Science of Security Lablets
Progress on Hard Problems
August 2015

Introduction

The National Security Agency sponsors the Science of Security (SoS) Initiative to promote the underpinning of cybersecurity with a discipline built on scientific foundations. A component of this initiative funds small multi-disciplinary labs at four Universities; Carnegie Mellon, North Carolina State, Illinois-Urbana Champaign, and Maryland-College Park, to discover foundational research, promote rigorous scientific principles and grow a SoS community. In 2012 as a measure to establish the beginnings of a common language and gauge progress, the lablet Principle Investigators (PIs) developed in collaboration with NSA Five Hard Problems in security needing science to advance. A paper defining the hard problems was publically released in November 2012.[1] These five are not all encompassing of cybersecurity, but are specific challenges for cybersecurity that need research for advancement and the lablet research is focused in these areas. The five hard problems are: *Scalability and Composability*; *Policy-Governed Secure Collaboration; Security Metrics Driven Evaluation, Design, Development, and Deployment*; *Resilient Architectures*; and *Understanding and Accounting for Human Behavior*. This document details how SoS lablet (SoSL) research has improved the challenges of the hard problems.

Scalability and Composability

The hard problem of *scalability and composability* involves techniques for constructing and analyzing software that scale to large systems, and furthermore that allow analysis of the entire system to proceed by analyzing its parts independently. The two properties are closely related, because in many cases compositionality is the key to achieving scalability in practice. Lablet research has produced fundamental advances in our ability to perform compositional security-related analyses on both software and data:

- Before our work, there was no way to compositionally guarantee safety when executing unknown code provided by an adversary. We developed a theory of compositional security called "adversary-aware assume guarantee," developed an accompanying program logic that allows compositional proofs of safety for programs that execute adversary-supplied code, and applied the approach to hypervisors and trusted computing systems.
- Before our research, integrating syntax from two languages within the same program was difficult to do compositionally; it was only possible by following arcane and/or verbose syntactic conventions. As a result, developers in today's programming systems tend to use strings to write domain-specific syntax (e.g. for SQL queries) within a host language, leading to command injection attacks. We developed a way to compose

[1] The 5 Hard Problem Paper can be found on the Science of Security Virtual Organization Webpage: http://www.sos-vo.org

domain-specific syntaxes within a single program, allowing developers to write programs in more natural ways while at the same time mitigating command injection attacks.

- Before our work, approaches to information flow security did not have a semantic basis for authorizing exceptions to a non-interference property. We use a novel linear epistemic logic to analyze the execution traces of programs and enforce information flow constraints using a type system based on the lax modality. This approach is inherently compositional, and the connection to logic allows us to generate proofs in the authorization policy that provides a justification for violating a too rigid non-interference property.
- Before our research, the best-known ways to do certain attack surface analysis had combinatorial complexity. Our work has found a way to do this analysis with only linear complexity, and we can characterize the scale of the information that is given up in this tradeoff.
- Our work has also improved the scalability of security-based planning and learning algorithms, and applied graph clustering algorithms to detect insider attacks in more scalable ways.

Policy-Governed Secure Collaboration

The hard problem of *Policy-Governed Secure Collaboration* seeks to develop 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 SoSL efforts, we have deepened the scientific foundations that help us address key limitations of the state of the art.

- Prior to the SoSL efforts, although policy approaches existed that handled authentication and authorization of users for performing data operations based on attribute or role-based credentials, they did not adequately and explicitly characterize the correctness requirements for secure collaboration and their impact on security. We have advanced our understanding of how to express and validate normative requirements that are left implicit in previous research. Specifically, we have developed elements of a formal language that helps capture various subtleties of secure collaboration requirements, including priorities between them. We are developing mathematical models for determining whether those requirements are mutually consistent. We are developing methods to determine whether participants are interacting in a way that complies with the stated requirements or deviates from the requirements only when necessary to satisfy higher priority requirements.
- Previously, there was inadequate scientific basis for judging if seurity policies were comprehensible. We have developed a model of policy complexity (as perceived by humans), yielding a modular form for firewall policies that we have empirically shown helps improve user understanding.
- Previously, there was inadequate study of how security policies are developed and whether they capture stakeholder requirements. We have empirically evaluated a method for capturing design rationales for security policies and shown empirically (for

firewall policies) that it helps enhance comprehension, thereby improving maintenance and reducing errors.

- Previous research did not study 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.
- Previously, there was inadequate understanding of how security requirements may interact to impact security and risk scoring methods did not operationalize defense-in-depth, a general theory of security. We have developed a method for computing the extent to which two or more security requirements interact to impact security based on actual analyst threat perceptions; this method operationalizes defense-in-depth.

Security Metrics Driven Evaluation, Design, Development, and Deployment

The hard problem of *security metrics* involves techniques for effectively measuring and quantifying the extent to which a given system satisfies a particular set of security properties. Lablet work in this area is making progress on several aspects of the problem, for example, performing statistical analyses of real-world datasets to understand and quantify factors leading to vulnerabilities and exploits; developing metrics (whether software-, network-, or system-based) to predict vulnerabilities and/or measure effectiveness of countermeasures; and measuring humans' perceptions of various security measures. Highlights of this work include the following:

- Before our work, one could only speculate about the real-world effect of new security technologies that have been introduced. Empirical studies using the WINE dataset have been used to give evidence as to how the cyber threat landscape has changed following the introduction of various security technologies. This work found, for example, evidence that exploits decreased following the introduction of sandboxing techniques in IE7 and Adobe Reader 10.
- Before our work, there were hundreds of disparate publications about intrusion-detection systems, each with varying methods and evaluation approaches. Our work has led to a taxonomy to compare those studies and to systematize that knowledge.
- Before starting this project, only a few approaches existed for analyzing the attack surface of software systems. Those approaches were more at the system level and did not provide actionable feedback to software engineers as they develop code. Today, we have discovered that there are metrics that can predict vulnerabilities at the method level. By simulating the sequence of actions that an attacker might take when exploiting a vulnerability, we are able to estimate the areas of a system where vulnerabilities are likely to be found. Recent results from a large, open-source project show that our metrics increase just before a vulnerability is found, and decrease after a vulnerability is fixed, giving empirical evidence that our new metrics are useful predictors of vulnerabilities.
- Before our work, models to predict the presence of vulnerabilities and the resilience of systems were not accurate enough to make them actionable by practitioners. We have

made progress in developing new metrics that can be used to more accurately evaluate the probability that a given host is vulnerable and, if so, whether it might be exploited.

- Before our work, it was unknown how quickly software patches were applied in the real world, or what techniques would be most beneficial for incentivizing faster patching. Lablet work has empirically demonstrated (using both the WINE dataset as well as network measurements on PKI revocation data following the Heartbleed incident) that software patches for known vulnerabilities are either not applied in a timely fashion, or are applied incorrectly. Beyond characterizing the rate of software patching, the work also aims to determine factors that influence this rate.
- Other lablet work is investigating techniques for representing real-world security incidents, and designing sound methods for preemptive detection of such events. This work looks at relations between available evidence (e.g., a sequence of observable events) and hidden system states to determine the most probable sequence of state transitions consistent with the evidence. This, in turn, is used to automatically assess, based on the evidence, whether a system or a user account has been compromised. The resulting tool was validated with real-world incidents collected over a 6-year period.
- Finally, lablet researchers are looking at quantifying users' perceptions of security, with current research focused on the graphical password system used by Android smartphones. By understanding what makes users perceive something as secure, researchers hope to design systems whose actual security aligns with those perceptions, thus influencing better choices on the part of users.

Resilience

The hard problem of *resiliency* has several attributes, with different emphases depending on the context and community. One attribute captures the notion of robustness. That is, the ability of the system to statically withstand attack, e.g. through diversity in implementation. Another attribute captures the notion that a system can continue to deliver essential services (albeit potentially at a diminished level) in the midst of an attack. Yet another attribute stress how quickly a system can be restored to full functionality following an attack.

- Before our work began, means to specify resiliency properties and requirements were not sufficiently precise or detailed to serve as a basis for rigorous systems engineering. We have developed a formal mathematical framework to enable more precise specification of the full range of properties of affordability, reliability, availability, safety, usability, scalability, evolvability, and resilience.
- Previously, network policy enforcement was primarily deployed statically at network ingress points. We have now found that a top-down strategy can be used to deploy policy enforcement across a network with greater efficiency and scalability than traditional, ingress-only deployments. This supports policy enforcement that can better absorb and adapt to adversarial traffic patterns.
- Previously, the properties and tradeoffs among different security isolation techniques had not been made explicit. Based upon a literature survey, we have found that the

security isolation problem comprises a large design space consisting of many design dimensions. Existing approaches fall short in terms of adaptability and measurability.

- Previously, there existed no way to reason about how robust a cyber-physical system might be to disruption (robustness being one attribute of 'resilience'). Understanding was particularly lacking in how to approach reasoning about how an attack on the cyber component might be effected by manipulation of the physical component. Lablet research has broken ground in developing practical mathematical frameworks that support this reasoning, within the context of a hybrid system model that conjoins the continuous control description that interacts with the physical realm, and the cyber component by which the control is implemented. Using this framework we have developed algorithms that measure bounds on "how close" a physical disturbance can push a cyber-physical system near deleterious states.
- Previously, analytic approaches to resiliency emphasized the robustness attribute of resiliency. We have extended the foundations of analyzing resilient architectures by providing a reasoning approach to self-protection that uses stochastic multiplayer games to reason about human involvement, self-protection latency, and uncertainty, accommodating a much more dynamic viewpoint of resiliency.

Understanding and Accounting for Human Behavior

The hard problem of *human* aids in the handling of the unpredictability of human actors in the security of systems. 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. A variety of research projects were dedicated to developing models and insights of human behaviors that enable the design, modeling, and analysis of systems with specified security properties.

- The Security Behavior Observatory is giving us unique insights into the human behaviors that lead to security vulnerabilities on personal computers. For example, we are able to trace the steps that our participants have taken that have led to malware infections and to participants uninstalling anti-virus software.
- Before we began our work, security analysts nominally scored security requirements (e.g., control lists) to determine their independent impact on security. Today, we have a new method for computing the extent to which two or more security requirements interact to impact security based on actual analyst threat perceptions. Our method implements defense-in-depth, a general theory of security that had been spoken to, but not operationalized by, traditional risk scoring methods.
- Before our project began, there was ample speculation but little empirical evidence about how or why type-like specifications affect the productivity and accuracy of human programmers. We studied this question in the setting of application programming interfaces (APIs) that define protocols of interaction that API clients must follow--a problem with significant security and reliability implications in practice. Using laboratory studies and experiments, we found that leveraging type-like protocol specifications as

documentation can increase programmer productivity and reduce programmer errors, both by a factor of 2 or more, when developers are using particularly challenging APIs.

- Before our work began, researchers knew that phishing scams were problematic but they did not understand why people fell prey to them. We have now found that social factors such as trust and cognitive factors such as attention and impulsiveness influence the likelihood of falling prey to social engineering when phishing emails are received. We have worked to classify the message content of hundreds of archived phishing emails to determine how classical persuasion research predicts the likelihood of data loss. Our goal is to understand how social engineering occurs and to develop techniques to combat these tactics being used by cyber-criminals.
- Before our project began, biometrics research had found ways of processing low-level interaction with a computer--the mouse movements and the keys typed on a keyboard--to authenticate individual users, potentially as a substitute for typing passwords. Our work took on a related question: could such interaction patterns be used to distinguish ordinary users from malicious users, behaving deceptively? The algorithms we have developed can detect patterns that co-occur with different usage motivations, in the laboratory settings we have evaluated. These algorithmic models, which are informed by work in cognitive psychology, have accuracy above 80% to 90% in our experimental user studies.