2018 Science of Security and Privacy (SoS) Kickoff Meeting

The Science of Security and Privacy (SoS) Initiative held its kickoff meeting for the 3rd generation of Lablets on 13-14 March 2018 at the Laboratory for Telecommunications Sciences in College Park, MD. The 3rd generation of SoS Lablets will focus on twenty specific projects that address some of the most significant cybersecurity research challenges aligned against the five Hard Problems. The Lablets are Carnegie Mellon University (CMU), the International Computer Science Institute (ICSI), North Carolina State University (NCSU), the University of Illinois at Urbana-Champaign (UIUC), the University of Kansas (KU), and Vanderbilt University (VU).

Following a meeting between NSA personnel and Lablet Principal Investigators (PIs) to discuss expectations under the new contract, Adam Tagert, SoS Technical Director, welcomed the attendees and introduced the NSA SoS initiative leadership team and the Lablet PIs.

CMU: Bill Scherlis, Jonathan Aldrich
ICSI: Serge Egelman
KU: Perry Alexander
NCSU: Munindar Singh, Laurie Williams
UIUC: William Sanders, David Nicol
VU: Xenofon Koutsoukos

Dr. Tagert noted that the 3rd generation of Lablets has the same goals as the prior two generations, and he emphasized building up the foundational aspects of SoS anchored in scientific methods, models, and approaches. He is expecting research breakthroughs and the development of new technologies and tools under the initiative. NSA SoS leadership is seeking better engagement between NSA and the Lablets, and he wants to facilitate tech transfer arising from Lablet research.
Following the welcome and introductions, approximately 65 attendees from government and academia heard presentations on NSA Research perspectives related to Privacy, Cyber-Resiliency, and Cyber-Physical Systems/Internet of Things (CPS/IoT) in order to set the stage for understanding the SoS challenges. After a presentation by George Coker, Chief Information Assurance Research, on an Information Assurance Research overview, Lablet representatives briefed the audience on the individual research projects to include their goals and plans for the upcoming year. Summaries all of the presentations are provided below, and selected presentations can be found at https://cps-vo.org/node/36719/browser

NSA Research Perspectives

Privacy
Travis Breaux (NSA—Secure Systems Architecture and Analysis)

Dr. Breaux identified the concepts of privacy to include confidentiality and secrecy as well as personhood and autonomy and also addressed self-determination and self-discovery. He focused on the project of Protected Autonomy which shifts the focus to data integrity, noting that assured decision making depends on trusted algorithms for making predictions and curating information for human consumption. One of the research questions being addressed under this project is how to compose multiple non-symbolic programs (ML models) into a single simulation. He identified technical challenges as realistic datasets and algorithms, simulation infrastructure design, and mitigation design and evaluation. Other projects include Privacy Enhanced Architecture (use-based privacy and private information retrieval); Compliance Assistance (querying policy to improve legal and engineering coordination, emergent compliance—discovering rules from norms, and organizational practices to design privacy into software) and Mobile and IoT privacy. See https://cps-vo.org/node/54406

Cyber-Resiliency
Jim Holt (NSA—Adaptive Cyber-Defense Systems)

The speaker's presentation focused on what his team is trying to accomplish and how the Lablets can participate. In addressing autonomous cyber defense, he noted that cyber attackers have an asymmetric advantage over defenders and that another detector, sensor, or tool doesn't help; rather, the challenge is how to change the balance of power. To address the challenge, the team hypothesizes that the balance can be changed through increasingly autonomous cyber defense and the strategic use of deception in cyber defense. Human decision-makers can't respond quickly enough, but they can achieve speed and scale through autonomous cyber defenses. They acknowledge that autonomy is a spectrum, but that humans can delegate decision-making and actions to the system. Strategic use of deception deals with crafting answers to attackers to

influence their decision-making, since virtually every success an attacker is able to have is possible because the networks provide correct answers to attacker questions.
A key concept in this strategy is cyber resilience: anticipate, withstand, recover, and evolve. The speaker addressed building autonomous cyber defenses and then demonstrating their effectiveness scientifically. Challenges in doing so include the following: system complexity; defining success; differing missions, values, goals, and environments; undefined, unbounded, and evolving threats; and the fact that full-loop execution requires many components which are interdependent. He addressed some of the goals in addressing the challenges and focused on their relationship to Science of Security including well-defined measures of success, reproducible experiments, and collaboration. The challenges to implementing the approach include incomplete, uncertain, and/or untrusted data and testing. The decision-making piece includes looking at tools and techniques in AI, ML, planning systems, game theory, robotics and optimization, with the eventual choice likely being a hybrid of these elements. Their plan is to start simply and evolve, building simple versions to demonstrate the framework's extensibility. They have also considered building an open-source, shareable testbed, and the speaker addressed multiple ways for the Lablets to participate and collaborate. See https://cps-vo.org/node/54407

Cyber-Physical Systems/IoT
Raj Pal (NSA—Trust Mechanisms)

Dr. Pal's presentation was entitled "Building towards a Trustworthy IoT Ecosystem" and he said that while the IoT team has been focusing on applied research, they are looking to partner with the Lablets to expand foundational research. He noted that the ideal system would have end-to-end trust, root of trust, be remotely attestable, and have trustworthy integrity verification. He defined IoT as a set of network capable products when platform integrity cannot be verified with confidence and said it was important to be able to measure these devices in a trustworthy way. He believes that the research challenge, in a resource and functionally constrained platform, is how to overcome hardware and functional barriers and incorporate the mechanisms identified for richer platforms into the constrained device. Relevant Information Assurance Research motivations included emerging technologies to further the mission and blending trusted and untrusted devices. He identified IoT research interests as 8,16, and 32 bit platforms; trusted computing base, trustworthy integrity verification, memory separation and P2P networks. The desired impact of the research, he concluded, was to develop solutions to shape the security landscape and influence industry. See https://cps-vo.org/node/54405

Information Assurance Research Overview
George Coker (NSA—Information Assurance Research Group)

Dr. Coker's presentation, "Cybersecurity: Effects at Scale" focused on a science-based approach to cybersecurity. He described the evolution of Information Assurance Research that has led to a focus on cyber and noted that resilient and scalable solutions are now the crux of the issue. The speaker addressed the five Hard Problems as well as CPS, IoT, and privacy. He concluded by noting that reversing the asymmetric advantage attackers now enjoy will require the achievement

of defensive effects at scale which, in turn, necessitates a science-based approach in order to understand if we're making a difference at scale. See https://cps-vo.org/node/54404

Lablet Project Presentations

Dr. Tagert introduced the Lablet project presentations, grouped by Hard Problem. Each presentation provided an overview of the project, its objectives, and project PI contact information. The project presentations were designed to inform attendees (NSA researchers and other Lablet researchers) of the specific areas of research and to make connections to government researchers interested in the work. Dr. Tagert noted that all projects being presented are funded and encouraged the NSA attendees to meet with the Lablet researchers to discuss mutual needs.

Hard Problem: Policy-Governed Secure Collaboration

Uncertainty in Security Analysis (UIUC)

Frank Nguyen

The speaker noted that there are many models in security analysis of computer systems, all of which require information about devices, interconnections, services, configurations, attacker and defender. The problem, he said, is that in practice, the information about the model is incomplete, which leads to either making simplifying assumptions or explicitly modeling the incomplete information as an input uncertainty. The goal is to develop techniques for expressing uncertainty in the input of the security models and for assessing uncertainty in the model output. He addressed the domain, the threat model, the attacker's goal, and security metrics. He talked about the early work done in this area (presented at HoTSoS 2017) including how topological uncertainty networks impact reachability analysis and Extended Uncertain Graph (EUG). Theoretical results show that EUG is capable of describing any joint distribution of edge existence and that uncertainty analysis in EUG is tractable if the Boolean functions are monotone. The technical approach includes formalisms for expressing uncertainty in model input, analysis techniques for assessing model output uncertainty, the UQ framework for scientific evaluation of outcomes, and demonstration on large-scale real-life attack graphs. He tied the research to Policy-Governed Secure Collaboration by noting that the attacker's ability to harm a system depends on the policy used to protect it—there might be uncertainty in knowing exactly what policy is used. The Year 1 milestone is to develop the formalisms. See https://cps-vo.org/node/54408

Analytics for Cyber-Physical System Cybersecurity (VU)

Nazli Chouri

This research is focused on analytics to better understand the structure of policy, and the focus is to construct methods-sequence to extract value for policy guidelines for cyber security. The issue was cast as a policy problem: guidelines encourage passive compliance rather than active performance and policy documents are framed as "stand alone" and unconnected to related

documents. The research uses NIST reports on the cybersecurity of the smart grid and researchers employ analytic methods to capture full value of cybersecurity policies and guidelines. The overall objective is to integrate smart grid cybersecurity policies with a research approach that undertakes a multi-method modular investigation of cybersecurity policy documents in order to create coherent and verifiable analytics for SoS. The speaker identified three dimensions of SoS contributions

- Policy analytics: replicable methods for analysis of systems and enterprise-wide cybersecurity
- Research: multi-methods for deep analysis of cybersecurity, policies and framework
- Education: demonstrate use of multi-methods for dynamic analysis of cybersecurity

See https://cps-vo.org/node/54401

Operationalizing Contextual Data (ICSI)
Serge Egelman

The speaker presented work on mobile device apps that has led up to what they plan to do in the future, addressing privacy as contextual integrity. He noted that inappropriate data flows violate contextual information norms; contextual information norms are modeled using data subjection, data sender, data recipient, information type, and transmission principle (constraints). In questioning what this means for user-centered design, he suggested that an app should only provide notice when reasonable privacy expectations are expected to be violated. He described studies done on permission requests when a phone was inactive (a training exercise), addressed the use of ML to detect when context has changed from expected data use to unexpected, and then described a second experiment done in real-time that confirmed earlier findings. The next steps to determine what parameters are actually important to users are:

- Phase 1: Factorial vignette studies (interviews, surveys; randomly generated scenarios based on controlled parameters)
- Phase 2: Observational studies (instrument phones, detect parameters and resulting behaviors)

Principles of Secure Bootstrapping for IoT—NCSU
Ninghui Li, Purdue University

The speaker noted that this research builds upon work begun several years ago, citing the motivation as the fact that IoT devices need trust and secure communication—trust between devices and trust between device and users. Constraints, however, limit options and deployment scenarios determine resource availability, including power supply, computing resources, and serviceability. The research goal is to develop a lexicon and principles to model the different IoT security bootstrapping scenarios and tools to help developers. He described a five-step research plan:

- Determine how it works today in different application domains
- Develop conceptual framework and vocabulary
- Analyze device interactions from the perspective of a single device

- Analyze combinations of adversary model, capability, resource, protocols and security goals
- Develop tool to aid developers

Metrics include the number and importance of protocols classified by the framework, the number of vulnerabilities and the percentage of failed protocols. The success criteria include being able to see the developed lexicon and develop the most important IoT bootstrapping tool. He also addressed the envisioned scientific contributions.

Contextual Integrity for Computer Systems—ICSI
Michael Tschantz

The speaker described the overall goal of the research as converting the philosophical theory of contextual integrity into terms computer scientists can use. He noted that there is no agreement on what a context is: philosophers and computer scientists have different understandings, with philosophers focusing on abstract spheres of life and computer scientists focusing on the concrete. The goal is to develop models of context and contextual integrity that meet computer scientists on their own truth. Relevant research questions include accounting for privacy in the design of multi-use computer systems that cut across contexts; modeling the adaptation of contexts to changes in technologies; and determining how contextual integrity relates to differential privacy. The current organizing hypothesis is that contexts are defined by a purpose. He noted that the privacy norms of a context promote the purpose and that purpose restrictions are ubiquitous. He proposed several possible models including game models, Markov decision process models, partially observable Markov decision process models, and multi-agent influence diagrams. Some of the challenges are that contexts don't exist in a vacuum, contexts might be in competition, privacy is multifaceted, and people often disagree. He identified potential outcomes as progress on defining privacy, further accountability for big data systems that cut across contexts and enabling policy governed privacy with respect to collaboration.

Obsidian Language for Blockchain—CMU
Joshua Sunshine
Jon Bell, George Mason University

This research is focusing on Etherium and Hyperledger platforms and is also focusing on smart contracts. The researchers are designing a program language—the goal is to do human-centric language design so that resultant programs are secure—and it will be obsidian language. They noted that smart contracts have forced ordering, and Typestate language enforces ordering constraints. They addressed secure collaboration in scientific research, noting the wide variety of artifacts and the desire for a decentralized mechanism to share across organizations. They also addressed the issue of complex access controls expected by different parties. A strawman approach is a trusted third party, but it is not ideal because of the temporal nature of saving. They noted that BitLedger allows researchers to share and that it uses blockchain system for access controls. They reported on controlled experiments of programmers, and the speed and security of writing Obsidian language smart contracts. They are also looking for other blockchain applications and addressed collaboration as key to their approach.

Scalable Trust Semantics and Infrastructure—KU
Perry Alexander

The speaker noted that although KU is a new Lablet, the university has a history working trust issues with Information Assurance Research and predecessor organizations. He described the criteria for when you should trust a system as the following: you know its identity; you know it's built from good parts; you know it's behaving as expected. He also addressed semantic remote attestation--he presented a simplistic model and then explained why it's more complicated than the model presented. In developing a science of trust there are five tasks:

- Semantics of trust
- Measurement, attestation, and appraisal:
- Roots of trust
- Attestation protocols
- Implementing and scaling infrastructure

This research is currently focused in two areas: Development of attestation protocol semantics (under way now with Information Assurance Research, MITRE, John Hopkins University Applied Physics Lab); and Soundness and sufficiency of measurements.

Governance for Big Data—ICSI
Serge Egelman

In introducing the topic, the speaker suggested that the risk in governance for big data is that access control does not capture privacy requirements. He addressed sensitive inferences and reidentification, noting that it is difficult to redact sensitive information from rich data sets and that often sensitive data can be reidentified using additional information outside the data set or proxies. He suggested that Machine Learning will find such correlations automatically; binary allow/deny access control fails to capture this well. In discussing limiting sensitive inferences, he pointed out several related issues, including differential privacy, encryption and access control, and fairness issues. A new data governance approach focuses on accountability and relates more to accounting and auditing. This project aims to synthesize computer science abstractions with governance goals. The first step is to develop a design methodology from all different approaches and mechanisms, and then validate the design methodology by working with practitioners and building case studies for generalizable design patterns.

Designing for Privacy—ICSI
Serge Edelman

The project focuses on designing for privacy holistically: from "privacy by design" to "privacy with design", i.e., designing with privacy throughout whole life cycle. The speaker noted that design interventions for privacy can occur at a lot of stages and levels, and that the goal of the project is to develop a new toolbox of techniques and help designers understand when best to apply tools. He addressed defining privacy in contextual, situational, and relational ways, and

identified its dimensions as theory, protection, harm, provision, and scope. The goal over the next year is to put together design card activities, design workbooks and privacy design patterns. He also plans to hold privacy design workshops to address engineering practices, methods, and tools, bringing together practitioners, researchers, and policy-makers.

Hard Problem: Resilient Architectures

Foundations of CPS Resilience—Vanderbilt
Xenofon Koutsokos

The speaker addressed the need to develop a systematic body of knowledge with strong theoretical and empirical underpinnings to inform the engineering of secure and resilient CPS that can resist unanticipated attacks. The foundation of CPS resilience includes developing principles and methods for designing and analyzing resilient CPS architectures that deliver required service utility in the face of compromised components; integrating redundancy, diversity, hardening methods for designing passive resilience methods that are inherently robust against attacks; and developing active resilience methods that allow response to attacks including optimal control and reconfiguration. He discussed automated and connected vehicles within the context of the resiliency of intelligent transportation models. The research seeks to develop models to form hypotheses for simulations. Model components include diversity, redundancy, and hardening, and integrating those components for designing passive and active resilience methods (passive--robust against attacks; active--allow responses). He addressed how to improve structural robustness in networks citing the need for more than redundancy by adding diversity and hardening. He discussed game-theoretic formulation to find optimal resiliency and optimal defense strategy in the face of attacks and used examples of attack and defense in the transportation network. System models include configuration, attack, detection, mitigation, and responsive attack. The goal is to develop a model that allows finding optimal resilient configurations of CPS by integrating redundancy, diversity and hardening in the face of strategic attacks. He concluded by nothing that considering attacks in CPS in all their insidious variety creates a massive challenge can't be neglected due to potential consequences.
See https://cps-vo.org/node/54403

Coordinated Machine Learning-Based Vulnerability & Security Patching for Resilient Virtual Infrastructures—NCSU
Helen Gu

The presentation focused on Docker security and addressed attack surfaces and vulnerabilities in Linux kernel, Docker engine, and container applications. The existing approach is static security analysis and scheduled patching. In the researchers' experiments, this approach fails to detect 90% of vulnerabilities, displays high false alarms, and shows memory inflation caused by unnecessary security patching. Their proposal is runtime vulnerability detection using online machine learning methods and just-in-time security patching. Just-in-time security patching includes applying patches intentionally after attacks are detected, enforcing update validation,

making intelligent decisions on update vice rebuild, and adhering to system operational constraints.

Model-Based Explanation for Human-in-the-Loop Security—CMU
David Garlan

The speaker provided context for the research by noting that automation is becoming increasingly important for modern systems, and many systems require combinations of automated and human involvement to handle security attacks. The problem is how to create effective coordination.
The solution is simple in explanation but difficult in practice--the system needs to understand what humans can do, and humans need to understand the system. There need to be decisions on which tasks are to be allocated to the system vice humans, and humans must be able to trust in automated actions. Automation is improved by learning based on what humans do. Prior research included the adoption of a control systems view of system autonomy and led to the development of RAINBOW framework; earlier work also looked at humans as actuators who effect changes. Current research addresses putting the human in the planning area. A key idea associated with this work is to use formal models for planning as the basis of human-understandable explanation. Technical challenges include explaining a plan that is computed from a probabilistic system model and determining the basis for selecting the best alternative. He concluded by noting that system resilience and security can be enhanced through automation, but autonomous decision making is often opaque. We need better transparency through an explanation of the models used for planning, and we can inform system autonomy by allowing a system to learn by example from expert user behavior.

Predicting the Difficulty of Compromise through how Attackers Discover Vulnerabilities—NCSU
Andy Meneely

This project focuses on the attack surface based on the notion that pathways into the system enable attackers to discover vulnerabilities. This knowledge is important to software developers, architects, system administrators, and users. The speaker noted that a literature review to classify attack surface definitions led to six clusters of definitions which differ significantly (methods, avenues, flows, features, barriers, and vulnerabilities). He further discussed the methodology used to discover the attack surface (mining stacktraces from thousands of crash reports) and what the attack surface meant within the context of metric actionability, evolving the models for risky walk and deploying a human-in-the-loop study. Future activities include incorporating risky systems calls, architectural decisions, risky developer activity and human-in-the-loop. One of the goals of the project is how to turn the attack surface into a number to be able to provide actionable feedback. The researchers want to develop metrics that are useful and improve the metric formulation based on qualitative and quantitative feedback.

Formal Approaches to the Ontology and Epistemology of Resilience—UK
John Symons

The speaker began by identifying the epistemic challenge as "what is the best way to understand resilience?" and the ontological challenge as "what is resilience and how does it emerge?" He noted that this work contributes to the establishment of interdisciplinary Science of Security by focusing on its most important concept at the fundamental level—formalism—and attention to neglected aspects. He said that with respect to cybersecurity, the view is that the network model is valuable but incomplete. The speaker addressed the definition and aspects of resilience, noting that a system can be said to be resilient if it is prepared for attack or disruption. maintains its identity, isn't compromised to the point of not being itself, bounces back, and learns from past disruptions or attacks and adapts. He also pointed out non-network aspects of resilience including the resilience of the mechanisms underlying functions, the functions themselves, the distinction between robustness (static) and resilience (dynamic), and the conditions underlying the emergence and persistence of the systems in question. With respect to the ontological aspect (the nature of resilience), he stated that the emergence of resilient norms, for example, is not amenable to network theoretic treatment but essential to security. He continued by noting that epistemic logicians and philosophy and theoretical computer science have modeled common knowledge, which is presupposed for models and norms. Plans for foundational research for science of resilience include existing foundational research and exploring the formalism. He identified open questions as: what are the constraints and factors that allow for resilience to emerge; how do we understand the role of emergent norms in the Science of Security; trust (roots of trust); and common knowledge. The researchers plan to run a series of cross-disciplinary seminars and build on KU network model for work.

Hard Problem: Metrics

Multi-model Test Bed for the Simulation-based Evaluation of Resilience—VU
Peter Volgyesi

The speaker described the existing cloud-based testbed environment for CPS developed under the Science of SecUre and REsilisant CPS (SURE) project, and proceeded to discuss new directions for future research. Areas to be explored include:

- New CPS domains (smart grid; IoT)
- Streamlined infrastructure for the Traffic CPS
- Different abstraction levels
- Hardware in the loop
- RF domain
- Transactive energy domain

He summarized the program goals as follows: integrate proven best-of-class simulators for CPS domains; add cyber security aspects (attack/defense programs); multiple levels of abstractions; collaborative design environment with versioning and libraries; and cloud-based simulation and analysis.

Safety Critical Machine Learning Algorithms—CMU
Matt Frederickson

The speaker noted that Machine Learning is ubiquitous and that it works in many applications, sometimes outperforming humans. He discussed the Deep Neural Network (DNN) model for image classifying and addressed the of an adversary that can change the features (pixels in images) that are given to the model and thereby change the outcomes (evasion attack). He raised the questions of whether attacks work if they have to be physically realizable and inconspicuous, and whether attacks can be robust to training and model selection. He presented a target attack centered around face recognition, addressing impersonation, dodging, and implementing attacks with physical changes. The challenge, he noted, is building models that are resilient to physical attacks. In addressing vulnerability, the researchers have looked at which parts of the DNN model were most susceptible to attack. They are seeking to leverage explainable features in classification to make models more resilient.

Hard Problem: Scalability and Composability

Automated Synthesis Framework for Network Security and Resilience—UIUC
Matthew Caesar

This project builds on earlier work and is focused on building a rigorous method for Science of Security, developing techniques for performing and integrating security analyses to automatically and rigorously study hypotheses about the end to end security of a network. The Automated Synthesis Framework (ASF) goal is a new network architecture for resilience with a focus on network data flow security. The approach is to leverage network synthesis to automate experiments and then apply results. The speaker identified the following three tasks:

- Network control syntheses—develop algorithms and systems that perform automated synthesis
- Network software analysis and modeling—develop frameworks for writing secure network control programs
- Resilient and self-healing network applications

With respect to their technical approach, the ASF consists of a network model, controller, policy, verification engineering, and correction engine. The project is representing network state with a policy model, and the speaker cast the problem as an optimization problem.

Monitoring, Fusion, and Response for Cyber Resilience—UIUC
Mohammad Noureddine

This project continues work done earlier, and the speaker identified the three components of the research as:

- Monitor deployment and compromise detection--monitor placement done in earlier phase, new phase looking at dealing with monitor compromise
- Rich data fusion for improved detection--prior work started at host level and incorporated more diverse data sources; since they don't know whether correlation chains are

malicious or administrative, they added new data sources from outside the network to address that question

- Automated response and recovery—the motivation is to lessen the burden on system administrators and enable response by designing autonomous agents to monitor the activity and respond; earlier work dealt with lateral movement and modeled zero-sum game and formulated same problem as control theory problem while the new work addresses puzzle difficulty selection and applies science

In the future, they plan to develop adaptive techniques to combat large-scale volumetric attacks with the goal being to push insights from control and game-theory into the reactive security realm.

Cloud-Assisted IoT Systems Privacy—KU
Fengjun Li

The speaker noted that the privacy problem is amplified in IoT because of the long and complex value chain and the large number of stakeholders included in data processing. The goal of this research is to develop a privacy threat analysis and protection framework to provide a systematic methodology for modeling and mitigating privacy threats in cloud-assisted IoT systems. Challenges include identifying which information is considered privacy and needs to be protected since privacy protection is subjective; is subjective, not all users are aware of privacy risk, and there is privacy leakage due to big data analytics. The speaker addressed privacy threats: including information disclosure, identifiability, profiling, and information linkage. The speaker identified Privacy-Enhancing Technologies (PET) as a potential solution but raised the issue of how to select and combine appropriate PETs to address identified privacy threats, with acceptable performance, within hardware, software, and data constraints. The research plan is a pilot project focused on privacy-preserving classification for cloud-assisted IoT applications. The desired research outcomes are a privacy threat analysis framework and a privacy protection framework.

Side-Channel Attack Resistance—KU
Heechul Yun

The speaker addressed the needs for Intelligent CPS and System On a chip (SOC). The speaker noted that micro-architectural side-channels in advanced embedded computing hardware are serious security threats in CPS and can compromise spatial and temporal isolation needed to implement secure and safe CPS. The project will investigate new abstractions, OS, and architecture designs for side-channel attack resistant computing platforms for CPS. The project goal is to develop micro-architectural side-channel attack resistant OS and architecture enhancements. By focusing on critical memory, the high cost of supporting strong isolation can be minimized. Tasks include critically and side-channel aware OS-level memory management on existing hardware, and new abstractions in both hardware and OS.