follow, and (3) provides a way to centrally manage all updates for a desktop. A human-subject evaluation showed that personalizing update interfaces, minimizing update interruptions, improving update information, and centralizing update information can improve software updating interfaces.

  - Arunesh Mathur, Josefine Engel, Sonam Sobti, Victoria Chang, and Marshini Chetty, "'They Keep Coming Back Like Zombies': Improving Software Updating Interfaces," in *Proceedings of the Twelfth Symposium on Usable Privacy and Security (SOUPS),* Denver, CO, 2016.

- Traditional approaches for differentiating legitimate human users from bots require that users "prove" they are human through either a Human Interactive Proof (HIP), i.e., explicit action by the user as in a CAPTCHA, or a Human Observational Proof (HOP), i.e., observation of user tendencies, e.g., by examining the spatial signature of mouse click locations. HIPs offer higher accuracy than HOPs but at the cost of cognitive burden and disruption to users. We have improved our understanding of an approach we proposed called Human Subtlety Proof (HSP) that is based on probing human cognition through subtle task modification. This understanding can facilitate the development of interfaces that yield improved effectiveness and improved usability.

  - Robert St. Amant and David L. Roberts, "Natural Interaction for Bot Detection," *IEEE Internet Computing*, vol. 20, no. 4, 2016.

- Before last year, we knew little about how users actually practice security aside from what they report in online surveys and interviews. Today, we have deployed a security behavior observatory with over 150 active users who provide continuous data about their risk-taking activities on the Internet. The project looks specifically at whether perceived protections (e.g., anti-virus software) induce greater risk-taking behavior by users.

  - Alain Forget, Sarah Pearman, J. Thomas, Alessandro Acquisti, Nicolas Christin, Lorrie Faith Cranor, Serge Egelman, Marian Harbach, and Rahul Telang, "Do or Do Not, There is No Try: User Engagement May Not Improve Security Outcomes," in *Proceedings of the Twelfth Symposium on Usable Privacy and Security (SOUPS)*, Denver, CO, 2016.

Carnegie Mellon University

PI William Scherlis

The Carnegie Mellon University (CMU) Science of Security (SoS) Lablet is led by Principal Investigator, William Scherlis, and includes subLablets Cornell University, University of California at Irvine and Berkeley, University of Pittsburgh, Wayne State University, University of Nebraska, and University of Texas, San Antonio. In 2016, the CMU Lablet had various projects involving approximately fifty faculty, postdoctoral, and PhD student researchers working at eight campuses. At Carnegie Mellon, researchers were drawn from six different academic departments in three different colleges. The Lablet team also worked with industry, including Bosch, Google, and Microsoft.

The CMU Lablet focused primarily on two of the five Hard Problem areas: Scalability and Composability, and Understanding and Accounting for Human Behavior. There was significant effort in two other hard problem areas: Policy-Governed Secure Collaboration, and Resilient Architectures.

FUNDAMENTAL RESEARCH

PROJECT: Security Reasoning for Distributed Systems with Certainty
PI: André Platzer
Participating Sublablet: Cornell University
Hard Problems: Scalability and Composability; Resiliency

Our research project investigates using artificial intelligence and optimization techniques for security problems related to systems with physical behavior. Many real-world devices–such as industrial plants and autonomous vehicles–use complicated control policies to guide the underlying dynamics of the system to desirable states. These control policies are, in general, difficult to analyze and prove free of security defects. Our research project is focused on improving capabilities of automatic synthesis and verification of the control of such systems, especially in the face of both benign uncertainty (e.g., sensor noise) and adversarial action. In particular, we are solving an open problem in this field: designing controllers using approximate techniques while retaining guarantees about the resulting behavior quality. As an alternative to exact approaches, we use approximate techniques that can be faster but need guarantees on their errors. In 2015 and 2016, we made substantial progress on approximately solving a class of optimization problem related to both anomaly detection and security policy synthesis. Anomaly detection and security policy synthesis are essential capabilities for security in complicated systems where uncertainty is present and must be handled in a principled fashion. We have developed a robust implementation of an approximate solver for such problems. This includes solving a wealth of issues associated with numerical stability and basis generation. Current work includes developing a novel variant of Monte Carlo Tree Search (MCTS). This MCTS variant allows us to use our approximate solver to tackle continuous real-world policy synthesis problems that are difficult to solve using other methods because of the exponential runtime scaling of conventional methods.

Carnegie Mellon University

This project has recently focused on approximate optimization techniques with specific application to security planning and policy setting. These techniques involve solving an optimization problem within a restricted basis of functions. This technique, called Galerkin approximation, allows problems to be

solved rapidly but with some loss in solution quality. We are focused on applying these methods to policy synthesis questions, but they can also be applied to anomaly detection problems. Many anomaly detection and policy synthesis questions can be cast as instances of a general problem called the Linear Complementarity Problem (LCP) (Zawadzki, Gordon, and Platzer, A Projection Algorithm for Strictly Monotone Linear Complementarity Problems, 2013).

We are currently working on LCP instances based on Markov Decision Processes (MDPs). MDPs are a popular framework for making sequential decisions in uncertain environments. Synthesizing good security polices is an example of this kind of task. For example, we may want to generate automatically a good policy for when access control restrictions should be escalated in response to anomalous system behavior, while still allowing legitimate users to work productively. Our MDP approximation techniques will improve policy synthesis capability for many of these large security problems.

Our current work ties into our earlier work on #E-SAT solving (Zawadzki, Platzer, and Gordon, A Generalization of SAT and #SAT for Robust Policy Evaluation, 2013) which offered new capabilities to analyze the robustness of security policies under uncertainty. Together, these methods represent two important links in a chain of tools for generating and verifying policies in a number of different security applications that exhibit uncertainty. For example, one can first synthesize a policy using our MDP techniques, and then can analyze the policy's robustness to random failure using our #E-SAT methods.

PROJECT: Highly Configurable Systems
PIs: Christian Kästner, Jürgen Pfeffer
Hard Problem: Scalability and Composability

We address scalability of assurances for highly configurable systems with exponentially growing configuration spaces. A compositional analysis of options will allow us to scale the analysis; it is therefore important to investigate how options are implemented and how they interact. In addition, modular and timely recertification of changes and variations is essential to make security judgments scale in practice. The team evaluated variational call graphs to predict vulnerabilities. Because vulnerable and non-vulnerable function measures can come from multiple, different statistical populations, the team adopted a bootstrap from the entire population of total functions to yield an empirical distribution of t-statistics upon which comparisons could be made.

We analyzed how data flows within applications are influenced by configuration options. We developed a measurement framework and measurement infrastructure (using a variational execution engine), finding that interactions on data flow are common but their analysis is tractable due to common sharing characteristics. This work can lead to a better understanding of interactions and toward analysis tools that can exploit typical interaction characteristics to identify security issues caused by interactions (such as leaking of data in specific combinations of configuration options).

We have interviewed sixteen developers, reviewers, and policy makers regarding software security and safety certification practices using Common Criteria and DO178, classifying challenges and opportunities, especially with regard to security-related recertification and compositional certification. We coded and classified issues and are preparing a discussion regarding security-related recertification and compositional certification.

We studied attack vectors and mitigation through malicious package updates in software ecosystems of open source software, particularly Node.js/npm. We further measured the attack surface and evaluated the feasibility of soundly verifying updates of small packages as free of certain attacks.

With our collaborators at Passau, we started a project to study the social dynamics of developer networks to better understand possible coordination problems or blind spots that cause failures to successfully address vulnerabilities in highly configurable systems. Using data from GitHub and email networks used in the development process, we have so far investigated whether we can explain the growth of developer networks through preferential attachment, replicating a proposed network growth model but using a more rigorous, maximum, likelihood-based approach for statistical inference. We also have begun to explore specifications in Stochastic Actor-Oriented Models, a class of statistical network models that are on the cutting edge of being able to model the co-evolution of behavior and network structure. These will allow us to consider preferential attachment as a proposed explanation for network growth, along with controls for reciprocity, transitivity, and node attributes.

Carnegie Mellon University

PROJECT: A Language and Framework for Development of Secure Mobile

PI: Jonathan Aldrich
Hard Problem: Scalability and Composability

Composability is a primary concern of the project. The goal is to produce a language and framework that enables the construction of secure mobile applications with known security properties. Restricting the authority of untrusted code is difficult in today's systems because, by default, code has the same authority as the user running it. Object capabilities are a promising way to implement the principle of least authority, but take away many conveniences provided by today's module systems. This year we developed a module system design that is capability-safe, yet preserves most of the convenience of conventional module systems. We demonstrated how to ensure key security and privacy properties of a program as a mode of use of our module system.

PROJECT: Race Vulnerability Study and Hybrid Race Detection

PI: Jonathan Aldrich.
Participating Sublablet: University of Nebraska
Hard Problems: Scalability and Composability, Resiliency

Today's software ecosystems present many challenges to program analysis including the need to identify problems that occur due to (1) interactions among applications, (2) mismatching interfaces between applications and underlying frameworks that change due to frequent updates, and (3) increased use of dynamic features to load code at runtime. As researchers, engineers, and analysts struggle to cope with these emerging challenges, they are forced to create complex workflows based on existing tools to meet their needs. Often, however, such approaches result in inefficient, ineffective workflows that are not practical for real-world use. As a result, defective and vulnerable applications are released. Even when defects and vulnerabilities are identified, long delays are often encountered before they can be corrected. To address these problems we developed JITANA, a new hybrid analysis framework for Android applications. JITANA has been designed to overcome the aforementioned challenges while maintaining the best parts of existing program analyses. With JITANA, existing analysis techniques continue to be effective but are significantly more efficient. To highlight the benefits of JITANA, we provide three use cases in which we show that JITANA can (1) detect shared Inter-Application Communication (IAC) channels among real-world applications, (2) support the creation of a real-time visualization engine to provide real-time feedback on code coverage, and (3) analyze large sets of applications on a device simultaneously to detect IAC channels.

Carnegie Mellon University

PROJECT: Usable Formal Methods for the Design and Comp of Security and Privacy Policies.

PI: Travis Breaux
Participating Sublablet: University of Texas, San Antonio
Hard Problems: Scalability and Composability, Policy-Governed Secure Collaboration, Human Behavior

Our research evaluates security pattern selection and application by designers in response to attack patterns. There are two threads of effort. The first area of research is based on theory in psychology concerning how designers comprehend and interpret their environment and how they plan and project solutions into the future, with the aim of better understanding how these activities exist in designing more secure systems. These are not typical models of attackers and defenders, but models of developer behavior, including our ability to influence that behavior with interventions. The evaluation is based on formal models of attack scenarios that are used to measure security risk and promote risk reduction strategies based on assurance cases constructed by the analyst. The aim is to improve the usability of formal methods for studying security design and composition. The second area focuses on tracing vulnerabilities from requirements to code, building on knowledge about information type categories. Crowdsourcing was used to collect and classify information type categories while using inter-rater reliability statistics to iteratively measure agreement during ontology construction. The results were used to find privacy policy violations in mobile applications.

PROJECT: Geo-Temporal Characterization of Security Threats

PI: Kathleen Carley
Hard Problems: Policy-Governed Secure Collaboration, Human Behavior

The goal of this project is to provide an empirical basis for understanding global security issues based on network models of interactions among actors of various kinds. Results could assist in understanding what nation-states are most threatening and may need special policies. We developed a data-informed global map of which country was purportedly attacking which, using what form of cyber-attack. Results showed wide variations in infrastructure that complicate the development of new procedures to enhance security at the global level. Results also showed that countries with high corruption and unsophisticated IT support are likely to be used by others as the apparent source of attacks. Results also showed that most global-scale attacks were between the US and China, or between former members of the USSR and Russia. However, in the past year, attacks between US and China have decreased. These results suggest the possibility of empirically-driven policy creation and evaluation.

In addition, the global cyber-attack map and the associated information and analysis from this project enable us to identify capability and IT gaps at the nation-state level. This is important in improving the selection of intervention policies and in prioritizing intervention strategies. This analysis is also critical for developing a global predictive attack model, which is the next research step, as well as providing an empirical basis for assessing human and organizational variability in capability to thwart and to engage in attacks at the nation-state level. Results provide insight into how to determine whether attacks that appear to be coming from a country are being directed out for malicious intent or whether that country is being inadvertently used by other countries and is so appearing malicious.

The results suggest that it should be possible to discriminate cyber-attacks that are primarily motivated by economic factors from those that are part of a large political attack. In particular, this analysis shows that attacks are more likely to occur from one country to another if the receiving country has greater GDP than the sender, if the receiving country and the sender are allies, and if the sending country has a higher level of corruption. In contrast, there are instances of attacks where the timing appears consistent with other hostile activities. This project was completed in March 2016.

PROJECT: USE: User Security Behavior

PI: Lorrie Cranor,
Participating SubLablets: University of California, Berkeley; University of Pittsburgh
Hard Problem: Human Behavior

The Security Behavior Observatory addresses the hard problem of "Understanding and Accounting for Human Behavior" by collecting data directly from people's own home computers, thereby capturing people's computing behavior "in the wild." This data is the closest to the ground truth of the users' everyday security and privacy challenges that the research community has ever collected. We expect the insights discovered by analyzing this data will profoundly impact multiple research domains, including but not limited to behavioral sciences, computer security and privacy, economics, and human-computer interaction. Preliminary results suggest that the extent to which people care about security is not strongly correlated with the actual level of security (or insecurity) that we observed on participants' machines. A question that then arises is whether users are more likely to engage in "risky" behavior when they believe they are operating securely.

We have successfully deployed version 5.0 and are currently deploying version 6.0 of our client data collection software. Within each of these releases, we added additional sensors to collect more information from our participants that we can apply to future research. Most notably, we added a file hash sensor to calculate hash values for downloaded files and executables and a user interface sensor to record user interactions with visual elements of the applications on their computer, such as system tray notifications.

We completed development on the next version of our client software and are deploying it to our study participants in phases over the next few months. This version of the client software primarily focused on improving the stability and performance of the client sensors and data transmission to our servers. Additionally, we included significant changes to the server applications to more easily

Carnegie Mellon University

handle a larger scale of participants, both in terms of computer resources consumed and administrative effort required to enroll and manage participants.

PROJECT: Secure Composition of Systems and Policies

PIs: Anupam Datta, Limin Jia
Hard Problems: Scalability and Composability, Resiliency

At a high-level, our work addresses the scalability and composability problem. The reasoning principles we are developing will allow both sequential and parallel composition of an arbitrary number of verified and unverified system program modules in varying configurations. Our program logic takes into consideration that components execute in a potentially adversarial environment (e.g., untrusted system modules), and therefore, the compositionality of components in the presence of such adversaries is built into our semantic model.

PROJECT: Multi-Model Run-time Security Analysis

PIs: David Garlan, Bradley Schmerl, Jürgen Pfeffer
Hard Problems: Scalability and Composability, Resiliency

This last year, we addressed composability through composing multiple semantic models (here, architectural, organizational, and behavioral), for the detection of anomalies in software systems. We made progress on these fronts: We developed software that generates large-scale architectures (in the order of a thousand software elements) from descriptions of architecture styles. We have applied this to a description of cloud-based Amazon web service systems. This is used as the basis of simulations for insider threat scenarios to test analysis of large-scale systems. We extract user traces from these simulations. We developed an algorithm to find anomalous users leveraging these traces. We used a model-based approach to cluster user sequences and find outliers. Such a technique can be useful in finding potentially anomalous users, insiders, or compromised accounts. The algorithm can also extract out the different roles that generated the traces automatically from the clusters.

Carnegie Mellon University

In a masters-level student project, we built a prototype that integrates software simulations like the one described above with social network simulations to model misinformation flow through a cyber-social network. The project uncovered a number of challenges in doing this–namely, that software simulations are predominantly deterministic, not accumulative, and do not change much over time; whereas, social network simulations are non-deterministic, accumulate state at different time steps, and change over time as agents or interactions are created or destroyed. This set of challenges helps define the agenda for how to analyze misinformation flow in cyber-social systems.

Designing secure Cyber-Physical Systems (CPS) is a particularly difficult task since security vulnerabilities stem not only from traditional cybersecurity concerns, but also from physical ones as well. Many of the standard methods for CPS design make strong and unverified assumptions about the trustworthiness of physical devices, such as sensors. When these assumptions are violated, subtle, inter-domain vulnerabilities are introduced into the system model. In this paper we propose to use formal specification of analysis contracts to expose security assumptions and guarantees of analyses from reliability, control, and sensor security domains. We showed that this specification allows us to determine where these assumptions are violated or ignore important failure modes that open the door to malicious attacks. We demonstrated how this approach can help discover and prevent vulnerabilities in a self-driving car example.

The MSIT project concluded, delivering software to convert software architecture descriptions into network descriptions that are suitable to simulation in social network analysis and simulation. We have generalized our simulator for insider anomalies to generate multiple anomalies in order to better evaluate our path-based anomaly detection algorithm.

PROJECT: Science of Secure Frameworks

PIs: David Garlan, Bradley Schmerl, Jonathan Aldrich
Participating SubLablets: University of California, Irvine; Wayne State University
Other Collaborator: George Mason University
Hard Problems: Scalability and Composability, Resiliency

This past year we investigated approaches for using software architecture coupled with static

and dynamic analysis to build scalable and composable frameworks. In cooperation with the University of California, Irvine (UCI) we developed lightweight formal verification techniques to detect several Inter-Component Communication (ICC) vulnerabilities on Android devices, and integrated them into COVERT. The approach is composable and scalable in that reevaluation of the analysis when a new app is installed is only re-run on affected areas. We also made progress on formalizing the Android permission protocol, with an aim to exhaustively analyzing and uncovering flaws in the protocols. We are in the process of updating this specification to conform to the most recent version of Android (Marshmallow).

In collaboration with UCI, we have integrated the tool for static analysis of Android ICC to generate an Acme architecture instance that can be further analyzed with architectural tools. We began development of Raindroid, a probing infrastructure using the open source Xposed framework for Android. Raindroid probes intent communication between apps in Android, and communicates with Rainbow, our self-adaptive framework, that can effect changes. We began integrating this with COVERT to identify potentially vulnerable communications, and then effect certain context-sensitive changes to the communication (such as denying, querying the user, or letting it pass). We explored this in the context of applications that dynamically load code. This work addresses the resiliency hard problem by combining information from static analysis (COVERT) with run time mitigations of vulnerabilities (Raindroid/Rainbow).

In collaboration with Wayne State University, we continued development and evaluation of Scoria, a tool that provides a semi-automated algorithm for iteratively refining global hierarchical graphs of programs for understanding their architectural structures and data flows. We also developed a dynamic web interface to begin evaluation of the tool with professionals.

Additionally, the CMU Lablet team made progress on our research assuring architectural control through a capability-based module system. In particular, we completed specifying a type system that enforces capabilities and drafted a basic soundness proof for the language. We also specified a refined concrete syntax for the module system and made progress on a prototype implementation, completed the first prototype implementation, demonstrated the idea of architectural control through a small concrete example, and made progress on extending our basic soundness proof with a proof of capability safety.

In collaboration with Bosch, we began investigating how to design frameworks to support checking of security policies, targeted at an Internet of Things framework being developed by Bosch, built on Android. We developed a tool that enforces policies on inter-app communication and tested it with three proprietary apps from Bosch, and showed that it is possible to define and enforce custom security policies within a framework.

The final thread of activity focuses on the Java sandbox, which was developed to securely encapsulate such components, but is primarily used to secure applications launched from the web while treating all of their components uniformly. We developed and evaluated MAJIC, a tool that assists developers in recovering and refining a security policy from Java bytecode and that uses the policy to sandbox targeted application components. MAJIC automates complex sandboxing tasks previously performed manually by developers. The overhead introduced by MAJIC on large, real world applications is low, with 6.6% overhead at borders between sandboxed and unsandboxed code and an additional 6.9% overhead across the application due to the use of the sandbox.

We developed a tool called Separ that combines static analysis with lightweight formal methods to automatically infer security-relevant properties from a bundle of apps. It then uses a constraint solver to synthesize possible security exploits, from which fine-grained security policies are derived and automatically enforced to protect a given device. In our experiments with over 4,000 Android apps, Separ has proven to be highly effective at detecting previously unknown vulnerabilities as well as preventing their exploitation.

We also enhanced a semi-automated algorithm and tool for iteratively and interactively refining a global hierarchical graph while maintaining its soundness, thus tackling the scalability hard problem associated with using the Scoria approach.

Due to the guarantees that they afford, formal approaches for assessing the security properties of software systems are highly

Carnegie Mellon University

desired. An example of a formal specification and analysis technique that has shown to be quite useful in security analysis of software is Alloy. However, analyzing Alloy specifications with the help of a constraint solver is known to be extremely expensive. Thus, its applications at runtime and during the evolution of a software system for evaluating its changing security properties have been limited. We have developed a new method of analyzing Alloy specifications that is substantially faster than all prior techniques for repeated analysis of an evolving specification. The tool realizing this approach is called Titanium. It provides a systematic method of reusing solutions calculated in previous analyses of an Alloy specification, thereby drastically improving the time and resources required to analyze evolving Alloy specifications. We are now exploring ways in which Titanium can be used at runtime for evaluating security properties of software. As a first step in this direction, we are experimenting with applications of Titanium to evaluate security properties of the changing configuration of Android apps installed on a mobile device.

PROJECT: Real-time Privacy Risk Evaluation and Enforcement
PI: Travis Breaux
Hard Problems: Security Metrics and Models, Human Behavior

Our research investigates new methods to measure privacy risk based on how systems collect and share personal information. Our research applies theory from psychology, judgment, and decision science concerning how individuals perceive benefits, assess risks, and make decisions about sharing cybersecurity information.

The project produced an empirically validated framework for measuring perceived privacy risk. The framework consists of a factorial vignette survey designed for collecting privacy risk measures from individuals given the benefits of sharing cybersecurity information to respond to cyber threats, and an algorithm for computing predicted privacy risk scores for independent information types. The research found that, while individuals can perceive increased risk with increased likelihood, the contribution to overall risk perception is sub-linear: there are greater perceived differences among the risks of sharing different information types, than the differences due solely to the increased likelihood of a privacy harm for a single information type. Moreover, the research shows that individuals are more willing to share information about what they do, than they are willing to share information about who they are. This indicates that privacy risk may increase non-linearly when identifiable information is combined with sensitive information types. With respect to scalability, we are currently investigating techniques to scale the information type ontology, to investigate the effect of data aggregation, and to identify cost-effective ways to re-sample privacy risk measures from individuals.

Carnegie Mellon University

PUBLICATIONS

PROJECT: Security Reasoning for Distributed Systems with Certainty

- Erik Zawadzki, "Estimating and Approximating Monotone Linear Complementarity Problems," Ph.D. Thesis, Computer Science Dept., Carnegie Mellon University, Pittsburgh, PA. Defense Date: August 3, 2016 - PENDING.

PROJECT: Highly Configurable Systems

- Gabriel Ferreira, Momim Malik, Christian Kästner, Jürgen Pfeffer, and Sven Apel, "Do #ifdefs Influence the Occurrence of Vulnerabilities? An Empirical Study of the Linux Kernel, in *Proceedings of the 20th International Systems and Software Product Line Conference (SPLC '16)*, Beijing, China, September 16-23, 2016, ACM Press, New York, NY, 2016, pp. 65-73.

- Jens Meinicke, Chu-Pan Wong, Christian Kästner, Thomas Thüm, and Gunter Saake, "On Essential Configuration Complexity: Measuring Interactions in Highly-Configurable Systems," in *Proceedings of the 31st IEEE/ACM International Conference on Automated Software Engineering (ASE 2016)*, Singapore, September 3-7, 2016, ACM Press, New York, NY, 2016, pp. 483-494.

- Christopher Bogart, Christian Kästner, James Herbsleb, and Ferdian Thung, "How to Break an API: Cost Negotiation and Community Values in Three Software Ecosystems," in *Proceedings*

of the 2016 24th ACM SIGSOFT International Symposium on the Foundations of Software Engineering (FSE 2016), Seattle, WA, November 13-18, 2016, ACM Press, New York, NY, 2016, pp. 109-120.

- Waqar Ahmad, Christian Kästner, Joshua Sunshine, and Jonathan Aldrich, "Inter-app Communication in Android: Developer Challenges," in *Proceedings of the 13th International Conference on Mining Software Repositories (MSR)*, Austin, TX, May 14-15, 2016, ACM Press, New York, NY, 2016, pp. 177-188. Presented by W. Ahmad.

- Flávio Medeiros, Christian Kästner, Márcio Ribeiro, Rohit Gheyi, and Sven Apel, "A Comparison of 10 Sampling Algorithms for Configurable Systems,"in *Proceedings of the 38th International Conference on Software Engineering (ICSE 2016)*, Austin, TX, May 14-22, 2016, ACM Press, New York, NY, 2016, pp. 643-654. Presented by F. Medeiros, M. Malik, J. Pfeffer, G. Ferreira, and C. Kästner.

- Momin M. Malik, Jürgen Pfeffer, Gabriel Ferreira, and Christian Kästner, "Visualizing the variational callgraph of the Linux Kernel: An approach for reasoning about dependencies," poster, *Proceedings of the Symposium and Bootcamp on the Science of Security (HoTSoS '16)*, Pittsburgh, PA, April 19-21, 2016, ACM Press, New York, NY, 2016, pp. 93-94. Presented by M. Malik.

- Jim Herbsleb, Christian Kästner, and Christopher Bogart, "Intelligently Transparent Software Ecosystems," *IEEE Software (IEEE-Softw.)*, vol. 33, no. 1, January 2016, pp. 89-96. DOI: 10.1109/MS.2015.156. Journal publication, no presentation.

- Waqar Ahmad, Joshua Sunshine, Christian Kästner, and Adam Wynne, "Enforcing Fine-Grained Security and Privacy Policies in an Ecosystem within an Ecosystem," in *Proceedings of the 3rd International Workshop on Mobile Development Lifecycle (MobileDeLi 2015)*, October 26, 2015, held in conjunction with Splash 2015, Pittsburgh, PA, ACM Press, New York, NY, 2015, pp. 28-34. Presented by W. Ahmad.

PROJECT: A Language and Framework for Development of Secure Mobile Application

- Cyrus Omar, Ian Voysey, Michael Hilton, Jonathan Aldrich, and Matthew Hammer, "Hazelnut: A Bidirectionally Typed Structure Editor Calculus," to appear in the *Proceedings of the ACM SIGPLAN Symposium on Principles of Programming Languages (POPL 2017)*, 2017.

- Cyrus Omar and Jonathan Aldrich, "Programmable Semantic Fragments," in the *Proceedings of the ACM SIGNPLAN International Conference on Generative Programming: Concepts and Experience (GPCE , '16)*, Amsterdam, Netherlands, October 31 - November 1, 2016, ACM Press, New York, NY, 2016.

- Filipe Militão, Jonathan Aldrich and Luís Caires, "Composing Interfering Abstract Protocols," in the *Proceedings of the European Conference on Object-Oriented Programming (ECOOP 2016)*, eds.: Shriram Krishnamurthi and Benjamin S. Lerner, vol. 56, article no. 16, Dagstuhl Publishing, Germany, 2016.

- Zack Coker, Michael Maass, Tianyuan Ding, Claire Le Goues, and Joshua Sunshine, "Evaluating the Flexibility of the Java Sandbox," in *Proceedings of the 31st Annual Computer Security Applications Conference (ACSAC 2015)*, Los Angeles, CA, December 7-11, 2015, ACM Press, New York, NY, 2015, pp. 1-10. Presented by Z. Coker.

PROJECT: Race Vulnerability Study and Hybrid Race Detection

- Lichao Sun, Zhiqiang Li, Qiben Yan, Witawas Srisa-an, and YuPan, "SigPID: Significant Permission Identification for Android Malware Detection," in *Proceedings of the 2016 11th International Conference on Malicious and Unwanted Software (MALWARE)*, Fajardo, Puerto Rico, October 18-22, 2016.

- Zhiqiang Li, Lichao Sun, Qiben Yan, Witawas Srisa-an, and Zhenxiang Chen, "Droid Classifier: Efficient Adaptive Mining of Application-Layer Header for Classifying Android Malware," to appear in *Proceedings of the*

Carnegie Mellon University

International Conference on Security and Privacy in Communication Networks (SecureComm 2016), Guangzhou, People's Republic of China, October10-12, 2016.

- Supat Rattanasuksun, Tingting Yu, Witawas Srisa-an, and Gregg Rothermel, "RRF: A Race Reproduction Framework for Use in Debugging Process-Level Races," in *Proceedings of the 2016 IEEE 27th International Symposium on Software Reliability Engineering (ISSRE)*, Ottawa, ON, Canada, 2016, pp. 162-172.

- Michael Coblenz, Joshua Sunshine, Jonathan Aldrich, Brad Myers, Sam Weber, and Forrest Shull, "Exploring Language Support for Immutability," in *Proceedings of the 38th International Conference on Software Engineering (ICSE 2016*), Austin, TX, May 14-22, 2016, ACM Press, New York, NY, 2016, pp. 736-747. Presented by M. Coblenz.

- Junjie Qian, Witawas Srisa-an, Sharad Seth, Hong Jiang, Du Li, and Pan Yi, "Exploiting FIFO Scheduler to Improve Parallel Garbage Collection Performance," in *Proceedings of the 12th ACM SIGPLAN/SIGOPS International Conference on Virtual Execution Environments (VEE '16)*, Atlanta, GA, April 2-3, 2016, ACM Press, New York, NY, 2016, pp. 109-121. DOI: 10.1145/2892242.2892248. Presented by J. Qian.

- Junjie Qian, Witawas Srisa-an, Du Li, Hong Jiang, Sharad Seth, and Yaodong Yang, "SmartStealing: Analysis and Optimization of Work Stealing in Parallel Garbage Collection for Java VM," in *Proceedings of the International Conference on Principles and Practices of Programming on the Java Platform (PPPJ '15)*, Melbourne, FL, September 8-11, 2015, ACM Press, New York, NY, 2015, pp. 170-181. Presented by J. Qian.

PROJECT: Usable Formal Methods for the Design and Comp of Security and Privacy Policies

- Hanan Hibshi, Travis D. Breaux, and Christian Wagner, "Improving Security Requirements Adequacy: An Interval Type 2 Fuzzy Logic Security Assessment System," in the *2016 IEEE Symposium Series on Computational Intelligence (SSCI)*, Athens, Greece, December 6-9, 2016, pp. 1-8.

- Hanan Hibshi, Travis D. Breaux, Maria Riaz, and Laurie Williams, "A Grounded Analysis of Experts' Decision-Making During Security Assessments,"*Journal of Cybersecurity*, vol. 2, no. 2, Oxford University Press, published online 4 October 2016. DOI: 10.1093/cybsec/tyw010.

- Jaspreet Bhatia, Morgan C. Evans, Sudarsham Wadkar, and Travis D. Breaux, "Automated Extraction of Regulated Information Types Using Hyponymy Relations," in *Proceedings of the 2016 IEEE 24th International Requirements Engineering Conference Workshops (REW): Third International Workshop on Artificial Intelligence and Requirements Engineering (AIRE '16)*, Beijing, China, September 2016, pp. 19-25.

- Rocky Slavin, Xiaoyin Wang, Mitra Bokaei Hosseini, William Hester, Ram Krishnan, Jaspreet Bhatia, Travis D. Breaux, and Jianwei Niu, "Toward a Framework for Detecting Privacy Policy Violation in Android Application Code," in *Proceedings of the 38th International Software Engineering Conference (ICSE 2016*), Austin, Texas, May 14-22, 2016, ACM Press, New York, NY, 2016, pp. 25-36. Presented by R. Slavin.

- Thomas MacGahan, Claiborne Johnson, Armando L. Rodriguez, Mark Appleby, Jianwei Niu, and Jeffery von Ronne, " Extended Abstract: Towards Verified Privacy Policy Compliance of an Actor-based Electronic Medical Record Systems," in *Proceedings of the 5th International Workshop on Programming Based on Actors, Agents, and Decentralized Control (AGERE! 2015)*, held in conjunction with Splash 2015, Pittsburgh, PA, October 25-30, 2015. Presented by T. MacGahaen.

PROJECT: USE: User Security Behavior

- Casey Canfield, Baruch Fischoff, Alex Davis, Alain Forget, Sarah Pearman, and Jeremy Thomas, "Comparing Phishing Vulnerability in the Lab to the Real World," submitted to *CHI 2017*, pending review.

- Sarah Pearman, Arnab Kumar, Nicholas Munson, Charu Sharma, Leeyat Slyper, Jeremy Thomas, Lujo Bauer, Nicolas Christin, and Serge Egelman, "Risk Compensation in Home-User Computer Security Behavior: A Mixed-Methods Exploratory Study," poster and extended abstract, presented at the *12th Symposium on Usable Privacy and*

Carnegie Mellon University

Security (SOUPS 2016), Denver, CO, June 22-24, 2016. Presented by S. Pearman. Received a SOUPS 2016 Distinguished Poster Award.

- Alain Forget, Sarah Pearman, Jeremy Thomas, Alessandro Acquisti, Nicolas Christin, Lorie Faith Cranor, Serge Egelman, Marian Harbach, and Rahul Telang, "My Daughter Fixes All My Mistakes: A Qualitative Study on User Engagement and Computer Security Outcomes,' in *Proceedings of the 12th Symposium on Usable Privacy and Security (SOUPS 2016)*, Denver, CO, June 22-24, 2016. Presented by A. Forget.

- Serge Egelman, Marian Harbach, and Eyal Peer, "Behavior Ever Follows Intention? A Validation of the Security Behavior Intentions Scale (SeBIS)," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '16)*, Santa Clara, CA, May 22-27, 2016, pp. 5257-5261. Presented by S. Egelman. Received a CHI '16 Best Paper Honorable Mention Award.

- Alain Forget, Sarah Pearman, Jeremy Thomas, Alessandro Acquisti, Nicolas Christin, Lorrie Cranor, Serge Egelman, Marian Harbach, and Rahul Telang, "Do or Do Not, There is No Try: User Engagement May Not Improve Security Outcomes," in *Proceedings of the 12th Symposium on Usable Privacy and Security (SOUPS 2016)*, Denver, CO, June 22-24, 2016.

- Serge Egelman, Marian Harbach, and Eyal Peer, "Behavior Ever Follows Intention? A Validation of the Security Behavior Intentions Scale (SeBIS)," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '16)*, San Jose, CA, May 7-12, 2016, pp. 5257-5261. Presented by S. Egelman. Received a CHI '16 Best Paper Honorable Mention Award.

PROJECT: Secure Composition of Systems and Policies

- Amit Vasudevan, Sagar Chaki, Petros Maniatis, Limin Jia, and Anupam Datta, "UberSpark: Enforcing Secure Object Abstractions for Automated Compositional Security Analysis of a Hypervisor," in *Proceedings of the USENIX Security Symposium*, Austin, TX, 2016, pp. 87-104.

- Limin Jia, Shayak Sen, Deepak Garg, and Anupam Datta, "A Logic of Programs with Interface-confined Code," in *Proceedings of the 2015 IEEE 28th Computer Security Foundations Symposium (CSF)*, Verona, Italy, July 13-17, 2015, pp. 512-525.

PROJECT: Multi-Model Run-time Security Analysis

- Hemank Lamba, Thomas J. Glazier, Bradley Schmerl, Javier Cámara, David Garlan, and Jürgen Pfeffer, "A Model-based Approach to Anomaly Detection in Software Architectures," poster, *Proceedings of the Symposium and Bootcamp on the Science of Security (HoTSoS '16)*, Pittsburgh, PA, April 19-21, 2016, ACM Press, New York, NY, 2016, pp. 69-71. T. Glazier presented.

- Ivan Ruchkin, Ashwini Rao, Dionisio De Niz, Sagar Chaki, and David Garlan, "Eliminating Inter-Domain Vulnerabilities in Cyber-Physical Systems: An Analysis Contracts Approach," in *Proceedings of the First ACM Workshop on Cyber-Physical Systems Security and/or PrivaCy (CPS-SPC)*, Denver, Colorado, October 16, 2015, pp. 11-22. I. Ruchkin presented.

- Thomas J. Glazier, Javier Cámara, Bradley Schmerl, and David Garlan, "Analyzing Resilience Properties of Different Topologies of Collective Adaptive Systems," in *Proceedings of the 3rd FoCAS Workshop on the Fundamentals of Collective Adaptive Systems*, Boston, MA, September 21, 2015. T. Glazier presented.

PROJECT: Science of Secure Frameworks

- Hamid Bagheri and Sam Malek, "Titanium: Efficient Analysis of Evolving Alloy Specifications," in *Proceedings of the 2016 24th ACM SIGSOFT International Symposium on the Foundations of Software Engineering (FSE 2016)*, Seattle, WA, November 13-18, 2016, ACM Press, New York, NY, 2016, pp. 27-38.

- Bradley Schmerl, Jeffrey Gennari, Alireza Sadeghi, Hamid Bagheri, Sam Malek, Javier Cámara, and David Garlan, "Architecture Modeling and Analysis of Security in Android Systems," in *Proceedings of the 10th European Conference*

Carnegie Mellon University

on Software Architecture (ECSA 2016), vol. 9839 of Lecture Notes in Computer Science, Springer, Copenhagen, Denmark, 30 November - 2 December 2016.

- Javier Cámara, David Garlan, Gabriel A. Moreno, and Bradley Schmerl, "Evaluating Trade-Offs of Human Involvement in Self-Adaptive Systems," in *Managing Trade-Offs in Adaptable Software Architectures*, eds.: Ivan Mistrik, Nour Ali, Rick Kazman, John Grundy, and Bradley Schmerl, Elsevier, published online 11 August 2016, chap. 7.

- Hamid Bagheri, Alireza Sadeghi, Reyhaneh Jabbarvand Behrouz, and Sam Malek, "Practical, Formal Synthesis and Autonomic Enforcement of Security Policies for Android," in *Proceedings of the 46th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN 2016)*, Toulouse, France, June 2016, pp. 514-525. DOI: 10.1109/DSN.2016.53

- Nariman Mirzaei, Joshua Garcia, Hamid Bagheri, Alireza Sadeghi, and Sam Malek, "Reducing Combinatorics in GUI Testing of Android Applications," in *Proceedings of the 38th International Conference on Software Engineering (ICSE 2016)*, Austin, TX, May 14-22, 2016, ACM Press, New York, NY, 2016, pp. 559-570. Presented by N. Mirzaei.

- Bradley Schmerl, Jeffrey Gennari, Javier Cámara, and David Garlan, "Raindroid — A System for Run-time Mitigation of Android Intent Vulnerabilities," poster, *Proceedings of the Symposium and Bootcamp on the Science of Security (HoTSoS '16)*, Pittsburgh, PA, April 19-21, 2016, ACM Press, New York, NY, 2016, pp. 115-117. Presented by B. Schmerl and J. Cámara.

- Marwan Abi-Antoun, Ebrahim Khalaj, Radu Vanciu, and Ahmad Moghimi, "Abstract Runtime Structure Reasoning about Security," poster, *Proceedings of the Symposium and Bootcamp on the Science of Security (HoTSoS '16)*, Pittsburgh, PA, April 19-21, 2016, ACM Press, New York, NY, 2016, pp. 1-3. Presented by E. Khalaj http://www.cs.wayne.edu/~mabianto/papers/16-hotsos.pdf/

- Eric Yuan and Sam Malek, "Mining Software Component Interactions to Detect Security Threats at the Architectural Level," in *Proceedings of the 13th Working IEEE/IFIP Conference on Software Architecture (WICSA)*, Venice, Italy, 2016, pp. 211-220. Presented by E. Yuan.

- Michael Maass, "A Theory and Tools for Applying Sandboxes Effectively," Carnegie Mellon University Doctoral Dissertation, CMU-ISR-16-105, March 2016. Thesis dissertation was defended by Michael Maass on March 7, 2016.

- Naeem Esfahani, Eric Yuan, Kyle R. Canavera, and Sam Malek, "Inferring Software Component Interaction Dependencies for Adaptation Support," *ACM Transactions on Autonomous and Adaptive Systems (TAAS)*. vol. 10, no. 4, Article 26. February 3, 2016.

- Michael Maass, Adam Sales, Benjamin Chung, and Joshua Sunshine, "A Systematic Analysis of the Science of Sandboxing," *PeerJ Computer Science*, vol. 2, no. 3, January 27, 2016.

- Nariman Mirzaei, Hamid Bagheri, Riyadh Mahmood, and Sam Malek, "SIG-Droid: Automated System Input Generation for Android Applications," in *Proceedings of the 26th IEEE International Symposium on Software Reliability Engineering (ISSRE)*, Gaithersburg, MD, November 2-5, 2015. Presented by S. Malek.

- Ivan Ruchkin, Ashwini Rao, Dionisio De Niz, Sagar Chaki, and David Garlan, "Eliminating Inter-Domain Vulnerabilities in Cyber-Physical Systems: An Analysis Contracts Approach," in *Proceedings of the First ACM Workshop on Cyber-Physical Systems Security and/or PrivaCy*, Denver, Colorado, October 12-16, 2015, pp. 11-22. Presented by I. Ruchkin.

- Waqar Ahmad, Joshua Sunshine, Christian Kästner, and Adam Wynne, "Enforcing Fine-Grained Security and Privacy Policies in an Ecosystem within an Ecosystem", in *Proceedings of the 3rd International Workshop on Mobile Development Lifecycle (MobileDeLi 2015)*, October 26, 2015, held in conjunction with Splash 2015, Pittsburgh, PA, ACM Press, New York, NY, 2015, pp. 28-34. Presented by W. Ahmad.

- Gabriel A. Moreno, Javier Cámara, David Garlan and Bradley Schmerl, "Proactive Self-

Carnegie Mellon University

Adaptation Under Uncertainty: a Probabilistic Model Checking Approach," in *Proceedings of the 2015 10th Joint Meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on the Foundations of Software Engineering*, Bergamo, Italy, August 30 - September 4, 2015. Presented by G. Moreno.

- Ebrahim Khalaj and Marwan Abi-Antoun, "What You See Is What You Get Object Graphs," Technical Report, Wayne State University, Detroit, MI, May 2016, 21 pages. http://www.cs.wayne.edu/~mabianto/inprep/16-iwaco.pdf

EDUCATIONAL

PROJECT: Highly Configurable Systems

- Course 15-313 Foundation of Software Engineering, Fall 2015

PROJECT: Usable Formal Methods for the Design and Comp of Security and Privacy Policies

- 08-605 Engineering Privacy, Spring 2016. In this class, we introduced new training on trading privacy for data utility. We also began developing a new course module on developing synthetic data for privacy design analyses using the Eddy language.

PROJECT: Geo-Temporal Characterization of Security Threats

- Course 08-801, 08-640, 19-640. Dynamic Network Analysis, Spring 2016

- Developed new module on tracking and analyzing big data, and had some students do projects using SoS related data – e.g., insider threat work, NetFlow analysis, and Arbor data assessment

PROJECT: USE: User Security Behavior

- A student group in Nicolas Christin & Lujo Bauer's Spring 2016 Usable Privacy & Security course used Security Behavior Observatory data extensively in their semester project, which explored the application of risk homeostasis theory to end-user security behaviors.

- A group of CMU master's degree students worked extensively with the SBO data for their final capstone project in Spring 2016. They generated a report that examined various security behaviors, including risky browsing behaviors that we will build on for future research.

COMMUNITY ENGAGEMENTS

WORKSHOPS: SĀF|ART|INT:

Carnegie Mellon University (CMU) held an "Exploratory Technical Workshop on Safety and Control for Artificial Intelligence," on June 27, 2016. Results from this one-day workshop informed the discussions at a workshop held the following day that was co-sponsored by the White House Office of Science and Technology Policy (OSTP) and Carnegie Mellon. Workshop presenters came from government, industry, and academia, including the White House OSTP, the Defense Advanced Research Projects Agency (DARPA), the National Science Foundation (NSF), the National Security Council (NSC), the Intelligence Advanced Research Projects Agency (IARPA), Google, Microsoft, the Future of Life Institute, Uber, ZF-TRW, Brown University, Vanderbilt University, Tufts University, Oregon State University, the University of California, Berkeley, and Carnegie Mellon University.

Approximately 130 participants attended the Carnegie Mellon Exploratory Workshop. Speakers included Ed Felten (White House OSTP), William Scherlis (CMU, chair), Eric Horvitz (Microsoft), Andrew Moore (CMU), Richard Mallah (Future of Life Institute), Tom Mitchell (CMU), Dario Amodei (Google Brain), Claire Le Goues (CMU), and Robert Rahmer (IARPA).

Speakers at the "Public Workshop on Safety and Control for Artificial Intelligence (SĀF|ART|INT) on June 28, 2016, co-sponsored by the White House OSTP and Carnegie Mellon University, included the

Carnegie Mellon University

following:

Ed Felten (White House OSTP)
William Scherlis (CMU)
Manuela Veloso (CMU)
John Launchbury (DARPA)
Jason Matheny (IARPA)
Andrew Grotto (White House NSC)
Claire Tomlin (UC Berkeley)
Thomas Dietterich (Oregon State)
Emma Brunskill (CMU)
Michael Littman (Brown)
Jeannette Wing (Microsoft)
Kathleen Fisher (Tufts)
Anupam Datta (CMU)
Sarah Loos (Google)
Michael Wagner (CMU)
Drew Bagnell (Uber)
Reid Simmons (NSF)
Brian Murray (ZF-TRW)
Douglas Schmidt (Vanderbilt)

Most talks from the workshops are available by video at the combined workshops website www.cmu.edu/safartint/.

PROGRAM: CMU ISR REU

CMU Lablet faculty leads Joshua Sunshine and Christian Kästner led a summer program, "Research Experience for Undergraduates (REU)," sponsored by the National Science Foundation (NSF). This was a highly competitive program–more than one hundred students from diverse universities applied for eighteen summer positions. These students were resident on campus working with faculty in the CMU Institute for Software Research (ISR), with many assigned to Lablet projects.

PROJECTS

PROJECT: Highly Configurable Systems

- Invited Keynote: Christian Kästner, "Quality Assurance for Highly Configurable Systems," VACE Workshop (Workshop on Variability and Complexity in Software Design) at ICSE, May 2016.
- Keynote: Christian Kästner, "Parsing Unpreprocessed C Code - The TypeChef Experience," Parsing@SLE Workshop at SPLASH 2015.
- Keynote: Christian Kästner, "Understanding Feature Interactions: From Bugs to Performance Surprises," Brazilian Symposium on Software Components, Architectures, and Reuse (SBCARS), September 2015.
- Talk: Christian Kästner, "Differential Testing for Variational Analyses," FOSD 2016 Working meeting on Feature-Oriented Software Development, May 2016.
- Talk: Gabriel Ferreira,"Software Certification: Composition, Evolution and Reuse Challenges," FOSD 2016 - Working meeting on Feature-Oriented Software Development, May 2016.
- Industrial Collaboration: We collaborated with Itemis GmbH.

PROJECT: Usable Formal Methods for the Design and Comp of Security and Privacy Policies

- Tutorial: Hanan Hibshi and Travis Breaux, "Grounded Analysis," 2015 IEEE International Requirements Engineering Conference, August 2015. Twenty-eight faculty, students, and practitioners registered and attended this tutorial.
- Talk: Kevin Baldor, "Declarative Deadlines in Functional-Reactive Programming," Workshop on Reactive and Event-based Languages & Systems (REBLS) at SPLASH 2015, Pittsburgh, PA, October, 2015. Seven-page paper was co-authored by Jianwei Niu.
- Panel: Jianwei Niu, "Security and Privacy in the Age of Internet of Things: Opportunities and Challenges,"Proceedings of the 21st ACM Symposium on Access Control Models and Technologies (SACMAT'16), pp. 49-50. Shanghai China, June 2016. Two-page paper was co-authored by Y. Jin, A. J. Lee, R. Sandhu, W. Xu, and X. Zhang.

Carnegie Mellon University

PROJECT: Geo-Temporal Characterization of Security Threats

Keynotes/Invited Talks:

- Kathleen Carley, "Social Media Network Analytics," NATO ACT-Sprint, Kraków, Poland, Plenary, April 2016.
- Kathleen Carley, "Dark Networks and Dynamic Network Analytics," Dark Web Conference, Naval Post Graduate School, Monterey, CA, Plenary, March 2016.
- Ghita Mezzour, "A Socio-Technical Approach to Global Cyber Security," 11th International Conference on Information Assurance and Security, Marrakesh, Morocco, December 2015.
- Ghita Mezzour, "A Big Data Approach to Studying International Cyber Security," Science and Culture Week, Ibn Tofail University, Kénitra, Morocco, May 2016.

Presentations

Ghita Mezzour, "A Big Data Approach to Studying International Cyber Security," Cyber Security Day, International University of Rabat, Rabat, Morocco, May 2016.

Sumeet Kumar, "More Noise does not Mean Effective Policy: Understanding the Impact of US Cyber Policies on Cyber-Attacks Trend," Poster for and entry for the SBP-BRiMS 2016 Challenge event, SBP-BRiMS 2016 Conference, Washington DC, May 2016.

PROJECT: USE: User Security Behavior

- Presentation: Alain Forget, "Do or Do Not, There is No Try: User Engagement May Not Improve Security Outcomes." 12th Symposium on Usable Privacy and Security (SOUPS 2016), Denver, CO, June 22-24, 2016. Fifteen-page paper was co-authored by S. Pearman, J. Thomas, A. Acquisti, N. Christin, L.F. Cranor, S. Egelman, M. Harbach, and R. Telang.
- Presentation: Serge Egelman, "Behavior Ever Follows Intention?: A Validation of the Security Behavior Intentions Scale (SeBIS)." SIGCHI Conference on Human Factors in Computing Systems (CHI '16), San Jose, CA, May 7-12, 2016. Five-page paper was co-authored by Marian Harbach and Eyal Peer. Received a CHI '16 Best Paper Honorable Mention award.
- Poster: Sarah Pearman, "Risk Compensation in Home-User Computer Security Behavior: A Mixed-Methods Exploratory Study," 12th Symposium on Usable Privacy and Security (SOUPS 2016), Denver, CO, June 22-24, 2016. Poster was co-authored by A. Kumar, N. Munson, C. Sharma, L. Slyper, J. Thomas, Lujo Bauer, Nicolas Christin, and Serge Egelman.

PROJECT: Multi-Model Run-time Security Analysis

- Keynote: David Garlan, "Reflections on the Past, Present, and Future of Software Architecture,"First International Workshop on Bringing Architecture Design Thinking into Developers' Daily Activities (Bridge,'16) at the 38th International Conference on Software Engineering, Austin TX, May 14 - 22, 2016.
- Co-organizer: 2nd International Workshop on Software Engineering for smart Cyber-Physical Systems (SEsCPS '16) at the 38th International Conference on Software Engineering, Austin TX, May 16, 2016. Security of cyber-physical systems was an emerging theme.
- Industrial Collaboration: We collaborated with General Dynamics | Viz on a project to integrate social- and cyber-network simulations.

PROJECT: Science of Secure Frameworks

- Industrial Collaboration: We collaborated with Bosch on transitioning Michael Maass's sandboxing technique to their IoT platform Bezirk.

Carnegie Mellon University

NORTH CAROLINA STATE UNIVERSITY

PI Laurie Williams

North Carolina State University's (NCSU) Science of Security (SoS) Lablet, led by Principal Investigator Laurie Williams, has continued to support the National Security Agency's (NSA's) vision for the Science of Security and the SoS community. We have emphasized data-driven discovery and analytics to formulate, validate, evolve, and solidify the theory and practice of security. Efforts in our current lablet have yielded significant findings, providing a deeper understanding of users' susceptibility to deception, developers' adoption of security tools, and how trust between people relates to their commitments. Motivated by NSA's overarching vision for SoS, and building on our experience and accomplishments, we are (1) developing a science-based foundation for the five Hard Problems that we previously helped formulate; and (2) fostering a SoS community with high standards for reproducible research. Our approach involves a comprehensive, rigorous perspective on SoS, including an integrated treatment of technical artifacts, humans (both stakeholders and adversaries) along with relationships and processes relevant to the Hard Problems. Continual evaluation of our research and community development efforts is key to our approach.

Team Overview

We have formed teams to conduct scientific research and evaluate progress on Hard Problems: Security Metrics and Models; Humans; Policy; and Resilient Architectures. The Scalability and Composability Hard Problem has no explicit team since we address it as a secondary hard problem in several of our projects. Each hard problem team is composed of three or four projects researching complementary aspects of the hard problem. We also have additional teams for Research Methods, Community Development and Support, and for Evaluation.

- **Research Methods, Community Development and Support:**
  University of Alabama - Jeff Carver
  NC State University - Lindsey McGowen, Jon Stallings, Laurie Williams, David Wright

- **Evaluation:**
  NC State University - Lindsey McGowen, Jon Stallings, David Wright, University of Alabama - Jeff Carver

FUNDAMENTAL RESEARCH

PROJECT: Attack Surface and Defense-in-Depth Metrics

PIs: Andy Meneely, Laurie Williams, Munindar P. Singh
Participating Sublablet: Rochester Institute of Technology
Hard Problems: Security Metrics and Models, Scalability and Composability, Resiliency

Our main technical accomplishment this year has been mapping the attack surface using stack traces.

Laurie Williams and Christopher Theisen were able to show that attack surface data can be used to predict vulnerabilities in binaries that crashed and had stacktraces reported. Beyond that, we have developed a prediction model that uses random walks on call graphs to predict vulnerabilities at the method level. This prediction model out-performs the literature in vulnerability prediction.

PROJECT: Systematization of Knowledge from Intrusion Detection Models

PIs: Huaiyu Dai, Andy Meneely,
Participating Sublablet: Rochester Institute of Technology,
Hard Problems: Security Metrics and Models, Resiliency, Scalability and Composability, Humans

Our main technical accomplishment has been the data collection, processing, and analysis of Intrusion Detection System (IDS) literature for a systematic literature review. Publication of these results is forthcoming. Specifically, we have discovered that IDS literature does not use consistent evaluation metrics. Researchers will often pick the evaluation metrics that cast their own research in the best light, and ignore trade-offs in their analysis. For example, in samples we examined that use "precision" as an evaluation metric, less than half used "recall" despite the widespread use of presenting both precision and recall together in statistical communities. These inconsistencies span multiple categories of metrics, such as failing to report space metrics when speed was being evaluated, or only evaluating precision and recall metrics without metrics of speed and space. Worse yet, this problem has persisted over the years and has not improved as the number of IDS papers per year has increasingly grown. These inconsistencies make systematization a considerably challenging task in the intrusion detection research community.

In a second area, we initiated the investigation of game-theoretic approaches to addressing the dynamic interplay between the intruders and IDS. In particular, we achieved two accomplishments in this direction. First, we tackled the challenging dynamic IDS configuration problem under the scientific framework of stochastic game with incomplete information, and proposed a new algorithm, Bayesian Nash-Q learning, to solve it. We then explored the benefits of collaboration among IDS systems. The problem is formulated as a two-layer game: the first layer game models the interplay between each IDS and its corresponding attackers, while the second-layer game models the collaboration among the IDSs. We finished the algorithm design for collaborative IDS configuration, and tested it through extensive simulations. Simulation results indicate that the proposed scheme can facilitate effective resource-sharing among IDSs, leading to significant gain in detection performance.

PROJECT: Vulnerability and Resilience Prediction Models

PIs: Mladen Vouk, Laurie Williams,Ä®Hard Problems: Security Metrics and Models, Resiliency, Scalability and Composability

Resilience of software to attacks is an open problem. Resilience depends on the science behind the approach used, as well as on our engineering abilities. The scope includes recognition of attacks through metrics and models we use to describe and identify software vulnerabilities, and the models we use to predict resilience to attacks in the field. It also depends on the software (and system) architecture(s) used, and their scalability. For example, if one has a number of highly attack-resilient components and appropriate attack sensors, is it possible to compose a resilient system from these parts? And, how does that solution scale and age? Cyber-attacks and breaches are often detected too late to avoid damage. "Classical" reactive cyber defenses usually work only if we have some prior knowledge about the attack methods and "allowable" patterns. Properly constructed redundancy-based anomaly detectors can be more robust and adaptable, and often they are able to detect even zero day attacks. Mitigation and management of detected issues can then follow in a number of ways. In the world where Internet of Things (IoT) elements are a routine component of a workflow, security will be orders of magnitude more difficult unless we make those elements security aware and self-defending from the start. During the last year, we have shown through both experimental and theoretical work that redundancy-based security anomaly detectors are viable and have considerable ability to recognize some high-risk and difficult to detect attacks (including zero-day attacks) on web servers - likely management interface for many IoT elements.

In parallel, we have been investigating

security of cloud-based application chains that may also benefit from pro-active resilience. We find that (1) three security properties (i.e., input validation, remote access validation, and data integrity) are essential for making such workflows more secure; and (2) use of a security-aware provenance collection can help secure such chains. We are working on a model that integrates IoT based attack detectors into workflow (application chain) resilience solutions.

PROJECT: Warning of Phishing Attacks: Supporting Human Information Processing, Identifying Phishing Deception Indicators, and Reducing Vulnerability

PIs: Christopher Mayhorn, Emerson Murphy-Hill,
Hard Problem: Human Behavior

In the past year, we completed two behavioral studies associated with phishing susceptibility and we are in the process of beginning data collection on a third study. In our initial study, we explored how the mental models of security experts and novices differed with regard to phishing-related terms. The terms were divided into three categories: prevention of phishing, trends and characteristics of phishing attacks, and the consequences of phishing. Expert mental models were more complex with more links between concepts. Specifically, experts had sixteen, thirteen, and fifteen links in the networks describing the prevention, trends, and consequences of phishing, respectively; however, novices only had eleven, nine, and nine links in the networks describing prevention, trends, and consequences of phishing, respectively. In our second study, two reviewers assessed a cache of eight hundred eighty-seven phishing emails from Arizona State University, Brown University, and Cornell University by examining them for attributes consistent with Cialdini's six principles of persuasion: authority, social proof, liking/similarity, commitment and consistency, scarcity, and reciprocation. A correlational analysis of email characteristics by year revealed that the persuasion principles of commitment and consistency and scarcity have increased over time, while the principles of reciprocation and social proof have decreased over time. Authority and liking/similarity revealed mixed results with certain characteristics increasing and others decreasing. Most recently, we applied for and received IRB approval from NCSU and the NSA to conduct a third study that will explore the interaction between persuasive attributes in phishing emails and user personality characteristics. Stimulus development and programming is well under way with data collection to begin this summer.

NC STATE UNIVERSITY

PROJECT: A Human Information-Processing Analysis of Online Deception Detection

PIs: Robert W. Proctor, Ninghui Li
Participating Sublablet: Purdue University
Hard Problem: Human Behavior

We completed the field study of a phishing warning Chrome extension in which we carried out a simulated phishing attack that bypassed the currently deployed defenses and reached almost all participants. Our results demonstrate the warning extension's ability to protect users against phishing, and the importance of combining skill training with understandable warning messages. We also evaluated the influence of domain highlighting in two experiments in which participants judged the legitimacy of web pages. Instructions to attend to the address bar improved detection of fraudulent web pages, whereas domain highlighting had little influence. In the second experiment, analysis of eye-gaze fixation measures showed that people attend to the highlighted domain on the address bar but this did not impact their judgments. This outcome implies that users lack knowledge of webpage security cues, or how to use those cues.

PROJECT: Leveraging the Effects of Cognitive Function on Input Device Analytics to Improve Security
PIs: David L. Roberts, Robert St. Amant
Hard Problem: Human Behavior

We developed a set of cognitive labels that are based on a first-pass at segmenting typing data collected during our studies. The labels include descriptions of motor processing, visual processing, and memory cognitive phenomena. Ultimately, the sequence, duration, and intervals between these labels will serve as descriptions of "normal"or "expected" behaviors for HSPs. To supplement those data, we collected eye tracking data during task interactions to provide additional insight into cognition. The eye motion data enables us to become much more detailed in labeling segments of time-series data with lower-level cognitive processes. We have processed the eye tracking data to compute gaze fixations and saccades. Fixations and saccades are extremely

useful for identifying points in the data that are reflective of the different cognitive processes we're interested in modeling. Visualization tools have also been invaluable in hypothesis generation for features of the log data that reflect the cognitive processes we're interested in detecting. Using our visualization tool and the data we collected, we have discovered the need to identify with a high-degree of confidence the "perceptual segmentation" of words during transcription typing. Our data suggest (and existing literature backs up) that users read words on the screen in a way that reflects dividing words into smaller, more recognizable and easily-spelled chunks. For the tasks we're focusing on, characterizing these chunks and the way users identify them is critical to success. Accordingly, we refined our visualization tool to facilitate a more detailed exploration of user task performance and enable a realtime comparison of the empirical data to an arbitrary cognitive model. Our tool now supports applying cognitive labels to sequences of task data based on human annotation in addition to model-based annotations. The model-based annotations can be integrated from existing cognitive modeling tools like ACT-R. Using this ACT-R integration, we have identified parameters that, when tuned, explain the differences in the typing speed based on familiarity of word being typed. Taken together, the visualization, data modeling, and cognitive modeling have laid a strong foundation for HSPs moving forward. Our efforts have resulted in tools and techniques to identify, describe, and probe specific cognitive functions during typing data tasks, and our next steps will be to implement primitive HSPs in more complex environments.

PROJECT: Understanding Effects of Norms and Policies on the Robustness, Liveness, and Resilience of Systems
PIs: Emily Berglund, Jon Doyle, Munindar Singh
Hard Problems: Policy-Governed Secure Collaboration, Resiliency

This year we have focused on: (1) identification of mathematical concepts relevant to expressing precise measures of system stability and resilience; (2) development of simple yet realistic abstract models of research organizations self-governed by self-enforced productivity and security norms; and (3) evaluation of hypotheses about these models using both computer simulations and experiments involving humans playing games based on the models.

PROJECT: Formal Specification and Analysis of Security-Critical Norms and Policies
PIs: Rada Y. Chirkova, Jon Doyle, Munindar P. Singh
Hard Problems: Policy-Governed Secure Collaboration, Scalability and Composability

Our work has focused on: (1) development of a formal language, formal semantics, and methodology for expressing, assessing, and verifying practicable statements of security, technical, regulatory, and social norms, including information about what norms govern action when norms conflict; and (2) development of studies and experiments to assess the usability and comparative usability of this language and methodology.

PROJECT: Scientific Understanding of Policy Complexity
PIs: Ninghui Li, Robert W. Proctor, Emerson Murphy-Hill
Participating Sublablet: Purdue University
Hard Problems: Policy-Governed Secure Collaboration, Human Behavior

We studied the reasons why large firewall policies are often too complex to understand and are error-prone, and identified three factors: (1) firewall rules may conflict with each other; (2) policies expressed in ACL-based languages are monolithic; and (3) complex policies require a large number of rules. A monolithic policy can only be understood as a whole. This becomes infeasible as the policy gets large, since most people are unable to put a large amount of information in the working memory. To reduce the complexity, we decided to tackle the monolithic nature of current firewall policies and developed an approach to specify modular policies. We identified five requirements for a successful modularization approach (i.e., logical partitioning into modules, isolation among components, flexible partitioning structure, human-computable policy slicing, and capability to be readily deployed). We introduced our Tri-modular approach of modularizing firewall policies. This includes identifying a primary attribute, which is either the source IP, the destination IP, or the service, and organizing a policy into three kinds of modules: primary, auxiliary, and template. Beyond making policies more modular and easier to understand, our approach also supports policy refactoring, either by distilling templates

NC STATE UNIVERSITY

from recurring patterns, or by breaking up a large module into multiple smaller ones, each covering a subset of the IP range. To support legacy firewall policies, we have defined a 5-step process and introduced algorithms for converting them into their modularized form. We have also implemented an automated tool for this purpose. By utilizing the tool, we have converted several real-world firewall policies into their modularized form, and found that the process consistently improved the understanding of a policy, and the benefit is much more significant when the policy is large and/or when it has substantial usage of both permit and deny rules.

PROJECT: Privacy Incidents Database
PI: Jessica Staddon
Hard Problem: Policy-Governed Secure Collaboration

Under this project we are building the first comprehensive encyclopedia and database of privacy incidents. Most privacy incidents, such as cyber-bullying, slander, stalking, revenge porn, social media oversharing, data reidentification, and surveillance do not involve a security breach. Therefore, such incidents are not represented in current security incident databases. This lack of a centralized resource leads to widely varying measurements: for example, the Privacy Rights Clearinghouse data finds fewer than 400 data breach incidents in 2014; conversely, based on proprietary data from 70 companies and organizations, Verizon finds over 2000 breaches in the same year. Our publicly accessible database will enable the privacy technology and policy communities to reach consensus around patterns in privacy incidents.

PROJECT: Resilience Requirements, Design, and Testing
PIs: Kevin Sullivan, Mladen Vouk, Ehab Al-Shaer
Participating SubLablets: University of Virginia, University of North Carolina at Charlotte
Hard Problems: Resiliency, Security Metrics and Models

During this year, we developed models and tools for resilient system verification. Specifically, we developed a language to model the DDoS and Worms propagation attacks. To specify various attack scenarios, we defined a language that the users can use to define the attack models to be used to verify resiliency. We also defined two resiliency metrics: (1) the isolation metric to quantify the counter-measure resistance on any path (source to destination) based on the network access controls including firewalls, IPSec, IDS, proxy; and (2) the diversity metric to quantify the required attack vector by adversary based on the different disjoint attack surface due to OS and application diversity in the attack path. We then integrated both metrics to consider the optimal isolation-diversity combination for multi-stage attack scenarios.

PROJECT: Redundancy for Network Intrusion Prevention Systems (NIPS)
PI: Michael Reiter
Participating Sublablet: University of North Carolina,
Hard Problem: Resiliency

While network optimization is central to many SDN applications, few efforts have attempted to make it accessible. The goal of our research this year was a general, efficient framework for expressing and solving network optimizations in support of a wide range of network management and security applications. Our framework, SOL, achieves both generality and efficiency via a path-centric abstraction, in which network managers express policies in terms of the characteristics of the forwarding paths that should be permissible for different classes of traffic (including the network functions that must be applied to each class). We have shown that SOL can concisely express applications with diverse goals (traffic engineering, offloading, topology modification, service chaining, etc.), and SOL yields optimal or near-optimal solutions with often better performance than custom formulations. Thus, SOL can lower the barrier to entry for novel SDN network optimization applications.

PROJECT: Smart Isolation in Large-Scale Production Computing Infrastructures
PIs: Xiaohui (Helen) Gu, William Enck,
Hard Problem: Resiliency

In the past year, we have progressed our characterization of existing security isolation techniques. The characterization highlights aspects that support resiliency, and our survey of existing literature reveals relatively few instances of proactive or dynamic isolation. Based on this survey study, we started to explore primitives for proactive and dynamic security isolation techniques. We decided to start from the Docker

isolation technology, as it recently has become the de facto isolation solution in industry. We have downloaded and analyzed over 40K Docker images from DockerHub to study the vulnerability of existing Docker images. Our results show that the vulnerability issue is prevalent and an intelligent vulnerability detection and containment technique is a must. When studying existing literature, we also found that primitives such as capabilities and Information Flow Control (IFC) enable a form of dynamic isolation. That is, the isolation protection domain changes based on runtime events. Historically, IFC systems have limited practical applications. Our research invited the notion of "lazy polyinstantiation" to make IFC more practical by limiting label explosion while maintaining compatibility with legacy software.

PROJECT: Automated Synthesis of Resilient Architectures
PI: Ehab Al-Shaer
Participating Sublablet: University of North Carolina at Charlotte
Hard Problem: Resiliency

We use hypothesis testing to improve our automated synthesis of security configuration frameworks in order to determine the optimal fine-grain isolation between any two hosts in the network. Our approach enables the users of the system to incrementally and automatically create and validate null hypotheses until the configuration for statistically optimal isolation is found. We also created this year a formal model to enable resilient-by-construction development of cyber system. Our preliminary model has two components: (1) resiliency metrics based on Cyber Resilience Engineering Framework (CREF), and (2) formal verification to investigate and extend the cyber model in order to exhibit the desired resilient properties in the design phase.

PUBLICATIONS

PROJECT: Attack Surface and Defense-in-Depth Metrics

- Christopher Theisen, "Automated attack surface approximation," in *Proceedings of the 2015 10th Joint Meeting on Foundations of Software Engineering (ESEC/FSE 2015),* ACM, New York, NY, USA, pp. 1063-1065. DOI=http://dx.doi.org/10.1145/2786805.2807563

- Christopher Theisen and Laurie Williams, "Risk-Based Attack Surface Approximation," poster, *Proceedings of the Symposium and Bootcamp on the Science of Security (HoTSoS '16)*, Pittsburgh, PA, April 20-21, 2016, pp.123-126.

- Nuthan Munaiah and Andrew Meneely, "Beyond the Attack Surface: Assessing Security Risk with Random Walks on Call Graphs," in *Proceedings of the 2016 ACM Workshop on Software PROtection (SPRO '16)*, Vienna, Austria, pp. 3-14.

- Özgür Kafali, Munindar P. Singh, and Laurie Williams, "NANE: Identifying Misuse Cases Using Temporal Norm Enactments," *2016 IEEE 24th International Requirements Engineering Conference (RE)*, Beijing, 2016, pp. 136-145.

- Christopher Theisen, "Reusing Stack Traces: Automated Attack Surface Approximation," in *Proceedings of the 38th International Conference on Software Engineering Companion (ICSE, '16)*, Doctoral Symposium, Austin, TX, 2016, ACM, New York, NY, USA, pp. 859-862.

- Christopher Theisen, Laurie Williams, Kevin Oliver, and Emerson Murphy-Hill, "Software Security Education at Scale," in Proceedings of the 38th International Conference on Software Engineering Companion (ICSE '16), Austin, TX, 2016, ACM, New York, NY, USA, pp. 346-355.

- Özgür Kafali, Munindar P. Singh, and Laurie Williams, "Toward a Normative Approach for Forensicability: Extended Abstract," Proceedings of the Symposium and Bootcamp on the Science of Security (HoTSoS '16), Pittsburgh, PA, April 20-21, 2016, pp. 65-67.

PROJECT: Systematization of Knowledge from Intrusion Detection Models

- Xiaofan He, Huaiyu Dai, Peng Ning, and Rudra Dutta, "Dynamic IDS Configuration in the Presence of Intruder Type Uncertainty," in *Proceedings of*

NC STATE UNIVERSITY

the IEEE Global Conference on Communications (GLOBECOM), San Diego, CA, Dec. 2015.

- Richeng Jin, Xiaofan He, and Huaiyu Dai, "Collaborative IDS Configuration: A Two-layer Game-Theoretical Approach," in *Proceedings of the IEEE Global Conference on Communications (GLOBECOM)*, Washington, DC, 2016. URL: http://www4.ncsu.edu/~hdai/GC16-RJ.pdf

- Nuthan Munaiah, Andrew Meneely, Benjamin Short, Ryan Wilson, and Jordan Tice, "Are Intrusion Detection Studies Evaluated Consistently? A Systematic Literature Review," 2016. URL: http://scholarworks.rit.edu/article/1810

PROJECT: Vulnerability and Resilience Prediction Models

- Roopak Venkatakrishnan and Mladen A. Vouk, "Using Redundancy to Detect Security Anomalies: Towards IoT Security Attack Detectors," ACM Ubiquity, vol. 2016, no. January (2016), pp. 1-19. http://ubiquity.acm.org/article.cfm?id=2822881

- Donghoon Kim and Mladen A. Vouk, 'Securing Software Application Chains in a Cloud," 2nd International Conference on Information Science and Security (ICISS 2015), Dec 14-16, 2015, Seoul, Korea, pp. 1-4. doi:10.1109/ICISSEC.2015.7371032

PROJECT: Warning of Phishing Attacks: Supporting Human Information Processing, Identifying Phishing Deception Indicators, and Reducing Vulnerability

- Olga A. Zielinska, Allaire K.Welk, Christopher B. Mayhorn, and Emerson Murphy-Hill, "The Persuasive Phish: Examining the Social Psychological Principles Hidden in the Phishing Email Message," in Proceedings of the Symposium and Bootcamp on the Science of Security (HoTSoS '16), Pittsburgh, PA, April 20-21, 2016, pp. 126-126.

- Olga A. Zielinska, Allaire K.Welk, Christopher B. Mayhorn, and Emerson Murphy-Hill, "A Temporal Analysis of Persuasion Principles in Phishing Emails,"in Proceedings of the Human Factors and Ergonomics Society 60th Annual Meeting, Santa Monica, CA, Human Factors and Ergonomics Society, 2016. URL: http://doi.acm.org/10.1145/2898375.2898382

- Carl J. Pearson, Allaire K. Welk, and Christopher B. Mayhorn, "In Automation We Trust: Identifying Varying Levels of Trust in Human and Automated Information Sources," in Proceedings of the Human Factors and Ergonomics Society 60th Annual Meeting, Washington, DC, 2016.

- Carl J. Pearson, Allaire K. Welk, William A. Boettcher, Roger C. Mayer, Sean Streck, Joseph M. Simons-Rudolph, and Christopher B. Mayhorn, 'Differences in Trust Between Human and Automated Decision Aids,' in Proceedings of the Symposium and Bootcamp on the Science of Security (HoTSoS '16), Pittsburgh, PA, April 20-21, 2016, pp. 95-98.

- Christopher B. Mayhorn, Emerson Murphy-Hill, Olga A. Zielinska, and Allaire K. Welk, 'The Social Engineering Behind Phishing,' The Next Wave, vol. 21, no. 1, 2015, pp. 24-31.

PROJECT: A Human Information-Processing Analysis of Online Deception Detection

- Aiping Xiong, Weining Yang, Ninghui Li, and Robert W. Proctor, "Improving Detection of Phishing Attacks by Directing Users' Attention to Domain Highlighted URLs,"Paper presented at the *45th Annual Meeting of the Society for Computers in Psychology*, Chicago, IL, November 2015.

- Aiping Xiong, Weining Yang, Ninghui Li, and Robert W. Proctor, "Ineffectiveness of Domain Highlighting as a Tool to Help Users Identify Phishing Webpages," Talk presented at the *60th International Annual Meeting of Human Factors and Ergonomics Society (HFES),* Washington, DC, September 2016 (nominated for the for the HFES 2016 Marc Resnick Best Paper Competition).

- Jing Chen, Aiping Xiong, Ninghui Li, and Robert W. Proctor, "The Description-Experience Gap in the Effect of Warning Reliability on User Trust, Reliance, and Performance in a Phishing Context," Talk presented at the *7th International*

NC STATE UNIVERSITY

Conference on Applied Human Factors and Ergonomics (AHFE), Orlando, FL, July, 2016.

- Aiping Xiong, Ninghui Li , Wanling Zou, and Robert W. Proctor, "Tracking Users' Fixations When Evaluating the Validity of a Web Site," Talk presented at the *7th International Conference on Applied Human Factors and Ergonomics (AHFE)*, Orlando, FL, July, 2016.

- Jing Chen, Christopher. S. Gates, Zach Jorgensen, Weining Yang, Aiping Xiong, Ninghui Li, Ting Yu, and Robert W. Proctor, "Effective Risk Communication for End Users: A Multi-granularity Approach," poster, *Women in CyberSecurity (WiCyS)* Conference, Atlanta, GA, 2015.

PROJECT: Leveraging the Effects of Cognitive Function on Input Device Analytics to Improve Security

- Ignacio X. Domínguez, Prairie Rose Goodwin, David L. Roberts, and Robert St. Amant, "Human Subtlety Proofs: Using Computer Games to Model Cognitive Processes for Cybersecurity," *International Journal Of Human-Computer Interaction*, Online, 2016.

- Ignacio X. Domínguez, Jayant Dhawan, Robert St. Amant, and David L. Roberts, "Exploring the Effects of Different Text Stimuli on Typing Behavior," in *Proceedings of the 14th International Conference on Cognitive Modeling (ICCM 2016)*, University Park, PA, USA, 2016, pp. 175-181.

- Ignacio X. Domínguez, Jayant Dhawan, Robert St. Amant, and David L. Roberts, "JIVUI: JavaScript Interface for Visualization of User Interaction," in *Proceedings of the 14th International Conference on Cognitive Modeling (ICCM 2016)*, University Park, PA, USA, 2016, pp. 125-130.

- Robert St. Amant and David L. Roberts, "Natural Interaction for Bot Detection," *IEEE Internet Computing* vol. 20, no. 4, July-Aug. 2016, pp. 69-73.

PROJECT: Understanding Effects of Norms and Policies on the Robustness, Liveness, and Resilience of Systems

- Luis G. Nardin, Tina Balke-Visser, Nirav Ajmeri, Anup K. Kalia, Jaime S. Sichman, and Munindar P. Singh, "Classifying Sanctions and Designing a Conceptual Sanctioning Process for Socio-Technical Systems," *The Knowledge Engineering Review*, vol. 31, no.2 March 2016.

PROJECT: Formal Specification and Analysis of Security-Critical Norms and Policies

- Nirav Ajmeri, Jiaming Jiang, Rada Chirkova, Jon Doyle, and Munindar P. Singh, "Coco: Runtime Reasoning about Conflicting Commitments," in *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, New York, NY, July 2016, pp. 17-23.

- Jiaming Jiang, Nirav Ajmeri, Rada Chirkova, Jon Doyle, and Munindar P. Singh, "Expressing and Reasoning about Conflicting Norms in Cybersecurity," poster, *Proceedings of the Symposium and Bootcamp on the Science of Security (HoTSoS '16)*, Pittsburgh, Pennsylvania, April 19-21, 2016.

- Özgür Kafali, Nirav Ajmeri, and Munindar P. Singh, "Formal Understanding of Tradeoffs among Liveness and Safety Requirements," *Proceedings of the 3rd International Workshop on Artificial Intelligence for Requirements Engineering (AIRE)*, 2016, pp. 17-18.

- Özgür Kafali, Nirav Ajmeri, and Munindar P. Singh, "Normative Requirements in Sociotechnical Systems," in *Proceedings of the 9th International Workshop on Requirements Engineering and Law (RELAW)*, 2016, pp. 259-260.

- Mehdi Mashayekhi, Hongying Du, George F. List, and Munindar P. Singh, "Silk: A Simulation Study of Regulating Open Normative Multiagent Systems," in *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI-16)*, New York, New York, USA, July 2016, pp. 373-379.

- Amit K. Chopra and Munindar P. Singh, "Custard: Computing Norm States over Information Stores," in *Proceedings of the 15th International Conference on Autonomous Agents and*

NC STATE UNIVERSITY

MultiAgent Systems (AAMAS), Singapore, May 2016, pp. 1096-1105.

- Amit K. Chopra and Munindar P. Singh, "From Social Machines to Social Protocols: Software Engineering Foundations for Sociotechnical Systems," in *Proceedings of the 25th International World Wide Web Conference*, Montreal, Canada, April 2016, pp. 903-914. doi: 10.1145/2872427.2883018

- Pradeep K. Murukannaiah, Nirav Ajmeri, and Munindar P. Singh, "Engineering Privacy in Social Applications," IEEE Internet Computing, vol. 20, no. 2 March - April 2016, pp. 72-76.

- Anup K. Kalia, Zhe Zhang, and Munindar P. Singh, "Güven: Estimating Trust from Communications," *Journal of Trust Management*, vol. 3, no. 1, January 2016.

PROJECT: Scientific Understanding of Policy Complexity

- Haining Chen, Omar Chowdhury, Ninghui Li, Warut Khern-am-nuai, Suresh Chari, Ian Molloy, and Youngja Park, "Tri-Modularization of Firewall Policies," in *Proceedings of the 21st ACM Symposium on Access Control Models and Technologies (SACMAT)*, Shanghai, China, June 2016, pp. 37-48.

PROJECT: Privacy Incidents Database

- Pradeep Murukannaiah, Jessica Staddon, Heather R. Lipford, and Bart P. Knijnenburg, "PrIncipedia: A Privacy Incidents Encyclopedia," Privacy Law Scholars Conference, 2016
- URL: https://drive.google.com/file/d/0B-is7Sqpwv0bZy1UTEQ2UnZkVUE/view

PROJECT: Resilience Requirements, Design, and Testing

- Mohammad Ashiqur Rahman, Abdullah Al Faroq, Amarjit Datta, and Ehab Al-Shaer, "Automated Synthesis of Resiliency Configurations for Cyber Networks," **IEEE Conference on Communications and Network Security (CNS)**, Philadelphia, Pennsylvania, USA, October 2016.

- Mohammad Ashiqur Rahman, A. H. M. Jakaria, and Ehab Al-Shaer, "Formal Analysis for Dependable Supervisory Control and Data Acquisition in Smart Grids," in *Proceedings of the 46th IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, Toulouse, France, June 2016, pp. 263-274.

- Yasir Imtiaz Khan, Ehab Al-Shaer, and Usman Rauf, "Cyber Resilience-by-Construction: Modeling, Measuring & Verifying," in *Proceedings of the 2015 Workshop on Automated Decision Making for Active Cyber Defense (SafeConfig '15)*, ACM, New York, NY, USA, 2015, pp. 9-14.

- Muhammad Qasim Ali, Ayesha Binte Ashfaq, Ehab Al-Shaer, and Qi Duan, "Towards a Science of Anomaly Detection System Evasion," *IEEE Conference on Communications and Network Security (CNS)*, Florence, Italy, 2015, pp. 460-468.

- Jafar Haadi Jafarian, Ehab Al-Shaer, and Qi Duan, "Adversary-aware IP address randomization for proactive agility against sophisticated attackers," *2015 IEEE Conference on Computer Communications (INFOCOM)*, Kowloon, 2015, pp. 738-746.

PROJECT: Redundancy for Network Intrusion Prevention Systems (NIPS)

- Victor Heorhiadi, Shriram Rajagopalan, Hani Jamjoom, Michael K. Reiter, and Vyas Sekar, "Gremlin: Systematic Resilience Testing of Microservices," *2016 36th IEEE International Conference on Distributed Computing Systems*, Nara, Japan, 2016, pp. 57-66.

- Victor Heorhiadi, Michael K. Reiter, and Vyas Sekar, "Simplifying Software-defined Network Optimization Using SOL," in *Proceedings of the 13th USENIX Symposium on Networked System Design and Implementation (NSDI'16)*, Santa Clara, CA, March 2016, Pages 223-237.

- Victor Heorhiadi, Michael K. Reiter, and Vyas Sekar, "Accelerating the Development of Software-Defined Network Optimization Applications Using SOL," Technical Report, arXiv preprint, 2015.

NC STATE UNIVERSITY

PROJECT: Smart Isolation in Large-Scale Production Computing Infrastructures

- Rui Shu, Peipei Wang, Sigmund A. Gorski, III, Benjamin Andow, Adwait Nadkarni, Luke Deshotels, Jason Gionta, William Enck, and Xiaohui Gu, "A Study of Security Isolation Techniques," *ACM Computing Surveys (CSUR)*, vol. 49, no. 3, December 2016.

- Adwait Nadkarni, Benjamin Andow, William Enck, and Somesh Jha, "Practical DIFC Enforcement on Android," in *Proceedings of the 25th USENIX Security Symposium*, Austin, TX, 2016, pp. 1119-1136.

PROJECT: Automated Synthesis of Resilient Architectures

- Mohammad Ashiqur Rahman, Abdullah Al Faroq, Amarjit Datta, and Ehab Al-Shaer, "Automated Synthesis of Resiliency Configurations for Cyber Networks," *IEEE Conference on Communications and Network Security (CNS)*, Philadelphia, Pennsylvania, USA, October 2016.

- Mohammad Ashiqur Rahman and Ehab Al-Shaer, "Automated Synthesis of Distributed Network Access Controls: A Formal Framework with Refinement," *IEEE Transactions on Parallel and Distributed Computing (TPDC)*, June 2016.

- Yasir Khan and Ehab Al-Shaer, "Property-Based Verification of Evolving Petri Nets," in *Proceedings of the Tenth International Conference of Software Engineering Advances (ICSEA 2015)*, Barcelona, Spain, 2015, pp. 301-306.

- Mohammad Ashiqur Rahman and Ehab Al-Shaer, "Formal Synthesis of Dependable Configurations for Advanced Metering Infrastructures," in 2015 *IEEE International Conference on Smart Grid Communications (SmartGridComm)*, 2015, pp. 289-294.

Research/Evaluation Team

- Jeffrey C. Carver, Morgan Burcham, Sedef Akinli Kocak, Ayse Bener, Michael Felderer, Matthias Gander, Jason King, Jouni Markkula, Markku Oivo, Clemens Sauerwein, and Laurie Williams, "Establishing a baseline for measuring advancement in the science of security: an analysis of the 2015 IEEE security & privacy proceedings," in *Proceedings of the Symposium and Bootcamp on the Science of Security (HoTSoS '16)*, Pittsburgh, PA, April 20-21, 2016, pp. 38-51. DOI: http://dx.doi.org/10.1145/2898375.2898380

EDUCATIONAL

- We further enhanced our research guidelines, creating a version for (1) empirical evaluation of real-world data; (2) analytical studies that use mathematical proofs; and (3) build-then-evaluate studies of security solutions. The research teams are guided in their plans through these guidelines. Students learn to critique others' work through the use of the guidelines.

- We developed a new rubric and web surveys to assist lablet participants in organizing their comments to reviews of exploratory ideas, research plans, and manuscripts prior to publication, which can lead to more focused comments and discussions before a research effort is considered complete.

- Our Summer Workshop included a lecture on the pitfalls of statistical analysis as well as a group exercise on how to promote scientific practices and writing, and how to evaluate the scientific clarity of published papers in cybersecurity.

- We have identified a seed list of publication venues where Science of Security research appears. We have been engaging the community (at other Lablets) on refining and ranking a list of venues.

NC STATE UNIVERSITY

- We continued our weekly seminar series in the Fall 2015 and Spring 2016 semesters with supported students and Principal Investigators (PIs). Students presented their research plans and publications to obtain feedback on their work. We have developed a research proposal outline to help researchers organize their thoughts and ensure they are conducting their research in a scientifically defensible manner. This outline is also the foundation for our feedback and evaluation instruments used during the seminars.

Project-Specific Educational and Curriculum Outcomes

Project: Attack Surface and Defense-in-Depth Metrics

One side benefit of this work has been the collection and aggregation of vulnerability data across multiple, large open source case studies (Wireshark and FFmpeg). This data is now being made available to Rochester Institute of Technology (RIT) students who study it in the classroom and discuss how the vulnerabilities could have been prevented. Understanding vulnerability history is a crucial step in developing the attacker mindset that students need when they enter the workforce.

Project: Systematization of Knowledge from Intrusion Detection Models

This work serves as a valuable case study to graduate students who are looking to provide new research in the IDS field. By seeing how a lack of evaluation can lead to poor systematization, graduate students have learned how to conduct empirical evaluations with proper evaluation.

Project: Vulnerability and Resilience Prediction Models

Some of the techniques examined in this research, and lessons learned, in particular pro-active run-time resilience to attacks based on redundancy, were added as topics into NC State's graduate course on cloud computing technologies.

NC STATE UNIVERSITY

Project: Warning of Phishing Attacks: Supporting Human Information Processing, Identifying Phishing Deception Indicators, and Reducing Vulnerability

We have noticed an increased emphasis on the social sciences within security research. Not too many years ago, there were few researchers interested in why humans behave in a certain manner when faced with security-related decisions. At recent conferences such as HoTSoS and Human Factors, research that investigates the human role in cybersecurity has been actively promoted. While we have not seen any specific changes to curricula to date, these changes are certainly on the horizon.

Projects: A Human Information-Processing

Analysis of Online Deception Detection; Scientific Understanding of Policy Complexity
We held a weekly research seminar involving graduate and undergraduate students and faculty members from Computer Science, Psychological Sciences, and Industrial Engineering. In this seminar, we discussed various security-related projects on which we were working, and the students got to play a major role in the design and conduct of the research.

Graduate student members of the research team prepared papers and posters presented at various conferences. Two members presented talks to the Cognitive Area colloquium in the Department of Psychological Sciences aimed at educating faculty members and students of ways in which basic psychological knowledge can be applied to research on security and privacy.

Dr. Proctor presented a talk, entitled "Cybersecurity: A human problem," at the Cybersecurity Awareness Day, University of Wisconsin, on October 22, 2015, explaining how human factors contribute to the science of cybersecurity.

Project: Understanding Effects of Norms and Policies on the Robustness, Liveness, and Resilience of Systems

Discussions of experimental and data-analytic results now introduce and emphasize the need for rigor and the dangers of simple statistical recipes, using examples patterned after ones given in works by Ioannidis and by Gigerenzer.

Project: Formal Specification and Analysis of Security-Critical Norms and Policies

NCSU CSC 503, Computational Applied Logic includes formal logical treatments of program and system specifications, obligations and norms, and knowledge and privacy, using security concerns to illustrate these ideas.

A new NCSU Computer Science senior-level undergraduate course on computability being developed introduces hybrid automata and systems along with ideas about verifying safety properties of hybrid computer/physical/social systems.

Project: Resilience Requirements, Design, and Testing

Many of the materials developed via this project were used in graduate level classes to measure and verify resiliency techniques and systems. For example, in ITIS 6230, which is Cyber Risk Determination and Mitigation, we study Industry Control System resiliency which is material developed in this project.

Project: Redundancy for Network Intrusion Prevention Systems (NIPS)

This project is unusual in that it identifies connections between research in different domains of network security and management, and then it distills from them primitives and tools that can be used to express and solve challenges across these multiple domains.
This has not resulted in changes to curriculum, per se, as of yet, but it has certainly caused us to think more deeply and rigorously about the fundamental nature of these different challenges.

Project: Smart Isolation in Large-Scale Production Computing Infrastructures

In Spring 2016, PI Enck taught CSC 574, Introduction to Computer and Network Security. As part of this graduate level class, students are required to conduct a novel research project. New for this semester, students were required to develop a research plan that described the threat model, methodology, and evaluation techniques for the project. This milestone was motivated by the SoS Lablet efforts at NCSU.

COMMUNITY ENGAGEMENTS

Over the last year, NCSU researchers have established collaborations with several industry and Government organizations, including IBM, Cisco, and the National Institute of Standards and Technology (NIST).

On February 2-3, 2016, the NCSU Lablet hosted the Science of Security Quarterly Meeting. A highlight of the meeting was a talk by Dr. Henry Petroski, the Aleksandar S. Vesic Professor of Civil Engineering at Duke University, on the paradoxical relationship between success and failure in design.

NCSU conducted a 2016 Summer Workshop on the theme “Translating Science of Security Research to Industry,” which focused on promoting impact, especially through technology transfer. We included a session on strategies to achieve technology transfer, a keynote by a Health and Human Services (HHS) researcher, reviews of previous attempts to measure impact in computer science, and a lively industry panel on the particular challenges of technology transfer in computer science.

We developed interview guides and surveys to assess the impact of lablet participation on various lablet stakeholder groups, which can provide greater visibility into potential benefits of our systematic approach to the science of security.

We refined our rubric to use for characterizing the main elements of science as reflected in published research in cybersecurity. Such a rubric may lead to improved classification of existing research and potentially guidelines for future publications that bring out their scientific contributions.

Our evaluation of papers appearing in ACM CCS 2015 showed that several papers omitted details of their study designs, including objectives, rationale for selection of cases, and threats to validity.

NC STATE UNIVERSITY

UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN

PI DAVID NICOL

In 2016, the University of Illinois at Urbana-Champaign (UIUC) Science of Security (SoS) Lablet, led by researchers David Nicol and William Sanders, contributed broadly to the development of security science while leveraging Illinois expertise in resiliency, which in this context means a system's demonstrable ability to maintain security properties even during ongoing cyber attacks. Collaborating on several projects with subLablets, including the Illinois Institute of Technology, Newcastle University, University of Southern California, University of Pennsylvania, Dartmouth College, and Rice University, the Lablet's work draws on several fundamental areas of computing research. Some ideas from fault-tolerant computing can be adapted to the context of security. Strategies from control theory are being extended to account for the high variation and uncertainty that may be present in systems when they are under attack. Game theory and decision theory principles are being used to explore the interplay between attack and defense. Formal methods are being applied to develop formal notions of resiliency. End-to-end system analysis is being employed to investigate resiliency of large systems against cyber attacks. The Lablet's work also draws upon ideas from other areas of mathematics and engineering as well.

The Science of Security has many attributes that range from use and development of scientific techniques in experimental security work, to modeling and mathematical foundations of systems where security and security properties are the object of the reasoning. UIUC contributes principally to the latter category with research that also supports the former category. We study how security properties are shaped by policy at different layers of the network stack, and use that to help define hypotheses that might be empirically tested. We are defining models of cyber-physical systems that allow us to analyze how closely the system is allowed to skirt disaster, a measure of the system's resilience to disturbance. We are developing mathematical models of systems under attack, the attackers, and the defenders, to better understand how well the system is able to maintain required service levels through the attack, and to aid defensive decision-makers. We are applying sophisticated stochastic modeling techniques to describe vast volumes of data within which there are attacks; the models describe correlations between observations that might suggest attacks, and unobservable state that describes the attack. Finally, we are developing models of human behavior that seek to explain the how and why of humans circumventing security mechanisms. In short, the UIUC Science of Security research is exploring foundational mathematical modeling formalisms that quantitatively describe security attributes, and seek to predict those attributes as a function of context and environment.

ILLINOIS
UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN

FUNDAMENTAL RESEARCH

Project: A Hypothesis Testing Framework for Network Security
PI: P. Brighten Godfrey
Participating Sublablet: Illinois Institute of Technology
Hard Problems: Security Metrics and Models, Scalability and Composability

This project is developing the analysis methodology needed to support scientific reasoning about the security of networks, with a particular focus on information and data flow security. The core of this vision is Network Hypothesis Testing Methodology (NetHTM), a set of techniques for performing and integrating security analyses applied at different network layers, in different ways, to pose and rigorously answer quantitative hypotheses about the end-to-end security of a network. To realize NetHTM, we need the effective evaluation methodologies designed scale to large and complex systems. A key challenge is that the real-world network configuration constrains the level of model abstraction. We made advances in scalable evaluation methodology and platform via the marriage of emulation (within which real network and application protocol stacks are directly executed) and simulation. We used the hybrid platform to realize the network models and the verification algorithms we developed in year one, and also studied the impact of various cyber-attacks on network behavior. In addition, we also made progress on developing new metrics and models, such as congestion-freedom property, for modeling and enforcing correct behavior in dynamic networks. We used our model and verification algorithms to develop dynamic control algorithms to preserve those specified properties across time. Several manuscripts describing our work in those two topics have been submitted, including five papers that were published in the leading journal and conference in modeling and simulation in year two. We received a best paper award, a best poster award, and best paper candidate recognition in the 2015 ACM SIGSIM Conference on Principles of Advanced Discrete Simulation (PADS). In addition, an undergraduate student, Mr. Adnan Haider (the second author of the PADS '15 best paper) was named finalist for prestigious CRA Outstanding Undergraduate Researcher Award.

Project: Data-Driven-Model-Based Decision-Making
PIs: William Sanders, Masooda Bashir, David Nicol, and Aad Van Moorsel
Participating Sublablet: Newcastle University, UK
Hard Problems: Security Metrics and Models, Human Behavior

Predictive security metrics on complex systems require an analysis of every facet of a system. Typically, this involves studying, among other things, the security policies, attack paths, and autonomous system behavior through rigorous mathematical models that simulate the complete system. An often overlooked, but critical, component in a model such as this is the human element. We have developed and refined the HITOP modeling formalism to create models of human behaviors and decision-making. By considering each user's opportunity, willingness, and capability to perform individual tasks throughout their daily routine, we can examine the pivotal role that good human users play in secure systems. A sound model requires valid model parameters to gather useful results. We have also developed an observed data collection strategy that ensures that useful metrics can be obtained.

Project: Data Driven Security Models and Analysis
PI: Ravi Iyer
Hard Problems: Security Metrics and Models, Resiliency, Human Behavior

We have continued work on development of scientifically sound data-driven methods and tools for early attack detection. This year we built, and integrated with our security testbed, a framework for generation, replay (in an isolated environment), and analysis of real-world attack variants, i.e., attacks that achieve the same objective as a known attack while bypassing the existing detection mechanisms. We assessed the framework on three real-world attacks for which we generated a total of 648 unique attack variants. Using the generated attack variants we evaluated the detection efficiency of: (1) signature-based detection, using a file hash of known malicious files; (2) anomaly-based detection, using high-frequency events observed in past attacks as an indicator of future attacks; and (3) detection based using AttackTagger, a factor graph based approach for preemptive detection of attacks

ILLINOIS
UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN

we developed. The results show that AttackTagger detects more than half of the attack variants (up to 75%), whereas the signature-based approach detects 25%, and the frequency-based approach detects up to 33%. Additionally, we employed a multi-stage security game with learning to model multi-stage attacks. An attack model is derived from the study on security incident data (to demonstrate the value of reinforcement learning) and accommodates the limited observation of the defender on the attack in progress. A simulation-based experiment demonstrated that despite the relaxed assumptions and restrictions on formulating the game, Naive Q-Learning shows performance comparable to Markov Game when played against players with relatively weak rationality (non-Markov).

Project: Science of Human Circumvention of Security
PI: Tao Xie
Participating SubLablets: University of Southern California, University of Pennsylvania, Dartmouth College
Hard Problems: Human Behavior, Scalability and Composability, Policy-Governed Secure Collaboration, Security Metrics and Models

We continue to study people's approaches to cyber security, and their use of authentication methods for accessing websites, their organization's databases, and the Internet. We focus especially on passwords as a prime method in the context of this trust (or suspicion or distrust). Use of passwords, adherence to password guidelines, and circumvention of password rules (e.g., sharing, writing them down on available files) are also excellent reflections of people's understanding, misunderstandings, and beliefs about personal and organizational efforts to protect individual and enterprise-level information. In addition, we are building and testing DASH agent models and beginning to test a mechanical Turk experiment/simulation to further examine users' use of passwords, workarounds, cyber trust, and strategies–measurements from the Turk experiment provide base calibration for the DASH model. We have developed a new version of DASH in python that improves ease of development. We are also working with researchers at the University of Pennsylvania who have developed methods to learn agent behavior from observational data. To date, our results include constructing a semiotic framework for circumvention, validating our basic DASH model by reproducing behavior found in ground-truth human surveys, and duplication in our simulation of a version of "uncanny descent." in which making constraints on passwords more complex can decrease overall security. Last, we continue to administer two surveys: one on users' understanding of cybersecurity processes and their modes of circumvention; and one on security administrators' understanding of cybersecurity processes and their rationales for security policies and decisions. Also, to study people's trust in cyber security, especially mobile app security, we focus on collecting and analyzing UI text information faced by mobile app users to enable them to make informed decisions on mobile app security.

Project: Static-Dynamic Analysis of Security Metrics for Cyber-Physical Systems
PI: Sayan Mitra
Participating Sublablet: Rice University
Hard Problems: Scalability and Composability, Security Metrics and Models

In this collaborative project, we have formulated the general problem of controller synthesis in the presence of resource constrained adversaries; namely, given an adversary of a certain class, parametrized according to the quantifiable resources available to them, we are creating a methodology to assess the worst-case potential impact and performance degradation of a control system from a threat of this class. We have developed a sound and complete algorithm for solving this problem, for the special case of control systems with linear and monotonic dynamics and adversary resources characterized by their signal energy. The approach used to develop the algorithms brings together ideas from robust control and recent developments in syntax-guided program synthesis. Using our algorithms, we are able to synthesize controllers that are provably resilient to certain threat classes; in addition, we are also able to characterize the states of the systems in terms of their vulnerability levels. We have also continued our work on characterizing the trade-off between privacy and performance in cyber-physical systems, particularly in cases where strategic preferences that govern dynamics are to be protected.

Project: Anonymous Messaging
PI: Pramod Viswanath
Hard Problem: Scalability and Composability

Anonymity is a basic right and a core aspect of the Internet. Recently, there has been tremendous interest in anonymity and privacy in social networks, motivated by the natural desire to share one's opinions without the fear of judgment or personal reprisal (by parents, authorities, and the public). We propose to study the fundamental questions associated with building such a semi-distributed, anonymous messaging platform, which aims to keep anonymous the identity of the source who initially posted a message as well as the identity of the relays who approved and propagated the message.

PUBLICATIONS

PROJECT: A Hypothesis Testing Framework for Network Security

- Ning Liu, Xian-He Sun, and Dong Jin, "On Massively Parallel Simulation of Large-Scale Fat-Tree Networks for HPC Systems and Data Centers," in *Proceedings of the 3rd ACM SIGSIM Conference on Principles of Advanced Discrete Simulation*, London, UK, 2015. Best Poster Award.

- Ning Liu, Adnan Haider, Xian-He Sun, and Dong Jin, "FatTreeSim: Modeling a Large-scale Fat-Tree Network for HPC Systems and Data Centers Using Parallel and Discrete Even Simulation," in *Proceedings of the 3rd ACM SIGSIM Conference on Principles of Advanced Discrete Simulation*, London, UK, 2015, pp. 199-210. Best Paper Award.

- Jiaqi Yan and Dong Jin, "A Virtual Time System for Linux-container-based Emulation of Software-defined Networks," in *Proceedings of the 3rd ACM SIGSIM Conference on Principles of Advanced Discrete Simulation*, London, UK, 2015, pp. 235-246. Finalist for Best Paper Award.

- Jiaqi Yan and Dong Jin, "VT-Mininet: Virtual-time-enabled Mininet for Scalable and Accurate Software-Define Network Emulation," in *Proceedings of the 1st ACM SIGCOMM Symposium on Software Defined Networking Research (SOSR 2015)*, Santa Clara, CA, 2015.

- Dong Jin and David Nicol, "Parallel Simulation and Virtual-Machine-Based Emulation of Software-Defined Networks," ACM Transactions on Modeling and Computer Simulation, vol. 26, no. 1, December 2015.

- Anduo Wang, Xueyuan Mei, Jason Croft, Matthew Caesar, and P. Brighten Godfrey, "Ravel: A Database-Defined Network," in *Proceedings of the ACM SIGCOMM Symposium on Software Defined Networking Research (SOSR 2016)*, Santa Clara, CA, 2016.

- Christopher Hannon, Jiaqi Yan, and Dong Jin, "DSSnet: A Smart Grid Modeling Platform Combining Electrical Power Distribution System Simulation and Software Defined Networking Emulation," in *Proceedings of the 2016 Annual ACM SIGSIM Conference on Principles of Advanced Discrete Simulation (SIGSIM-PADS '16)*, Banff, Canada, 2016, pp. 131-142, Best Paper Nomination.

- Dong Jin, Zhiyi Li, Christopher Hannon, Chen Chen, Jianhui Wang, and Mohammad Shahidehpour, "Towards a Resilient and Secure Microgrid Using Software-Defined Networking," *IEEE Transactions on Smart Grid, Special Issue: Special Section on Smart Grid Cyber-Physical Security*. Short paper accepted on April 2016. Invited for full paper submission.

- Christopher Hannon, Jiaqi Yan, and Dong Jin, "DSSnet: A Smart Grid Modeling Platform Combining Electrical Power Distribution System Simulation and Software Defined Networking Emulation," 2016, Best Poster Award.

- Dong Jin, Jiaqi Yan, Xin Lin, Christopher Hannon, Hui Lin, Zbigniew Kalbarczyk, Ravishankar K. Iyer, Chen Chen, Jianhui Wang, and Cheol Won Lee, "Towards a Secure and Resilient Industrial Control System with Software-defined Networking," *Workshop on Science of Security through Software-Defined Networking, June 2016. Best Poster Award.*

- Jiaqi Yan and Dong Jin, "A Lightweight Container-based Virtual Time System for Software-defined Network Emulation," *Journal of Simulation*, 2016, pp. 1-14. https://github.com/littlepretty/VirtualTimeKernel

ILLINOIS
UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN

- Ning Liu, Adnan Haider, Dong Jin, and Xian-He Sun, "A Modeling and Simulation of Extreme-Scale Fat-Tree Networks for HPC Systems and Data Centers," *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, 2016, To Appear.

- Xin Liu and Dong Jin, "ConVenus: Congestion Verification of Network Updates in Software-defined Networks," *Winter Simulation Conference (WSC)*, Washington, DC, December 2016. https://bitbucket.org/ksj0609/convenus/src

PROJECT: Data-Driven Model-Based Decision-Making

- John C. Mace, Charles Morisset, and Aad van Moorsel, "Impact of Policy Design on Workflow Resiliency Computation Time," in *Proceedings of the 12th International Conference on Quantitative Evaluation of Systems (QEST 2015)*, vol. 9259, Madrid, Spain, 2015, pp. 244-259.

- John C. Mace, Charles Moisset, and Aad van Moorsel, "Resiliency Variance in Workflows with Choice," *Proceedings of the 7th International Workshop on Software Engineering for Resilient Systems (SERENE 2015)*, vol. 9274, Paris, France, 2015, pp. 128-143.

- Ken Keefe and William H. Sanders, "Reliability Analysis with Dynamic Reliability Block Diagrams in the Möbius Modeling Tool," *Proceedings of the 9th EAI International Conference on Performance Evaluation Methodologies and Tools (VALUETOOLS 2015)*, Berlin, Germany, 2015, pp. 164-170.

PROJECT: Data Driven Security Models and Analysis

- Key-whan Chung, Charles A. Kamhoua, Kevin A. Kwiat, Zbigniew T. Kalbarczyk, and Ravishankar K. Iyer, "Game Theory with Learning for Cyber Security Monitoring,' in *Proceedings of the 2016 IEEE 17th International Symposium on High Assurance Systems Engineering Symposium (HASE 2016)*, Orlando, FL, January 7-9, 2016, pp. 1-8.

- Phuong Cao, Eric C. Badger, Zbigniew Kalbarczyk, and Ravishankar K. Iyer, "A Framework for Generation, Replay, and Analysis of Real-World Attack Variants," in *Proceedings of the Symposium and Bootcamp on the Science of Security (HoTSoS '16)*, Pittsburgh, PA, April 20-21, 2016, pp. 28-37.

- Hui Lin, Homa Alemzadeh, Daniel Chen, Zbigniew Kalbarczyk, and Ravishankar K. Iyer, "Safety-critical Cyber-physical Attacks: Analysis, Detection, and Mitigation," in Proceedings of the Symposium and Bootcamp on the Science of Security (HoTSoS '16), Pittsburgh, PA, April 20-21, 2016, pp. 82-89.

- Zachary J. Estrada, Cuong Pham, Fei Deng, Lok Yan, Zbigniew Kalbarczyk, and Ravishankar K. Iyer, "Dynamic VM Dependability Monitoring Using Hypervisor Probes," in *Proceedings of the 11th European Dependable Computing Conference-Dependability in Practice (EDCC 2015)*, Paris, France, September 7-11, 2015, pp. 61-72.

PROJECT: Science of Human Circumvention of Security

- Benjamin Andow, Adwait Nadkarni, Blake Bassett, William Enck, and Tao Xie, "A Study of Grayware on Google Play," *Workshop on Mobile Security Technologies (MoST 2016), held in conjunction with 2016 IEEE Symposium on Security and Privacy Workshops (SPW 2016)*, San Jose, CA, May 26, 2016.

- Sihan Li, Xusheng Xiao, Blake Bassett, Tao Xie, and Nikolai Tillmann, "Measuring Code Behavioral Similarity for Programming and Software Engineering Education," in *Proceedings of the 38th International Conference on Software Engineering Companion (ICSE '16)*, Software Engineering Education and Training track, Austin, TX, May 14-22, 2016, pp. 501-510.

- Harold Thimbleby and Ross Koppel, "The Healthtech Declaration," in *IEEE Security and Privacy*, vol. 13, no. 6, November/December 2015, pp. 82-84.

- Xusheng Xiao, Nikolai Tillmann, Manuel Fahndrich, Johnathan de Halleux, Michal Moskal, and Tao Xie, "User-Aware Privacy Control via Extended Static-Information-Flow Analysis,"

Automated Software Engineering Journal, vol. 22, no. 3, September 2015, pp. 333-366.

- Huoran Li, Xuan Lu, Xuanzhe Liu, Tao Xie, Kaigui Bian, Felix Xiaozhu Lin, Qiaozhu Mei, and Feng Feng, "Characterizing Smartphone Usage Patterns from Millions of Android Users," in *Proceedings of the 2015 Internet Measurement Conference (IMC 2015)*, Tokyo, Japan, October 28-30, 2015, pp. 459-472.

- Bruno Korbar, Jim Blythe, Ross Koppel, Vijay Kothari, and Sean W. Smith. "Validating an Agent-Based Model of Human Password Behavior," *The AAAI-16 Workshop on Artificial Intelligence for Cyber Security (AICS)*, Phoenix, AZ, February 2016.

- Ross Koppel, Jim Blythe, Vijay Kothari, and Sean W. Smith, "Beliefs about Cybersecurity Rules and Passwords: A Comparison of Two Survey Samples of Cybersecurity Professionals Versus. Regular Users," *12th Symposium on Usable Privacy and Security (SOUPS 2016) Security Fatigue Workshop*, Denver, CO, June 22-24, 2016.

- Xia Zeng, Dengfeng Li, Wujie Zheng, Fan Xia, Yuetang Deng, Wing Lam, Wei Yang, and Tao Xie. "Automated Test Input Generation for Android: Are We Really There Yet in an Industrial Case?" in *Proceedings of the 24th ACM SIGSOFT Symposium on the Foundations of Software Engineering (FSE 2016)*, Industrial Track, Seattle, WA, November 2016, pp. 987-992.

- Pierre McCauley, Brandon Nsiah-Ababio, Joshua Reed, Faramola Isiaka, and Tao Xie, "Preliminary Analysis of Code Hunt Data Set from a Contest," in *Proceedings of the 2nd International Code Hunt Workshop on Educational Software Engineering (CHESE 2016)*, Seattle, WA, November 2016, pp. 7-8.

PROJECT: Static-Dynamic Analysis of Security Metrics for Cyber-Physical Systems

- Zhenqi Huang, Yu Wang, Sayan Mitra, Geir E. Dullerud, and Swarat Chaudhuri, "Controller Synthesis with Inductive Proofs for Piecewise Linear Systems: an SMT-based Algorithm," in *Proceedings of the IEEE International Conference on Decision and Control (CDC 2015)*, Osaka, Japan, December 15-18, 2015, pp. 7434-7439.

- Zhenqi Huang, Yu Wang, Sayan Mitra, and Geir E. Dullerud, "Analyzing Cost of Securing Control Systems," The Next Wave, The National Security Agency's Review of Emerging Technologies Journal, vol. 21, no. 1, 2015, pp. 12-17.

- W. P. M. H. Maurice Heemels, Geir E. Dullerud, and Andrew R. Teel, "L2-Gain Analysis for a Class of Hybrid Systems with Applications to Reset and Event-Triggered Control: A Lifting Approach," in *IEEE Transactions on Automatic Control*, vol. 61, no. 10, October 2016, pp. 2766-2781.

- W. P. M. H. Maurice Heemels, Geir E. Dullerud, and Andrew Teel, "A Lifting Approach to L2-gain Analysis of Periodic Event-triggered and Switching Sampled-data Control Systems,' in *Proceedings of the 2015 IEEE Conference on Decision and Control (CDC 2015)*, Osaka, Japan, December 15-18, 2015, pp. 4176-4182.

- Zhenqi Huang, Yu Wang, Sayan Mitra, and Geir E. Dullerud, "Controller Synthesis for Linear Time-varying Systems with Adversaries," *Proceedings of the Symposium and Bootcamp on the Science of Security (HoTSoS '16)*, Pittsburgh, PA, April 20-21, 2016, pp. 53-62.

- Zhenqi Huang, Chuchu Fan, and Sayan Mitra, "Bounded Invariant Verification for Time-Delayed Nonlinear Networked Dynamical Systems," Nonlinear Analysis: Hybrid Systems, vol. 23, February 2017, pp. 211-219. http://authors.elsevier.com/sd/article/S1751570X16300279

PROJECT: Anonymous Messaging

- Giulia Fanti, Peter Kairouz, Sewoong Oh, Kannan Ramchandran, and Pramod Viswanath, "Metadata-Conscious Anonymous Messaging," in *Proceedings of the International Conference on Machine Learning -Volume 48 (ICML 2016)*, New York, NY, June 19-24, 2016, pp. 108-116.

- Giulia Fanti, Peter Kairouz, Sewoong Oh, Kannan Ramchandran, and Pramod Viswanath, "Rumor Source Obfuscation on Irregular Trees,"

ILLINOIS
UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN

Proceedings of the 2016 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Science, Antibes Juan-les-Pins, France, June 14-18, 2016, pp. 153-164.

EDUCATIONAL

P. Brighten Godfrey co-taught a Coursera online course on Cloud Networking this past fall, and the next offering was to relaunch in a new continuous operation mode on October 24, 2016. When this was taught last fall, roughly 30,000 students enrolled. The course included a segment on network security for the cloud, particularly with respect to network virtualization.

A set of notes summarizing the Bit Coin networking protocols is being developed, with the goal of using them in an upcoming privacy and anonymity course at the graduate level.

Five students completed research projects as part of the Science of Security Summer Internship Program that ended on July 29. Each student presented a poster on the last day of the internship. The students came from Tennessee State University, North Texas University, and the University of Illinois at Urbana-Champaign (UIUC). They also attended seminars on educational topics in conjunction with other internship programs within the UIUC College of Engineering.

COMMUNITY ENGAGEMENTS

As a part of the community development activities, the Science of Security (SoS) Lablet team at the University of Illinois at Urbana-Champaign (UIUC) organized the first "Science of Security for Cyber-Physical Systems (CPS) Workshop." This workshop was part of the "Cyber-Physical Systems Week 2016," an event held in Vienna, Austria in April. The workshop generated vibrant discussion and was attended by approximately 20 researchers from around the world. Leading experts from several universities, including MIT, Carnegie Mellon, UCLA, Oxford, and the KTH Royal Institute of Technology, presented technical lectures about challenges in securing cyber-physical systems and participated in panel discussions. It is expected that several new collaborative initiatives will result from this workshop.

The UIUC SoS Lablet team continued to put forth outreach efforts throughout the Science of Security community. We changed our bi-weekly research-meeting format by sharing the time slot with the Information Trust Institute (ITI) Trust and Security Seminar series, thus increasing to a university-wide audience. We brought in seven distinguished speakers from both academia and industry. We sponsored the "SoS Cyber-Physical Systems Workshop,Äù in conjunction with CPS Week 2016."

UIUC ITI Trust and Security and SoS Lablet Research Seminars

- September 15, 2015, P. Brighten Godfrey, "A Hypothesis Testing Framework for Network Security"
- October 6, 2015, Eric Badger, "Scalable Data Analytics Pipeline for Real-Time Attack Detection; Design, Validation, and Deployment in a Honey Pot Environment"
- October 20, 2015, Mohammad Noureddine, "Accounting for User Behavior in Predictive Cyber Security Models"
- November 3, 2015, Zhenqi Huang and Yu Wang, "SMT-Based Controller Synthesis for Linear Dynamical Systems with Adversary"
- January 26, 2016, Tao Xie,"User Expectations in Mobile App Security"
- March 1, 2016, Wing Lam, "Towards Preserving Mobile Users' Privacy in the Context of Utility Apps"
- March 15, 2016, Dong (Kevin) Jin, "Towards a Secure and Resilient Industrial Control Systems with Software-Defined Networking"

ILLINOIS
UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN

- April 26, 2016, Zhenqi Huang and Yu Wang, "Static-Dynamic Analysis of Security Metrics for Cyber-Physical Systems"
- May 3, 2016, Phuong Cao, "Preemptive Intrusion Detection–Practical Experience and Detection Framework"

SoS Quarterly Meetings

- July 2015, NSA SoS Quarterly Meeting, Ross Koppel, "Promises of the Group Studying Passwords and Cyber Security Circumvention"
- July 2015, NSA SoS Quarterly Meeting, William H. Sanders "Accounting for User Behavior in Predictive Cyber Security Models"
- October 2015, NSA SoS Quarterly Meeting, Sayan Mitra, "Model-based Analysis and Synthesis for Security of Control Systems"
- October 2015, NSA SoS Quarterly Meeting, Eric Badger, "Scalable Data Analytics Pipeline for Validation of Real-Time Attack Detection; Design, Validation, and Deployment in a Honey Pot Environment,"
- February 2016, NSA SoS Quarterly Lablet Meeting, Tao Xie, "User Expectations in Mobile App Security"
- July 2016, SoS Quarterly Meeting, William H. Sanders, "A Quantitative Methodology for Security Monitor Deployment"
- July 2016, SoS Quarterly Meeting, poster session, Jim Blythe, Vijay Kothari, Ross Koppel, and Sean Smith, "Modeling Human Security Behavior: Recent Results on Understanding Compliance"
- July 2016, SoS Quarterly Meeting, poster session, Ross Koppel, Jim Blythe, Vijay Kothari, and Sean Smith, "Beliefs about Cybersecurity Rules and Passwords: A Comparison of Two Survey Samples of Cybersecurity Professionals Versus Regular Users"
- July 2016, SoS Quarterly Meeting, poster session, Sean Smith, Ross Koppel, Jim Blythe, and Vijay Kothari, "Reasons for Cybersecurity Circumvention: A Study and a Model"
- July 2016, SoS Quarterly Meeting, poster session, Wing Lam, Dengfeng Li, Wei Yang, and Tao Xie, "User-Centric Mobile Security Assessment"
- July 2016, SoS Quarterly Meeting, poster session, P. Brighten Godfrey, Matthew Caesar, David Nicol, and William H. Sanders, "A Hypothesis Testing Framework for Network Security"
- July 2016, SoS Quarterly Meeting, poster session, Zachary Estrada, Phuong Cao, Zbigniew Kalbarczyk, and Ravishankar Iyer, "Detection of Malicious Keyloggers in Virtual Desktop Environments"
- July 2016, SoS Quarterly Meeting, poster session, Hui Lin, Homa Alemzadeh, Daniel Chen, Zbigniew Kalbarczyk, and Ravishankar Iyer, "Safety-critical Cyber-physical Attacks: Analysis, Detection, and Mitigation"
- July 2016, SoS Quarterly Meeting, poster session, Zhenqi Huang, Chuchu Fan, Alexandru Mereacre, Sayan Mitra, and Marta Kwiatkoska, "Automatic Safety Verification of Implantable Medical Devices"
- July 2016, SoS Quarterly Meeting, poster session, Yu Wang, Zhenqi Huang, Sayan Mitra, and Geir E. Dullerud, "Differentially Private and Efficient Sequential Learning Algorithms"
- July 2016, SoS Quarterly Meeting, poster session, Peter Kairouz, Sewoong Oh, Kannan Ramchandran, Giulia Fanti, and Pramod Viswanath, "Metadata-Conscious Anonymous Messaging"
- July 2016, SoS Quarterly Meeting, poster session, Ken Keefe and William H. Sanders, "ADVISE ,Äî ADversary View Security Evaluation: Practical Metrics for Enterprise Security Engineering"
- July 2016, SoS Quarterly Meeting, poster session, John C. Mace, Nipun Thekkummal, and Aad van Moorsel, "Sensitivity Analysis of Probabilistic Workflow Models with Security Constraints"

ILLINOIS
UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN

SoS Speaker Series

- September 24, 2015, Patrick McDaniel, Pennsylvania State University, "Seven Years of Mobile Smpartphone Security"
- October 16, 2015, Paul Barford, University of Wisconsin, Madison, "Methods and Characteristics of Fraud in Online Advertising"
- November 13, 2015, Niels Provos, Google, Inc., "Security at Scale"
- February 9, 2016, J. Alex Halderman, University of Michigan, "Logjam: Diffie-Hellman, Discrete Logs, the NSA, and You"
- March 30, 2016, Srdjan Capkun, ETH Zürich Institute of Information Security, "Secure Positioning: From GPS to IoT Applications"
- April 12, 2016, Patrick Traynor, University of Florida, "Who Do I Think You Are? Challenges and Opportunities in Telephony Authentication"

July 2016, SoS Summer Intern Poster Session

- Quentin Mayo and Tao Xie, "OpenSSL: Diving Deeper in Vulnerability Causing Patterns and Reporting Practices using Static Analysis"
- Abhiram Kothapalli, Andrew Miller, and Nikita Borisov, "Smart TRB: An Incentive Compatible Consensus Protocol Utilizing Smart Contracts"
- Andrew Marturano and Masooda N. Bashir, "Exploring the Human Aspect of Computer Security: A Review of Literature"
- Esther M. Amullen, Hui Lin, Zbigniew Kalbarczyk, and Lee Keel, "Multi-Agent System for Detecting False Data Injection Attacks Against the Power Grid"
- Kelly Greeling, Alex Withers, and Masooda Bashir, "Intrusion Detection: Separating the Human from the Program"

Other Presentations

- September 2015, New England Security Day, University of Massachusetts, Amherst, Vijay Korthari, "Mismorphism and Circumvention"
- October 2015, invited tutorial, NSF Interdisciplinary Workshop on Statistical NLP and Software Engineering, Tao Xie, "Software Mining and Software Datasets"
- October 2015, tutorial, 2015 Annual ACM Conference on Systems, Programming Languages, and Applications: Software for Humanity, Tao Xie, "Software Analytics: Achievements and Challenges"
- November 2015, demonstration of the security testbed for attack replay and testing of attack detection techniques at The International Conference for High Performance Computing, Networking, Storage and Analysis (SC2015)
- November 2015, St. Lawrence University, Sean Smith, "Circumvention: Why Do Good People Do Bad Things, and What Can We Do About It"
- November 2015, invited talk, Washington State University, Tao Xie, "Text Analytics for Mobile App Security and Beyond"
- December 2015, Holy Cross College, Sean Smith, "Security Circumvention: why do good people do bad things, and what can we do about it?"
- December 2015, Society for Risk Analysis, Jim Blythe, "A Toolkit for Exploring the Impact of Human Behavior on Cybersecurity through Multi-agent Simulations"
- December 2015, IEEE International Conference on Decision and Control, Zhenqi Huang, "Controller Synthesis for Linear Time-varying Systems with Adversaries"
- December 2015, Technical University of Vienna, Sayan Mitra, "Parametrized Verification of Distributed Systems"
- December 2015, PhD preliminary examination, Zhenqi Huang, "Compositional Verification and Security of Cyber-physical Systems"

ILLINOIS
UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN

- January 2016, University of Central Florida, Department of Computer Science Spring 2016 Distinguished Speaker Series, David Nicol, "Quantitative Analysis of Stepping Stone Access to Cyber-Physical Assets"

- April 2016, Symposium and Bootcamp on the Science of Security (HoTSoS 2016), invited tutorial, Tao Xie and William Enck,"Text Analytics for Security"

- May 2016, 38th International Conference on Software Engineering (ICSE 2016), Austin, TX, Tao Xie, "Measuring Code Behavioral Similarity for Programing and Software Engineering Education"

- June 2016, Department of Computer Science, University of Central Florida, Orlando, FL, invited speaker, Tao Xie, "User Expectations in Mobile App Security"

- June 2016, European Control Conference (ECC), Aalborg, Denmark, Geir E. Dullerud, "Lyapunov Constructions, Formal Proof Frameworks, and Computation-Based Verification for Complex Systems"

- July 2016, 36th International Workshop on Operator Theory and Applications (IWOTA 2016), St. Louis, MO, Geir E. Dullerud, "Operators and Feedback Control Theory: Linear Switched Systems"

- August 2016, Fermilab Computing Techniques Seminar, Batavia, IL, invited seminar, Kevin Jin, "Uncertainty-Aware Network Verification in Software-Defined Networks"

- September 2016, Assured Cloud Computing Weekly Research Seminar, University of Illinois at Urbana-Champaign, Key-whan Chung, "An Indirect Attack on Computing Infrastructure through Targeted Alteration on Environmental Control"

Workshops

- Science of Security for Cyber-Physical Systems (SoSCYPS 2016) Workshop, April 11, 2016 in conjunction with CPS Week 2016, in Vienna, Austria. The workshop agenda and a list of speakers and topics can be found at the SoSCYPS webpage. http://publish.illinois.edu/science-of-security-lablet/science-of-security-for-cyber-physical-systems-workshop/

- Workshop on Science of Security through Software-Defined Networking (SoSSDN 2016), June 16-17, at the Illinois Institute of Technology in Chicago, IL. The goal of the workshop was to identify opportunities and challenges in using SDNs to advance the science of security. We brought together leaders from academia, industry, national laboratories, and government agencies in the areas of SDN and security in this two-day workshop. The workshop consisted of 11 invited talks (2 keynotes), a poster session (10 accepted posters) and a panel on "How to Make Security for SDN a Science?" Covered topics included SDN principles that support formal and experimental analysis of security; metrics for SDN security; identifying hard open problems for academic research in SDN security; SDN-based testbeds and cyber-infrastructures in security research; success and failures in designing for resilient and secure networks; and identifying tools and techniques that can advance networks/systems security research. http://publish.illinois.edu/science-of-security-lablet/workshop-on-science-of-security-through-software-defined-networking

ILLINOIS
UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN

UNIVERSITY OF MARYLAND

PI JONATHAN KATZ

The University of Maryland (UMD) Science of Security (SoS) Lablet, led by co-PIs Jonathan Katz and Michel Cukier, has ten research projects that address different aspects of the five Hard Problems, with specific focus on the areas of metrics, policy-governed secure collaboration, and human behavior.
The UMD lablet consists of twenty faculty from both UMD and partner institutions. The fifteen UMD faculty are drawn from five different departments across campus, including computer science, electrical and computer engineering, information studies, criminology, and reliability engineering. The collaborators hail from the United States Naval Academy, Virginia Polytechnic Institute and State University (Virginia Tech), The University of Texas at Austin, Indiana University, and The George Washington University.

FUNDAMENTAL RESEARCH

PROJECT: Understanding Developers' Reasoning about Privacy and Security

PIs: Michelle Mazurek, Charalampos Papamanthou, Mohit Tiwari
Participating Sublablet: University of Texas at Austin
Hard Problem: Human Behavior

Researchers on this project are developing an Information Flow Control (IFC) platform called Blox that is intended to make it easier and more intuitive for users to control how their personal information is used, shared, and disseminated. The major highlights of the past year are completing a functional prototype of Blox and conducting a user study on access control in practice. The prototype includes a user interface, similar to Dropbox, for managing data as well as a backend that implements IFC on model-view-controller web applications. Additionally, the team has written a secure templating language that enables the viewing of cross-folder information on the platform without compromising security. This allows developers to retain all their application's functionality when porting an application to Blox, something that other IFC platforms for web applications are unable to do. The team analyzed the porting effort required to add both new and existing applications to the Blox platform. A team of undergraduate students has written an in-browser integrated development environment and calendar application on Blox, illustrating previously missing functionality in the templating language. This work is complemented by a user study analyzing how compatible the Blox platform is with a typical user's workflow. This study pulls information from Mechanical Turk participants' Google Drive, Gmail, and Google Calendar accounts, and uses machine learning to create logical groupings of access-control decisions of data. The accuracy of these groupings is then tested by asking participants to answer a series of questions. The team plans to submit two papers this coming year based on these results.

PROJECT: Measuring and Improving the Management of Today's PKI
PI: Dave Levin
Hard Problems: Security Metrics and Models, Human Behavior

This project focuses on metrics by means of large-scale measurements of the existing Public-Key

Infrastructure (PKI) used in today's web. Over the past year, the team has investigated the roles that administrators, Content Distribution Networks (CDNs), and browsers play in the PKI through these measurement studies. One such study focused primarily on certificate revocation, and found that a surprisingly large fraction (8%) of served certificates has been revoked, yet no browsers fully check for revocations, and mobile browsers perform no revocation checks whatsoever. Another study investigated invalid certificates in the Web's PKI, and found that they constitute a shocking 88% of all certificates in a four-years-long dataset. Further investigation found that these arise mostly from end-user devices such as home routers, IPTVs, and printers. The team developed techniques that allow them to track devices by their certificates to observe user mobility and IP address reassignment policies. Finally, the team has been investigating the role that content distribution networks (CDNs) and web-hosting services play in the Web's PKI, having observed that they have access to their customers' private keys (in contrast to what is typically assumed). As part of this work, the team has developed techniques that allow them to "anti-alias" domain names by determining, through machine learning over network data and "who is" data, whether two domain names correspond to the same administrative entity. This study is still in progress, but initial results indicate extensive key sharing in the Web's PKI, which has serious implications on the security of popular Web services.

PROJECT: Trust, Recommendation Systems, and Collaboration

PIs: John S. Baras, Jennifer Golbeck
Hard Problems: Policy-Governed Secure Collaboration, Scalability and Composability, Human Behavior

Work done as part of this project is primarily directed toward the hard problem of policy-governed secure collaboration. The goal of the project is to develop a transformational framework for a science of trust, and its impact on local policies for collaboration, in networked multi-agent systems. The framework will take human behavior into account from the start by treating humans as integrated components of these networks, interacting dynamically with other elements. Work so far has shown novel results regarding the evolution of opinions (or beliefs) over a social network modeled as a signed graph; new models and analytical methods for the investigation of consensus dynamics with both collaborative and non-collaborative node interactions; and new probabilistic models of multi-domain crowdsourcing tasks. The team has also formalized the problem of trust-aware task allocation in crowdsourcing and developed a principled way to solve it. The formulation models the workers' trustworthiness and the costs based on both the question and the worker group. In other work, the team performed an experimental study on the concerns, knowledge, and perceptions on privacy among users of the Internet. The primary focus of this work was to understand whether users are continuing to share once they are aware of the privacy risks and have made an informed choice about what they are comfortable sharing, or whether they operating under false assumptions or without the knowledge they need to make an informed choice.

PROJECT: User-Centered Design for Security

PIs: Jennifer Golbeck, Adam J. Aviv
Participating Sublablet: United States Naval Academy
Hard Problems: Security Metrics and Models, Human Behavior

Researchers on this project have made significant progress in measuring and applying metrics of security to mobile authentication, particularly graphical password systems on Android, and in using those metrics to design systems to improve the human factor in password selection, such as password meters. The team has also completed a paper on user perception and understanding of privacy issues related to personal information sharing in apps. This work focused on Facebook apps and set out to understand how concerned users are about privacy and how well-informed they are about what personal data apps can access. We found that initially, subjects were generally under-informed about what data apps could access from their profiles. After viewing additional information about these permissions, subjects' concern about privacy on Facebook increased. Subjects' understanding of what data apps were able to access increased, although even after receiving explicit information on the topic many subjects still did not fully understand the extent to which apps could access their data.

PROJECT: Reasoning about Protocols with Human Participants

PIs: Jonathan Katz, Poorvi L. Vora
Participating Sublablet: George Washington University
Hard Problem: Human Behavior

The goal of this project is to study protocols–in particular, electronic-voting protocols–in which humans are explicitly modeled as participants. In the last year, we have been working on two major goals: finalizing the Remotegrity voting protocol and developing proofs of its security properties, and using aspects of Remotegrity to develop a new protocol that improves the dispute-resolution properties of Helios, assuming an honest registrar (credential provider). In Helios currently, the voting terminal can change a vote after it obtains the voter's credential, and although the voter will be aware that her vote was changed, she will be unable to prove it to a third party. In addition to other changes, the new protocol uses Remotegrity's lock-in code to prevent the vote from being changed.

PROJECT: Empirical Models for Vulnerability Exploits

PI: Tudor Dumitras
Hard Problem: Security Metrics and Models

This team is researching more-informative metrics to quantify security of deployed systems. This year, the team explored machine-learning techniques for preventing global malware dissemination. Today, few malware families have the ability to propagate autonomously. Instead, they rely on malware-delivery networks, which specialize in helping malware infect millions of hosts worldwide. These malware-delivery techniques largely rely on two key components: (1) drive-by download exploits, which enable the initial malware insertion and (2) downloader Trojans, which retrieve additional malware from the Internet. Two main results were obtained. First, the team designed and implemented a system for forecasting which vulnerabilities would be exploited in the wild, using information mined from Twitter. This system can be used to prioritize responses to vulnerability disclosures, or to model risk for cyber-insurance applications. Second, the project introduced a downloader-graph abstraction, which captures the client-side activity of malware delivery networks. Properties of downloader graphs were in turn used to train a classifier that has the potential to expose large parts of the malware-download activity that may otherwise remain undetected.

PROJECT: Human Behavior and Cyber Vulnerabilities

PI: V.S. Subrahmanian
Participating Sublablet: Virginia Tech
Hard Problems: Human Behavior, Security Metrics and Models

Researchers on this project are using an empirical approach to study factors that affect the rate at which security patches are deployed. When a vulnerability is exploited, software vendors often release patches fixing the vulnerability. However, prior research has shown that some vulnerabilities continue to be exploited more than four years after their disclosure. Why? There are both technical and sociological reasons for this. On the technical side, it is unclear how quickly security patches are disseminated, and how long it takes to patch all the vulnerable hosts on the Internet. On the sociological side, users/administrators may decide to delay the deployment of security patches. The goal of this task is to validate and quantify these explanations. Specifically, it seeks to characterize the rate of vulnerability patching, and to determine the factors–both technical and sociological–that influence the rate of applying patches.

PROJECT: Does the Presence of Honest Users Affect Intruders' Behavior?

PIs: Michel Cukier, David Maimon
Participating Sublablet: University of Texas at Austin
Hard Problem: Human Behavior

As part of this project, Michel Cukier and David Maimon are applying criminological techniques to develop a better understanding of attacker behavior. One particular highlight of the past year is the examination of previously uninvestigated experimental data–an experiment that randomly assigned infiltrated target computers to have a certain type (administrative or non-administrative) and number (1 or 10) of users to appear on the system at the same time as the system trespasser. Using this data, the team examined whether the number and type of users present on a system reduced the seriousness and frequency of

trespassing. Results indicated that the presence of an administrative user (regardless of the number of users) significantly reduced the number of system trespassing events. Additionally, with 10 users present, the presence of an administrative user significantly reduced the total amount of time attackers spent on the compromised system. Interestingly, comparing between conditions with different numbers of users, it was found that the number of users present on the system has no effect on the number of trespassing events or total time spent on the system. These findings together indicate that presence of an administrative user can produce a deterrent effect on system trespassers, while the number of users present on the system has no effect on system trespasser actions.

PROJECT: Understanding How Users Process Security Advice

PI: Michelle Mazurek
Hard Problem: Human Behavior

This project addresses the hard problem of human behavior from the perspective of educational efforts. People encounter a tremendous amount of cybersecurity advice. It would be impossible to follow all the available advice, so people pick and choose which advice to follow and which to ignore in different circumstances. But the advice they pick is not always the most correct or useful. This project examines where people encounter security advice, how they evaluate its trustworthiness, and how they decide which advice to follow or reject. This year, the team first completed a qualitative study on this topic. Key findings include that participants evaluate digital security advice based on the trustworthiness of the advice sources, as compared to evaluating physical security advice based on their intuitive assessment of the advice content, and that participants reject advice for a variety of reasons including that it contains marketing material or threatens their privacy. More recently, the team also completed data collection and began data analysis for a large-scale quantitative survey designed to validate the findings from the qualitative study, and began development and testing of security behavior interventions based on the findings.

PROJECT: Trustworthy and Composable Software Systems with Contracts

PIs: David Van Horn, Jeffrey S. Foster, Michael Hicks, Sam Tobin-Hochstadt

Participating Sublablet: Indiana University
Hard Problem: Scalability and Composability

As part of this project, researchers are investigating compositional-verification techniques using language-based mechanisms called contracts for specifying and enforcing program properties. Results confirm that behavioral properties of programs can be verified using this approach, and the team members are now trying to scale the approach to cover multi-language programs and more-complex security properties.

PUBLICATIONS

PROJECT: Measuring and Improving the Management of Today's PKI

- Frank Cangialosi, Dave Levin, and Neil Spring, "Ting: Measuring and Exploiting Latencies Between All Tor Nodes," in *Proceedings of the 2015 Internet Measurement Conference (IMC 2015)* Tokyo, Japan, October 28-30, 2015, ACM Press, New York, NY, 2015, pp. 289-302.

PROJECT: Trust, Recommendation Systems, and Collaboration

- Xiangyang Liu, He He, and John S. Baras, "Trust-Aware Optimal Crowdsourcing with Budget Constraint," in *Proceedings of the 2015 IEEE International Conference on Communications (ICC 2015)*, London, UK, June 8-12, 2015, pp. 1176-1181.

- Xiangyang Liu, He He, and John S. Baras, "Crowdsourcing with Multi-Dimensional Trust," in *Proceedings of the 18th IEEE International Conference on Information Fusion (FUSION 2015)*, Washington DC, 2015, pp. 574-881.

- Peixin Gao, John S. Baras, and Jennifer Golbeck, "Semiring-Based Trust Evaluation for Information Fusion in Social Network Services," in *Proceedings of the 18th IEEE International Conference on Information Fusion (FUSION 2015)*, Washington DC, 2015, pp. 590-596.

- Guodong Shi, Alexandre Proutiere, Mikael Johansson, John S. Baras, and Karl Henrik Johansson, "Emergent Behaviors Over Signed Random Dynamical Networks: State-Flipping Model," in *IEEE Transactions on Control of Network Systems (IEEE TCNS*), vol. 2, no. 2, pp. 142-153, June 2015.

- Hui Miao, Peixin Gao, Mohammad Tagji HajiAghayi, and John S. Baras, "HyperCubeMap: Optimal Social Network Ad Allocation Using Hyperbolic Embedding," in Proceedings of the 2015 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM 2015), Paris, France, 2015, pp. 357-362.

- Xiangyang Liu and John S. Baras, "Trust-Aware Crowdsourcing with Domain Knowledge," in *Proceedings of the 54th IEEE Conference on Decision and Control (CDC 2015)*, Osaka, Japan, 2015, pp. 2913-2918.

- Evripidis Paraskevas, Tao Jiang, and John S. Baras, "Trust-Aware Network Utility Optimization in Multihop Wireless Networks with Delay Constraints," in *Proceedings of the 24th Mediterranean Conference on Control and Automation (MED*), Athens, Greece, 2016, pp. 593-598.

- Eirini Eleni Tsiropoulou, John S. Baras, Symeon Papavassiliou, and Gang Qu, "On the Mitigation of Interference Imposed by Intruders in Passive RFID Networks," in: Zhu Q., Alpcan T., Panaousis E., Tambe M., Casey W., eds., *Decision and Game Theory for Security*. GameSec 2016. Lecture Notes in Computer Science, vol. 9996, pp. 62-80, 2016, Springer, Cham, Switzerland.

- Pexin Gao, Hui Miao, John S. Baras, and Mohammad Taghi HajiAghayi, "Social Network Ad Allocation and Optimization: A Hyperbolic Embedding-Based Approach," *53rd IEEE Conference on Decision and Control,* 2014, in Social Network Analysis and Mining Journal (SNAM), 1/2016, Springer.

PROJECT: User-Centered Design for Security

- Adam J. Aviv, Justin Maguire, and Jeanne Luning Prak, "Analyzing the Impact of Collection Methods and Demographics for Android's Pattern Unlock," in *Proceedings of the Workshop on Usable Security (USEC)*, February 21, 2016, San Diego, CA, 2016.

- Abdullah Ali, Ravi Kuber, and Adam J. Aviv, "Developing and Evaluating a Gestural and Tactile Mobile Interface to Support User Authentication," in Proceedings of the iConference 2016, March 2016.

- Jennifer Golbeck and Matthew Louis Mauriello, "User Perception of Facebook App Data Access: A Comparison of Methods and Privacy Concerns," Future Internet, vol. 8, art. 2, pp. 9, 2016.

- Flynn Wolf, Ravi Kuber, and Adam J. Aviv, "Preliminary Findings from an Exploratory Qualitative Study of Security-Conscious Users of Mobile Authentication," in *Proceedings of the Workshop on Security Information Workers at SOUPS* 2016, Denver, CO, June 22-24, 2016.

- Adam J. Aviv, Markus Dürmuth, and Payas Gupta, "Position Paper: Measuring the Impact of Alphabet and Culture on Graphical Passwords," in *Proceedings of the Who Are You?! Adventures in Authentication Workshop*, Denver, CO, June 22-24, 2016.

- Flynn Wolf, Ravi Kuber, and Adam J. Aviv, "Towards Non-Observable Authentication for Mobile Devices," poster, *12th Symposium on Usable Privacy and Security (SOUPS 2016)*, Denver, CO, June 22-24, 2016.

- Susanna Heidt and Adam J. Aviv, "Refining Graphical Password Strength Meters," poster, *12th Symposium on Usable Privacy and Security (SOUPS 2016)*, Denver, CO, June 22-24, 2016.

- Adam J. Aviv, Justin Maguire, and Jeanne Luning-Prak., "Analyzing the Impact of Collection Methods and Demographics for Android's Pattern Unlock," in P*roceedings of the Workshop on Usable Security (USEC)*, February 21, 2016, San Diego, CA, 2016.

- Abdullah Ali, Ravi Kuber and Adam J. Aviv, "Developing and Evaluating a Gestural and Tactile Mobile Interface to Support User Authentication," in *Proceedings of the iConference* 2016, March 2016.

- Jennifer Golbeck and Matthew Louis Mauriello, "User Perception of Facebook App Data Access: A Comparison of Methods and Privacy Concerns," Future Internet, vol., 8, no. 2, 2016.

- Adam J. Aviv and Ravi Kuber, "Do Privacy Attitudes on Mobile Devices Impact the Strength of Unlock Authentication?" submitted to *CHI '16*.

- Flynn Wolf, Ravi Kuber, and Adam J. Aviv, "An Empirical Study Examining the Perceptions and Behaviors of Security Conscious Users of Mobile Authentication," submitted to *CHI '16*.

PROJECT: Reasoning about Protocols with Human Participants

- Richard T. Carback et al., "The Scantegrity Voting System and its Use in the Takoma Park Elections," in Real-World Electronic Voting: Design, Analysis and Deployment, Feng Hao and Peter Y. A. Ryan, eds., 2016, Auerbach Publications.

PROJECT: Empirical Models for Vulnerability Exploits

- Carl Sabottke, Octavian Suciu, and Tudor Dumitras, "Vulnerability Disclosure in the Age of Social Media: Exploiting Twitter for Predicting Real-World Exploits," in *Proceedings of the 24th USENIX Security Symposium (USENIX Security '15)*, Washington, DC, August 12-14, 2015, pp. 1041-1056.

- Bum Jun Kwon, Jayanta Mondal, Jiyong Jang, Leyla Bilge, and Tudor Dumitras, "The Dropper Effect: Insights into Malware Distribution with Downloader Graph Analytics," in *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security (CCS '15)*, Denver, CO, October 12-15, 2015, ACM Press, New York, NY, 2016, pp. 1118-1129.

- Soyumya Indela, Mukul Kulkarni, Kartik Nayak, and Tudor Dumitras, "Helping Johnny Encrypt: Toward Semantic Interfaces for Cryptographic Frameworks," in *Proceedings of the 2016 ACM International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software (Onward!* 2016), Amsterdam, The Netherlands, ACM Press, New York, NY, 2016, pp. 180-196.

PROJECT: Human Behavior and Cyber Vulnerabilities

- Arunesh Mathur, Josefine Engel, Sonam Sobti, Victoria Chang, and Marshini. Chetty, "They Keep Coming Back Like Zombies: Improving Software Updating Interfaces," in *Proceedings of the 12th Symposium on Usable Privacy and Security (SOUPS 2016)*, Denver, CO, June 22-24, 2016, pp. 43-58.

- Chanhyun Kang, Noseong Park, B.Aditya Prakash, Edoardo Serra, and V.S. Subrahmanian, "Ensemble Models for Data-Driven Prediction of Malware Infections," in *Proceedings of the 9th ACM International Conference on Web Search and Data Mining (WSDM '16)*, San Francisco, CA, February 22-25, 2016, pp. 583-592.

- Sushil Jajodia, Noseong Park, Edoardo Serra, and V.S. Subrahmanian, "Using Temporal Probabilistic Logic for Optimal Monitoring of Security Events with Limited Resources," *Journal of Computer Security,* vol. 24, no. 6, pp. 735-791, IOS Press, Amsterdam, The Netherlands, December 2016.

- V.S. Subrahmanian et al., "The DARPA Twitter Bot Challenge," *Computer*, vol. 49, no. 6, pp. 38-46, IEEE Computer Society, June 2016.

PROJECT: Understanding How Users Process Security Advice

- Elissa M. Redmiles, Amelia Malone, and Michelle L. Mazurek, "I Think They're Trying to Tell Me Something: Advice Sources and Selection for Digital Security," in *Proceedings of the 2016 IEEE Symposium on Security and Privacy (SP)*, San Jose, CA, May 22-28, 2016, pp. 272-288.

- Elissa M. Redmiles, Sean Kross, and Michelle L. Mazurek. "How I Learned to Be Secure: A Census-Representative Survey of Security Advice Sources and Behavior," in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security (CCS '16)*, ACM Press, New York, NY, 2016, pp. 666-677.

PROJECT: Trustworthy and Composable Software Systems with Contracts

- Phúc C. Nguyên, Sam Tobin-Hochstadt, and David Van Horn, "Higher Order Symbolic Execution for Contract Verification and Refutation," *Journal of Functional Programming*, vol. 27, 2017. Published online: 21 December 2016.

- Thomas Gilray, Steven Lyde, Michael D. Adams, Matthew Might, and David Van Horn, "Pushdown Control-Flow Analysis for Free," in *Proceedings of the 43rd ACM SIGPLAN-SIGACT Symposium on Principles in Programming Languages (POPL 2016)*, St. Petersburg, Florida, January 20-22, 2016, pp. 691-704.

EDUCATIONAL

Dave Levin added content on Pubic Key Infrastructure (PKI) revocation and Content Delivery Network (CDN) to his undergraduate course on computer and network security.

Mohit Tiwari at the University of Texas, Austin, has introduced a three-semester sequence in cybersecurity where first- and second-year undergraduates read security papers, write reviews, and work on projects with graduate students (in addition to doing standard homework). The goal is to introduce students to both systems and theory aspects of cybersecurity.

John Baras introduced concepts, models, and algorithmic evaluation of trust in graduate courses on multi-agent control. Since 2010, he has been teaching a capstone course entitled ENES 489P: "Hands-on Projects in Systems Engineering." In this project-oriented course, groups of undergraduates (3–4 students) work on projects inspired from important practical challenges. In the last two years, several of these projects addressed security-related questions and challenges.

Adam Aviv has incorporated elements of his research as capstone projects in his courses. He has also involved several undergraduates in his research. He is developing a senior-level elective on cybersecurity, as well as one focusing on usable security.

Jennifer Golbeck, Mike Hicks, and Jonathan Katz have developed a series of related courses as Massive Open Online Courses (MOOCs) on Coursera focusing on usable security, software security, and cryptography, respectively. The courses are part of an interdisciplinary MOOC specialization in cybersecurity, with courses covering programming-language security, cryptography, and usable security.

Michelle Mazurek developed a new graduate course, CMSC 414: "Human Factors in Security and Privacy," that was taught in Spring 2015 and Spring 2016. As a project for Mazurek's Spring 2016 course, five students planned and conducted a participatory design workshop for developing entertaining, relatable, and educational videos to convince viewers to accept software updates. This was inspired directly by the results of her qualitative study that suggested that relatable fiction is a strong vehicle for learning security behaviors. The workshop served as a pilot study. A follow-up study that applies participatory design to users recruited in pairs is currently underway. Her goal is to develop a high-quality 5–10 minute storyboard and contract with the Maryland Filmakers Club to produce it as a video. She will then evaluate its usefulness for security education.

David Van Horn has given tutorials on the formal-method tools and techniques he developed in the hopes of advancing the best practices used to develop high-assurance software. The content of his research has been incorporated into the University of Maryland graduate class "Program Analysis and Understanding." Van Horn presented a tutorial on this material at the ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL) in St. Petersburg, Florida, in January 2016. He has incorporated his lablet research into his graduate class on "Program Analysis and Understanding." He is also working to incorporate this into the pedagogically oriented programming environment accompanying his textbook, "How to Design Programs." A web interface is being investigated for the system so that users can experiment with the system without needing to install specialized software.

Michel Cukier leads the Advanced Cybersecurity Program Experience for Students (ACES) undergraduate honors program in cybersecurity.

The program incorporates a holistic approach to cybersecurity covering technical, policy, and behavioral aspects of the problem.

COMMUNITY ENGAGEMENTS

Dave Levin presented the results of the study on PKI revocation at several academic institutions, including Georgia Tech, UC Riverside, the University of Jordan in Amman, Jordan, and the Max Planck Institute for Software Systems. He also presented to a wide range of non-academic audiences, including the NMEA Conference, the RTCM Conference, the CyberSci Summit, the ICF International Conference, and to Cloudflare, one of the largest CDNs to host HTTPS content. The audiences consisted of a wide range of practitioners who are influential in developing communication policies at both institutional and international levels. He also presented these results to graduate and undergraduate students at the University of Maryland (UMD), as well as students and faculty at several other universities.

John Baras participated heavily in the work of the Transatlantic Summit Project developing frameworks for collaboration and joint funding in the area of cyber-physical systems (CPS). Security challenges, especially including human and technological networked systems, were part of his work. He was also deeply involved in the National Institute of Standards and Technology (NIST)-organized Cyber-Physical Systems Public Working Group, and in particular, with the subgroup working on security problems and formulations for CPS. Baras gave several invited talks about his work at venues including the 2015 Cyber-Security and Privacy Winter School (October 21-23, 2015, Stockholm, Sweden), the University of Cambridge, the University of Oxford, and the 2015 International Symposium on Industrial Control Systems & SCADA Cyber Security Research (September 17-18, 2015, Ingolstadt, Germany). In addition, he delivered an invited plenary address at the 35th Chinese Control Conference (CCC 2016), July 27, 2016, in Chengdu, China, entitled "Networked Cyber-Physical Systems (Net-CPS)," in which he included his work on trust supported by the NSA Science of Security (SoS) Lablet grant.

Adam Aviv was workshop and tutorial chair for SOUPS 2016 and was on the program committee of the Privacy Enhancing Technologies Symposium (PETS 2016). He gave invited talks about his work at UC Berkeley, and the International Computer Science Institute (March 2016). Aviv served on the program committee of the Computer Security Applications Conference (ACSAC '16), and he is on the program committee of PETS 2017. He is a steering committee member of the Advances in Computer Security Education (ASE) Workshop.

Jennifer Golbeck gave several invited talks about her work to audiences in industry as well as to K-12 educators and students.

Poorvi Vora gave an invited talk and participated in a panel, at the Remote Voting Conference in July 2016. The meeting was organized by the Government of India to discuss the challenges of, and possible solutions for, remote voting by Indian citizens. She is also part of the technical team for the end-to-end verifiable internet voting (E2E VIV) project (examining the feasibility of secure internet voting) of the overseas vote foundation (OVF). She has been contributing to a description of end-to-end, independently-verifiable voting systems meant for non-technical readers, including election officials. She also testified to the Maryland State Board of Elections in September 2016, regarding its online ballot-delivery system, and participated in a panel about voting irregularities for the monthly seminar of the Washington Statistical Society on October 4, 2016.

Jonathan Katz gave an invited talk on "Usable Cryptography,"at the Workshop on Human Factors in Cybersecurity Design (Hebrew University, March 2016), which included a discussion of security proofs for electronic-voting protocols. Katz serves as program chair for Crypto 2016-2017, as well as program co-chair for HoTSoS 2017, and he is a member of the steering committee for the IEEE Cybersecurity Initiative and the Maryland Cybersecurity Council.

Tudor Dumitras gave invited talks about his research at Qualcomm Research at UNC Charlotte, Boston University, Symantec Research Labs, and the AT&T Security Research Center.

Tudor Dumitras and Michelle Mazurek co-led a discussion session

at HotSec '15 on security research conducted with non-public data and the impact these research methods have on the science of security.
Michelle Mazurek co-led a tutorial at SOUPS 2016 on the science of password research.
David Maimon presented a poster at the annual meeting of the American Society of Criminology (ASC).

David Van Horn gave a tutorial titled "An Introduction to Redex with Abstracting Abstract Machines," at the 2016 Annual Symposium on Programming Languages (POPL). He also gave an invited talk about his work at the Dagstuhl Seminar on Language Based Verification Tools for Functional Programs in March 2016. Van Horn presented his work at the National Institute of Informatics (NII), Japan special meeting on "Higher-Order Model Checking." He has been invited to present a tutorial at the 2016 ACM SIGPLAN International Conference on Functional Programming.

Michael Hicks served as program chair for the 2016 IEEE Security Development Conference (SecDev), whose goal was to encourage and disseminate ideas for secure system development among both academia and industry. He also serves on the IDA/CCS program review committee. He has been blogging about programming-language security at pl-enthusiast.net

Marshini Chetty gave talks on her research results at the Center for Information and Technology Policy (CITP) at Princeton University and to the HCI group at the Jacobs Technion-Cornell Institute in November.

Graduate student Elissa Redmiles received a "data grant" from the Data and Society Research Institute to study security habits of low-socioeconomic status (low-SES) Americans, in part due to her work as part of the UMD SoS Lablet.

Science of Security Quarterly Meetings

Winter 2016 Quarterly: North Carolina State University

The winter 2016 quarterly Science of Security (SoS) Lablet meeting, sponsored by the National Security Agency (NSA), was hosted at North Carolina State University (NCSU) on February 2 and 3, 2016. Laurie Williams and Munindar Singh, Principal Investigators (PIs) at NCSU, hosted the event. NSA and the four Lablets, Carnegie Mellon University (CMU), North Carolina State University, University of Illinois at Urbana-Champaign (UIUC), and the University of Maryland (UMD), provided speakers, shared current research, presented interim findings, and stimulated thought and discussion about the Science of Security.Panel discussions and focus groups provided an opportunity for researchers to interact both with each other and with guests from the government and industry to address the Hard Problems of cyber security. The importance of good design was a theme that ran through the presentations.

The keynote by Henry Petroski, noted civil engineer, author, and professor from Duke University, addressed the paradox between success and failure in design. Illustrating his point with historic failures in bridge design and construction, he showed how success, over time, leads to complacencies that in turn lead to failure. Conversely, failure stimulates revisions in design that can produce successes. From the 1850s experience to the present, the paradox of design is that anticipating failure leads to success and successful designs evolve into failures.

Peter Loscocco of NSA presented a related keynote addressing, "An Approach to Secure Design." His approach is to look at the design process and develop a methodology. He documents a methodical process for design that can be easily taught, produces suitable designs that have been analyzed, captures reasoning behind the design decisions, and enables understanding for consequences of modifications. One big challenge is documentation. Using a design tree provides a tangible artifact of the design process and allows the use of threat models as assumptions.

Guided discussions and breakout groups addressed the security metrics and human aspects in security Hard Problems. The security metrics discussion addressed the importance of measurement and asked the questions: "Context always matters, so how do we protect against attacks that haven't been thought of yet? Can metrics help? Many current metrics are on the negative side, measuring, for example, attacks, and failures, so can we develop the positive? The human factors workshop determined that different traits make people susceptible to threats in different ways, and that cognitive modeling can help understand human interaction with security.

A panel of leading researchers from the four Lablets and guest speakers Warren Grunbok from IBM and Andrew Porter from Merck provided their views on how to transfer technology and the value of Science of Security research into the private sector. Communication and an iterative approach seemed to offer the greatest opportunities for success as a consensus of the group.

Technical research presentations included papers by each Lablet. Tao Xie, UIUC, presented his study on "AppContext: Differentiating Malicious and Benign Mobile App Behavior under Context." Jonathan Aldrich, CMU, presented "Capability-Based Architectural Control." A study of user-generated pattern passwords was presented by University of Maryland-affiliated researcher Adam Aviv from the U.S. Naval Academy. Robert Proctor, an NCSU cognitive psychologist, addressed ways for people to detect phishing attacks.

Adam Tagert, Science of Security Technical Director,

NSA Research Directorate, spoke on human subject research procedures at the Department of Defense (DoD), and how to coordinate with university institutional review boards. Beth Richards, NCSU Laboratory for Analytic Sciences (LAS), described LAS as an NSA lab using non-traditional data and approaches from open sources to get to "anticipating," that is, to move from reaction or observation of threats and attacks to anticipation, to get ahead of the foreign adversary, and run at scale and speed since the nature of the threat requires a real-time response.

Updates on progress in measuring advancement in SoS were presented by Jeffrey Carver, University of Alabama, and on evaluation of the research and research publications by Lindsey McGowan, NCSU. Carver's talk reviewed a rubric-based method of evaluating the scientific content of articles published in IEEE Security & Privacy.

More than a dozen excellent student poster presentations provided an opportunity to see a range of Science of Security research and discuss issues, methods, and findings.

During the business sessions of the Winter 2016 Lablet Quarterly meeting, the NSA Research Directorate presented the, "Science of Security (SoS) Initiative Annual Report 2015," to the Lablet PIs. Shown on the left are Bill Scherlis and Laurie Williams, Principal Investigators at CMU and NCSU, holding the report.

Spring 2016 Quarterly: In lieu of a Spring 2016 Quarterly, the SoS Community met at HoTSoS, April 19-21, 2016, hosted by Carnegie Mellon University in Pittsburgh, PA.
See page 82 for full Hot SoS details.

Summer 2016 Quarterly: University of Illinois at Urbana-Champaign

The summer 2016 quarterly Science of Security (SoS) Lablet meeting, sponsored by NSA, was hosted at the University of Illinois at Urbana-Champaign July 26 and 27, 2016. David Nicol, Principal Investigator at UIUC, chaired the event. Lablet researchers shared current research and stimulated thought and discussion about the Science of Security hard problem of resiliency. A workshop, five technical papers, a keynote, and an industry panel provided an opportunity for researchers to interact with each other and with guests from the government and industry.

The keynote by Roger Hill, CTO of Veracity Security Intelligence, addressed "An Industry Approach to a Foundational Security Platform Through the Use of Software Defined Networking for Industrial Networks." He began with an overview of industrial control systems layers: the Enterprise zone; DMZ; and manufacturing, security, and safety zones. Identified problems include visibility of the systems (unknown assets, no active scanning), complexity (long life cycles/mixed modes, lack of SMEs), and risk (misconfiguration, misuse, emerging and targeted threats). Challenges that include operational requirements at odds with cybersecurity needs, limited security controls at the switch, no auditing, frequent misconfigurations, and redundancy increase complexity. "Leveraging SDN technologies, [his company] gives industrial operators complete situational awareness for asset identification and management, client-defined security zones, and threat-based incident response in near real-time. Our approach improves the reliability, efficiency, and security of OT networks without adding additional layers of complexity to the network," he said.

The Workshop on Cyber Resilience with James Holt, Steve Danko, and Ahmad Ridley, all from DoD, was enthusiastically interactive. Panelists and the audience offered a range of comments about how to produce research that disrupts the adversary's ability to get in, stay in, and act within our systems and networks. The targeted impact of this research should, according to the panel, reverse the asymmetry between attack and defense, develop enterprise-level "sense making" using disparate data, reduce operator overload, redefine productivity in terms of enabling proactive actions versus reactive information sifting, and fully engage the research

community. Their recommended approach: to advance multi-year research that builds foundational science to "human-aided, system-driven, and automated response."

Jim Lenz, John Deere; Mark Scott, Forcepoint; and David Greve, Rockwell Collins offered an industry panel providing their views and interacting with the SoS research audience. Asked what gaps they see in advancing security that a science of security could help solve, they answered they are looking for more research analytics for threat detection, modelling, biometrics, quantum crypto, MANETs, VMs, locking down desktop environments at the hypervisor level, multimedia, general malware and ransomware, and hardening the code they write.

Technical research presentations included papers by each lablet. David Garlan, CMU, addressed, "Improving Resilience through Synthesis of Adaption Strategies." "Measuring Perceived Privacy Risk in Cybersecurity Information Sharing," was offered by Travis Breaux, also from CMU. Bill Sanders, UIUC, presented "A Quantitative Methodology for Security Monitor Deployment." Michael Reiter, University of North Carolina at Chapel Hill spoke on "Simplifying Software-defined Network Optimization using SOL [SDN Optimization Layer]." "Evidence-Based Cyber Security: Suggestions and Recommendations for Building Cyber Resiliency against System Trespassing Events," was presented by David Maimon, UMD.

Adam Tagert spoke on the basics of the Science of Security program–how to move to a scientifically sound approach. "The goal," he said, "is to promote rigorous, generalizable, predictable, foundational, and replicable research into the Science of Security."

There were more than twenty student poster presentations.

Fall 2016 Quarterly: National Security Agency

The fall 2016 quarterly Science of Security (SoS) Lablet meeting was hosted, for the first time, by NSA in the Emerson V Auditorium on November 2-3, 2016. The agenda included panels on progress in the Science of Security and NSA research Hard Problems. The quarterly meeting also included the presentation ceremony for the 4th Best Scientific Cybersecurity Research Paper, presentations from the Lablets on their current research projects, and more than twenty poster presentations.

Dr. Deborah Frincke, Director of NSA's Research Directorate, hosted the event. This was the first quarterly meeting held at NSA, offering more agency participation and giving the researchers a better opportunity to see the sponsoring agency's interests in the Science of Security. Dr. Frincke noted the importance of rigorous science to the agency's mission, and that this importance is demonstrated by the appointment of a Science Advisor working with the Director. "Science," she said, "is a way to get the government, academia, and industry working together to make better, smarter decisions about security investment."

Adam Tagert provided a program overview and described progress to date. Researchers have produced more than 370 scientific papers dealing with the five Hard Problems of scalability and composability, predictive metrics, resilient architectures, policy-based governance, and human factors. The four Lablets now have 25 sub-Lablets and 106 research collaborators.

The panel on progress in the Science of Security was moderated by Brad Martin and consisted of the four Lablet Principal Investigators: Laurie Williams, NCSU; David Nicol, UIUC, Michel Cukier, UMD; and Bill Scherlis, CMU. Each lablet introduced its structure and progress to date.

NCSU's Laurie Williams described their structure as 15 faculty members from computer science, electrical engineering, psychology, and computer engineering departments, and 6 external collaborators. Their work addresses verification and revision of normative specifications of privacy, enumeration of potential misuse cases from requirements, attack surfaces, risk-based attack surface approximation, literature reviews and updates on intrusion detection, an automated synthesis of resilient configurations, information flow control in Android, SDN optimization, natural interactions for bot detection, and phishing patterns and users' personalities.

Michel Cukier, UMD, spoke of 20 faculty members from computer science and electrical and computer engineering, along with 5 external researchers focused on human behavior issues: understanding how users process security advice, and whether to trust the source or the context. They are applying experience in criminology to cybersecurity with regard to routine activities, rational choice, and

deterrence theory. Finally, they are looking at empirical models of vulnerabilities and attack surfaces.

David Nicol, UIUC, identified 20 faculty members from 9 universities as part of the UIUC lablet. Their projects include hypothesis testing for network security, data driven diversity models and model-based decision making, human circumvention of security, static-dynamic analysis of security metrics for CPS, and anonymous messaging. For large-scale problems, they are using factor graphs to deal with complexity and to develop predictive metrics. They have developed a testbed and library of attacks for resilient architectures that focus on attack detection that use rule-based machine learning for anomaly detection.

Carnegie Mellon has 15 faculty and 6 external collaborators according to Bill Scherlis. Their project matrix shows emphasis on scalability and composability and human factors research, and touches on the other Hard Problems. Their projects are looking at cognitive and psycho-social factors, as well as technologies, to determine factors that influence susceptibility to phishing. Their Security Behavior Observatory provides a way to observe actual security behavior by users.

Following these brief descriptions, the discussion became enthusiastically interactive with the audience. Panelists and the audience offered a range of comments about how to produce research that disrupts the adversary's ability to get in, stay in, and act within our systems and networks. The NSA panel moderator began the interactive session by asking about the impact of "science envy" and the areas where the Lablets have improved knowledge. Bill Scherlis offered the view that cybersecurity is so broad that it can't advance by advancing only one piece, that it requires a multi-disciplinary approach, and that "this war is being waged on a synthetic landscape." Science is growing and must continue to grow to deal with the complexity. David Nicol compared cybersecurity to biology and said our [cybersecurity] laws are not so immutable." As we learn, our "laws" change and we need validation of relationships. Michel Cukier suggested the need for a systems approach with the spotlight on metrics and the human side. Laurie Williams directly addressed the concept of "science envy" by defining it as the impulse to apply the right approach and model to a myriad of problems. The SoS Community is making progress in the Science of Security specifically because we are pushing the science model into cybersecurity research, and our work is improving as a result. We are gaining systematization of knowledge, doing better with predictability, and getting better over time. Bill Scherlis rejoined with an analogy to being "under the streetlamp"; we academics are seeing what we see under the direct light and now need to move on into the shadows to make greater gains.

An audience member asked if people would be willing to deal better with risk if they were aware of the risk. David Nicol stated that most users don't know what risk is. Bill Scherlis said that work at CMU's human Security Behavior Observatory confirms David Nicol's statement. Michel Cukier said it is a complex problem to understand how people understand security online and cited patches as an example of lack of knowledge and understanding. Bill Scherlis concluded that risk in cybersecurity is no longer actuarially tabulated.

Other issues that came up from audience questions included a discussion of human factors. One audience member observed that human behavior is challenging because when one measures human behaviors, people change those behaviors. In the past in physics, new theoretical models predicted new areas and approaches. Bill Scherlis said that model building is an important part of the Science of Security, that it is an aspirational thing to think about designing new systems. David Nicol concurred.

The panel moderator asked about other hard problem areas, including the supply chain, which the panel agreed is an important area of concern. Laurie Williams concluded the discussion with the observation that science heightens the need for accurate metrics. The SoS Community has a need for better security metrics, as well as performance metrics.

The agenda included a series of speakers about NSA research objectives. Speakers for the Cross Domain Services Management Office, Office of Computing and Analytic Research (R6), the Laboratory for Telecommunications Sciences, and the Trusted Research Division (R2) described the Hard Problems they face in mission fulfillment. A lively interactive discussion followed the presentations.

The Unified Cross Domain Services Management Office (UCDSMO) addressed the need to develop safe, secure, cross domain solutions that enable

secure transfer of intelligence on separate networks with different levels of classification. The UCDSMO is an interagency office between DoD and the intelligence community (IC) to coordinate development efforts. The goal for the program is to reach out to encourage innovative cross domain technologies and services that satisfy DoD and IC requirements. The hard problem is to safeguard data against intrusion by adversaries since the cross domain path is a major attack surface. Academe may be a source of ideas and approaches that can develop possible solutions. Currently, there is defense in depth, but every major weapons system has a cross domain issue that needs to be addressed. The adversary must be prevented from accessing the low side to be able to exfiltrate from the high side.

The representative from the Office of Computing and Analytic Research spoke about his office's role in protecting both information and signals. His challenge is to build very large-scale systems and make data consistent across those systems. For commodity computers, the challenge is to debug thousands of nodes and parallel systems and terabytes of data. Data processing analytics now need to move the analysis to the data in distributed systems and do remote evaluations. The analytics challenge is to fuse heterogeneous data together, analyze it, and put it into formats usable by decision makers. He is also looking for scalable approaches to visualization, natural language processing, modelling protection scenarios, and eliminating cognitive biases. The volume of information also suggests the need to develop tools that flag items of interest for analysts to look at more closely. Specific technologies and tools are judged by their mission impact.

The technical expert from the Laboratory for Telecommunications Sciences described his role as looking at networks, telecommunications, and computer research collaboratively in order to protect electronic signals generally. Specific research areas of interest include computer network operations, vulnerability exploitation and mitigation, network research including network mapping and measurement, the Internet, radio frequency, and high performance computing. Some key areas for collaborative research with industry and academe are quantum key distribution, the Internet of Things (IoT), blockchains, protocols, swarms, vehicles, and smart cities. The Internet of Things is a particular area of concern. For IoT, the questions are how much security is necessary or practical? How does one authenticate in the IoT? How do we approach data aggregation and storage in the IoT? Augmented reality and software defined networks completed his list of research concerns.

The Director of the Office of Trusted Research spoke on the five Hard Problems in the Science of Security and asked whether there is a sixth hard problem, and if so, what is it? He observed that no zero-day attack has been used in the largest hacks, that adversaries have been taking advantage of poorly secured systems. After they gain a foothold, they take advantage of poor security to elevate their privileges and exfiltrate data. The challenge is to ask why the current security landscape is failing. The audience described humans in the loop, system complexity, and asymmetry as the largest parts of the problem. Patch Tuesday is generating more than 5000 CPEs per year. This volume points up the software assurance problems related to lack of tools, lack of scalability, lack of sustainability and poor metrics–we are addressing small problems in isolated components instead of holistic processes and procedures. Given that determined adversaries cannot be stopped, another challenge is to control their behavior and deal with resilient systems that respond to attacks with well-understood effects, enhance the detection threshold, eliminate tradecraft, and improve metrics. The audience response to his question about a sixth hard problem varied– participants identified uncertainty, more and better tool development, and that metrics is the fundamental research problem.

Science of SecUre and REsilient Cyber-Physical Systems (SURE)

PI Xenofon Koutsoukos

The SURE project, led by Principal Investigator (PI) Xenofon Koutsoukos of Vanderbilt University, was initiated in 2014 to develop foundations and tools for designing, building, and assuring Cyber-Physical Systems (CPS) can maintain essential system properties in the presence of adversaries. CPS are devices that deeply integrate physical and computational components together, such as smart grids, traffic control, medical monitoring, and autonomous automobiles. SURE is the National Security Agency (NSA) Research Directorate-funded project aimed at improving scientific understanding of resiliency in cyber-physical systems. The project addresses the question of how to design systems that are resilient despite significant decentralization of resources and decision-making. An integral part of the research program includes a sustained effort to create a new generation of engineers who are comfortable with understanding, exploiting, and managing security and resilience in the context of integrated, computational, physical phenomena interacting with human designers and operators. The SURE research team is drawn from four universities: Vanderbilt; University of Hawaii; University of California, Berkeley; and the Massachusetts Institute of Technology.

VANDERBILT
UNIVERSITY®

SURE Research Thrusts

Hierarchical Coordination and Control

Cyber risk analysis and incentive design—aimed at developing regulations and strategies at the management level

Resilient monitoring and control of the networked control system infrastructure

Science of Decentralized Security
Aimed at developing a framework that will enable reasoning about the security of all the integrated constituent CPS components

Reliable and Practical Reasoning about Secure Computation and Communication in Networks
Aimed at contributing a formal framework for reasoning about security in CPS

Evaluation and Experimentation
Aimed at using modeling and simulation for the integration of cyber and physical platforms that directly interface with human decision-making

Education and Outreach
Aimed at educating the next generation of researchers in the field of security and resilience of CPS

In November 2016, researchers from the SURE project met with members of NSA's Trusted Research Directorate for their annual review. Accomplishments for this year included better tools for addressing CPS resilience and security and decentralized security; resilient dynamic adaptation for recovering from cyber attacks; integrative methods for resilient CPS operations under complex attacks; sequential games; improvements in the SURE testbed; and increased scalability.

Demonstrations of SURE's cloud-based testbed were presented. The first was decentralized security in adversarial settings. Péter Völgyesi presented "Adversary Models in Designing and Evaluating CPS Security and Resilience." Using a well-known CPS domain, traffic monitoring and control and existing simulation tools, his research project used machine learning to develop a stochastic map-level demand model, and it was able to detect unexpected patterns using Gaussian multivariate distributions. His work enabled him to study both integrity and distributed denial of service attacks on the transportation infrastructure using the Vanderbilt campus road system as a model.

The second demonstration illustrated the effects of tampering with traffic signal control's evaluation of vulnerability of transportation networks. The overall utility of the testbed is to provide experimental validation of attack models. Himanshu Neema's "C2WT: A Model-based and Scalable Integrated Simulation Testbed for Science of Security," described the architecture and use of the Command-and-Control Wind Tunnel (C2WT), a model-based, scalable integrated testbed. Based on C2WT simulation models, configuration files are generated for the different simulation components that configure how the component is connected to the simulation using data flow, timing, and parameters. Federates have a common data model to be able to share data. The data model can be imported from domain-specific models and domain-specific models can be generated from data models. A library of models has been developed that includes DDoS and integrity attacks, delays, data corruption, and network manipulation.

Additional research projects on resiliency were presented covering both behavioral and technical subjects. Yevgeniy Vorobeychik covered "Resilient Closed Loop Control in Adversarial Settings." Aron Laszka presented "Detection of Cyber-Attacks Against Traffic Control Integrity." Bradley Potteiger described "A Hardware in the Loop Testbed for Evaluating and Measuring Security and Resilience in CPS." Other participants included Professor Janos Sztipanovits, government representatives from the Air Force Research Labs, and representatives from the sponsoring agency.

In their 2015 annual review, SURE participants demonstrated the project developed cloud-based testbed for evaluating and measuring resiliency through modeling. The 2015 review also included presentations on research products and a presentation on the next generation CPS-VO platform. Five additional research projects on resiliency were presented. These covered both behavioral and technical subjects including the use of inversion to train robust machine learning models, reasoning about security in cyber-physical systems, resilience in the wake of disruptions, hierarchical control of incident management, and decentralized security in adversarial settings.

VANDERBILT UNIVERSITY®

PUBLICATIONS

- Himanshu Neema et al., "Demo: Abstract: SURE: An Experimentation and Evaluation Testbed for CPS Security and Resilience," *ACM/IEEE 7th International Conference on Cyber-Physical Systems (ICCPS 2016)*, Vienna, Austria, April 12-14, 2016, pp.1-1.

- Aron Laszka, Bradley Potteiger, Yevgeniy Vorobeychik, Saurabh Amin, and Xenofon Koutsoukos, "Vulnerability of Transportation Networks to Traffic-Signal Tampering," *ACM/IEEE 7th International Conference on Cyber-Physical Systems (ICCPS 2016)*, Vienna, Austria, April 12-14, 2016, pp. 1-10.

- Aron Laszka, Waseem Abbas, S. Shankar Sastry, Yevgeniy Vorobeychik, and Xenofon Koutsoukos, "Optimal Thresholds for Intrusion Detection Systems," in Proceedings of the Symposium and Bootcamp on the Science of Security (HoTSoS '16), Pittsburgh, PA, April 19-21, 2016, ACM Press, New York, NY, 2016, pp. 72-81.

- Jian Lou and Yevgeniy Vorobeychik, "Decentralization and Security in Dynamic Traffic Light Control," in *Proceedings of the Symposium and Bootcamp on the Science of Security (HoTSoS '16)*, Pittsburgh, PA, April 19-21, 2016, ACM Press, New York, NY, 2016, pp. 90-92.

- Bradley Potteiger, Goncalo Martins, and Xenofon Koutsoukos, "Software and Attack Centric Integrated Threat Modeling for Quantitative Risk Assessment," in *Proceedings of the Symposium and Bootcamp on the Science of Security (HoTSoS '16)*, Pittsburgh, PA, April 19-21, 2016, ACM Press, New York, NY, 2016, pp. 99-108.

- Yue Yin, Yevgeniy Vorobeychik, Bo An, and Noam Hazon, "Optimally Protecting Elections," in *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence (IJCAI-16)*, New York, NY, July 9-15, 2016, pp. 538-545.

- Amin Ghafouri, Waseem Abbas, Yevgeniy Vorobeychik, and Xenofon Koutsoukos, "Vulnerability of Fixed-Time Control of Signalized Intersections to Cyber-Tampering," in *Proceedings of the 9th International Symposium on Resilient Control Systems*, Chicago, IL, Aug. 16-18, 2016.

- Xenofon Koutsoukos et al., "Performance Evaluation of Secure Industrial Control System Design: A Railway Control System Case Study," in 2016 *Resilience Week (RWS)*, Chicago, IL, 2016, pp. 101-108.

- Goncalo Martins, Arul Moondra, Abhishek Dubey, Anirban Bhattacharjee, and Xenofon Koutsoukos, "Computation and Communication Evaluation of an Authentication Mechanism for Time-Triggered Networked Control Systems," *Sensors, Special Issue on Real-Time and Cyber-Physical Systems*, Sep. 2016, vol. 16, no. 8, p. 1166.

- Lina Sela Perelman, Waseem Abbas, Xenofon Koutsoukos, and Saurabh Amin, "Sensor placement for fault location identification in water networks: A minimum test cover approach," *Automatica*, vol. 72, pp. 166-176, October 2016.

- Amin Ghafouri, Waseem Abbas, Aron Laszka, Yevgeniy Vorobeychik, and Xenofon Koutsoukos, "Optimal Thresholds for Anomaly-Based Intrusion Detection in Dynamical Environments," *2016 Conference on Decision and Game Theory for Security (GameSec 2016)*, New York, NY, Nov. 2-4, 2016.

- Mingyi Zhao, Aron Laszka, Thomas Maillart, and Jens Grossklags, "Crowdsourced Security Vulnerability Discovery: Modeling and Organizing Bug-Bounty Programs," *HCOMP16: Workshop on Mathematical Foundations of Human Computation*, Austin, TX, November 3, 2016.

VANDERBILT
UNIVERSITY®

- Bo Li, Yining Wang, Aarti Singh, and Yevgeniy Vorobeychik, "Data Poisoning Attacks on Factorization-Based Collaborative Filtering," in *Proceedings of the 30th Annual Conference on Neural Information Processing Systems (NIPS 2016)*, Barcelona, Spain, 2016.

- Bo Li, Yevgeniy Vorobeychik, Muqun Li, and Bradley Malin, "Scalable Iterative Classification for Sanitizing Large-Scale Datasets," in *IEEE Transactions on Knowledge and Data Engineering*, 2017, vol., 29, no. 3, pp. 698-711.

- Aron Laszka, Yevgeniy Vorobeychik, and Xenofon Koutsoukos, "A Game-Theoretic Approach for Integrity Assurance in Resource-Bounded Systems," *International Journal of Information Security*, Springer, 2017.

- Nika Haghtalab, Aron Laszka, Ariel D. Procaccia, Yevgeniy Vorobeychik, and Xenofon Koutsoukos, "Monitoring Stealthy Diffusions," *Knowledge and Information Systems*, Springer, 2017, to appear.

- Jian Lou, Andrew M. Smith, and Yevgeniy Vorobeychik, "Multidefender Security Games," in *IEEE Intelligent Systems*, vol. 32, no. 1, pp. 50-60, Jan.-Feb. 2017.

- Aron Laszka, Waseem Abbas, and Xenofon Koutsoukos, "Scheduling Battery-Powered Sensor Networks for Minimizing Detection Delays," *IEEE Communications Letters*, 2017, to appear.

- Jiarui Gan, Bo An, Yevgeniy Vorobeychik, and Brian Gauch, "Security Games on a Plane," *AAAI Conference on Artificial Intelligence (AAAI '17)*, March 2017, to appear.

- Bo Li, Kevin Roundy, Chris Gates, and Yevgeniy Vorobeychik, "Large-Scale Identification of Malicious Singleton Files," *7th ACM Conference on Data and Application Security (CODASPY 2017)*, Scottsdale, AZ, March 22-24, 2017, to appear.

- Andrew M. Smith, Jackson Mayo, Vivian Kammler, Robert C. Armstrong, and Yevgeniy Vorobeychik, "Using Computational Game Theory to Guide Verification and Security in Hardware Designs," *IEEE International Symposium on Hardware Oriented Security and Trust (HOST 2017)*, McLean, VA, May 2017, to appear.

- Waseem Abbas, Aron Laszka, Yevgeniy Vorobeychik, and Xenofon Koutsoukos, "Improving Network Connectivity Using Trusted Nodes and Edges," *2017 American Control Conference (ACC)*, Seattle, WA, May 2017, to appear.

VANDERBILT
UNIVERSITY®


Best
Scientific


RESEARCH DIRECTORATE
ADVANCING INTELLIGENCE THROUGH SCIENCE
NSA Research Directorate


To Register: Go SoS Wiki
4TH ANNUAL
BEST SCIENTIFIC CYBERSECURITY PAPER
COMPETITION
Winning Paper:
Nomad: Mitigating Arbitrary Cloud Side Channels via Provider-Assisted Migration
by Soo-Jin Moon, Vyas Sekar and Michael Reiter
from Carnegie Mellon University and University of North Carolina.
Recognition Ceremony:
Emerson 5 Auditorium
November 2nd at 1330


NATIONAL SECURITY AGENCY
UNITED STATES OF AMERICA

Section 2

Promoting Rigorous Scientific Principles

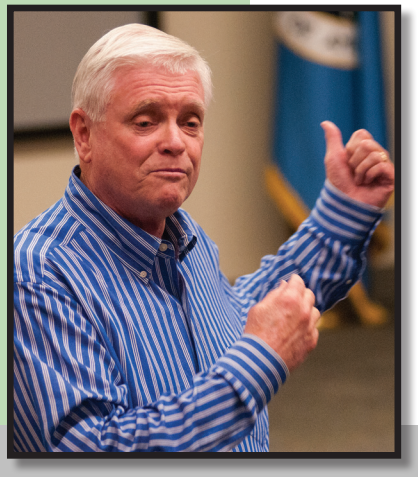

The Science of Security program promotes rigorous scientific principles through an annual Best Scientific Cybersecurity Paper Competition, and through sponsorship of an award at the Intel International Science and Engineering Fair (Intel ISEF). At this year's paper competition, the NSA Research Directorate selected a winner and two honorable mentions from over 50 nominations. At the Intel ISEF, there were approximately 100 cybersecurity-related projects, 24 of which were selected for further NSA review. The winner and two runners-up received cash awards and visited NSA to meet with cybersecurity researchers. Details on the paper competition and the Intel ISEF follow.

ANNUAL BEST CYBERSECURITY

The Best Scientific Cybersecurity Paper Competition is sponsored annually by NSA's Research Directorate and reflects the Agency's desire to increase scientific rigor in the cybersecurity field. The competition was established to recognize current research that exemplifies the development of scientific rigor in cybersecurity research. Science of Security (SoS) is a broad enterprise, involving both theoretical and empirical work across a diverse set of topics. While there can only be one best paper, no single paper can span the full breadth of SoS topics. Nevertheless, work in all facets of security science is both needed and encouraged.

The fourth Annual NSA Competition for Best Scientific Cybersecurity Paper recognized the best scientific cybersecurity paper published in 2015 from among 54 nominations. Papers were solicited for nomination via announcements on the NSA home page and SoS web page, in cooperation with the four SoS Lablets, and informally throughout the SoS community. The paper competition was also advertised in the *IEEE Security & Privacy Magazine*, online and print copies, and in conjunction with NSA recruiting efforts. Papers published in journals, peer-reviewed magazines, or technical conferences were eligible for nomination. Nomination statements described the scientific contribution of the paper and explained why this paper merited the award. The nominator could not be an author or co-author of the nominated paper. Papers were reviewed by a panel of Distinguished Experts as well as by NSA personnel. Over the four years of the competition, almost 180 papers have been nominated with four winners and seven honorable mentions selected. Past winning papers were "The Science of Guessing: Analyzing an Anonymized Corpus of 70 Million Passwords," "Memory Trace Oblivious Program Execution," and "Additive and Multiplicative Notions of Leakage and Their Capacities." Winners for all four years of the competition were announced on the NSA homepage.

The winning paper for 2015 was "Nomad: Mitigating Arbitrary Cloud Side Channels via Provider-Assisted Migration." It was written by Soo-Jin Moon, Vyas Sekar, and Michael Reiter from Carnegie Mellon University and the University of North Carolina at Chapel Hill, and was presented at the Association for Computer Machinery (ACM) Conference on Computer and Communications Security (CCS 15).

The paper discusses a new system called Nomad that demonstrates a general and immediate defense against side-channel attacks as opposed to the current method of developing attack-specific fixes. This side-channel defense applies to attacks that come from another co-resident virtual machine. Conceptually, computers can simulate the appearance of multiple physical computers but in reality be just one computer. These simulated computers should be completely separated and act as if they were two physically different devices, but side-channel attacks break this separation. This is particularly relevant to cloud service providers where each virtual machine could be under the control of different people.

Nomad was selected as the winning paper because it provided several scientific advances and then tested its conclusions. It developed a threat model and information leakage model. It also developed and tested several algorithms for how to move around the virtual machines within the cloud to reach its goals. In summary, the paper's proposed defense is very simple and yet very powerful. It addresses a relevant problem and provides a pragmatic solution with the science to support it.

The first honorable mention paper was "Quantum-Secure Covert Communication on Bosonic Channels." This paper in *Nature Communications* was written by a team of researchers at the University of Massachusetts, Amherst, and Raytheon BBN Technologies: Boulat A. Bash, Andrei H. Gheorghe, Monika Patel, Jonathan L. Habif, Dennis Goeckel, Don Towsley, and Saikat Guha.

This paper explored the limits of how much information can be transmitted on a Bosonic Channel where an attacker cannot determine that the transmission has even occurred. The paper presented quantum communications as well as real world examples and proofs of concept. The paper received an honorable mention recognition because it was a strong and impactful paper with fresh ideas. Most importantly, the paper is being held up as an example of how effective scientific exposition should be

PAPER COMPETITION

organized in scientific and engineering disciplines that employ heavy mathematical analysis.

The second honorable mention paper was "Increasing Cybersecurity Investments in Private Sector Firms." It was written by a research group at the University of Maryland, College Park by team members Lawrence Gordon, Martin Loeb, William Lucyshyn, and Lei Zhou. It was published in a new journal on cross discipline cybersecurity issues, Journal of Cybersecurity.

The paper developed an economics-based framework for evaluating governmental approaches to increase private sector investment in cybersecurity. Increased cybersecurity investment is needed because of the cost of externalities, the effects on others, and this paper helps informs policy makers on the impact of policies. This paper was chosen for recognition as it has meticulous methodology and insights that can be applied in the real world to improve security.
A ceremony recognizing the winners was held on 2 November. Following a luncheon for the winners and reviewers hosted by Dr. Deborah Frincke, NSA's Director of Research, the winners received plaques noting their award and then presented their papers to attendees. Attendees included personnel from across NSA and from the SoS Lablets.

NSA Competition Leads

Dr. Deborah Frincke – Director of Research, NSA

Dr. Adam Tagert – Science of Security Technical Director, NSA Information Assurance Research

Distinguished Expert Reviewers

Dr. Whitfield Diffie – Cybersecurity Advisor

Dr. Daniel Earl Geer Jr., Sc.D. – Chief Information Security Officer at In-Q-Tel, Inc.

Dr. John D. McLean – Superintendent of the Naval Research Laboratory's Information Technology Division

Professor M. Angela Sasse – Professor of Human-Centered Technology and Head of Information Security Research, Department of Computer Science at University College London, UK

Professor Fred B. Schneider – Samuel B. Eckert Professor of Computer Science at Cornell University

Mr. Philip J. Venables – Chief Information Risk Officer at Goldman Sachs

Professor David A. Wagner – Professor, Department of Electrical Engineering and Computer Sciences, Computer Science Division at the University of California, Berkeley

Dr. Jeannette Wing – Vice President, Head of Microsoft Research International

NSA Reviewers

The papers were also reviewed by a team of experts drawn from various backgrounds across NSA including the Research Directorate, the former National Threat Operations Center, and the former Information Assurance Directorate

Intel International Science and Engineering Fair

For the second consecutive year, the NSA Research Directorate's Science of Security Initiative sponsored special awards at the Intel International Science and Engineering Fair (Intel ISEF). Intel ISEF, held annually, is the world's largest high school science fair with approximately 1,700 entrants from more than 75 countries, regions, and territories. The fair has 22 categories including behavioral and social sciences, biochemistry, embedded systems, engineering mechanics, mathematics, robotics and intelligent machines, and systems software. There are 600 awarded finalists and approximately $4 million worth of prizes.

Of approximately 100 projects that were relevant to cybersecurity at Intel ISEF 2016, 24 were chosen for interviews with NSA award judges. The Research Directorate-sponsored awards were judged on their demonstrated advancement in the science to secure and safeguard cyberspace. NSA's participation is part of an effort to encourage more high school students to pursue cybersecurity education, research, and careers. The results of both years' competitions were featured on NSA's web page (www.nsa.gov).

The winner received a $3,000 first prize award, and there were two $1,000 runner-up prizes. In addition to the monetary prizes, NSA recognized the winners with a two-day visit to NSA headquarters. During their visit, the awardees presented their work to NSA researchers, received overviews of the Agency's mission, and toured various operations centers and the National Cryptologic Museum.

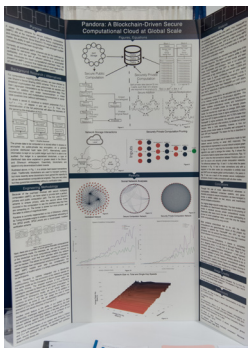

First Place was awarded to **Charles Noyes**, 17, of Villa Park, California, for his project, "Efficient Blockchain-Driven Multiparty Computation Markets at Scale." His project addressed a long-standing search within the intersecting fields of computer science, cryptography, and game theory for faster, more efficient, secure multiparty computation (sMPC). Noyes developed a novel scheme that combines blockchains (data structures that serve as the backbone for secure banking applications such as Bitcoin) with homomorphic computations (which allow computations on encrypted data such as search) and verification schemes (which guarantee authenticity). On visiting the Agency, he noted, "It has been fantastic. Throughout the entire visit, I was surprised by the things being done at the Agency. I thought it was exclusively digital defense, but I didn't realize there was 3-D printing or carpentry."

Rucha Joshi, 16, of Austin, Texas, received a runner-up award for her project, "Determining Network Robustness Using Region Based Connectivity." Joshi developed a method to test a network's resilience using region-based connectivity rather than the traditional node and edge connectivity. Typically, in real world scenarios, a node failure is due to a regional issue, such as natural or man-made disasters, and not just a single point. Her research was inspired by her aunt's 20-hour drive from Houston to Austin before Hurricane Katrina. "I wanted to find alternate paths so something like that doesn't happen again." Although Joshi has been entering science fairs for years, she only knew about NSA from what she learned in school. "I was excited to learn that I won the award from NSA. It has been a great two days," she said.

Karthik Yegnesh, 16, of Eagleville, Pennsylvania, received a runner-up award for his project, "Cosheaf Theoretical Constructions in Networks and Persistent Homology." He applied persistent homology, an algebraic method for measuring the topological feature of shapes and functions, to analyze the data flow in financial, social, and biological networks. Yegnesh's project could facilitate smooth transmission of data. After touring the agency and meeting employees, he admitted that he didn't know much about NSA. "I was so impressed by the goals of the organization. And how human it was. Regular people work here with regular desks."


SCIENCE OF SECURITY
Hard Problems
METRICS
POLICY
HUMAN BEHAVIOR
RESILIENCY
SCALABILITY


Phipps Conservatory and Botanical Gardens
Where Beauty Comes Naturally


Carnegie Mellon

Section 3

Growing the Science of Security

The Science of Security program's premier event, the 2016 3rd Annual Hot Topics in Science of Security (HoTSoS) Symposium and Bootcamp, attracted participants from academia, industry, and government for the two-day workshop held at Carnegie Mellon University. HoTSoS provides an opportunity for researchers to exchange approaches for applying rigorous scientific principles and methods to cybersecurity challenges. The Science of Security Virtual Organization (SoS-VO), with over 1100 members, encourages the exchange of information and provides a centralized location for cybersecurity research, upcoming events, and news. The Science of Security program leverages the lablet research, the SURE program, the Best Cybersecurity Paper Competition, the Intel ISEF awards, HoTSoS, and the SoS-VO to engage other researchers across academia, industry, and government, both nationally and internationally, to expand the Science of Security community. Details on HoTSoS, the SoS-VO, and other outreach efforts are found in the following section.

HoTSoS 2016

HoTSoS, Science of Security's annual symposium and bootcamp, convenes a broad research community focused on addressing the fundamental problems of security in a principled manner. Co-sponsored by the Association for Computer Machinery (ACM) and the National Security Agency (NSA), the symposium brings together researchers from multiple academic fields for presentations demonstrating methodical, rigorous, and scientific approaches to identify, prevent, and remove cyber threats. Details about HoTSoS can be found at http://cps-vo.org/group/HoTSoS.

The 2016 Symposium and Bootcamp on the Science of Security (HoTSoS), the fourth annual event, was held April 19-21 in Pittsburgh, Pennsylvania at Carnegie Mellon University. A major focus of the symposium was on the advancement of scientific methods, including data gathering and analysis; experimental methods; and mathematical models for modeling and reasoning. The HoTSoS 2016 Proceedings have been published by ACM and are available online at the ACM Digital Library at: http://dl.acm.org/citation.cfm?id=2898375&picked=prox.

HoTSoS 2016 was co-chaired by William L. Scherlis and David Brumley of Carnegie Mellon University. The program committee was made up of the following individuals, representing Science of Security (SoS) Lablets, subLablets, and NSA:

Marwan Abi-Antoun, Carnegie Mellon University
Alessandro Acquisti, Carnegie Mellon University
Jonathan Aldrich, Carnegie Mellon University
Michael Bailey, University of Illinois at Urbana-Champaign
Lujo Bauer, Carnegie Mellon University
Travis Breaux, Carnegie Mellon University
Kathleen Carley, Carnegie Mellon University
Michel Cukier, University of Maryland
Serge Egelman, University of California, Berkeley
David Garlan, Carnegie Mellon University
Limin Jia, Carnegie Mellon University
Jonathan Katz, University of Maryland
Michael Maass, Carnegie Mellon University
Stuart Krohn, National Security Agency
Sam Malek, University of California, Irvine
David Malcolm Nicol, University of Illinois at Urbana-Champaign
André Platzer, Carnegie Mellon University
William H. Sanders, University of Illinois at Urbana-Champaign
Munindar P. Singh, North Carolina State University
Witawas Srisa-An, University of Nebraska-Lincoln
Joshua Sunshine, Carnegie Mellon University
Adam Tagert, National Security Agency

Conference co-chair William Scherlis, co-Principal Investigator for the Carnegie Mellon University Lablet, welcomed the group by reiterating the opportunity for researchers to share their ideas about how better to address cybersecurity challenges using strong scientific principles and rigorous methods.

Four outside speakers addressed Science of Security from the perspectives of consumer interest, government policy, and industry, and the value of a large graph method of analysis. Research papers, presentations, tutorials, and poster sessions rounded out the agenda.

Research Papers

Nine research papers covering a range of issues related to the five Hard Problems were presented, including studies about intrusion detection, threat modeling, anomaly detection, and attack variants.

1.Safety-critical Cyber-physical Attacks: Analysis, Detection, and Mitigation
Hui Lin, Homa Alemzadeh, Daniel Chen, Zbigniew Kalbarczyk, and Ravishankar Iyer

Today's Cyber-Physical Systems (CPSs) can have very different characteristics in terms of control algorithms, configurations, underlying infrastructure, communication protocols, and real-time requirements. Despite these variations, they all face the threat of malicious attacks that exploit the vulnerabilities in the cyber domain as footholds to introduce safety violations in the physical processes. In this paper, we focus on a class of attacks that impact the physical processes without introducing anomalies in the cyber domain. We present the common challenges in detecting this type of attacks in the contexts of two very different CPSs (i.e., power grids and surgical robots). In addition, we present a general principle for detecting such cyber-physical attacks, which combine the knowledge of both cyber and physical domains to estimate the adverse consequences of malicious activities in a timely manner.

2.Optimal Thresholds for Intrusion Detection Systems
Aron Laszka, Waseem Abbas, S. Shankar Sastry, Yevgeniy Vorobeychik, and Xenofon Koutsoukos

Intrusion-detection systems can play a key role in protecting sensitive computer systems since they give defenders a chance to detect and mitigate attacks before they could cause substantial losses. However, an oversensitive intrusion-detection system, which produces a large number of false alarms, imposes prohibitively high operational costs on a defender since alarms need to be manually investigated. Thus, defenders have to strike the right balance between maximizing security and minimizing costs. Optimizing the sensitivity of intrusion detection systems is especially challenging in the case when multiple interdependent computer systems have to be defended against a strategic attacker, who can target computer systems in order to maximize losses and minimize the probability of detection. We model this scenario as an attacker-defender security game and study the problem.

3.Intrusion Detection in Enterprise Systems by Combining and Clustering Diverse Monitor Data
Atul Bohara, Uttam Thakore (presenters, shown), and William Sanders

Intrusion detection using multiple security devices has received much attention recently. The large volume of information generated by these tools, however, increases the burden on both computing resources and security administrators. Moreover, attack detection does not improve as expected if these tools work without any coordination.
In this work, we propose a simple method to join information generated by security monitors with diverse data formats. We present a novel, intrusion detection technique that uses unsupervised clustering algorithms to identify malicious behavior within large volumes of diverse security monitor data. First, we extract a set of features from network-level and host-level security logs that aid in detecting malicious host behavior and flooding-based network attacks in an enterprise network system. We then apply clustering algorithms to the separate and joined logs and use statistical tools to identify anomalous usage behaviors captured by the logs. We evaluate our approach on an enterprise network data set, which contains network and host activity logs.

Our approach correctly identifies and prioritizes anomalous behaviors in the logs by their likelihood of maliciousness. By combining network and host logs, we are able to detect malicious behavior that cannot be detected by either log alone.

4.Establishing a Baseline for Measuring Advancement in the Science of Security – an Analysis of the 2015 IEEE Security & Privacy Proceedings
Jeffrey Carver, Morgan Burcham, Sedef Akinli Kocak, Ayse Bener, Michael Felderer, Matthias Gander, Jason King, Jouni Markkula, Markku Oivo, Clemens Sauerwein, and Laurie Williams

In this paper we aim to establish a baseline of the state of scientific work in security through the analysis of indicators of scientific research as reported in the papers from the 2015 IEEE Symposium on Security and Privacy.

To conduct this analysis, we developed a series of rubrics to determine the completeness of the papers relative to the type of evaluation used (e.g., case study, experiment, proof). Our findings showed that while papers are generally easy to read, they often do not explicitly document some key information like the research objectives, the process for choosing the cases to include in the studies, and the threats to validity. We hope that this initial analysis will serve as a baseline against which we can measure the advancement of the science of security.

5.Software and Attack Centric Integrated Threat Modeling for Quantitative Risk Assessment
Bradley Potteiger, Goncalo Martins, and Xenofon Koutsoukos

Threat modeling involves understanding the complexity of the system and identifying all of the possible threats, regardless of whether or not they can be exploited. Proper identification of threats and appropriate selection of countermeasures reduces the ability of attackers to misuse the system.

This paper presents a quantitative, integrated threat modeling approach that merges software and attack centric threat modeling techniques. The threat model is composed of a system model representing the physical and network infrastructure layout, as well as a component model illustrating component specific threats. Component attack trees allow for modeling specific component contained attack vectors, while system attack graphs illustrate multi-component, multi-step attack vectors across the system. The Common Vulnerability Scoring System (CVSS) is leveraged to provide a standardized method of quantifying the low level vulnerabilities in the attack trees. As a case study, a railway communication network is used, and the respective results using a threat modeling software tool are presented.

6.Security is about Control: Insights from Cybernetics
Antonio Roque, Kevin Bush, and Christopher Degni

Cybernetic closed loop regulators are used to model sociotechnical systems in adversarial contexts. Cybernetic principles regarding these idealized control loops are applied to show how the incompleteness of system models enables system exploitation. We consider abstractions as a case study of model incompleteness, and we characterize the ways that attackers and defenders interact in such a formalism. We end by arguing that the science of security is most like a military science, whose foundations are analytical and generative rather than normative.

7.Controller Synthesis for Linear Dynamical Systems with Adversaries
Zhenqi Huang, Yu Wang, Sayan Mitra, and Geir Dullerud

We present a controller synthesis algorithm for a reach-avoid problem in the presence of adversaries. Our model of the adversary abstractly captures typical malicious attacks envisioned on cyber-physical systems such as sensor spoofing, controller corruption, and actuator intrusion. After formulating the problem in a general setting, we present a sound and complete algorithm for the case with linear dynamics and an adversary with a budget on the total L2-norm of its actions. The algorithm relies on a result from linear control theory that enables us to decompose and compute the reachable states of the system in terms of a symbolic simulation of the adversary-free dynamics and the total uncertainty induced by the adversary. With this decomposition, the synthesis problem eliminates the universal

quantifier on the adversary's choices and the symbolic controller actions can be effectively solved using an SMT solver. The constraints induced by the adversary are computed by solving second-order cone programmings. The algorithm is later extended to synthesize state-dependent controller and to generate attacks for the adversary. We present preliminary experimental results that show the effectiveness of this approach on several example problems.

8.Differences in Trust between Human and Automated Decision Aids
Carl Pearson, Allaire Welk, William Boettcher, Roger Mayer, Sean Streck, Joseph Simons-Rudolph, and Christopher Mayhorn

Humans can easily find themselves in high cost situations where they must choose between suggestions made by an automated decision aid and a conflicting human decision aid. Previous research indicates that humans often rely on automation or other humans, but not both simultaneously. Expanding on previous work conducted by Lyons and Stokes (2012), the current experiment measures how trust in automated or human decision aids differs along with perceived risk and workload. The simulated task required 126 participants to choose the safest route for a military convoy; they were presented with conflicting information from an automated tool and a human. Results demonstrated that as workload increased, trust in automation decreased. As the perceived risk increased, trust in the human decision aid increased. Individual differences in dispositional trust correlated with an increased trust in both decision aids. These findings can be used to inform training programs for operators who may receive information from human and automated sources. Examples of this context include air traffic control, aviation, and signals intelligence.

9.A Framework for Generation, Replay and Analysis of Real-World Attack Variants
Phuong Cao, Eric Badger, Zbigniew Kalbarczyk, and Ravishankar Iyer

This paper presents a framework for (1) generating variants of known attacks, (2) replaying attack variants in an isolated environment and, (3) validating detection capabilities of attack detection techniques against the variants. Our framework facilitates reproducible security experiments. We generated 648 variants of three real-world attacks (observed at the National Center for Supercomputing Applications at the University of Illinois). Our experiment showed the value of generating attack variants by quantifying the detection capabilities of three detection methods: a signature-based detection technique, an anomaly-based detection technique, and a probablistic graphical model-based technique.

Keynotes

1.Adventures in Usable Privacy and Security: From Empirical Studies to Public Policy.
Lorrie Cranor

Lorrie Cranor, Carnegie Mellon University professor and SoS lablet project PI on loan to the Federal Trade Commission (FTC), addressed "Adventures in Usable Privacy and Security: From Empirical Studies to Public Policy." Topics of interest to FTC are quantifying privacy interests, disclosures, financial technologies, attack trends, improving complaint reporting, tools to automate tracking, targeted advertising, cross device tracking, fraud, and emerging scams. "Good warnings help users determine whether they are at risk," she said, but "people ignore poor warnings that put the onus of calculating risk on the end user." She described a study that shows that password expiry is counterproductive.

2.Trustworthy Cyberspace: Strategic Plan for the Federal Cybersecurity Research & Development Program
Greg Shannon

Greg Shannon, Carnegie Mellon University professor currently working at the White House Office of Science and Technology (OSTP), spoke on the science challenges in "Trustworthy Cyberspace: Strategic Plan for the Federal Cybersecurity Research & Development Program." The strategic plan, originally issued in December 2011, has been updated and expanded. The revised plan, the "Federal Cybersecurity Research and Development Strategic Plan," was released in February 2016. In the plan, OSTP is looking at the science and technology issues that affect policy and the policy issues that impact technology. Keys to addressing the problems include creating a Commission

on Enhancing National Cybersecurity, appointing a Federal Chief Information Security Officer to take the lead on policies, oversight, and strategy; budgeting a $3.1 billion IT modernization fund; and working with industry to encourage broader use of security tools such as multi-factor authentication. The plan's goals are to counter and reverse adversaries' asymmetrical advantages; achieve science and technology advantages to achieve effective deterrence; enable the cybersecurity research, development, and operations community to quickly design, develop, deploy, and operate effective new cybersecurity technologies and services; minimize and simplify cybersecurity tasks for users; and deter adversaries from launching malicious cyber activities.

3.Anomaly Detection in Large Graphs
Christos Faloutsos

Christos Faloutsos, Professor of Electrical and Computer Engineering at Carnegie Mellon University, spoke on "Anomaly Detection in Large Graphs." Citing recent research from several projects, he demonstrated the value of using graph theory to identify patterns that are otherwise hidden. Motivating his research are the problems of identifying such patterns for fraud detection and patterns in time-evolving graphs/tensors. He uses Hadoop to search for many clusters in parallel, starting with random seeds, updates sets of pages and like times for each cluster, and then repeats until convergence is achieved. This approach has been deployed at Facebook as CopyCatch where it was determined that most clusters (77%) come from hard to detect real-but-compromised users and that easier to detect fake accounts are only 22%.
For eBay fraud detection, the technique has been used on the non-delivery scam. Using NetProbe allows detection of the scam by identifying the groups of people who are cross rating each other to appear honest. The fraudulent nodes look trustworthy.His conclusions: patterns and anomalies go hand in hand; large data sets reveal patterns and outliers that are otherwise invisible.

4.A View from the Front Lines with M-Trends
Matthew Briggs

FireEye Inc.'s Matt Briggs gave "A View from the Front Lines with M-Trends." Continuing trends in corporate cybersecurity show that spear phishing remains the most common entry point. Most hacks are leveraging trust relationships to get in, especially leveraging IT outsourcing.New trends are described as "David v. Goliath": the rise of business disruption attacks that are politically and/or financially motivated and lead to data leaks to embarrass the company and destroy critical systems. He noted that in the last five years the median number of days before discovery has been reduced from 416 days to to 146.

Posters

The fifteen posters presented at HoTSoS 2016 addressed all five Hard Problems and represented contributions from all four of the Lablets and multiple subLablets as well. Many related directly to fundamental research projects while others did not tie directly to SoS lablet projects. "Characterizing Cybersecurity Jobs," for example, introduced a way to address cybersecurity personnel requirements; "Modules in Wyvern" provided updates to work on a new programming language; and "Decentralization and Security in Dynamic Traffic Light Control" addressed work ongoing under the SURE program.

1.Abstract Runtime Structure for Reasoning about Security
Marwan Abi-Antoun, Ebrahim Khalaj, Radu Vanciu, and Ahmad Moghim

We propose an interactive approach where analysts reason about the security of a system using an abstraction of its runtime structure, as opposed to looking at the code. They interactively refine a hierarchical object graph, set security properties on abstract objects or edges, query the graph, and investigate the results by studying highlighted objects or edges or tracing to the code. Behind the scenes, an inference analysis and an extraction analysis maintain the soundness of the graph with respect to the code.

2.Efficient Solving of String Constraints for Security Analysis
Clark Barrett, Cesare Tinelli, Morgan Deters, Tianyi Liang, Andrew Reynolds, and Nestan Tsiskaridze

The security of software is increasingly more critical for consumer confidence, protection of privacy, protection of intellectual property, and even national security. As threats to software security have become more sophisticated, so too have the techniques developed to ensure security. One basic technique that has become a fundamental tool in static security analysis is symbolic execution. There are now a number of successful approaches that rely on symbolic methods to reduce security questions about programs to constraint satisfaction problems in some formal logic (e.g., [4, 5, 7, 16]). Those problems are then solved automatically by specialized reasoners for the target logic. The found solutions are then used to construct automatically security exploits in the original programs or, more generally, identify security vulnerabilities.

3.Characterizing Cybersecurity Jobs: Applying the Cyber Aptitude and Talent Assessment Framework
Susan G. Campbell, Lelyn D. Saner, and Michael F. Bunting

Characterizing what makes cybersecurity professions difficult involves several components, including specifying the cognitive and functional requirements for performing job-related tasks. Many frameworks that have been proposed are focused on functional requirements of cyber work roles, including the knowledge, skills, and abilities associated with them. In contrast, we have proposed a framework for classifying cybersecurity jobs according to the cognitive demands of each job and for matching applicants to jobs based on their aptitudes for key cognitive skills (e.g., responding to network activity in real-time). In this phase of research, we are investigating several cybersecurity jobs (such as operators vs. analysts), converting the high-level functional tasks of each job into elementary tasks, in order to determine what cognitive requirements distinguish the jobs. We will then examine how the models of cognitive demands by job can be used to inform the designs of aptitude tests for different kinds of jobs. In this poster, we will describe our framework in more detail and how it can be applied toward matching people with the jobs that fit them best.

4.Expressing and Reasoning about Conflicting Norms in Cybersecurity
Jiaming Jiang, Nirav Ajmeri, Rada Y. Chirkova, Jon Doyle, and Munindar P. Singh

Secure collaboration requires the collaborating parties to apply the right policies for their interaction. We adopt a notion of conditional, directed norms as a way to capture the standards of correctness for a collaboration. How can we handle conflicting norms? We describe an approach based on knowledge of what norm dominates what norm in what situation. Our approach adapts answer-set programming to compute stable sets of norms with respect to their computed conflicts and dominance. It assesses agent compliance with respect to those stable sets. We demonstrate our approach on a healthcare scenario.

5.Toward a Normative Approach for Forensicability
Özgür Kafali, Munindar P. Singh, and Laurie Williams

Sociotechnical systems (STSs), where users interact with software components, support automated logging, i.e., what a user has performed in the system. However, most systems do not implement automated processes for inspecting the logs when a misuse happens. Deciding what needs to be logged is crucial, as excessive amounts of logs might be overwhelming for human analysts to inspect. The goal of this research is to aid software practitioners to implement automated forensic logging by providing a systematic method of using attackers' malicious intentions to decide what needs to be logged. We propose Lokma: a normative framework to construct logging rules for forensic knowledge. We describe the general forensic process of Lokma, and discuss related directions.

6.Modules in Wyvern: Advanced Control over Security and Privacy
Darya Kurilova, Alex Potanin, and Jonathan Aldrich

In today's systems, restricting the authority of untrusted code is difficult because, by default, code has the same authority as the user running it. Object capabilities are a promising way to implement the principle of least authority, but

being too low-level and fine-grained, take away many conveniences provided by module systems. We present a module system design that is capability-safe, yet preserves most of the convenience of conventional module systems. We demonstrate how to ensure key security and privacy properties of a program as a mode of use of our module system. Our authority safety result formally captures the role of mutable state in capability-based systems and uses a novel non-transitive notion of authority, which allows us to reason about authority restriction: the encapsulation of a stronger capability inside a weaker one.

7.A Model-based Approach to Anomaly Detection in Software Architectures
Hemank Lamba, Thomas J. Glazier, Bradley Schmerl, Javier Cámara, David Garlan, and Jürgen Pfeffer

In an organization, the interactions users have with software leave patterns or traces of the parts of the systems accessed. These interactions can be associated with the underlying software architecture. The first step in detecting problems like insider threat is to detect those traces that are anomalous. Here, we propose a method to end anomalous users leveraging these interaction traces, categorized by user roles. We propose a model based approach to cluster user sequences and find outliers. We show that the approach works on a simulation of a large scale system based on an Amazon Web application style.

8.Decentralization and Security in Dynamic Traffic Light Control
Jian Lou and Yevgeniy Vorobeychik

Complex traffic networks include a number of controlled intersections, and, commonly, multiple districts or municipalities. The result is that the overall traffic control problem is extremely complex computationally. Moreover, given that different municipalities may have distinct, non-aligned, interests, traffic light controller design is inherently decentralized, a consideration that is almost entirely absent from related literature. Both complexity and decentralization have great bearing both on the quality of the traffic network overall, as well as on its security. We consider both of these issues in a dynamic traffic network. First, we propose an effective local search algorithm to efficiently design system-wide control logic for a collection of intersections. Second, we propose a game theoretic (Stackelberg game) model of traffic network security in which an attacker can deploy denial-of-service attacks on sensors, and develop a resilient control algorithm to mitigate such threats. Finally, we propose a game theoretic model of decentralization, and investigate this model both in the context of baseline traffic network design, as well as resilient design accounting for attacks. Our methods are implemented and evaluated using a simple traffic network scenario in SUMO.

9.Visualizing the Variational Callgraph of the Linux Kernel: An Approach for Reasoning about Dependencies
Momin M. Malik, Jürgen Pfeffer, Gabriel Ferreira, and Christian Kästner

Software developers use #ifdef statements to support code configurability, allowing software product diversification. But because functions can be in many executions paths that depend on complex combinations of configuration options, the introduction of an #ifdef for a given purpose (such as adding a new feature to a program) can enable unintended function calls, which can be a source of vulnerabilities. Part of the difficulty lies in maintaining mental models of all dependencies. We propose analytic visualizations of the variational callgraph to capture dependencies across configurations and create visualizations to demonstrate how it would help developers visually reason through the implications of diversification, for example through visually doing change impact analysis.

10.Security Practices in DevOps
Akond Ashfaque Ur Rahman, and Laurie Williams

DevOps focuses on collaboration between different teams in an organization to achieve rapid deployment of software and services to end-users by automating the software delivery infrastructure. We summarize the contributions of this study as follows:

- A list of security practices and an analysis of how they are used in organizations that have adopted DevOps to integrate security.

- An analysis that quantifies the levels of collaboration amongst the development teams, operations teams, and security teams within organizations that are using DevOps.

11.Raindroid –A System for Run-time Mitigation of Android Intent Vulnerabilities
Bradley Schmerl, Jeffrey Gennari, Javier Cämara, and David Garlan

Modern frameworks are required to be extendable as well as secure. However, these two qualities are often at odds. In this poster, we describe an approach that uses a combination of static analysis and run-time management, based on software architecture models, that can improve security while maintaining framework extendability. We implement a prototype of the approach for the Android platform. Static analysis identifies the architecture and communication patterns among the collection of apps on an Android device and which communications might be vulnerable to attack. Run-time mechanisms monitor these potentially vulnerable communication patterns, and adapt the system to either deny them, request explicit approval from the user, or allow them.

12.Modeling, Analyzing, and Consistency Checking Privacy Requirements Using Eddy
Daniel Smullen and Travis D. Breaux

Eddy is a privacy requirements specification language that privacy analysts can use to express requirements over data practices; to collect, use, transfer and retain personal and technical information. The language uses a simple SQL-like syntax to express whether an action is permitted or prohibited, and to restrict those statements to particular data subjects and purposes. Eddy also supports the ability to express modifications on data, including perturbation, data append, and redaction. The Eddy specifications are compiled into Description Logic to automatically detect conflicting requirements and to trace data flows within and across specifications. Conflicts are highlighted, showing which rules are in conflict (expressing prohibitions and rights to perform the same action on equivalent interpretations of the same data, data subjects, or purposes), and what definitions caused the rules to conflict. Each specification can describe an organization's data practices, or the data practices of specific components in a software architecture.

13.Risk-Based Attack Surface Approximation
Christopher Theisen and Laurie Williams

Proactive security review and test efforts are a necessary component of the software development lifecycle. Since resource limitations often preclude reviewing, testing and fortifying the entire code base, prioritizing what code to review and test can improve a team's ability to find and remove more vulnerabilities that are reachable by an attacker. One way that professionals perform this prioritization is the identification of the attack surface of software systems. However, identifying the attack surface of a software system is non-trivial. The goal of this poster is to present the concept of a risk-based attack surface approximation based on crash dump stack traces for the prioritization of security code rework efforts. For this poster, we will present results from previous efforts in the attack surface approximation space, including studies on its effectiveness in approximating security relevant code for Windows and Firefox. We will also discuss future research directions for attack surface approximation, including discovery of additional metrics from stack traces and determining how many stack traces are required for a good approximation.

14.The Persuasive Phish: Examining the Social Psychological Principles Hidden in Phishing Emails
Olga Zielinska, Allaire Welk, and Christopher B. Mayhorn

Phishing is a social engineering tactic used to trick people into revealing personal information [Zielinska, Tembe, Hong, Ge, Murphy-Hill, & Mayhorn 2014]. As phishing emails continue to infiltrate users' mailboxes, what social engineering techniques are the phishers using to successfully persuade victims into releasing sensitive information?

Cialdini's [2007] six principles of persuasion (authority, social proof, liking/similarity, commitment/consistency, scarcity, and reciprocation) have been linked to elements of phishing emails [Akbar 2014; Ferreira, & Lenzini 2015]; however, the findings have been conflicting. Authority and scarcity were found as the most common persuasion

principles in 207 emails obtained from a Netherlands database [Akbar 2014], while liking/similarity was the most common principle in 52 personal emails available in Luxembourg and England [Ferreira et al. 2015]. The purpose of this study was to examine the persuasion principles present in emails available in the United States over a period of five years.

Two reviewers assessed eight hundred eighty-seven phishing emails from Arizona State University, Brown University, and Cornell University for Cialdini's six principles of persuasion. Each email was evaluated using a questionnaire adapted from the Ferreira et al. [2015] study. There was an average agreement of 87% per item between the two raters.

Spearman's Rho correlations were used to compare email characteristics over time. During the five year period under consideration (2010–2015), the persuasion principles of commitment/consistency and scarcity have increased over time, while the principles of reciprocation and social proof have decreased over time. Authority and liking/similarity revealed mixed results with certain characteristics increasing and others decreasing.

The commitment/consistency principle could be seen in the increase of emails referring to elements outside the email to look more reliable, such as Google Docs or Adobe Reader ($rs(850) = .12$, $p = .001$), while the scarcity principle could be seen in urgent elements that could encourage users to act quickly and may have had success in eliciting a response from users ($rs(850) = .09$, $p = .01$). Reciprocation elements, such as a requested reply, decreased over time ($rs(850) = -.12$, $p = .001$). Additionally, the social proof principle present in emails by referring to actions performed by other users also decreased ($rs(850) = -.10$, $p = .01$).

Two persuasion principles exhibited both an increase and decrease in their presence in emails over time: authority and liking/similarity. These principles could increase phishing rate success if used appropriately, but could also raise suspicions in users and decrease compliance if used incorrectly. Specifically, the source of the email, which corresponds to the authority principle, displayed an increase over time in educational institutes ($rs(850) = .21$, $p < .001$), but a decrease in financial institutions ($rs(850) = -.18$, $p < .001$). Similarly, the liking/similarity principle revealed an increase over time of logos present in emails ($rs(850) = .18$, $p < .001$) and decrease in service details, such as payment information ($rs(850) = -.16$, $p < .001$).

The results from this study offer a different perspective regarding phishing. Previous research has focused on the user aspect; however, few studies have examined the phisher perspective and the social psychological techniques they are implementing. Additionally, they have yet to look at the success of the social psychology techniques. Results from this study can be used to help to predict future trends and inform training programs, as well as machine learning programs used to identify phishing messages.

15.Operation-Level Traffic Analyzer Framework for Smart Grid

Wenyu Ren, Klara Nahrstedt, and Tim Yardley

The Smart Grid control systems need to be protected from internal attacks within the perimeter. In Smart Grid, the Intelligent Electronic Devices (IEDs) are resource-constrained devices that do not have the ability to provide security analysis and protection by themselves. And the commonly used industrial control system protocols offer little security guarantee. To guarantee security inside the system, analysis and inspection of both internal network traffic and device status need to be placed close to IEDs to provide timely information to power grid operators. For that, we have designed a unique, extensible, and efficient operation-level traffic analyzer framework. The timing evaluation of the analyzer overhead confirms efficiency under Smart Grid operational traffic.

Tutorials

Two tutorials about ways to add rigor to relatively soft data analysis were provided.

1.Systematic Analysis of Qualitative Data in Security
Hanan Hibshi, Carnegie Mellon University

Hanan Hibshi's tutorial introduced participants to Grounded Theory, a qualitative framework to discover new theory from an empirical analysis of data. It is useful when analyzing text, audio, or video artifacts that lack structure, but contain rich descriptions. She framed Grounded Theory in the context of qualitative methods and case studies, which complement quantitative methods, such as controlled experiments and simulations. She then contrasted the approaches developed by Glaser and Strauss, and introduced coding theory–the most prominent qualitative method for performing analysis to discover Grounded Theory. Topics included coding frames, first- and second-cycle coding, and saturation. She used examples from security interview scripts to teach participants how to develop a coding frame, code a source document to discover relationships in the data, develop heuristics to resolve ambiguities between codes, and perform second-cycle coding to discover relationships within categories. Then, participants learned how to discover theory from coded data, and further learned about inter-rater reliability statistics, including Cohen's and Fleiss' Kappa, Krippendorf's Alpha, and Vanbelle's Index. Finally, she reviewed how to present Grounded Theory results in publications, including how to describe the methodology, report observations, and describe threats to validity.

2.Text Analytics for Security
Tao Xie, University of Illinois at Urbana-Champaign, and William Enck, North Carolina State University

Computing systems that make security decisions often fail to take into account human expectations. This failure occurs because human expectations are typically drawn from textual sources (e.g., mobile applications, descriptions, and requirements documents) and are hard to extract and codify. Recently, researchers in security and software engineering have begun using text analytics to create initial models of human expectation. In this tutorial, Xie and Enck provided an introduction to popular techniques and tools of Natural Language Processing (NLP) and text mining, and shared experiences in applying text analytics to security problems. They also highlighted the current challenges of applying these techniques and tools for addressing security problems. The tutorial concluded with a discussion of future research directions.

NOTES

The Science of Security Virtual Organization (SoS-VO)

The Science of Security Virtual Organization (SoS-VO) was established to provide a focal point for information about ongoing activities related to cybersecurity science and as a repository for significant research results. It emphasizes community development, information sharing, and interaction among researchers in the field. The goal of the SoS-VO is to help establish and support true collaboration in advancing cybersecurity science. The SoS-VO enables those interested in cybersecurity science to survey current research, stay current on news in the field, find out about events related to cybersecurity, collaborate and share work, and access educational resources contributed by members. The SoS-VO now has over 1100 members. New members are encouraged and can join by signing up via the SoS VO website at http://cps-vo.org/group/SoS.

In addition to general cybersecurity news, upcoming events, Science of Security activities, and research being done by the Lablets and the SURE consortium, the SoS-VO has published a newsletter, which includes research papers and other items of interest to the SoS community. The 24 newsletters published thus far have over 550 bibliographies with an average of 15-20 papers each and an index to almost 700 unique entries. This year the Newsletter added "Cyber Scene" articles which are intended to provide an informative, timely backdrop of events, thinking, and developments that feed into technological advancement of SoS Cybersecurity collaboration and extend its outreach.

For more information scan the link or visit http://SoS-VO.org

NOTES

Word cloud of key terms found in titles of "Publications of Interest" topics.
Source: Permission granted by WordItOut at http://worditout.com for use of the generated word cloud.

OUTREACH ACTIVITIES

Efforts to grow the Science of Security (SoS) community are not limited to HotSoS, the SoS-VO, and the Annual Best Paper Competition. In addition to the 4 Lablets, 25 subLablets, 106 recognized collaborators, and the SURE program, there is a growing network of additional collaborators from academia, industry, and government who contribute to SoS community growth. Each of the Lablets sponsors community engagements that draw participants from the local area, including researchers from related disciplines within the lablet university or nearby universities, as well as industry. The Lablets also sponsor summer internships so researchers can work with Lablets on SoS research projects. Additionally, presentations by lablet researchers at international conferences serve to expose a broad audience to the Science of Security.

When the SoS-VO Newsletter cites an article, the authors are being notified that their work has been cited and are provided a link to the newsletter. During the past year, this outreach effort expanded to include earlier issues of the newsletter. Authors also receive recognition for having been selected for inclusion in the newsletter. With more than 10,000 authors having been published thus far, this on-going effort to connect with these published researchers has increased the awareness of Science of Security worldwide.

Also, there is a Science of Security prototype Facebook page that can be found at https://www.facebook.com/groups/SoSConnect/. This is a new effort to engage the Science of Security community via this popular social networking site.

The NSA Research Directorate has been active in involving other parts of NSA in Science of Security initiatives, briefing SoS activities to information assurance professionals and researchers from outside the cybersecurity field. This year, the Research Directorate presented a Science of Security briefing at the Information Assurance Symposium in Washington DC.

SoS sponsorship of special awards at the Intel International Science and Engineering Fair (Intel ISEF) provides an opportunity for interaction with Intel ISEF participants and facilitates a discussion of Science of Security principles with the next generation of researchers.

NOTES

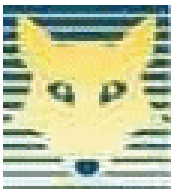

Produced by Cyber Pack Ventures, Inc.

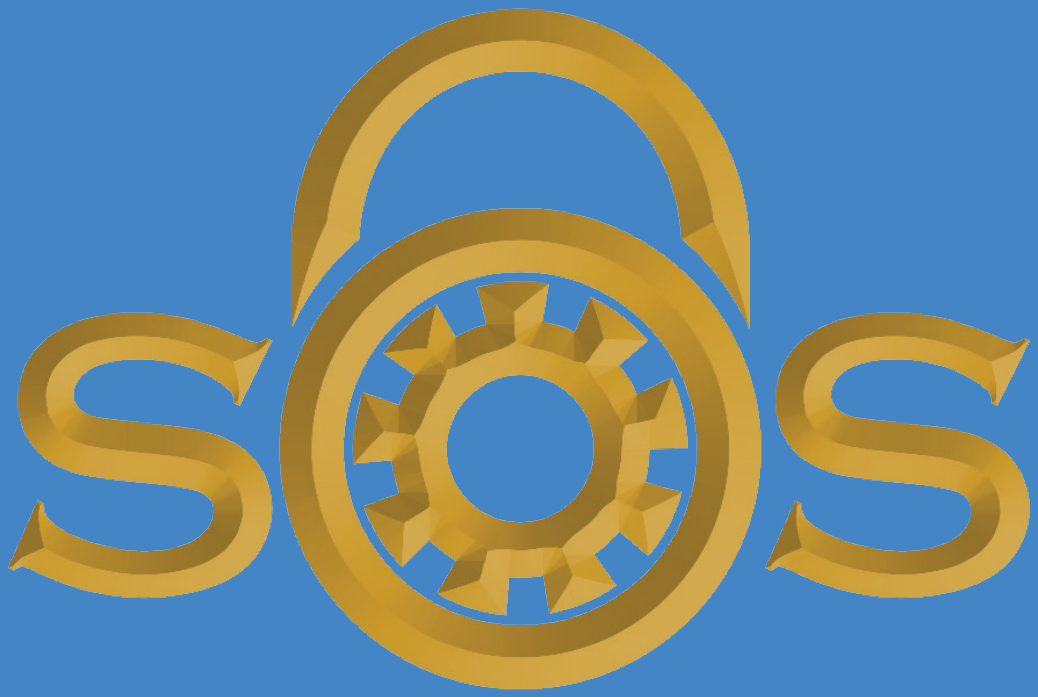

SCIENCE OF SECURITY AND PRIVACY