TRUST

Team for Research in Ubiquitous Secure Technology

Autumn 2008 Conference

November 11 – 12, 2008

Nashville Marriott at Vanderbilt University
Nashville, Tennessee

TRUST is funded by the National Science Foundation
(award number CCF-0424422)

Berkeley UNIVERSITY OF CALIFORNIA
Carnegie Mellon
Cornell University
MILLS COLLEGE
San José State UNIVERSITY
SMITH COLLEGE
STANFORD UNIVERSITY
VANDERBILT UNIVERSITY

CONTENTS

WELCOME MESSAGE

It is with great pleasure that we welcome you to the TRUST Autumn 2008 Conference in Nashville, Tennessee.

This is one of two major conferences each year that highlights activities of the TRUST Center. Specifically, work of the Center is focused on:

- Advancing a *leading-edge research agenda* to improve the state-of-the art in cyber security and critical infrastructure protection;
- Developing a *robust education plan* to teach the next generation of computer scientists, engineers, and social scientists; and
- Pursuing *knowledge transfer* opportunities to transition TRUST results to end users within industry and the government.

This conference provides an opportunity to hear firsthand about recent research results and future plans of TRUST faculty and students across all TRUST-affiliated universities. We hope you will find the conference educational, engaging, and insightful.

We are honored to have a unique keynote speaker, Federal Bureau of Investigation Supervisory Special Agent Scott E. Augenbaum. Scott heads the FBI's Cyber Crime Squad Memphis Division and will discuss the FBI's role and responsibilities for handling the emerging threats from computer intrusions and cyber crime. Scott will share with us some insight into the FBI's activities to combat these threats that are increasing in activity and sophistication and posing a larger criminal and national security threat—raising concerns from the government, the private sector, and overseas partners.

For those of you not affiliated with TRUST, or new to TRUST, I encourage you to use this conference to meet the TRUST team and find out more about the center and its projects.

Sincerely,

S. Shankar Sastry
Director, Team for Research in Ubiquitous Secure Technology
Dean of Engineering, University of California, Berkeley

TRUST OVERVIEW

The **Team for Research in Ubiquitous Secure Technology (TRUST)** is focused on the development of cyber security science and technology that will radically transform the ability of organizations to design, build, and operate trustworthy information systems for the nation's critical infrastructure. Established as a National Science Foundation Science and Technology Center (STC), TRUST is addressing technical, operational, legal, policy, and economic issues affecting security, privacy, and data protection as well as the challenges of developing, deploying, and using trustworthy systems.

TRUST activities are advancing a leading-edge *research* agenda to improve the state-of-the art in cyber security; developing a robust *education* plan to teach the next generation of computer scientists, engineers, and social scientists; and pursuing *knowledge transfer* opportunities to transition TRUST results to end users within industry and the government.

TRUST is addressing technical, operational, privacy, and policy challenges via interdisciplinary projects that combine fundamental science and applied research to deliver breakthrough advances in trustworthy systems in three "grand challenge" areas:

Financial Infrastructures – Creation of a trustworthy environment that links and supports commercial transactions among financial institutions, online retailers, and customers.

Health Infrastructures – Technology that advances "Healthcare Informatics" to enable engaged patients, personalized medicine, providers as coach-consultants, and agile evidence-based care.

Physical Infrastructures – Advances that support Next Generation Supervisory Control and Data Acquisition (SCADA) and control systems, including power, water, and telecommunications.

TRUST is led by the University of California, Berkeley with partner institutions Carnegie Mellon University, Cornell University, Mills College, San Jose State University, Smith College, Stanford University, and Vanderbilt University. TRUST projects have a holistic view that addresses computer security, software technology, analysis of complex interacting systems, and economic, legal, and public policy issues. As such, TRUST draws on researchers is such diverse fields as Computer Engineering, Computer Science, Economics, Electrical Engineering, Law, Public Policy, and the Social Sciences.

More information on TRUST is available at **http://www.truststc.org**.

CONFERENCE AGENDA

TUEDAY, NOVEMBER 11, 2008

TIME	*TOPIC*
0800 – 0845	**Breakfast**
0845 – 0900	Conference Welcome
0900 – 1000	Keynote Address – The FBI and Emerging Threats of Computer Intrusions and Cyber Crime *Scott E. Augenbaum (Cyber Crime Squad, Federal Bureau of Investigation)*
1000 – 1020	**Break**
1020 – 1040	Towards a Scalable System for Distributed Management of Private Information *Michael Siegenthaler (Cornell), Ken Birman (Cornell)*
1040 – 1100	A Model-Integrated Approach to Implementing Individualized Patient Care Plans Based on Guideline-Driven Clinical Decision Support and Process Management - A Progress Report *Jason B. Martin (Vanderbilt University Medical Center), Janos L. Mathe (Vanderbilt), Peter Miller (Vanderbilt HealthTech Laboratory), Akos Ledeczi (Vanderbilt), Liza Weavind (Vanderbilt University Medical Center), Anne Miller (Vanderbilt University Medical Center), David J. Maron (Vanderbilt HealthTech Laboratory, Vanderbilt University Medical Center), Andras Nadas (Vanderbilt), Janos Sztipanovits (Vanderbilt)*
1100 – 1120	Integration of Clinical Workflows with Privacy Policies on a Common Semantic Platform *Jan Werner (Vanderbilt), Bradley Malin (Vanderbilt), Yonghwan Lee (Vanderbilt), Akos Ledeczi (Vanderbilt), Janos Sztipanovits (Vanderbilt)*
1120 – 1140	Automatic Detection of Policies from Electronic Medical Record Access Logs *John M. Paulett (Vanderbilt), Bradley Malin (Vanderbilt)*
1140 – 1200	Fault Tolerant Sensor Network Routing for Patient Monitoring *Shanshan Jiang (Vanderbilt), Annarita Giani (Berkeley), Allen Yang (Berkeley), Yuan Xue (Vanderbilt), Ruzena Bajcsy (Berkeley)*
1200 – 1300	**Lunch**
1300 – 1320	Redundancy Minimizing Techniques for Robust Transmission in Wireless Networks *Anna Kacewicz (Cornell), Stephen Wicker (Cornell)*
1320 – 1340	DexterNet: An Open Platform for Heterogeneous Body Sensor Networks and Its Applications *Philip Kuryloski (Cornell), Annarita Giani (Berkeley), Roberta Giannantonio (WSN Lab Berkeley), Katherine Gilani (University of Texas at Dallas), Ville-Pekka Seppa (Tampere University of Technology), Edmund Seto (Berkeley), Raffaele Gravina (WSN Lab Berkeley), Victor Shia (Berkeley), Curtis Wang (Berkeley), Posu Yan (Berkeley), Allen Y. Yang (Berkeley), Jari Hyttinen (Tampere University of Technology), Shankar Sastry (Berkeley), Stephen Wicker (Cornell), Ruzena Bajcsy (Berkeley)*

CONFERENCE AGENDA (cont.)

TUEDAY, NOVEMBER 11, 2008 (continued)

TIME	*TOPIC*
1340 – 1400	An Intrusion Detection System for Wireless Process Control Systems *Adrian P. Lauf (Vanderbilt), Jonathan Wiley (Vanderbilt), Tanya Roosta (Berkeley), William H. Robinson (Vanderbilt), Gabor Karsai (Vanderbilt)*
1400 – 1420	On the Connectivity of Finite Wireless Networks with Multiple Base Stations *Sergio Bermudez (Cornell), Stephen Wicker (Cornell)*
1420 – 1440	A Security Standard for Smart Power Meters *Coalton Bennett (Cornell), Darren Highfill (Enernex Corporation), Stephen Wicker (Cornell)*
1440 – 1500	Online Information Security Education through Anchored Instruction *Eric Imsand (Memphis), Larry Howard (Vanderbilt), Ken Pence (Vanderbilt), Mike Byers (SPARTA), Dipankar Dasgupta (Memphis)*
1500 – 1520	**Break**
1520 – 1540	Bootstrapping Trust in a "Trusted" Platform *Bryan Parno (Carnegie Mellon)*
1540 – 1600	Secure Control and the Analysis of Denial of Service Attacks *Saurabh Amin (Berkeley), Alvaro Cardenas (Berkeley), Alexandre Bayen (Berkeley), Shankar Sastry (Berkeley)*
1600 – 1620	Detecting Forged TCP Reset Packets *Nicholas Weaver (International Computer Science Institute), Robin Sommer (International Computer Science Institute), Vern Paxson (Berkeley)*
1620 – 1640	Expressing and Enforcing Flow-Based Network Security Policies *Timothy Hinrichs (Chicago), Natasha Gude (Stanford), Martin Casado (Stanford), John Mitchell (Stanford), Scott Shenker (Berkeley)*
1640 – 1700	Open Problems in the Security of Learning *Marco Barreno (Berkeley), Peter L. Bartlett (Berkeley), Fuching Jack Chi (Berkeley), Anthony D. Joseph (Berkeley), Blaine Nelson (Berkeley), Benjamin I. P. Rubinstein (Berkeley), Udam Saini (Berkeley), J. D. Tygar (Berkeley)*
1700 – 1815	**TRUST Poster/Demonstration Session**
1830	**Conference Attendee Dinner** *Acorn Restaurant* *114 28th Avenue North, Nashville, TN 37203* *615-320-4399 (phone)* www.theacornrestaurant.com

CONFERENCE AGENDA (cont.)

WEDNESDAY, NOVEMBER 12, 2008

TIME	*TOPIC*
0800 – 0840	**Breakfast**
0840 – 0900	Automatic Patch-Based Exploit Generation is Possible: Techniques and Implications *David Brumley (Carnegie Mellon), Pongsin Poosankam (Carnegie Mellon), Dawn Song (Berkeley), Jiang Zheng (Pittsburgh)*
0900 – 0920	Programming with Live Distributed Objects *Krzysztof Ostrowski (Cornell), Ken Birman (Cornell), Danny Dolev (Hebrew), Jong Hoon Ahnn (Cornell)*
0920 – 0940	Smoke and Mirrors: Shadowing Files at a Geographically Remote Location Without Loss of Performance *Hakim Weatherspoon (Cornell), Lakshmi Ganesh (Cornell), Tudor Marian (Cornell), Mahesh Balakrishnan (Cornell), Ken Birman (Cornell)*
0940 – 1000	Quantitative Information Flow as Network Flow Capacity *Stephen McCamant (Berkeley), Michael D. Ernst (MIT)*
1000 – 1020	Comparison of Blackbox and Whitebox Fuzzers in Finding Software Bugs *Marjan Aslani (George Washington), Nga Chung (San Jose State), Jason Doherty (San Jose State), Nichole Stockman (Mills), William Quach (San Jose State)*
1020 – 1040	**Break**
1040 – 1100	Verifiable Functional Purity in Java *Matthew Finifter (Berkeley), Adrian Mettler (Berkeley), Naveen Sastry (Berkeley), David Wagner (Berkeley)*
1100 – 1120	An Algorithmic Approach to Authorization Rules Conflict Resolution in Software Security *Weider D. Yu (San Jose State), Ellora Nayak (San Jose State)*
1120 – 1140	Verifying the Safety of User Pointer Dereferences *Suhabe Bugrara (Stanford), Alex Aiken (Stanford)*
1140 – 1200	The TRUST-SCADA Experimental Testbed: Design and Experiments *Annarita Giani (Berkeley), Gabor Karsai (Vanderbilt), Aakash Shah (Carnegie Mellon), Bruno Sinopoli (Carnegie Mellon), Jon Wiley (Vanderbilt)*
1200 – 1205	Wrap Up / Conference Closing
1205	**Conference End *(Box Lunch Provided)***

PRESENTATION ABSTRACTS

Tuesday, November 11

0900 – 1000 **Keynote Address – The FBI and Emerging Threats of Computer Intrusions and Cyber Crime**

Scott E. Augenbaum (Cyber Crime Squad, Federal Bureau of Investigation)

1020 – 1040 **Towards a Scalable System for Distributed Management of Private Information**

Michael Siegenthaler (Cornell University)
Ken Birman (Cornell University)

In the healthcare industry and others, there is a need to electronically share privacy-sensitive data across distinct organizations. We demonstrate how this can be done while allowing organizations to keep their legacy databases and maintain ownership of the data that they currently store. Data sharing is indeed possible without sending or mirroring data to any trusted, centralized entity. In this paper we outline how queries can be executed against distributed data while revealing nothing more than the answer of the query to the organization who asked for it. Additionally, our system avoids revealing the query to the organization where the data is stored, and keeps the identity of each organization secret.

1040 – 1100 **A Model-Integrated Approach to Implementing Individualized Patient Care Plans Based on Guideline-Driven Clinical Decision Support and Process Management – A Progress Report**

Jason B. Martin (Vanderbilt University Medical Center)
Janos L. Mathe (Institute for Software Integrated Systems)
Peter Miller (Vanderbilt HealthTech)
Akos Ledeczi (Institute for Software Integrated Systems)
Liza Weavind (Vanderbilt University Medical Center)
Anne Miller (Vanderbilt University Medical Center)
David J. Maron (Vanderbilt HealthTech, Vanderbilt University Medical Center)
Andras Nadas (Institute for Software Integrated Systems)
Janos Sztipanovits (Institute for Software Integrated Systems)

Standardizing the care of patients with complex problems in hospital settings is a challenge for physicians, nurses and other medical professionals. In acute care settings such as intensive care units, the inherent problems of stabilizing and improving vital patient parameters is complicated by the division of responsibilities among different individuals and teams. The use of evidence-based guidelines for managing complex clinical problems has become the standard of practice. Computerized support for implementing such guidelines has tremendous potential. The use of model-based techniques for specifying and implementing guidelines as coordinated asynchronous processes is a promising new methodology for providing advanced clinical decision support. Combined with visual dashboards, which show the status of the implemented guidelines, a new approach to computer-supported care is possible. These techniques are being applied to the management of sepsis in acute care settings at Vanderbilt University Medical Center.

1100 – 1120 **Integration of Clinical Workflows with Privacy Policies on a Common Semantic Platform**

Jan Werner (Vanderbilt University)
Bradley Malin (Vanderbilt University)
Yonghwan Lee (Vanderbilt University)
Akos Ledeczi (Vanderbilt University)
Janos Sztipanovits (Vanderbilt University)

As healthcare organizations (HCOs) migrate to electronic systems, they must ensure compliance with complex data protection legislation, such as the Health Insurance Portability and Accountability Act (HIPAA). Legislation specifies rules that must be enforced, but regulatory Language is often imprecise, forcing HCOs to define local policies and procedures, as well as specific enforcement technologies. It is difficult for HCOs to ensure requirements are correctly translated across the enterprise, a problem compounded by the constant growth and evolution of deployed information technology (IT), such as clinical information systems (CISs). The consequence is that HCOs frequently rely on ad hoc IT configurations, which are unverified and potentially conflict with an HCO's policy. We introduce a solution to these challenges by integrating HIPAA policy rules with a domain specific model-integrated computing suite, tailored to the clinical enterprise. We present a detailed description of the policy-modeling process, the enforcement mechanism, and illustrate how to implement several policies, including mandatory access control and emergency access. All policies are formally specified through Prolog, but their enforcement is dependent on when their compliance can be evaluated. Static policies are enforced at design-time by mapping them to the structural constraints of system models. In contrast, dynamic policy rules, enforced at run-time, are loaded into a Prolog-based Policy Decision Point and Policy Enforcement Point, our extension to the standard SOA execution platform, which controls access to all services reliant upon protected health information.

1120 – 1140 **Automatic Detection of Policies from Electronic Medical Record Access Logs**

John M. Paulett (Vanderbilt University)
Bradley Malin (Vanderbilt University)

Healthcare organizations (HCOs) are increasingly adopting clinical information systems for managing patients' electronic medical records (EMRs). To support these activities, various model-based software platforms, such as Vanderbilt's Model-Integrated Clinical Information System (MICIS) have been proposed to assist in the rapid development and evaluation of formal systems based on service oriented architectures. At the same time, these systems have integrated robust privacy and security policy specification and validation languages, such as Stanford's logic based on contextual integrity. However, a significant remaining question is "what policies should be specified for data protection?" This question is difficult to address because healthcare environments are inherently dynamic, such that system have fuzzy underspecified rules, and both users and patients are constantly moving in and out of the system. This paper describes a software tool to automatically assist healthcare organizations in discovering and defining policies for access to their clinical information systems. The Healthcare Organizational Network Extraction Toolkit (HORNET) is an organization-nonspecific Java-based program that mines HCO EMR access logs to determine the underlying workflows and relationships in the system. HORNET performs this task by extracting a social network of users from the access logs and then generating association rules to indicate probabilities and strengths of associations. The system is heavily optimized to handle large networks, such as interactions between thousands of care providers. HORNET leverages novel statistical mechanisms, based on reciprocity in networks, to discover relationships between users and rules across a hospital's departments.

We evaluated HORNET with five months of access logs from the Vanderbilt University Medical Center. The sample started in January 2006 and included 9940 unique care providers and 350,889 unique patients, resulting in over 7.5 million access events. Our findings show that the network, at an individual level is highly volatile over time—82% of relationships no longer exist after 1 week and 90% no longer exist after 5 months. At a global level, though, the network remains stable, as we see a high degree of stability in our rules. We evaluated the rules for their existence and variability over time, in order to discover meaningful rules that can form the basis of defining what is normal for more advanced auditing. This duality quantifies the difficulty with which security administrators have in defining strict access policies and shows that a data mining approach can likely generate stable rules. We have successfully generated association rules which show logical and expected relationships as having high confidence and support. Our research demonstrates the feasibility of mining HCO access logs to discover underlying relationships and workflows in a dynamic setting.

1140 – 1200 **Fault Tolerant Sensor Network Routing for Patient Monitoring**

Shanshan Jiang (Vanderbilt University)
Annarita Giani (University of California, Berkeley)
Allen Yang (University of California, Berkeley)
Yuan Xue (Vanderbilt University)
Ruzena Bajcsy (University of California, Berkeley)

The use of wireless sensor networks as a means for providing remote healthcare has the potential to radically change our nation's approach to healthcare, spurring a transition from a clinic-oriented, centralized service provider paradigm to a patient-oriented, distributed healthcare system. There remain, however, significant obstacles to the use of sensor networks in a medical context. In particular, the ability to recover from failures and maintain an acceptable level of service degradation despite of failures is a crucial aspect in the design of remote healthcare systems. This paper investigates the fault tolerate routing and restoration solutions that can provide the performance assurance for the life-critical patient monitoring service despite of various failures. In particular, we consider three strategies, namely global restoration, end-to-end restoration, and local restoration, which can support a range of tradeoffs between the restoration latency and network utilization. We formulate the routing restoration problem under these three strategies as linear programming problems and build their solutions into CareNet, a two-tier wireless sensor network environment for remote healthcare. Extensive experiments are conducted to validate and evaluate our design.

1300 – 1320 **Redundancy Minimizing Techniques for Robust Transmission in Wireless Networks**

Anna Kacewicz (Cornell University)
Stephen Wicker (Cornell University)

Wireless networks are being increasingly used to download/transmit a large amount of data. This data is subject to deterioration because it is sent through the air rather than a more reliable medium. Not only are wireless networks susceptible to noise, but are also more vulnerable to malicious or un-cooperating nodes. The messages sent across a wireless network should be received without error in a short amount of time regardless of the increased disturbance levels. Hence, it is useful to devise coding methods and algorithms to ensure message integrity across wireless channels. While redundancy increases network robustness, it also reduces network efficiency. This suggests the importance of minimally increasing the message length as to guarantee an aspired degree of network reliability.

In this paper we consider a multiple path wireless network through which a source-destination pair would like to communicate. Moreover, the presence of adversaries on some paths may cause information to be lost or corrupted. We model the presence of

a malicious node as an erasure channel. The erasure channel is a type of a wireless channel, in which packets forwarded through that channel are either fully received or erased with a certain probability. For security purposes, we assume that the message is encrypted by a standard encryption technique. The erasure channel assumption is feasible since the encryption allows the destination node to validate data it receives, meaning that if the node receives corrupted data then that data will be thrown out and treated as an erasure. Malicious nodes may also steal messages, causing the receiving node to observe an erasure. The path erasure probabilities are associated with the degree to which they can be trusted.

Redundancy is vital in erasure channels since it allows perfect decoding even with some erasures. Our goal is to devise a coding method to guarantee robustness in a network with known path distributions. Wireless networks such as mobile ad hoc networks or sensor networks have frequent changes in topology due to link failures, physical obstructions, network intrusions, etc. Dependability on these sorts of networks requires dynamic algorithms which quickly determine routing, redundancy, etc. for deviations in the network. Thus, it is important that our robustness algorithm determine parameters dynamically.

Our main contribution includes the design of two algorithms MRAET (exponential time) and MRAPT (polynomial time) to determine minimum redundancy and optimal symbol allocation to attain a probability of success. We compare the performance of the algorithms with respect to each other and the desired success level. Also, we design MDS, LT, and Raptor code parameters to work in the algorithms. We implement them using the MRAET algorithm and evaluate their performance.

1320 – 1340 **DexterNet: An Open Platform for Heterogeneous Body Sensor Networks and Its Applications**

Philip Kuryloski (Cornell University)
Annarita Giani (University of California, Berkeley)
Roberta Giannantonio (WSN Lab Berkeley)
Katherine Gilani (University of Texas at Dallas)
Ville-Pekka Seppa (Tampere University of Technology)
Edmund Seto (University of California, Berkeley)
Raffaele Gravina (WSN Lab Berkeley)
Victor Shia (University of California, Berkeley)
Curtis Wang (University of California, Berkeley)
Posu Yan (University of California, Berkeley)
Allen Y. Yang (University of California, Berkeley)
Jari Hyttinen (Tampere University of Technology)
Shankar Sastry (University of California, Berkeley)
Stephen Wicker (Cornell University)
Ruzena Bajcsy (University of California, Berkeley)

We design and implement a novel platform, called DexterNet, for heterogeneous body sensor networks. The system is motivated by shifting research paradigms in body sensor networks to support real-time, persistent human monitoring in both indoor and outdoor environments. The platform adopts a three-layer, hierarchical architecture to control heterogeneous body sensors. The first layer, called the body sensor layer (BSL), deals with design of different wireless body sensors and their instrumentation on the body. We detail two custom-built body sensors, one measuring body motions and the other measuring the ECG and respiratory patterns. At the second layer, called the personal network layer (PNL), the wireless body sensors on a single subject communicate with a mobile computer station. The mobile station can be either a computer or a smart phone that supports Linux OS and the IEEE 802.15.4 protocol. It issues control commands to the body sensors and receives and processes sensor

data measured from the body sensors. These functions are abstracted and implemented as an open-source software library, called Signal Processing In Node Environment (SPINE). A DexterNet network is scalable, and can be reconfigured on-the-fly via SPINE. At the third layer, called the global network layer (GNL), multiple PNLs communicate with a remote Internet server to permanently log the sensor data and support higher-level applications in both indoor and outdoor environments. We demonstrate the versatility of the DexterNet platform via three applications: avatar visualization, human activity recognition, and integration of DexterNet with global positioning sensors and air pollution sensors for asthma studies.

1340 – 1400 **An Intrusion Detection System for Wireless Process Control Systems**

Adrian P. Lauf (Vanderbilt University)
Jonathan Wiley (Vanderbilt University)
Tanya Roosta (University of California, Berkeley)
William H. Robinson (Vanderbilt University)
Gabor Karsai (Vanderbilt University)

Wireless sensor networks employed in Supervisory Control and Data Acquisition (SCADA) networks, such as power plants, oil and gas pipelines, and industrial applications, can suffer from an inadequate provision of security resources to monitor intrusions which may threaten the normal operation of SCADA networks. While most SCADA infrastructures are equipped with an enterprise-grade firewall at the highest-level data infrastructure, recent publications suggest that such measures can be insufficient to protect the sensor and actuator networks from attacks that could render individual nodes inoperative, or cause a general failure in the overall control network.

Within the domain of the sensory and actuation device network, some nodes are being connected with low-power wireless network protocols, such as the ubiquitously-implemented 802.15.4 standard. This represents an evolution from wired networking to a mesh-networked wireless communications protocol. We begin with the assumption that such technologies are favorable in new and existing implementations, as they reduce wiring and location costs for sensory and actuation hardware. The sensors and actuators are controlled via embedded computers that receive and transmit control loop set points, and status information, respectively. The actual communication may be performed by the sensor/actuator node itself (in the case of the Intelligent Electronic Device) or by a remote terminal unit (RTU) to which the sensors are directly linked. In the case of the intelligent devices, we can use the networking capabilities of these devices to act as mesh routers for transmitting information between the farthest system node and a central access point. All communications between nodes are also performed via mesh networking.

Mesh routing allows the sensory and actuation infrastructure to be placed according to topological necessity, and given an efficient and reliable routing protocol, can ensure redundancy of data link paths among nodes. Methods such as Wireless HART exist to guarantee a high degree of link-level stability, ensuring virtually no data loss occurring on these network types. Furthermore, multiple encryption levels exist at various layers in the communications protocol that can ensure data integrity and confidentiality. However, this cannot and will not protect against all types of attacks.

Assuming that all encryption standards remain unchallenged (an assumption that will be discarded later on), jamming attacks, networking disruptions, and application-layer attacks can still be performed on the nodes without triggering any exception in the Wireless HART or similar protocol. For this reason, an intelligent monitoring system must be implemented to allow the system to identify intrusions based on deterministic information that presents itself during normal operation.

To meet this challenge, we have proposed and designed an intrusion detection system (IDS) mechanism that monitors key factors in the operation of the wireless mesh-routing network. Statistical information, such as: (1) the number of status updates, (2) mean packet size (remains nearly constant while using Wireless HART), (3) number of health packets, (4) link stability, (5) radio power usage, (6) MAC authentication failures, or (7) a difference in a packet's expected absolute serial number (ASN) can be logged and identified. This information can provide evidence for abnormal behavior, and a possible intrusion based upon a reasonable deviation of a prescribed policy. By concentrating on data specific to the networking protocol, and less on the actual plant operations, protocol-specific exploits can be more easily identified without the need for collecting data on plant operations which will vary naturally.

The IDS strategy itself relies on specialized monitoring applications that collect information on the networking protocol's operation from individual nodes and aggregates it into a logging system onboard the entity responsible for monitoring intrusions on the network. This logger is equipped with a policy management engine that reads in policy files stating acceptable operation on the network. Operation of the detection engine is separated into learning and monitoring phases, during which the system either "learns" acceptable network behavior, or begins to monitor and identify any behaviors that would signal the violation of a specified system policy. The policy and logging system itself is implemented in Java and runs on an optimized JVM designed for operation in low-power, low-resource embedded system devices.

We are currently in the process of adapting an existing SCADA client model to work with the intrusion detection mechanism. This involves taking a static pre-existing system, adding mesh-networking capabilities, and then instantiating the IDS where applicable. Data already exists for normal operation of the plant; attacks must be simulated by assessing the most critical weak points that can be targeted (i.e., which points can be jammed for maximum network disruption, or which sensors can be tampered with to cause a deviation in normal operating procedures as the result of altered data or control routines). When complete, it is expected that our system will add an extra layer of security and protection for physical infrastructures.

1400 – 1420 **On the Connectivity of Finite Wireless Networks with Multiple Base Stations**

Sergio Bermudez (Cornell University)
Stephen Wicker (Cornell University)

Connectivity is a basic quality of a network since it is required to obtain useful information out of the network. Research literature has mainly focused on the connectivity of wireless networks having a single connected network. However, due to diverse factors, it is feasible that a wireless network will have multiple base stations or sinks, i.e. it will have infrastructure. In this paper we analyze the connectivity of wireless networks with infrastructure and provide formulas to calculate the probability of having a connected network considering finite number of node platforms. With infrastructure the requirements for having a connected network may be relaxed since there can be more than one single-connected component of a network. Using combinatorics arguments we analyze the connectedness of a one-dimensional network over a line segment as a function of the available infrastructure—base stations and wireless nodes. This paper contains definitions for connected subnetworks that form a connected communication network; it also presents a one-dimensional network model and provides probability formulas for network connectivity for any given base stations arbitrary locations and the number of nodes. Through simulation we show the large increase in the probability of having a connected network when comparing a deployment with two base stations versus one with just one base station. This is of strong relevance when designing a wireless network.

1420 – 1440 **A Security Standard for Smart Power Meters**

Coalton Bennett (Cornell University)
Darren Highfill (Enernex Corporation)
Stephen Wicker (Cornell University)

There is a growing interest in 'Smart Grid' technologies in both industry and academic circles. Few attempts have been made to develop a written specification consummated with standards agreed upon by members of both coteries, due to lack of government support. Utilities in the state of California are obligated, by state legislature, to create a more efficient, reliable, and intelligent electric power system. This initiative along with Florida Power & Lighting's 'Smart Grid' pilot program has created a sense of exigency within the industry regarding smart grid technologies and standardizations. Their accomplishments are beginning to shape the policies and standards with marginal input from academic societies, ushering in a very lopsided, and business acclimatized set of standards. We will present, simulate, and analyze, a SCE 'Smart Grid' use case, in which the utilities back office applications interact with the customer's meter, and provide technical recommendations for system security improvements.

1440 – 1500 **Online Information Security Education through Anchored Instruction**

Eric Imsand (University of Memphis)
Larry Howard (Vanderbilt University)
Ken Pence (Vanderbilt Vanderbilt)
Mike Byers (SPARTA Inc.)
Dipankar Dasgupta (University of Memphis)

The Internet is unquestionably the most extensive and accessible resource for information and commerce in history. But it is also providing a medium for new forms of crime, espionage, and even terror, targeting organizations and individuals alike. Broad awareness of vulnerabilities and defenses is needed to protect against all types of cyber attacks. While online learning environments provide a great opportunity to train large numbers of people, they have yet to demonstrate effectiveness in high-stakes situations. In an effort to better prepare cyberspace defenders, we are developing a multidisciplinary training program that encompasses topics from computer science, management information systems, and legal and ethical studies, using state-of-the-art online learning methods and technology. This paper describes the Adaptive Cyber-security Training (ACT) Online program, giving details of its targeted training population, curriculum, and instructional design strategy. We further report pilot testing results from two recently developed courses that show significant learning gains following this cyber security training.

1520 – 1540 **Bootstrapping Trust in a "Trusted" Platform**

Bryan Parno (Carnegie Mellon University)

For the last few years, many commodity computers have come equipped with a Trusted Platform Module (TPM). Existing research shows that the TPM can be used to establish trust in the software executing on a computer. However, at present, there is no standard mechanism for establishing trust in the TPM on a particular machine. Indeed, any straightforward approach falls victim to a cuckoo attack. In this work, we propose a formal model for establishing trust in a platform. The model reveals the cuckoo attack problem and suggests potential solutions. Unfortunately, no instantiation of these solutions is fully satisfying, and hence, we pose the development of a fully satisfactory solution as an open question to the community.

1540 – 1600

Secure Control and the Analysis of Denial of Service Attacks

Saurabh Amin (University of California, Berkeley)
Alvaro Cardenas (University of California, Berkeley)
Alexandre Bayen (University of California, Berkeley)
Shankar Sastry (University of California, Berkeley)

Control systems are computer-based systems that monitor and control physical processes. These systems represent a wide variety of networked information technology (IT) systems connected to the physical world. Depending on the application, these control systems are also called Process Control Systems (PCS), Supervisory Control and Data Acquisition (SCADA) systems (in industrial control or in the control of the critical infrastructures), or Cyber-Physical Systems (CPS) (to refer to embedded sensor and actuator networks).

Several control applications can be labeled as safety-critical: their failure can cause irreparable harm to the physical system being controlled and to the people who depend on it. SCADA systems, in particular, perform vital functions in national critical infrastructures, such as electric power distribution, oil and natural gas distribution, water and waste-water treatment, and transportation systems. They are also at the core of health-care devices, weapons systems, and transportation management. The disruption of these control systems could have a significant impact on public health, safety and lead to large economic losses. Control systems have been at the core of critical infrastructures and industrial plants for many decades, and yet, there have been very few confirmed cases of cyberattacks. Control systems, however, are more vulnerable now than before to computer vulnerabilities for many reasons, such as the use of commodity IT solutions, corporate network interconnections, they are more and more ubiquitous and there is an increasing number of people capable of launching computer attacks on control systems with different motivations including disgruntled employees, cyber crime, extortion, terrorism etc.

While it is clear that the security of control systems has become an active area in recent years, we believe that, so far, no one has been able to articulate what is new and fundamentally different in this field from a research point of view compared to traditional IT security. In general, information security has developed mature technologies and design principles (authentication, access control, message integrity, separation of privilege, etc.) that can help us prevent and react to attacks against control systems. However, research in computer security has focused traditionally on the protection of information. Researchers have not considered how attacks affect the estimation and control algorithms -and ultimately, how attacks affect the physical world. We argue that while the current tools of information security can give necessary mechanisms for the security of control systems, these mechanisms alone are not sufficient for the defense-in-depth of control systems.

We believe that by understanding the interactions of the control system with the physical world, we should be able to (1) better understand the consequences of an attack: so far there is no research on how an adversary would select an strategy once it has obtained unauthorized access to some control network devices; (2) design novel attack-detection algorithms: by understanding how the physical process should behave based on our control commands and sensor measurements, we can identify if an attacker is tampering with the control or sensor data; and (3) design new attack-resilient algorithms and architectures: if we detect an attack we may be able to change the control commands to increase the resiliency of the system.

To this end, we provide a taxonomy of attacks to control systems at the "systems level". We then focus our technical work in the analysis of Denial of Service Attacks.

Our two main contributions are (1) we formulate a new problem of denial of service attacks on **predictive control** of linear systems that need to satisfy certain **safety constraints**. We chose predictive control because it is one of the most used control algorithms in industrial settings. We also included safety constraints to describe security specifications. (2) we show how by analyzing a wider class of attack options than the ones assumed in networked control, an adversary can create novel attacks that have higher negative effects on the performance of the control system. We also show how to design an optimal control policy for these worst-case attacks.

Other technical contributions of this work are (1) we consider a set of affine state feedback policies. This class generalizes previous work; (2) we derive the Kalman filter equations for predictive control when there is a denial of service attack; (3) we prove that with an error feedback parameterization the closed loop response is affine in the control parameters; and (4) we show how our problem can be posed as a convex problem. In particular, we show that the optimal control can be found as the solution of a semi-definite program. This is an improvement over previous work which has solved similar problems using dynamic programming.

1600 – 1620 **Detecting Forged TCP Reset Packets**

Nicholas Weaver (International Computer Science Institute)
Robin Sommer (International Computer Science Institute)
Vern Paxson (University of California, Berkeley)

Several off-the-shelf products enable network operators to enforce usage restrictions by actively terminating connections when deemed undesirable. While the spectrum of their application is large—from ISPs limiting the usage of P2P applications to the "Great Firewall of China"—many of these systems implement the same approach to disrupt the communication: they inject artificial TCP Reset (RST) packets into the network, causing the endpoints to shut down communication upon receipt. In this work, we study the characteristics of packets injected by such traffic control devices. We show that by exploiting the race-conditions that out-of-band devices inevitably face, we not only can detect the interference but often also fingerprint the specific device in use. We develop an efficient injection detector and demonstrate its effectiveness by identifying a range of disruptive activity seen in traces from four different sites, including termination of P2P connections, anti-spam and anti-virus mechanisms, and the finding that China's "Great Firewall" has multiple components, sometimes apparently operating without coordination. We also find a number of sources of idiosyncratic connection termination that do **not** reflect third-party traffic disruption, including NATs, load-balancers, and spam bots. In general, our findings highlight that (1) Internet traffic faces a wide range of control devices using injected RST packets, and (2) significant care is required to reliably detect RST injection while avoiding misidentification of other types of activity.

1620 – 1640 **Expressing and Enforcing Flow-Based Network Security Policies**

Timothy Hinrichs (University of Chicago)
Natasha Gude (Stanford University)
Martin Casado (Stanford University)
John Mitchell (Stanford University)
Scott Shenker (University of California, Berkeley)

While traditional network security policies have been enforced by manual configuration of individual network components such as router ACLs, firewalls, NATs and VLANs, emerging enterprise network designs and products support global policies declared over high level abstractions. We further the evolution of simpler and more powerful network security mechanisms by designing, implementing, and testing a flow-based network security policy language and enforcement infrastructure. Our policy language, FSL, expresses basic network access controls, directionality in communication

establishment (similar to NAT), network isolation (similar to VLANs), communication paths, and rate limits. FSL supports modular construction, distributed authorship, and efficient implementation. We have implemented FSL as the primary policy language for NOX, a network-wide control platform, and have deployed it within an operational network for over 10 months. We describe how supporting complex policy objectives and meeting the demanding performance requirements of network-wide policy enforcement have influenced the FSL language design and implementation.

1640 – 1700 **Open Problems in the Security of Learning**

Marco Barreno (University of California, Berkeley)
Peter L. Bartlett (University of California, Berkeley)
Fuching Jack Chi (University of California, Berkeley)
Anthony D. Joseph (University of California, Berkeley)
Blaine Nelson (University of California, Berkeley)
Benjamin I. P. Rubinstein (University of California, Berkeley)
Udam Saini (University of California, Berkeley)
J. D. Tygar (University of California, Berkeley)

Machine learning has become a valuable tool for detecting and preventing malicious activity. However, as more applications employ machine learning techniques in adversarial decision-making situations, increasingly powerful attacks become possible against machine learning systems. In this paper, we present three broad research directions towards the end of developing truly secure learning. First, we suggest that finding bounds on adversarial influence is important to understand the limits of what an attacker can and cannot do to a learning system. Second, we investigate the value of adversarial capabilities—the success of an attack depends largely on what types of information and influence the attacker has. Finally, we propose directions in technologies for secure learning and suggest lines of investigation into secure techniques for learning in adversarial environments. We intend this paper to foster discussion about the security of machine learning, and we believe that the research directions we propose represent the most important directions to pursue in the quest for secure learning.

Wednesday, November 12

0840 – 0900 **Automatic Patch-Based Exploit Generation is Possible: Techniques and Implications**

David Brumley (Carnegie Mellon University)
Pongsin Poosankam (Carnegie Mellon University)
Dawn Song (University of California, Berkeley)
Jiang Zheng (University of Pittsburgh)

The automatic patch-based exploit generation problem is: given a program P and a patched version of the program P', automatically generate an exploit for the potentially unknown vulnerability present in P but fixed in P'. In this paper, we propose techniques for automatic patch-based exploit generation and show that our techniques can automatically generate exploits for five Microsoft programs based upon patches provided via Windows Update. Although our techniques may not work in all cases, a fundamental tenant of security is to conservatively estimate the capabilities of attackers. Thus, our results indicate that automatic patch-based exploit generation should be considered practical. One important security implication of our results is that current patch distribution schemes which stagger patch distribution over long time periods, such as Windows Update, may allow attackers who receive the patch first to compromise the significant fraction of vulnerable hosts who have not yet received the patch.

0900 – 0920 **Programming with Live Distributed Objects**

Krzysztof Ostrowski (Cornell University)
Ken Birman (Cornell University)
Danny Dolev (Hebrew University)
Jong Hoon Ahnn (Cornell University)

A component revolution is underway, bringing developers improved productivity and opportunities for code reuse. However, whereas existing tools work well for builders of desktop applications and client-server structured systems, support for other styles of distributed computing has lagged. In this paper, we propose a new programming paradigm and a platform, in which instances of distributed protocols are modeled as "live distributed objects". Live objects can represent both protocols and higher-level components. They look and feel much like ordinary objects, but can maintain shared state and synchronization across multiple machines within a network. Live objects can be composed in a typesafe manner to build sophisticated distributed applications using a simple, intuitive drag and drop interface, very often without writing any code or having to understand the intricacies of the underlying distributed algorithms.

0920 – 0940 **Smoke and Mirrors: Shadowing Files at a Geographically Remote Location Without Loss of Performance**

Hakim Weatherspoon (Cornell University)
Lakshmi Ganesh (Cornell University)
Tudor Marian (Cornell University)
Mahesh Balakrishnan (Cornell University)
Ken Birman (Cornell University)

The Smoke and Mirrors File System (SMFS) mirrors files at geographically remote datacenter locations with negligible impact on file system performance at the primary site, and minimal degradation as a function of link latency. It accomplishes this goal using wide-area links that run at extremely high speeds, but have long round-trip-time latencies—a combination of properties that poses problems for traditional mirroring solutions. In addition to its raw speed, SMFS maintains good synchronization: should the primary site become completely unavailable, the system minimizes loss of work, even for applications that simultaneously update groups of files. We present the SMFS design then evaluate the system on Emulab. Intended applications include file sharing in wide-area settings and remote backup for disaster recovery.

0940 – 1000 **Quantitative Information Flow as Network Flow Capacity**

Stephen McCamant (University of California, Berkeley)
Michael D. Ernst (Massachusetts Institute of Technology)

We present a new technique for measuring information flow: the extent to which a program's outputs reveal information about or are influenced by its inputs. In contrast to previous techniques based on reachability from secret inputs (tainting), it achieves a more precise quantitative result by computing a maximum flow of information between the inputs and outputs. The technique uses static control-flow regions to soundly account for implicit flows via branches and pointer operations, but operates dynamically by observing one or more program executions and giving numeric flow bounds specific to them (e.g., "17 bits"). We performed case studies on 5 real C, C++, and Objective C programs, 3 of which had more than 250K lines of code. The tool checked multiple security policies, including one that was violated by a previously unknown bug. In addition to the maximum flow technique for computing upper bounds, we'll also describe a decision-procedure technique for computing corresponding lower bounds.

1000 – 1020 **Comparison of Blackbox and Whitebox Fuzzers in Finding Software Bugs**

Marjan Aslani (George Washington University)
Nga Chung (San Jose State University)
Jason Doherty (San Jose State University)
Nichole Stockman (Mills College)
William Quach (San Jose State University)

Both blackbox and whitebox fuzzing techniques have been widely used to uncover security vulnerabilities in software applications, but there have been few studies comparing each technique. Our approach was to use Zzuf, a blackbox fuzzer, and Catchconv, a whitebox fuzzer, to generate test cases that were then run on open source and commercial software to compare both fuzzers efciency in terms of the number of unique bugs found per test case. An analysis of our results showed that Zzuf found an average of 2:69 unique errors per 100 unique test cases, while Catchconv found an average of 2:63 unique errors per 100 test cases. In terms of unique errors per total errors, 22 percent of the total errors found by Catchconv were unique, while 0:05 percent of the total errors found by Zzuf were unique. From the analysis of the data we collected, we identified metrics which we suggest for future comparison between fuzzers, but we did not collect enough information to evaluate in our study.

1040 – 1100 **Verifiable Functional Purity in Java**

Matthew Finifter (University of California, Berkeley)
Adrian Mettler (University of California, Berkeley)
Naveen Sastry (University of California, Berkeley)
David Wagner (University of California, Berkeley)

Proving that particular methods within a code base are functionally pure—deterministic and side-effect free—would aid verification of security properties including function invertibility, reproducibility of computation, and safety of untrusted code execution. Until now it has not been possible to automatically prove a method is functionally pure within a high-level imperative language in wide use, such as Java. We discuss a technique to prove that methods are functionally pure by writing programs in a subset of Java called Joe-E; a static verifier ensures that programs fall within the subset. In Joe-E, pure methods can be trivially recognized from their method signature.

To demonstrate the practicality of our approach, we refactor an AES library, an experimental voting machine implementation, and an HTML parser to use our techniques. We prove that their top-level methods are verifiably pure and show how this provides high-level security guarantees about these routines. Our approach to verifiable purity is an attractive way to permit functional-style reasoning about security properties while leveraging the familiarity, convenience, and legacy code of imperative languages.

1100 – 1120 **An Algorithmic Approach to Authorization Rules Conflict Resolution in Software Security**

Weider D. Yu (San Jose State University)
Ellora Nayak (San Jose State University)

Conflicts can occur in access control models due to many reasons. Conflict resolution in real-time can be a challenging task due to the complexity of the access control rules. For a large number of access rules for a given service, detecting exactly which rules caused conflicts can be a daunting task. This paper describes an algorithm to resolve conflicts using the ARSL (Authorization Rule Specification Language) model. The algorithm is designed to work on the complexity issues of conflict resolution by preventing the occurrences of conflicts. The algorithm is based on priority of the authorization rules for a given resource and the priority is based on the sequence of occurrences of the authorization rules specified in the ARSL input file.

1120 – 1140 **Verifying the Safety of User Pointer Dereferences**

Suhabe Bugrara (Stanford University)
Alex Aiken (Stanford University)

Operating systems divide virtual memory addresses into kernel space and user space. The interface of a modern operating system consists of a set of system call procedures that may take pointer arguments called user pointers. It is safe to dereference a user pointer if and only if it points into user space. If the operating system dereferences a user pointer that does not point into user space, then a malicious user application could gain control of the operating system, reveal sensitive data from kernel space, or crash the machine. Because the operating system cannot trust user processes, the operating system must check that the user pointer points to user space before dereferencing it. In this paper, we present a scalable and precise static analysis capable of verifying the absence of unchecked user pointer dereferences. We evaluate an implementation of our analysis on the entire Linux operating system with over 6.2 million lines of code with false alarms reported on only 0.05% of dereference sites.

1140 – 1200 **The TRUST-SCADA Experimental Testbed: Design and Experiments**

Annarita Giani (University of California, Berkeley)
Gabor Karsai (Vanderbilt University)
Aakash Shah (Carnegie Mellon University)
Bruno Sinopoli (Carnegie Mellon University)
Jon Wiley (Vanderbilt University)

The TRUST SCADA Testbed is an affordable software-hardware infrastructure that supports experimentation with systems-level security technologies for SCADA systems. The testbed, in its current implementation, consists of (1) a realistic plant simulator, which uses Simulink/Stateflow dynamic system models for real-time simulation of physical plants (e.g., chemical manufacturing processes, power generation and distribution systems, oil refineries), (2) low-cost SCADA RTU emulator boards that run low-level regulator/SCADA software and are connected to the plant simulator, and (3) affordable, networked SCADA host emulator boards that could run higher-level control and optimization algorithms which provide setpoints and control commands to the RTU emulators. The presentation will describe the current status of the testbed, the chemical plant model that we used for test examples, and the initial experiments we have performed.

KEYNOTE SPEAKER BIOGRAPHY

Scott E. Augenbaum

Supervisory Special Agent, Cyber Crime Squad
Federal Bureau of Investigation

Supervisory Special Agent Scott E. Augenbaum is the supervisor of the FBI's Memphis Division Cyber Squad, based in Nashville, TN. SSA Augenbaum manages the Memphis Crimes Against Children Task Force made up of two FBI Special Agents, an Investigator from the Memphis Police Department and Shelby County Sheriff's Office, and a Special Agent from the Immigrations and Customs Enforcement (ICE). As Supervisory Special Agent, Augenbaum formed a Joint Cyber Crime Task Force with the Tennessee Bureau of Investigation and Franklin, TN Police Department to investigate Computer Intrusion, Online Child Exploitation, Intellectual Property Rights and Internet Fraud.

Augenbaum started his career in the FBI's New York Office in 1988, spending six years as an Operations Assistant and Accounting Technician for the Foreign Counter Intelligence Branch. Augenbaum worked on his first Cyber investigation in the fall of 1995 as part of the Innocent Images Online Child Exploitation initiative. For nine years, Augenbaum was the Case Agent on numerous Internet Fraud, Computer Intrusion, and Innocent Images investigations. In this role, he established and maintained liaison with industry, academia, and law enforcement and formed a working group with over 150 individuals/agencies committed to protecting the critical Infrastructure in upstate New York from Cyber attack. After the September 11th attacks, Augenbaum created an ad-hoc task force of Federal, State and local law enforcement agencies which lead to the formation of the Joint Terrorism Task Force in Syracuse, New York.

In December 2003, Augenbaum was promoted to a Supervisory Special Agent at FBIHQ, Cyber Division. He was the Program Manager of the FBI's Cyber Task Force Program—an initiative that enabled the FBI to partner with Federal, State and local law enforcement to combat the emerging threat of Computer Intrusion/Cyber Crime—and was responsible for the program within the FBI's 56 Field Offices. Augenbaum also had Program Management responsibilities for the FBI's Intellectual Property Rights Program and maintained liaison with the industry partners on IPR issues.

Augenbaum graduated from the City College of New York in 1992 and worked on his Master of Business Administration (MBA) in Information Technology and Finance at Fordham University's Lincoln Center Campus, New York, NY.

SPEAKER BIOGRAPHIES

Saurabh Amin

University of California, Berkeley

Saurabh is a fifth year graduate student in the Civil and Environmental Engineering systems engineering program. Before coming to UC Berkeley, Saurabh studied civil engineering at the Indian Institute of Technology, Roorkee (formerly, University of Roorkee). He then studied transportation engineering at the University of Texas at Austin, where he obtained an M.S.E. in the area of infrastructure systems. He moved to Berkeley in fall 2004 for his Ph.D. studies and is currently interested in control of hybrid systems, robust optimization, boundary control of hyperbolic PDEs, and reachability analysis for stochastic systems. His research has been supervised by Prof. S. Shankar Sastry and Prof. Alexandre M. Bayen.

Marjan Aslani

George Washington University

Marjan Aslani is a senior in electrical engineering at the George Washington University and her area of interest is design and implementation of large scale integration circuits and microelectronic systems. During the summer of 2008, Marjan had the opportunity to do research at the TRUST Center as a participant in the Summer Undergraduate Program in Engineering Research at Berkeley (SUPERB). This internship provided an excellent opportunity to work in a multidisciplinary team with a diverse mixture of skills and knowledge and learn about different aspects of cyber security and the challenges of building and operating trustworthy systems.

Sergio Bermudez

Cornell University

Sergio Bermudez is a graduate student in the School of Electrical and Computer Engineering at Cornell University. He received his B.S. in Electrical and Communications Engineering from Monterrey Institute of Technology, Mexico. His work experience is in the front-end of manufacturing as a test, process, and project engineer. His research interests include the analysis of topology, mobility, and security of wireless sensor networks.

Coalton Bennett

Cornell University

Coalton Bennett earned the B.S.E.E. degree from Howard University in May of 2005. He graduated with honors from Howard University and is a member of the Tau Beta Pi international honor society. He received both a Sloan and Cornell Graduate fellowship. He is currently a fourth year M.S./Ph.D. student in the Electrical and Computer Engineering Department at Cornell University working under the tutelage of Dr. Stephen B. Wicker. Coalton's research interests include SCADA systems, Wireless Sensor Actuator Networks (WSANs) and their applications to SCADA systems, as well as security in wireless sensor networks. Currently he is interested in the integration of WSANs and Smart Grid communication systems for: commercial, industrial, and residential demand response systems. In particular the development of reliable protocols, for communication between wirelessly enabled Smart Meters and utility back office computer systems, which are compliant with the ANSI C12.22 and ANSI C12.19 standards. Coalton has also worked closely with Dr. Judith Cardell of Smith College, in the development of a testbed for wirelessly controlled appliances using wireless sensor networks.

Suhabe Bugrara

Stanford University

Suhabe Bugrara is a fourth year Ph.D. student in the Computer Science Department at Stanford University. His research interest is automated tools for software reliability. He graduated from the Massachusetts Institute of Technology in 2005 with a B.Sc. in computer science.

Katherine Gilani

University of Texas at Dallas

Katherine Gilani is an Electrical Engineering Student at the University of Texas at Dallas. She spent one year as a Research Assistant for the Embedded Systems and Signal Processing Lab at UTD; followed by a short summer at the University of California, Berkeley researching as a member of TRUST. She is currently the Secretary of IEEE at UTD, a member of SWE and ACM, as well as a student leader in the Student Advisory Council. She believes in making the most of her academic career and is engaged in the research involved.

Timothy Hinrichs

University of Chicago

Tim Hinrichs received a B.S. in Computer Science from the University of Illinois at Urbana-Champaign in 2001 and a Ph.D. in Computer Science from Stanford University in 2007. He is currently a postdoctoral researcher at the University of Chicago focusing on Logic and Automated Reasoning, specifically on language design and compilation.

Larry Howard

Vanderbilt University

Larry Howard is a Senior Research Scientist at Vanderbilt University's Institute for Software Integrated Systems (ISIS), where he explores technologies for individually scaffolded learning via the Internet. His project has designed the TRUST Academy Online (TAO) portal to support the Center's web-based dissemination of educational materials. Under a grant from the U.S. Department of Homeland Security, the project has contributed to creating online training resources that address cyber-preparedness for risk managers, IT professionals, and general computer users.

Shanshan Jiang

Vanderbilt University

Shanshan Jiang received her B.S. degree in 2002 and M.S. degree in 2005 both from the Department of Computer Science, University of Science and Technology of China. Currently she is a Ph.D. candidate in the Department of Electrical Engineering and Computer Science and the Institute for Software Integrated Systems (ISIS) at Vanderbilt University. Her research area includes distributed and networking systems, with a focus on distributed service development, deployment, and communication in wireless sensor and mobile networks with quality of service and reliability support.

Anna Kacewicz

Cornell University

Anna received her undergraduate degree in electrical engineering from University of Texas at Austin in 2005. She came to Cornell University in fall 2005 and began research in coding and information theory focusing on networks. She is a member of Professor Wicker's Wireless Intelligence Systems Lab where she is actively advancing the application of coding theory to robustness and security in wireless networks.

Adrian Lauf

Vanderbilt University

Adrian Lauf is a Graduate Student in the Department of Electrical Engineering and Computer Science at Vanderbilt University. He received his B.E. in Computer Engineering at Vanderbilt University in 2005, his M.S. in Electrical Engineering in 2007 at Vanderbilt University, and is continuing his studies for a Ph.D. in Electrical Engineering. His research thrust is harbored within the Institute for Software Integrated Systems (ISIS), a research institute that's part of the Electrical Engineering and Computer Science department at Vanderbilt University. As part of TRUST, he researches computer security and intrusion detection mechanisms applicable to networked, embedded device platforms. Such work can be applied to both civilian and military applications, scalable from a small network size of less than 10 agents to large collections spanning thousands of nodes. He is a student member of the IEEE.

Janos Mathe

Vanderbilt University

Janos L. Mathe received his M.Sc. degree in Computer Science at the Technical University of Budapest in 2004. He continued his studies by enrolling to the Department of Electrical Engineering and Computer Science at Vanderbilt University where he is currently pursuing his Ph.D. under the guidance of Janos Sztipanovits. Janos is interested in applying Model-Integrated Computing techniques to address the security and privacy requirements in healthcare settings. His current research focuses on the model-based development of clinical information systems where he investigates how modeling, validation, verification and deployment of treatment protocols can be performed by using the example of sepsis management.

Stephen McCamant

University of California, Berkeley

Stephen McCamant is a postdoctoral scholar at the University of California, Berkeley, working with Professor Dawn Song and her research group on binary-level program analysis techniques for software security. From 2002 to 2008, he was a graduate student at the Massachusetts Institute of Technology's Computer Science and AI Lab. His research at MIT applied dynamic and static analysis techniques to the challenges of software engineering and security with Professor Michael Ernst and colleagues in the Program Analysis Group, including predicting incompatible software upgrades and measuring leaks of secret information.

Blaine Nelson

University of California, Berkeley

Blaine Nelson received his B.S. in Computer Science from the University of South Carolina at Columbia in 2003 and was admitted to UC Berkeley's Computer Science graduate program in that year. He did internships in Duke's Research Experience for Undergraduates program in 2002 and at HP Labs in 2005. He received his M.S. in Computer Science from Berkeley in 2005 and is currently pursuing his Ph.D. working with Anthony Joseph, Doug Tygar, and other researchers at UC Berkeley in studying the behavior of machine learning algorithms in security sensitive environments where data may be influenced by an adversary.

Krzysztof Ostrowski

Cornell University

Krzys Ostrowski is a Postdoctoral Associate at the Computer Science department at Cornell University. He received his B.Sc. in Mathematics and M.Sc. in Computer Science from Warsaw University, Poland in 1998 and 2001, and Ph.D. in Computer Science from Cornell University in 2008. His research focuses on scalable architectures for distributed multi-party protocols and services, modeling reliability and other types of strong guarantees in ways that permit scalable implementations, and on language, type system and runtime support for scalable distributed architectures. His ultimate vision is a strongly-typed programming environment in which the Internet plays the role of a virtual machine, where distributed protocol instances and their constituent components can be treated as casually as threads and variables in Java, and in which a cleanly structured protocol code implementing scalable agreement, atomic commit, or service replication would fit on a monitor screen and be understandable at a glance.

Bryan Parno

Carnegie Mellon University

Bryan Parno is a Ph.D. student in the Electrical & Computer Engineering Department at Carnegie Mellon University. Bryan works with his advisor, Adrian Perrig, and is supported by a National Defense Science and Engineering Graduate (NDSEG) Fellowship, the Department of Homeland Security, and the National Science Foundation. Bryan studied Computer Science at Harvard University and is currently focus on network and host-based security as well as applied cryptography.

John Paulett

Vanderbilt University

John Paulett is a Masters of Science candidate at Vanderbilt University's Department of Biomedical Informatics. He graduated from the University of Pennsylvania with a Bachelor of Science in Engineering, majoring in Bioengineering. His current research is focused on access log mining for workflow analysis and anomaly detection. He additionally works on care-provider notification systems for Vanderbilt University Medical Center. He previously worked on improving radiology speech recognition for radiologists and RadLex, a terminology for radiologists.

Michael Siegenthaler

Cornell University

Michael Siegenthaler is a third year Ph.D. student in Computer Science at Cornell University. He received his B.S. degree in Electrical Engineering from the University of California, Davis in 2006. Michael works with Professor Ken Birman on privacy and security aspects of distributed systems.

Dawn Song

University of California, Berkeley

Dawn Song is an Assistant Professor at University of California, Berkeley. She obtained her Ph.D. in Computer Science from UC Berkeley (2002). Prior to joining UC Berkeley, she was an Assistant Professor at Carnegie Mellon University from 2002 to 2007. Her research interest lies in security and privacy issues in computer systems and networks. She is the author of more than 70 research papers in areas ranging from software security, networking security, database security, distributed systems security, to applied cryptography. She is the recipient of various awards including the NSF CAREER Award, the IBM Faculty Award, the George Tallman Ladd Research Award, the Sloan Award, the Okawa Foundation Research Grant Award, and Best Paper Awards in top security conferences.

David Wagner

University of California, Berkeley

David Wagner is an Associate Professor in the Computer Science Division at the University of California, Berkeley, working in the areas of computer security and electronic voting. He and his Berkeley colleagues are known for discovering a wide variety of security vulnerabilities in various cellphone standards, 802.11 wireless networks, electronic voting systems, and other widely deployed systems. Last year, he helped lead a comprehensive review commissioned by California Secretary of State Debra Bowen to examine three California e-voting systems. David is a member of the Election Assistance Commission's Technical Guidance Development Committee, the federal advisory board charged with helping to draft future voting standards.

Hakim Weatherspoon

Cornell University

Hakim Weatherspoon is an Assistant Professor in Computer Science at Cornell University. His research interests cover various aspects of distributed systems: fault-tolerance, reliability, security, and performance of Internet-scale systems with decentralized—autonomous, federated, multi-organizational, and cooperative—control. He graduated from the University of California at Berkeley with a Ph.D. in Computer Science, supported by a Intel Foundation PhD Fellowship. His full CV and publication list can be found at http://www.cs.cornell.edu/~hweather.

Nicholas Weaver

International Computer Science Institute

Nicholas Weaver is a researcher at the International Computer Science Institute in Berkeley, California, specializing in network attacks, intrusion detection, and malcode, after having received his Ph.D. from the University of California, Berkeley in Computer Architecture in 2003. He also possesses a very devious mind.

Jan Werner

Vanderbilt University

Jan Werner is a Ph.D. student in the Department of Electrical Engineering and Computer Science at Vanderbilt University. His research area includes Model Integrated Computing with a focus on privacy and security policy models. He conducts research on integrating security privacy and security in workflows for next generation clinical information systems. He received M.S. degree in Computer Science and B.S. in Applied physics from Nicolaus Copernicus University, Poland in 2005.

Jonathan Wiley

Vanderbilt University

Jon Wiley is a second year graduate student studying Electrical Engineering at Vanderbilt University. He received his B.E. in Computer Engineering from Vanderbilt University in 2007. His research interests include securing, modeling, and simulating real-time systems, particularly SCADA systems.

Posu Yan

University of California, Berkeley

Posu is currently a staff engineer in the Electrical Engineering and Computer Sciences department at the University of California, Berkeley. After graduating from EECS at UC Berkeley and receiving an M.S. in computer science at UCLA, he is now mainly involved in the areas of wireless camera motes and wireless body sensors. Aside from engineering, he spends his free time pursuing music and other artistic endeavors.

Weider Yu

San Jose State University

Dr. Weider D. Yu is an associate professor in the Computer Engineering Department at San Jose State University, San Jose (Silicon Valley), California. He received an M.S. in Computer Science from the State University of New York at Albany, and a Ph.D. from Northwestern University in Electrical Engineering and Computer Science. He also attended the MBA program in the Graduate School of Business at University of Chicago and received a certificate in information security engineering from Carnegie Mellon University. Prior to the university, Dr. Yu was a Distinguished Member of Technical Staff at Bell Laboratories and an adjunct associate professor in the department of Electrical Engineering and Computer Science, University of Illinois at Chicago.

Dr. Yu performs his research and teaching in the areas of distributed software systems, experimental software engineering, wireless mobile and web based software systems, software security, quality and reliability related software processes, service oriented software engineering, and systems performance factors. Dr. Yu has publications on Bell Labs Technical Journal, AT&T Technical Journal, IEEE Journal of Selected Areas in Communications, and various international IEEE conferences.

NOTES

NOTES (cont.)