TRUST

Team for Research in Ubiquitous Secure Technology

Autumn 2007 Conference

October 10-11, 2007

Ithaca, New York

CONTENTS

CONTENTS 3

WELCOME MESSAGE 4

TRUST OVERVIEW 5

CONFERENCE AGENDA 6

PRESENTATION ABSTRACTS 8

KEYNOTE SPEAKER BIOGRAPHY 21

SPEAKER BIOGRAPHIES 22

NOTES 29

WELCOME MESSAGE

It is with great pleasure that we welcome you to the TRUST Autumn 2007 Conference in Ithaca, New York. I would like to thank TRUST partner institution Cornell University for hosting this event.

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

- Advancing a *leading-edge research agenda* to improve the state-of-the art in cybersecurity 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 past results and future plans of TRUST faculty and students in each area. We hope you will find the conference educational, engaging, and insightful.

We are honored to have as a keynote speaker Bob Sullivan, an author, blogger, and technology writer for MSNBC. Bob will share his unique insight into issues associated with security and privacy and will discuss recent events affecting companies and individuals.

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 role and penetration of computing systems and networks in our societal infrastructure continues to grow, and their importance to societal safety and the security has never been greater. Beyond mere connection to the Internet and access to global resources, information systems are now used for controlling critical infrastructures for electricity, healthcare, finance, and medical networks. As society uses computers, systems, and networks in increasingly important ways, the underlying technology provided often does not meet the desired level of trust and many critical infrastructure systems remain untrustworthy. Viruses and worms sweep the Internet and exhibit increasing virulence and a rate of speed that is directly proportional to their growing ease of deployment. Privacy and security remain poorly understood, poorly supported, and generally inadequate. Broader issues of software usability, reliability, and correctness remain challenging as does understanding how users interact with computers and ways in which systems can be designed to influence users to behave in a more secure manner.

The **Team for Research in Ubiquitous Secure Technology** (TRUST) is addressing these challenges of developing, deploying, and using trustworthy systems. TRUST, a National Science Foundation sponsored Science and Technology Center (STC) is focused on the development of cybersecurity science and technology that will radically transform the ability of organizations to design, build, and operate trustworthy information systems for our state and nation's critical infrastructure.

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, interdisciplinary view that address 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.

TRUST is addressing fundamental problems and advancing the state-of-the-art in a number of areas:

- Security and privacy issues associated with the rapidly increasing use of electronic media for the archival and access of patient medical records.
- Web authentication, end-user privacy, next-generation browser security, malware detection, and improved system forensic techniques to combat online attacks.
- Application defenses for network-level intrusions and attacks including compromised and malfunctioning legacy applications, viruses, worms, and spyware.
- Incentives for research, investment, policies, and procedures for technology that enhance system security, privacy, and trustworthiness.
- Secure embedded sensor networks for large-scale applications critical to the nation’s economy, energy, security, and health.
- Techniques that ensure trustworthy computing by securing hardware, improving software robustness, and increasing the survivability of critical systems.

More information on TRUST is available at www.truststc.org.

CONFERENCE AGENDA

WEDNESDAY, OCTOBER 10, 2007	
TIME	*TOPIC*
0800 – 0830	**Breakfast [Owego/Brindley Room]**
0830 – 0900	**Conference Welcome** *Shankar Sastry (Director, TRUST and Dean of Engineering, UC Berkeley)*
0900 – 0930	**Selective Sensing and the 4th Amendment** *Deirdre Mulligan (UC Berkeley), Stephen Wicker (Cornell University)*
0930 – 1000	**Competition and Fraud in Online Advertising Markets** *Bob Mungamure (Stanford University), Stephen Holt (Google Inc.)*
1000 – 1030	**Network Security and the Need to Consider Provider Coordination in Network Access Policy** *Aaron J. Burstein (UC Berkeley), Fred B. Schneider (Cornell University)*
1030 – 1100	**Break**
1100 – 1200	**Keynote Speech – Consumers and Privacy: Who Cares?** *Bob Sullivan, MSNBC*
1200 – 1330	**Lunch [Owego/Brindley Room]**
1330 – 1400	**Probabilistic Opaque Byzantine Quorum Systems** *Michael G. Merideth (Carnegie Mellon University), Michael K. Reiter (University of North Carolina, Chapel Hill)*
1400 – 1430	**The Building Blocks of Consensus** *Yee Jiun Song (Cornell University), Robbert van Renesse (Cornell University), Danny Dolev (The Hebrew University of Jerusalem)*
1430 – 1600	**Live Distributed Objects: Enabling the Active Web** *Krzysztof Ostrowski (Cornell University), Ken Birman (Cornell University), Danny Dolev (The Hebrew University of Jerusalem)*
1500 – 1530	**Break**
1530 – 1800	**A Modeling Environment for Patient Portals** *Sean Duncanson (Vanderbilt University), James Noble (Vanderbilt University), Jon Werner (Vanderbilt University), Bradley A. Malin (Vanderbilt University), Alvin Ledeal (Vanderbilt University), James Schlemmer (Vanderbilt University)*
1800 – 1830	**Privacy and Utility in Business Processes** *Adam Barth (Stanford University), Anupam Datta (Carnegie Mellon University), John C. Mitchell (Stanford University), Sharada Sundaram (Tata Consultancy Services and Stanford University)*
1830 – 1700	**Don't Sweat Your Privacy: Using Humidity To Detect Human Presence** *Jun Han (Carnegie Mellon University), Abhishek Shah (Carnegie Mellon University), Mark Luk (Carnegie Mellon University), Adrian Perrig (Carnegie Mellon University)*
1700 – 1730	**A Model-Based Intrusion Detection System for Wireless Process Control Systems** *Tanya Roosta (UC Berkeley)*
1830	**Conference Attendee Dinners [Various Local Restaurants]**

NOTE: Unless otherwise noted, conference events will be held in the ***Seneca/Tioga Room*** of the Hilton Garden Inn Ithaca.

CONFERENCE AGENDA (cont.)

THURSDAY, OCTOBER 11, 2007	
TIME	*TOPIC*
0800 – 0830	Breakfast [Owego/Brindley Room]
0830 – 0855	ARSL: A Language for Authorization Rule Specification in Software Security *Weider D. Yu (San Jose State University), Ellora Nayak (San Jose State University)*
0855 – 0920	Towards Automatic Discovery of Deviations in Binary Implementations with Applications to Error Detection and Fingerprint Generation *David Brumley (Carnegie Mellon University), Juan Caballero (Carnegie Mellon University), Cody Hartwig (Carnegie Mellon University), Zhenkai Liang (Carnegie Mellon University), James Newsome (Carnegie Mellon University), Dawn Song (UC Berkeley)*
0920 – 0945	Using Social Network Theory Towards Development of Wireless Ad Hoc Network Trust *Sameer Pai (Cornell University), Tanya Roosta (UC Berkeley), Stephen Wicker (Cornell University), Shankar Sastry (UC Berkeley)*
0945 – 1010	Automatic Generation of XSS and SQL Injection Attacks with Goal-Directed Model Checking *Michael C. Martin (Stanford University), Monica S. Lam (Stanford University)*
1010 – 1030	Break
1030 – 1055	The TRUST Home Medical Sensing Testbed *Roozbeh Jafari (UC Berkeley), Sameer Iyengar (UC Berkeley), Ruzena Bajcsy (UC Berkeley), Philip Kuryloski (Cornell University), Sameer Pai (Cornell University), Stephen Wicker (Cornell University), Shanshan Jiang (Vanderbilt University), Yuanhuan Cao (Vanderbilt University), Yuan Xue (Vanderbilt University)*
1055 – 1120	Portcullis: Protecting Connection Setup from Denial-of-Capability Attacks *Bryan Parno (Carnegie Mellon University), Dan Wendlandt (Carnegie Mellon University), Elaine Shi (Carnegie Mellon University), Adrian Perrig (Carnegie Mellon University), Bruce Maggs (Carnegie Mellon University), Yih-Chun Hu (University of Illinois at Urbana-Champaign)*
1120 – 1145	Inductive Proofs of Computational Secrecy *Arnab Roy (Stanford University), Anupam Datta (Carnegie Mellon University), Ante Derek (Stanford University), John C. Mitchell (Stanford University)*
1145 – 1210	Characterizing the Remote Control Behavior of Bots *Elizabeth Stinson (Stanford University), John C. Mitchell (Stanford University)*
1210 – 1215	Wrap Up / Conference Closing
1215 – 1315	TRUST Student Poster Review *(Box Lunches Provided)*

NOTE: Unless otherwise noted, conference events will be held in the ***Seneca/Tioga Room*** of the Hilton Garden Inn Ithaca.

PRESENTATION ABSTRACTS

Wednesday, October 10

0900 – 0930 **Selective Sensing and the 4th Amendment**

Deirdre Mulligan (University of California, Berkeley)

Stephen Wicker (Cornell University)

The evolution of sensor networks able to detect various physical phenomena presents an opportunity to reduce the dependence of law enforcement and intelligence agencies on visual surveillance as a means of policing and intelligence gathering. The use of specific sensing technologies has the capacity to produce more salient data to support specific policing activities. It has the added benefit of reducing the effect of surveillance on certain forms of privacy by limiting the collection of data that reveals or eases the revelation of identity, limiting the ease with which data can be repurposed, and limiting the misuse and abuse of systems to track and monitor subsets of the population based on animus, bias, or voyeurism. This paper examines the existing rules that govern police/citizen interactions and uses them to construct a mock sensor network that would collect more relevant data for policing while limiting the unintended negative effects on privacy.

0930 – 1000 **Competition and Fraud in Online Advertising Markets**

Bob Mungamuru (Stanford University)

Stephen Weis (Google Inc.)

Advertising fraud, particularly click fraud, is a growing concern to the online advertising industry. Unlike many online security threats, ad fraud is primarily motivated by financial gain. Successfully committing ad fraud yields direct monetary gains for attackers at the expense of the victims. Thus, it is natural to consider online ad fraud in an economic context.

The online advertising market can be modeled as a dynamic game between three classes of players: publishers, advertisers, and ad networks. Publishers produce content that draws traffic and sell advertising space to advertisers. Ad networks act as intermediaries who connect publishers and advertisers. In practice, there are many publishers, many advertisers, and relatively few ad networks. Often a single organization may play two roles in the system, for instance, being both a publisher and advertiser or being both a publisher and ad network.

In our work, we first pose the ad fraud problem in detail. We then make a sequence of simplifications, leading to a tractable model that isolates the strategic interactions between competing ad networks. In particular, each ad network makes a single decision per period - how "aggressively" to filter for ad fraud. If an ad network filters more aggressively, advertisers realize a higher return on their investment, but the publishers and the ad network incur a higher false positive rate. Conversely, if an ad network filters less aggressively, the publishers and the ad network see a short term gain, but advertisers realize lower returns. The publishers and advertisers are assumed to observe the past history and play their best responses. We characterize the equilibria of this game, study convergence properties, and discuss the effect of reputation.

1000 – 1030 **Network Security and the Need to Consider Provider Coordination in Network Access Policy**

Aaron J. Burstein (University of California, Berkeley)

Fred B. Schneider (Cornell University)

The policy debate over how to govern access to broadband networks has largely ignored the objective of network trustworthiness—a set of properties (including security) that guarantee that a network will behave as expected. Instead, the terms of the network access debate have focused on whether imposing a nondiscrimination, or

network neutrality, obligation on service providers is justified by the condition of competition among last-mile providers. Some argue that, in the absence of a nondiscrimination obligation, service providers will discriminate against content, applications, and services that they (or their affiliates) do not provide. Others argue that this kind of discrimination is unlikely and that a nondiscrimination obligation would reduce incentives to invest in improving networks and developing new applications and services.

One point of agreement is that any nondiscrimination obligation must allow network providers to take measures to protect network security. This agreement, however, is rather abstract. Legislative, regulatory, and scholarly proposals have set forth substantially different security exceptions to nondiscrimination rules; but there has been little analysis of how these exceptions would affect the corresponding rule. Just as importantly, there has been little analysis of whether various exceptions allow sufficient room to defend against modern-day attacks. Moreover, the question of how network access policy affects other elements of trustworthiness, such as privacy, have gone unexamined. Put simply, network trustworthiness and network neutrality are closely related technologically and through network access policy. Decisions about technology or policy that are based on either trustworthiness or network neutrality principles in isolation pose the risk of affecting the other area in unexpected and undesirable ways.

This paper seeks to expand the network access policy debate to include both trustworthiness and neutrality. Our analysis leads to three principal conclusions. First, network providers need leeway to block or degrade traffic within their own subnets, as well as traffic exchanged between providers' subnets, in order to offer guarantees against certain kinds of attacks. Some currently proposed security exceptions to network neutrality requirements fail to allow such blocking. Second, some trustworthiness guarantees that are within technical reach, such as routing guarantees, would require service providers not to refuse to interconnect. The potential competitive effects of service provider coordination—which is critical in establishing these guarantees—warrant further study. Finally, individual providers are well situated to provide stronger privacy and confidentiality guarantees, without either coordinating with other providers or awaiting new technology. Drawing greater attention to the competitive dimensions of these elements of trustworthiness would likely help induce service providers to strengthen these guarantees.

1100 – 1200 **Consumers and Privacy: Who Cares?**
Bob Sullivan (MSNBC)

1330 – 1400 **Probabilistic Opaque Byzantine Quorum Systems**
Michael G. Merideth (Carnegie Mellon University)
Michael K. Reiter (University of North Carolina, Chapel Hill)
Byzantine-fault-tolerant service protocols like Q/U and FaB Paxos that optimistically order requests can provide increased efficiency and fault scalability. However, these protocols require n > 5b servers (where b is the maximum number of faults tolerated), owing to their use of opaque Byzantine quorum systems; this is 2b more servers than required by some non-optimistic protocols. In this paper, we present a family of probabilistic opaque Byzantine quorum systems that require substantially fewer servers. Our analysis is novel in that it assumes Byzantine clients, anticipating that a faulty client may seek quorums that maximize the probability of error. Using this as motivation, we present an optional, novel protocol that allows probabilistic quorum systems to tolerate Byzantine clients. The protocol requires only one additional round of interaction between the client and the servers, and this round may be amortized over multiple operations. We consider actual error probabilities introduced by the probabilistic approach for concrete configurations of opaque quorum systems, and

	prove that the probability of error vanishes with as few as n > 3.15b servers as n and b grow.
1400 – 1430	**The Building Blocks of Consensus** ***Yee Jiun Song (Cornell University)*** ***Robbert van Renesse (Cornell University)*** ***Danny Dolev (The Hebrew University of Jerusalem)*** Consensus is an important building block for building replicated systems, and many consensus protocols have been proposed. In this paper, we investigate the building blocks of consensus protocols and use them to design a skeleton algorithm that can be configured to produce, among others, three well-known consensus protocols: Paxos, Chandra-Toueg, and Ben-Or. Although each of these protocols specify only one quorum system explicitly, we show that all employ a second quorum system. We use the skeleton algorithm to implement a replicated service, allowing us to compare the performance of these consensus protocols under various workloads and failure scenarios.
1430 – 1500	**Live Distributed Objects: Enabling the Active Web** ***Krzysztof Ostrowski (Cornell University)*** ***Ken Birman (Cornell University)*** ***Danny Dolev (The Hebrew University of Jerusalem)*** Although we've been building distributed systems for decades, it remains remarkably difficult to get them right. Distributed software is hard to design and the tools available to developers have lagged far behind the options for building and debugging non-distributed programs targeting desktop environments. At Cornell, we're trying to change this dynamic. The first part of this talk will describe "Live Distributed Objects", a new and remarkably easy way to create distributed applications, with little or no programming required. Supporting these kinds of objects forced us to confront a number of scalability, security and performance questions not addressed by prior research on distributed computing platforms.
1530 – 1600	**A Modeling Environment for Patient Portals** ***Sean Duncavage (Vanderbilt University)*** ***Janos Mathe (Vanderbilt University)*** ***Jan Werner (Vanderbilt University)*** ***Bradley A. Malin (Vanderbilt University)*** ***Akos Ledeczi (Vanderbilt University)*** ***Janos Sztipanovits (Vanderbilt University)*** Many health-care organizations have migrated from paper-based to Electronic Medical Records (EMR), which have been shown to increase both staff productivity and patient safety. Expanding on the success of EMRs, Clinical Information Systems (CIS) incorporate a wide range of the informational and organizational components of the health-care environment. Local and federal regulations concerning the management of patient information present challenges for CIS design and implementation. The Health Insurance Portability and Accountability Act (HIPAA) specifically grants patients the right to access their medical records and requires healthcare organizations to provide security protection for protected health information. Patient Portals are one method to provide patients with a simple method to access their medical records, disclosures, and audits. Designing such a system optimally to protect patient confidentiality and respect health-care providers' rights is an open problem. We begin to address this challenge by casting patient portals, a key portion of CIS, onto a Service-Oriented Architecture (SOA). We developed a domain-specific modeling environment called Model-based Design Environment for Clinical Information Systems (MODECIS) with which we create formal models of healthcare services and

features for detailed analysis. Our initial research with MODECIS successfully demonstrates that patient portals can be modeled as SOA. The development of critical modeling abstractions adds the feature of scalability to our tool. Although MODECIS is a work-in-progress, it has been used to create high-fidelity models of the MyHealth@Vanderbilt patient portal.

SOA has been previously proposed for the design of formally-composed CIS environments; however, current implementations are limited by the fact they do not model patient-provider interactions. In this paper, we show how SOA can be applied to a specific patient-associated environment. We propose to use the web service standards defined by OASIS, which includes the Business Process Execution Language (BPEL) for web service orchestration and the Extensible Access Control Markup Language (XACML) for policy representation.

Workflows provide a representation of the manner by which data is accessed, handled, and shared. Without formal representations of daily business processes and their interrelationships within the healthcare environment, it is not clearly evident why a patient's medical record is accessed or how the interactions between patient and provider are managed. Both underspecified and ad hoc workflow design can lead to malformed policies with unanticipated consequences, and even seemingly routine business processes can lead to serious privacy compromises when taken in combination. Taking this into account, formal workflow models are a starting point for the development and analysis of policy-driven operations supporting privacy and security.

This inspired the creation of the building of our tool suite, MODECIS, where the formal basis of our approach allows for the extension, reuse, and evolution of clinical information system. MODECIS has three main components: a) a graphical design environment for capturing the business logic of CIS through workflows, b) an analysis tool, which allows for the analysis of information flows and the exploration of security and privacy properties of a CIS system modeled with the graphical design environment, and finally c) a model translator that maps the CIS-specific workflows to BPEL, WSDL and XACML. By translating the domain models onto these SOA standards, the underlying alternative implementations of SOA platforms for the standards become applicable. This radically simplifies the fast prototyping, integration and testing tasks.

By capturing the appropriate level of abstraction, it is possible to satisfy utility, security, and policy requirements for CIS. In MODECIS, workflows provide us with this abstraction layer, which is suitable for patient-centered clinical information representation and management. It will allow us to perform vulnerability, security and privacy analyses through model verification and simulation-based testing tools. Additionally, model-based design provides the tools for automated system generation directly from the models.

At the heart of our approach, the domain-specific modeling language captures the system from multiple aspects. The *workflow models* can be thought of as a graphical equivalent of a simplified BPEL representation. They capture the orchestration logic with graphs that describe control, which specify the sequence of service invocations and data flows that represent the movement of information within a CIS system. One aspect allows for the orchestration of control flows that are defined as a composition of service invocations – which can either be asynchronous or synchronous – and the typical control structures – such as switch, join, while, and catch – which allows for the definition of arbitrary workflow logic. Anotherr aspect of workflow modeling describes the flow of data elements: how these elements are exchanged, processed and stored between and within various processes. This way each workflow model can be thought

of as an available service with well-defined interfaces. Since the workflow models only describe how data elements are used, we have created the view for building *datatype models* making the language to be strongly typed.

Workflows in general allow system architects to follow the information traveling between entities and can represent diverse entities interacting with the system, such as physical databases or people. For this reason MODECIS incorporates two more types of models for the integration of workflows with the underlying architectures and physical entities. This means that a complex, explicitly represented social and technical architecture can be constructed that the services build on.

The creation of *organizational models* allows for the human coordination within CIS. These models are used to specify the architecture of the enterprise itself, such as the roles of different people. Organizational models reflect inter- and intradepartmental interactions, as well as people's roles within departments specifying tasks and groups to whom these tasks are assigned to. This enables the specification of policies to facilitate role-based access control, for example.

While organizational models relate human-based workflow (i.e. workflows that describe expected behavior of and tasks preformed by the human players in CIS), *deployment models* specify the organization of computer servers, their conjunctive networks and interface with workflows in a similar manner to organizational models. They are often referred to as the network architecture (ex: they depict hospital servers and workstations along with the services they provide).

The final abstraction captures policy statements that crosscut workflow, organizational and deployment models. They place restrictions on accessing certain services and information.

MODECIS includes a model translator capable of mapping domain-specific models to executable BPEL code. Despite its wide acceptance, BPEL provides no support for the detection of a) possible deadlocks or b) process paths that are not viable. MODECIS plans to capitalize on existing technologies, such as Petri Nets, the SPIN model checker, Process Algebras, and Abstract State Machines, for (BPEL) model verification to identify such erroneous workflows.

As a final system-integration step to guarantee correct flow of logic captured by the domain models, the tool suite interfaces with an execution engine, which manages the multiple instances of workflows after deployment. Specifically, the engine organizes and executes the services required by the CIS entities (e.g., a patient, primary care provider, and patient portal) and enforces policies.

The MODECIS tool suite provides a domain-specific, graphical design environment for precisely describing organizational, deployment, service, and data models in relation to patient portals. Through our collaboration with the Vanderbilt University Medical Center (VUMC), we were able to create a modeling language capable of representing a functional patient portal. The VUMC group was also able to confirm the expressiveness and correctness of our patient portal workflow models, which we have begun to deploy on the Oracle BPEL execution engine.

Although MODECIS is a work-in-progress, models created with the tool suite serve as formal system specifications that can be mapped onto various SOA execution platforms for simulation. Consistency and wellformedness checking is already supported by MODECIS; support for policy verification and vulnerability and security analysis of the models is our next step, which will be supported through the use of existing analysis tools.

	MODECIS provides a scalable tool to evaluate design decisions and system changes before deploying costly healthcare infrastructure. The creation of patient portal models and simulations is one step toward designing robust CIS that are able to take into account the diverse privacy and security concerns of stakeholders.
1600 – 1630	**Privacy and Utility in Business Processes** ***Adam Barth (Stanford University)*** ***Anupam Datta (Carnegie Mellon University)*** ***John C. Mitchell (Stanford University)*** ***Sharada Sundaram (Tata Consultancy Services and Stanford University)*** Privacy is an increasingly important business concern in health care, financial services, and other organizations. Hospitals, clinics, banks, credit card clearing houses, customer support centers, and academic institutions all maintain databases with sensitive information. These databases are used regularly by employees to carry out business critical tasks. Organizations that collect and use personal information face the growing challenge of conducting their business effectively while managing privacy risks and compliance requirements. The risks are very real, with the theft of 26 million veteran records in May 2006 demonstrating how easily sensitive information can fall into unauthorized hands. In the United States, privacy legislation, such as HIPAA for the health care sector and GLBA for financial institutions, has spurred many business, including 68% of the Direct Marketing Association member companies as of 2001, to appoint Chief Privacy Officers whose primary job is privacy issues and policies. One of the biggest problems that privacy-sensitive organizations face is designing their internal activities and information practices to simultaneously serve their customers effectively and manage risks from disclosure of sensitive information. This fundamental problem arises in hospitals and clinics, where personal health information must be used to provide effective health care, but must also be protected from indiscriminate sharing to respect the privacy of patients—a requirement made more precise by HIPAA. Financial institutions use sensitive financial information to decide whether to grant loans, for example, and suffer direct loss and brand erosion if sensitive information is lost. Retail enterprises use credit card details in resolving charge-back disputes (where the privacy concerns are exacerbated by the common practice of outsourcing this task). College admissions officers review confidential letters of recommendation and transcripts. In all of these situations, the organization must carefully design the way it processes and uses information to balance the competing goals of privacy and the usefulness, or utility, of the business process. Business process designs involve instructing individuals how and when to access and use information, coupled with access and use policies embedded in information processing systems. Because considering utility or privacy alone does not provide enough information to make meaningful management decisions, our goal is to develop a framework and model for designing, evaluating, and auditing business processes to achieve utility goals while minimizing privacy risks. We propose an abstract model of business processes, utility, and privacy, present some specific results, and illustrate our concepts usingMyHealth@Vanderbilt, a web-based patient portal built and used at the Vanderbilt Medical Center. Examining the MyHealth portal led to many insights captured in our general theory. Our approach builds on contextual integrity, a conceptual framework for understanding privacy expectations and their implications developed in the literature on law, public policy, and political philosophy. The primary tenant of contextual integrity is that people interact in society not simply as individuals in an undifferentiated social world, but as individuals in certain capacities or roles, in distinctive social contexts (e.g., health care or banking). For example, the individuals in MyHealth act as patients, doctors, nurses,

or secretaries, according to a specific workflow for scheduling appointments, viewing lab results, and asking and answering heath questions.

Each context is characterized by its business objectives, or utility goals, and its norms of transmission. For example, one utility goal for MyHealth is to respond to health questions from patients. The norms of transmission identify conditions under which personal information can be communicated from one party to another. These norms are represented by the privacy goals of a workflow. A privacy goal for MyHealth is to restrict health information to doctors and nurses, the health care providers. Using a model of actions that transmit personal information from a sender in one role to a receiver in a possibly different role, agents may accumulate and send different types of personal information they receive. These messages represent emails, web forms, database entries, workflow data structures, and arguments to actions. We assume that messages have associated tags (e.g., "health information") to indicate their contents, but consider business processes in which human agents may tag messages incorrectly. Since agents may act independently, with different motives, we express privacy and utility goals using a form of alternating time temporal logic, which we call the Logic of Privacy and Utility (LPU), interpreted over the concurrent game structure of agent actions. In this logical setting, privacy is a trace property expressible in LTL, while utility requires that agents have strategies to achieve certain useful outcomes, and is therefore expressed naturally using the stronger ATL* path quantifiers. We also formulate workflows in temporal logic, by associating a responsibility to each agent role. For example, in the patient portal workflow, doctors are responsible for answering health questions and secretaries are responsible for scheduling appointments. We consider both a general class of workflows presented abstractly by logical formulas and a more concrete subclass of practical workflows presented as a labeled graph or automata. Within this setting, we formulate and address design-time and run-time questions about whether a give nworkflow achieve its privacy and utility goals, without assuming that human agents always follow their assigned responsibilities.

1. Does a given workflow achieve privacy and utility if all agents act responsibly? We present algorithms for answering this question. Specifically, privacy properties may be evaluated using standard LTL model-checking over the traces generated by responsible executions of the concurrent game structure. Evaluating utility is more involved because of the ATL* path quantifiers and, in general, is undecidable because agents learn only of messages they send or receive. Because of this limitation, we present a sound decision procedure for a restricted, but useful, class of formulas.

2. Can irresponsible agents be detected and held ac- countable for violations? If the execution of a workflow satisfying our design criteria actually violates privacy, then some agent must have caused the violation by acting irresponsibly. These violations can be caught at run time, and the accountable agent determined using auditing algorithms we present. These algorithms are not fully automatic (or else they could be used for enforcement), but require an oracle (such as a human auditor) to determine the accuracy of message tags. We seek to minimize the number of oracle calls (reducing the human auditor's work) by using classical causality ideas in distributed computing and a new notion of "suspicious events."

Privacy advocates often recommend reconciling the competing interests of privacy and utility with the principle of minimum necessary disclosure: disclose the minimum information necessary to achieve the utility goal. This principle is included expressly in several influential privacy policies, including the HIPAA Privacy Rule and the Markle Connecting for Health Common Framework. We leverage our unified model of privacy and utility to provide a formal definition of this principle. We apply these concepts to the MyHealth patient portal and recommend several design changes to the MyHealth developers at Vanderbilt. Message tags are themselves one such suggestion, enabling

	finer grained message routing. Our auditing algorithms were developed in response to the MyHealth developers' concern about incorrectly tagged messages. In this paper, we suggest further privacy improvements in the MyHealth workflow based on tagging and illustrate our auditing methods using a hypothetical execution of MyHealth with an irresponsible agent.
1630 – 1700	**Don't Sweat Your Privacy: Using Humidity To Detect Human Presence** ***Jun Han (Carnegie Mellon University)*** ***Abhishek Shah (Carnegie Mellon University)*** ***Mark Luk (Carnegie Mellon University)*** ***Adrian Perrig (Carnegie Mellon University)*** Sensor nodes are increasingly deployed in many environments. Most of these nodes feature onboard sensor chips to measure environmental data such as humidity, temperature and light. In this paper, we show that seemingly innocuous and non-sensitive data such as humidity measurements can disclose private information such as human presence. We conduct several experiments using Telos motes running TinyOS to justify our claims and discuss research to investigate mechanisms to prevent the leakage of private information.
1700 – 1730	**A Model-Based Intrustion Detection System for Wireless Process Control Systems** ***Tanya Roosta (University of California, Berkeley)*** A recent trend in the process control system (PCS) is to deploy sensor networks in the hard-to-reachareas. Using wireless sensors greatly decreases the wiring costs and increases the volume of data gathered for plant monitoring. However, ensuring the security of the deployed sensor network, which is part of the overall security of PCS, is of crucial importance. In this paper, we design a model-based intrusion detection system (IDS) for sensor networks used for PCS. Given PCS tends to have regular traffic patterns and a well-defined request-response communication, we can design an IDS that defines the model of normal behavior of the entities and detects attacks when there is a deviation from this model. Model-based IDS can prove useful in detecting unknown attacks.

Thursday, October 11

0830 – 0855	**ARSL: A Language for Authorization Rule Specification in Software Security** ***Weider D. Yu (San Jose State University)*** ***Ellora Nayak (San Jose State University)*** Web services constitute an important part of distributed applications, providing flexibility in the development of distributed applications. One of the key challenges in Web Service security is to determine whether an authenticated user has access to only those services for which he has authorization. Since all authorization patterns for accessing resources cannot be anticipated and hence the access rules cannot be defined beforehand, implementing authorization becomes a concern. This paper describes an approach aimed at a more generalized and reusable solution which provides the flexibility to handle authorization rule updates in real time. The authorization framework is complemented by ARSL (Authorization Rule Specification Language), which is based on predicate logic.
0855 – 0920	**Towards Automatic Discovery of Deviations in Binary Implementations with Applications to Error Detection and Fingerprint Generation** ***David Brumley (Carnegie Mellon University)*** ***Juan Caballero (Carnegie Mellon University)*** ***Cody Hartwig (Carnegie Mellon University)*** ***Zhenkai Liang (Carnegie Mellon University)***

	James Newsome (Carnegie Mellon University) ***Dawn Song (University of California, Berkeley)*** Different implementations of the same protocol specification usually contain deviations, i.e., differences in how they check and process some of their inputs. Deviations are commonly introduced as implementation errors or as different interpretations of the same specification. Automatic discovery of these deviations is important for several applications. In this paper, we focus on automatic discovery of deviations for two particular applications: error detection and fingerprint generation. We propose a novel approach for automatically detecting deviations in the way different implementations of the same specification check and process their input. Our approach has several advantages: (1) by automatically building symbolic formulas from the implementation, our approach is precisely faithful to the implementation; (2) by solving formulas created from two different implementations of the same specification, our approach significantly reduces the number of inputs needed to find deviations; (3) our approach works on binaries directly, without access to the source code. We have built a prototype implementation of our approach and have evaluated it using multiple implementations of two different protocols: HTTP and NTP. Our results show that our approach successfully finds deviations between different implementations, including errors in input checking, and differences in the interpretation of the specification, which can be used as fingerprints.
0920 – 0945	**Using Social Network Theory Towards Development of Wireless Ad hoc Network Trust** ***Sameer Pai (Cornell University)*** ***Tanya Roosta (University of California, Berkeley)*** ***Stephen Wicker (Cornell University)*** ***Shankar Sastry (University of California, Berkeley)*** The evolution and existence of stable trust relations have been studied extensively in the context of social theory. However, reputation systems or trust schemes have only been recently used in the domain of wireless ad hoc networks. It has been shown that these schemes provide positive results as self-policing mechanisms for the routing of data in wireless ad hoc network security. Much work has gone into the development of ad hoc wireless networks of sensors. One serious fear that still exists, however, is that of exposure of the network to misbehaving nodes during routing. Misbehaving nodes are defined to be those nodes that are incorrectly functioning or malicious. The exposure of the network to misbehaving nodes would jeopardize its function. Consequently, the data gathered by the network would be insufficient or even incorrect, which would result in the failure of the application for which the network is deployed. More specifically, this paper examines trust in the context of routing and reliable forwarding of data in wireless ad hoc networks. The lack of a reliable infrastructure or a central authority in wireless ad hoc networks means that nodes must cooperate to route data from point to point. When a source node transmits data, if the intermediary nodes fail to cooperate and route the data, due to any subset of them misbehaving, energy and other resources in the network are wasted. Moreover, if these misbehaving nodes fail to transmit the correct information or re-route the data to the wrong nodes, data integrity and/or confidentiality could be compromised. Therefore, nodes need a way to distinguish behaving nodes from those that misbehave. Naturally, nodes can make this assessment by associating every other node and potential router with a trust value or belief in the other nodes ability to successfully route data. Multiple distributed schemes to compute trust values and to help eliminate misbehaving nodes have been suggested. These schemes, however, are slow to assess misbehavior in the network and at the same time they are prone to eliminating

behaving nodes in the presence of benign interaction failures (e.g., those failures arising from an error-prone wireless channel for transmitting data). Our work reviews these existing schemes while mapping them to similar models for trust assessment in the social network theory. Most importantly, a refined model of trust evaluation in social networks is constructed using insights from Network Balance Theory and mapped to a new trust scheme for wireless ad hoc networks. The new trust scheme introduces two fundamental changes to existing schemes:

1. nodes calculate direct trust values based on direct observations and mutual transitive observations in their neighborhood; and
2. trust value thresholds, below which nodes determine not to route data to untrustworthy nodes, are heterogeneous and set autonomously be each node

The new trust scheme is analyzed and shown to outperform existing schemes using scenario analysis and simulation. The new trust scheme is shown (i) to be faster at detecting misbehaving nodes and (ii) better at preserving the network connectivity in the presence of the benign failures.

0945 – 1010 **Automatic Generation of XSS and SQL Injection Attacks with Goal-Directed Model Checking**

Michael C. Martin (Stanford University)
Monica S. Lam (Stanford University)

Cross-site scripting (XSS) and SQL injection errors are two prominent examples of taint-based vulnerabilities that are responsible for a large number of security breaches in recent years. This paper presents QED, a goal-directed model-checking system for locating software defects including these in large Java web applications.

The user provides a specification of undesirable behavior and the application to test. Based on the specification and the structure of the web application, QED harnesses the application for model checking and also performs a static analysis to remove unnecessary portions of the search space. The remaining state space can then be systematically checked. This checking is complete enough to also discover defects that only manifest themselves after multiple requests have been made to a web application.

Not only does QED identify these vulnerabilities automatically, it proves the existence of errors by providing an example input of an attack and a program trace showing how the code is compromised. QED does make it easy for the application maintainer to recognize the errors and to make the necessary fix. In other words, unlike many other static analysis tools, QED does not generate any false-positive warnings. For a class of applications, QED can guarantee that it has found all the potential bugs in the program.

QED combines techniques from advanced pointer analysis, dynamic monitoring, and model checking, combining the advantages of all three: QED provides a bound on errors in both directions, and can readily demonstrate how to trigger the bugs it finds.

We have run QED over a set of 3 Java web applications that use the popular Struts framework. In 130,000 lines of code, we found 10 SQL injections and 13 cross-site scripting errors.

1030 – 1055 **The TRUST Home Medical Sensing Testbed**

Roozbeh Jafari (University of California, Berkeley)
Sameer Iyengar (University of California, Berkeley)
Ruzena Bajcsy (University of California, Berkeley)
Philip Kuryloski (Cornell University)
Sameer Pai (Cornell University)

	Stephen Wicker (Cornell University) ***Shanshan Jiang (Vanderbilt University)*** ***Yanchuan Cao (Vanderbilt University)*** ***Yuan Xue (Vanderbilt University)*** Remote patient monitoring has great potential to improve many aspects of patient care. In one area alone, the monitoring for adverse events such as falls, the quality of care for elderly patients can be significantly improved, while allowing for greater independence and improved quality of life. Early detection of health conditions can also reduce the severity of illness and reduce the need for and potential duration of hospitalization. TRUST researchers are developing a first stage system that focuses on patient motion. The system includes sensors, wearable wireless transmitters, local data relays, and long-range wireless networking. A two-level networking scheme has been developed to provide connectivity throughout the patient's residence, data processing to recognize significant events, and a reporting scheme that routs data to the local care facility. The research program, a collaborative effort between UCB, Cornell, UT Dallas, and Vanderbilt, has three critical components: the sensor platform, data processing and networking, and security and privacy. The wearable motion sensor consists of a moteiv tmote sky equipped with a 3-axis accelerometer and 3-axis gyroscope. Algorithms for recognizing specific problematic motions have been developed and tied to event detection and classification algorithms. Ongoing network research focuses on the reliability, robustness, and throughput of a two-tiered wireless transport network composed of the tmotes, Crossbow Stargate boards, and portable personal computers. Privacy & security research focuses on the methods for event notification and data display. This includes the use of virtual rather than standard video or photographic representations of patients, as well as other data filtering techniques. Experiments using a system prototype have been conducted at Vanderbilt Home Care Services. We will report on our system design, deployment and experiment experiences and challenges.
1055 – 1120	**Portcullis: Protecting Connection Setup from Denial-of-Capability Attacks** ***Bryan Parno (Carnegie Mellon University)*** ***Dan Wendlandt (Carnegie Mellon University)*** ***Elaine Shi (Carnegie Mellon University)*** ***Adrian Perrig (Carnegie Mellon University)*** ***Bruce Maggs (Carnegie Mellon University)*** ***Yih-Chun Hu (University of Illinois at Urbana-Champaign)*** Systems using capabilities to provide preferential service to selected flows have been proposed as a defense against large-scale network denial-of-service attacks. While these systems offer strong protection for established network flows, the Denial-of-Capability (DoC) attack, which prevents new capability-setup packets from reaching the destination, limits the value of these systems. Portcullis mitigates DoC attacks by allocating scarce link bandwidth for connection establishment packets based on per-computation fairness. We prove that a legitimate sender can establish a capability with high probability regardless of an attacker's resources or strategy and that no system can improve on our guarantee. We simulate full and partial deployments of Portcullis on an Internet-scale topology to confirm our theoretical results and demonstrate the substantial benefits of using per-computation fairness.
1120 – 1145	**Inductive Proofs of Computational Secrecy** ***Arnab Roy (Stanford University)*** ***Anupam Datta (Carnegie Mellon University)*** ***Ante Derek (Stanford University)*** ***John C. Mitchell (Stanford University)***

	Secrecy properties of network protocols assert that no probabilistic polynomial-time distinguisher can win a suitable game presented by a challenger. Because such properties are not determined by trace-by-trace behavior of the protocol, we establish a trace-based protocol condition, suitable for inductive proofs, that guarantees a generic reduction from protocol attacks to attacks on underlying primitives. We use this condition to present a compositional inductive proof system for secrecy, and illustrate the system by giving a modular, formal proof of computational authentication and secrecy properties of Kerberos V5.
1145 – 1210	**Characterizing the Remote Control Behavior of Bots** ***Elizabeth Stinson (Stanford University)*** ***John C. Mitchell (Stanford University)*** A botnet is a collection of bots, each generally running on a compromised system and responding to commands over a "command-and-control" overlay network. We investigate observable differences in the behavior of bots and benign programs, focusing on the way that bots respond to data received over the network. Our experimental platform monitors execution of an arbitrary Win32 binary, considering data received over the network to be tainted, applying library-call-level taint propagation, and checking for tainted arguments to selected system calls. As a way of further distinguishing locally-initiated from remotely-initiated actions, we capture and propagate "cleanliness" of local user input (as received via the keyboard or mouse). Testing indicates behavioral separation of major bot families (agobot, DSNXbot, evilbot, G-SySbot, sdbot, Spybot) from benign programs with low error rate. Botnets have been instrumental in distributed denial of service attacks, click fraud, phishing, malware distribution, manipulation of online polls and games, and identity theft. As much as 70% of all spam may be transmitted through botnets and as many as 25% of all computers may be participants in a botnet. A bot master (or "botherder") directs the activities of a botnet by issuing commands that are transmitted over a command-and-control (C&C) overlay network. Some previous network-based botnet detection efforts have attempted to exploit this ongoing C&C behavior or its side effects. Our work investigates the potential for host-based behavioral bot detection. In particular, we test the hypothesis that the behavior of installed bots can be characterized in a way that distinguishes malicious bots from innocuous processes. We are not aware of any prior studies of this topic. Each participating bot independently executes each command received over the C&C network. A bot command takes some number of parameters (possibly zero) – each of a particular type – in some fixed order. For example, many bots provide a web-download command, which commonly takes two parameters; the first is a URL that identifies a remote resource (typically a file) that should be downloaded, and the second is the file path on the host system at which to store the downloaded data. A botnet constitutes a remotely programmable platform with the set of commands it supports forming its API. Many parameterized bot commands are implemented by invoking operating system services on the host system. For example, the web-download command connects to a target over the network, requests some data from that target, and creates a file on the host system; all of these actions (connect, network send and receive, and file creation) are performed via execution of system calls. Typically, a command's parameters provide information used in the system call invocation. For example, the connect system call takes an IP address argument, which identifies the target host with which a connection should be established. Implementations of the web-download command obtain that target host IP from the given URL parameter. Thus, execution of many parameterized commands causes system call invocations on arguments obtained from those parameters.

In this paper, we test the experimental hypothesis that the remote control of bots through parameterized commands separates bot behavior from normal execution of innocuous programs. We postulate that a process exhibits external or remote control when it uses data received from the network (an untrusted source) in a system call argument (a trusted sink). We test our hypothesis via a prototype implementation, BotSwat, designed for the environment in which the vast majority of bots operate: home users' PCs running Windows XP or 2000. BotSwat can monitor execution of an arbitrary Win32 binary and interposes on the run-time library calls (including system calls) made by a process. We consider data received over the network to be tainted and track tainted data as it propagates via dynamic library calls to other memory regions. We identify execution of parameterized bot commands when tainted arguments are supplied to select gate functions, which are system calls used in malicious bot activity.

Our experimental results suggest that the presence of network packet contents in selected system call arguments is an effective indicator for malicious Win32 bots, including tested variants of agobot, DSNXbot, evilbot, G-SySbot, sdbot, and Spybot. Bots from these families constitute 98.2% of malicious bots seen in the wild. While these bots may implement commands in significantly different ways, similarities in the way they respond to external control allow a single approach to identify them. Additionally, the thousands of variants of each such family generally differ in ways that will not affect our ability to detect them; this is in contrast to traditional anti-malware signature scanners which may require a distinct signature for each variant. Moreover, our generic approach does not rely on a particular command-and-control communication protocol (e.g., IRC) or botnet structure (e.g., centralized or peer-to-peer).

KEYNOTE SPEAKER BIOGRAPHY

Bob Sullivan (MSNBC)

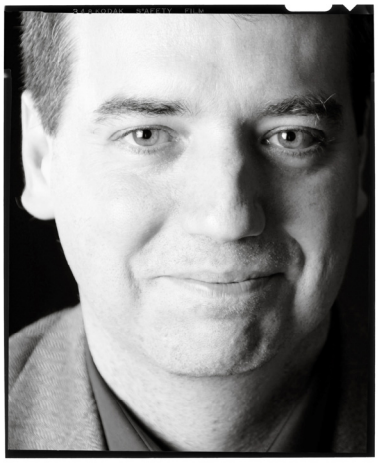

Bob Sullivan is an author and technology writer for MSNBC who has concentrated on technology crime and consumer fraud. He is the nation's leading journalist covering identity fraud having written more than 100 articles on the subject since 1996. His work appears on several MSNBC partner sites including MSN.com, Wall Street Journal Interactive, and ZDNet.com. His first book, *Your Evil Twin: Behind the Identity Theft Epidemic*, introduced the nation to the complex crime of identity theft. His second book, *Gotcha Capitalism*, will expose the work of companies that trick consumers into pay hidden fees and surcharges for nearly every good and service they buy. It furthers the work of his popular blog, The Red Tape Chronicles. It will be available from Random House beginning this January.

Among his many scoops, Bob was the first to tell the world about the existence of Magic Lantern, and top-secret Trojan horse program designed by FBI researchers to steal encryption passphrases; he was also the first to describe the data theft at ChoicePoint, the first of what would become an avalanche of stories about stolen and lost personal information. Sullivan also appears on air on MSNBC, CNBC, NBC Nightly News, the Today show, and various local NBC affiliates. He is the winner of the prestigious 2002 Society of Professional Journalists Public Service Award for his series of articles on online fraud. He has spoken before trade and government groups including the National Association of Attorney Generals. He lives in Maltby, Washington with his golden retriever, Lucky.

MSNBC

THE RED TAPE CHRONICLES

http://redtape.msnbc.com/

SPEAKER BIOGRAPHIES

Ruzena Bajcsy

University of California, Berkeley

Ruzena Bajcsy is a pioneering researcher in machine perception, robotics and artificial intelligence. She is a professor in the Electrical Engineering and Computer Science Department at Berkeley. She was the founding director of the University of Pennsylvania's General Robotics and Active Sensory Perception (GRASP) Laboratory, which she founded in 1978. Bajcsy has done seminal research in the areas of human-centered computer control, cognitive science, robotics, computerized radiological/medical image processing and artificial vision. She is highly regarded, not only for her significant research contributions, but also for her leadership in the creation of a world-class robotics laboratory, recognized world wide as a premiere research center. She is a member of the National Academy of Engineering, as well as the Institute of Medicine. She is especially known for her wide-ranging, broad outlook in the field and her cross-disciplinary talent and leadership in successfully bridging such diverse areas as robotics and artificial intelligence, engineering and cognitive science.

Ken Birman

Cornell University

Ken Birman is Professor of Computer Science at Cornell University. He currently heads the QuickSilver project, which is developing a scalable and robust distributed computing platform. Previously he worked on fault-tolerance, security, and reliable multicast. In 1987 he founded a company, Isis Distributed Systems, which developed robust software solutions for stock exchanges, air traffic control, and factory automation. The author of several books and more than 200 journal and conference papers, Dr. Birman was Editor in Chief of ACM Transactions on Computer Systems from 1993-1998 and is a Fellow of the ACM.

Aaron J. Burstein

University of California, Berkeley

Aaron Burstein, J.D (Boalt Hall), is the TRUST and ACCURATE Research Fellow for the Berkeley Center for Law & Technology (BCLT) and the Samuelson Clinic at the University of California, Berkeley. His current research interests include legal and economic approaches to improving computer and network security and electronic voting. Prior to returning to Boalt, Burstein was an attorney in the U.S. Department of Justice Antitrust Division.

Anupam Datta

Carnegie Mellon University

Anupam Datta is a Research Scientist at Carnegie Mellon University. He obtained MS and Ph.D. degrees from Stanford University and a BTech from IIT Kharagpur, all in Computer Science. Dr. Datta's research interests are in security, cryptography and privacy. He has authored over 20 papers in topics including security analysis of network protocols, theory of cryptography, languages for privacy policy specification and enforcement, and software system security. Dr. Datta is the General Chair for the 2008 IEEE Computer Security Foundations Symposium and has served on program committees for a number of security conferences including the 2007 IEEE Symposium on Security and Privacy.

Cody Hartwig

Carnegie Mellon University

Cody Hartwig is a Ph.D. student in Carnegie Mellon's Computer Science Department. He received his BSEs in computer science and computer engineering from the University of Michigan in 2006. Cody's research interests include computer security with specific attention to binary analysis.

Phillip Kuryloski

Cornell University

Philip Kuryloski is a graduate student in the School of Electrical & Computer Engineering at Cornell University in Ithaca, NY. Philip works in Dr. Stephen Wicker's Wireless Intelligent Systems Laboratory, where his research focus is wireless sensor networks for medical applications. Philip is currently an exchange scholar at the University of California, Berkeley. There he is working with Dr. Ruzena Bajcsy as well as other researchers at University of Texas, Dallas and Vanderbilt University on a project using wearable wireless sensors for motion and activity analysis and classification.

Akos Ledeczi

Vanderbilt University

Ákos Lédeczi is a Research Assistant Professor at the Electrical Engineering and Computer Science Department at Vanderbilt University and a Senior Research Scientist at the Institute for Software Integrated Systems (ISIS). Dr. Lédeczi received his M.Sc. degree in Electrical Engineering from the Technical University of Budapest in 1989. He received his Ph.D. degree in Electrical Engineering from Vanderbilt University in 1995.

Mark Luk

Carnegie Mellon University

Mark Luk is a Ph.D. candidate in Electrical and Computer Engineering at Carnegie Mellon University. He completed his Masters degree in Information Security at Carnegie Mellon University under Prof. Adrian Perrig. Prior to that, he received his Bachelors degree in Computer Science from the University of California at Berkeley. His research interests revolve around sensor network security and systems security.

Michael C. Martin

Stanford University

Michael Martin is a Ph.D. student from Stanford University working under Professor Monica Lam. His primary research interest is in custom runtimes and trace-based pattern matching to efficiently track program executions for safety or program understanding purposes. He is the primary designer and maintainer of the PQL pattern language.

Deirdre K. Mulligan

University of California, Berkeley

Deirdre K. Mulligan is the director of the Samuelson Law, Technology & Public Policy Clinic and a clinical professor of law at the UC Berkeley School of Law (Boalt Hall). Before coming to Boalt, she was staff counsel at the Center for Democracy & Technology in Washington.

Through the clinic, Mulligan and her students foster the public's interest in new computer and communication technology by engaging in client advocacy and interdisciplinary research, and by participating in developing technical standards and protocols. The clinic's work has advanced and protected the public's interest in free expression, individual privacy, balanced intellectual property rules, and secure, reliable, open communication networks.

Mulligan writes about the risks and opportunities technology presents to privacy, free expression, and access and use of information goods. Recent publications about privacy include: "Storing Our Lives Online: Expanded Email Storage Raises Complex Policy Issues," with Ari Schwartz and Indrani Mondal (2005), I/S: A Journal of Law and Policy for the Information Society; and, "Reasonable Expectations in Electronic Communications: A Critical Perspective on the Electronic Communications Privacy Act," 72 Geo. Wash. L. Rev. 1557 (2004).

Mulligan was a member of the National Academy of Sciences Committee on Authentication Technology and Its Privacy Implications; the Federal Trade Commission's Federal Advisory Committee on Online Access and Security, and the National Task Force on Privacy, Technology, and Criminal Justice Information. She was a vice-chair of the California Bipartisan Commission on Internet Political Practices and chaired the Computers, Freedom, and Privacy (CFP) Conference in 2004. She is currently a member of the California Office of Privacy Protection's Advisory Council and a co-chair of Microsoft's Trustworthy Computing Academic Advisory Board. She serves on the board of the California Voter Foundation and on the advisory board of the Electronic Frontier Foundation.

Bob Mungamuru

Stanford University

Bob Mungamuru is a Ph.D. student in Electrical Engineering and Computer Science at Stanford University. His advisor is Hector Garcia-Molina in the Stanford InfoLab.

Sameer Pai

Cornell University

Sameer Pai received the B.S. degree (with highest honors) in Electrical and Computer Engineering at Rutgers University in 2003. From 2002 to 2003 he was employed as an Engineer in the Personal Communications Sector at Motorola Inc. Since Fall 2004, he has been a Ph.D. candidate at Cornell University working with Professor Stephen B. Wicker in the Wireless Intelligent Systems Laboratory. From Summer of 2006 until Summer 2007, Sameer was an exchange scholar at UC Berkeley working with Professor Shankar Sastry and Professor Deirdre Mulligan. Sameer's broad research interests include wireless ad hoc networks, game theory, social networks and public policy on privacy. His current focus is on the mitigation of misbehavior in wireless sensor networks using trust and reputations as well as the impacts of security and public policy decisions on privacy issues in these networks. Sameer has been awarded a Rutgers Undergraduate Research Fellowship, a Cornell Graduate Fellowship, a NSF Graduate Fellowship Honorable Mention, a NSF IGERT Fellowship, and the UC Berkeley Phi Kappa Sigma Erdelatz Prize. Sameer is a student member of Tau Beta Pi, Eta Kappa Nu and the IEEE.

Adrian Perrig

Carnegie Mellon University

Adrian Perrig is an Associate Professor at Carnegie Mellon University with appointments in the departments of Electrical and Computer Engineering; Engineering and Public Policy; and Computer Science. He is also a member of CyLab, Carnegie Mellon's Cybersecurity Laboratory.

Adrian's research focuses on networking and systems security, security for mobile computing and sensor networks. Other research interests are in human interfaces for security, networking, operating systems, and applied cryptography.

Michael K. Reiter

University of North Carolina at Chapel Hill

Michael Reiter is the Lawrence M. Slifkin Distinguished Professor in the Department of Computer Science at the University of North Carolina at Chapel Hill (UNC). He received the B.S. degree in mathematical sciences from UNC in 1989, and the M.S. and Ph.D. degrees in computer science from Cornell University in 1991 and 1993, respectively. He joined AT&T Bell Labs in 1993 and became a founding member of AT&T Labs – Research when NCR and Lucent Technologies (including Bell Labs) were split away from AT&T in 1996. He then returned to Bell Labs in 1998 as Director of Secure Systems Research. In 2001, he joined Carnegie Mellon University as a Professor of Electrical & Computer Engineering and Computer Science, where he was also the founding Technical Director of CyLab. He joined the faculty at UNC in 2007.

Dr. Reiter's research interests include all areas of computer and communications security and distributed computing. He regularly publishes and serves on conference organizing committees in these fields, and has served as program chair for the flagship computer security conferences of the IEEE, the ACM, and the Internet Society. He currently serves as Editor-in-Chief of ACM Transactions on Information and System Security and on the Board of Visitors for the Software Engineering Institute. He previously served on the editorial boards of IEEE Transactions on Software Engineering, IEEE Transactions on Dependable and Secure Computing, and the International Journal of Information Security, and as Chair of the IEEE Technical Committee on Security and Privacy.

Tanya Roosta

University of California, Berkeley

Tanya Roosta is in the last year of her Ph.D. in Electrical Engineering and Computer Science at the University of California, Berkeley, after having received her B.S. in EECS there in 2000 and her M.S. there in 2004.

Tanya's research interests include sensor network security at the application layer, fault detection, data integrity, reputation systems, sensor correlation modeling, power saving methods, and privacy issues associated with the application of sensors at home. Tanya is also interested in ad-hoc wireless networks, specifically the design of low power protocols at the network and MAC layer as well as robust statistical methods, outlier detection models, statistical modeling and model validation, and the application of game theory to sensor network design.

Arnab Roy

Stanford University

Arnab Roy is a Ph.D. candidate at the Stanford University Department of Computer Science and a Siebel Scholar, Class of 2006. Prior to joining Stanford, he completed undergraduate studies at Indian Institute of Technology Kharagpur obtaining the Prime Minister of India Gold Medal for graduating at the top of his class.

Yee Jiun Song

Cornell University

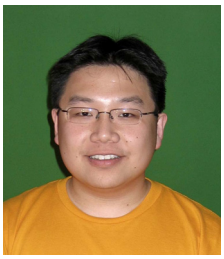

Yee Jiun Song is a Ph.D. student in the department of computer science at Cornell University. He is interested in the scalability and reliability of distributed systems. He is currently working on the design and understanding of consensus protocols and their performance in real world applications. Yee Jiun is also interested in peer-to-peer systems and self-organizing systems. Yee Jiun holds a B.S. and an M.S. degree in computer science from UC Berkeley and Stanford University, respectively.

Elizabeth Stinson

Stanford University

Liz Stinson is a Ph.D. student in the Computer Science Department of Stanford University, working with John Mitchell. She is interested in systems, networks, and security.

Stephen Wicker

Cornell University

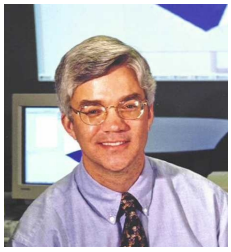

Stephen B. Wicker is a Professor of Electrical and Computer Engineering at Cornell University, and a member of the graduate fields of Computer Science and Applied Mathematics. Professor Wicker was awarded the 1988 Cornell College of Engineering Michael Tien Teaching Award and the 2000 Cornell School of Electrical and Computer Engineering Teaching Award. As of early 2007, he has supervised thirty doctoral dissertations.

Professor Wicker is the author of *Codes, Graphs, and Iterative Decoding* (Kluwer, 2002), *Turbo Coding* (Kluwer, 1999), *Error Control Systems for Digital Communication and Storage* (Prentice Hall, 1995) and *Reed-Solomon Codes and Their Applications* (IEEE Press, 1994). He has served as Associate Editor for Coding Theory and Techniques for the *IEEE Transactions on Communications*, and is currently Associate Editor for the *ACM Transactions on Sensor Networks*. He has served two terms as a member of the Board of Governors of the IEEE Information Theory Society, and chaired the Technical Program Committee for the Fifth International Conference on Information Processing in Sensor Networks (IPSN 2006).

Professor Wicker teaches and conducts research in wireless information networks, digital systems, self-configuring systems, and artificial intelligence. His current research focuses on the use of probabilistic models and game theory in the development of highly distributed, adaptive sensor networks. He is also conducting joint research with the Berkeley School of Law on privacy policy and the impact of the deployment of sensor networks in public spaces. Professor Wicker is the Cornell Principal Investigator for the TRUST Science and Technology Center – a National Science Foundation center dedicated to the development of technologies for securing the nation's critical infrastructure.

Professor Wicker heads the Wireless Intelligent Systems Laboratory, whose focus is on the field of wireless networks, including traditional cellular networks, ad-hoc networks, and sensor networks.

Weider D. Yu

San Jose State University

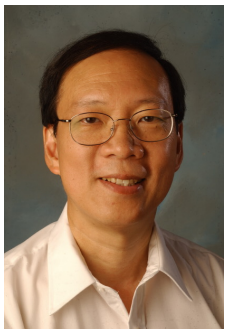

Dr. Weider D. Yu is an associate professor in the Computer Engineering Department at San Jose State University, San Jose (Silicon Valley), California. He received an M.S. in Computer Science from the State University of New York at Albany, and a Ph.D. from Northwestern University in Electrical Engineering 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.)