Editors: Patrick McDaniel, mcdaniel@cse.psu.edu | Sean W. Smith, sws@cs.dartmouth.edu

Toward a Science of Secure Environments

Patrick McDaniel | Pennsylvania State University
Brian Rivera and Ananthram Swami | Army Research Laboratory

One of the challenges of systems security is its relative immaturity. Whereas other disciplines have well-established theories and design principles, a universal theory ensuring that our personal and professional activities remain secure simply doesn't exist.[1,2] As many in the security community have observed, there's a lack of a fundamental science underlying current security practices. Even the best guesses as to the nature of such a science are often lacking. What we need is a new science of securing not just systems, but environments.

Systems security has been studied for many decades. Concepts such as mandatory and role-based access control, formal policy, and information flow provide powerful abstractions on which we can build secure systems. Systems and models that faithfully adhere to these abstractions have guaranteed behavior with respect to well-defined security properties—thereby providing a proof of security or compliance in the scope of that system.

However, applying these models to real-world environments has been more difficult in practice. Existing models focus on one system aspect and often fail to account for the complex interactions and dependencies among systems, networks, and users. Although these models provide strong guarantees in protected systems, they begin to break down when faced with complex and heterogeneous environments, buggy implementations, uncertainty, and human error.

There's room for another science of security that builds on these models and theories to extend to diverse, heterogeneous, and unpredictable environments containing many systems and users performing a multitude of functions.

Another View of the Science of Security

A traditional view of the science of security stems from formal reasoning about modeled systems or domains. This approach can be generalized in the following way: Given a known initial system state and a known set of system dynamics and behaviors, we can make inferences about all possible future system states. In such an approach, any perfectly modeled system that can be shown to never arrive in an insecure system state will be secure in perpetuity.

Inherent to such an approach is the assumption that the system and possible inputs can be characterized well—that is, models of dynamic threats and adversarial behaviors are known or predictable. Such isn't often the case in real, diverse networks. The gap between the idealization of an environment and its reality is fertile ground for vulnerability and exploitation.

Modeling (and understanding) security is complicated by diversity, uncertainty, and complexity. Real environments contain a diversity of people and devices, each with its own functions, implementation, failure modes, and limitations.

Users and adversaries exist in the environment—potentially for short periods—and fulfill roles that might change over time. Understanding the complete state of the environment and users is nearly always impossible. Most often, all we can do is approximate the environment state. The difference between this approximation and reality—the uncertainty—hampers any attempt to secure the environment. More important, the environment isn't static. All this leads to the fundamental problem that environments are complex places in which organizations attempt to secure intractably numerous and unknowable interactions among people, devices, and networks.

This situation has led many in the technical community to believe that the environment's complexity and uncertainty prevent any chance of building a comprehensive science of security.[3] We feel differently.

Copublished by the IEEE Computer and Reliability Societies
1540-7993/14/$31.00 © 2014 IEEE

Suppose we reformulate the reasoning system in the following way: Given a fixed initial approximate system and user state, risk assessments, and probable behavior sets, what course of action will likely lead to the best outcome? Here, the goal is to provide the rigorously derived optimal reaction to a set of circumstances based on our best but inaccurate and incomplete knowledge of the environment and people, including adversaries.

This is precisely the science that the new Cyber-Security Collaborative Research Alliance initiative (http://cra.psu.edu)—a 10-year project spanning five universities and the Army Research Laboratory—is attempting to form. We consider different aspects of this new science, in particular, models for users, risk detection, and decision-making in partially known environments.

Understanding Users Is Essential to Security

Security models often omit users entirely or simply view them as random, untrustworthy actors. In real environments, this limited view blinds us to the opportunity to engage and use them to inform security decisions. Moreover, failure to understand users prevents us from predicting and reacting to their behavior. Users are nuanced and context-sensitive actors and must be treated as such. Would a better understanding of users', malicious insiders', and other adversaries' motivations, goals, and techniques lead to better data security? Almost certainly.

Key to good decision-making is understanding the users' state and predicting how they will react to a set of stimuli. We refer to this as a *user model*. However, user models can be as diverse as the people who populate computing environments. An important step is determining which aspects to include in a user model. Accept that users can subvert even the best security apparatus through negligence, by accident, or through malicious action. Such behaviors might be more pronounced during times of stress, exhaustion, or other mental or physical states. Thus, we can predict the expected behavior only if we understand a user's state in the moment.

Users' training and experience heavily influence their reactions to security prompts and events. One view of experience suggests that advanced users tend to interpret security prompts and signals more quickly and correctly, whereas novices tend to have trouble. Yet, some studies suggest quite the opposite. Furthermore, users often act within groups and communities that influence their behavior. What's clear is that experience impacts the cognitive processes that lead to decision-making. Thus, we must integrate notions of experience into our models.

Given a fixed initial approximate system and user state, risk assessments, and probable behavior sets, what course of action will likely lead to the best outcome?

Action Calibrated by Risk

Risk is the chance that something of value will be lost or impacted, be it data, trust, availability, time, or another resource. For example, there's a nonzero risk of a user's browser being compromised when surfing websites. Security in complex environments is often about trading risk for availability and usability. Users constantly make these tradeoffs often without knowing it, for instance, trading website access with potential browser compromise. Users make these security decisions at least in part via a risk calculation, even if they aren't conscious of it. Such decision-making is at the heart of much of practical security, and thus it's important to integrate it into a science aimed at governing it.

Integrating risk into a science of security poses two challenges. First, recognizing and assessing risk are fairly hard. Like beauty, risk is in the eye of the beholder. One user's perceived risk might differ significantly from another user's. Context also matters. For example, all other things being equal, the risk of browser compromise on a critical system is likely to be very different from that on a low-value netbook.

The second challenge involves weighing risk relative to other needs and risks. Simply recognizing that particular risks exist isn't very informative unless you have a coherent means to compare them. When is it appropriate to browse the network? Is the information critical to some other need in the network? When do functional needs outweigh risk? What's the composite risk in a set of sequential or parallel actions? What are the significant sociocognitive factors that impact a user's risk assessment?

Fortunately, the technical community has put forth significant effort to develop taxonomies and risk assessments in cybersystems.[4] These works have led to initial frameworks for assessing and managing risk and serve as a grounding mechanism for a science of security.

Understanding the Environment State

Knowing the environment state is also essential to good decision-making. However, developing perfect knowledge of the state of a nontrivial environment is infeasible owing to its complex and interconnected nature as well as the systems'

inherent randomness, inaccuracies in threat detection, and dynamics of user interactions—both adversary and defender. Thus, developing accurate estimators of the environment state is important.

Decision-making systems based on state estimation are well-practiced in engineering. For example, the principle of physical control systems enables the design of actuator inputs based on approximated physical states and stochastic models of sources of randomness and unmodeled dynamics ("noise"). Indeed, many controls that ensure power grid safety and operation are built on this model. It's important that the approximation is accurate (within some bounds) and self-correcting over time.

State estimation might also lead to new methods of discovering and characterizing anomalous activity, potentially offering new approaches that diverge from contemporary detection systems. An open problem in this effort is determining state estimation's applicability in the presence of adversarial action. Could adversaries use such estimation as a tool against the protection mechanisms, or can we design these mechanisms to prevent malicious manipulation? The answers to these questions are as yet unknown.

One way to reconstruct the environment state is to collect information from many sources and at different system levels. Correlating these sensor inputs into an internally consistent model of the system or environment under inspection is key to developing accurate models. Updating the model with new information while remaining consistent with all available information will help ensure that the state estimate is self-correcting.

What information should we collect to make this state estimation? The answer lies in the purposes for and timescales at which it will be used. If we're making decisions about incident response, we need to understand the nature of each incident and the potential and real impact it has on users and systems around it.

A Science of Optimization

So, given an approximate environment state, how do we determine the action that produces the best outcome? Put another, more practical way: Given an imperfect approximation of the current state, what do we do to best ensure the environment's security?

Suppose that the science can accurately capture the users, risks, and system state. How do we use this information to determine a best course of action? This requires finding a way to determine what *best* means, and there might be many metrics of interest. If we understand the desired outcomes and the means used to arrive at them, we can weigh different actions' relative risks and expected outcomes. The desired outcome is the network's target end state.

One course of action here is to maneuver through this state space; each point in the trajectory is associated with risk and utility. Note that the space of maneuvers across which we can optimize is quite large. We can change security policy, allocate new resources, disable or modify services, increase sensor collection, or perform any other alteration that would improve the probability of a positive outcome.

Ultimately, the new science of security seeks to minimize risk and optimize utility. It models security as a continuous optimization of the environment given imperfect information and incomplete models of future behaviors. If successful, this complementary departure from traditional cybersecurity science will broaden the scope of investigation to allow users to be more secure in realistic, highly dynamic environments and unknown to the people and organizations securing them. ■

References

1. *Science of Cyber-Security*, report no. JSR-10-102, JASON Program Office, MITRE; www.fas.org/irp/agency/dod/jason/cyber.pdf.
2. F.B. Schneider, "Blueprint for a Science of Cybersecurity," *Next Wave*, vol. 19, no. 2, 2012; www.cs.cornell.edu/fbs/publications/SoS.blueprint.pdf.
3. T. Longstaff, "Barriers to Achieving a Science of Cybersecurity," *Next Wave*, vol. 19, no. 4, 2012; www.nsa.gov/research/tnw/tnw194/article5.shtml.
4. D. Liu et al., "Security Risk Management Using Incentives," *IEEE Security & Privacy*, vol. 9, no. 6, 2011, pp. 20–28.

Patrick McDaniel is a computer science and engineering professor at Pennsylvania State University. Contact him at mcdaniel@cse.psu.edu.

Brian Rivera is a branch chief in the Army Research Laboratory's Network Science Division. Contact him at brian.m.rivera.civ@mail.mil.

Ananthram Swami is the senior research scientist in the Army Research Laboratory's Network Science Division. Contact him at a.swami@ieee.org.

cn *Selected CS articles and columns are also available for free at http://ComputingNow.computer.org.*

Have an idea for a future article? Email editors Patrick McDaniel (mcdaniel@cse.psu.edu) and Sean W. Smith (sws@cs.dartmouth.edu).