Cyber Security Policy and Research Institute

THE GEORGE WASHINGTON UNIVERSITY

Secure and Trustworthy Cyberspace First PI Meeting

Summary Report

Carl Landwehr
Lance J. Hoffman

March 25, 2013

Report GW-CSPRI-2013-01

Support for this research was provided through National Science Foundation Division of Computer and Network Systems Award 1243386

Secure and Trustworthy Cyberspace First PI Meeting Summary Report

November 27-29, 2012
Gaylord National Hotel and Convention Center
National Harbor, Maryland

Carl Landwehr
Lance J. Hoffman

Abstract

This report describes a Secure and Trustworthy Cyberspace (SaTC) conference that involved a wide range of researchers and Principal Investigators (PI). This meeting provided an important opportunity to broaden the perspectives of researchers from all of the fields involved by hearing from a series of experts and viewing a wide variety of posters. Additionally, it helped them begin to form new partnerships and collaborations through a series of short 10-minute conversations entitled Cross-Disciplinary Conversations. Finally, it showed researchers how their discoveries and advances may be transitioned into practice in systems such as those operated by the Office of Cyberinfrastructure (OCI). Some additional feedback from conference participants and ideas for future work suggested by the conference's organizers are also described.

Table of Contents

1. Introduction

Cybersecurity is an important challenge in today's world. Corporations, agencies, national infrastructure and individuals have been victims of cyber-attacks. Addressing this problem requires multi-disciplinary expertise in human, statistical, mathematical, computational, and computer sciences and ultimately the transition of new concepts and technologies to practice.

2. Secure & Trustworthy Cyberspace Program

The Secure and Trustworthy Cyberspace (SaTC) program of the National Science Foundation (NSF) is an interdisciplinary project including technologists, social scientists, and educators from programs sponsored by NSF CISE, SBE, and EHR directorates. This PI meeting encompassed all of these perspectives on cybersecurity through plenary talks, breakout sessions, posters, and informal Birds of a Feather gatherings. The technology portion of SaTC replaced the Trustworthy Computing (TC) and Cyber Trust (CT) programs, so former TC and CT PIs are now SaTC PIs.

The SaTC program seeks proposals that address cybersecurity from one or more of three perspectives: *Trustworthy Computing Systems, Social, Behavioral and Economics,* and *Transition to Practice,* as well as proposals that combine multiple perspectives. Proposers are invited to submit proposals in three project classes, which are defined below:

- Small projects - up to $500,000 in total budget, with durations of up to three years
- Medium projects - $500,001 to $1,200,000 in total budget, with durations of up to four years
- Frontier projects - $1,200,001 to $10,000,000 in total budget, with durations of up to five years

Projects with Trustworthy Computing Systems and/or Social, Behavioral and Economic perspectives can include an optional *Transitions* phase, described in a supplemental document of no more than five pages that lays out how successful research results are to be further developed, matured and experimentally deployed in organizations or industries, including in networks and end systems used by members of the NSF science and engineering communities. Proposals for small or medium projects with a Transitions phase can exceed the above-stated maximums up to $167,000 in the small category and as much as $400,000 in the medium category.

3. Relationship of SaTC to NSF and to U. S. Cybersecurity Research

The Secure and Trustworthy Cyberspace program now involves several NSF Directorates: Computer and Information Science and Engineering (CISE), which leads and integrates the program, Social, Behavioral, and Economic Sciences (SBE), Mathematical and Physical Sciences (MPS), Engineering (ENG), and also the Office of Cyberinfrastructure (OCI). Broad as the program is, research related to cybersecurity in different respects can also be found in a variety other NSF research programs and activities, and budgetary reports from NSF aiming for comprehensive reporting include those additional research activities. Outside NSF, many other government agencies conduct research in cyber security and information assurance, including the Departments of Defense, Energy, Homeland Security, and others. NSF and other agencies coordinate their research programs through the Cyber Security and Information Assurance (CSIA) Interagency working group under the auspices of the National Information Technology Research and Development (NITRD) program, which reports to the National Science and Technology Council in the Office of Science and Technology Policy (OSTP).

4. Conference Events

SaTC held a two-and-a-half-day conference of its principal investigators (PIs) from Tuesday morning, November 27, to Thursday noon, November 29, 2012 at the Gaylord National Hotel and Convention Center in National Harbor, MD (in the Washington, DC area).

a. Agenda

All of the events listed below took place at the Gaylord National Resort and Convention Center in National Harbor, Maryland. A digital copy of the agenda can be found at http://cps-vo.org/group/satc/program, and the PowerPoint slides for each presenter are available at http://cps-vo.org/node/3488/browser.

Tuesday, November 27

Time	Event
7:30 am	Breakfast
8:15 am	The New Era of Science and Engineering Subra Suresh, Director of NSF *Introduction and Session Chair: Farnam Jahanian, Assistant Director of NSF for Computer & Information Science & Engineering*
9:15 am	Keynote: Eric Grosse *Vice President for Security Engineering, Google* *Introduction and Session Chair: Keith Marzullo, NSF*
10:00 am	*Break*
10:15 am	Crossing the Line: Recent research results that cross disciplines *Michael Byrne, Rice: Voting Machines and Human Behavior* *Fabian Monrose, UNC-CH: Understanding Encrypted Speech* *Vern Paxson, ICSI: Spam Economics* *Dan Boneh, Stanford: Implicitly Learned Passwords* Discussion: How do Cross-Disciplinary Efforts Get Started? *Session Chair: Jeremy Epstein, NSF*
11:45 pm	Plenary Address: Multi-Disciplinary Aspects of Cyber Security: *Angela Sasse, University College London* *Introduction and Session Chair: Keith Marzullo, NSF*
12:30 pm	Lunch pickup and return to meeting room
1:15 pm	The Federal Cybersecurity R&D Strategic Plan: What is it, What Gets Funded, and What's the Future? *Part 1: What is it* *Tomas Vagoun, NITRD and Bill Newhouse, NIST, NITRD Cyber Security and Information Assurance Interagency Working Group* *Part 2: What Gets Funded* *Keith Marzullo, NSF* *Brad Martin, ODNI* *Steve King, OSD*

Douglas Maughan, DHS

Part 3: What's The Future – An Open Discussion
Session Chair: Jeremy Epstein, NSF

3:00pm Break

3:30 pm Cross Disciplinary Conversations
Pre-arranged, focused 1-1 meetings between researchers with expertise in different disciplines (see separate sheet for details and badge for assignments)
Introduction and Procedures: Sam Weber, NSF
Technical Coordinators: Apu Kapadia, Indiana University, and Elaine Shi, University of Maryland

5:30 pm Poster Room opens
Coordinator: Micah Sherr, Georgetown University

5:30 pm Rump / BoF sessions

7:00 pm Adjourn

7:00 pm Dinner (on own)

Wednesday, November 28

Time	Event
7:30am	Breakfast
8:30 am	Welcome and Introductions: Myron Gutmann, Assistant Director of NSF for Social, Behavioral & Economic Sciences Alan Blatecky, Office Director for the NSF Office of Cyberinfrastructure
8:50 am	Transition to Practice: How to Identify Ideas Ready for Transition and What to Do Next *Introduction and session chair: Kevin Thompson, NSF* *Ron Perez, Cyber Security Research Alliance (CSRA)* *Doug Maughan, DHS* *Becky Bace, University of South Alabama* *Vern Paxson, ICSI* *Paul Barford, University of Wisconsin*
10:10 am	Teaching and Learning: Competitions and Cybersecurity *Nick Weaver, ICSI: Skills Competitions vs. "Build-It" Competitions* *Ben Cook, Sandia: Starting a "Build-It" Competitio:* *Ron Dodge, USMA: Learning more from Skills Competitions* *Session Chair: Victor Piotrowski, NSF*
10:55 am	Break
11:25 am	Massively Open Online Courses (MOOCs) and Cybersecurity *Introduction and Session Chair: Victor Piotrowski, NSF*

John Mitchell, Stanford (remote)

12:10 pm	Lunch Pickup and return to meeting room
12:45 pm	Discussion Session Charge: Jeremy Epstein, NSF *Session Coordinators: Daniel Weitzner, MIT and Michael Reiter, University of North Carolina Chapel Hill*
1:15 pm	Discussion Sessions Convene
3:00 pm	Break
3:30 pm	Discussion Sessions Continue, Develop Out-Briefs
5:30 pm	Poster Session *Coordinator: Micah Sherr, Georgetown University*
5:30 pm	Rump / BOF sessions
7:00 pm	Adjourn
7:00 pm	Dinner (on own)

Thursday, November 29

Time	Event
7:30 am	Breakfast
8:30 am	Opening Announcements: NSF Leadership
8:45 am	Out-Briefs from Discussion Sessions: *Discussion Group Leaders* *Session Co-Chairs: Daniel Weitzner, MIT and Michael Reiter, University of North Carolina Chapel Hill*
10:30 am	Break
11:00 am	Closing Plenary: Ignorance: How it Drives Science *Stuart Firestein, Columbia University* *Introduction and Session Chair: Jeremy Epstein, NSF*
12:00 pm	Closing remarks *Keith Marzullo, NSF*
12:10 pm	SaTC PI Meeting Adjourns; Box Lunches Available
1:00 pm	Science of Security Community Meeting Opens (for those remaining)

b. Discussion Sessions

One key aspect of the conference was splitting up the participants to answer 19 challenging questions related to cybersecurity. Members met for several hours and produced a presentation that was shared with the rest of the conference. Their slides can be found at http://cps-vo.org/node/6522, and a summary of their findings and answers to the questions listed below can be found in Appendix I of this report.

1. How can we teach, and encourage and evaluate the teaching of, safe programming practices to reduce the vulnerability of future software systems?
2. What threat models should guide future SaTC research?
3. What are the characteristics of SaTC ideas/technologies that are ready for transition to practice, and what are the success paths and pitfalls for different approaches to transition?
4. What are the barriers to creating systems with security and privacy properties that users can understand and use?
5. What might a building code for critical infrastructure software/hardware look like?
6. What models are being used in cybersecurity research, and what models are needed?
7. Where does deconstructive security research belong in the research portfolio?
8. What policies and norms should govern the Internet commons in an era of cyberwar?
9. How do research methods vary across the disciplines involved in cybersecurity?
10. What modeling techniques should we use to account for the role of humans in complex cyber systems?
11. Predicting the next "flash crash" or blackout: What methods are available for evaluating the stability/trustworthiness of complex digital infrastructure systems?
12. Anonymity and accountability: how do we enable tradeoffs?
13. What policies and technologies would be required to enforce the expiration of data?
14. How can we assure provenance, integrity, longevity of scientific records?
15. Identity management: why don't we have it and do we actually need it?
16. How can we leverage R&D work done to improve cybersecurity education?
17. How can the nation best build and sustain an appropriately sized and qualified cybersecurity workforce?
18. What issues are unique to cyber warrior education (compared to other members of the cyber workforce? [CANCELLED]
19. What incentives, norms, attitudes, habits, cognitive limits, or other mechanisms present the most important obstacles to cybersecurity, and how might such factors be utilized to benefit cybersecurity?
20. What are the group, organization, institutional, and policy obstacles to cybersecurity?

c. Cross Disciplinary Conversations

A primary objective of the meeting was to stimulate cross-disciplinary and interdisciplinary research leading to a more secure and trustworthy cyberspace than we have today. To facilitate new cross-disciplinary research collaborations, a session of "cross disciplinary conversations" was organized and held the first afternoon of the meeting. The idea was simply to create new connections among participants with different backgrounds early in the meeting.

This session required substantial planning and organization but, according to many of the participants, was a highlight of the meeting. A subcommittee of the Steering Committee (Susan Landau and Stefan Savage) helped develop the initial plan for the session. Chairs (Apu Kapadia and Elaine Shi) were recruited to organize and conduct the session. After considering the facilities available at the conference venue, the session chairs developed a simplified plan that worked well at the meeting.

In advance of the meeting, a set of categories of expertise was developed in cooperation with the Steering Committee, Chairs, and NSF staff that was intended to cover the breadth of the SaTC program. The categories were:

cognitive science
computer science/engineering - foundations and theory
computer science/engineering - systems
computer science/engineering - networks
computer science/engineering - information and applications
cyber physical systems
cybersecurity
economics
behavioral economics
education
mathematics and statistics
political science
psychology
public policy
science of organizations
sociology

As part of the automated registration process, each participant was asked to identify up to four categories in which they were expert and up to four more categories in which they would like to find conversation partners. A program was developed by Zahid Rahman, a graduate student at Indiana University to take the registration information and create a schedule of conversations using algorithms originally invented to solve the stable marriage problem.

A period of two hours was set aside during the first afternoon of the meeting to accommodate 8 10-minute conversations for each participant, with 5 minutes between discussions to locate the next partner. A seating arrangement was organized so that for each discussion session, the

matched participants would be seated across from each other (see Figure 1). Simple procedures were developed to deal with "no-shows" as well.

A suggested organization for each discussion was provided:

(i) Each person spends two minutes introducing themselves (four minutes);

(ii) Each person suggests someone in their field that it would be good for the other one to meet; those proposals are sent in real time to a coordinator to facilitate possible conversations for rounds seven and eight (one minute);

(iii) They spend six minutes brainstorming about connections between their fields, i.e., trying to think of research projects that two such people could do together.

Informal discussions with the participants as well as comments received on forms reflected widespread satisfaction with the session. A few concerns were also raised: many felt that six(rather than eight) conversations would have been sufficient, and the noise level in the discussion room was rather high during the conversations. Subsequently, one of the participants suggested that by creating a "co-author" graph using public resources like Google Scholar and the ACM Digital Library, future organizers could help avoid scheduling conversations between people who were already collaborating.

Figure 1. Cross Disciplinary Conversations at SaTC PI Meeting

d. Posters Presented

The following 95 posters were presented on November 28-29 at the conference. Some of the posters are available online at http://cps-vo.org/node/3488/browser.

	PI Name	Poster Title	Poster Session Presenter(s)
1	Bonnie Anderson (Brigham Young Univ.)	Using Neuroscience to Explain User Responses to Malware Warnings	Bonnie Anderson
2	Terrence August (Univ. of California, San Diego)	A Comparative Analysis of Software Liability Policies	Terrence August
3	Clark Barrett (New York Univ.)	CVC4: A tool for automatically reasoning about programs and systems	Clark Barrett
4	Jim Basney (NCSA)	Distributed Web Security for Science Gateways	Jim Basney
5	Elisa Bertino (Purdue Univ.)	A Comprehensive Model for Provenance	Elisa Bertino
6	Alexandra Boldyreva (Georgia Institute of Technology)	Efficiently Searchable Symmetric Encryption	Alexandra Boldyreva
7	Wayne P. Burleson (Univ. of Massachusetts Amherst)	E-cash for Intelligent Public Transportation	Wayne P. Burleson
8	Justin Cappos (NYU Poly)	Seattle: An Educational Platform for System Classes	Justin Cappos
9	Sandra Carpenter (Univ. of Alabama in Huntsville)	Psychological Attacks and Mitigations	Sandra Carpenter
10	Hao Chen (Univ. of California, Davis)	Detecting Cloned Applications on Android Markets	Hao Chen
11	Yan Chen (Northwestern Univ.)	Real-time Private Information Leakage Detection on Android without System Modification	Yan Chen
12	Yingying Chen (Stevens Institute of Technology)	CAREER: EASE: Enhancing the Security of Pervasive Wireless Networks by Exploiting Location	Yingying Chen
13	Yu Cheng (Illinois Institute of Technology)	Intrusion Detection for Multimedia Communications over 802.11 Based Wireless Networks: An Analytical Approach	Yu Cheng
14	Kimberly Claffy (CAIDA, Univ. of California, San Diego)	Detection and Analysis of Large-scale Internet Infrastructure Outages	Alberto Dainotti
15	Chris Clifton (Purdue Univ.)	Generalizing Text to Protect Privacy	Chris Clifton
16	Reza Curtmola (NJIT)	Remote Data Integrity Checking for Cloud Storage	Reza Curtmola
17	Ram Dantu (Univ. of North Texas)	Another Free App: Does It Have the Right Intentions?	Ram Dantu

	PI Name	Poster Title	Poster Session Presenter(s)
18	Yingfei Dong (Univ. of Hawaii)	Experimental Study of Accountability in Existing Anonymous Networks	Yingfei Dong, Zhenhai Duan
19	Wenliang Du (Syracuse Univ.)	Web Security: Re-Designing the Web's Security Infrastructure	Wenliang Du
20	Tudor Dumitraş (Symantec Research Labs)	WINE: Data-Intensive Experiments in Security	Tudor Dumitraş
21	Shantanu Dutt (Univ. of Illinois at Chicago)	Trusted FPGA Design Using Non-Integrated and Fully-Integrated Embedding of Check Structures	Shantanu Dutt
22	Sonia Fahmy (Purdue Univ.)	Scaling Network Security Experiments	Sonia Fahmy
23	Nelly Fazio (Graduate Center & City College of CUNY)	Anonymous, Secure and Robust Multi-Recipient Communication	Nelly Fazio
24	Wu-chang Feng (Portland State Univ.)	kaPoW Web Plug-ins	Wu-chang Feng
25	Errin Fulp (Wake Forest Univ.)	An Evolutionary-Inspired Approach for Moving Target Defenses	Errin Fulp
26	Johannes Gehrke (Cornell Univ.)	Controlling Disclosure in App Ecosystems	Johannes Gehrke
27	Sharon Goldberg (Boston Univ.)	A Strategy for Transitioning to BGP Security	Sharon Goldberg
28	Luanne Goldrich (JHU/APL)	DART3/Integrated Demonstration	Luanne Goldrich, Tanner Allen
29	Venu Govindaraju (SUNY Buffalo)	Integrating Privacy Preserving Biometric Templates and Efficient Indexing Methods	Venu Govindaraju
30	Guofei Gu (Texas A&M Univ.)	Using Enemies' Strength against Them	Guofei Gu
31	Mina Guirguis (Texas State Univ.)	Securing Mobile CPSs against Stealthy Attacks	Mina Guirguis
32	Andreas Haeberlen (Univ. of Pennsylvania)	Secure Network Provenance	Andreas Haeberlen
33	Drew Hamilton (Auburn Univ.)	A Digital Forensics Cyberinfrastructure Workforce Training Initiative for America's Veterans	Drew Hamilton
34	Matthew Hashim (U. of Arizona)	Collaboration, Interdependency, and Transfer Pricing	Matthew Hashim
35	Haibo He (Univ. of Rhode Island)	Secure the Electrical Power Grid: Smart Grid versus Smart Attacks	Haibo He
36	Raquel Hill (School of Informatics and Computing, Indiana Univ.)	Privacy Risks in Social Science Datasets	Raquel Hill
37	Adele Howe (Colorado State Univ.)	Computer Security for the Home User	Adele Howe

	PI Name	Poster Title	Poster Session Presenter(s)
38	Shuyuan Mary Ho (Florida State Univ.)	Organic Social Firewall: A Human-Computational Study of Trustworthy Communications	Shuyuan Mary Ho
39	Ted Huffmire (Naval Postgraduate School)	3D Security	Ted Huffmire
40	Trent Jaeger (Pennsylvania State Univ.)	Comprehensive System Verification in Cloud Computing Environments	Trent Jaeger
41	Niraj Jha (Princeton Univ.)	Enhancing the Safety and Trustworthiness of Medical Devices	Niraj Jha and Anand Raghunathan
42	Murat Kantarcioglu (Univ. of Texas at Dallas)	Efficient Similarity Search over Encrypted Data	Murat Kantarcioglu
43	Sang Wu Kim (Iowa State Univ.)	Securing Wireless Network Coding against Pollution Attack at the Physical Layer	Sang Wu Kim
44	David Kotz (Dartmouth College)	Trustworthy Information Systems in Healthcare	Denise Anthony
45	Adam Lee (Univ. of Pittsburgh)	Foundations of Application-Sensitive Access Control Evaluation	Adam Lee
46	Janne Lindqvist (Rutgers Univ.)	Redesigning Mobile Privacy	Janne Lindqvist
47	Anna Lysyanskaya (Brown Univ.)	Reconciling Privacy and Accountability	Anna Lysyanskaya
48	Dan Massey (Colorado State Univ.)	BGP Monitoring and Security	Dan Massey
49	Nasir Memon (Colorado State Univ.)	Investigating Multi-touch Gesture as a Novel Biometric Modality	Nasir Memon
50	Jelena Mirkovic (USC/ISI)	Critter@home: Content-Rich Traffic Trace Repository from Real-Time Anonymous, User Contributions	Jelena Mirkovic
51	Antonio Nicolosi (Stevens Institute of Technology)	Provable Security from Group Theory and Applications	Antonio Nicolosi
52	Leon Osterweil (Univ. of Massachusetts Amherst)	Process Model-Based Continuous Improvement of Election Process Quality and Robustness	Leon Osterweil
53	Xinming Ou (Kansas State Univ.)	Classification of UDP Traffic for DDoS Detection	Alexandru G. Bardas
54	Xinming Ou (Kansas State Univ.)	Prioritizing Intrusion Analysis Using Dempster-Shafer Theory	Loai Zomlot
55	Roberto Perdisci (Univ. of Georgia)	Malware Defense via Download Provenance Classification	Roberto Perdisci
56	Ping Ji (City University of New York (CUNY) – John Jay College)	Security Monitoring for Wireless Network Forensics	Ping Ji
57	Corin Pitcher (DePaul Univ.)	Compositional Declarative Forensics	Corin Pitcher

	PI Name	Poster Title	Poster Session Presenter(s)
58	Donald Porter (Stony Brook Univ.)	Hardware Isolation from an Untrustworthy OS	Donald Porter
59	Michael Pozmantier (DHS S&T Cyber Security Division)	Transition to Practice	Michael Pozmantier
60	Atul Prakash (Univ. of Michigan)	Information Confinement on Commodity Systems	Atul Prakash
61	Lisa PytlikZillig (Univ. of Nebraska Public Policy Center)	Designing Cyber-Spaces & Promoting Cyber-Cultures to Reduce SPEC-Inspired Hacktivism	Lisa M. PytlikZillig, Tonya K. Bernadt, Michael J. Hayes, Ashok Samal, Leen-Kiat Soh, Alan J. Tomkins, Shiyuan Wang
62	H. Raghav Rao (Univ. of Buffalo)	Distrust of the Internet in Older Adults	H. Raghav Rao
63	Narasimha Reddy (Texas A&M University)	Detecting Spammers on Twitter	Narasimha Reddy
64	Kui Ren (SUNY Buffalo)	Privacy-preserving Search and Computation for Cloud Data	Kui Ren
65	Norman Sadeh (Carnegie Mellon Univ.)	User-Controllable Learning of Security and Privacy Policies	Norman Sadeh
66	Andre Scedrov (Univ. of Pennsylvania)	Declarative Privacy Policy: Finite Models and Attribute-Based Encryption	Andre Scedrov
67	Andre Scedrov (Univ. of Pennsylvania)	Towards an Automated Assistant for Clinical Investigations	Andre Scedrov
68	Patrick Schaumont (Virginia Tech)	Foundations for Future On-chip Fingerprints	Patrick Schaumont
69	Simha Sethumadhavan (Columbia Univ.)	Trustworthy Hardware from Untrustworthy Hardware	Simha Sethumadhavan
70	Zhong Shao (Yale Univ.)	Making OS Kernels Crash-Proof by Design and Certification	Zhong Shao
71	Radu Sion (Stony Brook Univ.)	NFSv4: An Extensible Security Layer for Network Storage	Radu Sion
72	Adam Smith (Pennsylvania State Univ.)	Analyzing Graphs with Node-level Differential Privacy	Adam Smith
73	Jon Solworth (Univ. of Illinois at Chicago)	Why Johnny's Application Is Not Secure	Jon Solworth
74	Gary Stoneburner (JHU APL)	New Journal for Non-Public Cybersecurity Research	Gary Stoneburner
75	Aaron Striegel (Univ. of Notre Dame)	On Cell Phones and Punishment: Encouraging Secure Mobile Behavior Through Morality	Aaron Striegel
76	Berk Sunar (Worcester Polytechnic Institute)	Homomorphic Encryption for Cloud Privacy	Berk Sunar
77	Patrick Tague (CMU)	Constrained Adaptive Jamming and Anti-Jamming	Patrick Tague

	PI Name	Poster Title	Poster Session Presenter(s)
78	Gang Tan (Lehigh Univ.)	Rocksalt: a formally verified machine code security checker	Gang Tan
79	Shambhu Upadhyaya (Univ. of Buffalo)	Designing Mission Survivable Systems Using Proactive Schemes	Shambhu Upadhyaya
80	Joseph Urban (Texas Tech Univ.)	An Innovative Interdisciplinary Cybersecurity Education Program for Protecting Critical Infrastructure	Vittal Rao
81	Anthony Vance (Brigham Young Univ.)	Deterring Unauthorized Access by Insiders: Raising Perceptions of Accountability in End Users Through User Interface Artifacts	Anthony Vance
82	Weichao Wang (UNC Charlotte)	Exploring the Security Capabilities of Physical Layer Network Coding (PNC) in Wireless Networks	Weichao Wang
83	Susanne Wetzel (Stevens Institute of Technology)	Privacy-preserving Reconciliation Protocols on Ordered Sets	Susanne Wetzel
84	Tilman Wolf (Univ. of Massachusetts Amherst)	Securing the Router Infrastructure of the Internet	Tilman Wolf
85	Matthew Wright (Univ. of Texas at Arlington)	Persea: A Sybil-Resistant Social DHT	Matthew Wright
86	Rebecca Wright (Rutgers Univ.)	Accountability and Identifiability	Aaron Jaggard and Rebecca Wright
87	Li Xiong (Emory Univ.)	Real-time Aggregate Monitoring with Differential Privacy	Li Xiong
88	Heng Xu (Pennsylvania State Univ.)	Privacy-by-ReDesign: Alleviating Users' Privacy Concerns for Third-Party Applications	Heng Xu
89	Yuan Xue (ISIS/Vanderbilt Univ.)	iSEE: integrated Simulation and Emulation platform for cyber-physical system security Experimentation	Yuan Xue
90	Yaling Yang (Virginia Tech)	SDR Spectrum Monitoring Through Power Finger Printing	Yaling Yang
91	Danfeng (Daphne) Yao (Virginia Tech)	Storytelling Security: Semantic and Structural Causal Analysis	Danfeng (Daphne) Yao
92	Heng Yin (Syracuse Univ.)	Virtualization and Binary Centric Approach to Malware Analysis	Heng Yin
93	Ben Zhao (UC Santa Barbara)	Social Turing Tests: Crowdsourcing Sybil Detection	Ben Zhao

	PI Name	Poster Title	Poster Session Presenter(s)
94	Huiyang Zhou (North Carolina State Univ.)	Architecting against Software Cache-based Side Channel Attacks	Huiyang Zhou
95	Lina Zhou (Univ. of Maryland Baltimore County)	Online Deception Behavior and its Automatic Detection	Lina Zhou

5. Participant Feedback

Participants were asked to complete a survey form at the close of the meeting. Approximately 100 responses were received, reflecting well over a third of the registrants. The detailed responses have been provided to NSF, but a few summary points are indicative:

- Nearly half of those responding (45 of 98) reported they met someone at the meeting with whom they may collaborate on a future proposal.
- About 90% rated the meeting as "Very Good" or "Good" overall
- About 80% said they would recommend colleagues attend a future (optional) meeting of a similar type held in two years.
- The Cross Disciplinary Conversations and the Plenary talk sessions were rated as Valuable, Highly Valuable or Extremely Valuable by the largest fractions (about 60%) of the participants. "Crossing the Line" talks, Poster sessions, and Breakout discussions were next with a bit over 50% rating each of these sessions as Valuable or better.

Thus, in the eyes of the participants (as well as the organizers) the meeting was quite successful overall.

6. Acknowledgements

This event could not have taken place without the hard work of individuals on three committees. They are acknowledged below.

Steering Committee

Alessandro Acquisti (Carnegie Mellon University)
Bill Arbaugh (University of Maryland)
Matt Bishop (University of California Irvine)
Elisa Bertino (Purdue University)
Joan Feigenbaum (Yale University)
Deb Frincke (National Security Agency)
Cormac Herley (Microsoft Research)
Susan Landau (Guggenheim Scholar, Privacyink.org)
Deirdre Mulligan (University of California Berkeley)
Stefan Savage (University of California San Diego)
Fred Schneider (Cornell University)

Organizing Committee

Lance Hoffman, Principal Investigator (George Washington University)
Carl Landwehr, Organizing Committee Liaison (George Washington University)
Micah Sherr, Poster Session Chair (Georgetown University)
Costis Toregas, University Liaison (George Washington University)
Elaine Shi, Cross-Disciplinary Conversations co-chair (University of Maryland)
Apu Kapadia, Cross-Disciplinary Conversations co-chair (Indiana University)
Zahid Rahman, Cross-Disciplinary Conversations programmer (Indiana University)

Local Arrangements Committee

Katie Dey (Vanderbilt University)
Anne Dyson (Vanderbilt University)
Frankie King (Vanderbilt University)
Xenofon Koutsoukos (Vanderbilt University)

Support for this research was provided through National Science Foundation Division of Computer and Network Systems Award 1243386.

Appendix: Discussion Group Summaries

DG1: How Can We Teach, and Encourage and Evaluate the Teaching of, Safe Programming Practices to Reduce the Vulnerability of Future Software Systems*?*

Discussion Leaders: Bill Pugh, Matt Bishop

The first step is to convince people that the change is beneficial, that change will enhance the reputation or ability of the academic institution in some way. For example, if teaching safe programming made graduates more attractive to prospective employers, especially those who are already engaged with the institution, then the institution could entice more students to come, increasing both the number of attendees and the resources that their extra tuition would bring. These benefits would encourage academic administrators to sustain the commitment to those teaching resources.

Faculty need to co-operate. In some cases, faculty would teach safe programming, but many faculty either are not skilled in it or don't know how to teach it. In this case, funding for faculty learning or for assistance in research or teaching could provide the faculty member with time to develop lessons and teaching material, and in doing so to become intimately familiar with the subject matter. Release time, allowing them a wider choice of courses to teach, and professional development opportunities (such as sending them to a SIGCSE workshop on safe programming) could also entice faculty.

Part of such faculty development is providing resources to help the faculty teach. A repository of educational units, associated exercises, and ancillary material such as videos and tools that users could adapt for their courses with minimal effort could help faculty build on the work of others. While this is an old idea, previous efforts in the realm of safe coding have failed, principally because these resources have been hard to locate, and they are not maintained or updated. Funding small efforts, where specific practitioners or faculty develop some educational units and, as part of the funding requirement, maintain them to ensure they were current, well documented, and easy to use or adapt, could be more fruitful than one large effort. A server that housed these educational units that is easy to use and robust could provide the visibility necessary for these units. Mechanisms for soliciting and collecting feedback from users—a rating system and commenting facility such as Amazon and other e-vendors use—could guide both developers those who are considering the use of the module.

One approach to encouraging adoption of robust programming techniques relies on this. Creating a set of educational units, and getting these endorsed by various major computer companies like Google, Apple, Microsoft, and others, and government agencies such as the Department of Homeland Security and the National Security Agency would indicate the importance of the topic and provide guidance to academic institutions and faculty of what others believe need to be taught.

Ultimately, the group felt that development of a "security mindset" would be the most important aspect of teaching safe programming. This mindset requires students to think holistically about the assumptions underlying the software. This need not be adversarial; but it must allow the student to consider what errors could happen if the assumptions prove to be incorrect. Once this mode of thinking is understood, taking steps to counter such errors—safe programming—follows naturally.

DG2: What threat models should guide future SaTC research?

Discussion Leader: Carl Gunter

Threat models are often viewed it as having three components: the attacker, the targeted asset, and the vector the attacker aims to use to gain access to the asset. The attacker's goal may be to damage integrity and/or availability of the asset, to learn its secrets, or to control its operation. If any of the elements of the triad are missing then the security threat generally disappears. If attacker, vector, or asset is lacking, the resource owner need not be concerned. This very general threat model remains acceptably robust as a foundation for SaTC research, but it is important to note that much of the current innovation arises from fresh ideas about what entities one sees as occupying one or more of these roles. Ideally, research proposals should cover as much of the model as possible, given the fact that a loss results only if all three elements are present. However, in many cases the threat is clear enough with only one or two of the elements being clearly specified, since the third may reasonably be conjectured. For instance, if an attacker has an exploit (vector) that compromises the microphone on a cell phone, the threat is well-defined even if the specific nature and goal of the attacker remains poorly defined because of the wide potential range of parties that might be interested in obtaining this capability.

New types of assets are getting new or greater attention. Indeed, this is possibly the greatest area of growth in new research by SaTC participants. Examples include embedded devices, such as networked computers in cars, the power grid, and healthcare systems. They also include large bodies of personal information and credentials held by certain types of entities, like companies with repositories of credit card numbers. Here it is important to note that there are often two asset owners involved, namely the manager of the repository and the provider of the personal data; it is important to view this as an aspect of the model in formulating the concept of loss to best estimate risk. Cloud services that host data are another example, as is the radio spectrum offered through white space programs. Advanced Persistent Threats (APTs) are creating risks for high-value assets like intellectual property.

Most types of attackers have been present for some time, but there are shifts in the scale of the attacker type. For instance, the growth of black markets to monetize information like credit card numbers has introduced an influential class of professional attackers for financial profit. Other groups like political activists ("hacktivists") and state actors have generated increasing concern. Insiders remain a major class of attackers for whom countermeasures are difficult. There are also new types of attackers like ISPs that modify or block traffic for financial or political reasons.

Attack vectors are also changing, although possibly at a slower rate than assets since many familiar vectors are found to work against new types of assets. For example, a buffer overflow may be exploited to compromise an embedded system. The attack is not new but the deployment of software in the embedded system creates a new threat. However, there are novel vectors as well as adaptations of old vectors to new assets. For instance, systems that once ran sequentially have gone to multi-core and this opens vulnerabilities from race conditions. App stores create a new vector for attack as do new side channels and sensors, especially on cell phones. Social networks offer new threats as well.

DG3: What are the characteristics of SaTC ideas/technologies that are ready for transition to practice? What are the success paths and pitfalls for different approaches to transition?

Discussion Leader: Douglas Maughan

Transition to practice is important, but without the research, there is nothing to transition. NSF should encourage transition but must not sacrifice research in favor of transition. There is some confusion and concern in the research community about the standards for evaluating proposals addressing transition and the competence of most review panels for this task.

Alternative approaches to transition include:

Make research products available under some form of open source license. This permits wide adoption, including commercial adoption, and adopters can enhance the software through their own contributions. Evangelists may champion the software and encourage others to adopt it. However, there is a continuing maintenance burden, and quality assurance can be an issues. Whether open source software tends to be more vulnerable or less vulnerable than closed source remains an open question.

Commercialize directly (startup companies). Companies become advocates for technology precisely because there is potential profit for them. They are motivated to find the value proposition for the technology, create marketing plans, and encourage uptake. Profits can also support Q/A and maintenance. On the other hand, enhancements to the technology are likely to be proprietary and therefore may not feed back into the research community.

Release results into the open literature, export students. This is the traditional approach for many research projects (transition through "broader impact"). Research ideas are made available and can be adopted by anyone who sees a beneficial use. Students hired by industry are often a key part of this process, as they bring the ideas with them. Benefits include advances to the research community generally as other researchers learn from open publications. Drawbacks include the difficulty of recognizing when a transition actually occurs and of assessing the value of the research publications.

Foster transition to practice through community building. Sometimes it is possible to build a community around a new idea or infrastructure, and that community can foster adoption of research results. The Internet Measurement Conference provides an example. Curricula and training materials can be a component of this approach. Difficulties including developing champions and building a critical mass of interested parties.

Release products under commercial licensing arrangements ("Idea transfer"). This approach can provide a shortcut to direct commercialization and can ease development of pilot implementations. It tends to be suited for smaller-scale efforts. Pitfalls include the risk that no one will license the ideas, sometimes weak support from inexperience university offices, loss of control of the idea, and trying to license intellectual property that is of the wrong type.

DG 4: What are the barriers to creating systems with security and privacy properties that users can understand and use?

Discussion Leaders: Alessandro Acquisti, Angela Sasse

Each group member identified up to 3 barriers, and each then voted for the barriers they considered most and least significant. At the end of this process, the top 4 barriers identified were:

1. **We don't know enough about the users**

Problem: We still have insufficient knowledge on what users are willing and able to do to for security and privacy. There is knowledge about human capabilities in human factors, and human behavior in psychology and economics, but it is difficult to make valid predictions because behavior depends on specific characteristics of users, the device they are using, the primary task they are trying to complete, and the physical and social context. We need more empirical research on security and privacy behaviors on specific tasks and contexts. Current studies are often so short-term and limited in terms of sample composition, and, sometimes, sample size, that designers cannot build on them. **Way forward**: We need generalizable results through replicable studies, as well as large-scale, long-term studies producing reliable findings that can be accessed by the research community. We need a testbed to carry out such studies – a 'PlanetLab+Humans' or DETER type facility. Having a large pool of potential participants, and standard briefing/consent/IRB procedures would facilitate better research.

2. **Misaligned incentives for system owners**

Problem: The group agreed that for privacy in particular, there is a gap between what service providers say they do, and what they actually do: "Privacy policies ≠ Privacy Practices." Users are not offered a genuine choice – they can either accept Terms & Conditions or not get a service. Once users have signed up, there is little feedback and no way for users to audit what has been promised**. Way forward**: (Apart from Shakespeare: Henry The Sixth, Part 2 Act 4, scene 2, 71–78). Move on from purely legalistic privacy policies to expressions that users can understand when they sign up, and audit during use. Provide technology that offers value (e.g. digital wallets) and use social pressure on service providers to implement them, e.g. unionize users and orchestrate campaign against coercive systems that offer no real choice or feedback

3. **Security and privacy depend …**

Problem: Even if a single stakeholder (user, designer, system owner) makes all the right choices, in practice, security and privacy depend on others stakeholders' expertise, behavior, motivation, diligence . For instance, one service provider leaking passwords or backup credentials can compromise other accounts. **Way Forward**: Develop robust, usable authentication without need for backup, and systems that don't disclose information about people without consent - even if data are published by others. Develop processes for negotiating (or learning) expected acceptable behaviors.

4. **Complexity of Systems**

Problem: Even when all stakeholders are well-intentioned, security and privacy are hard to get right – even by technical people *(e.g., TPM).* **Way Forward***:* Develop better abstractions and training, design patterns for security mechanisms. And the same spirit and practices of sharing and improving seen in the Open Source Community. A side benefit would be a single learning curve for end user. From an economic perspective we need a paradigm shift and to provide rewards and incentives for reuse.

DG5: What might a building code for critical infrastructure software/hardware look like?

Discussion Leader: Bill Scherlis

Despite the adoption of standards and best practices embodied in (for example) NIAP/CC, DO 178B->C, DO 133, FDA CDRH 514(k), the complexity, scale, and rapid evolution of many critical infrastructure systems thwart our attempts to assure security is built into them. Is it possible to create a "building code" for such systems that could provide or at least improve this assurance?

In considering the analogy with building codes, we found five areas of common concern: (1) *Engineering constraint*. Codes constrain engineering decisions and reduce options relating to both product structure and process model. (2) *Quality outcomes*. The constraints are intended to enable certain system qualities (electrical safety, fire resistance, energy consumption) to be predicted and enhanced. (3) *Visible evidence*. Code conformance additionally facilitates direct evaluation of artifacts by an inspector or evaluator both during development and after-the-fact. (4) *Support for response*. Despite code conformance, adverse events can occur. Support can include standard interfaces for responder systems, instrumentation, and logging. (5) *Support for ongoing evolution*. Building codes themselves must be able to evolve as materials, engineering , evaluation, and other practices advance, so that out-of-date practices are not mandated. This evolution is the product of negotiation, compromise, and consensus among diverse stakeholders (architects, builders, suppliers, safety officers).

A building code for critical infrastructure software and hardware systems might also need to differ in some significant ways from the conventional concept of a building code. Here are five significant ways where the normative concept of building code could be adapted to critical infrastructure systems:

(1) Rapid pace of technological change and diversity of sourcing. This factor suggests that codes should be based more on product-derived evidence that the developed system meets particular security requirements and relatively less on metrics of compliance with a prescribed process. Codes should motivate developers and evaluators to collaborate early in the development cycle to produce necessary evidence, including adapting development practices and product structures to accommodate this need.

(2) System scale and complexity. A focus on scale, complexity, and use of rich ecosystems unavoidably leads to composition as a dominant concern. Composition, in turn, often leads to more focus on rigorous mathematically-based approaches, often involving fragmentary specifications at component interfaces.

(3) Diverse and interacting quality and security attributes. While each attribute requires appropriate models, analyses, and composition approaches, the attributes also interact and can support each other in developing cases to support assurance claims. Quality attributes can support security outcomes.

(4) Hardware characteristics. Variances in hardware manufacture can necessitate evaluation of individual produced units, yet these units are often opaque to inspection. There may be opportunities to rethink the concept of trusted hardware and to address issues related to root of trust, extraneous functionality, and compromised reliability.

(5) Economics drivers of codes and compliance. Ideally, a code should provide incentives for compliance, in which the cost of complying is proportional to the importance of the attribute (in mission assurance, confidentiality, integrity) that compliance assures. As noted above, quality and security attributes often have complex interactions that make it difficult to apportion cost/value appropriately. Requiring a positive and traceable case to be developed in support of quality and security claims is a helpful way forward.

DG6: Models in cybersecurity research: what's used? What's needed?

Discussion Leader: Joshua D. Guttman

A model is a recipe for abstraction; or, equivalently, a set of decisions about what to ignore in a real world problem. A model carves out sets of entities and properties to study, and codifies explanatory principles. Models are therefore required to be able to acquire data, to make predictions, and to test hypotheses. Without models, there is no science, and precious little engineering either, since a model also furnishes the engineer with list of the measurable outcomes to achieve (or avoid), and of the techniques that lead to success.

Models for physical systems often take the form of differential equations. In cybersecurity, a threat model often focuses on defining some capabilities of the adversary, meaning that it codifies assumptions about what the adversary cannot do, and may also define the adversary's goals – discovering a secret or delivering a forged message. Game theoretic models allow exploring sequences of actions between an adversary and a system, helping to identify strategies that may be favorable to one party or the other. We also use the word "model" in more inclusive, less precise, but often extremely important ways. A model may be a tractable summary of a mass of data and experience; a set of guidelines for design (e.g. "identify and reduce the attack surface"); or a set of norms that predict human behavior and expectations, forming the basis of law or ethics.

The "big models" have largely organized cybersecurity research over the last few decades. These are notions like non-interference (information flow); access control represented in various frameworks (RBAC being a successful instance); the computational model of cryptography in terms of probabilistic, polynomial-time algorithms; and the symbolic, Dolev-Yao model of protocol behavior in the face of an adversary. Apart from the last, these big models played a very limited role in the discussion, and this was surprising and suggestive.

On the other hand, the importance of "little models" – single-use models designed to extract the right information to allow specific problems to be solved – was strongly recognized. Little models are we face. Examples include work on BGP configurations, to analyze route-flapping, in which rewrite rules transform a large network to a smaller one, such that the latter permits route flapping if and only if the original did. Continuous time state-machine models used to analyze stability of electricity distribution networks were also discussed. Work on authenticity and confidentiality in IPsec configurations uses a special purpose model. Special purpose models are also needed to represent the design choices of experts in setting up secure systems. Reducing a complicated natural-language text to a set of rules also means constructing a special-purpose model. Thus, constructing special-purpose models that enable extraction of key facts seems to be at the heart of cybersecurity research in many domains.

Security analysis, we conclude, means constructing a tissue of models. Systems have layers, and these layers often need layered models, often using quite different concepts. Composition of system components is similar although horizontal, in that it may require combining several models for the relevant components. In these systems, using a single big model is a likely to be too coarse a representation of the relevant facts of the components and layers.

There is also a potential risk in adapting small models for fine-grained security analysis. Attackers seek unexpected behaviors by looking for the points in a system where different models overlap and possibly conflict. These join points frequently offer security flaws. Thus we also encourage investigation of methods to ensure smooth welds at the points where little models interact.

DG7: Where does deconstructive security research belong in the research portfolio?

Discussion Leader: Fred B. Schneider

The term *deconstructive research* refers to investigations where the focus is on applying known attacks in order to compromise the security of a real system. Such work might lead to new insights that could be useful in subsequent computer security research and/or might call attention to risks that the artifact, which is being shown insecure, creates. Research is best disseminated through conferences and journals; risks to society are better disseminated in the popular press or communicated directly to some organization that can repair the problem that was found.

If we use the term *research* to refer to work that is intended to have a broad audience and to remain relevant for a long time, then there are clear instances where retargeting extant attacks to real systems constitutes research. These include:

- work that shows the need for new kinds of defenses (and, in the ideal case, new defenses would be proposed along with the attacks),
- work that illustrates new classes of vulnerabilities (perhaps due to new requirements or properties that need to be satisfied), being mindful that humans are often part of the system and can themselves be transformed into vulnerabilities,
- work that extends our understanding of the applicability for a class of attacks or and/or defenses (acknowledging that researchers are not always diligent about describing the applicability when proposing attacks or defenses), and
- work where the effort to retarget involves novel insights and is itself valuable research.

Presumably, the retargeting of an extant attack to a new system is a means to answer some question. Such work is valuable as research to the extent that it answers an important question and/or the answer is novel, but the work also should articulate a threat model and provide enough detail to be reproduced.

For societal impact, a report of some system's insecurity should not only articulate a threat model but should give a risk assessment that is objective and defensible or should provide information that allows a reader with enough knowledge of the deployment context to perform that assessment.

Some technologies already have organizations whose mission is to evaluate artifacts for vulnerabilities—Underwriters Laboratories is an example. Such an organization does not exist for computer security, and the academic research community has assumed that role. It is an important function, and we should be mindful of that reality. Yet we should also be careful not to conflate in our conferences and literature research efforts with work that primarily is concerned with identifying risks so that others will undertake appropriate action.

DG8: What policies and norms should govern the Internet commons in an era of cyberwar?

Discussion Leaders: Herb Lin, Alan Friedman

Unlike some other groups, our discussion raised very real policy questions before an expert group that did not, for the most part, have experience thinking about the theory and practice of international conflict and national security policy. However, the participants engaged wholeheartedly, and the exercise demonstrated the value of bringing technical experts into the policy discussion. Conversation ranged from whether a truly secure network was technically possible to how to define 'cyberwar,' but the discussion did produce some contributions towards future research.

After a discussion on what makes a Commons, the group agreed cyberspace that has some of the properties of a commons, including a lack of centralized control, no universal ownership structure and a high degree of cooperation. However, it was felt that the metaphor of the commons was less useful when thinking specifically about security. While a Commons can operate resiliently when a small number of actors defect against pro-social behavior, cyber attacks by a very small number of actors can undermine the global network. Thus, despite the power of a cybercommons for shared norms, it offers little to understand international conflict.

Like many conversations about cyberconflict, the group regularly returned to the idea of attribution. A discussion of the different types of attribution--the machine, the human, the political actor--highlighted a number of instances in which complete attribution may not be critical. Context matters a great deal. While false flag operations are an obvious threat from an untraceable attack, how would the response of Estonia or Georgia change if they knew concretely that the Russian state was or was not behind the denial of service attacks? The group explored how attribution might be improved, including nontechnical means such as the organizational structure of an attack. Perhaps the most important contribution computer scientists could offer to work with social scientists would be a better understanding of the time/utility trade off for forensics research. What does an additional hour of investigation reveal, and how is it useful? This could guide future forensics research.

Forensics is also useful in understanding the impact of an attack, or the damage assessment. On the offensive side, we know little about the certainties of a 'blast radius' of a cyber weapon, which can limit our ability to use them effectively. On the defensive side, the legal community lacks clear definitions and red lines to guide escalatory responses. Advanced modeling of cyber attacks might allow us to better understand attacks in technical terms, instead of using visible damage.

Another key realization of this group was a recognition that computer scientists excel as organizational innovators. In dealing with evolving threats, security researchers regularly interact and develop new institutions that are ad hoc, international, cross-domain, focused on specific issues while still remaining informal. A short term example might be the Conficker working group, which banded together to address a specific malware threat. More enduring organizations include the Anti-Phishing Working Group (APWG) which brings together relevant stakeholders in an evolving role.

DG9: How do research methods vary across the disciplines involved in cyber security?

Discussion Leader: Roy Maxion

PROCESS: The group established ground rules for the discussion, established a definition of "method" and discussed the posed question systematically, concluding with recommendations.

GROUND RULE: The chair should represent the group's view rather than his/her own.

DEFINITION OF "RESEARCH METHOD." Almost any empirical investigation (1) draws a conclusion, (2) demonstrates that a system (or other entity) does a particular thing, or (3) makes a claim. We will term all three of these as claims. Evidence in support of a claim almost always involves measurements of the form: this is better/worse; this is so; this is more/less; etc. The methods used to make the measurement are critical, and are (or should be) described in the "method section" of a paper. This section contains a detailed accounting of everything done to assemble and demonstrate the measurements that support the claim, as well as any factors that might have affected those measurements. A "method" in this context is any procedure used to measure or assess the phenomenon of interest. It's the method section of a paper that enables readers to judge the validity of the work.

ENUMERATE RANGE OF METHODS: (1) observational (which are mostly descriptive), (2) inductive (which comprise classical hypothesis testing and the scientific method), and (3) deductive (mathematical, including formal proofs). Each of these methods – approaches, really – has its own set of canonical evaluation procedures (or methods), but none were regarded to be common across disciplines; rather they had commonalities with respect to the kinds of problems being addressed. This observation pertains to formal and experimental methods, as well.

ENUMERATE DISCIPLINES IN CYBER SECURITY: A range of "disciplines" in cyber security was suggested (e.g., intrusion detection, malware detection, trustworthy hardware / software, behavioral biometrics, usability, privacy, crypto, data collection, criminology, forensics, economics, etc.), but these did not coalesce into obvious and unambiguous groups, or have methodologies particular to any one of them. Most of the areas in cyber security are inherently interdisciplinary, and hence share research methods that are particular to the problem being addressed or the phenomenon being assessed, rather than having a given method more or less uniquely associated with a discipline.

DISCUSS RECONSIDER QUESTION AND VIEWPOINTS: After much discussion aiming to discern which disciplines align with which methods, the group concluded that particular research methods are well matched with different phenomena under investigation, and that the same kinds of phenomena could be investigated across a range of disciplines. From that point of view, no clear answer could emerge for the question put to the group. While there truly are different research methodologies, the term seems to pertain more to the method by which a phenomenon of interest is measured or characterized.

MAKE RECOMMENDATIONS TO NSF & to PIs: The group felt strongly that research methods, in the sense defined above, are lacking in most cyber security research. The group observed that courses in research methods are typically not offered in disciplines such as computer science, software engineering, and the like. The group recommended that: (1) a straw research methodology, that can be adapted to individual situations, should be provided to the community as a model to follow; (2) research methods should be a required course (as it is in many disciplines outside the computer sciences); (3) publication traditions should be changed to require a method section in every empirical paper (and review criteria should reflect this).

DG10: What modeling techniques should we use to account for the role of humans in complex cyber systems?
Discussion Leader: Rick Wash

Defining modeling of humans in cyber security. The group began by discussing what modeling means in cyber security. Modeling the role of humans in complex cyber systems can be thought of in several ways as follows. [Note: For the purposes of our discussion, we included four types of *humans*: (1) computer users (individuals and organizations), (2) attackers (individuals and organizations), (3) software producers, and (4) security professionals. We focused on computer users and attackers.]

Models can also be thought of as being developed in several ways: (1) a model developed based on theory (e.g., model models, cognitive processes), (2) a model developed based on experimental data, or (3) a model developed based on observational data. Further, a model that is developed based on theory can be tested or validated with experimental / observational data, and models developed with experimental or observational data can be used to develop new theories.

Models can be static and dynamic. Models can be focused on the describing status quo/current behavior or focused on predicting factors involved in behavior change. The research goal should be clear and should result in a model of one of the following types: descriptive, normative, prescriptive, generative, or introspective.

Examples of types of cyber security modeling of humans. We talked about a variety of specific models that have been used and should continue to be used to study the role of humans in cyber systems: (1) cognitive/mental models as a way of thinking – users, attackers, developers, etc.; (2) models of individual agents/actors (micro) – both rational choice theory models and behavioral economic models (e.g., bounded rationality, loss aversion, prospect theory); (3) models of a system of agents/actors (macro); (4) models of the interface between technology and humans; (5) models of network traffic aggregated to the individual user level; and (6) model of types of threats. Models can be qualitative (e.g., cognitive or mental models), quantitative (e.g., computational models), or a combination.

Several participants noted that models often do not clearly state the assumptions that they are making or the primary research questions of interest. Both assumptions and research questions should be clearly stated. Further, there was concern that model results are being used to make decisions even though the models are not sufficiently sophisticated; others noted that the model results may provide better input to decisions that not have the model results.

Recommendations for future research.

- More modelers should be encouraged to reuse their models, build on them, share them, and combine them with other researchers' models.
- Modelers should leverage other disciplines – e.g., climate change and epidemiological models.
- New models should be developed of sophisticated attackers who conduct their own research (e.g., what data do they collect before attacking?).
- New models should be developed of the value of victim characteristics to attackers (e.g., bots).
- More models should be developed of the status quo / base case (for all actors/systems/threats).
- New models should be developed of users' ideal online activities (design security around users).
- More models should account for the evolutionary nature of threats and actors.

DG11: Predicting the next "flash crash" or blackout: What methods are available for evaluating the stability/trustworthiness of complex digital infrastructure systems (CDISs)?
Discussion Leader: Bill Sanders

CDISs are heterogeneous systems-of-systems under large-scale distributed control involving large volumes of data and significant human-in-the-loop presence. They have the potential for highly dynamic, stochastic and chaotic behavior. Mostly autonomous, CDISs leverage the Internet and/or substantial private Intranets for communication. Examples include the smart grid (spanning the physical power system, its marketplace, and its communication/control infrastructure (SCADA)), financial systems, and smart transportation systems (e.g., smart cars, airline systems).

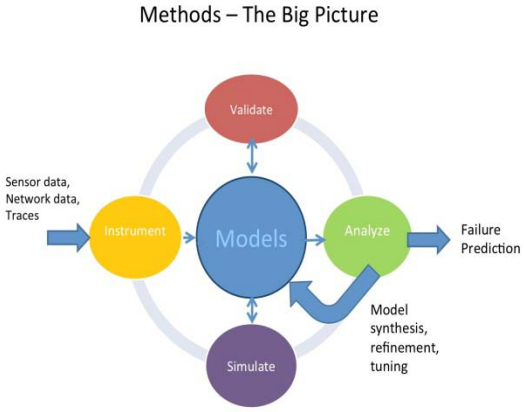

A naïve answer to the challenge to the question is that there are no methods. The reality is, however, is that practitioners are engaged in detecting and predicting stability and/or trustworthiness of CDISs all the time, e.g., through active monitoring of the transmission portion of the power grid and conducting thus informed, and there is a strong community of practice to build upon to create better and more accurate methods. Failures and security breaches in the recent past show there is substantial work still to be done. The group addressed issues of instrumentation (e.g., placement strategies for lightweight sensors into legacy, aging infrastructure), model synthesis (hybrid discrete and continuous time and discrete-event, including feedback control of a hierarchically decomposed structure), model solution by analysis (e.g., formal models, differential equation solvers, graph-theoretic analysis) and simulation (including trace-derived experiments for stress testing), and producing predictions and validating them. A table of providing details regarding these approaches can be found to below.

We identified three research-challenge areas where existing methods are limited. 1) Scaling techniques for isolated systems to complex systems-of-systems, in particular issues of: scalability of predictions, scalability of the computational infrastructure, and integrating federated models from different disciplines (CS, Econ., Finance, Social). 2) CDIS model creation spanning a very large space of available techniques (thus begging the question, Which is most applicable to the specific complex system?) and identifying critical variables from the very large available space. 3) Model validation, because hypothesis driven experimentation is difficult without a science of complex systems, and validation on a full-scale CDIS is not possible; to this end, past catastrophic events can be recorded and used for validation.

Instrumentation	• Placement strategy for lightweight sensors • Use of mobile sensor agents • Privacy preserving collection and storage • In-line processing
Model Synthesis	• Hybrid state-space (cont. and/or discrete-time) with discrete-events • Known "normal" operating range and well-defined decision functions for anomalies, attacks and faults (known vulnerabilities). • Incorporate automated feedback control (including containment actions), human factors and economic considerations (e.g., through game theoretic framework) • Hierarchical, decomposed structure
Analytical Model Solution	• Solvers of formal models including both hybrid state-spaces and discrete-events • Differential equation solvers • Graph-theoretic models, link analysis of interconnectedness (suitably weighted to account for trust/reliability, priority), etc.
Simulation Model Solution	• "Brute" force simulation at scale for, e.g., Chinese power system and financial markets – cf. combinatorially big computational challenges • Test how systems would handle hypothetically abnormal (e.g., attack/defense) situations • Trace-driven experiments for stress testing on-line/off-line at system level
Analysis (Producing Predictions)	• Decision makers act on mixture of experts: model, simulation, human operators
Validation of Predictions	• Cross-check simulation and analytical modeling result in terms of behavioral details • Expert prediction consensus, comparisons against "real world" historical data

DG 12: Anonymity and Accountability: How Do We Enable Tradeoffs?

Discussion Leaders: Jeannette Wing, Rebecca Wright

Anonymity and accountability often seem to be opposing requirements. At different times and in different contexts, either anonymity or accountability requirements may dominate, or some intermediate ground may be sought. This discussion group addressed how we can enable such tradeoffs, as well as when and whether the apparent conflict can be circumvented. A motivating scenario is that of a Healthcare chat room. People should be able to participate anonymously in the chat room in order to enable safe and comfortable discussion their health concerns. Furthermore, sometimes posting photos can help convey information. However, this creates opportunities for inappropriate users and uses, such as people posting child porn photos in order to attract potential customers and other kinds of spam messages. In such a setting, it would desirable to have accountability to deter such misbehavior, without compromising anonymity for appropriate use.

We noted that accountability and anonymity need not always be in conflict. In particular, through the use of cryptographic techniques with revocable anonymity, it is possible to design systems so that participants can participate anonymously in general, but their anonymity is revoked if they break particular rules. A challenge in actually using such systems is recognizing that anonymity can be broken in a variety of ways in practice. Even if the application layer provides anonymity, if the network or the hardware device being used does not, then the overall system does not provide anonymity. Anonymous communication systems (notably Tor) can provide network-layer anonymity. In some cases, the existence of multiple layers can be used to provided anonymity where it is desired while requiring authentication — for example, allowing a user to reach his or her Gmail account without disclosing his or her network location to Gmail or to the network, but still authenticating to Gmail.

We identified a number of open research questions:

- Definitions and metrics for anonymity, accountability, accountable anonymity, and anonymous accountability.
- How do we balance accountability and privacy when different kinds of participants have different constraints/policies (e.g., voting systems, online privacy policies)?
- How much anonymity can you get as a function of the power of the adversary to control the network? (and possibly of efficiency of solution)
- efficient cryptographic solutions for a larger class of policies, particularly more flexible policies, dynamic policies, and policies that may depend on private data.
- Retaining anonymity or unlinkability at all in today's world (where identity can be leaked or partially leaked by devices, applications, network, users…).
- How strong are the accountability and anonymity properties that can be achieved, relative to the cost of obtaining an identity (or a credential)? Can we change the costs to effectively dissuade bad behavior?
- Design anonymity system that enables data mixing for utility goals (e.g., deviation detection, pricing, targeted advertising while preserving privacy)
- How to enable users to make informed decisions about anonymity-accountability tradeoffs? Can the informed consent model be useful here?
- Where must regulatory frameworks supplement technical approaches?
- Do users care, and in what contexts? How much do people value their identity?

Finally, we noted a need to better educate policy makers and system designers about what is technologically possible.

DG13: What policies and technologies would be required to enforce the expiration of data?

Discussion Leader: Tadayoshi Kohno

We began by discussing metaphors for data deletion. The goal was to develop an understanding of where everyone was coming from and what data deletion means to them. This metaphor-generation phase also served as a foundation for helping define the goals and parameter space for data deletion. We then moved on to explicitly defining use cases for data deletion, and exploring opportunities for and challenges with data deletion.

The group concluded that data deletion is a very challenging problem. There was no consensus on a technical or policy approach that would completely solve the data deletion problem. However, the group did produce two main contributions:

1. A set of explicit situations (data use cases) for which we may desire data expiration;

2. A list of "axes" for the problem / solution space.

The group argued that progress for any one of these use cases at specific points along these axes could be a valuable contribution to research and society.

Elaborating on (1), the identified data use cases include: Corporate email (internal to a company); corporate email (between companies); web mail (for public users); laptop and phone data (e.g., data on a lost device); social network data; captured public data (e.g., drones, ATM cameras); health records; financial records; childhood records (e.g., some records may disappear when a person turns 18); death (e.g., some records may disappear or reappear when a person dies); deletion as a solution for account compromise (e.g., all data disappears after an account is compromised); sexting; individuals' digital media (e.g., personal photos); industry-produced digital media (e.g., published movies); and criminal uses (criminals may wish for data to disappear too).

Elaborating on (2), the axes include: Consumer versus corporate data; clean versus comingled data (e.g., consider an email that contains some information that must persist for years *and* some information that should disappear immediately); structured versus unstructured data; trust in second parties (e.g., Google), third parties (e.g., ad networks that buy data from Google), and other first parties (e.g., Alice and Bob, if they're communicating via Gmail); prevention vs. auditing (the group discussed whether there would be value in *not* requiring data providers to prove to users that they have deleted data but instead have third-party entities regularly audit data providers); desired lifetime of data (e.g., seconds, days, years, forever); type of data (data about users that is not shared with those users, data authored by users, and data that is shared with users); who implements the mechanism (client only, client + cloud, cloud); incentives and disincentives for adoption (e.g., economic, government, ethical); role of policy in the solution (e.g., none or an explicit law); and the degree to which people's preferences will change over time.

The slide deck for DG13 also includes some possible directions that the group thought might be worthy of additional study. For example, one person wondered whether it might be possible to enact a law that says that if a user clicks on "delete," then the data is legally deleted even if the data technically persists on a server.

DG14: How can we assure provenance, integrity, longevity of scientific records?

Discussion Leaders: Ravi Sandhu, Elisa Bertino

Provenance is a relatively new topic for the SATC community. The discussion group brought together participants with diverse backgrounds and perspectives on this emerging topic. The result was a lively and informative discussion with some very pragmatic conclusions.

It was noted early in the discussion that an ideal solution for preserving provenance would be a time machine that allows us to go back and replay the past, so to speak. Given the impossibility of this solution, the next best approximation would be to record all context with every new data item, whether the context is relevant or not. Even if one could theoretically capture all context in its entirety, this approach would be highly impractical and potentially drown us in vast amounts of irrelevant provenance information, just in case it might be useful. What then is practical? The consensus was that provenance-aware systems should be designed to capture what is relevant for the purpose we want to use provenance information. This would require provenance to be considered all through a system's lifecycle staring with its initial conception.

Turning to the issue of provenance with respect to scientific data, the group identified three major threats. First, and most significant, is the use of bad data resulting in bad science without any bad intent as such. Reducing this threat alone would justify a major effort in making data driven scientific research provenance-aware. The second threat is deliberate scientific fraud by insiders. While the vast majority of scientists adhere to high ethical and professional norms in their research, there are enough bad apples out there to make this a serious threat. The third threat is deliberate mischief by malicious outsiders.

Finally, the group identified challenges to building provenance-aware systems for the scientific enterprise as follows.

- Scientific data manipulation processes are complex and computer scientists need to be careful not to oversimplify in abstractions.
- Provenance data is big, and bigger than the underlying data.
- Usability and adoption by scientists of provenance-aware systems is necessary for their success but not easy to achieve due to cultural gaps between different communities.
- Provenance capture must be automated including human-in-the-loop when appropriate.
- Sensitive data will require some degree of anonymization, e.g., medical data, sociological data and cyber security data.

In conclusion, the group agreed that provenance of scientific data is major research opportunity for the NSF SATC community. The principal goal is to enable good data-based science which is one the central goals of NSF looking forward. It is inherently interdisciplinary involving scientists from almost any discipline as all disciplines evolve in a data-driven direction. Computer scientists can make a material difference in the conduct of science at large by enabling development of provenance-aware systems with strong computer science foundations. The group had strong consensus that center-scale funding was appropriate to advance provenance-aware systems for data-driven science.

DG15: Identity Management: Why So Slow?

Discussion Leader: Susan Landau

Identity management provides convenience for the user, increased security (one strong authentication method instead the user storing multiple weak passwords), increased ability of enterprise to control and understand use of its resources, and the ability of systems such as Twitter or Craigslist to separate “real" responses from manufactured ones. Do we actually need identity management? Yes. Adoption of identity-management systems, which began in the 1990s with the Department of Defense (DOD) Common Access Cards (CAC), has been slow: why?

Environment. The CAC system is successful in a constrained environment: DOD. InCommon operates identity federation for universities and related partners at "Internet" scale. But InCommon is also in a constrained environment, one in which there are contracts between all participants and the federation operator (InCommon). At the other end, there are Internet-scale identity providers, e.g., Facebook or Google. Successful as businesses, these systems fall far short in providing strong identity assurance. The U.S. government is seeking a flourishing set of Internet-scale identity management systems. Support includes funding for the development of privacy infrastructure for federated identity management systems. What will it take for Internet-scale systems to succeed outside constrained environments?

There is a natural conflict between accountability and anonymity. At one level, this conflict is a policy concern. At a finer grain, one can see that while, for example, an IP address does not identify a person, it may be “good enough" for some purposes. Deception and obfuscation will always be used in identity systems, including those with high levels of assurance; how can/should an identity system handle the issues raised. Or does it need to? Reputations are sufficiently important that while anonymized handles provide some anonymity, they nonethless provide a fair bit of accountability. Understanding those tradeoffs is a research question at the intersection of technology and social science.

We need to understand the value of data better in order to better understand the levels of assurance that should be required in data-sharing environments (e.g., use of resources at supercomputing centers). NIST has done some work in this area in its 800 series, but more work is needed to determine how to perform scientifically rigorous risk assessments, an issue that transcends identity management.

We need to determine what system design enables federated systems to please *all* stakeholders, critical for broad acceptance of robust identity systems in unconstrained environments. What are the economics of anonymous credentials? We know the cost of privacy spill, but we need to understand the other side of the equation. What do I, and the system, get from anonymous credentials? Is adoption of federated systems necessarily a set of tradeoffs (economic, privacy, and political/policy)? If so, what set of incentives will enable broader acceptance?

Critical to adoption is making systems genuinely usable, including making it simple for users to control the distribution of information about personal attributes. User interface issues are key. What does the user need to know about attribute release, especially in order to control these? And finally, how does the system trust the user (what the user does to get trust)?

Identity information acts almost like a layer “below" applications. In a world where metadata sharing has become increasingly rich and increasingly dynamic, what changes does that force on identity federation? That is a both a technical issue and a social one – as are most of the problems of deploying federated identity management at Internet scale.

DG 16: How Can We Leverage R&D Work Done to Improve Cybersecurity Education?

Discussion Leader: David Balenson

Questions discussed included: **"How to identify key concepts during the research phase?"** We identified some guidelines, though not specific means, for identifying key concepts, including the desire to be open and inclusive, to allow for niche topics as well as broader concepts, and to include both fundamental and applied concepts. Not all research projects will yield concepts that need to be taught. The ACM/IEEE Computer Science Curriculum provides a taxonomy researchers might use to identify relevant concepts. Key concepts should be useful in more than one domain. Security trends and industry practice could be a guide to form concepts that might be of use or interest in the classroom.

"How to make developing educational materials an integral part of R&D activities?" We recommend creating an "open source" community for educational materials in which materials produced from research are freely available for use by others. They would peer-reviewed, cross-tested, and maintained. NSF could offer additional funds for contributed material that is adopted by others and/or for adopting and evaluating material provided by others. The funding should be in addition to the research grant, in order to maximize the incentive to participate. A (semi-) standard format for the materials would provide guidance and structure, without limiting creativity, and a unified set of platforms for lab exercises would help.

"What types of educational materials could be produced?" Types include lecture materials (including reading materials, lecture notes, slide decks, videos, animations, etc.); hands-on student labs and exercises (with instructor manuals); and tests and quizzes for assessing student learning. Other materials might include data sets, source code, and case studies.

"Where and how to collect educational materials to make them available to others?" Educational materials from research projects will likely fit best in graduate and upper-level undergraduate classes. Materials tailored to other educational environments, including lower-level undergraduate classes, K-12, games and competitions, workforce development, and online learning are also needed. We recommend NSF support a community wiki or repository to collect and disseminate educational materials, including a catalog with metadata describing various educational materials. In addition to source materials, the repository could provide links to other repositories (e.g., to exercises collected on the DETERlab educational portal). NSF and the community would publicize the repository and promote its use.

"How to evaluate the quality of education material or learning based on the material?" In addition to normal feedback from users, we envision "crowd sourcing" evaluation approaches to leverage the "community" aspect of the wiki/repository, such as reputation-based scoring (e.g., 1-5 stars), +1/-1 from users like Amazon reviews, online reviews, etc. The repository could collect various metrics regarding contributions and adoptions of educational materials. To evaluate the quality of the learning from education materials, students could be tested both before and after the learning process, or the evaluation process could be integrated into the materials and learning process itself (feedback-based learning). We also imagined surveys of both educators and students, including following students and conducting surveys some number of years (e.g., 1-2 years) after taking a class with research-based educational materials, to gauge the impact it had on their ultimate knowledge.

DG 17: How can the nation best build and sustain an appropriately sized and qualified cybersecurity workforce?

Discussion Leader: Deborah Frincke

Theme 1: How to prepare students for the cybersecurity workforce?

- Balance the need for breadth vs depth
- Train for specific roles, non-specialist + specialist
- Approaches to scale the workforce (all needed depending on specific role)
  - Big cyber programs vs. Pervasive programs vs. Experts + high level
  - Delivery options for scale: cost issues for delivery, online courses, sharing of materials/tools; MOOCs; alternatives to articulation agreements (encourage diversity + class sharing)
  - The skills/background needed is broader than STEM:
  - Need to consider critical thinking, adversarial thinking, risk analysis, economic models, communication with domain experts
  - "fixing" skills vs. "breaking" skills vs "transforming" skills
  - practical + formal/theoretical skill

Theme 2: How to raise cybersecurity awareness across disciplines? Pervasive approach requires adding security / security literacy courses to basic undergraduate curriculum (just like required math and communication courses).

- Issue: universities resist; seems can only be done at the expense of other courses like programming
- Can this be added to CAE (Center of Academic Excellence) requirements?
- Need to get ABET on board if you want to experiment with the curriculum; ABET metrics don't support new courses; no way to mark course as "in progress" or "experimental"
- Cybersecurity awareness effort should start earlier (K-12)

Theme 3: New generation of students and new workforce leads to the question: will we age "out of" or "in to" the cyber problem?

- Familiarity breeds complacency (new set of "worst practices")
- **Recommend:** Revisit body of knowledge from new perspectives (through NSF programs?) in the context of what makes sense in a curriculum (e.g., capacity building grads)

Theme 4: Diversity of pipeline/graduation/hiring

- Effects of "competitions" on particular populations: it seems to disproportionately draw in only certain populations and repel other populations
- **Recommend:** Identify and invest in a parallel program that counterbalances these themes. Needs to be something exciting and intriguing.
- There is anecdotal evidence that being part of a team that participates in a competition seems to help (rather than individual participation)

Additional Recommendations:

- Attract best and brightest to "SFS" program through: 1) Reward based on potential; 2) Fellowships + postdocs; 3) "National merit scholars" for high potential STEM then take this to a program with a security aspect. No service required. Optional summer programs?; 4) Internships + co-ops; 5) MS funding for research (not just PhDs and ugrads); 6) Forgivable education loans for diversity candidates choosing "service" (similar to Teacher Corps program)
- Invest in support for faculty in teaching (+developing) security literacy
  - Ready modules are good here to provide to instructors
  - TA support (this could serve as incentives to help adopt security literacy courses as part of basic undergrad curriculum)
- Create "Security Guru" program, similar to adult education Master Gardner program

DG18. What issues are unique to cyber warrior education (compared to other members of the cyber workforce)?

[Session canceled due to travel restrictions placed on leader]

DG 19: What incentives, norms, attitudes, habits, cognitive limits, or other mechanisms present the most important obstacles to cybersecurity, and how might such factors be utilized to benefit cybersecurity?

Discussion Leaders: Kevin McCabe, Sandra Carpenter

Most users of cyber-devices are typically engaged in a task during which they do not expect a hazard, and their task is more salient than security or privacy. Habits are therefore likely to influence their interaction with such devices. In order to increase cautious behavior, developers need to be cognizant of economic and of social and cognitive psychological processes that may impact behavior. Current models in social or cognitive psychology and in economics can therefore be leveraged to increase the trustworthiness of cyberspace, with the following caveat: much of the extant research shows that behavior is context-specific and domain-specific. Thus, age, knowledge, experience, culture, and other factors influence security and privacy behaviors. *This context-specificity of human behavior is potentially the most important obstacle to cybersecurity*. We therefore need research that focuses on identifying the degree to which current theories can inform cybersecurity developers.

With respect to attitudes, security and privacy attitudes are related to trust and perceived risk. Most of the time, however, the actual risk of disclosure is unknown. What *are* the costs of revealing my name and address online? Companies sell consumer information to other entities, but how these entities might use this information is not clear. Research needs to identify conditions under which trust can be called into question and under which perceptions of risk are accurate. People may be willing to take risks if incentives are involved. From social psychology and economic perspectives, therefore, the range of incentives for engaging in potentially risky behaviors needs to be identified. These incentives can be monetary, but could also be related to users' goals (e.g., acquiring information or services). Social norms can strongly impact behavior, and norms can change across time. For example, few people currently throw trash from the windows of their cars while driving. Potentially, children and adolescents could be targeted for learning to behave more appropriately with respect to security and privacy.

Research on users' cognitive abilities and limitations can provide designers with information about how to design privacy and security systems. Human factors research can address, through usability analyses, how users interact with these systems. It is unlikely that most users' mental models of security and privacy match extant threat models. Users may have poor or incomplete knowledge of the types of risk they may face, the likelihood of those risks, and the severity of the risks. Users' cognitive load and stress levels can reduce their attention to details other than are relevant to the task they are attempting to undertake, such that security and privacy issues are ignored. Moreover, users may lack the ability and/or willingness to attend to, comprehend, or remember security or privacy mechanisms (e.g., long passwords). Thus, designers should take into consideration an array of cognitive processes when designing systems.

The SBE literature, as it relates to cybersecurity and cyberprivacy, is not well-organized or accessible. Relevant publications in social science may appear in "marginal" journals, such that the visibility of this research may be low. The creation of an infrastructure to warehouse, constantly update, and disseminate SBE research conducted in the cyber-domain (e.g., CyLab at CMU; Ross Anderson's web pages) would aid researchers in this area and make the existence of this research more salient to social scientists. To recruit SBE researchers to this research area, the Federal Cybersecurity R&D Strategic Plan could also explicitly call for social and cognitive psychology research, in addition to the economic research that it currently explicitly addresses.

DG 20: What are the group, organizational, institutional, and policy obstacles to cybersecurity?

Discussion Leader: Andrew Whinston

Obstacle	NSF Research & Education Directions	Sample Research Questions
Cyber-risk has not been quantified	Cyber-insurance	What markets exist (or how can they be developed) to insure un- or under-mitigated risks? (e.g., HMO model?)
	Risk modeling	What is the cybersecurity equivalent of actuarial tables for life insurance? Can we introduce dynamic warranties that cover a portfolio of security tools provided by the "HMO"?
	Standards and certification	Can they be leveraged to support risk assessment and guarantee insurability or lower insurance premiums?
	Policy, regulation, and liability	Should cybersecurity insurance be regulated in ways similar to banks and auto insurance? What is the legal framework for liability and accountability? How do the insurance and liability models change as software and data move into the cloud?
A gap in understanding of privacy and security user/behavioral models exists	Development of behavioral models	What are the psychological and sociological models of risk recognition, assessment, and related decisions?
	Understanding of usage patterns	What is the interplay between how/why a system is used and the dynamics of individual risk decisions (e.g., Facebook vs. online banking)?
	Characterization of individual and group differences	What individual/group factors influence risk-related behavior? Age, gender, race, education, …?
Insufficient formal cybersecurity education	Optimal educational delivery mechanisms, e.g. general education requirement, interdisciplinary education	How should cybersecurity education be extended to "the masses"? (e.g., general ed requirements in cybersecurity?)
	Online education, e.g. Stanford model	How can the MOOC model be adapted to initial and continuing education for cybersecurity professionals and others?
	Assessment, evaluation, and effectiveness of the educational design, content, and delivery	How do we assess and improve the effectiveness of cybersecurity education (e.g., within a MOOC framework)?
Misinformation on cybersecurity and lack of understanding of the consequences of inaction.	Communicating with policy makers	How do we increase awareness where it doesn't exist and promote deeper understanding where needed?
	Awareness for a wider audience, e.g. YouTube	What are the most effective mechanisms for reducing public vulnerability to cybersecurity attacks (e.g., a video entitled **Blood on the Information Superhighway**)

Carl Landwehr

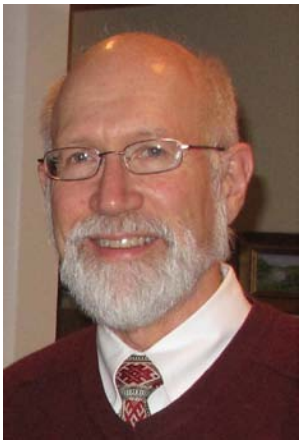

Carl Landwehr is a Lead Research Scientist at the Cyber Security Policy and Research Institute at The George Washington University in Washington, D.C. Dr. Landwehr has more than 35 years experience in computer science research, focusing primarily on issues in computer security, information assurance, and trustworthy computing, and in research management, research funding and program management.

Dr. Landwehr has assisted, developed and managed research programs in cybersecurity for the National Science Foundation (NSF), DARPA, and the Intelligence Advanced Research Activity (IARPA) and its predecessor organizations, the Advanced Research and Development Activity (ARDA) and the Disruptive Technology Office (DTO). In addition, Dr. Landwehr previously served for 23 years as a researcher and section head at the U.S. Naval Research Laboratory, managing a small group of researchers developing concepts and prototypes in security modeling, high assurance software, secure system development, database management system security, and token-based authentication.

Dr. Landwehr's professional activities with the IEEE include serving as Chair of the Technical Committee on Security and Privacy, Chair of the IEEE Symposium on Security and Privacy, and a four-year term as Editor-in-Chief of IEEE Security and Privacy Magazine, the leading peer-reviewed technical magazine in this field. He is a Fellow of the IEEE and was in the first class of eleven people inducted into the National Cyber Security Hall of Fame in 2012.

Lance Hoffman

Lance J. Hoffman is Distinguished Research Professor of Computer Science and Director of the Cyber Security Policy and Research Institute at The George Washington University in Washington, D. C.

Professor Hoffman developed the first regularly offered course on computer security at the University of California, Berkeley in 1970 after serving on the Advisory Committee to the California Assembly Committee on Statewide Information Policy. He has authored or edited numerous articles and five books on computer security and privacy. His teaching innovations include multidisciplinary courses on electronic commerce and network security and the development of a portable educational network for teaching computer security. He co-directed the workshop that led to the Collegiate Cyber Defense Competition which now has over 100 participating universities. He pioneered a holistic, multidisciplinary approach to teaching and research in cybersecurity, and currently directs the Department of Homeland Security, Defense Department, and National Science Foundation computer security scholarship programs at GW.

A Fellow of the Association for Computing Machinery (ACM), Dr. Hoffman institutionalized the ACM Conference on Computers, Freedom, and Privacy. He has served on a number of Advisory Committees including those of Federal Trade Commission, the Department of Homeland Security, the Center for Democracy and Technology, and IBM. He has chaired the Information Security Subcommittee of the IEEE Committee on Communications and Information Policy and is a Member of the Subcommittees on Law, and Security and Privacy of the U. S. Public Policy Council of the ACM.