

TRUST

Team for Research in Ubiquitous Secure

Technology

An NSF Science and Technology Center

2005-2006 Annual Report

1. GENERAL INFORMATION	5
1.1. Contact Information.....	5
1.2. Biographical Information for New Faculty Members (by institution)	6
1.3. Executive Summary	10
2. RESEARCH	13
2.1. Research Objectives.....	13
2.2. Current and Anticipated Problems	14
2.3. Research Thrust Areas	14
2.3.1. Thrust Area I: Model-Based Integration of Trusted Components	14
2.3.2. Thrust Area II: Secure Information Management Software Tools.....	18
2.3.3. Thrust Area III: Integrative Projects	22
2.3.4. Thrust Area IV: Trusted Platforms and Trustworthy Systems.....	26
2.3.5. Thrust Area V: Secure Sensor Networks.....	27
2.3.6. Thrust Area VI: Network Defense.....	33
3. EDUCATION	42
3.1. Educational Objectives	42
3.2. Current and Anticipated Problems	43
3.3. Internal Educational Activities	43
3.4. External Educational Activities	48
3.5. Student Participation in Professional Development Activities	48
3.6. Integration of Research and Education	48
3.7. Future Plans for Internal and External Educational Activities.....	48
4. KNOWLEDGE TRANSFER	51
4.1. KT Objectives.....	51
4.2. KT Current and Anticipated Problems.....	52
4.3. KT Activities	53

4.4.	Other KT Activities	57
4.5.	KT Future Plans	57
5.	PARTNERSHIPS	59
5.1.	Partnership Objectives	59
5.2.	Current and Anticipated Problems for Partnerships	59
5.3.	Partnership Activities	59
5.4.	Other Partnership Activities.....	60
5.5.	Future Plans for Partnerships	61
6.	DIVERSITY	62
6.1.	Diversity Objectives	62
6.2.	Demographics.....	62
6.3.	Current and Anticipated Problems for the Center's Diversity Goals.....	66
6.4.	Contributions.....	66
6.5.	Future Plans for Enhancing Diversity.....	67
7.	MANAGEMENT	69
7.1.	Organizational Strategy	69
7.2.	Current and Anticipated Management Problems	70
7.3.	Management and Communication Systems	70
7.4.	External Advisory Board	70
7.5.	Changes to the Strategic Plan	71
8.	CENTER-WIDE OUTPUTS AND ISSUES	72
8.1.	Publications.....	72
8.2.	Presentations	77
8.3.	Other Dissemination Activities	82
8.4.	Awards & Honors	84
8.5.	TRUST's Class of 2006	85
8.6.	Outputs of Knowledge Transfer Activities	85
8.7.	TRUST's Participants.....	86
8.8.	TRUST's Institutions & Partners	91
8.9.	TRUST in the News	91
8.10.	Center Employment and Turnover	92

9. BUDGET

Error!

Bookmark not defined.

- 9.1. Current Year Budget Summary..... **Error! Bookmark not defined.**
- 9.2. Plans for Unobligated Funding..... **Error! Bookmark not defined.**
- 9.3. FY 2006-2007 Projections **Error! Bookmark not defined.**
- 9.4. Annual Levels of Support..... **Error! Bookmark not defined.**
- 9.5. Cost Sharing **Error! Bookmark not defined.**

1. GENERAL INFORMATION

1.1. Contact Information

Date submitted: April 20, 2006

Reporting period: June 1, 2005 – May 31, 2006

Name of the Center: Team for Research in Ubiquitous Secure Technology

Name of the Center Director: Shankar Sastry

Lead Institution:

University of California, Berkeley
Shankar Sastry
337 Cory Hall
Berkeley, CA 94720-1774
TEL: 510-642-5883
FAX: 510-642-2718
Sastry@eecs.berkeley.edu
<http://trust.eecs.berkeley.edu>

Subaward Institutions:

Carnegie Mellon University
Prof. Mike Reiter
2123 Collaborative Innovation Center
Pittsburgh, PA 15213
TEL: 412-268-1318
FAX: 412-268-6779
reiter@cmu.edu

Cornell University
Prof. Stephen Wicker
386 Rhodes Hall
Ithaca, NY
TEL: 607-255-8817
FAX: 607-255-9072
wicker@ece.cornell.edu

Mills College
Prof. Almudena Konrad
CPM 204
Oakland, CA 94613
TEL: 510-430-2201
FAX: 510-430-3314
akonrad@mills.edu

San Jose State University
Prof. Sigurd Meldal
ENGR 284
San Jose, CA
TEL: 408-924-4151
FAX: 408-924-4153
smeldal@email.sjsu.edu

Team for Research in Ubiquitous Secure Technology
2005-2006 Annual Report

Smith College
Joseph O'Rourke
McConnell Hall 212
Northampton, MA 01063
TEL: 413-585-3673
FAX: 413-585-3786
orourke@cs.smith.edu

Stanford University
Prof. John Mitchell
Gates Building 4B-476
Stanford, CA 94305-9045
TEL: 650-723-8634
FAX: 650-725-7411
mitchell@cs.stanford.edu

Vanderbilt University
Janos Sztipanovits
2015 Terrace Place
VU Station B 356306
Nashville, TN 37235-6306
TEL: 615-343-7572
FAX: 615-343-6702
Janos.Sztipanovits@vanderbilt.edu

1.2. Biographical Information for New Faculty Members (by institution)

Brief biographical information (one page or less) for each new faculty member by institution.

John Canny

John is the Paul and Stacy Jacobs Distinguished Professor of Engineering at the University of California, Berkeley. He is currently working on research with The Berkeley Institute of Design, Flexonics, Multiview, Ubicomp Privacy, Activity-Based Computing, Glaze, PACT, TinyMotion, English Language Learning, SmartSpace, Livenotes, and the Virtual Development Center.

Dawson Engler

I am a joint EE/CS associate professor. Before that, I was an irresponsible graduate student in Frans Kaashoek's PDOS group at MIT's Lab for Computer Science, where I co-founded the exokernel operating system project, which formed the basis of my thesis work. I am a joint EE/CS associate professor. Before that, I was an irresponsible graduate student in Frans Kaashoek's PDOS group at MIT's Lab for Computer Science, where I co-founded the exokernel operating system project, which formed the basis of my thesis work.

Johannes Gehrke

Johannes Gehrke is an Associate Professor in the Department of Computer Science at Cornell University and as Associate Director of the Cornell Theory Center. He obtained his Ph.D. in computer science from the University of Wisconsin-Madison in 1999.

Johannes' research interests are in the areas of data mining, search, data privacy, and applications of database and data mining technology to marketing and the sciences. Johannes has received a National Science Foundation Career Award, an Arthur P. Sloan Fellowship, an

Team for Research in Ubiquitous Secure Technology
2005-2006 Annual Report

IBM Faculty Award, the Cornell College of Engineering James and Mary Tien Excellence in Teaching Award, and the Cornell University Provost's Award for Distinguished Scholarship. He is the author of numerous publications on data mining and database systems, and he co-authored the undergraduate textbook *Database Management Systems* (McGrawHill (2002), currently in its third edition), used at universities all over the world.

Johannes was co-Chair of the 2003 ACM SIGKDD Cup, Program co-Chair of the 2004 ACM International Conference on Knowledge Discovery and Data Mining (KDD 2004), and he will be Program Chair of the 33rd International Conference on Very Large Data Bases (VLDB 2007). At Cornell, Johannes teaches in the Department of Computer Science, the Information Science Program, and in the Johnson Graduate School of Management. He has given courses and tutorials on data mining and data stream processing at international conferences and on Wall Street, and he has extensive industry experience as technical advisor.

Dan Harkey

Dan is co-founder of the Enterprise Software Technologies program (formerly the Client/Server Program) at SJSU. He is co-author of the award-winning and best-selling books: *Client/Server Programming with Java and CORBA* and *Client/Server Survival Guide*. His most recent publishing venture is *Wireless Java Programming for Enterprise Applications*. Dan has over 23 years of experience in industry at IBM and academia. He has spoken at numerous conferences on distributed objects, components, and Java technologies. Dan holds a BS in Electrical Engineering from New Mexico State University and a MS in Computer Science from Santa Clara University.

Current Research and Interests:

Distributed object and component architectures.
Wireless technologies for enterprise applications.
Tool-based software development.

Stephen Maurer

Academic Research. Acting Director, Goldman School Project on Information Technology and Homeland Security. Original published research on open source biology (*Public Library of Science - Medicine*), R&D incentives for drug development (forthcoming), database policy (*Nature, Science*), patent law (*Economica*), and academic/industry transactions (*Research Policy*). Co-leader of National Science Foundation-funded study of how California judges apply precedent.

Teaching. University of California (Berkeley) , Lecturer, presenting graduate-level courses on internet law and economics and technology (Cyberlife, Science Policy., and Information Technology and Public Policy). Invited speaker at intellectual property conferences hosted by US National Academy of Sciences, US National Institutes of Health, US Department of Transportation, The Human Genome Organization, Duke University Law School, Stanford University, and The University of California at Berkeley.

Intellectual Property and Litigation Attorney. Practiced intellectual property and high technology litigation at leading law firms since 1982. Handled complex, high-value cases for clients including IBM (computer hardware), Apple Computer and Symantec (software), ROLM (computerized telephone systems), UTC (advanced composite irrigation pipe), Zilog (semiconductor chip design), Tegal Corporation (microchip fabrication tools), Aerojet General Corporation (rocket engines), and The Navajo Nation (boundary dispute). Responsible for preparing and/or examining over one dozen trial witnesses in \$150 million insurance case.

Team for Research in Ubiquitous Secure Technology
2005-2006 Annual Report

Member of the California Bar.

Policy Analysis and Consulting. Performed sponsored research for The US National Academy of Sciences (academic/industrial research agreements; database protection legislation) and Industry Canada (US and European database policies). Performed consulting services for Diversified Risk Management (designed novel insurance policy for intellectual property); Mutations Database Initiative (negotiated \$2.3 million collaboration between academic scientist organization and Incyte Pharmaceutical Company); and Virtual Physics Associates (co-leader of group seeking to build advanced nuclear physics database at Lawrence Berkeley Laboratory).

Deirdre Mulligan

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, forthcoming 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.

Vern Paxon

I received my M.S. and Ph.D. degrees from the University of California, Berkeley, and have been (and continue to be) a staff scientist with the Lawrence Berkeley National Laboratory's Network Research Group for many years. I began at the ICIR group of the International Computer Science Institute (ICSI) in 1999.

My main active research projects are Bro, CCIED (the NSF-sponsored *Collaborative Center for Internet Epidemiology and Defenses*, a joint effort with UC San Diego), DETER, and PREDICT. Much of the current emphasis of CCIED is on **Internet worms**, including our network telescope project. A significant milestone with our PREDICT efforts is the public release of enterprise header traces: anonymized packet headers of traffic recorded inside LBNL. My professional activities include: Vice chair of ACM SIGCOMM; Program co-chair for IEEE

Team for Research in Ubiquitous Secure Technology
2005-2006 Annual Report

Security and Privacy 2006 (as well as IEEE Security and Privacy 2005); and Program committee member for SRUTI 2005, RAID 2005, ACSAC 2005, and USENIX/ACM NSDI '05.

Simon Shim

Simon Shim is an Associate Professor in Computer Engineering department at San Jose State University. He received M.S. from Rensselaer Polytechnic Institute, and Ph.D. from University of Minnesota both in Computer Science. His expertise includes High speed security server, Internet Computing, Multimedia Database and SAN. He is a co-director of the Internet Technology Laboratory which is supported by grants from Intel, Microsoft, Wytec, and Informix corporation. He has authored and co-authored more than 35 technical publications in IEEE Computer, IEEE Transaction on Knowledge and Data Engineering, IEEE Concurrency, ACM MONET, and Journal of Multimedia Tools and Applications (Kluwer Academic Publishers). His work was recognized by IBM Rochester as a significant Research Contribution in the study of SAN with 4 colleagues. He has been served as the workshop chair of IEEE E-Commerce conference, co-chair of IEEE workshop on Mobile Commerce and Services, and co-chair of IEEE workshop on Data Engineering Issues in E-commerce. He served as the lead guest editor of the special issue on high-speed security for IEEE Computer, June 2004. He has served in many international conferences/workshops as a technical committee member.

Einar Vollset

I work in the Systems group, primarily with Ken Birman and Robbert van Renesse, as well as some extremely bright PhD students. I got my PhD in 2005 at the University of Newcastle, UK, in the area of fault-tolerance in wireless ad-hoc networks, but more broadly I'm interested in distributed systems and ubiquitous computing, particularly as these relates to scalability and reliability.

Currently I'm the lead on an effort to build a system for settings where traditional end-to-end connectivity is not necessarily present. You can envisage such systems as consisting of most wireless devices you have and use today (your laptop, your pda, your phone, your mobile rocket launcher..). With such devices, even if you have direct connectivity to the internet (through for example the GSM card in your laptop), such connectivity may not be the only or necessarily the best way to communicate. What I'm currently thinking about is how you can build a trustworthy, scalable systems which exploit these multiple forms of connectivity; what kind of applications become possible once you can say "I want to share this file with everyone I know in this room for the next 5 minutes" or "This message should get to all robots in search party Mars East" and when adding devices to the network *increases* usability?

Yuan Xue

I am an assistant professor of Computer Science at Vanderbilt University. My research area includes networking and distributed systems with a focus on wireless networks, mobile systems, and network security. My doctoral research applies optimization theory and micro-economic theory to the design of optimal resource allocation mechanisms in wireless networks. At Vanderbilt, I will conduct research on providing protection, performance optimization, and quality of service support for next generation networking and distributed systems.

Weider Yu

I started my adjunct teaching in the department of Electrical Engineering and Computer Science, University of Illinois at Chicago, in 1986, after I joined Bell Laboratories. After teaching for approximately 14 years, I have enjoyed each and every course I have taught. The time I had in the classes with my students was very valuable. I like to interact with students. The

interactions with them always made me energetic and satisfied.

I have had strong enthusiasm and interest in research activities since my graduation from Northwestern University. Due to my concentration in Software Engineering in the Ph.D. program, I felt that it was important to gain advanced and practical working experience in the field of industrial applied research.

For the past 18 years with Bell Laboratories, I have had great opportunities to lead various teams in researching many innovative and challenging projects and to interact with many reputable researchers in Bell Labs. My accumulated work essentially spanned the entire domain of software engineering. The majority of my work at Bell Labs has been in the area of applied software engineering research and advanced technologies.

The work environment involves R&D software and hardware embedded development for some very large and complex advanced digital and packet switching real-time projects. There are thousands of scientists and engineers working on the projects. The environment has given me numerous excellent opportunities to work on multiple areas in Software Engineering.

I have worked in the following areas:

- software systems and development methodology and process (product planning analysis, requirement specification, architecture, design, implementation, verification and validation)
- distributed software engineering technology
- distributed communication software systems
- software development environment
- software estimation technology
- software metrics (all kinds of product and process metrics, and organizational/enterprise metrics)
- requirement traceability engineering
- requirement specifications for distributed systems
- high level specification language design
- knowledge based systems
- data warehousing and data mining
- C programming language guidelines and techniques
- relational database software specification language automation
- software faults (fault flow and fault removal) analysis
- software quality assurance
- software tools
- software fault prevention tools and techniques
- software process engineering goals and organizational entrepreneurial/strategic goals
- software speed, cost and quality improvement techniques

I am very interested in performing research in the majority of the Software Engineering and closely related areas. With my many years of experience in applied research and high tech industry, I am also interested in applying my experience and insight in an integrated and cohesive manner to the state of the art in Software Engineering research in the emerging areas, such as mobile systems, wireless systems and embedded systems.

1.3. Executive Summary

The Executive Summary provides a brief overview of the vision, goals, plans, and performance and management indicators for the Center. Any significant changes from the original plans for the Center

Team for Research in Ubiquitous Secure Technology
2005-2006 Annual Report

should be described. This section also reports on progress toward meeting the goals set for the Center (described in detail in the remaining sections) and provides an overview of significant accomplishments (one of the four pages should highlight these in nugget form) during the reporting period.

The Team for Research in Ubiquitous Security Technology (TRUST) was created in response to a growing sense of urgency in dealing with all aspects of cybersecurity as it affects society. First, 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. Second, and somewhat contradictorily, many such control systems remain untrustworthy. Waves of viruses and worms sweep the Internet, and exhibit increasing virulence and rate of spread that are also directly proportional to their growing ease of deployment. Privacy remains poorly understood and poorly supported; security is generally inadequate, and some speak of a “market failure” in the domain. Broader issues of software usability, reliability and correctness remain challenging. Industry stakeholders are unable to recruit new employees adequately trained in these technologies. Society is placing computers into critical roles, although they do not meet the requirements of trust.

TRUST is composed of several universities—Berkeley, Carnegie Mellon, Cornell, San Jose State, Stanford and Vanderbilt—which have joined forces to organize a multifaceted response. The Team for Research in Ubiquitous Secure Technology (TRUST) represents the strongest and most diverse engagement of the issue of trusted systems ever assembled. TRUST is the first to recognize the breadth of the problem and to combine fundamental science with a broader multidisciplinary focus on economic, social and legal considerations and a substantial educational mission. TRUST will enable dialog with stakeholders whose needs simply cannot be approached in a narrower and purely technical manner, or by any single research group. TRUST seeks to be an intermediary between the policy makers and society at large on the one hand, and the researchers, academics, and industrial providers of services and technology on the other.

TRUST seeks to achieve its mission through research as well as a global policy for engaging in education of society as a whole. This, the first annual report of TRUST details the experience of the center along many dimensions—research, industrial outreach and knowledge transfer, education and diversity outreach.

In research, TRUST has achieved success along several fronts—in model-based integration of trusted components and co-design of networked embedded systems, in the creation of new software tools for monitoring and controlling large sensor infrastructures, creation of integrative testbeds for critical infrastructures, in understanding privacy and other legal issues surrounding identity theft, and designing tools for ant-phishing technology, etc. All these are reported in detail in the research thrusts area of this report.

In education, TRUST is leveraging an existing learning technology infrastructure to quickly enable TRUST courseware and material to be assembled, deposited in a repository, and adapted for wide web-based content dissemination. In addition to developing special courses for undergraduate and graduate curricula, and regular seminars in all campuses as well as webcasts, TRUST has hosted a series of workshops on sensor networks, privacy, identity theft, electronic medical records. The major thrust in the second year will be the TRUST Academy Online (TAO) and the Education Community Development efforts. Again, all these are reported below in the section on education.

Team for Research in Ubiquitous Secure Technology
2005-2006 Annual Report

We have begun an aggressive program of knowledge transfer to industry (from bug reports of open source software to tools such as Spoofguard and various consulting activities) and active engagement with governmental agencies such as DHS, AFOSR, DoD, and DoE concerned with issues of security. Also, TRUST has a large and growing set of industrial partners such as Microsoft, Sun, Intel, UTRC, etc, with whom we are beginning to engage in collaborations of mutual interest. For example, Telecom-Italia will harvest the incipient technology that comes out of TRUST in the healthcare sector to better understand and build upon its own base.

TRUST has an ambitious goal of reaching a diversity goal of 30% of women in its faculty and students, and 10% of researchers from underrepresented communities, and has been proactive in this regard. Several activities for enhancing diversity are reported in the corresponding section.

Overall, we are happy to report that the center is making excellent progress towards its goals, its participants are actively engaged, and the outlook is positive.

2. RESEARCH

2.1. Research Objectives

In the current reporting period, how have the Center's overall research objectives and plans changed from the previous reporting period? What performance and management indicators has the Center developed to assess progress in meeting its research objectives?

This is the first reporting period of TRUST.

The TRUST vision is to provide a unique opportunity for a wide range of cybersecurity issues to be addressed from many points of view—technological, scientific, social, policy and legal. Of paramount importance to TRUST is the creation of a science that will simultaneously address the imperatives of all these points of view and allow scientists and technology developers, policy makers, and social scientists to make informed and rigorous decisions, with the full understanding of tradeoffs involved. We think that this new science, though exciting and far-reaching, will of necessity come about from an evolution in more traditional areas that impinge on this “science of TRUST” as theory and praxis of these areas coevolve. In particular, the primary areas of new science creation include new cryptographic protocols and supporting systems, high confidence software science, security functionality, policy and management, and complex interconnected networked systems. Furthermore, TRUST will have strong, well proven ties with IT vendors and infrastructure providers, which will serve to both ground its research in real-world problems and ensure avenues for knowledge transfer. TRUST will have a significant impact at a national scale, as its research results will lead to new concepts and doctrine for application to: public policy issues around privacy, access control, and security; technology for protecting and preventing information security breaches; and increased protection of the nation's critical infrastructures, most notably in the areas of telecommunication, healthcare, electric power, finance and Department of Defense networks.

The overall research objectives are as outlined in the Strategic and Implementation Plan (April 2006). TRUST projects are both continuously and periodically monitored for meeting the research objectives. Periodic monitoring consists of biannual meetings of TRUST as a whole where progress in each research thrust area is formally reviewed. Continuous monitoring consists of monitoring both by the project leaders in Research thrust areas as well as by the Executive Board. (The reports of thrust area project leaders are reproduced below.) The evaluation metrics are as indicated at length in the Strategic and Implementation Plan, and are outlined in the table below.

Objective	Metric	Frequency
Scientific Impact	Publications, Presentations, Recognition	Annual
Technological Impact	Transitions, Industry interest	Annual
Timeliness	Milestone completion	Semi-annual
Social Impact	Policy Papers, Legal Policy	Annual

2.2. Current and Anticipated Problems

Discuss any problems the Center has encountered in making progress toward its research goals during the reporting period as well as any problems anticipated in the next period. Include plans for addressing these problems.

No significant problems were encountered.

2.3. Research Thrust Areas

Briefly describe the research thrust areas at the Center. Provide basic information (thrust name, principal investigator's name, participants' names and status, funding in the reporting year from NSF and other sources, funding anticipated for the next year from NSF and other sources) for each thrust area and details of significant accomplishments (i.e., publications, presentations, summary of student and postdoctoral researchers' involvement in research activities of the Center) during the reporting period. For each thrust area, a narrative should describe the goals, activities, and outcomes and/or impacts in the current reporting period; plans for the next reporting period with attention to any major changes in research direction or level of activity; and how the activities enable the Center to meet its goals.

TRUST has been organized into several project areas. During the first year, we started with an organization, as described in the Strategic and Implementation Plan, of project areas into 11 challenge areas. Later this organization was modified, for reasons explained in the revision of the Strategic Plan, resulting in a renaming of project areas and division of labor in the research.

The following report of thrust areas represents an intermediate stage of activity, but from the detailed descriptions below the mapping to thrust areas mentioned in the S&I Plan should be clear.

The 7 thrust areas reported below have the following leaders:

1. Model-based Integration of Trusted Components, Janos Sztipanovits
2. Secure Information Management Tools, Ken Birman
3. Integrative Projects, Janos Sztipanovits and John Mitchell
4. Trusted Platforms and Trustworthy Systems, Mike Reiter
5. Secure Sensor Networks, Steve Wicker
6. Network Defense, Anthony Joseph
7. Privacy and Information Forensics, Doug Tygar

2.3.1. Thrust Area I: Model-Based Integration of Trusted Components

Co-design Environment for Networked Embedded Systems

We started the project by discussing challenge problems with industry stakeholders. Security of networked embedded systems is a major concern in the automotive and process industry and in defense. General Motors, Honeywell, Raytheon and Boeing researchers helped us in defining a physical experimental platform and related real-life design challenges. We consider the project described below a seedling for an integrative project in following years.

Experimental Platform

In order to facilitate the experimentation with security-enabled embedded systems, we have designed and constructed an experimental platform, shown on the figure below.

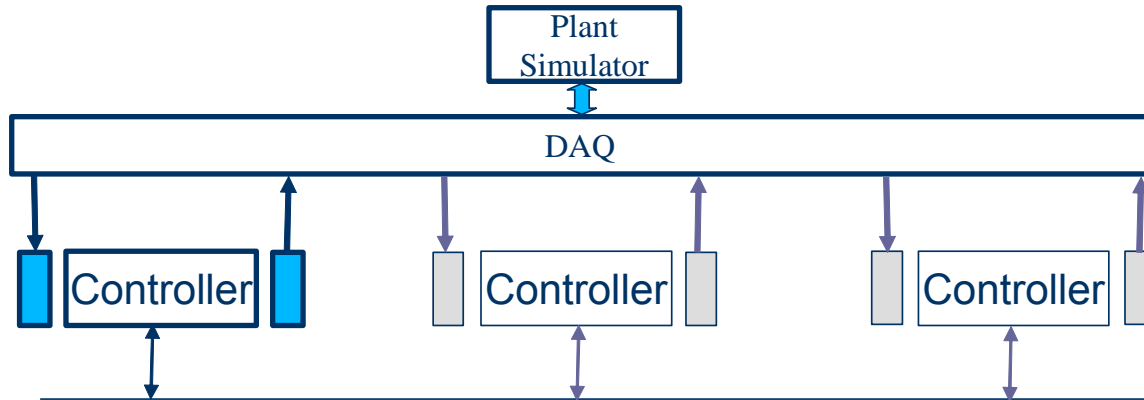


Figure 1: Experimental platform for security experiments on embedded systems

The platform consists of a plant simulator and several, networked controllers whose secure behavior is of interest. The plant simulator is a high-end PC, equipped with analog and digital input and output channels, running a real-time simulation of some physical process. The hardware interfaces of the simulator are such that they are indistinguishable from the real plant interfaces from the viewpoint of controllers.

For controllers, we have purchased SBC4495 boards manufactured by Micro/Sys. This board has a 486 compatible processor running at 133 MHz with 64 MB of RAM, and it has a number of analog and digital interfaces. The 10/100BASE-T Ethernet and PCMCIA slot allows flexibility to easily integrate the SBC4495 into wired or wireless embedded system applications. The controllers are connected through their analog and digital interfaces to the data acquisition board of the plant simulator, thus receiving and sending real-time, electrical signals, just as they would do in a real plant environment. The controllers have their own (wired and wireless) LAN connections, for controller communication and coordination.

The experimental platform above provides opportunity to configure it according to the characteristics of a wide range of embedded applications, such as SCADA systems or Electronic Control Units in cars.

Co-Design Environment

As an initial effort under TRUST, we have taken an industry-standard architecture modeling language the Architecture Analysis and Design Language (AADL) and extended it with security aspects. AADL, standardized by SAE, is a language for modeling relevant aspects of the architecture of embedded systems (like avionics and automotive control systems). While it provides a visual and textual syntax to capture architectural views, it does not have any support for incorporating security aspects into the design. Related, existing work proposes adding security aspects to UML, the Unified Modeling Language, however no such proposals exist yet for comparable languages in embedded systems.

The AADL security extensions address security in two aspects: (1) Role-based Access Control (RBAC), and (2) Partitions. Using the MIC toolsuite as foundation we have created a metamodel the AADL and instantiated a domain-specific modeling environment for AADL using GME. Next we have extended the basic AADL metamodel with a metamodel for RBAC, effectively extending the modeling language with support for RBAC. Additionally, we have added support for ARINC 653-style partitioning to the modeling language such component could be placed into protected and isolation partitions (if the run-time platform provides them). The RBAC models

allow specification access control on the level of components, and thus system-level security properties could be checked easily on the models.

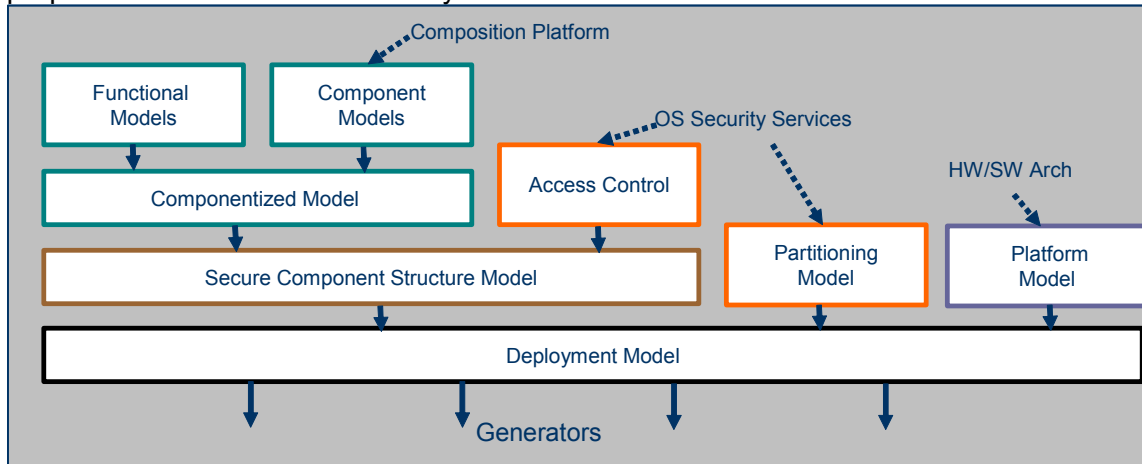


Figure 2: AADL Extended with Security Aspects

The figure above shows the resulting extended AADL modeling language. The baseline AADL supports functional and component modeling and HW/SW architectural modeling. The newly added access control models support the modeling of the secure component structure, while the partitioning model extends the existing deployment models. AADL equipped with the security extensions is supported by the Generic Modeling Environment: a graphical modeling tool. We have created a number of example models for illustrating the concepts of the security extensions. We have also built software generators that generate C/C++ (glue-) code and configuration files for a secure embedded platform for experimentation. Functional code has to be provided through other means. The modeling environment and the generators constitute the first prototype of an Embedded System Security Co-Design Environment.

Experiments

We were able to run initial experiments with the testbed described above that allowed us to study the impact of network attacks on the behavior of controllers (and thus on the controlled “plant”), and the impact of the overhead of security extensions on the control algorithms. In these experiments we have used Linux as the OS for the controllers. We are in the process of running more complex experiments on the testbed, including experiments with a DO-178B-compatible separation kernel (LynxOS).

Trusted Integration Platform for Enterprise Distributed Real-Time Embedded Systems
Enterprise distributed real-time and embedded (DRE) systems, such as those found in aerospace, defense, telecommunications and healthcare domains, have stringent, simultaneous QoS requirements in terms of performance attributes such as latency and throughput, and dependability attributes, such as resilience to failures, availability, reliability and security. Despite rigorous software design, testing, and certification processes, however, it is hard to eliminate (a) unplanned failures, including hardware, software, or network link failures, (b) vulnerabilities including unauthorized access, and (c) performance degradations due to resource exhaustion or denial of service attacks. Critical to survivable DRE systems is a powerful set of (1) security features, such as integrity, confidentiality, authentication, authorization, and delegation, and (2) dependability features, such as availability, reliability and fault tolerance. Integrity and confidentiality technologies are readily available and usable. However, the latter three security features are not. Without them, developers of DRE systems must rely on proprietary equivalents that may not be interoperable nor can be proven.

Dependability management techniques for high availability typically use redundancy via replication of critical components as an approach to addressing problems caused by failures and faults.

Implementing survivability schemes in DRE systems using third-generation languages is hard, however, and can distract application developers from their primary job of creating business logic. Component middleware platforms, such as Enterprise Java Beans and the CORBA Component Model (CCM), have become popular for conventional distributed systems because they provide effective reuse of the core intellectual property (i.e., the "business logic") of an enterprise. Although these middleware platforms have several desired characteristics, they are not yet suitable for survivable enterprise distributed real-time and embedded (DRE) systems, however, due to their lack of support for the following DRE systems requirements:

Survivable heterogeneous service assemblies, where end-to-end missions in enterprise DRE systems are realized by assembling services offered by the different subsystems executing on (potentially) heterogeneous platforms/machines. Hence survivability mechanisms must no longer be restricted to a single node, but needed to be coordinated across different machines in which the services execute. Such survivability management considerations must account for semantic and operational compatibilities among the interacting services for dependability attributes including (a) replica deployments, configuration, and failure management and (b) resource exhaustion and denial of service attacks, and security attributes including (a) assembly-wide integrity and confidentiality and (b) assembly-wide authentication and authorization. Resolving these challenges requires appropriate deployment assignments for the services, accurate dependability, trustworthiness and consistency support for the services including diverse replication schemes, failover granularities among the participating services, and services-wide state synchronization and authentication mechanisms.

Mission-driven quality of service (QoS), where different missions in the system will have different importance thereby requiring the survivability of the most mission critical elements. Middleware support is needed for provisioning different survivability requirements that may be heavyweight, in the case of stateful applications, and lightweight, in the case of stateless applications. Different missions require different strategies, such as replication requirements and authentication across individual or groups of services, their failover strategies, the quorum of replicas required to grant access or for the correctness of the mission, and style of replication used in individual subsystems. Resolving these challenges requires design-time specification of survivability requirements, deployment time provisioning of middleware support that is needed to provide the required survivability needs and subsequent enforcement of these requirements during runtime. The middleware can thereby be configured to be heavyweight or lightweight depending on the needs of the missions.

Handcrafting solutions to these problems within middleware platforms do not scale to larger DRE systems nor are they reusable across different DRE systems and domains. Model-Driven Engineering (MDE) is a promising approach to address the survivability challenges described above by raising the abstraction of system design to a level higher than that provided by third-generation programming languages. Our goal is to apply MDE to design and deploy survivable DRE systems so that these concerns can be considered earlier in the software development lifecycle rather than as an afterthought. MDE approaches also provide a basis for easier reasoning of the system properties. In particular, our R&D focuses on modeling survivability concerns of the system focusing on (1) replication by means of redundancy or clustering, which improves reliability (2) attributes like failover granularities and degree of replication and replication styles being used, which improves fault tolerance, (3) group-wide authentication and

authorization, which serve to maintain system integrity and confidentiality, and (4) collecting and satisfying requirements involving the placement decisions and actual deployment of replicas into appropriate target nodes of the system, which helps improve availability.

At the heart of our MDE approach to survivable enterprise DRE systems is a domain-specific modeling language (DSML) that incorporates dependability concerns at the design and deployment phases of Lightweight CCM-based DRE systems, including specification, assembly, and packaging. Enterprise DRE systems are constructed in the CCM world as monolithic components or assembly of components, and allow exchange of data and control among the participating components so that end to end services can be composed out of disparate services. The important components of services could be individual monolithic components, or a partial assembly of components within the service, or the whole assembly of components within the service.

Our research on deployment and runtime solutions for survivable systems focuses on the following novel ideas: (1) exploring novel architectures to deploy and manage groups of services for the purposes of replication, synchronization, consistency management, and failover granularities, (2) extension of OMG deployment and configuration (D&C) metadata to capture the different dependability and security requirements of DRE services to provide configurable deployment and runtime support for replica management like rejuvenation and reconstitution, (3) use of publish/subscribe services to provide topics-based distributed failure notifications or unauthorized accesses of processes and the hosted services, thereby ensuring scalability as well as controlling notification explosions, (4) extension of deployment, configuration, and lifecycle management capabilities to provision operating environments (for example, state synchronization transports, credential verification mechanisms, membership data and control transports) tailored for the services, their replicas and their requirements, and (5) extension of FT CORBA mechanisms, such as interoperable object group references (IOGRs), to enhance dependability in component based enterprise DRE systems.

Our next set of R&D tasks will involve incorporating security attributes discussed earlier. In particular we will incorporate our earlier work on Common Secure Interoperability (CSI) and Authorization Token Layer Acquisition Service (ATLAS) within the Lightweight CCM middleware. Additionally, we will enhance our MDE tools to incorporate the security dimensions. We have also begun collaborating with Ken Birman's group at Cornell on applying service placement algorithms to their research on scalable services architecture, which provides high availability mechanisms for hosting applications, such as web services, in cluster environments. Additionally our middleware-based fault tolerance mechanisms will benefit from the low-latency multicast protocols, such as Ricochet, being developed by Ken's group. Together we will collaborate on enhancing our synergistic research to enhance security and thereby survivability of large distributed systems.

2.3.2. Thrust Area II: Secure Information Management Software Tools

Description of Effort

The TRUST software tools effort focuses on the development of new software tools for monitoring and controlling large sensor infrastructures. Few "robust" communications architectures are known for scalability. We are finding that by adopting probabilistic goals, we can break through this barrier. Our new approach combines peer-to-peer protocols with what are called epidemic or gossip algorithms. By demonstrating a new generation of robust software platforms that scale extremely well, combine rigorous semantics with good performance, and have user-friendly API's, we can enable the creation of a tremendous variety

of new control and monitoring solutions for nationally critical infrastructure.

Examples of Specific Projects

Fireflies: This was funded in part by other NSF grants, and by DARPA, AFRL and AFOSR. Fireflies is an effort to develop a new generation of extremely robust overlay networks to support content distribution in potentially large-scale settings subject to attacks up to and including Byzantine attack scenarios. (Cornell Principal Research Scientist Van Renesse, PhD student Maya Haridasan, Havard Johansen).

Quicksilver: This work, funded in part by DARPA, AFRL and AFOSR, explores scalability for a publish-subscribe style of event notification platform, using peer-to-peer techniques and other methods. The platform is now operational and achieves a true breakthrough in scalability and performance; a series of papers are in preparation to discuss the mechanisms by which this was achieved. We are also extending Quicksilver with a strong type system and with a fault-tolerance and consistency model; these steps will offer an exceptionally flexible, robust and scalable framework within which type checking can play a role as part of a stronger security architecture. (Birman, PhD candidate Krzysztof Ostrowski)

Ricochet and Slingshot: This work was funded in part by Intel, DARPA, AFRL and AFOSR. Ricochet is a new protocol for time-critical data replication in clusters and data center computing platforms. It introduces the concept of lateral error correction and with it, demonstrates three orders of magnitude better delivery for use in settings requiring time-critical multicast or data updates. Slingshot was an earlier protocol in which we first introduced the notion of lateral error correction, but within a single multicast group at a time. Thus, Ricochet extends Slingshot to multigroup settings. Ricochet achieves far better scalability in the numbers of groups than in any prior work. (Birman, PhD students Mahesh Balakrishnan and Amar Phanishayee). We have begun to collaborate with Vanderbilt (Doug Schmidt) on aspects of this work.

Tempest: This work was funded in part by Intel, DARPA, AFRL and AFOSR. Tempest is a new platform that runs over Ricochet and automates most aspects of developing new scalable and robust services to run on data centers and clusters. Tempest provides automated data replication, query load-balancing, fault-tolerance and data repair after faults that introduce inconsistency. (Birman, PhD students Tudor Marian, Mahesh Balakrishnan and Amar Phanishayee). We have begun to collaborate with Vanderbilt (Doug Schmidt) on aspects of this work.

Beehive: Beehive is a high-performance distributed hash table. A novel optimization technique enables Beehive to respond to queries quickly, tolerate denial of service attacks, and balance load. We have used Beehive to build new, resilient infrastructure services for the Internet. CoDoNs is a safety net and a replacement for the Domain Name System that provides strong security, performance, and fast dynamic updates for existing Internet names. CobWeb is an Akamai-like open-access content distribution network. CorONA is a high-performance publish-subscribe system for web micronews. (Sirer, with graduate student Venu Ramasubramanian).

Octant: Octant is a system for determining the physical location of Internet hosts. Given a host, Octant determines the boundaries of the region in which the node is likely to lie. Behind the scenes, Octant consists of two parts: (1) a comprehensive framework for efficiently representing and combining a system of constraints, and (2) a set of mechanisms for extracting useful and tight constraints on where nodes are likely to be, without resulting in an overconstrained system. (Sirer, with PhD student Saikat Guha and undergraduate Rohan Murty).

Team for Research in Ubiquitous Secure Technology
2005-2006 Annual Report

Meridian: Meridian is a peer-to-peer overlay network for performing location-aware node and path selection in large-scale distributed systems. It is simple to deploy, robust to churn, and can accurately find the nearest node, pick the most centrally placed node, and find a node that fits latency constraints.

Credence: Credence is a reputation system for peer-to-peer networks, designed to provide an accurate metric for the trustworthiness of labels associated with shared files. It differs from previous work in that it derives its trust metric from votes on objects (since voting on peers is not feasible in a p2p system with anonymous participants), has a completely distributed design with no fully trusted nodes, and a concrete implementation. The Credence implementation is free, open-source, and backwards compatible with Gnutella. (Sirer with PhD student Kevin Walsh).

CorSSO: CorSSO is a distributed authentication service that provides network identities that span multiple application services, also known as single sign-on. It enables authentication functionality to be factored out of application services and delegated to combinations of authentication servers. It uses threshold cryptography for efficiency, fault tolerance and resilience against attackers.

MagnetOS: MagnetOS is an operating system for ad hoc networks. It makes the entire network appear as a single Java virtual machine. It enables applications to be constructed easily and to execute efficiently. (Sirer with multiple student collaborators).

Nexus: Still under development, Nexus is a microkernel for exploiting hardware trusted computing technology to perform secure attestation. It enables applications to obtain hardware-based signatures for data. For example, suppose that a camera is used to photograph a holdup. Today we have little help for verifying that the image was really taken on at the time and place claimed. With Nexus, a GPS unit can produce hardware attested time and date stamps and these can be overlaid on the digital image, with the camera signing both. The resulting image is one in which the hardware attests to the time and location at which this specific image was taken. (Sirer and Schneider with graduate student Dan Grossman).

The Component Integrated ACE ORB (CIAO): This work, funded in part by DARPA, Raytheon, Lockheed Martin, Symantec, and Siemens, provides a powerful component-based abstractions using the specification, validation, packaging, configuration, and deployment techniques defined by the Lightweight CORBA Component Model (CCM) and Deployment and Configuration specifications. CIAO integrates the Lightweight CCM capabilities with Real-time CORBA features, such as thread-pools and client-propagated and server-declared policies. CIAO is being used on commercial and military projects, including major programs at Lockheed Martin and Raytheon. (Schmidt, PhD candidate Jaiganesh Balasubramanian)

The Resource Allocation and Control Engine (RACE): This work, funded in part by DARPA, Raytheon, and Lockheed Martin, provides a component framework for managing the use of various nodal system resources (such as network bandwidth, CPU, and memory) by selectively applying algorithms designed to meet application QoS requirements. RACE is being used in R&D programs at Lockheed Martin and Raytheon. (Schmidt, PhD candidate Nishanth Shankaran)

Component Synthesis with Model Integrated Computing (CoSMIC): This work, funded in part by DARPA, Raytheon, and Lockheed Martin, provides a suite of domain-specific modeling languages and their associated analysis/synthesis tools that support various phases of component-based distributed real-time and embedded system development, assembly,

Team for Research in Ubiquitous Secure Technology
2005-2006 Annual Report

configuration, and deployment. CoSMIC been used in R&D efforts at Lockheed Martin, Raytheon, and Siemens and is being transitioned to the Navy's DD(X) program. (Schmidt, PhD candidate Krishnakumar Balasubramanian). We are beginning to collaborate with Cornell (Ken Birman) on aspects of this work.

TAO Data Distribution Service (DDS): This work, funded in part by DARPA, AFRL, and AFOSR, provides a scalable publish/subscribe platform that enables tactical information management systems to specify and enforce performance requirements between different parts of tactical information management systems using quality of service (QoS) parameters that (1) configure the networks, operating systems, and middleware and (2) establish contracts that precisely specify a wide variety of QoS properties. (Schmidt, PhD candidate Jeff Parsons). We are beginning to collaborate with Cornell (Ken Birman) and CMU (Mike Reiter) on aspects of this work.

The Tele-Immersive System: We are developing a peer-to-peer system that enables geographically distributed people to meet (visually), and interact in the cyberspace. In addition to the capabilities the people interaction, the users can simultaneously also interact with any 3D data, such as radiological data sets, and/or nay scientific data sets. The goals of this project is to develop an end to end system of video, audio streamed data in real time utilizing off the shelf equipment and standard Networking facilities such as Internet2 and Lambda Rail. The outcomes are development of video and audio capture of dancing people. The research issues relevant to TRUST program are guaranteeing privacy of the streamed data. (R. Bajcsy, Samuel Johnston, Ross Diankov, Edgar Lobaton, and Klara Nahrstedt (UIUC)) (This work has been supported in part by grants from HP and NSF.)

2.3.3. Thrust Area III: Integrative Projects

Description of Effort

The central goal of TRUST is the deeper understanding of and scientific foundation for analyzing the interaction between security, systems science and economic policy. Integrative projects play an important role in achieving this goal. Specifically, integrative projects:

- connect research efforts to real-life, national-scale challenges,
- provide context for integrating research discipline-oriented research efforts,
- help validating research results and
- facilitate technology transitioning toward National stakeholders.

A crucial tenet of our approach is to use carefully selected testbeds, which translate the national needs to specific research challenges. These research challenges contribute to defining and updating our technical agenda, providing a framework for combining core scientific areas such as cryptography, secure software design, modeling, distributed and embedded systems, economics, and information management into an integrated foundation for understanding and managing security functions and vulnerabilities.

During the past 8 months, we have refined our plans in identifying integrative projects and started the efforts described below.

Examples of Specific Projects

Electronic Medical Records (EMR)

Societal Context: Computer technology, patient sensors, and networking are revolutionizing several aspects of healthcare and medical information processing. Small wireless sensors will free many patients from managed care facilities, while providing timely medical assistance when needed. At the other end of the spectrum, virtually all patients will soon gain greater control over their records and treatment options through web portals. The TRUST EMR project addresses the complex security and privacy issues emerging from the rapidly increasing use of electronic media for the archival and access of patient records. This change is driven and strongly influenced by the Health Insurance Portability and Accountability Act of 1996 (HIPAA). EMR has become an area where technology, public policy and individual interests intersect and conflict, making the development of information systems for EMR archiving and access a very challenging problem. There is clear evidence that without a detailed understanding of the relevant issues on all sides, an acceptable solution cannot and will not emerge.

Integrative Testbed: The integrative project will leverage a cooperative relationship established with the Informatics Institute and the Biomedical Informatics Department of the Vanderbilt University Medical Center (VUMC). The MyHealthAtVanderbilt system – a functioning experimental patient portal – is a unique resource that will be used as the basis for experimentation and interaction through real-life deployment scenarios. The MyHealth portal has enrolled over 8,000 patients and is growing at the rate of more than 1,000 new enrollees per month, making it one of the largest operational healthcare portals in the world. MyHealthAtVanderbilt gives patients secure messaging with their providers, the ability to make appointments online, see the contents of their medical records, and request changes to care plans and records. Experience with this portal has highlighted the inadequacy of our current understanding of the interdependences among aspects of security and privacy, systems design and policy.

Team for Research in Ubiquitous Secure Technology
2005-2006 Annual Report

Research Collaboration: External collaborators in the EMR project are:

- VUMC: Their contribution to the project are (a) access to the MyHealthAtVanderbilt patient portal and (b) collaborating research team working on medical informatics aspects and clinical trial.
- Information Technology for Assisted Living at Home project at Berkeley: Contribution to the project are (a) smart sensing technologies that enable alert monitoring and long-term out-patient biometric data and (b) tools to integrate sensor systems into EMR.

Internal collaborations in the EMR project are extensive, all partner Universities participate in one or more research areas in the project.

Activities: Our work focused on identification of challenges, establishing collaborative relationship with VUMC and planning.

- Dan Masys (VUMC): “Electronic Medical Records and Secure Patient Portals as an application domain for Team Research in Ubiquitous Secure Technologies,” White Paper, October 2005
- Dan Masys (VUMC) and Janos Sztipanovits VU-ISIS co-chairs: Design Workshop for an Integrative Project related to Patient Portals Vanderbilt Center for Better Health, Nashville, TN, December 16, 2005
- Janos Sztipanovits, Ruzena Bajcsy, Mike Eklund and Shankar Sastry: “TRUST Electronic Medical Record Proposal,” March 2006

Trusted Federated Sensor Networks (TuFNet)

Societal Context: Threats and vulnerabilities continuously emerge, morph, and change in complex, dynamic societies and economies. This integrative project focuses on technology foundations required for sensor-based dynamic threat analysis and disaster prevention. Threat characterization, vulnerability and risk analysis in dynamic, changing environments requires networked, real-time, sensor-based observations and usually demand multi-level security cutting across jurisdictional boundaries. Real-time sensor data is expected to be received from federated sensor networks deployed in targeted areas of critical infrastructure systems. The dynamic system state space required for vulnerability and threat analysis is much richer than that provided by real-time sensor data observations and many state variables are not even observable. It means that data streams must be interfaced with federated systems of simulators. This heterogeneity has large implications on the network and computing architecture. In this integrative project we will pursue an solution that not only considers the need for information protection of sensor networks, but also provides an end-to-end seamless security and privacy support over the federated network.

Integrative Testbed: The project will leverage a cooperative relationship established with the Computer Science and Mathematics Directorate (CSMD) of ORNL. Sensor networks for homeland security applications are a major focus for ORNL. Combining multi-modal real-time sensor input, sensor fusion and multi-resolution simulation of critical infrastructures have the potential for dynamic vulnerability analysis. ORNL has one of the largest high-performance computing center in the Nation and has ongoing program on sensor network testbeds. In this integrative project TRUST research groups and ORNL will create the technology foundation for the federation of secure sensor networks. The research will focus on:

- attack modeling and vulnerability analysis:

Team for Research in Ubiquitous Secure Technology
2005-2006 Annual Report

- security agreement in network federation:
- performance and security tradeoffs:.
- privacy and market issues:.

Research Collaboration: External collaborator in the TuFNet project:

ORNL: Their contribution to the project are (a) experimental testbeds and infrastructure for linking federating sensor networks to ORNL simulation centers (b) challenge problems with homeland security relevance.

Internal collaborations in the TuFNet project include Vanderbilt, UC Berkeley and Cornell.

Activities: Our work focused on identification of challenges, establishing collaborative relationship with ORNL, planning and building a real-life demonstration system as an initial experimental platform.

Dirty Bomb Detection and Localization system was designed and implemented by Vanderbilt and ORNL researchers and demonstrated in the Vanderbilt Stadium on April 20, 2006. The actual demonstration is part of the official program of the Information Processing in Sensor Networks (IPSN 06) conference to be held in Nashville, TN on April 19-21, 2006. This integrated demonstration showcases important technologies and potential homeland security applications of sensor networks:

- highly accurate positioning and tracking (ISIS-VU),
- sophisticated radiation detection capabilities (ORNL),
- secure sensor network architecture (ISIS-VU / TRUST),
- early example of the application of federated sensor networks (ISIS-VU and ORNL),
- highly accurate fine grained camera control (ORNL),
- modular micro-operating system for sensor networks (UC Berkeley),
- low-power mote design (UC Berkeley, Crossbow, OSU)

Demonstration web-site is: <http://www.isis.vanderbilt.edu/projects/rips/>

Frank Denap (ORNL), Shankar Sastry (UC Berkeley), Janos Sztipanovits (VU-ISIS) and Steve Wicker (Cornell) co-chair a research workshop on Trusted Federated Sensor Networks to be held in Nashville, TN, on April 21, 2006

Web Authentication and Identity Theft

Societal Context: Internet use has become extremely widespread, ranging from entertainment, to casual browsing, to information gathering, to Internet commerce of different sorts. In interactions between individuals with web browsers and commercial sites that manage money or process web-initiated transactions, the user must authenticate herself or himself to the commercial site. Although many authentication methods have been developed by researchers over past decades, the vast majority of web sites accessible to ordinary individuals use password-based authentication. Unfortunately, passwords are subject to various forms of subversion, including phishing attacks, keystroke logging, and related methods. These attacks have been carried out on a huge scale in recent years, with law enforcement and large private companies estimating annual losses over \$1 billion. In addition, password theft and identity theft have pierced the public consciousness.

Since current attacks involve tricking human users by presenting replicas of trusted interfaces,

Team for Research in Ubiquitous Secure Technology
2005-2006 Annual Report

there are substantial social science issues involved, including legal issues related to the responsibility of financial and other institutions that use web authentication, and human factors questions about how users are fooled into entering sensitive data into malicious web sites, or fooled into installing spyware that then carries out malicious activity on the users computing platform. Because this problem is highly visible and affects many users beyond the research community, this area presents an excellent opportunity for outreach, education, and technology transfer. We believe that practical methods, delivered directly to concerned individuals or concerned enterprises, will have broad societal impact and will reflect positively on the TRUST center and its sponsors.

Integrative Testbed: This collaborative TRUST project, involving faculty and students from computer science and law, will examine the social and legal context of identity theft, develop improved technology to combat phishing, spyware, botnets, and related threats, pursue technology transfer opportunities, and study the policy and legal implications of intrusive activities and possible defensive measures. Participation to date has come from computer science departments and law schools at Berkeley, CMU, and Stanford; additional participation from other universities or TRUST industrial partners will be welcomed. The identity theft thrust has four primary objectives.

- Understand how users perceive their vulnerability to identity theft attacks and how well they understand the privacy threats associated with installed software.
- Develop mechanisms for detecting potential and actual loss of personal data from computers.
- Build and demonstrate active systems that prevent identity theft.
- For each of our proposed mechanisms for detecting identity theft and preventing identity theft, understand the policy implications and legal implications

Research Collaboration: Berkeley participants Tygar and Dhamija have developed dynamic skins anti-phishing technology, have performed user studies of anti-phishing methods, and have analyzed other identity theft techniques such as acoustic emanations. Berkeley Law member Deidre Mulligan has experience with legal issues related to identity theft, as does Stanford Law member Jennifer Granick, co-leader of a recent study of spyware technology and legal issues at Stanford. CMU faculty Perrig and Song and their students have worked on botnet detection and enhanced web authentication methods. Stanford faculty Boneh and Mitchell, with their students, have developed a series of software browser extensions that combat identity theft and collaborated with Granick on spyware study. Stanford professor Rosenblum designed virtualization methods that are central to the planned SpyBlock effort.

External collaborators in the ID Theft project include:

- DHS/SRI Identity Theft Technology Council, a group that includes representatives from financial service companies, leading auction sites, and related organizations.
- RSA Securities has engaged with Stanford, under DHS sponsorship, to transition a password hashing method to their commercial identity management product.
- Google has hosted a Stanford intern who has helped incorporate Stanford SpoofGuard technology into a Google toolbar release in spring, 2006.

Activities: Our accomplishments to date are decentralized, although collaborative discussion began at the June TRUST kickoff meeting and has continued through the 2006-2007 planning process. We plan to compare methods from different campuses in our evaluative studies, and

integrate compatible methods in future software distributions.

- SpoofGuard is a browser extension designed to help prevent phishing by detecting attacks in progress. The extension, which predates the TRUST center, is freely available at <http://crypto.stanford.edu/SpoofGuard/>. Integration into a Google toolbar was completed in the past year.
- Dynamic skins automatically customize secure windows and provide visual cues that help prevent phishing attacks.
- A “phoolproof phishing prevention” effort produced an authentication protocol that leverages Bluetooth-enabled cell phones to circumvent malicious code on a client station.
- PwdHash, developed in part before TRUST funding was received, produces a custom passwords for each site, in a manner that combats phishing and other password attacks. The software is freely available at <http://crypto.stanford.edu/PwdHash/>
- Initial progress on botnet detection has been completed, with planned conference submission for summer 2006.
- A spyware project studied commonly deployed spyware and developed technical foundation for legal action, in collaboration with Stanford Law School.
- A semantics-based malware detection method has been developed and tested (at CMU).
- Education and outreach efforts include course modules on identity theft (to be completed spring 2006) and interaction with law enforcement (Secret Service, FBI Infragard) and private companies concerned with identity theft (e.g. PassMark Security).

2.3.4. Thrust Area IV: Trusted Platforms and Trustworthy Systems

Carnegie-Mellon University has developed a system called Bump-in-the-Ether (BitE), an approach for preventing user-space malware from accessing sensitive user input and providing the user with additional confidence that her input is being delivered to the expected application. Rather than preventing malware from running or detecting already-running malware, we facilitate user input that bypasses common avenues of attack. User input traverses a “trusted tunnel” from the input device to the application. This trusted tunnel is implemented using a trusted mobile device working in tandem with a host platform capable of attesting to its current software state.

Based on a received attestation, the mobile device verifies the integrity of the host platform and application, provides a trusted display through which the user selects the application to which her inputs should be directed, and encrypts those inputs so that only the expected application can decrypt them.

A paper on this work has been accepted to the 2006 USENIX Annual Technical Conference. Quorum systems underlie numerous approaches for implementing intrusion-tolerant distributed services. A quorum system over a universe of logical elements is a collection of subsets (quorums) of elements, any two of which intersect. In implementations of intrusion-tolerant distributed services, the elements of the universe reside on the nodes of a physical network and the participants access the system by contacting every element in some quorum.

We have initiated a research program to study the network-centric costs that these quorum accesses induce. Specifically, this year we studied algorithms to place universe elements on

the nodes of a physical network so as to minimize the network congestion that results from quorum accesses, while also ensuring that no physical node is overloaded by access requests from clients. We considered two models, one in which communication routes can be chosen arbitrarily and one in which they are fixed in advance. We showed that in either model, the optimal congestion (with respect to the load constraints) cannot be approximated to any factor (unless $P = NP$). However, we showed that at most doubling the load on nodes allows us to achieve a congestion that is close to this optimal value. We also provided initial steps to elucidate the extent to which element migration can reduce congestion in this context.

A paper on this work has been accepted to the 2006 ACM Symposium on Principles of Distributed Computing.

2.3.5. Thrust Area V: Secure Sensor Networks

Introduction

The TRUST Secure Sensor Networks initiative is focused on the development and use of secure embedded sensor networks in a variety of large-scale applications. Applications include the protection and monitoring of critical infrastructure, rapid response systems for homeland defense, and the remote monitoring of individuals for clinical purposes, whether living at home or in group facilities. Recent developments in the field of sensor and networking technology have made such networks possible; our initiative drives the further development of the requisite deployment, network configuration, data recovery, and security technologies, while continuing to develop the theoretical foundations for this field. The TRUST Secure Sensor Networks initiative also considers the privacy issues arising from the use of sensor networks, and the ways in which embedded sensor networks affect the experience and use of public spaces. The economic issues surrounding the development of markets in sensed data will be considered in the coming year. A significant educational and outreach component is also being developed with the joint objective of increased diversity in the ongoing development of these technologies and an increase in public awareness of the surrounding technical, legal, economic, and social issues.

The first year of effort in the Secure Sensor Networks initiative included collaborative activity between Cornell, Vanderbilt, and Berkeley. At Cornell, the primary research focus lay in the development of algorithms and technologies that support the deployment of very large arrays of randomly distributed sensors in the critical infrastructure. We first considered the boundaries of what was possible, developing an information-theoretic approach to sensor allocation, as well as a theory for trading off network robustness vs. efficiency [10], [16]. A foundation for scalable distributed MAC protocols was then developed using the tools of game theory [13], [14], [15]. The problem of localization of randomly distributed sensors was also considered, resulting in the development of a lightweight, anchor-free algorithm based on a small number of simple beacons [17], [18]. In the coming year, we will work with Vanderbilt in applying these algorithms and technologies to large “federations” of sensor networks. We have also begun a collaborative effort with the Vanderbilt Medical School to develop medical sensor networks that use some of these ideas. We have also been very actively involved in considering privacy issues, an area of great interest to scholars at the Berkeley School of Law, as will be discussed below.

Vanderbilt’s research on secure network embedded systems has focused on integrated security analysis and support for sensor networks and sensor-based federated networking systems. While much research on secure sensor networks has been focused within the sensor network itself (e.g. key distribution mechanisms for sensor networks [1][2][3][4][5][6] and link-level security architecture for sensor networks [7]), we note that the sensor network marks not the

end, but the beginning of the journey of the sensed content. Sensor networks are now at the crucial stage of evolving from an isolated system to an integral component of the global information infrastructure, which may compose of WiFi wireless LAN, broadband wireless network, global Internet, and cellular network etc. We call such a sensor-based information infrastructure a sensor network federation. Based on this observation, our work has targeted not only secure transport for each sensor network, but more importantly, establishing a *trusted sensor network federation* (TuFNet).

One of Berkeley's primary contributions at the outset lay in the social sciences. Collaboration with Pamela Samuelson and Deirdre Mulligan at the Berkeley Law School has opened up a wide range of research topics that focus on the interface between new surveillance technologies and the public's expectations of privacy and its use and perception of public spaces. Two workshops were held and a joint research initiative, described below, is now underway. Expected deliverables in the coming year include a privacy "roadmap" that correlates new technical developments with effective regulations for maintaining the public interest in privacy.

In what follows, we discuss several focus areas that were developed in the first year of the TRUST center.

Mutual Authentication in Sensor Networks

This part of the TRUST secure sensor networks initiative is based on a real-world sensor network application: "Dirty Bomb Detection and Localization", built jointly by Vanderbilt-ISIS and Oak Ridge National Laboratory (ORNL) as the initial version of a TRUST Integrative Project. Based on systematic analysis of its security vulnerabilities, we have focused on one particular security issue in this system – the lack of mutual authentication among sensors. Towards this problem, our contributions are two-fold. Theoretically, we have developed a novel light-weighted symmetric-key-set-based authentication mechanism for sensor networks. Practically, we have integrated our security solution into the "Dirty Bomb Detection and Localization" demonstration and validated the feasibility of such security solution based on real field system testing.

The operation of "Dirty Bomb Detection and Localization" system relies on the localization initiation from the right sensor node. In the system design, this sensor (called master node) upon detecting the bomb initiates the localization procedure by sending out a localization message to the rest of sensors in the network. Without appropriate authentication mechanism, an attacker could forge such a message, interrupting the system operation and wasting the scarce sensor network energy resource. Thus the master node needs to prove its identity to the rest of the sensor nodes. In a traditional networking environment, such authentication could be achieved by applying public key cryptography. In the setting of our sensor networking system, existing public key algorithms consume too much CPU cycles and memory spaces to be practical. To address this issue, we propose a key-set-based authentication mechanism that utilizes symmetric key cryptography, thus placing a relatively limited computational burden on the platform. In this mechanism, each sensor in the network is assigned a unique subset of symmetric keys from a key pool. The key assignment is computed in such a way that (1) two different nodes will only share a portion of common keys, and (2) for any node, its key set will be fully covered by the union of several other nodes' key sets. To authenticate a message from the master node, the message authentication code is first computed by the master node based on its subset of the keys. Upon receiving the message, each sensor node could provide a partial verification of the message. The whole message will be authenticated by the collective effort of the sensor nodes whose key set union covers the key set of the master node. We implement this authentication mechanism as a component on TinyOS, and integrated it with the "Dirty Bomb Detection and Localization" system. The real system measurement shows that our

solution is feasible and efficient.

Digital Rights Management for Sensor Network Federation

For sensor network federation, we have developed a digital right management framework to protect the sensor data [8]. Sensor network federation enables sensor data to be a public information source on the Internet, many urgent content protection issues arise, mainly due to the sensitivity and the privacy natures associated with the sensor content.

This motivates the crucial necessity of enforcing DRM (Digital Rights Management) for the sensor data in its networking federation. The most salient advantages of DRM of sensor network federation are as follows: (A) Protecting Privacy [9], where DRM can effectively isolate its protected data from unauthorized access. For example, in patient monitoring, the access to any footage including the patient's appearance should be absolutely limited to his/her care-takers and family members; (B) Promoting Economical Incentives, where DRM provides a set of solutions for the trading, accounting, and transaction processing of sensor content as commodities, since many types of sensor content hold significant commercial values. (C) Protecting Ownership, where DRM assists to identify and attest the abuse of access right and illegal distribution of the sensor content, hence protect ownership and customer rights.

We focus on four major challenges to apply DRM in sensor network federation: (1) DRM for the composed and co-dependent multi-sensor content, (2) differentiated privacy and security for different parts of the composed multi-sensor content, (3) key management for secure content creation and aggregation in multi-sensor network environment, and (4) DRM-aware access control and composed multi-sensor content delivery to legal customers. *Theoretically*, we propose a DRM-enabled service architecture for sensor network federation and its DRM-based algorithms and protocols. *Practically*, we deploy a video sensor network as the enabling case study of our solution to validate the novel DRM architecture for sensor network federation.

Game Theoretic MAC Protocol Design

In this work, we addressed the fundamental question of whether or not there exist stable operating points in a network in which selfish nodes share a common channel, and consider how the nodes behave at these stable operating points. We begin with a wireless communication network in which n nodes (agents), which might have different perceived utilities, contend for access on a common, wireless communication channel. We characterize this distributed multiple access problem in terms of a oneshot random access game, and then analyze the behavior of the nodes using the tools of game theory.

We have completely characterized the Nash equilibria of this game for all $n \geq 2$. We have also shown that all centrally controlled optimal solutions are a subset of this game theoretic solution and almost all (w.r.t Lebesgue measure) transmission probability assignments chosen by a central authority are supported by the game theoretic solution.

After establishing the Nash equilibria of our one-shot random access game and analyzing the behavior of the nodes at Nash equilibria, we pursued an asymptotic analysis of the system as the number of selfish players goes to infinity. By means of the asymptotic analysis, we study the behavior of the dense wireless networks containing selfish nodes. When all nodes are identical, we can give the best possible convergence bounds on the asymptotic distribution of the packet arrivals and asymptotic channel throughput.

After the asymptotic analysis of the homogenous case, we extend our results to the heterogenous case in which nodes may have different costs. We first showed that asymptotic analysis of the heterogenous case reduces to asymptotic analysis of the homogeneous case.

This means qualitatively that if there exists a Nash equilibria at which number of nodes contending for the access of the common wireless channel goes to infinity, then all of the transmitters expect for finitely many of them have cost of unsuccessful transmission lying in an arbitrarily small region, and our results in the previous section characterize the behavior of these infinitely many transmitters. We finally put this intuitive idea into a formal proof showing that the total number of packet arrivals can be approximated by a sum of a Poisson random variable and finitely many Bernoulli random variables up to an error term.

The practical importance of this work lies in its demonstration that scalable, distributed MAC protocols can provide the same performance as the standard centrally-controlled network. The results have been published (or are under review) in [13], [14], and [15].

Localization of Randomly Distributed Sensors

Work is ongoing at Cornell on localization and location-based routing for wireless sensor networks. We considered a wireless sensor network consisting of a single fusion point at the center of a field of randomly distributed sensors. An anchor-free node localization algorithm was developed in which the sink node imparts radial location information through the phased-array transmission of a series of beacons, and the individual sensors use knowledge of received beacons as well as information from nearest neighbors to identify the sub-sector in which they reside [17]. A routing algorithm was developed which uses the localization results in the selection of relay nodes to reduce energy consumption as well as to extend network lifetime [18]. We showed that our localization algorithm works well for high-density networks and is robust against measurement noise. We also show that the combined localization and location-based routing algorithms make power-efficient routing decisions at low cost.

Modeling Network Lifetime

A challenge in the design of sensor networks is choosing simple operating strategies that attempt to minimize energy consumption in the network and maximize the lifetime of the sensor nodes. The lifetime of a sensor node is a non-deterministic function of the rate at which energy is consumed by communication and processing events. Therefore, it is highly useful to develop a stochastic model for energy consumption and node lifetime based on the behavior of the node and the network that can subsequently be used to optimize node behavior. We propose an applied probability model to predict energy consumption and lifetime of a sensor node based on a renewal-reward process. There exists a law of large numbers and a central limit result for this process, which we show provide good and easy to use approximations for the the expected value and the distribution of node lifetime. The advantage of this model is that it only requires the calculation of first and second moments of the energy and time variables, which enables the model to take into account many energy consuming factors without making the calculations intractable. We apply this model to a specific scenario to calculate the optimal active ratio amongst a set of sensor nodes with sleep capabilities, based on the traffic rates and behavior of the nodes, and derive a logarithmic relationship between the optimal active ratio and the traffic load.

We have also developed models for network lifetime based order statistics, which use the distribution of sensor node lifetime to calculate the distribution of the time of the critical node failure in the network. By being able to carefully predict node and network lifetime, our models serve to determine appropriate deployment and replacement strategies for wireless sensor networks. Several publications are in process.

Network Security

Team for Research in Ubiquitous Secure Technology
2005-2006 Annual Report

In related work based at Berkeley, active insider attacks are being addressed through the formulation of the problem as a game between the sensor network and the adversary. Using tools from game theory, we are analyzing the attacks by the adversary and will minimize the security breaches.

A central authority (i.e. base station) keeps track of a state variable for each node in the network. A decision rule is used by the central authority to assign each node a probability of trustworthiness (the "reputation"). The decision rule is based on the fact that local neighborhoods of nodes transmit highly correlated data. The reputation tracks the degree to which each node sends information that agrees with or disagrees with the information sent by adjacent nodes. If there is a substantial or prolonged disagreement by a node, it is rendered suspect and given a low reputation (untrustworthy).

The strategy of the adversary is to maximally corrupt the sensor network. This includes how to corrupt the data that is transmitted as well as how many nodes to compromise. The objective of studying this game is to know in what circumstances the central authority can find a state variable and decision rule that deters attacks from adversary.

In order to deter the adversary, the central authority must collect redundant information so that a reputation value can be constructed as a state variable for a large collection of randomly chosen nodes. This creates transmission costs. On the other hand, the adversary could learn the central authority's decision rule, and corrupt more nodes if necessary. At some point compromising more nodes becomes too costly for the adversary, and it will be deterred. The point is to characterize the cost considerations that will make it feasible to deter the adversary in this way.

In the coming year the theoretical and practical issues presented by this game-theoretic approach to security will be considered in detail. Game theorists at Berkeley and Cornell will work together to characterize the Nash equilibria of this game, thus relating performance to network parameters in a systematic form that supports algorithm design.

Privacy and Social Science Issues

We are developing a map of privacy rights and expectations and current legal, technological, and social methods for protecting and managing them. This involves boundaries defined through legal precedents as well as considerations of evolving social awareness. We are considering how, in general, the advance of sensing/surveillance, as driven by application requirements, creates novel privacy concerns. When are such concerns outweighed, if ever, by needs being met by the surveillance technology? We will continue to explore the way in which sensor networks may influence the experience and use of public places. In particular we will consider whether sensor networks will change the capacity of public spaces to support the traditional functions for which they have been valued. Given the nexus between a public spaces appeal to terrorists or other attackers and the symbolic and functional value of a public space, the effect of embedding sensor networks deserves considerable attention. This work builds off of current efforts to identify the privacy and security issues unique to sensor networks.

Specific issues to be considered in the development of a privacy roadmap include a mapping of sensor capabilities and network mission into deployment and data use rules.

The proposed map will expand privacy "requirements" to incorporate Constitutional roots of privacy and then consider options for extending this broader understanding of privacy along with fair information practices to cover sensor networks through technical, regulatory and legal

interventions.

As a possible means for providing notice, a component of privacy protection, we propose to study user interfaces that enables people to query a sensor network about the personally identifiable information collected. Giving people this information will heighten awareness of data collection and may help reduce the concerns of privacy-sensitive individuals in some sensor networks. We will also study public policy aspects of this problem to figure out whether a viable standardization approach exists that industry can adopt, or whether other approaches such as legislation are needed to achieve privacy-preserving sensor network environments.

Two testbeds are under development for use in the Sensor Networking/Privacy initiative. The first is being deployed in the Herbert F. Johnson Art Museum on Cornell's campus, and is designed to provide near real-time information on the movement and presence of museum patrons. This is a joint project with Cornell's Human Computer Interaction Group (HCI). The floor space will be divided into a number of "zones" for which the system will be continuously estimating the number of persons present, although specific individuals will not be identified by the system. Information about the rate at which individuals move throughout the floor space will also be reported. Data from the system may also be used by museum curators to estimate the popularity of specific art pieces and the effectiveness of the museum layout. The system uses Crossbow Technology's MicaZ wireless sensor motes, for which we have developed software to achieve these goals. We are also implementing software which aggregates and processes data from the mote network to provide the types of data mentioned above. This testbed provides a real-time means for exploring the public perception of privacy in public spaces.

The second project is a medical wireless sensor network designed for continuous monitoring of congestive heart failure patients, allowing them to be kept at the minimum level of care. This is a joint project through TRUST with the Vanderbilt Center for Better Health and the Samuelson Law Clinic at Berkeley. Due to the sensitive nature of medical data, security and privacy are critical concerns of the system and will be addressed at all levels in the system. The system will interface with the Vanderbilt Center for Better Health's web-based patient portal, which provides doctors, patients, and other members of a patient's care network with varying degrees of access to different types of medical data. Data from the patient's sensor network will be collected and processed by a PC in the home. This PC will arbitrate access to the collected data according to the roles of various users requesting data through the patient portal.

As an initial effort to create a cross-university research community, and to familiarize the social scientists with technical issues and vice versa, we held two workshops at Cornell in the past year.

Sensor Networks and Privacy, Cornell and Berkeley - Tuesday, March 28, 2006. There were ~30 attendees over the course of the day. The focus of the workshop was the state of the art of sensor networking and the privacy concerns concomitant with their deployment in public spaces. Engineers, computer scientists, social scientists and legal scholars discussed the need and potential for developing public policies that are deeply coupled to the advancing technologies.

Sensor Networking Workshop, Cornell and New York Department of Health - Tuesday, October 11, 2005. There were ~50 attendees over the course of the day. This workshop focused on the potential uses for sensor networks as a means for protecting critical infrastructure. Attendees included electrical and civil engineers, computer scientists, molecular biologists, psychologists, and chemists who discussed the potential for developing bio-toxin specific sensor networks to

protect water supplies. Attendees also discussed potential uses for sensor networks in public spaces as a means for understanding how the public uses public spaces, and in particular how they interact with publicly displayed pieces of art. The latter was intended to illustrate the extremely broad range of uses for sensor networks in a research context.

2.3.6. Thrust Area VI: Network Defense

The bulk of the work done this year in the area of network defense uses the testbed DETER supported on a related NSF project. The Cyber Defense Technology Experimental Research network (DETER network) provides necessary infrastructure- networks, tools, and supporting processes-to support national-scale experimentation for research and advanced development on emerging security technologies, without disturbing the production Internet. The DETER project has achieved the goal of creating, operating, and supporting a researcher- and vendor-neutral experimental infrastructure that is open to a user community including academia, government, and industry. An operational capability was achieved within the first six months of funding by deploying two geographically distributed clusters, one at the University of California, Berkeley, (U.C. Berkeley) and another at the University of Southern California's Information Sciences Institute (USC/ISI), interconnected by the California High Performance Research Network.

The DETER network project's on-going successes include continual expansion of hardware and software capabilities from the initial deployments, implementation of software changes that have reduced network hardware costs by an order of magnitude, security enhancements, and support for an expanding user community. More than 30 user organizations (100 users) from government, academia, and industry (including several start-up companies with emerging technology) have been approved to use the testbed for their security experiments. The experiments span the range of cyber defense technologies in the areas of Distributed Denial of Service attacks (DDoS), worms/viruses, and routing protocols, and they have yielded numerous published papers in refereed journals. The testbed was used in support of a Cyber Storm exercise in spring 2006, providing both emulation of a simulated attack and near-real time graphical visualization and analysis of the data generated from the simulated attack to the exercise participants. The testbed was also used in a new undergraduate security course at U.C. Berkeley in fall of 2005. The project has conducted three large-scale experiment demonstrations (June 2004, October 2004, and September 2005) to audiences composed of government, academia and industry members. A final large-scale experiment demonstration is planned for June 2006.

DETER network is more than a passive research instrument; it also serves as a center for interchange and collaboration among security researchers. In particular, it supports research in the areas of information system infrastructure, including networking technologies and the analysis of architectures that support secure interaction between humans and information systems, as well as secure networking, network attack recognition and identification, and the ability to continue operations in the presence of a successful attack. More broadly, the DETER network is a shared laboratory in which researchers, developers, and operators can experiment with potential cyber-security technologies under realistic conditions, with the aim of accelerating research, development, technology transition, reference prototype implementation, and deployment of effective defenses for U.S.-based computer networks.

The DETER testbed is an experimental facility to safely support a broad range of cyber security research projects, including those experiments that involve malicious or "risky" code. Most of these experiments could not be performed in the real Internet, because of

Team for Research in Ubiquitous Secure Technology
2005-2006 Annual Report

the type and volume of traffic and the risk of escape. Development and operation of the DETER testbed is currently funded by the National Science Foundation and the Department of Homeland Security under awards # ANI-0335298 (DETER) [September 2003 - August 2006] and CNS-0454381 (DECCOR) [August 2005 - September 2007]. The DETER grant has funded testbed design, testbed control software development, the basic testbed hardware infrastructure, and testbed operation. The early successes under this effort and requests from the sponsors have led to an expanded hardware deployment under the DECCOR grant that has provided additional infrastructure to better meet the over DETER project goals and serve more users. The project is transitioning from its current R&D phase to an Operational and Maintenance (O&M) capability that begins in GFY 2007, and provides options for funding of the O&M functions beyond the transition. The parallel NSF/DHS project EMIST (award # ANI-0335241) funds a small group of cyber-security researchers in seven organizations to collaborate with DETER and to use the DETER testbed. The result is a rich synergy between experimenters and testbed builders.

In brief, the DETER testbed is composed of two linked clusters of experimental nodes, controlled by a version of Utah's Emulab [White02] software. Experimenters can use these nodes in arbitrary combinations as network traffic sources, routers, traffic shapers, and sinks. The Emulab software provides sharing of testbed resources among multiple concurrent experiments when enough nodes are available. However, the current sizes of the clusters are a serious limitation to such sharing. In addition, the current operations staff is limited in number, limiting the types and amount of user assistance and support that can be provided. The proposed transition effort will thus provide enhanced support for a greater number of users – security researchers and advanced technology developers - from an expanded user community.

The testbed provides pre-built support for the following operating system software on experiment nodes: RedHat Linux 7.3, FreeBSD 4.9 and Microsoft Windows XP. However, users can load arbitrary operating system code on to experiment nodes, as users are given full "root" access to their allocated experiment nodes. The testbed uses a secure process to replace the operating system on each experiment node after each experiment finishes. We are developing and optional disk scrub operation that would be automatically performed after experiments that use live malware or are sensitive (i.e., they involve confidential data or applications). To isolate the experiment nodes, there is no direct network path into experimental network from the Internet. Instead users can set up encrypted tunnels across Internet using SSL, SSH, and IPSec tunnels.

On TRUST the major use of the testbed has been for the following sets of experiments

Detecting viral/worm e-mail

Protecting against computer worms and viruses is an active research area; however, the most commonly deployed anti-virus defenses based on signature-matching have not changed significantly over time. While these methods generally benefit from high accuracy, there is significant lag time between the release of a new threat and public availability of signatures for that threat, a gap which malware writers exploit to infect machines. In 2004, an average of over 30 new Windows worm and virus variants was documented each day. However, the average lag time among major anti-virus vendors for the release of a new signature was several hours. In several cases, signatures were not available for over a day from when the threat was first released. Similar response times were observed for some of the most damaging e-mail worms seen recently (e.g., major vendors took more than 8 hours to release signature files for MyDoom.A, the worst worm in the first half of 2004). The delay for new signatures depends upon the time it takes for a vendor to: receive a malware sample, create and test a signature,

and distribute it to end users. Our approach augments this process to eliminate the delays. Furthermore, most vendors distribute new signatures between 1 and 10 times per week.

Many e-mail servers and clients allow users to define filtering rules for e-mail attachments that can be used to block them from being downloaded and opened. However, it is difficult to create an exhaustive list of potentially harmful attachment types a priori. For example, seemingly harmless JPEG image format files can be used to exploit vulnerabilities in certain image renderers. Similarly, ZIP files are used for exchanging information, but some worms use password protected ZIP files to propagate. Moreover, certain worms do not rely on attachments to propagate at all. Instead, they use embedded scripts that auto-execute on e-mail clients, or entice users to click on hyperlinks which in turn exploit browser vulnerabilities.

Thus, traditional signature-based methods or e-mail filters are insufficient for containing the spread of novel e-mail worms. The key to stopping worm outbreaks is to quickly, adaptively learn user and viral behavior, and quarantine users that exhibit viral behavior. Our approach is based on four key ideas:

1. Learning users' outgoing e-mail behavior:

Our system evaluates several features on outgoing e-mail to develop viral and normal e-mail behavioral models that can be used to counter viral activity. Incoming e-mail consists of messages from many distinct users (including spam and worm traffic), and by itself can be ineffective at profiling individual user behavior. However, outgoing e-mail consists solely of messages sent by a single user, and can hence be far more effective in profiling individual user behavior and effectively discriminating between normal and viral e-mail behavior. Users whose outgoing e-mail exhibits significant deviations from their normal behavior (indicating viral contamination) can be quarantined.

2. Using a multi-tiered approach:

Previous work using statistical learning to detect worm propagation has often suffered from excessive amounts of normal e-mail being classified as viral (i.e., false positives). Because classifying viral e-mails as normal (i.e., false negatives) is a highly undesirable occurrence, most systems are intentionally configured to be overly sensitive (i.e., eliminating most false negatives at the cost of excessive false positives). Instead, we use a sensitive novelty detector trained only on clean e-mail to isolate most normal e-mail. We then apply a second layer classifier that learns on viral and clean e-mail, and filters most false positives. Our multi-tiered approach achieves close to 99% accuracy making it more effective than using our novelty detection or classification alone.

3. Leveraging existing solutions to improve results:

Our system has been designed to work in concert with existing solutions, specifically signature-based scanners, which are highly accurate once signatures have been propagated (i.e., once the "window of vulnerability" has passed). Since many worms remain prevalent after an initial outbreak, signature based systems can be effectively leveraged outside this window to prevent recurring infections. In addition, as an aid to automatic retraining of statistical models for adaptiveness to changing user/virus behaviors, scanners can be trusted to accurately label old e-mails as clean or viral. This results in a partially-labeled corpus, which we use to implement "semi-supervised learning."

4. Providing a containment-based approach:

Preventive solutions (e.g., virus scanners and e-mail filters) try to reduce the size of the

vulnerable population to limit the spread of a worm outbreak. However, newly identified vulnerabilities, unprotected users, and increasingly sophisticated worms continue to cause outbreaks, so we focus instead on quick detection and containment. There are three key reasons why this approach can be more effective than prevention-based techniques. First, the containment process can be completely automated since detecting and characterizing a worm can be far easier than understanding the worm itself or the vulnerability being exploited. Note that containment is more intrusive than preventive approaches, since all outbound e-mail is blocked, not just an incoming message. Second, since containment can be deployed in the network, it is possible to implement a solution without requiring universal host-based deployment. Finally, even though containment does not prevent every host from being infected, it can significantly slow down a worm's propagation rate and buy the crucial time needed to bring signature-based systems up-to-date. Simulated experiments with our containment strategy demonstrate its effectiveness, even in limited deployment scenarios.

A feature is a statistic measuring some aspect of a user's e-mail activity or behavior. While there is a large number of potential features that could be analyzed, we focus on those that can reveal the abnormal sending behavior caused by a worm infection or active spamming activity.

We implemented more than two dozen separate features with the underlying goal of obtaining a set of statistics that accurately distinguishes between normal and abnormal e-mail activity. Each feature returns either a continuous value or multiple binary values (multinomial). For example, a frequency calculation returns a number, whereas a feature involving types of e-mail attachments is represented as an array of bits, where each bit represents the presence/absence of a specific type of attachment.

Our features consist of those calculated on a single e-mail (i.e., single points in ongoing e-mail activity, such as the attachment type) and those that examine several e-mails over a fixed interval of time (i.e., trends in message characteristics, such as a running average of the number of characters in a single user's e-mail subject lines).

The features we use include:

- Whether the message is a reply or forward
- Presence of HTML
- Presence of certain HTML script tags/attributes
- Presence of remote/embedded images
- Presence of hyperlinks
- MIME types of file attachments
- Presence of binary, text attachments
- UNIX "magic number" of file attachments
- Total size of e-mail, including attachments
- Total size of files attached to the e-mail
- Number of files attached to the e-mail
- Number of words/characters in the subject and body
- Number of e-mails sent
- Number of unique e-mail recipients
- Number of unique sender addresses
- Average number of words/characters per subject, body; average word length
- Variance in number of words/characters per subject, body; variance in word length
- Ratio of e-mails with attachments

We performed analysis using the Enron dataset, a large, publicly available collection of messages composed of the internal e-mail from Enron Corp. This corpus contains 126,078 internal e-mails without attachments sent by 148 distinct users, each of whom sent between 3 and 8,926 e-mails. Most users sent under 1,000 e-mails.

Through empirical analysis of the Enron dataset, we observed that users could be grouped into common clusters enabling sets of users to be largely represented by a single behavioral model. By using our features to build a model of normal behavior for a given user, we can detect deviant behavior, such as worm propagation, with reasonable accuracy. This process is known as "novelty detection."

There are many different novelty detection techniques that could be used to detect worm propagation attempts, however our current approach uses a one-class Support Vector Machine (SVM). A one-class SVM applies a linear algorithm that attempts to maximally separate the "normal" data from the origin via a hyperplane boundary. This technique's properties enable it to be transformed into a non-linear algorithm by application of a similarity measure known as a kernel.

We employ a commonly used one-class SVM with a Gaussian kernel. Our one-class SVM was trained to allow only a small fraction, 0.1%, of outliers during training. Increasing the detector's sensitivity makes a point more likely to be classified as infected. Unfortunately, a significant drawback of using novelty detection is the difficulty in selecting a sensitivity that is sufficiently high enough to yield a low false negative rate, while also yielding a low false positive rate. Too many false negatives would allow an outbreak to propagate, while too many false positives could result in alerts being ignored by human supervisors, rendering the system useless. This selection problem is a common one for adaptive intrusion detection systems.

We address this issue by introducing a second layer to our overall model. The strength of sensitive novelty detection is its low false negative rate. This is a very desirable goal in worm infection detection, as each false negative could potentially cause another infection. Therefore, we train another model to reclassify the messages already flagged as abnormal by the SVM. The additional classification step reduces the false positives by attempting to filter out the actual worm messages from the output of the novelty detector.

While many different models could be employed in this role, our current approach uses a Naive Bayes parametric classifier as the second model. We chose this particular model for several reasons. First of all, it is a simple model that is very quick to train, both in scenarios where the data is completely labeled, such as initial training, and when there is data with undetermined labels present, such as during online operation. Second, parametric methods in general can perform useful classifications after having seen fewer data points than other methods. Finally, Naive Bayes classifiers leverage differences between feature distributions of e-mail data from separate users. Naive Bayes models exploit this disparity by maintaining separate parameter sets for feature distributions of virus and non-virus classes. These features are then used in the calculation of the posterior probability for the classes, which determine the final classification decision.

The models used in our system are initially trained separately. The novelty detector is trained only on known clean e-mail sampled from all users in our system, while the parametric classifiers learn on clean as well as a sample of viral e-mail. After their initial training, the models are continuously retrained to capture any changes in user and viral e-mail behavior that occurs over time.

Initial training of the one-class SVM entails learning of a hyperplane in high-dimensional space such that it creates a boundary around normal data in feature space and is highly sensitive to anomalous data. Initial training of the Naive Bayes classifiers involves a maximum likelihood estimation (MLE) on the training data to estimate viral and non-viral distribution parameters for each feature, which in our system, can simply be obtained by calculating the sample means and variances for the viral and non-viral training data distributions belonging to each feature.

For initial training, we mine clean e-mail data from existing users' 'Sent' mail folders. We address some privacy concerns with a tool that users can run to extract feature vectors from their e-mail. For viral e-mail, we use preset distribution parameters calculated from collected viral samples. Subsequent new users' initial training uses their 'Sent' mail folders if available. Otherwise, their classifiers could be initialized with parameters aggregated from data belonging to other users that share common e-mail sending behavior. We use two sources of feedback for retraining models: a virus scanner and the network administrator.

We leverage existing anti-virus technology by periodically passing all e-mail through a signature-based scanner to label known worms. If the scanner indicates that an e-mail is infected, it is labeled as such. Because there is an interval during which we do not trust the virus scanner (i.e., the time to update it with signatures for novel worms), we only allow the scanner to label an e-mail as clean if the message was sent before that interval. Otherwise, the message is left unlabeled. The virus scanner 'trust' interval is represented as d , and only e-mails labeled as viral by the virus scanner are trusted as such during this period. For messages older than d , all labels from the scanner are trusted. Moreover, once they are scanned at time d , no further virus scanning is necessary and we no longer retain them.

Once an infection is detected, we gain additional feedback by reporting it to an administrator, who either confirms or denies the problem. If an infection is confirmed by the administrator, a small number i randomly chosen messages from the infected user are shown to the administrator, who then confirms whether or not these messages are infected. To address privacy concerns, e-mails might instead only be shown to their original senders for confirmation.

The virus scanner and administrator feedback in the form of e-mail labels is used for periodic retraining of the per-user classifiers. The amount of time between each model retraining is configurable, and can be set in terms of wall-clock time or number of e-mails. The Naive Bayes models take on the order of a few seconds to train with several thousand e-mails, making frequent retraining of per-user classifiers feasible. In practice, we anticipate that retraining once every few days, or whenever a new virus infection is detected. However, frequent retraining of the one-class SVM is infeasible, as it can take significantly longer given that the training time increases roughly quadratically with the number of training points, in the worst case. Fortunately, since the SVM is common to all users, and we do not expect the cumulative e-mail sending behavior of all users to change frequently, retraining for the SVM can be done offline.

For each user, we store a buffer of their unlabeled feature vectors. The feature parser generates and stores the vectors in this buffer, and they are then assigned labels by the virus scanner or administrator. If the e-mail corresponding to the feature vector is identified as viral, the vector is transferred to a shared viral data buffer, that is common to all users. If the vector is labeled as non-viral, it is moved to the clean data queue belonging to the sender of that e-mail. The per-user clean data queues maintain a sliding window (in terms of wall-clock time or number of e-mails) of feature data per user. During retraining, the classifier for each user uses the labeled and unlabeled data for that particular user, along with the shared viral data to calculate the new

feature distribution parameters.

Using a separate buffer for unlabeled data can be useful if a user in the system actually gets infected. In that case, a worm might send enough viral e-mails to overflow the data buffer for that user, resulting in the loss of all known clean e-mail for that user. Hence, using a separate buffer for unlabeled data allows us to preserve known clean e-mail data even in the presence of a worm attack.

Since the data used for retraining is semi-supervised (or partially labeled), the standard maximum likelihood estimation technique used for initial training cannot be applied for retraining. Instead, we make use of the Expectation-Maximization (EM) algorithm during retraining to learn the new feature distribution parameters. Overall, experimental evaluation of our system shows that it accurately detects viral e-mail (low false negative rate), while also accurately classifying normal e-mail (low false positive rate).

References

- [1] W. Du, J. Deng, Y. S. Han, and P. K. Varshney, "A pairwise key predistribution scheme for wireless sensor networks," in *CCS '03: Proceedings of the 10th ACM conference on Computer and communications security*. New York, NY, USA: ACM Press, 2003, pp. 42–51.
- [2] H. Chan, A. Perrig, and D. Song, "Random key predistribution schemes for sensor networks," in *SP '03: Proceedings of the 2003 IEEE Symposium on Security and Privacy*. Washington, DC, USA: IEEE Computer Society, 2003, p. 197.
- [3] J. Lee and D. R. Stinson, "Deterministic key predistribution schemes for distributed sensor networks," in *Selected Areas in Cryptography*, 2004, pp. 294–307.
- [4] L. Zhou, J. Ni, and C. V. Ravishankar, "Efficient key establishment for group-based wireless sensor deployments," in *WiSe '05: Proceedings of the 4th ACM workshop on Wireless security*. New York, NY, USA: ACM Press, 2005, pp. 1–10.
- [5] D. Liu and P. Ning, "Location-based pairwise key establishments for static sensor networks," in *SASN '03: Proceedings of the 1st ACM workshop on Security of ad hoc and sensor networks*. New York, NY, USA: ACM Press, 2003, pp. 72–82.
- [6] L. Eschenauer and V. D. Gligor, "A key-management scheme for distributed sensor networks," in *CCS '02: Proceedings of the 9th ACM conference on Computer and communications security*. New York, NY, USA: ACM Press, 2002, pp. 41–47.
- [7] C. Karlof, N. Sastry, and D. Wagner, "Tinysec: a link layer security architecture for wireless sensor networks," in *SenSys '04: Proceedings of the 2nd international conference on Embedded networked sensor systems*. New York, NY, USA: ACM Press, 2004, pp. 162–175.
- [8] T. Wu, Y. Cui, and Y. Xue "Preserving Privacy in Wireless Mesh Networks", in *Proc. of IEEE WoWMoM*, 2006.
- [9] T. Wu, L. Dai, Y. Cui and Y. Xue, "Digital Right Management for Video Sensor Networks", submitted to *ACM NOSSDAV*, 2006.
- [10] X. Zhang, S. B. Wicker, "On the Optimal Distribution of Sensors in a Random Field", *ACM*

Transactions on Sensor Networks, March 2006.

[11] P. Samar and S.B. Wicker, "On the Behavior of Communication Links in a Multi-Hop Mobile Environment," *Frontiers in Distributed Sensor Networks*, S.S. Iyengar and R.R. Brooks (eds.), CRC Press, 2005.

[12] Xin Zhang, Jun Chen, Stephen B. Wicker, and Toby Berger, "Successive Coding in Multiuser Information Theory," accepted pending revision, *IEEE Trans. on Inform. Theory*.

[13] Inaltekin H., Wicker S. B., A one shot random access game for wireless networks - behavior of nodes at Nash equilibria, submitted to *IEEE/ACM Trans. on Networking*, 2005.

[14] Inaltekin H., Wicker S. B., A one shot random access game for wireless networks - asymptotic behavior of the system, submitted to *IEEE/ACM Trans. on Networking*, 2006.

[15] H. Inalkekin, S. B. Wicker, "A One-shot Random Access Game For Wireless Networks," The International Conference on Wireless Networks, Communications, and Mobile Computing (WirelessCom 2005), June 13-16, 2005, Maui, Hawaii.

[16] X. Zhang, S. B. Wicker "Robustness vs. Efficiency in Sensor Networks," *Information Processing in Sensor Networks* (ACM IPSN 2006), Berkeley, 2005.

[17] Hui Qu and Stephen Wicker, "Anchor-Free Localization for Wireless Sensor Networks", submitted to the IEEE Global Communication Conference, 2006.

[18] Hui Qu and Stephen Wicker, "A Combined Localization and Location-Based Routing Algorithm Design for Wireless Sensor Networks," submitted to ACM SECON, 2006.

[19] Martin, S.; Sewani, A.; Nelson, B.; Chen, K.; Joseph, A. D., Analyzing Behavioral Features for Email Classification, In the Proceedings of the IEEE Second Conference on Email and Anti-Spam (CEAS 2005), July, 2005.

[20] Benzel, T.; Braden, B.; Kim, D.; Neuman, C.; Joseph, A.; Sklower, K.; Ostrenga, R.; Schwab, S., Experience with DETER: A testbed for security research, In the Proceedings of the 2nd International IEEE/Create-Net Conference on Testbeds and Research Infrastructures for the Development of Networks and Communities, Barcelona, Spain, March, 2006.

[21] Barreno, M.; Nelson, B.; Sears, R.; Joseph, A. D.; Tygar, J. D., Can Machine Learning Be Secure? (Invited paper), In the Proceedings of the ACM Symposium on InformAtion, Computer and Communications Security (ASIACCS'06), Taipei, Taiwan, March, 2006.

[22] Barreno, M.; Nelson, B.; Sears, R.; Joseph, A. D., RUMBLE: Rapid User Modeling for Barring Loathsome E-mail, To appear in the Proceedings of the First ACM Workshop on Tackling Computer Systems Problems with Machine Learning Techniques (SysML), Saint-Malo, France, June, 2006.

[23] Nelson, B.; Joseph, A. D., An Analytic Approach to Assessing the Security of

Learning, To appear in the Proceedings of the First ACM Workshop on Tackling Computer Systems Problems with Machine Learning Techniques (SysML), Saint-Malo France, June, 2006.

Thrust Area VII: Privacy and Information Forensics

Our initial year saw significant research in the thrust area of Privacy and Information Forensics. One significant thrust, uniting efforts at multiple institutions, has been the in protecting against identity theft, in particular, through the use of "phishing," to fool users into providing information about financial accounts and other personal details. Efforts between Berkeley, Stanford, and CMU on this thrust have been closely coordinated. Here are some of the highlights:

UC Berkeley:

The first detailed user study of what makes phishing sites work. This research which has been featured on major news media (including a half-hour show on CNN) found that the most successful web sites are able to fool more than 90% of users. By identifying weak points in the ways users identify web sites, this research allows us to develop better tools for defending against web sites.

A framework, Dynamic Security Skins, for protecting against phishing attacks. Dynamic Security Skins have been widely discussed and several commercial firms have begun implementatoin of this structure. Initial tests indicated that Dynamic Security Skins are exceptionally powerful in deterring attacks.

Finally, through our Technology Law Clinic, students have been active in examining issues related to End User License Agreements and possible instillation of spyware

Stanford:

Stanford has developed a set of browser extensions that use powerful cryptographic techniques to limit the effect of password disclosure to a single site only. This eliminates the possibility of large-scale password hijacking. This research naturally integrates with tools such as Dynamic Security Skins and extends previous phishing detection browser mechanisms developed by the Stanford group.

CMU:

Research at CMU has focused on using auxiliary devices such as cell phones or PDAs to provide authentication services, eliminating the possibility of wide scale phishing attacks. This work is closely related to both the Dynamic Security Skins work and the Stanford browser extensions.

CMU has also started a major new workshop, SOUPS: Symposium on Usable Privacy and Security in which many TRUST members are presenting work related to phishing, identity theft, and digital forensics.

3. EDUCATION

3.1. Educational Objectives

Describe the Center's overall educational objectives. In the current reporting period, how have the Center's overall educational objectives and plans changed from the previous reporting period? What performance and management indicators has the Center developed to assess progress in meeting its educational objectives?

The TRUST Vision in Education

One of the drivers of this STC is the view that concerns regarding security must be consciously engineered into new and legacy critical infrastructure systems, and that to do so requires a rethinking of every component level of the system. To ensure that these concerns are shared and addressed by the next generation of computer scientists, engineers and social scientists, TRUST researchers will incorporate their findings and methods wherever possible into the standard. Thus, this project will result in a broad curriculum reform of existing computer science and engineering courses. We will develop a whole set of courses from the lower division to the advanced graduate level as the research on trust matures.

The center has distinct education constituencies – both undergraduate and graduate programs – for which there are distinct mechanisms for knowledge dissemination.

For undergraduates, the center has adopted a two-pronged approach. On the one hand, the center will have activities concerned with diffusing ideas of trustworthiness throughout the entire undergraduate curriculum. On the other hand, the center needs is working towards defining a modern “standard” computer security course at the undergraduate level.

For graduate students, the center finds that a series of summer schools on specific disciplines is where a significant impact can be made, in addition, of course, to developing topic specific customized courses. The summer schools are to be 1-2 week courses, where research leaders provide intensive short courses in areas of current research interest.

Beyond the above partition, the realization that

TRUST solutions = policy options + technology options

requires TRUST to bring together two communities of researchers: technology researchers and policy researchers. Technology done independent of policy risks irrelevance; policy done independent of the technology risks obsolescence or suppresses options.

From the marriage of policy and technology arises some horizontal partitions in addition to the ones by education level, and the TRUST center will engage the educational community to work towards:

A broader understanding of TRUST *technology* options as such among (future) *technologists*
A broader understanding of TRUST *technology* options as such among (future) *policy shapers*
A broader understanding of TRUST *policy* options as such among (future) *policy shapers*
A broader understanding of TRUST *policy options* as such among (future) *technologists*.

The center strategy for achieving this broad influence is through a combination of *push* and *pull* tactics: to *generate learning material* (such as learning modules, course syllabi, textbooks,

broader curricula), provide *effective dissemination structures* (such as on-line repositories, internet delivery mechanisms, summer schools, center-wide seminar series), and establishing *broad educator communities* (such as summer schools, education conference participation) that engage with the center in adopting and adapting the results of the center to their instructional context.

The TRUST Objectives in Research

To establish:

- (a) Learning Technology Infrastructure
- (b) Undergraduate Programs: generate best-practices material for *computer science* courses, security modules for *other engineering* programs and the *social sciences*, create a *signature new undergraduate trusted system course*, capstone experience for undergraduates
- (c) Graduate programs: specialized material for both engineering and policy
- (d) TRUST Summer Schools for Students, for Industry, for Instructors and for Researchers
- (e) A recurring and significant presence at key education conferences
- (f) A series of TRUST domain workshops

The first ten months of the center-wide activities in the education area have focused on c, d, e and f with a ramping up of the efforts related to a and b: on establishing the infrastructure for the learning modules repository, and on establishing a set of pilot course modules within this repository, bringing together material from the various TRUST partner institutions in an integrative learning material generation exercise. The first TRUST summer school will be offered the summer of 2006.

3.2. Current and Anticipated Problems

Discuss any problems the Center may have encountered in making progress toward its educational goals during the reporting period as well as any problems anticipated in the next period. Include plans for addressing these problems.

No problems were encountered.

3.3. Internal Educational Activities

Describe the Center's internal educational activities in the reporting period. Include the activity name, leader, intended audience, approximate number of attendees, and a narrative. The narrative should describe the activity and its goals, outputs, outcomes or impacts, and how the activity will enable the Center to meet its goals.

(a) Learning Technology Infrastructure

TRUST is leveraging an existing learning technology infrastructure for the development and online dissemination of its educational materials that was created by the NSF VaNTH Engineering Research Center for Bioengineering Educational Technologies (<http://www.vanth.org>). This infrastructure has three principal components:

A web-based dissemination portal/content management system

A repository-based authoring technology for adaptive web-based courseware (CAPE)

An online learning platform (eLMS)

The dissemination portal is based on an open source content management system (Plone, <http://www.plone.org>) that has been adapted for educational materials.

The CAPE and eLMS technologies primarily address online learning in blended learning environments. CAPE can additionally be used for curriculum modeling, where the elements can

be classroom-based, online, or blended. Online courseware authored with CAPE is delivered to learners using the eLMS learning platform. This standards-based platform can be used directly, or it can be used as a service from campus learning platforms such as Blackboard, WebCT, or Moodle.

CAPE is used to design online learning experiences involving static, interactive, and dynamic content elements created with conventional web authoring tools and within CAPE itself. The designs specify when, or under what circumstances, content elements are presented to a learner during the course of a learning experience. Interactive elements can elicit information from a learner, and the outcomes are available immediately to adaptations incorporated into designs. A data modeling facility enables capturing facts, including data defined abstractly by expression, for use in realizing adaptation schemes. Simple sequencing constructs can be extended with computational components for more advanced reasoning.

CAPE supports both elaborative (top-down) and integrative (bottom-up) approaches to design. Rapid prototyping of adaptation schemes can be performed prior to content development. Existing content and design elements can be readily incorporated into new designs. The environment supports design-time adaptation by providing abstraction facilities that can be used to capture invariants among families of designs and elements as *instructional design patterns*. While CAPE—as a general-purpose design tool—is pedagogically neutral, these design abstractions can be used to scaffold particular learning strategies that can then be shared with other authors through an integrated *web-based design repository*.

CAPE is built on open source technologies from ISIS—particularly, the Generic Modeling Environment (GME) and Meta-GME—and uses the open source Python dynamic programming language for realizing its extension components and for computational aspects of CAPE designs.

eLMS Learning Platform

eLMS is an adaptive learning platform that supports interoperation using web services, both in conjunction with enacting courseware designs and in managing domain-specific objects, such as classes, users, and courseware.

The platform automatically captures detailed instrumentation of these design enactments, and additional instrumentation—to support grading using custom rubrics, for example—can be incorporated into courseware designs with CAPE. The resulting *delivery records* can be queried by instructors and authors using an integrated *data mining facility*. These capabilities enable an intimate understanding of what learners actually do with on-line learning experiences, which is essential to making incremental improvements over time.

While eLMS can be employed directly to manage the use of CAPE-authored designs by classes of learners, it can also be transparently embedded into other learning platforms, such as WebCT, as well as non-commercial platforms, such as Moodle and (eventually) Sakai.

eLMS is built on open source technologies, including the Zope web application server and Apache, and is deployed on the secure open source OpenBSD operating system.

(b) Undergraduate Programs: generate best-practices material for computer science courses, security modules for other engineering programs and the social sciences, create a signature new undergraduate trusted system course, capstone experience for undergraduates

Team for Research in Ubiquitous Secure Technology
2005-2006 Annual Report

In preparation for the use of these learning technologies and the dissemination portal, during the period under review we have undertaken a set of pilot projects to better understand how to effectively employ the infrastructure and to determine what adaptations might be needed to the technologies themselves.

The objective was also to collect material and establish a set of learning modules for each of three domains – *Network Security*, *Computer Security* and *Chemical Network Plant Security* – and to use the VaNTH repository and the CAPE system to organize the material into a form suitable for re-use and easy adaptability into new course architectures.

The material was collected from across the TRUST partners, organized by personnel at Vanderbilt, San Jose State and Stanford, and collected into the VaNTH system.

For the Network Security collection Yuan Xue (Vanderbilt) and Xiao Su (San Jose State) drew upon material from Vanderbilt (CS291 Network Security), San Jose State (CmpE 209 Network Security) and Stanford (CS259 Security Protocols).

For these courses, we were interested in similarities and differences in terms of sequencing and course content (concepts taught). We were also interested in granularity and the extent to which elements of these courses could be offered to other instructors in units called *modules*, or in sequences of modules called *mosaics*. To conduct these investigations, the courses (in whole or part) were modeled using the CAPE environment. The representations capture how the course was organized into units and how these units were sequenced. Learning objectives for the units were expressed and a common curricular taxonomy developed by Yuan Xue was used to indicate the mapping of subjects to units. Finally, companion resources (typically, lecture notes) were associated with the units. These design representations were shared among the authors using the CAPE Repository.

For the Computer Security set Weider Yu (San Jose State) and Simon Shim (San Jose State) brought together materials from UC Berkeley (CS161 Computer Security, CS276 Cryptography) and Stanford (CS155 Computer and Network Security) for a similar exercise.

In addition to understanding the design of these courses and their relationships, we were also interested in the ability to generate information for the dissemination portal from these formal representations. We used CAPE to create a content generation wizard that assembled information about the course units using their structure, metadata, and taxonomic descriptions.

An additional pilot investigation was conducted of creating online courseware. An interesting dimension of this investigation concerned adapting concepts from information system security for teaching security in another domain: chemical plant security. This pilot is a collaboration between Ken Debelak of Chemical Engineering, Yuan Xue and Janos Sztipanovits of EECS, and Larry Howard of ISIS at Vanderbilt University. The concept for the project is to use role-based access control as a design and analysis approach to teach security concepts in a chemical process engineering capstone design course.

These pilot efforts have informed our thinking about adapting VaNTH's dissemination platform for the new TRUST Academy Online (TAO). In particular, TRUST presents issues of varying granularity that were less important to VaNTH. The pilots have also influenced changes to the CAPE authoring environment to support direct publishing to the dissemination portal.

Team for Research in Ubiquitous Secure Technology
2005-2006 Annual Report

The partners in the center are developing rich new material for courses offered locally and that will be prepared for broader systematic dissemination. The course material ranges from first year experience courses such as the first year experience course at Stanford (CS55N Ten Ideas in Computer Security and Cryptography), through more directly systems- or technically oriented courses such as *System Security* (Cornell), *Fault-tolerant Distributed Computer Systems* (Cornell), *Secure Software Systems* (CMU), *Secure Technologies* (San Jose State) to more policy oriented courses such as *ID-theft* (Stanford) and *Public Policy for Engineers* (UC Berkeley).

At Carnegie Mellon University, the TRUST faculty have developed a new introductory undergraduate security course, and a new course on usability and security.

(c) Graduate programs: specialized material for both engineering and policy

Aside from the curriculum development happening at each of the partner campuses, the center has instituted a roving, webcast-based seminar series. Organized by Bajcsy (UC Berkeley), the seminar provides a webcast to all the partnering campuses, with the instructor being drawn from each of the partners as the semester progresses. The audience is the graduate students of the center partner campuses, and serves to build the TRUST community as well as to serve as a knowledge dissemination vehicle.

Ken Birman (Cornell) has developed a comprehensive set of courseware around his textbook, *Reliable Distributed Systems Technologies, Web Services, and Applications*. The materials include two complete slide sets for instructors, approaching the book from different points of view. These seem very successful. At Cornell the class typically has 75 students, mostly MEng level, of which about 15 are typically women or other diversity students.

(d) TRUST Summer Schools for Students, for Industry, for Instructors and for Researchers

The summer of 2006 the TRUST Center will offer its first Summer School – WISE – the “Women Institute for Summer Education.” The summer school is organized by Ruzena Bajcsy (UC Berkeley), and the intended audience consists of graduate students and faculty who wish to learn about TRUST agenda. The summer school is open to exceptionally senior undergraduates as well. The summer school will enroll some 20 students, and will draw from across the TRUST university and industry partners for instructors.

The summer school will combine lectures by prominent researchers in the TRUST domain with hands-on problem exploration and problem solving. The summer school is expected to be an effective vehicle for the dissemination of recent results in cyber security and relevant policy research, to provide a platform for TRUST domain instructors and to further strengthen the network of collegial relationships among current and future leaders in the TRUST domain. Besides the education component, the summer school also addresses the need for outreach to under-represented communities, and in particular to women – of the 21 applicants, 20 are women with varied backgrounds both in qualifications as well as in ethnicity.

In 2005 Stephen Wicker (Cornell) organized a Networked Embedded Systems Summer Program: We are developing a sensor networking testbed in part to test existing algorithms and processors, as well as to project sensor networking technology into new arenas of research and public utility. As part of an interdisciplinary collaboration between the Cornell School of Electrical and Computer Engineering, Computer Science, the Human-Computer Interaction Group, and

Team for Research in Ubiquitous Secure Technology
2005-2006 Annual Report

the Art Department, we are deploying an iMote-based sensor network in Cornell's Herbert F. Johnson Museum of Art. The initial deployment consists of 26 Motes (donated by Intel through the Nets-NOSS program) deployed in the Asian Art Collection. In the first phase, the network will monitor the time-varying density of patrons across the collection. The data is being collected and processed in real time, and then displayed as part of a separate piece of art. This use of sensor data began as part of HCI research into "expressive AI," the goal being to exploring the interface between art and tools, while engaging users in an open interaction with a sensing system and their own perceptions of art.

In conjunction with the TRUST Center, this effort has expanded to focus on the privacy concerns that emerge from the use of sensors in a public arena. We are now working with the Berkeley Law School to develop policies regarding the use of sensing systems in public spaces. Initial results include development of the concept of limited acuity – carefully limiting sensor capabilities to the needs of task and no further. A workshop is being planned for January in which Cornell and Berkeley researchers will explore the various issues surrounding the use of sensors in public forums.

(e) A recurring and significant presence at key education conferences

The center has submitted a proposal for a TRUST panel session at the 2006 Frontiers in Education conference (October in San Diego). The panel members are drawn from the partnering campuses (San Jose State, Vanderbilt, Stanford and UC Berkeley) as well as from non-center organizations with a strong education mission. The panel title is "*Learning modules for security, privacy and information assurance in undergraduate engineering education.*" The objective is to strengthen the community of TRUST related educators, and to establish a broader community of contributors to the TRUST learning module repository.

TRUST (San Jose State) has together with SEI developed a proposal for funding for a two-year project within the IACB Program offering two summer programs to faculty members from minority universities, with follow-up workshops to share didactic experiences and to disseminate developing knowledge regarding instruction in the IA domain.

TRUST (Vanderbilt) participated in the NSF HBCU-UP conference in Baltimore on 2/2/06. This conference gave us an opportunity to advertise the TRUST center agenda to approximately 400 students and 130 faculty members.

(f) A series of TRUST domain workshops

Sensor Networks and Privacy, Cornell and Berkeley - Tuesday, March 28, 2006.

There were ~30 attendees over the course of the day. The focus of the workshop was the state of the art of sensor networking and the privacy concerns concomitant with their deployment in public spaces. Engineers, computer scientists, social scientists and legal scholars discussed the need and potential for developing public policies that are deeply coupled to the advancing technologies.

Sensor Networking Workshop, Cornell and New York Department of Health - Tuesday, October 11, 2005.

There were ~50 attendees over the course of the day. This workshop focused on the potential uses for sensor networks as a means for protecting critical infrastructure. Attendees included electrical and civil engineers, computer scientists, molecular biologists, psychologists, and chemists who discussed the potential for developing bio-toxin specific sensor networks to

protect water supplies. Attendees also discussed potential uses for sensor networks in public spaces as a means for understanding how the public uses public spaces, and in particular how they interact with publicly displayed pieces of art. The latter was intended to illustrate the extremely broad range of uses for sensor networks in a research context.

Cornell-Tsinghua Workshop on Information Technology, November 18, Tsinghua University, Beijing, China.

Attendees: 40 (faculty, students). The goal of this workshop is to present current research in the area of information technology and explore possibilities in future collaboration. Among the topics covered in this workshop include peer-to-peer systems, wireless sensor networks, learning theory, image and natural language processing.

TRUST Workshop on Social Security Numbers (jointly with PORTIA), Stanford – May 2006.

The purpose of the workshop will be to understand how SSN are used for income tax, social security, credit bureaus, and as an “authenticator” by many kinds of organizations.

In organizing the workshop, we will invite speakers from the financial services industry who can explain how they do business, and what they would need from an SSN alternative. A possible outcome would be a recommendation for revised practices that are acceptable for the financial services industry, and reduce the likelihood of identity theft.

3.4. External Educational Activities

Describe the Center’s external educational activities in the reporting period in a manner similar to 1.c.(3) of this document.

Included in section 3.3.

3.5. Student Participation in Professional Development Activities

Summarize the participation of Center students in professional development activities in the reporting period. Include in the narrative a discussion of how the various professional development activities enable the Center to meet its goals and produce meaningful results.

Included in section 3.3.

3.6. Integration of Research and Education

Describe and discuss the ways in which the Center integrated research and education in the reporting period, with examples as appropriate.

Included in section 3.3.

3.7. Future Plans for Internal and External Educational Activities

Describe the Center’s plans for internal and external educational activities for the next reporting period with attention to any major changes in direction or level of activity.

No major change of direction is anticipated – the activities initiated in the first year of the center will be expected to progress through the second year; a judicious combination of content creation, content organization support and content community construction. The major components of the second year efforts are the *TRUST Academy Online* and *Education Community Development*, besides the ongoing curriculum development and dissemination activities described above.

Team for Research in Ubiquitous Secure Technology
2005-2006 Annual Report

TRUST Academy Online (TAO)

Objectives

The main objective is to build on the pilot experience from year 1 and create a sustained and integrated infrastructure that will support the educational outreach mission of TRUST. The repository-based infrastructure will include tools and services for educators to create and publish learning resources for dissemination online. For consumers, the infrastructure will provide services for identifying, retrieving, and using teaching and learning resources, including support for courseware delivered directly to individual learners.

The dissemination portal of the VaNTH system will be adapted in appearance, organization, and services and workflows for the underlying content management system will be defined to support development and editorial processes identified for TRUST. The aim of these modifications will be to enable educators (and reviewers) to directly perform tasks associated with the preparation of new resources for dissemination and the evolution of already available resources. The TAO core project team will work directly with TRUST educators and end users to address usability aspects of the dissemination portal.

The project will provide direct support for TRUST educators in developing new learning modules using features of VaNTH's adaptive online learning technologies.

Deliverables

Tailoring of Plone CMS including content templates, workflows, site navigation and visual styling.

Online tutorials that support use of dissemination portal by TRUST users.

CAPE integration with dissemination portal.

Technology and pedagogy support for development of new learning modules.

Education Community Development (EDC)

Objectives

This focus has the following four objectives:

Development of re-targetable courseware modules made available through TAO

Development of prototype courseware and making it available through TAO

Demonstration and evaluation of the use of TAO

Establishing a broad community of educators that utilize (and contribute to) resources provided through the infrastructure

Approach

In the first year, content development played two roles. First, the developed components served as well documented examples demonstrating the development method, content granularity and packaging technology. Second, the delivered modules provided reusable assets for creating systems/security courses or enriching existing courses with security content. We expect that TRUST faculty will contribute to content development in a range of topics. The TAO team will provide coordination and support for the contributors to decrease the "time penalty" coming with the use of the infrastructure. The community building will proceed by organizing sessions and workshops associated with key education-oriented conferences, by organizing (or participating in the organization of) TRUST oriented education symposia and by working with educators to test the usability of the infrastructure established in (1) and (2).

TRUST (San Jose State) has partnered with SEI to propose for NSF funding a two-year project

Team for Research in Ubiquitous Secure Technology
2005-2006 Annual Report

for Information Assurance Capacity Building, which (if funded) will combine a month-long summer program for instructors working in the Information Assurance domain with a follow-up intense workshop focusing on the teaching of TRUST topics.

Deliverables

Security courseware modules based on the learning material of network security course at Vanderbilt University

Implementations of courseware modules and their related taxonomies in the CAPE system. The modules include

Network attacks

Cryptography

Authentication protocol

Network security standard and applications

Courseware modules to implement hardware support for security using field programmable gate arrays (FPGAs) based upon learning material from an FPGA Design course

Software and physical attacks on FPGAs

Multidisciplinary General Education modules suitable for use in lower division courses

Modules suitable for use in undergraduate business programs

Demonstration of applying several re-targetable modules (e.g., role-based access control) to the domain of chemical process control.

Demonstration of the use of the prototype infrastructure for course development.

Establish a broad-based TRUST oriented education community.

Establish a regular and repeating presence of TRUST at education oriented conferences.

Use such conferences to engage other campuses in conversations about TRUST didactics, bring them into the course material development, repositorying and use community. Try to engage conferences that extend beyond the CS domain, and conferences that engage the education-mission universities and colleges.

Establish a (periodic) security oriented education symposium/workshop. Tie the symposium to summer schools or other established gathering of experts in the field.

4. KNOWLEDGE TRANSFER

4.1. KT Objectives

Describe the Center's overall knowledge transfer objectives. In the current reporting period, how have the Center's overall knowledge transfer objectives and plans changed from the previous reporting period? What performance and management indicators has the Center developed to assess progress in meeting its knowledge transfer objectives?

In our NSF proposal to create TRUST we articulated a broad and strong vision for knowledge transfer, as follows. The structure of TRUST lends itself to a comprehensive approach to knowledge transfer. Since TRUST addresses well defined and long term societal needs, the results in computer security, privacy and critical infrastructure protection can be easily communicated to decision makers, policy makers, and government agencies. With respect to industry, the selected integrative testbeds represent focal points for interaction and dialog with major stakeholder industries: power, telecommunication and embedded systems. In fact, several integrative testbeds are being provided by the stakeholders, which offer significant leverage for the Center. To facilitate technology transfer from the research community to the industrial community a number of the investigators on this proposal, led by Sastry and Sztipanovits, have created a non-profit entity entitled ESCHER (Embedded Software Consortium for Hybrid Systems Research) for acting as a repository for the tools and algorithms developed by the researchers and for establishing case-studies for design. TRUST will utilize ESCHER as a repository for developed tools and reference solutions. Finally, TRUST researchers are leaders in their scientific communities. Their broad cooperation to achieve the TRUST objectives will serve as a catalyst to turn attention of the community toward the emerging science of secure systems.

TRUST comprises multiple institutions, technology vendors, and infrastructure users and providers. Broad participation from leading research universities, undergraduate colleges serving under-represented groups, computer vendors (IBM, HP, Intel, Microsoft, Cisco, Symantec), and infrastructure providers (Bellsouth, Raytheon, Boeing, Qualcomm, GM) will result in wide spread dissemination, adaptation and continued evolution of ubiquitous secure technology. Our research will learn and evolve with our results, using an iterative investigate-develop-educate-apply cycle. We will develop science, technology and proof of concept prototypes that will be tested through models that emerge from a series of analytical and case studies, experimentation and simulations. We plan to use periodic updates of living reports and community workshops throughout the life-cycle of TRUST. The research output of the Center will be disseminated in four ways: (1) publications in the open literature and on the web, (2) Short courses held at major ACM and IEEE conferences as well as Infrastructure Protection Meetings, (3) Public Lectures and Meetings with the general public concerned about security and privacy issues on the internet and critical infrastructure protection, and (4) curriculum development and courses taught at the partner institutions as well as the outreach institutions.

During the reporting period, we believe that TRUST has been solidly on track with respect to its entire transfer objective. Success is measurable in many ways: technologies that are being commercialized, TRUST researchers who are working hand-in-hand with industry and standards groups to help improve trustworthiness of major infrastructure systems, activities aimed at educating the public and exploring non-technical ramifications of TRUST themes, and development of significant TRUST spin-offs, such as the AF-TRUST-GNC center for the Air Force, the exploratory work on a center for research on trustworthy electronic health records,

and the TRUSTED Financial Systems center under discussion with Treasury.

4.2. KT Current and Anticipated Problems

Discuss any problems the Center has encountered in making progress toward its knowledge transfer goals during the reporting period as well as any problems anticipated in the next period. Include plans for addressing these problems.

Relatively few problems have arisen. Several researchers expressed concern about new government rules concerning “deemed export”, restrictions on involvement of foreign nationals in research on security-related topics, pre-publication review restrictions, intellectual property concerns (such as industry members who are unable to accept the constraints associated with IP development in academic research settings) and military projects that require clearances. Our community is thus denied access to certain kinds of projects and research problems. However, we are pleased to report that this has not risen to the level of posing a threat to the success of TRUST, and we also feel that for a group with our scope and prominence, issues of this nature were inevitable. Some individuals in TRUST, of course, hold clearances and play more direct roles in classified research, but this sort of information is not suitable for inclusion in the present report.

With respect to dialog with stakeholder industries, the level of dialog has been higher in some areas than in others, but this was to be expected. For example, we are finding that the Electric Power research community is in a degree of disarray caused in part by the increasingly commercial stance of the Electric Power Research Institute, which previously functioned as a neutral ground for exploration of research issues in power, but now is more and more product-oriented and operates many for-profit activities. Thus, whereas TRUST in the past might have been able to engage that industry through EPRI, today this is less practical. Yet even in this area, TRUST researchers have made progress, by organizing workshops around next-generation SCADA platforms that successfully included a diversity of researchers from industry, academic settings and government.

Similarly, we have been pursuing dialog with the financial community and are finding that there are barriers to progress, but not of a nature that cannot be surmounted. Our work has been undertaken with the help and guidance of the Deputy Secretary of the Treasury for Critical Infrastructure Protection (Scott Parsons), and so far has led to a valuable series of workshops and meetings. Treasury is not able to directly support research, and this is one challenge for us. Yet they are able to function in a match-making role and this is proving quite valuable. Industry players are fiercely competitive with one-another and not inclined to share their needs or agendas, but Treasury is able to help us overcome this understandable reticence. We anticipate that we'll create a small center as a spin-off from TRUST in 2006/2007 out of which a larger effort can grow.

Schmidt, who works with many government projects, points to a different kind of transitioning obstacle. Although he has been relatively successful with transitions, he often encounters inertia and risk aversion, particularly in the context of large DoD acquisition programs, where it takes a long time to convince DoD buyers that the technologies are mature enough to be trusted in mission-critical systems. A related problem is that the DoD's certification processes are not well-suited to for flexible/adaptive technologies (such as dynamic resource management for shipboard computing), which makes it harder to get them approved for weapons systems.

Thus while not every path that we are exploring has born fruit, we believe that TRUST is finding considerable success as it pursues avenues for dialog with industry and knowledge transfer.

4.3. KT Activities

List organizations with which knowledge transfer has occurred and the frequency and type of interactions. Describe each of the Center's knowledge transfer activities with goals and outputs or impacts in the current reporting period and discuss how they enable the Center to meet its goals. For each knowledge transfer activity, provide the name of the activity, leader, participants (organization name and state).

Alex Aiken (Stanford): Aiken's primary activities in these categories are posting bug reports for the open source software that my group analyzes. We have conservatively posted 1,000 bugs to developers in the last year and a half, many of which have security implications. The tools will be open sourced at some point, probably starting this summer. In this respect he points to one obstacle, namely the difficulty of finding both funding and adequate time to get these research platforms far enough past the research prototype stage so they are in form non-researchers can use. Aiken receives many spontaneous inquiries from developers asking for the tools, and is planning to go the open source route.

Venkat Anantharam (Berkeley): Anantharam reports that in October 2005, he was an invited speaker at a NATO Workshop sponsored through the "Security Through Science" programme. The topic of the workshop was "Network Security and Intrusion Detection". His NATO talks will be part of a volume to be published. I have also given talks on security related topics at the Tata Institute for Fundamental Research (India), the Indian Institute of Technology, Mumbai (India), and Ecole Normale Superieur (France) in 2005.

Ruzena Bajcsy (Berkeley): Bajcsy points to her role on the advisory board of Argonne Nat. lab, which has focused on security of facilities. It meets 3 times a year. She is also on a board for NIH/NIST which is planning standard for images as biomarkers. Here the privacy issues are of particular importance.

Ken Birman (Cornell): Birman has been extremely active in TRUST-related transitioning efforts during the reporting period. First, Birman participated in a series of high-profile studies for the Air Force that focused on TRUST themes that arise in connection with that organization's move to GIG and NCEC "SOA" standards. One study, for the Air Force CIO (Mr. Gilligan, later replaced by General Croom and Mr. Tillotson) focused on the implications of the deployment now underway in these areas; the second, Prometheus, was conducted for AFRL and explored options for aligning AFRL research on the Joint Battlespace Infosphere with AFRL priorities. Sastry participated in the CIO study, and Schmidt and Reiter were team-members on the Prometheus study. Both resulted in additional funding to the TRUST community. Additionally, he has worked with the US government both to develop a new national strategy for research in cyber security (this was part of an effort led by DHS but also involved participants from White House OSTP and NSF), and with the US Department of Treasury on the creation of a small center for research on TRUST issues in financial settings. A 2-day research topic on the subject helped refine a Treasury priorities and strategic vision document, and Birman is now teaming with developers of the eCavern remote backup and disaster recovery facility on replication techniques for their setting. As program committee chairman for the 20th ACM Symposium on Operating Systems Principles, Birman helped highlight many of the best results in the field by emphasizing TRUST topics in the call for papers, and also arranged a panel on peer-to-peer computing that focused on the real value and robustness of these new but highly controversial protocols and techniques. Birman's group has developed software that is in wide use; his latest efforts include Astrolabe (which runs Amazon.com's data centers), and Ricochet, which has just been released to the public in open-source form and slashes the latencies for time-critical computing systems. Birman also points to several publications aimed specifically at

Team for Research in Ubiquitous Secure Technology
2005-2006 Annual Report

educating the general public about TRUST issues, notably Ken Birman, Coimbatore Chandrasekaran, Danny Dolev, and Robbert van Renesse. How the Hidden Hand Shapes the Market for Software Reliability. Submitted to the First IEEE Workshop on Applied Software Reliability, June 2006; Ken Birman. The Untrustworthy Services Revolution. IEEE Computer (ISSN 0018-9162). Vol.39 No.2, Pgs. 98-100. February 2006; and Ken Birman. Can Web Services Scale Up? . IEEE Computer. Volume 38. Number 10. Pgs.107-110. October 2005. Birman's book has been widely adopted as the basis for MEng and PhD-level courses in reliability and trusted computing: *Reliable Distributed Systems Technologies, Web Services, and Applications*. Birman, Kenneth P. 2005, XXXVI, 668 p. 145 illus., Hardcover ISBN: 0-387-21509-3.

Dan Boneh (Stanford): Boneh has a long history of collaboration with the San Francisco offices of the US Secret Service and Oakland FBI cybersquad. He also consults for Microsoft on crypto research topics. Technology transition success stories from the reporting period include the adoption of technology from SpoofGuard anti phishing plug-in used by EBay toolbar, and in the Earthlink ScamBlocker. He also points to his PwdHash browser extension, which is being integrated into RSA SecurID server.

Sigurd Meldal (Vanderbilt): Consults with Morrison & Foerster on IP issues, many of which involve TRUST issues (in particular encryption and networking). Meldal is also offering a new graduate level course to Intuit employees on QA, which involves TRUST-related topics.

George Necula (Berkeley): Necula has consulted for Microsoft on the issue of language-based security mechanisms for Windows device drivers. In this context, he encountered one of the barriers to success mentioned earlier. A Microsoft technical manager became convinced that Necula's technology and tools might be helpful, but progress halted when lawyers got involved. The issue was that Berkeley's technology was packaged as an open-source tool (CCured) that would rewrite their code to ensure all sorts of safety properties. Microsoft's lawyers felt that they could not allow an open-source tool to rewrite Microsoft proprietary code because they have not studied what the legal status of the output of the code would be: open-source or proprietary? Over this concern, everything came to a halt. Apparently the issue rose to the highest levels: company lawyers were ultimately asked to present their concerns to Bill Gates, who pressed for a solution. Nonetheless, Microsoft was unable to resolve the matter and Necula's technical contacts eventually dropped the issue.

Joseph O'Rourke (Smith College): Has involved students in projects with a security and trust theme and helped some join in broader collaborations with TRUST faculty elsewhere.

Adrian Perrig (CMU): Collaborates with researchers in the following corporations: Bosch, Cisco, IBM, Intel, Microsoft on trusted computing issues. He was recently featured on CNBC to talk about click-fraud, spyware, adware, etc. Perrig plays roles in several standards communities, including the IETF MSEC working group for the standardization of the TESLA protocol. He is also active in the WiFi alliance, where he gave 2 presentations on key establishment and simple and secure access point configuration. Perrig's books include Adrian Perrig and J. D. Tygar. "Secure Broadcast Communication in Wired and Wireless Networks". Kluwer Academic Publishers, 2002. He also points to several papers on TRUST-related topics, notably Adrian Perrig, John Stankovic, and David Wagner. Security in Wireless Sensor Networks. In Communications of the ACM, 47(6):53-57, June 2004; Yih-Chun Hu and Adrian Perrig. A Survey of Secure Wireless Ad Hoc Routing. In IEEE Security & Privacy, 2(3):28-39, 2004; Haowen Chan and Adrian Perrig. Security and Privacy in Sensor Networks. In IEEE Computer

Team for Research in Ubiquitous Secure Technology
2005-2006 Annual Report

Magazine, pages 103-105, October 2003; Elaine Shi and Adrian Perrig. Designing Secure Sensor Networks. In *Wireless Communication Magazine*, 11(6):38--43, 2004. His work is being used, for example by Bosch, which is deploying secure sensor network protocols that he designed, and ZigBee, which based their security protocols on Adrian Perrig, Robert Szewczyk, Victor Wen, David Culler, and J. D. Tygar. "SPINS: Security Protocols for Sensor Networks". In *ACM Journal of Wireless Networks*, to appear, 2002.

Mike Reiter (CMU): Mike Reiter has monthly interactions with Robert Bosch Corporation; he collaborates with them for developing secure sensor network protocols. He also has less frequent interactions with the CyLab corporate partners. Reiter has substantial projects with AFRL, DARPA, ARO, and ONR. Through his roles in TRUST and in the CMU Security Institute, he interacts with a number of companies that are active in the CyLab corporate partners program. The list is quite extensive; it can be found at <http://www.cylab.cmu.edu/default.aspx?id=250>

Mendel Rosenblum (Stanford): Mendel Rosenblum works as the Chief Scientist of VMware Inc. (<http://www.vmware.com>) This company markets the industry leading product for virtualization, and is credited with having had an enormous impact on the reliability and security of systems of all kinds. Moreover, VMWare is key to such projects as the various Honeyfarms used to detect and deter virus attacks.

Doug Schmidt (Vanderbilt): Schmidt has a long and particularly successful track record of technology transition and dialog with industry. Recent noteworthy events include the following. First, he has consulted for several DoD acquisition programs, e.g., Navy's DD(X) program, Army's FCS program, Air Force's F-15, FA-18, and AV-8B programs. He has also worked closely with a number of DoD Service labs, e.g., AFRL/IF and ONR, and several DoD system integrators, e.g., Boeing, LMCO, Raytheon. On behalf of funding agencies, he chaired the NSF/NCO Workshop on New Research Directions in High Confidence Software Infrastructure for Distributed Real-time and Embedded (DRE) systems, July 13th, 2006, Fairfax VA. His team has developed middleware and modeling tools used by major IT companies, e.g., Qualcomm, Siemens, Cisco, Symantec, IBM. Schmidt was guest editor for IEEE Computer special issue on Model-Driven Engineering, Feb 2006. He works with the Object Management Group on middleware and modeling tools for distributed real-time and embedded systems, and is considered to be the "father" of the CORBA real-time and quality-of-service standards. Moreover, he has had many technology transition successes. His work on dynamic resource management algorithms, component deployment and configuration middleware for system integration, and model-driven tools for system execution modeling and performance analysis in the DARPA Adaptive and Reflective Middleware Systems (ARMS) program is transitioning to the Navy's DD(X) Destroyer program.. His work on Real-time CORBA middleware for distributed real-time and embedded systems has transitioned to manned/unmanned combat air vehicles, the Orbital Express, low earth orbit (LEO) satellite, telemetry and control framework, and the Ground Support System (GSS) for the X33 Single Stage To Orbit (SSTO) Reusable Launch Vehicle, and the USS Ronald Reagan, and the USAF Upgraded early warning radar system, among many others.

Fred Schneider (Cornell): Schneider's high profile both within TRUST and on the national scene have given him exceptional opportunities for transitioning activities and scholarship in the area of trusted computing systems and security. He has taught short courses at many workshops, given keynote talks on TRUST topics, and advises a great number of companies. Major trust-related activities include the following: Through the Cornell/AFOSR Information Assurance Institute, Schneider maintains a continuing and high-bandwidth dialog with the

Team for Research in Ubiquitous Secure Technology
2005-2006 Annual Report

airforce about their security needs. He is a member of the editorial boards of several journals, including IEEE Security and Privacy, November 2002--present (Associate Editor-in-Chief). IEEE Transactions on Dependable and Secure Computing, March 2004--present. He is a member of the National Research Council, Computer Science and Telecommunications Board, March 2002—present, National Science Foundation CISE Advisory Committee, March 2002--present. Committee on Improving Cybersecurity Research, Computer Science Telecommunications Board, National Research Council, National Academy of Sciences. June 2004--present. Schneider has several major roles with industry in TRUST related capacities, including FAST ASA, Chief Scientist-Security and Reliability, March 2000-present; Cigital, Technical Advisory Board, Nov. 2000--present. He is Co-Chair of Microsoft's Trustworthy Computing Academic Advisory Board, Feb 2003—present. And he serves Fortify Software, as a member of their Technical Advisory Board, Feb. 2004—present.

Gun Sireer (Cornell): Researchers at AFRL(Rome Research labs) have been collaborating with his team on identifying Air Force applications for Corona, a high-performance event monitoring and publish-subscribe system. He has worked with CNNIC, the name registrar for the Chinese (.cn) name space, to deploy and evaluate CoDoNS software as a DNS for Chinese computing nodes. He has deployed the CobWeb cache, which is an open access Akamai-like system for speeding up web browsing and protecting content providers from flash crowds, on PlanetLab, where it handles between 10-15 million requests per day. His Credence system, for determining the trustworthiness of peers and for identifying pollution in large scale filesharing networks, has been downloaded by over 10000 people. The Corona system has been deployed on PlanetLab and currently monitors a few hundred channels on behalf of its users. The Meridian system for locating nearby nodes has recently been deployed through a site called "closestnode.com". In addition to his many technical papers during the reporting period (tabulated elsewhere), Sireer has also written a paper specifically motivated by TRUST dialog and issues: Emin Gun Sireer. Heuristics Considered Harmful: Using Mathematical Optimization for Resource Management in Distributed Systems. IEEE Intelligent Systems, Special Issue on Self-Management through Self-Organization in Information Systems, Mar/Apr 2006.

Dawn Song (CMU): Song has worked closely with the Army Research Office on security issues, and also consults for Symantec, Microsoft on Malware Detection and Defense. She is co-editor on a book titled "Malware Detection and Defense". A number of technology transfer activities are currently in the pipeline.

Lang Tong (Cornell): working with Army Research Laboratory (ARL) to extend his TRUST work to military wireless ad hoc networks.

David Wagner (Berkeley): Consults for Fortify Software, a startup producing software security tools, on their security products.

Stephen Wicker (Cornell). Steve served as Program Committee Chair for ACM Conference on Information Processing in Sensor Networks (IPSN 2006), and as a Principal Investigator for NSF Nets_NOSS and ITR programs. In January he gave talks on sensor networking at the Indian Institute of Technology campuses in Kanpur and Delhi. In May he will give an invited talk on "Sensor Networking and Privacy – Conflicting Agendas" at the CMU Special workshop on Sensor Network Security. Steve I also organized the following two workshops. *Sensor Networks and Privacy*, Cornell and Berkeley - Tuesday, March 28, 2006. There were ~30 attendees over the course of the day. The focus of the workshop was the state of the art of sensor networking and the privacy concerns concomitant with their deployment in public

Team for Research in Ubiquitous Secure Technology
2005-2006 Annual Report

spaces. Engineers, computer scientists, social scientists and legal scholars discussed the need and potential for developing public policies that are deeply coupled to the advancing technologies. *Sensor Networking Workshop*, Cornell and New York Department of Health - Tuesday, October 11, 2005 (agenda attached). There were ~50 attendees over the course of the day. This workshop focused on the potential uses for sensor networks as a means for protecting critical infrastructure. Attendees included electrical and civil engineers, computer scientists, molecular biologists, psychologists, and chemists who discussed the potential for developing bio-toxin specific sensor networks to protect water supplies. Attendees also discussed potential uses for sensor networks in public spaces as a means for understanding how the public uses public spaces, and in particular how they interact with publicly displayed pieces of art. The latter was intended to illustrate the extremely broad range of uses for sensor networks in a research context. Steve also consults for Qualcomm, Nokia, Samsung, Verizon, Sprint, and Motorola.

Hal Varian (Berkeley): Varian has been one of the organizers for a series of Workshops exploring the Economics of Information Security. This group is having its 5th meeting in England in June. <http://www.cl.cam.ac.uk/~rja14/econsec.html>

Jeanette Wing (CMU): A wide range of relevant activities, including membership on the Microsoft Trustworthy Academic Advisory Board, the muSecurity advisory board and other Silicon Valley startups specific to security. Wing has written several papers that represent knowledge transition vehicles, such as J.M. Wing, "Beyond the Horizon: A Call to Arms," IEEE Security and Privacy, November/December 2003, pp. 62-67 and J.M. Wing "Computational Thinking," CACM, vol. 49, no.3 March 2006, pp. 33-35. She also receives additional funding for security-related research from the SEI, CMU's CyLab (ARO money), an ARO URI, and an NSF ITR. Wing is co-PI on the CMU SAFE Center funded by the NSF Cybertrust Program. She collaborates with people in the Idaho National Laboratory and is on the Idaho National Laboratory National and Homeland Security Strategic Advisory Committee.

4.4. Other KT Activities

Describe any other outcomes or impacts of knowledge transfer activities not listed above. Discuss, in particular, applications of Center research in industry, federal laboratories, or elsewhere.

To avoid repetition, we simply note that there are many such applications in the list given above, and this focuses only on the past year. Were we to cover a longer time period, we would add such applications as the computing infrastructure of the New York and Swiss Stock Exchanges (which use Birman's Isis technology), the US Navy AEGIS warship (again, Isis) and the French Air Traffic Control System (Isis, in two major subsystems). IBM Websphere and Microsoft's latest scalable clustering platform both use reliability and replication technologies developed jointly with Birman's group. And these are just a few of many examples that have arisen during the past few years.

4.5. KT Future Plans

Describe the Center's plans for knowledge transfer activities for the next reporting period with attention to any major changes in direction or level of activity.

Our hope is that in the coming year, TRUST dialog with major stakeholder communities will gain momentum and take on lives of their own. We have created AF-TRUST-GNC as a center for expertise and research on Air Force trusted computing needs, and hope to see similar centers arise in the areas of SCADA computing, electronic health care records, and trusted computing for financial applications. These centers will bring additional resources to the table and also

Team for Research in Ubiquitous Secure Technology
2005-2006 Annual Report

provide concrete application areas on which TRUST researchers can focus attention.

5. PARTNERSHIPS

5.1. Partnership Objectives

Describe the Center's overall partnership objectives. In the current reporting period, how have the Center's overall objectives and plans changed from the previous reporting period? What performance and management indicators has the Center developed to assess progress in meeting its partnership objectives?

One of the goals of the center is to serve as a trusted intermediary between academics, industry, and policy makers, while simultaneously addressing long term societal needs in its research and education activities, and pursuing knowledge transfer. To integrate these objectives together, TRUST has sought to partner with several representatives of the IT industry and national laboratories both for the sake of knowledge transfer as well as for guidance in its overall strategic planning and implementation through the External Advisory Board. Several performance indicators are used to track progress in meeting the overall metric of global impact of the center—number of partners, amount of funds donated by industrial partners to TRUST activities, number of collaborations in knowledge transfer activities, joint research activities with national laboratories, amount of interaction with policy making bodies and governmental agencies. On all these metrics, TRUST is making very steady and significant progress.

5.2. Current and Anticipated Problems for Partnerships

Discuss any problems the Center has encountered in making progress toward its partnership goals during the reporting period as well as any problems anticipated in the next period. Include plans for addressing these problems.

No significant problems have been encountered.

5.3. Partnership Activities

Describe and discuss the activities that are conducted as part of partnerships. Lists the organizations, domestic or international, with which your Center has established partnerships or those with which your Center has collaborated. Also list other organizations with which the Center may share equipment, facilities, and resources (even without being a formal partner); and describe how the Center or the organizations use the resources. If appropriate, include any activities covered above in the knowledge transfer category if they were part of a partnership agreement. Discuss how the partnership activities enable the Center to meet its goals.

First, TRUST has several academic partners who work closely together on research, education, and management activities as a center. The university partners are:

1. University of California, Berkeley (lead institution)
2. Carnegie-Mellon University
3. Cornell University
4. Mills College
5. Stanford University
6. San Jose State University
7. Smith College
8. Vanderbilt University

Second, the TRUST industrial partners to date who participate in knowledge transfer, serve on the External Advisory Board, or collaborate actively in research are:

1. Cisco Systems
2. ESCHER
3. Hewlett Packard
4. IBM
5. Intel Corporation
6. Microsoft Corporation
7. Pirelli
8. Qualcomm
9. Sun Microsystems
10. Symantec Corporation
11. Telecom Italia
12. United Technologies

Third, TRUST has established collaborations with two national laboratories, Oak Ridge National Laboratories and Lawrence Berkeley National Laboratories.

Several TRUST members serve on national boards and influence policy makers—these include National Science Foundation's committee of visitors, DARPA's ISAT study panel, AFRL's advisory board, etc. In addition, TRUST has sought and obtained supplemental funding from other governmental agencies that will leverage TRUST work. These include DHS, AFOSR, and AFRL at the moment.

5.4. Other Partnership Activities

Describe any other outcomes or impacts of partnership activities not listed elsewhere.

As part of TRUST's goals of disseminating results, we are eager to establish relationships with international programs where mutually beneficial opportunities exist. Our first large effort in this direction is with Taiwan. Our program has received significant attention from Taiwan, and funds for cooperating with TRUST have been approved the National Legislature (the Legislative Yuan) and a member of the Taiwanese Cabinet at the level of Minister of State has been assigned to oversee the program.

Taiwan is a leading player in the world of electronics and IT -- Taiwan has been expanding its scope from more narrowly focused areas in manufacturing and integrated circuit design to become an aggressive player in the world of IT services. Taiwan is by most accounts has the second or third largest penetration of broadband services (as of July 2005, with 10.5 million broadband users and 14.6 Internet users out of a total population of 22.8 million.)

Taiwan also faces unique challenges because of its relationship with mainland China, and both public and private institutions in Taiwan are under constant attack from mainland Chinese sources. Some of these are believed to be government sponsored.

Based on TRUST, Taiwan has set up an inter-university institute called TWISC (Taiwan Information Security Center), and has adopted an international collaboration center for research in computer security, directed by Dr. D. T. Lee, a former NSF program officer. TWISC is overseen by the cabinet level Science and Technology Advisory Group (run by a Minister of State). Major members include the "Taiwanese NSF" (NSC, the National Science Council); III, the Institute for Information Industry (a public/private software industry coordinating group); ITRI, the Industrial Technology Research Institute; major infrastructure groups (such as telecommunication companies); and government representatives from public safety and law enforcement.

Funds have been allocated effective April 1 for collaboration and final negotiations on funding are in process initially with Berkeley and CMU, with plans for expansion to other TRUST members. Total funding to TRUST from Taiwan is likely to be approximately US\$ 2 million/year.

We have considerable excitement in TRUST over the collaboration because of the outstanding quality of our Taiwanese research counterparts, their impact in the IT area, and because of the chance to observe the emerging patterns of cyber attack within Asia (and particularly emerging from mainland China) firsthand.

5.5. Future Plans for Partnerships

Describe the Center's plans for partnership activities for the next reporting period with attention to any major changes in direction or level of activity.

In 2006-2007, TRUST partnership will spread through numerous active collaborations with industry, educational and research activities such as the one described above with Taiwan, and joint research activities with national laboratories. No major changes are planned.

6. DIVERSITY

6.1. Diversity Objectives

Describe the Center's overall objectives related to increasing diversity at the Center. If there have been any changes in the Center's overall objectives and plans related to increasing diversity since the last reporting period, discuss these changes and the reasons behind them. What performance and management indicators has the Center developed to assess progress in meeting its diversity objectives?

The overall TRUST goal is to have no weak links left in the education of our society about the technical, compositional, privacy, economic and legal aspects of trusted information systems. To this end, we will begin locally but spread our outreach as far as we can along as many diverse axes as we can.

The specific objectives of our currently planned outreach activities are:

- Grades 6-12 outreach: educating children about cyber security 6-12 (Oakland and Pittsburgh) through the Berkeley Foundation for Opportunities in Information Technology (BFOIT)
- Summer Research in Information Assurance for HBCU/HSI Faculty (CMU, Berkeley, Cornell)
- Curriculum Development for HBCU/Hispanic Serving Institutions (CMU, SJSU)
- Summer Internship for HBCU Faculty in TRUSTed Embedded Systems (Vanderbilt)
- Women Only Universities Research (Mills, Smith)
- Community Outreach (all campuses)

6.2. Demographics

Provide demographic data for the Center with respect to gender, disability status, ethnicity, race, and citizenship.

All Sites

Constituency	Gender		Race				US citizen		Total
	M	F	White	African American	Asian	Hispanic	Y	N	
Faculty	45	7	36	2	13	1	51	1	52
	86.54%	13.46%	69.23%	3.85%	25.00%	1.92%	98.08%	1.92%	46.02%
Undergraduate students	1	3	0	1	1	2	4	0	4
	25.00%	75.00%	0.00%	25.00%	25.00%	50.00%	100.00%	0.00%	3.54%
Graduate students	32	13	17	3	25	0	24	21	45
	71.11%	28.89%	37.78%	6.67%	55.56%	0.00%	53.33%	46.67%	39.82%
Research scientists	3	0	3	0	0	0	3	0	3
	100.00%	0.00%	100.00%	0.00%	0.00%	0.00%	100.00%	0.00%	2.65%
Post Doctorates	5	0	3	0	2	0	1	4	5
	100.00%	0.00%	60.00%	0.00%	40.00%	0.00%	20.00%	80.00%	4.42%
Staff	3	1	4	0	0	0	4	0	4
	75.00%	25.00%	100.00%	0.00%	0.00%	0.00%	100.00%	0.00%	3.54%
Total	89	24	63	6	41	3	87	26	113
	78.76%	21.24%	55.75%	5.31%	36.28%	2.65%	76.99%	23.01%	100.00%

UC Berkeley

Team for Research in Ubiquitous Secure Technology
2005-2006 Annual Report

Constituency	Gender		Race				US citizen		Total
	M	F	White	African American	Asian	Hispanic	Y	N	
Faculty	18	3	16	1	4	0	20	1	21
	85.71%	14.29%	76.19%	4.76%	19.05%	0.00%	95.24%	4.76%	46.67%
Undergraduate students	1	3	0	1	1	2	4	0	4
	25.00%	75.00%	0.00%	25.00%	25.00%	50.00%	100.00%	0.00%	8.89%
Graduate students	6	6	5	1	6	0	8	4	12
	50.00%	50.00%	41.67%	8.33%	50.00%	0.00%	66.67%	33.33%	26.67%
Research scientists	1	0	1	0	0	0	1	0	1
	100.00%	0.00%	100.00%	0.00%	0.00%	0.00%	100.00%	0.00%	2.22%
Post Doctorates	3	0	2	0	1	0	1	2	3
	100.00%	0.00%	66.67%	0.00%	33.33%	0.00%	33.33%	66.67%	6.67%
Staff	3	1	4	0	0	0	4	0	4
	75.00%	25.00%	100.00%	0.00%	0.00%	0.00%	100.00%	0.00%	8.89%
Total	32	13	28	3	12	2	38	7	45
	71.11%	28.89%	62.22%	6.67%	26.67%	4.44%	84.44%	15.56%	100.00%

Carnegie Mellon University

Constituency	Gender		Race				US citizen		Total
	M	F	White	African American	Asian	Hispanic	Y	N	
Faculty	2	2	3	0	1	0	4	0	4
	50.00%	50.00%	75.00%	0.00%	25.00%	0.00%	100.00%	0.00%	36.36%
Undergraduate students	0	0	0	0	0	0	0	0	0
	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Graduate students	5	2	3	0	4	0	5	2	7
	71.43%	28.57%	42.86%	0.00%	57.14%	0.00%	71.43%	28.57%	63.64%
Research scientists	0	0	0	0	0	0	0	0	0
	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Post Doctorates	0	0	0	0	0	0	0	0	0
	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Staff	0	0	0	0	0	0	0	0	0
	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Total	7	4	6	0	5	0	9	2	11
	63.64%	36.36%	54.55%	0.00%	45.45%	0.00%	81.82%	18.18%	100.00%

Cornell University

Constituency	Gender		Race				US citizen		Total
	M	F	White	African American	Asian	Hispanic	Y	N	
Faculty	8	0	5	0	3	0	8	0	8
	100.00%	0.00%	62.50%	0.00%	37.50%	0.00%	100.00%	0.00%	33.33%
Undergraduate students	0	0	0	0	0	0	0	0	0
	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Graduate students	9	5	4	0	10	0	6	8	14

Team for Research in Ubiquitous Secure Technology
2005-2006 Annual Report

Research scientists	64.29%	35.71%	28.57%	0.00%	71.43%	0.00%	42.86%	57.14%	58.33%
	1	0	1	0	0	0	1	0	1
Post Doctorates	100.00%	0.00%	100.00%	0.00%	0.00%	0.00%	100.00%	0.00%	4.17%
	1	0	1	0	0	0	0	1	1
Staff	100.00%	0.00%	100.00%	0.00%	0.00%	0.00%	0.00%	100.00%	4.17%
	0	0	0	0	0	0	0	0	0
	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Total	19	5	11	0	13	0	15	9	24
	79.17%	20.83%	45.83%	0.00%	54.17%	0.00%	62.50%	37.50%	100.00%

Mills College

Constituency	Gender		Race				US citizen		Total
	M	F	White	African American	Asian	Hispanic	Y	N	
Faculty	0	1	1	0	0	0	1	0	1
	0.00%	100.00%	100.00%	0.00%	0.00%	0.00%	100.00%	0.00%	100.00%
Undergraduate students	0	0	0	0	0	0	0	0	0
	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Graduate students	0	0	0	0	0	0	0	0	0
	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Research scientists	0	0	0	0	0	0	0	0	0
	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Post Doctorates	0	0	0	0	0	0	0	0	0
	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Staff	0	0	0	0	0	0	0	0	0
	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Total	0	1	1	0	0	0	1	0	1
	0.00%	100.00%	100.00%	0.00%	0.00%	0.00%	100.00%	0.00%	100.00%

San Jose State University

Constituency	Gender		Race				US citizen		Total
	M	F	White	African American	Asian	Hispanic	Y	N	
Faculty	4	1	2	0	3	0	5	0	5
	80.00%	20.00%	40.00%	0.00%	60.00%	0.00%	100.00%	0.00%	100.00%
Undergraduate students	0	0	0	0	0	0	0	0	0
	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Graduate students	0	0	0	0	0	0	0	0	0
	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Research scientists	0	0	0	0	0	0	0	0	0
	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Post Doctorates	0	0	0	0	0	0	0	0	0
	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Staff	0	0	0	0	0	0	0	0	0
	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Total	4	1	2	0	3	0	5	0	5
	80.00%	20.00%	40.00%	0.00%	60.00%	0.00%	100.00%	0.00%	100.00%

Team for Research in Ubiquitous Secure Technology
2005-2006 Annual Report

Smith College

Constituency	Gender		Race				US citizen		Total
	M	F	White	African American	Asian	Hispanic	Y	N	
Faculty	1	0	1	0	0	0	1	0	1
Undergraduate students	0	0	0	0	0	0	0	0	0
Graduate students	0	0	0	0	0	0	0	0	0
Research scientists	0	0	0	0	0	0	0	0	0
Post Doctorates	0	0	0	0	0	0	0	0	0
Staff	0	0	0	0	0	0	0	0	0
Total	1	0	1	0	0	0	1	0	1

Stanford University

Constituency	Gender		Race				US citizen		Total
	M	F	White	African American	Asian	Hispanic	Y	N	
Faculty	6	0	5	0	0	1	6	0	6
Undergraduate students	0	0	0	0	0	0	0	0	0
Graduate students	7	0	2	2	3	0	3	4	7
Research scientists	0	0	0	0	0	0	0	0	0
Post Doctorates	1	0	0	0	1	0	0	1	1
Staff	0	0	0	0	0	0	0	0	0
Total	14	0	7	2	4	1	9	5	14

Vanderbilt University

Constituency	Gender		Race				US citizen		Total
	M	F	White	African American	Asian	Hispanic	Y	N	
Faculty	6	0	3	1	2	0	6	0	6
Undergraduate students	0	0	0	0	0	0	0	0	0

Team for Research in Ubiquitous Secure Technology
2005-2006 Annual Report

Graduate students	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
	5	0	3	0	2	0	2	3	5
Research scientists	100.00%	0.00%	60.00%	0.00%	40.00%	0.00%	40.00%	60.00%	41.67%
	1	0	1	0	0	0	1	0	1
Post Doctorates	100.00%	0.00%	100.00%	0.00%	0.00%	0.00%	100.00%	0.00%	8.33%
	0	0	0	0	0	0	0	0	0
Staff	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
	0	0	0	0	0	0	0	0	0
	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Total	12	0	7	1	4	0	9	3	12
	100.00%	0.00%	58.33%	8.33%	33.33%	0.00%	75.00%	25.00%	100.00%

6.3. Current and Anticipated Problems for the Center's Diversity Goals

Discuss any problems the Center has encountered in making progress toward its diversity goals during the reporting period as well as any problems anticipated in the next period. Include in the annual Implementation Plan the Center's strategy for addressing these problems.

No problems are anticipated.

6.4. Contributions

Describe and discuss Center contributions to the development of United States human resources in science and engineering at the postdoctoral, graduate, undergraduate, and pre-college levels (if applicable), with particular attention to accomplishments and activities that aim to attract, increase, and retain the participation of U.S. citizens, nationals, or lawfully admitted permanent residents of the United States, women, and underrepresented groups.

While we have given ourselves the ambitious goal of having 30% of women (faculty and students) and 10% under represented researchers in our center and considering that this is the first year report, we can say that we have been proactive in recruitment in this regard. Since the recruitment for new faculty and students is not finished, it is hard to put numbers out at this moment, but general indications seem to be positive.

In the following, we list the new efforts that we have made towards this goal:

1. We have increased our commitments to support underrepresented undergraduate summer students at all our sites (UCB:4; CMU: 5; Vanderbilt: 8)
2. We have increased our commitment to BFOIT (nurturing underrepresented high schools students and their teachers in engineering with focus on TRUST agenda) both financially as well via active participation of Professor Bajcsy.
3. We have started to actively participate in National conferences for underrepresented faculty and students. Dr. W.Robinson from Vanderbilt University attended the NSF Joint Annual Meeting HER, on March 16-17th, 2006 in Washington, DC.,see <http://www.edjassociates.com/jam06>
4. Meltem Errol from UCB attended HBCU conference in February, 2006 in Baltimore, Md. See: <http://www.hbcu-upconference.com/> Both these conference were used to advertise the TRUST agenda.

Team for Research in Ubiquitous Secure Technology
2005-2006 Annual Report

5. As planned, we are organizing a summer Institute one week long at UCB, called WISE which has registered 20 participants out of which 19 are women (graduate students and junior faculty).
6. We are starting to increase our visibility amongst underrepresented faculty and students. Concretely, Stanford will host this summer professor Mario Garcia from Texas A&M University –Corpus Christi. This visit is sponsored by NSF Quality Education for Minorities (QEM) Program, see: <http://qemnetwork.qem.org/STC.htm> .
7. Smith College and Mills College are to be active participants in TRUST. Judy Cardell from Smith College will participate in the research of the TRUST Sensor Networking project.
8. Last summer, Carnegie Mellon University and Professor Reiter of TRUST organized the Information Assurance Capacity Building Program (IACPB) with participation from Professor Weider Yu from San Jose State University (SJSU)
9. As mentioned in the section on Education, SJSU together with Mills college are using the initial learning material and testing it in their classes. At SJSU the majority of students are underrepresented. Mills College is a women’s college. Currently Professor Sigurd Meldal from SJSU is trying to propagate this material across all the California State universities.
10. Professor Bajcsy together with Professor Nahrsted from UIUC, Professor Wymur (UCB) and Professor Katherine Mezure from Mills college are building a cyber infrastructure for distributed dance performances in the Cyberspace and testing the issues of privacy. It is worth mentioning that all the PIs and most of their students are women.
11. We are engaged in continuous efforts of fundraising that would facilitate to increase and extend our outreach efforts. Professors Yelick, Graham and Bajcsy (UCB) have applied for NSF grant in the program Broadening participation, we did not make it but we plan to reapply. We are also participating in preparation of a proposal for Cyber infrastructure Team utilizing the Tele-immersive infrastructure.

6.5. Future Plans for Enhancing Diversity

Describe your plans for programs or activities to enhance diversity for the next reporting period with attention to any major changes in direction or level of activity. Discuss the impact of these programs or activities on enhancing diversity at the Center. Discuss how the planned activities will enable the Center to meet its goals.

While we can say that we need to stay the course because the effort of enhancing diversity is a long term activity, we are also planning some new efforts to maximize our effectiveness. The most natural effort in this direction is to increase our recruiting efforts of diverse population engaged in the TRUST agenda.

We shall try to increase the acceptance of underrepresented students into our graduate programs (currently the acceptance rate is approximately 50% across the various campuses). We shall actively recruit faculty from the available pool of underrepresented recently graduated students, by attending National conferences of HBCU.

We shall also make an active plan to visit certain departments from the HBCU to give colloquia

Team for Research in Ubiquitous Secure Technology
2005-2006 Annual Report

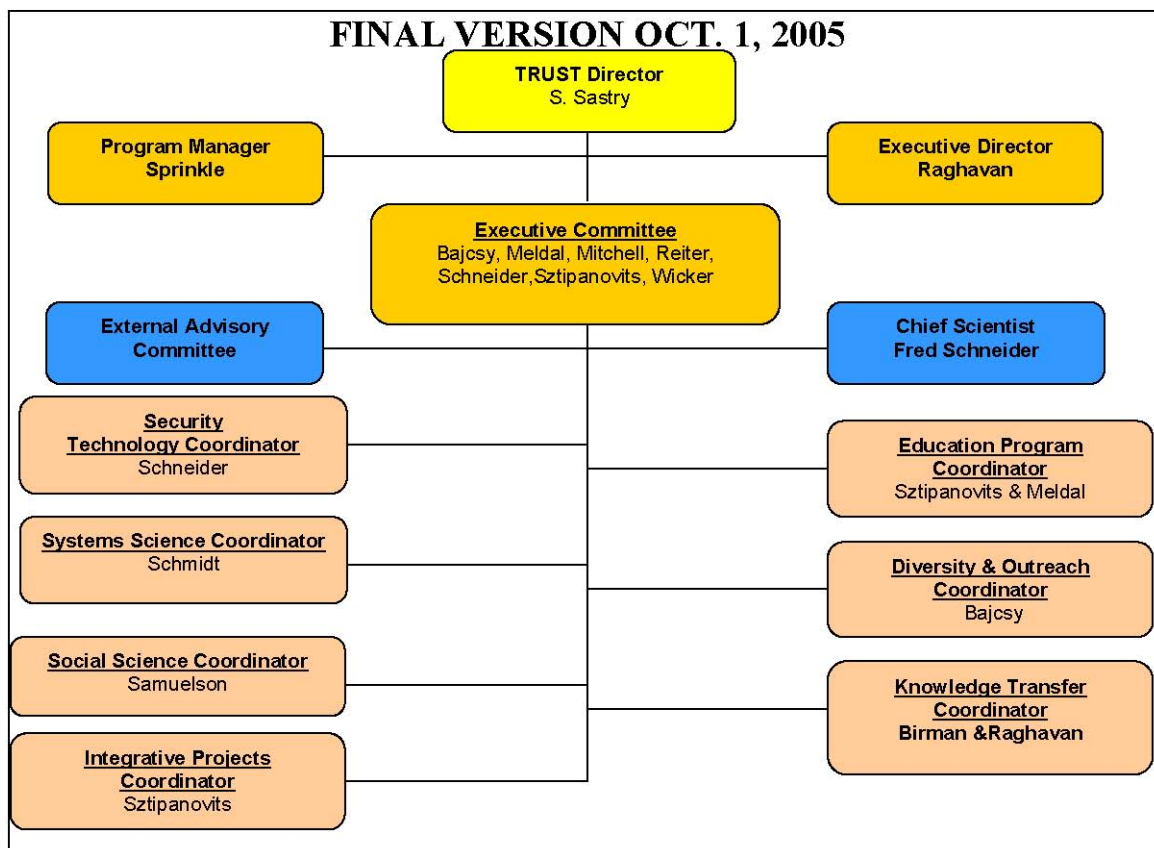
describing our work in the TRUST center. Following up on the WISE summer school of TRUST, we plan to have another summer school in 2007 on the East coast. Finally, we will make an extra effort to engage in our research more women and underrepresented students, such as in the Tele-immersive project, this time connecting with researchers at Cornell and Vanderbilt.

7. MANAGEMENT

7.1. Organizational Strategy

Describe the Center's organizational strategy and underlying rationale. To assist in your description, attach the organization chart of the Center during the reporting period. If there have been any changes in the Center's organization or management since the last reporting period, discuss these changes and the reasons behind them. What performance and management indicators has the Center developed to assess progress in organizational and management objectives?

The TRUST organization chart is shown below, The Executive Board which consists of the Center Director, the TRUST PIs, the Chief scientist, the Executive Director, and the Program Manager manage and execute the overall administration of the Center.



A top-down approach is used to review and create project areas, and a bottom-up approach is used to solicit broadly to populate project areas with research ideas from TRUST at large. Then, resource allocation is done along both dimensions of the matrix structure of TRUST—i.e., along the project areas as well as along the institutions.

Projects are reviewed by designated project leaders continuously as well as semi-annually during TRUST retreats.

The TRUST management depends on several key processes and agreements for its functioning.

First, the TRUST Center By-Laws (Appendix B3.1 of the Strategic Plan) govern the operation of the center. This was drafted and accepted into practice in the early part of the first year itself.

Second, the administration and allocation of funds from industrial partners are governed by a TRUST memorandum (Appendix B3.2), accepted into practice during the Washington DC Winter Retreat of TRUST in 2005, and is subject to the Industrial Partnership Agreement (Appendix B3.3).

Finally, the redefining of TRUST project areas and funds allocation are governed by the guidelines (Appendix B3.4) and process of funds allocation (Appendix B3.5).

Review of research, education, and outreach activities are done according to the metrics specified in the Strategic Plan, which include scientific and societal impact as well as center-wide emphasis on integration.

7.2. Current and Anticipated Management Problems

Discuss any problems (e.g., technical, personnel, communication) you may have encountered in realizing the Center's organizational strategy or management objectives in the reporting period as well as any problems anticipated in the next period. Include plans for addressing any problems.

No significant problems have been encountered.

7.3. Management and Communication Systems

Describe and discuss the management and communications systems used to develop a fully integrated STC as well as any problems encountered in achieving this integration.

The Trust website at <http://trust.eecs.berkeley.edu> is a comprehensive and secure website that provides email lists, collaborative workspaces, conference registration and access to publications and presentations. Industrial, governmental and academic participants have individual accounts and membership in multiple workspaces. Email lists and Usenet style newsgroups are gatewayed to each other providing easy access to discussion threads. Email is archived and searchable. Resources such as workgroups and publications have fine grained access control. The website provides workgroup web pages via participant supplied html and Wiki pages. The website is based on preexisting code developed for other projects. There have been no problems with the website.

In addition to these electronic systems, the TRUST Executive Board has a standing meeting every month to discuss the current status of projects, funding and other resource allocation, and other management issues. Ad hoc meetings are also arranged in addition to these regularly scheduled meetings, and several workshops have been organized around specific project areas.

7.4. External Advisory Board

Provide a list of names and affiliations of the Center's internal and external advisors or advisory bodies in the reporting period. Attach summary minutes of advisory committee minutes.

The current External Advisory Board consists of leaders from industry, academia, and government research labs, as indicated below.

Team for Research in Ubiquitous Secure Technology
2005-2006 Annual Report

Institution	Contact	Title
CAL-IT 2	Larry Smarr	Director, CAL IT 2
Cisco Systems	Ken Watson	Senior Manager, Critical Infrastructure Assurance Group
HP Laboratories	Wayne Johnson	Vice President University Relations Worldwide
IBM	Robert Morris	VP, Personal Systems & Storage & Director, IBM Almaden Research Center
Infineon Technologies.	Ulrich Ramacher	Professor
Intel	Andrew Chien	Professor, UC San Diego & VP, Director of Intel Research
Microsoft	Daniel Ling	Corporate VP, Microsoft Research
Nortel Networks	Phil Edholm	CTO, VP Network Architecture
Pirelli	Pieroguido Iezzi	EVP, Security
Qualcomm	John W Noerenberg	Principal Engineer, Consumer Products
SRI	William Mark	VP, Information and Computing Sciences (ICS) Division
Sun Microsystems	Emil Sarpa	Manager, External Research
Symantec	Steve Trilling	Acting Assoc. Dir. For Homeland Security Directorate
Telecom Italia	Giovanni Penna	Group Senior Vice President
UTRC	Jean Colpin	Director, Systems Development
Oak Ridge National Labs	Brian Worley	Director, Computational Science and Engineering
Cornell University	Don Greenberg	Professor, Director Cornell Program of Graphics

The External Advisory Board meeting is to take place on April 26, 2006. The minutes will be included in the final version of the Annual Report.

7.5. Changes to the Strategic Plan

Describe and discuss any changes to the Center's strategic plan since its last submission.

The revisions to the Strategic Plan for the second year (Research, Education, and Management) are detailed in Appendix B of the Strategic Plan.

8. CENTER-WIDE OUTPUTS AND ISSUES

8.1. Publications

List all Center publications in the reporting period using a standard citation format.

- V. Anantharam, "A technique to study the correlation measures of binary sequences," Submitted to "Discrete Mathematics," November 2005.
- V. Anantharam and V. Borkar, "Common randomness and distributed control; a counterexample," Submitted to *Systems and Control Letters*, September 2005.
- M. Balakrishnan and K. Birman, "Reliable Multicast for Time-Critical Systems," *1st Workshop on Applied Software Reliability*, June, 2006.
- M. Balakrishnan, K. Birman, A. Phanishayee, and S. Pleisch, "Ricochet: Low-Latency Multicast for Scalable Time-Critical Services," In Submission, 2005.
- M. Balakrishnan, S. Pleisch, and K. Birman, "Slingshot: Time-Critical Multicast for Clustered Applications," In *Proc. IEEE Network Computing and Applications 2005 (NCA 05)*. Boston, MA.
- K. Balasubramanian, A. Gokhale, G. Karsai, J. Sztipanovits, and S. Neema, "Developing Applications Using Model-Driven Design Environments," *IEEE Computer*, 39(2):33-40, February 2006.
- K. Balasubramanian, J. Balasubramanian, J. Parsons, A. Gokhale, and D.C. Schmidt, "A Platform-Independent Component Modeling Language for Distributed Real-time and Embedded Systems," *Elsevier Journal of Computer and System Sciences*, 2006.
- K. Balasubramanian, A.S. Krishna, E. Turkay, J. Balasubramanian, J. Parsons, A. Gokhale, and D.C. Schmidt, "Applying Model-Driven Development to Distributed Real-time and Embedded Avionics Systems," *The International Journal of Embedded Systems, Special Issue: Design and Verification of Real-Time Embedded Software*, April 2005.
- R. Bajcsy, K. Nahrstedt, and L. Wymore, "Humans in Real and Virtual Space: Studies of Interaction and Collaboration," *Mediated by IT, Special Issue: Women in Robotics*, 2006.
- A. Barth, D. Boneh, and B. Waters, "Private encrypted content distribution using private broadcast encryption," In *Proc. Financial Crypto (FC) '06*, 2006.
- K. Birman, "Can Web Services Scale Up?," *IEEE Computer*. Volume 38. Number 10. Pgs.107-110. October 2005.
- K. Birman, C. Chandrasekaran, D. Dolev, and R. van Renesse, "How the Hidden Hand Shapes the Market for Software Reliability," Submitted to the First IEEE Workshop on Applied Software Reliability, June 2006.
- K. Birman, "Reliable Distributed Systems Technologies, Web Services, and Applications," 2005, XXXVI, 668 p. 145 illus., Hardcover ISBN: 0-387-21509-3.

Team for Research in Ubiquitous Secure Technology
2005-2006 Annual Report

- K. Birman, "The Untrustworthy Services Revolution," *IEEE Computer* (ISSN 0018-9162). Vol.39 No.2, Pgs. 98-100. February 2006.
- D. Boneh, E. Shen, and B. Waters, "Strongly Unforgeable Signatures Based on Computational Diffie-Hellman," In *Proc. of PKC '06, LNCS 3958*, pp. 229-240, 2006.
- E. Brand, P. Walsh, J. Hall, and D.K. Mulligan. *ACCURATE (A Center for Correct, Usable, Reliable, Auditable and Transparent Elections), Public Comment on the 2005 Voluntary Voting System Guidelines Submitted to the United States Election Assistance Commission*. Technical report, ACCURATE, January, 2006.
- C. Clifton, D.K. Mulligan, and R. Ramakrishnan. *Privacy and Technologies of Identity: A Cross-Disciplinary Conversation*. Katherine Strandburg, Daniela Stan Raicu, 11, 191-208, Springer, 2006.
- C. Clifton, J. Werner, J. Mathe, and M. Eby, "Secure AADL in the Generic Modeling Environment," *ISIS Technical Report*. In preparation.
- L. Cranor, "Stopping Spyware at the Gate: A User Study of Privacy, Notice and Spyware," *SOUPS*, July, 2005.
- C. Gill, J.M. Gossett, D. Corman, J.P. Loyall, R.E. Schantz, M. Atighetchi, and D.C. Schmidt, "Integrated Adaptive QoS Management in Middleware: An Empirical Case Study," *The International Journal of Time-Critical Computing Systems*, Springer, Vol. 29, No. 2-3, pp. 101-130, March-April 2005.
- A. Gokhale, K. Balasubramanian, J. Balasubramanian, A. Krishna, G. T. Edwards, G. Deng, E. Turkay, J. Parsons, and D.C. Schmidt, "Model Driven Middleware: A New Paradigm for Deploying and Provisioning Distributed Real-time and Embedded Applications," *Elsevier Journal of Science of Computer Programming: Special Issue on Model Driven Architecture*, Edited by Mehmet Aksit, 2006.
- N. Good, J. Grossklags, D. Thaw, A. Perzanowski, D.K. Mulligan, and J. Konstan, "User Choices and Regret: Understanding Users' Decision Process about Consensually acquired Spyware," *I/S: A JOURNAL OF LAW AND POLICY FOR THE INFORMATION SOCIETY*, January 2006.
- T. He and L. Tong, "Detecting Encrypted Interactive Stepping-stone Connections," *IEEE Intl. Conf. Acoust. Speech and Signal Processing*, May 2006.
- T. He and L. Tong, "Detecting Encrypted Stepping-stone Connections," submitted to *IEEE Transactions on Signal Processing*, Feb. 01, 2006.
- T. He and L. Tong, "Robust Detection of Stepping-Stone Connections," submitted to 2006 *Recent Advances in Intrusion Detection (RAID) Symposium*, April 2006.
- T. He and L. Tong, "A Signal Processing Prospective to Stepping-stone Detection," *Conference on Information Sciences and Systems 2006 (CISS'06)*, Princeton, NJ, March, 2006.
- K.M.Hopkinson, R. Giovanini, X. Wang, K.P. Birman, D.V. Coury, J.S. Thorp, "EPOCHS:

Team for Research in Ubiquitous Secure Technology
2005-2006 Annual Report

Integrated Cots Software For Agent-Based Electric Power And Communication Simulation,"
To appear *IEEE Transactions on Power Systems*.

- H. Inaltekin, S. Wicker, "A One Shot Random Access Game for Wireless Networks - Behavior of Nodes at Nash Equilibria," Submitted to *IEEE/ACM Trans. on Networking*, 2005.
- C. Jackson, A. Bortz, D. Boneh, and J. Mitchell, "Protecting Browser State from Web Privacy Attacks," In *Proc. of WWW '06*, 2006.
- S. Jung and R. Bajcsy, "Learning Physical Activities in Immersive Virtual Environments," In *Proc. 4th IEEE Int. Conference on Computer Vision Systems*, New York City, January 5-7th, 2006.
- A.S. Krishna, N. Wang, B. Natarajan, A. Gokhale, D.C. Schmidt, and G.Thaker, "CCMPerf: A Benchmarking Tool for CORBA Component Model Implementations," *The International Journal of Time-Critical Computing Systems*, Springer, Vol. 29, No. 2-3, pp. 281-308, March-April 2005.
- A.S. Krishna, A. Gokhale, D.C. Schmidt, J. Hatcliff, and V.P. Ranganat, "Towards Highly Optimized Real-time Middleware for Software Product-line Architectures," *SIGBED Review*, Volume 3, No. 1, January 2006.
- D. Kostoulas, D. Psaltoulis, I. Gupta, K. Birman, and A. Demers, "Decentralized Schemes for Size Estimation in Large and Dynamic Groups," *IEEE Network Computing and Applications 2005 (NCA 05)*. October, 2005, Boston, MA.
- J. Liang, Z. Yang, B. Yu, Y. Cui, K. Nahrstedt, S. Jung, A. Yeap, and R. Bajcsy, "Experience with Multi-Camera Tele-Immersive Environment," *NSF Workshop on Pervasive Computing Experience*, Urbana Champaign, 2005.
- H. Liu, V. Ramasubramanian, and E.G. Sirer, "Client and Feed Characteristics of RSS, A Publish-Subscribe System for Web Micronews," In *Proc. Internet Measurement Conference (IMC)*, Berkeley, California, October, 2005.
- H. Liu, T. Roeder, K. Walsh, R. Barr, and E.G. Sirer, "Design and Implementation of a Single System Image Operating System," In *Proc. The International Conference on Mobile Systems, Applications, and Services (Mobisys)*, Seattle, Washington, June, 2005.
- T. Marian, K. Birman, and R. van Renesse, "A Scalable Services Architecture," In submission.
- G. Madl, S. Abdelwahed, and D.C. Schmidt, "Verifying Distributed Real-time Properties of Embedded Systems via Graph Transformations and Model Checking," *The International Journal of Time-Critical Computing Systems*, 2006.
- V. Naware and L. Tong, "Cross Layer Design for Multiaccess Communication over Rayleigh Fading Channels," *IEEE Trans. Wireless Communications*, In Submission, February 2006.
- S. Oh, P. Chen, M. Manzo, and S. Sastry, "Instrumenting Wireless Sensor Networks for Real-time Surveillance," In *Proc. International Conference on Robotics and Automation*, May, 2006.

Team for Research in Ubiquitous Secure Technology
2005-2006 Annual Report

- S. Oh, I. Hwang, K. Roy, and S. Sastry, "A Fully Automated Distributed Multiple-Target Tracking and Identity Management Algorithm," *AIAA Guidance, Navigation, and Control Conference*, August, 2005.
- S. Oh and S. Sastry, "An Efficient Algorithm for Tracking Multiple Maneuvering Targets," In *Proc. IEEE International Conference on Decision and Control*, December, 2005.
- K. Ostrowski and K. Birman, "Extensible Web Services Architecture for Notification in Large-Scale Systems," *International Conference on Web Services (ICWS 2006)*, IEEE, September, 2006.
- K. Ostrowski, K. Birman, and A. Phanishayee, "The Power of Indirection: Achieving Multicast Scalability by Mapping Groups to Regional Underlays," In submission.
- B. Pfaff, T. Garfinkel, and M. Rosenblum, "Virtualization Aware File Systems: Getting Beyond the Limitations of Virtual Disks," *3rd Symposium of Networked Systems Design and Implementation (NSDI)*, May, 2006.
- S. Pleisch, M. Balakrishnan, K. Birman, and R. van Renesse, "Mistral: Efficient Flooding in Mobile Ad-hoc Networks," To appear in *The Seventh ACM International Symposium on Mobile Ad Hoc Networking and Computing (ACM MobiHoc 2006)*. Florence, Italy May 2006.
- R. Pucella and F. Schneider, "*Independence From Obfuscation: A Semantic Framework for Diversity*," Technical report, Cornell University, TR2006-2016, January, 2006.
- V. Ramasubramanian, R. Peterson, and E.G. Sirer, "Corona: A High Performance Publish-Subscribe System for the World Wide Web," In *Proc. Networked System Design and Implementation*, San Jose, California, May, 2006.
- V. Ramasubramanian and E.G. Sirer, "Perils of Transitive Trust in the Domain Name System," In *Proc. Internet Measurement Conference (IMC)*, Berkeley, California, October, 2005.
- R. van Renesse and K. Birman, "Autonomic Computing - A System-Wide Perspective," *Autonomic Computing: Concepts, Infrastructure, and Applications*, ed. Manish Parashar and Salim Hariri, CRC press, 2006.
- R. van Renesse and H. Johansen, "Fireflies: Scalable Support for Intrusion-Tolerant Overlay Networks," In *Proc. EuroSys 2006*, Leuven, Belgium, April 2006.
- R van Renesse, "Using Randomized Techniques to Build Scalable Intrusion-Tolerant Overlay Networks," In *Proc. Workshop on Stochasticity in Distributed Systems (StoDis 2005)*, December 2005, San Jose, CA.
- B. Ross, C. Jackson, N. Miyake, D. Boneh, and J. Mitchell, "Stronger Password Authentication Using Browser Extensions," In *Proc. Usenix Security*, 2005.
- P. Samar and S. B. Wicker, "Link Dynamics and Protocol Design in a Multi-Hop Mobile Environment." *IEEE Transactions on Mobile Computing*, To appear in Spring 2006.
- P. Samar and S.B. Wicker, "On the Behavior of Communication Links in a Multi-Hop Mobile Environment." *Frontiers in Distributed Sensor Networks*, edited by S.S. Iyengar and R.R.

Team for Research in Ubiquitous Secure Technology
2005-2006 Annual Report

Brooks, CRC Press, 2005.

- R.E. Schantz, D.C. Schmidt, J.P. Loyall, and C. Rodrigues, "Controlling Quality-of-Service in Distributed Real-time and Embedded Systems via Adaptive Middleware," *The Wiley Software Practice and Experience Journal special issue on Experiences with Auto-adaptive and Reconfigurable Systems*, co-edited by Mehmet Aksit, Zied Choukair, and Tzilla Elrad, 2006.
- D.C. Schmidt, "Model-Driven Engineering," *IEEE Computer*, Vol. 39. No. 2, February 2006, pp. 41-47.
- F. Schneider, K. Hamlen, and G. Morrisett, "Computability classes for enforcement mechanisms," *TOPLAS* 28, 1 (January 2006), 175--205.
- F. Schneider and L. Zhou, "Implementing Trustworthy Services Using Replicated StateMachines," *IEEE Security and Privacy*, Volume 3, Number 5 (September/October 2005), 34--43.
- F. Schneider, L. Zhou, and R. van Renesse, "APSS: Proactive secret sharing in asynchronous systems," *ACM Transactions on Information and System Security* 8, 3 (August 2005), 259--286.
- A. Schwartz, D.K. Mulligan, and I. Monda, "Storing Our Lives Online: Expanded Email Storage Raises Complex Policy Issues," *I/S: A JOURNAL OF LAW AND POLICY FOR THE INFORMATION SOCIETY*, January 2005.
- E.G. Sizer, "Heuristics Considered Harmful: Using Mathematical Optimization for Resource Management in Distributed Systems," *IEEE Intelligent Systems*, Special Issue on Self-Management through Self-Organization in Information Systems, Mar/Apr 2006.
- Y.J. Song, V. Ramasubramanian and E.G. Sizer, "*Optimal Resource Utilization in Content Distribution Networks*," Technical report, Cornell University, TR2005-2004, November, 2005.
- P.A. Subrahmanyam, D. Wagner, U. Shankar, D.K. Mulligan, E. Jones, and J. Lerner, "Network Security Architecture for Demand Response/Sensor Networks," Technical report, On behalf of California Energy Commission, Public Interest Energy Research Group, January, 2005.
- V. Subramonian, G. Deng, C. Gill, J. Balasubramanian, L. Shen, W. Otte, D.C. Schmidt, A. Gokhale, and N. Wang, "The Design and Performance of Component Middleware for QoS-enabled Deployment and Conguration of DRE Systems," *Elsevier Journal of Systems and Software*, Special Issue on Component-Based Software Engineering of Trustworthy Embedded Systems, 2006.
- Y. Sung, S. Misra, L. Tong and A. Ephremides, "Cooperative Routing for Signal Detection in Large Sensor Networks," *IEEE JSAC*, Special Issue on Cooperative Communications and Networking, Submitted in February, 2006.
- A. Wagner and V. Anantharam, "Information Theory of Covert Timing Channels," Unpublished article, October, 2005; Document based on talks presented at the *2005 NATO/ASI Workshop on Network Security and Intrusion Detection*, Yerevan, Armenia, October 2005.

Team for Research in Ubiquitous Secure Technology
2005-2006 Annual Report

- K. Walsh and E.G. Sirer, "Evaluation of a Deployed, Distributed Object Reputation System for Peer-to-Peer Filesharing," In *Proc. Networked System Design and Implementation*, San Jose, California, May, 2006.
- K. Walsh, E.G. Sirer, "Fighting Peer-to-Peer SPAM and Decoys with Object Reputation," In *Proc. P2PECON Workshop*, Philadelphia, Pennsylvania, August, 2005.
- J. Werner, M. Eby, J. Mathe, G. Karsai, Y. Xue, and J. Sztipanovits, "Integrating Security Modeling in Embedded System Design," *IEEE Real-Time and Embedded Technology and Applications Symposium*, April, 2006.
- B. Wong and E.G. Sirer, "ClosestNode.com: An Open-Access, Scalable, Shared Geocast Service for Distributed Systems," In *SIGOPS Operating Systems Review*, 40(1), January 2006.
- B. Wong, A. Slivkins, E.G. Sirer, "Meridian: A Lightweight Network Location Service without Virtual," In *Proc. The ACM SIGCOMM Conference*, Philadelphia, Pennsylvania, August, 2005.
- T. Wu, Y. Xue, and Y. Cui, "Preserving Traffic Privacy in Wireless Mesh Networks," In *Proc. of WOWMOM 2006*, IEEE, June, 2006.
- Z. Yang, J. Liang, B. Yu, Y. Cui, K. Nahrstedt, S. Jung, and R. Bajcsy, "TEEVE: Tele-immersive Environment for Everybody," In *Proc. of IEEE International Symposium of Multimedia*, Irvine, California, December 12-14th, 2005.
- Xin Zhang, Jun Chen, Stephen B. Wicker, and Toby Berger, "Successive Coding in Multiuser Information Theory," Accepted pending revision, *IEEE Trans. on Information Theory*.
- X. Zhang and S. B. Wicker, "On the Optimal Distribution of Sensors in a Random Field," *ACM Transactions on Sensor Networks*, March 2006.
- L. Zhou, F. Schneider, and R. van Renesse, "APSS: Proactive Secret Sharing in Asynchronous Systems," *ACM Transactions on Information and System Security (TISSEC)*. Vol 8, Issue 3, August 2005.
- L. Zhuang, F. Zhou, J.D. Tygar, "Keyboard Acoustic Emanations Revisited," In *Proc. 12th ACM Conference on Computer and Communications Security*, November 2005.

8.2. Presentations

List all conference presentations in the reporting period using a standard citation format.

- F. Akopyan, R. Manohar, and A. Apsel, "A level-crossing Flash Asynchronous Analog-to-Digital Converter." In *Proc. 12th International Symposium on Asynchronous Circuits and Systems*, March 2006.
- Dzintars Avots, Michael Dalton, Benjamin Livshits, Monica S. Lam, Improving Software Security with a C Pointer Analysis. In *Proceedings of the 27th International Conference on Software Engineering (ICSE)*, May 2005

Team for Research in Ubiquitous Secure Technology
2005-2006 Annual Report

Mahesh Balakrishnan, Stefan Pleisch, Ken Birman. Slingshot: Time-Critical Multicast for Clustered Applications. Pleisch, Ken Birman. IEEE Network Computing and Applications 2005 (NCA 05). Boston, MA

Mahesh Balakrishnan and Ken Birman. Reliable Multicast for Time-Critical Systems. Submission to: 1st Workshop on Applied Software Reliability, June 2006.

Ken Birman. "Air Force Center for Research on GIG/NCES Challenges (AF-TRUST-GNC)." Talk or presentation, 10, January, 2006.

Ken Birman, Coimbatore Chandrasekaran, Danny Dolev, and Robbert van Renesse. How the Hidden Hand Shapes the Market for Software Reliability. Submitted to the First IEEE Workshop on Applied Software Reliability, June 2006.

Stephen Chong, Andrew C. Myers.

<<http://www.cs.cornell.edu/andru/papers/erasure.pdf>> Language-Based Information Erasure
<<http://www.cs.cornell.edu/andru/papers/erasure.pdf>> Proceedings of the 18th IEEE Computer Security Foundations Workshop <<http://www.lif.univ-mrs.fr/%7Eamadio/CSFW18/>> (CSFW'05), pages 241–254, June 2005.

Michael Clarkson, Andrew C. Myers, Fred B. Schneider.

<<http://www.cs.cornell.edu/andru/papers/InfoFlowBelief.pdf>> Belief in Information Flow
<<http://www.cs.cornell.edu/andru/papers/InfoFlowBelief.pdf>> Proceedings of the 18th IEEE Computer Security Foundations Workshop <<http://www.lif.univ-mrs.fr/%7Eamadio/CSFW18/>> (CSFW'05), pages 31–45, June 2005.

George Cybenko. "Process Detection in Secure and Reliable Computing." Talk or presentation, 20, October, 2005.

O. Dousse, C. Tavoularis and P. Thiran, "Delay of intrusion detection in sensor networks", Mobihoc'06.

Simson Garfinkle. "Johnny 2: A User Test of Key Continuity Management with S/MIME and Outlook ExpressPPT." Talk or presentation, 9, September, 2005.

Ken Goldberg and Deirdre Mulligan. "Too Close For Comfort: Free Speech, Privacy, and the Demonstrate Project." Talk or presentation, 15, September, 2005.

T. He and L. Tong, "Detecting Encrypted Interactive Stepping-stone Connections," IEEE Intl. Conf. Acoust. Speech and Signal Processing, May 2006.

T. He and L. Tong, "A Signal Processing Prospective to Stepping-stone Detection," Conference on Information Sciences and Systems 2006 (CISS'06), Princeton, NJ, March, 2006.

H. Inaltekin, and S.B. Wicker, A one shot random access game for wireless networks Symposium on Information Theory in Wirelesscom, 2005.

H. Inaltekin, and S.B. Wicker, The Analysis of a Game Theoretic MAC Protocol for Wireless Networks submitted to SECON'06, 2006.

H. Inaltekin, C. Tavoularis, S.B. Wicker, Coverage Analysis of Mobile Sensor Networks to be

Team for Research in Ubiquitous Secure Technology
2005-2006 Annual Report

submitted to INFOCOM'07, under preparation.

H. Inaltekin, S.B. Wicker, Interference Analysis for Wireless Networks to be submitted to INFOCOM'07, under preparation.

Hemant Jain. "Preventing rate-based attacks: Requirements, Architecture for a solution and Lessons from Field-Trials." Talk or presentation, 1, December, 2005.

Erin Jones. "Privacy and Security in Demand Response Energy Systems." Talk or presentation, 21, March, 2006.

Anthony Joseph. "Trust Winter Meeting Outbrief: Network Defenses." Talk or presentation, 9, January, 2006.

Dionysios Kostoulas, Dimitrios Psaltoulis, Indranil Gupta, Ken Birman, Al Demers. Decentralized Schemes for Size Estimation in Large and Dynamic Groups. IEEE Network Computing and Applications 2005 (NCA 05). October, 2005, Boston, MA.

Monica S. Lam, John Whaley, Benjamin Livshits, Michael Martin, Dzintars Avots, Michael Carbin, Christopher Unkel, Context-Sensitive Program Analysis as Database Queries. In Proceedings of Principles of Database Systems (PODS), Baltimore, Maryland, June 2005.

D.T. Lee. "TWISC—Taiwan Information Security Center." Talk or presentation, 10, January, 2006.

Ulf Lindqvist. "Securing Control Systems in the Oil and Gas Infrastructure: The I3P SCADA Security Research Project." Talk or presentation, 17, November, 2005.

Benjamin Livshits, Defining a Set of Common Benchmarks for Web Application Security. Position paper on Stanford SecuriBench for the Workshop on Defining the State of the Art in Software Security Tools, Baltimore, August 2005.

Benjamin Livshits and Monica S. Lam, Finding Security Vulnerabilities in Java Applications with Static Analysis. In Proceedings of the Usenix Security Symposium, Baltimore, Maryland, August 2005.

Michael Martin, Benjamin Livshits, and Monica S. Lam, Finding Application Errors and Security Flaws Using PQL: a Program Query Language. Presented at the 20th Annual ACM Conference on Object-Oriented Programming, Systems, Languages, and Applications, San Diego, California, October 2005.

Jonathan M. McCune, Adrian Perrig, and Michael K. Reiter, Bump in the Ether: A Framework for Securing Sensitive User Input, in Proceedings of Usenix Annual Technical Conference, June 2006.

Sigurd Meldal, Janos Sztipanovits, Ruzena Bajcsy. "Education and Outreach." Talk or presentation, 10, January, 2006.

John Mitchell. "Combating Online Identity Theft: Spoofguard, PwdHash, Spyware, Botnets." Talk or presentation, 10, January, 2006.

Team for Research in Ubiquitous Secure Technology
2005-2006 Annual Report

- John Mitchell. "Trust Winter Meeting Outbrief: Phishing, Identity Theft, and Related Issues." Talk or presentation, 9, January, 2006.
- Bryan Parno and Cynthia Kuo and Adrian Perrig, Phoolproof Phishing Prevention, In Proceedings of International Conference on Financial Cryptography and Data Security, February 2006.
- Ben Pfaff, Tal Garfinkel, Mendel Rosenblum. Virtualization Aware File Systems: Getting Beyond the Limitations of Virtual Disks. 3rd Symposium of Networked Systems Design and Implementation (NSDI), May, 2006.
- S. Pleisch, M. Balakrishnan, K. Birman, and R. van Renesse. Mistral: Efficient Flooding in Mobile Ad-hoc Networks. To Appear: The Seventh ACM International Symposium on Mobile Ad Hoc Networking and Computing (ACM MobiHoc 2006). Florence, Italy May 2006.
- Mike Reiter. "Trust Winter Meeting Outbrief: Trustworthy Systems." Talk or presentation, 9, January, 2006.
- Hui Qu and Stephen Wicker, "Anchor-Free Localization for Wireless Sensor Networks", submitted to the IEEE Global Communication Conference, 2006.
- Hui Qu and Stephen Wicker, "A Combined Localization and Location-Based Routing Algorithm Design for Wireless Sensor Networks," submitted to ACM SECON, 2006.
- Robbert van Renesse and Havard Johansen. Fireflies: Scalable Support for Intrusion-Tolerant Overlay Networks. EuroSys 2006, Leuven, Belgium, April 2006.
- Robbert van Renesse. Using Randomized Techniques to Build Scalable Intrusion-Tolerant Overlay Networks. StoDis 2005 -- Workshop on Stochasticity in Distributed Systems, December 2005, San Jose, CA.
- Shankar Sastry. "TRUST: Team for Research in Ubiquitous Secure Technologies, an overview." Talk or presentation, 10, January, 2006.
- Shankar Sastry. "TRUST: Team for Research in Ubiquitous Secure Technologies." Talk or presentation, 1, September, 2005.
- Mike Schiffman. "The Common Vulnerability Scoring System (CVSS)." Talk or presentation, 10, November, 2005.
- Doug Schmidt. "Trust Winter Meeting Outbrief: Software System Dependability Breakout Session." Talk or presentation, 9, January, 2006.
- Douglas Schmidt, Andy Gokhale, and Jaiganesh Balasubramanian: "Investigating Survivability Strategies for Enterprise Distributed Real-time & Embedded Systems," Poster presentation, NSF TRUST PI Meeting, Dec 2005
- Douglas Schmidt, Andy Gokhale, and Jaiganesh Balasubramanian: "Investigating Lightweight Fault Tolerance Strategies in Enterprise Distributed Real-time Embedded Systems," Submitted to OMG RTWS 2006, July 2006, Washington DC.

Team for Research in Ubiquitous Secure Technology
2005-2006 Annual Report

Douglas Schmidt, Andy Gokhale, and Jaiganesh Balasubramanian: "Model-Driven Engineering of Fault Tolerance in Enterprise Distributed Real-time and Embedded Systems," Submitted to OMG RTWS 2006, July 2006, Washington DC.

Douglas Schmidt, Andy Gokhale, and Jaiganesh Balasubramanian: "Dependability in Enterprise Distributed Real-time and Embedded Systems," To be submitted to Fast Abstracts at DSN 2006, Philadelphia, PA, June 2006.

Douglas Schmidt, Andy Gokhale, and Jaiganesh Balasubramanian: Paper on MDE approach for dependability to be submitted for GPCE 2006, Portland, OR, Oct 2006.

Douglas Schmidt, Andy Gokhale, and Jaiganesh Balasubramanian: Paper on Service Placement algorithms to be submitted to IEEE RTSS 2006, Rio De Janeiro, Brazil, Dec 2006.

Douglas Schmidt, Andy Gokhale, and Jaiganesh Balasubramanian: Paper on FT Middleware Solutions to be submitted to ICSOC 2006, Chicago, IL, Dec 2006.

Fred B. Schneider. "The TRUST Agenda: Convergence of Technical and Policy Issues." Talk or presentation, 10, January, 2006.

Umesh Shankar. "Better web browser privacy using automation." Talk or presentation, 7, March, 2006.

Eugene H. Spafford. "Testimony before the House Armed Services Committee Hearing on Cyber Security, Information Assurance and Information Superiorit." Talk or presentation, 27, October, 2005.

Terence Spies. "No More Alice to Bob: Reality-based Models for Message Encryption and Key Management." Talk or presentation, 29, September, 2005.

Mary Margaret Sprinkle. "TRUST:Team for Research in Ubiquitous Secure Technologies, an overview." Talk or presentation, 8, July, 2005.

Janos Sztipanovits. "Integrative Projects." Talk or presentation, 10, January, 2006.

Janos Sztipanovits. "Trust Winter Meeting Outbrief: Integrative Projects: Patient Portals." Talk or presentation, 9, January, 2006.

David Wagner. "Trust Winter Meeting Outbrief: Programming Language Techniques for Software Security." Talk or presentation, 9, January, 2006.

Jan Werner, Matthew Eby, Janos Mathe, Gabor Karsai, Yuan Xue, Janos Sztipanovits: Integrating Security Modeling in Embedded System Design, workshop presented at "Research Directions for Security and Networking in Critical Real-Time and Embedded Systems" of the 2006 IEEE Real-Time and Embedded Technology and Applications Symposium, San Jose, CA, 2006.

Stephen Wicker. "Securing Public Spaces with Sensor Networks: Science, Technology, and Privacy." Talk or presentation, 10, January, 2006.

Stephen Wicker. "Trust Winter Meeting Outbrief: Embedded Systems and Secure Sensor Networks." Talk or presentation, 9, January, 2006.

Junfeng Yang, Can Sar, Paul Twohey, Cristian Cadar and Dawson Engler, "Automatically Generating Malicious Disks using Symbolic Execution," to appear: IEEE Proceedings on Security and Privacy, 2006.

X. Zhang, S. B. Wicker "Robustness vs. Efficiency in Sensor Networks," Information Processing in Sensor Networks (ACM IPSN 2006), Berkeley, 2005.

Lantian Zheng, Andrew C. Myers. <<http://www.cs.cornell.edu/andru/papers/avail.pdf>> End-to-End Availability Policies and Noninterference <<http://www.cs.cornell.edu/andru/papers/avail.pdf>> Proceedings of the 18th IEEE Computer Security Foundations Workshop <<http://www.lif.univ-mrs.fr/%7Eamadio/CSFW18/>> (CSFW'05), pages 272–286, June 2005.

8.3. Other Dissemination Activities

Briefly describe any other dissemination activities not included above.

Ken Birman's group is distributing software Ricochet, with Quicksilver and Tempest to come soon:

Quicksilver: This work, funded in part by DARPA, AFRL and AFOSR, explores scalability for a publish-subscribe style of event notification platform, using peer-to-peer techniques and other methods. The platform is now operational and achieves a true breakthrough in scalability and performance; a series of papers are in preparation to discuss the mechanisms by which this was achieved. We are also extending Quicksilver with a strong type system and with a fault-tolerance and consistency model; these steps will offer an exceptionally flexible, robust and scalable framework within which type checking can play a role as part of a stronger security architecture. (Birman, PhD candidate Krzysztof Ostrowski)

Ricochet: This work was funded in part by Intel, DARPA, AFRL and AFOSR. Ricochet is a new protocol for time-critical data replication in clusters and data center computing platforms. It introduces the concept of lateral error correction and with it, demonstrates three orders of magnitude better delivery for use in settings requiring time-critical multicast or data updates. (Ricochet extends Slingshot to multigroup settings, and the work achieves far better scalability in the numbers of groups than in any prior work). (Birman, PhD students Mahesh Balakrishnan and Amar Phanishayee). We have begun to collaborate with Vanderbilt (Doug Schmidt) on aspects of this work.

Tempest: This work was funded in part by Intel, DARPA, AFRL and AFOSR. Tempest is a new platform that runs over Ricochet and automates most aspects of developing new scalable and robust services to run on data centers and clusters. Tempest provides automated data replication, query load-balancing, fault-tolerance and data repair after faults that introduce inconsistency. (Birman, PhD students Tudor Marian, Mahesh Balakrishnan and Amar Phanishayee). We have begun to collaborate with Vanderbilt (Doug Schmidt) on aspects of this work.

Adrian Perrig
IBM Faculty Fellowship, 2005

Team for Research in Ubiquitous Secure Technology
2005-2006 Annual Report

Carnegie Mellon University

Adrian Perrig
Sloan Faculty Fellowship, 2005
Carnegie Mellon University

F. Schneider,
Language-based security for malicious mobile code.
MURI Project Review.
Washington, DC. July 2005.

F. Schneider,
Implementing fault-tolerant and scalable storage services.
AFOSR PI Meeting.
Griffiss Institute, Rome, New York. August 2005.

F. Schneider,
Asynchronous proactive secret sharing.
Invited speaker.
DosCoVeri: Distributed Algorithms meet Concurrency Theory.
San Francisco, California. August 2005.

F. Schneider,
Progress Towards Trustworthy Services.
EECS Division Distinguished Lecture Series.
Ann Arbor, Michigan, November 2005.

F. Schneider,
Implementing Security and Fault-tolerance.
Keynote address, 2nd ITI Workshop on Dependability and Security.
Champaign, Ill.
December 2005.

F. Schneider,
Implementing Security and Fault-tolerance.
EECS Departmental Colloquium Distinguished Lecture Series.
U.C. Berkeley, Berkeley, California, December 2005.

F. Schneider,
The TRUST Agenda:
Convergence of Technical and Policy Issues.
TRUST 2006 Winter Meeting.
Washington, D.C. January 2006.

F. Schneider,
Cyber-terrorism: Yesterday, Today, and Tomorrow.
Cornell ALS 481 (Global Conflict and Terrorism).
Cornell University,
Ithaca, New York.
February 2006.

Team for Research in Ubiquitous Secure Technology
2005-2006 Annual Report

F. Schneider,
Non-Technical Impediments to Securing Cyberspace.
Symposium on Fostering International Collaborations in Information Security,
AAAS Annual Conference,
Saint Louis, Missouri.
February 2006.

Gun Sirer's CoDoNS system has been deployed in China, by CNNIC, the name registrar responsible for the .cn domain.

CobWeb cache has been deployed, which is an open access Akamai-like system for speeding up web browsing and protecting content providers from flash crowds, on PlanetLab, where it handles between 10-15 million requests per day.

Gun Sirer's Credence system was built for determining the trustworthiness of peers and for identifying pollution in large scale filesharing networks has been downloaded by over 10000 people.

Gun Sirer's Corona system has been deployed on PlanetLab and currently monitors a few hundred channels on behalf of its users.

Gun Sirer's Meridian system for locating nearby nodes has recently been deployed through a site called "closestnode.com".

8.4. Awards & Honors

List all awards and other honors with names of those honored and source in the reporting period.

F. Akopyan, R. Manohar, A. Apsel: Best paper award, ASYNC 2006.

Animashree Anandkumar, student paper contest finalists in IEEE Intl. Conf. Acoust. Speech and Sig. Proc. (ICASSP): "A Large Deviation analysis of detection over Multiaccess channels with random number of sensors."

Y.-W. Hong, B. Sirkeci-Mergen, A. Scaglione, R. Manohar: Best paper award (unclassified), MILCOM 2005.

R. Manohar: named to MIT Technology Review Magazine's TR35 (top 35 young innovators under 35).

L. Tong, Selected as Plenary Speaker, 2007 Signal Processing Advances in Wireless Communications, Intrusion Detection and Secure Wireless Transmission

P. Venkitasubramaniam for paper "Minimax Quantization for Distributed Estimation," Final selection is in May 2006.

8.5. TRUST's Class of 2006

List M.S. and Ph.D. students who graduated during the reporting period, with placements. Include the number of years taken since entering graduate school to complete the Ph.D. List postdoctoral associates who left the STC during the reporting period, with placements.

Name	Degree	Advisor	Graduation Date	Placement	No. of Years to Ph.D.
X. Zhang	ECE Ph.D.	Wicker	May 2006	Qualcomm	Unknown
H. Inaltekin	ECE Ph.D.	Wicker	May 2006	Postdoc at Cornell	Unknown
Z. Yang	ECE Ph.D.	Tong	May 2006	Marvell Semiconductor, Inc.	Unknown
C. Lee	ECE Ph.D.	Tong	May 2006	Unknown	Unknown
F. Akopyan	ECE M.S.	Manohar	May 2006	Continuing on to Ph.D.	Unknown

8.6. Outputs of Knowledge Transfer Activities

List, to the extent known, the general outputs of knowledge transfer activities since the last reporting period. Include patent names, numbers, application dates, and receipt dates; license names, numbers, licensed by, and dates; names of start-up companies, year, main product; and any other outputs of knowledge transfer activities not listed above.

F. Akopyan, R. Manohar, and A. Apsel. "A level-crossing Flash Asynchronous Analog-to-Digital Converter." Provisional patent filed, 3/2006.

Birman has been extremely active in TRUST-related transitioning efforts during the reporting period. First, Birman participated in a series of high-profile studies for the Air Force that focused on TRUST themes that arise in connection with that organization's move to GIG and NCES "SOA" standards. One study, for the Air Force CIO (Mr. Gilligan, later replaced by General Croom and Mr. Tillotson) focused on the implications of the deployment now underway in these areas; the second, Prometheus, was conducted for AFRL and explored options for aligning AFRL research on the Joint Battlespace Infosphere with AFRL priorities. Sastry participated in the CIO study, and Schmidt and Reiter were team-members on the Prometheus study. Both resulted in additional funding to the TRUST community. Additionally, he has worked with the US government both to develop a new national strategy for research in cyber security (this was part of an effort led by DHS but also involved participants from White House OSTP and NSF), and with the US Department of Treasury on the creation of a small center for research on TRUST issues in financial settings. A 2-day research topic on the subject helped refine a Treasury priorities and strategic vision document, and Birman is now teaming with developers of the eCavern remote backup and disaster recovery facility on replication techniques for their setting. As program committee chairman for the 20th ACM Symposium on Operating Systems Principles, Birman helped highlight many of the best results in the field by emphasizing TRUST topics in the call for papers, and also arranged a panel on peer-to-peer computing that focused on the real value and robustness of these new but highly controversial protocols and techniques. Birman's group has developed software that is in wide use; his latest efforts include Astrolabe (which runs Amazon.com's data centers), and Ricochet, which has just been released to the public in open-source form and slashes the latencies for time-critical computing systems. Birman also points to several publications aimed specifically at educating the general public about TRUST issues, notably Ken Birman, Coimbatore Chandrasekaran, Danny Dolev, and Robbert van Renesse. How the Hidden Hand Shapes the Market for Software Reliability.

Team for Research in Ubiquitous Secure Technology
2005-2006 Annual Report

Submitted to the First IEEE Workshop on Applied Software Reliability, June 2006; Ken Birman. The Untrustworthy Services Revolution. IEEE Computer (ISSN 0018-9162). Vol.39 No.2, Pgs. 98-100. February 2006; and Ken Birman. Can Web Services Scale Up? IEEE Computer. Volume 38. Number 10. Pgs.107-110. October 2005. Birman's book has been widely adopted as the basis for MEng and PhD-level courses in reliability and trusted computing: Reliable Distributed Systems Technologies, Web Services, and Applications. Birman, Kenneth P. 2005, XXXVI, 668 p. 145 illus., Hardcover ISBN: 0-387-21509-3.

8.7. TRUST's Participants

List all participants in Center activities alphabetically by category (undergraduate students, graduate students, faculty, visiting faculty, other research scientists, post-doctorates, pre-college students, teachers, educators, and other participants) and demographic characteristics (gender, disability status, ethnicity, race, and citizenship).

	Role	Name	University						
1	Faculty	Aiken, Alex	Stanford University						
2	Graduate Students	Akopyan, Filipp	Cornell University						
3	Faculty	Anantharam, Venkat	University of California, Berkeley						
4	Faculty	Bajcsy, Ruzena	University of California, Berkeley						
5	Graduate Students	Balakrishnan, Mahesh	Cornell University						
6	Graduate Students	Balasubramanian, Jaiganesh	Vanderbilt University						
7	Graduate Students	Barreno, Marco	University of California, Berkeley						
8	Graduate Students	Barth, Adam	Stanford University						
9	Undergraduate Students	Bhattacharjee, Tonmoy	SUNY at Stonybrook						
10	Faculty	Birman, Ken	Cornell University						
11	Faculty	Boneh, Dan	Stanford University						
12	Graduate Students	Bowers, Kevin	Carnegie Mellon University						
13	Other Participants	Brooks, Christopher	University of California, Berkeley						
14	Graduate Students	Cadar, Cristian	Stanford University						
15	Faculty	Canny, John	University of California, Berkeley						
16	Graduate Students	Chan, Haowen	Carnegie Mellon University						
17	Faculty	Culler, David	University of California, Berkeley						

Team for Research in Ubiquitous Secure Technology
2005-2006 Annual Report

18	Graduate Students	Datta, Anupam	Stanford University						
19	Post Doctorates	Datta, Anupam	Stanford University						
20	Graduate Students	Dhamija, Rachna	University of California, Berkeley						
21	Graduate Students	Eby, Matthew	Vanderbilt University						
22	Post Doctorates	Eklund, Mikeal	University of California, Berkeley						
23	Faculty	Engler, Dawson	Stanford University						
24	Undergraduate Students	Frye, Kaseima	North Carolina State University						
25	Other Participants	Gamble, Jessica	University of California, Berkeley						
26	Faculty	Garcia-Molina, Hector	Stanford University						
27	Faculty	Gehrke, Johannes	Cornell University						
28	Graduate Students	Goh, Eu-jin	Stanford University						
29	Graduate Students	Gohari, Amin Aminzadeh	University of California, Berkeley						
30	Graduate Students	Haridasan, Maya	Cornell University						
31	Faculty	Harky, Dan	San José State University						
32	Graduate Students	He, Ting	Cornell University						
33	Undergraduate Students	Hernandez, Sonny	University of Southern California						
34	Graduate Students	Ibrahim, Mahad	University of California, Berkeley						
35	Undergraduate Students	Jimenez, Jessica	University of Puerto Rico						
36	Graduate Students	Jones, Erin	University of California, Berkeley						
37	Faculty	Joseph, Anthony	University of California, Berkeley						
38	Graduate Students	Karlof, Chris	University of California, Berkeley						
39	Faculty	Karsai, Gabor	Vanderbilt University						
40	Faculty	Konrad, Almudena	Mills College						
41	Graduate Students	Kuo, Cynthia	Carnegie Mellon University						
42	Graduate Students	Kuryloski, Philip	Cornell University						

Team for Research in Ubiquitous Secure Technology
2005-2006 Annual Report

43	Research Scientists	Larry Howard	Vanderbilt University						
44	Faculty	Lee, Edward	University of California, Berkeley						
45	Other Participants	Lerner, Jack	University of California, Berkeley						
46	Graduate Students	Li, Yaping	University of California, Berkeley						
47	Faculty	Manohar, Rajit	Cornell University						
48	Graduate Students	Mathe, Janos Laszlo	Vanderbilt University						
49	Faculty	Maurer, Stephen M.	University of California, Berkeley						
50	Graduate Students	McCune, Jonathan M.	Carnegie Mellon University						
51	Faculty	McFadden, Dan	University of California, Berkeley						
52	Graduate Students	Meingast, Marci	University of California, Berkeley						
53	Faculty	Meldal, Sigurd	San José State University						
54	Graduate Students	Misra, Saswat	Cornell University						
55	Faculty	Mitchell, John	Stanford University						
56	Graduate Students	Modadugu, Nagendra	Stanford University						
57	Faculty	Mulligan, Deirdre	University of California, Berkeley						
58	Graduate Students	Mungamuru, Bob	Stanford University						
59	Faculty	Nacht, Michael	University of California, Berkeley						
60	Faculty	Necula, George	University of California, Berkeley						
61	Faculty	O'Rourke, Joe	Smith College						
62	Graduate Students	Ostrowski, Krzysztof	Cornell University						
63	Graduate Students	Pai, Sameer	Cornell University						
64	Graduate Students	Parno, Bryan	Carnegie Mellon University						
65	Faculty	Paxon, Vern	University of California, Berkeley						
66	Faculty	Perrig, Adrian	Carnegie Mellon University						

Team for Research in Ubiquitous Secure Technology
2005-2006 Annual Report

67	Graduate Students	Phanishayee, Amar	Cornell University						
68	Post Doctorates	Pleisch, Stefan	Cornell University						
69	Graduate Students	Qu, Hui	Cornell University						
70	Faculty	Raghavan, Vijay	Vanderbilt University						
71	Graduate Students	Ramasubramanian, Venugopalan	Cornell University						
72	Faculty	Reiter, Michael	Carnegie Mellon University						
73	Faculty	Robinson, William	Vanderbilt University						
74	Graduate Students	Roosta, Tanya	University of California, Berkeley						
75	Faculty	Rosenblum, Michael	Stanford University						
76	Graduate Students	Samar, Asad	Carnegie Mellon University						
77	Faculty	Samuelson, Pamela	University of California, Berkeley						
78	Faculty	Sastry, Shankar	University of California, Berkeley						
79	Faculty	Schmidt, Doug	Vanderbilt University						
80	Faculty	Schneider, Fred	Cornell University						
81	Faculty	Seshia, Sanjit A.	University of California, Berkeley						
82	Graduate Students	Shi, Elaine	Carnegie Mellon University						
83	Graduate Students	Shieh, Alan	Cornell University						
84	Faculty	Shim, Simon	San José State University						
85	Visiting Faculty	Shiuh-pyng Shieh	University of California, Berkeley						
86	Post Doctorates	Sinopoli, Bruno	University of California, Berkeley						
87	Faculty	Sirer, Gun	Cornell University						
88	Faculty	Song, Dawn	Carnegie Mellon University						
89	Other Participants	Sprinkle, Mary Margaret	University of California, Berkeley						
90	Faculty	Stoica, Ion	University of California, Berkeley						

Team for Research in Ubiquitous Secure Technology
2005-2006 Annual Report

91	Faculty	Su, Xiao	San José State University						
92	Post Doctorates	Subrahmanyam, P.A.	University of California, Berkeley						
93	Faculty	Sztipanovits, Janos	Vanderbilt University						
94	Graduate Students	Tavoularis Christina	Cornell University						
95	Other Participants	Terheggen, Sara	University of California, Berkeley						
96	Faculty	Tong, Lang	Cornell University						
97	Faculty	Tygar, Doug	University of California, Berkeley						
98	Research Scientists	VanRenesse, Robbert	Cornell University						
99	Faculty	Varian, Hal	University of California, Berkeley						
100	Faculty	Vollset, Einar	Cornell University						
101	Faculty	Wagner, David	University of California, Berkeley						
102	Graduate Students	Weatherspoon, Hakim	University of California, Berkeley						
103	Faculty	Weber, Steve	University of California, Berkeley						
104	Graduate Students	Werner, Jan	Vanderbilt University						
105	Faculty	Wicker, Steve	Cornell University						
106	Faculty	Wing, Jeannette	Carnegie Mellon University						
107	Graduate Students	Wu, Taojun	Vanderbilt University						
108	Graduate Students	Xie, Yichen	Stanford University						
109	Faculty	Xue, Yuan	Vanderbilt University						
110	Graduate Students	Yee, Ka-Ping	University of California, Berkeley						
111	Faculty	Yu, Weider	San José State University						
112	Graduate Students	Zhang, Xin	Cornell University						
113	Graduate Students	Zhuang, Li	University of California, Berkeley						

8.8. TRUST's Institutions & Partners

Provide a summary table with the following information for the Center: the number of participating institutions, the number of institutional partners, the total leveraged support, and the number of participants.

Participating Institutions	8
Institutional Partners	6
Total Leveraged Support	\$5,200,000
Total Participants	113

8.9. TRUST in the News

Describe any media publicity the Center received in the reporting period. Provide any appropriate media materials that can be used to disseminate information on Center accomplishments and activities to the public.

1. TRUST has been reported on in several national news outlets. The following is a list of some of the main ones, in reverse chronological order, that we are aware of.
2. February 19, 2006: Fred Schneider's presentation at the annual meeting of the American Association for the Advancement of Science was covered in Linux Electrons: Computer Security Lacks Accountability Says Cornell Expert.
3. February 7, 2006: The February 2006 IEEE Computer Magazine contains articles by a number of Trust Members including Kenneth Birman, Janos Sztipanovits, Gabor Karsai, and Douglas Schmidt.
4. January 27, 2006: Deirdre Mulligan was interviewed on Democracy Now, a Radio and TV program about "The Great Firewall of China: Internet Companies Censor Material at Chinese Government"
5. September 13 - 19, 2005 The Keyboard Sound Detection work of Professor Doug Tygar's group was covered in The San Francisco Chronicle, Scientific American, Slashdot and other media outlets. See Professor Tygar's publication page for a preprint.
6. August 4, 2005: The Credence project of Professor Emin Gun Sirer's group was featured on Slashdot and in the New Scientist in March. Credence is a distributed object reputation management scheme that counteracts content pollution in peer-to-peer filesharing systems.
7. 2 professors go fishing for phishers San Francisco Chronicle, July 25, 2005.
8. Stanford joins multi-institution center on research in cybersecurity and computer trustworthiness Stanford Report, April 14, 2005.
9. Campus to Direct New Research Center UC Berkeley to Lead Team In Pursuit of Internet Security The Daily Californian, April 14, 2005.
10. U.S. Grant Offered To Team Studying Computer Attacks Wall Street Journal, April 12, 2005.
11. U.C. Berkeley to head cybersecurity project NY Times, April 12, 2005.

Team for Research in Ubiquitous Secure Technology
2005-2006 Annual Report

12. Vanderbilt engineering part of national 'dream team', To design, develop new secure system design technologies Vanderbilt News Service, April 12, 2005.
13. Smith joins bid to thwart cyberattacks. Boston Globe (AP), April 12, 2005.
14. NSF establishes cybersecurity center ComputerWorld, April 12, 2005.
15. Cal picked to lead coalition to fortify network security Contra Costa Times, April 12, 2005.
16. Cal will lead effort against cyberattacks Berkeley to lead U.S. effort to foil cyberattacks Oakland Tribune, April 12, 2005.
17. U.C. Berkeley to head cybersecurity project ZDNet, April 12, 2005.
18. Universities, industry to fight hacker threat 5-year, \$19 million project intended to boost cybersecurity San Francisco Chronicle (AP), April 12, 2005. UC-Berkeley Leads Cybersecurity Consortium Washington Post (AP), April 12, 2005.
19. NSF established two new technology centers Washington Times (UPI) April 12, 2005.
20. UC-Berkeley Leads Cybersecurity Consortium Forbes, April 11, 2005.
21. Grant to research computer security San Jose Mercury News, April 11, 2005.
22. NSF launches \$19 million research program for computer security, Cornell University News Service, April 11, 2005.
23. Researchers Are Part of New NSF Center Studying Cybersecurity and Trustworthy Computing Carnegie Mellon Media Relations, April 11, 2005.
24. UC Berkeley to lead \$19 million NSF center on cybersecurity research UC Berkeley Campus News, April 11, 2005.
25. NSF Announces Intent to Establish Two New Science and Technology Centers National Science Foundation, April 11, 2005.

8.10. Center Employment and Turnover

NSF is to be notified if the Center experiences difficulties in filling any of the management positions named in the approved proposal as modified. The Center must make all reasonable efforts to have these unfilled Center management team personnel in place within 6 months of the vacancy. In the event of management team turnover, immediate notification of the vacancy to the NSF Lead Program Official shall include plans for filling the position and for allocating the duties of the position among existing staff members on an interim basis.

All positions are currently filled. There have been no vacancies in the Center management team for longer than 6 months.