THE Next Wave

The National Security Agency's review of emerging technologies

Vol. 19 | No. 4 | 2012

Building a national program for cybersecurity science

Globe at a Glance | According to the Experts | Pointers | Spinouts

THE Next Wave
The National Security Agency's review of emerging technologies

GUEST Editor's column

Frederick R. Chang, PhD

Considered by most to be the first computer worm ever, the Creeper worm was written over 40 years ago. Unlike today's worms and other malicious code, Creeper was not written with malicious intent, but rather as an experiment in self-replicating code. It spread through the ARPANET—a precursor to the modern Internet—by "jumping" from machine to machine, and it caused an infected system to display the message: "I'M THE CREEPER, CATCH ME IF YOU CAN." In response, the first antivirus program, Reaper (itself a computer worm), was created.

Back then it would have been nearly impossible to predict how dependent we would become on modern networking and computing infrastructure. As a sign of our increasing dependency on modern networking, this issue of *The Next Wave (TNW)* as well as future issues will be available primarily electronically instead of in print. As with commercial publishers, the federal government is finding the incentives to move from a print publication to an electronic publication irresistible—increased audience for lower cost.

It would also have been nearly impossible to predict the difficulty of defending the modern infrastructure. Early research on computer security had already begun by the time Creeper was spreading through the ARPANET. Yet, after over 40 years of research and development on computer and information security, we find ourselves searching for fundamental answers on how to secure systems in cyberspace. This existing research base has yielded important and significant findings through the decades, and computing systems are unquestionably more secure as a result. There is, however, an increasing awareness in the cybersecurity community that the research has not produced a consistent scientific understanding of cybersecurity and that such an understanding is now urgently required.

This issue of *TNW* is the second of two issues dedicated to the science of cybersecurity. The first issue, published in March of 2012, included contributions from experts primarily from academia and the private sector and offered an impressive collection of insights that touched on a wide range of perspectives on the problem, from technology to policy to strategy and more. This second issue includes contributions from experts within government (US and UK) and offers a wide array of perspectives on the problem as well as activities under way to develop and implement solutions.

There are some promising indications that a science of cybersecurity initiative is gaining momentum, including several workshops, conferences, and reports that point to the need for an interdisciplinary approach to addressing the problem. Most recently, in November of 2012, NSA sponsored the first annual Science of Security Community meeting to discuss issues foundational to the advancement of a science of cybersecurity. This issue of *TNW* provides additional detail on some other notable activities taking place both inside and outside of government.

The theme of interdisciplinarity is important. Indeed, there is evidence that scientific advances often occur at the boundaries of established but related fields, when scientists from different disciplines address a problem free from the ordinary constraints of working in a more intradisciplinary fashion. A science of cybersecurity offers many opportunities for advances based on a multidisciplinary approach, because, after all, cybersecurity is fundamentally about an adversarial engagement. Humans must defend machines that are attacked by other humans using machines. So, in addition to the critical traditional fields of computer science, electrical engineering, and mathematics, perspectives from other fields are needed. Cognitive science will help us understand adversarial intent and human decision making under uncertainty in cyberspace. Economics will illuminate how misaligned economic incentives hamper fundamental progress in cybersecurity. Biology will shed light on the extent to which it may be possible to transfer concepts from our understanding of the human immune system toward the conceptualization of a cyber immune system. Thinking from other scientific disciplines will offer perspectives that will trigger new, valuable ideas.

Progress in this new science will be unpredictable, uneven, and slower than we want. We will need to be patient. Cybersecurity research experts will have to resist the urge to focus their efforts on the cyberattack of the day. We will need our research scientists to help us understand not only what is possible, but also what is not possible. Indeed, a rigorous understanding of the limits of cybersecurity will be fundamental to the formation of the new science. We have learned much about how to defend computing systems since the first computer worm, but now we must advance our understanding through the creation of a disciplined and systematic science of cybersecurity. We cannot wait any longer; there is too much at stake.

Frederick R. Chang

Former Director of Research, NSA

Contents

The Next Wave is published to disseminate technical advancements and research activities in telecommunications and information technologies. Mentions of company names or commercial products do not imply endorsement by the US Government.

Vol. 19 | No. 4 | 2012

An introduction by General Alexander

A previous issue of NSA's *The Next Wave* magazine provided academic perspectives on what a cybersecurity science might look like. This follow-on issue focuses on the government's response to this topic by describing how various organizations, individually and collectively, are addressing the challenges of developing a true science for cybersecurity.

The past several decades have witnessed the phenomenon of a fledgling military computer network transform into an essential national and international information infrastructure that has fueled the growth of the global information age. This new infrastructure, often described as cyberspace, has already taken its place alongside long-established infrastructures, such as the national transportation system, in shaping society and reshaping governments.

The rapid acceptance and pervasiveness of this information technology, and cyber technology more generally, has come with a significant cost. We see evidence of that cost on almost a daily basis, and often with spectacular consequences. The ongoing cyber-thefts from the networks of public and private organizations, including Fortune 500 companies, represent the greatest transfer of wealth in human history.

While the need for cybersecurity is widely recognized, current views and definitions of security differ greatly. Commercial-world cybersecurity implements new security measures in reaction to new cyberattacks in an unending arms race. The discipline of security engineering implements best practices to build less vulnerable cyber systems, but security failures often arise in spite of compliance with best practices. Both approaches seek to secure known vulnerabilities of systems against attack. But, the systems and the cyber environment are dynamic, not static, and new vulnerabilities arise. Security fails in this dynamic environment when the adversary simply changes the game by exploiting new vulnerabilities. Adversaries have the easier job, and they can expand their methodologies and techniques to acquire significant power in cyberspace with relatively modest resources.

The ball is now in our court.

In recognition of cybersecurity as a national priority, the US Cyber Command was chartered to protect our national interests in cyberspace. Although support for this national initiative is gaining ground, it is imperative, going forward, that we broaden our understanding of the science that underpins cybersecurity. We must form collaborative public and private partnerships and devote more attention to understanding security science. And it must be a team effort with the DoD, FBI, and DHS working together for the benefit of the nation. For decades, NSA has invested heavily in cryptology, but because our nation's current security challenges involve so much more than cryptography and cryptanalysis, we will lead the effort to broaden our work in the science of security.

KEITH B. ALEXANDER
General, US Army
Commander, US Cyber Command
Director, NSA/Chief, CSS

Introducing the federal cybersecurity R&D strategic plan

Douglas Maughan,
Bill Newhouse,
and Tomas Vagoun

In December 2011, the White House Office of Science and Technology Policy (OSTP) released the document, "Trustworthy cyberspace: Strategic plan for the federal cybersecurity research and development program," [1] which provides a framework for a set of coordinated federal strategic priorities and objectives for cybersecurity research. The release of this strategic plan marked an important milestone by the federal government's research community. It expresses an understanding of key causes of cybersecurity deficiencies and presents research themes with high potential to significantly improve the security of cyber systems and infrastructure. The strategic plan is a culmination of many efforts within the federal government, most notably by the Cyber Security and Information Assurance (CSIA) Senior Steering Group for Cybersecurity Research and Development (R&D), the CSIA Interagency Working Group of the federal Networking and Information Technology Research and Development (NITRD) Program, and by the Special Cyber Operations Research and Engineering (SCORE) Interagency Working Group.

Leaping ahead on cybersecurity

Focused efforts to develop a federal cybersecurity R&D strategy gained momentum in 2008 with the Leap-Ahead Initiative, a component of the Comprehensive National Cybersecurity Initiative (CNCI) [2]. Pursuant to CNCI, OSTP tasked the NITRD Program with carrying out the R&D goals of this initiative—to coordinate and prioritize R&D efforts and to develop strategies for a portfolio of government R&D activities to pursue high-risk/high-payoff solutions to critical

NITRD Program coordinates federal R&D in computing and cybersecurity

Since 1991, the federal Networking and Information Technology Research and Development (NITRD) Program has been the forum for coordinating interagency research activities in networking, computing, software, cybersecurity, and related information technology areas. Cybersecurity research is coordinated among the agencies in the Cyber Security and Information Assurance (CSIA) Interagency Working Group.

The primary participants are representatives from the Defense Advanced Research Projects Agency (DARPA), the Department of Homeland Security (DHS) Directorate of Science and Technology, the Department of Energy (DOE), the Intelligence Advanced Research Projects Activity (IARPA), the National Institute of Standards and Technology (NIST), the National Security Agency (NSA), the National Science Foundation (NSF), the Office of the Secretary of Defense (OSD), and the DoD Service Research Organizations. Along with the CSIA Interagency Working Group, the Special Cyber Operations Research and Engineering (SCORE) Interagency Working Group coordinates research related to national security systems.

The NITRD CSIA R&D Senior Steering Group was established in 2008 in response to the Presidential Comprehensive National Cybersecurity Initiative to define, coordinate, and recommend strategic federal R&D objectives in cybersecurity and to provide a robust conduit for cybersecurity R&D information across the policy, fiscal, and research levels of the government. The CSIA Senior Steering Group is composed of senior representatives of agencies with national cybersecurity leadership positions, including the Office of the Director of National Intelligence, DoD, DHS, NSA, NSF, NIST, the White House Office of Science and Technology Policy, and the Office of Management and Budget.

cybersecurity problems. At the onset, the CSIA Senior Steering Group determined that a government-wide framework for cybersecurity research was needed to provide both the coordination mechanism and the strategic directions for R&D. It was also clear within the CSIA Senior Steering Group that in order to achieve high-payoff, transformational results in cybersecurity, the framework needed to embody the following principles: the research must focus on root causes of cybersecurity vulnerabilities (not symptoms); the research activities must bring together expertise from a range of disciplines, given that cybersecurity is a challenge with technological, social, and economic aspects; and we must develop enduring cybersecurity concepts to assure trustworthiness of our systems despite changes in technologies and cyber threats.

With these principles in mind, the CSIA Senior Steering Group issued three public requests for input from October 2008 through April 2009, canvassing industry and academia for game-changing ideas that could fundamentally change the cyber environment into one where the rightful users and owners have an advantage over attackers and illicit efforts. Two hundred and thirty-eight responses were received by the CSIA Senior Steering Group. (To view and download copies of the responses, see [3].) The Senior Steering Group's review of the responses gave rise to five prospective game-changing categories: hardware-enabled trust, cyber economics, moving target defense, digital provenance, and nature-inspired cyber health. In August 2009, the NITRD Program and OSTP held the National Cyber Leap Year Summit where some 150 researchers from industry, academia, and government met for four days to examine the five game-changing categories. The Summit provided a forum to review the prospective categories, elevate key ideas, and capture the output in the Co-Chairs' Report [4] and the Participants' Ideas Report [5].

Following the National Cyber Leap Year Summit, the CSIA Senior Steering Group synthesized the five game-changing category reports and established three initial cybersecurity R&D themes: tailored trustworthy spaces, moving target, and cyber economic incentives. These themes were announced [6] at a public event collocated with the 2010 Institute for Electrical and Electronic Engineers Symposium on Security & Privacy. Two months later, the White House released the Office of Management and Budget/Office of Science and Technology Policy's memo to the agency heads on science and technology priorities for the 2012 fiscal year budget [7], highlighting the three cybersecurity R&D themes and directing agencies to utilize the themes in prioritizing cybersecurity R&D budgets and programs. The release of the White House memo accelerated the creation of new programs to focus on the three cybersecurity R&D themes.

Cybersecurity R&D thrusts

With the successful release of the framework for cybersecurity game-changing R&D, the CSIA Senior Steering Group and the CSIA Interagency Working Group began developing the federal cybersecurity R&D strategic plan. Together with accelerating research in areas with game-changing potential, four areas (or thrusts) were defined by the strategic plan:

- **Inducing change—**utilizing game-changing themes to direct efforts toward understanding the underlying root causes of known threats with the goal of disrupting the status quo; the research themes in the strategic plan include tailored trustworthy spaces, moving target, cyber economic incentives, and designed-in security;
- **Developing scientific foundations—**developing an organized, cohesive scientific foundation to the body of knowledge that informs the field of cybersecurity through adoption of a systematic, rigorous, and disciplined scientific approach;
- **Maximizing research impact—**catalyzing integration across the game-changing R&D themes, cooperation between governmental and private-sector communities, collaboration across international borders, and strengthened linkages to other national priorities, such as health IT and Smart Grid; and
- **Accelerating transition to practice—**focusing efforts to ensure adoption and implementation of the powerful new technologies and strategies that emerge from the research themes and of the activities to build a scientific foundation so as to create measurable improvements in the cybersecurity landscape.

The strategic plan deliberately does not focus on specific technical challenges, such as more secure operating systems. Instead, the plan defines desired end states and future capabilities, which, if achieved, would overcome critical underlying causes of cybersecurity vulnerabilities. By defining the end states, the themes invite a diversity of approaches and encourage innovation across disciplines and sectors. The essence of the strategic plan is to express a vision for the research necessary to develop game-changing technologies that can neutralize the attacks on the cyber systems of today and lay the foundation for a scientific approach that better prepares the field to meet the challenges of securing the cyber systems of tomorrow. Altogether, the plan provides guidance for federal agencies, researchers, and the public on how to prioritize research activities to achieve the greatest impact.

Efforts to develop scientific foundations in cybersecurity

In conjunction with the process to formally release the strategic plan, the federal agencies with R&D activities in cybersecurity began to introduce programs to pursue the goals outlined within each of these thrusts. In support of the thrust embodying the development of scientific foundations are representative R&D activities such as:

- **The Air Force Office of Scientific Research (AFOSR) 2011 Science of Security (SoS) Multidisciplinary Research Program of the University Research Initiative (MURI).** The objective of the AFOSR 2011 SoS MURI is to begin the development of an architecture or first principle foundation to define cybersecurity. The intent is to discover and define basic system properties that compose system security and other useful attributes in a manner that allows system properties to be verified and validated through theoretical proof and/or experiment.
- **NSA SoS lablets.** NSA support to academic lablets is focused on the development of a science of cybersecurity and a broad, self-sustaining community effort to advance it. A major goal is the creation of a unified body of knowledge that can serve as the basis of a trust engineering discipline, curriculum, and rigorous design methodologies. The results of SoS lablet research are to be extensively documented and widely distributed through the use of a new, network-based collaboration environment. The intention is for that environment to be the primary resource for learning about ongoing work in security science and to be a place to participate with others in advancing the state of the art.
- **The Army Research Laboratory (ARL) science for cyber portfolio.** The goal of ARL's science for cyber research portfolio is to examine a number of issues underlying cybersecurity and to develop novel theoretical constructs on which future cybersecurity advances can be based. The program explores models for the representation of cybersecurity, develops ensemble techniques for

improved detection of attacks, and investigates behavior as a fundamental indicator in detection and analysis. In particular, the research program focuses on theories and models that will lead to more effective intrusion detection techniques.

- **The National Science Foundation (NSF) Team for Research in Ubiquitous Secure Technology (TRUST)/Secure and Trustworthy Cyberspace (SaTC) Program.** TRUST, established as an NSF Science and Technology Center, focuses on addressing technical, operational, privacy, and policy challenges via interdisciplinary projects that combine fundamental science and applied research to deliver breakthrough advances in trustworthy systems in “grand challenge” areas such as the science of cybersecurity. In this area, TRUST researchers are developing a science base for security, with hopes to ultimately leverage these views in revising course content and embodying this theory in tools for system developers. Similarly, NSF’s SaTC program is focused on making cyberspace secure and trustworthy. Research in cybersecurity must “change the game,” check the misuses of cyber technology, bolster education and training in cybersecurity, establish a science of cybersecurity, and transition promising cybersecurity research into practice. The program recognizes that cyberspace will continue to grow and evolve and that advances in the sciences and technologies must grow and evolve as well, creating new “leap-ahead” opportunities.

The research in support of the strategic plan thrusts represents an increasing portion of the CSIA R&D budgets across federal agencies. This also translates into greater support of national priorities, such as health IT or Smart Grid, where key cybersecurity challenges can be addressed by focusing R&D activities within the framework of the thrusts.

Going forward, the execution of the strategic plan continues to be a collaborative process among a group of stakeholders: OSTP, responsible for policy and budgets; the CSIA Senior Steering Group, responsible for strategic directions; the CSIA Interagency Working Group, responsible for coordinating R&D activities; the SCORE Interagency Working Group, responsible for coordinating with R&D for national security systems; the federal agencies with cybersecurity R&D responsibilities; and the private sector. After a deliberate and thoughtful process, the nation’s cybersecurity research community can focus its energy and resources on a shared vision of a trustworthy cyberspace.

About the authors

Dr. Douglas Maughan is the cybersecurity division director in the Homeland Security Advanced Research Projects Agency within the Science and Technology (S&T) Directorate of the Department of Homeland Security (DHS). Dr. Maughan has been at DHS since October 2003 and is directing and managing the cybersecurity R&D activities and staff at DHS S&T. His research interests and related programs are in networking and information assurance.

Prior to his appointment at DHS, Dr. Maughan was a program manager at the Defense Advanced Research Projects Agency (DARPA) in Arlington, Virginia. Prior to his appointment at DARPA, Dr. Maughan worked for NSA as a senior computer scientist and led several research teams performing network security research. Dr. Maughan received bachelor’s degrees in computer science and applied statistics from Utah State University, a master’s degree in computer science from Johns Hopkins University, and a PhD in computer science from the University of Maryland, Baltimore County.

Bill Newhouse is a cybersecurity program lead in the Computer Security Division, one of six divisions in the Information Technology Laboratory at the National Institute of Standards and Technology (NIST). Newhouse represents NIST in several collaborative efforts including (1) the National Initiative for Cybersecurity Education, (2) a partnership with DHS and the financial sector to develop and test innovative cybersecurity technologies and processes, and (3) as a member of federal interagency cybersecurity R&D committees.

Before coming to NIST in 2010, Newhouse spent five years in the Office of the Secretary of Defense where he focused on the cybersecurity and information assurance R&D portfolio, first with the assistant secretary of defense for research and engineering and then with the assistant secretary of defense for Networks and Information Integration (NII). While in NII, he championed Defense Venture Catalyst Initiative workshops to focus on cybersecurity solutions from innovative companies. He is an electrical engineering graduate of both the Georgia Institute of Technology and George Washington University and has been with the federal government for over 25 years, beginning his career as a cooperative education student at NSA in 1986.

Dr. Tomas Vagoun provides subject matter expertise, technical leadership and management, and guidance to the Networking and Information Technology Research and Development (NITRD) Program in the areas of cybersecurity and information assurance and cybersecurity research and development. He supports the groups' monthly meetings and workshops, to include identifying opportunities for coordination and collaboration within the NITRD Program. Dr. Vagoun also assists in the writing and editing of technical reports and planning documents (e.g., National Science and Technology Council federal plans, workshop reports, and the annual supplement to the president's budget [8]).

Before joining the National Coordination Office for the NITRD Program, Dr. Vagoun led the implementation of software development projects for major federal civilian and defense agencies. Dr. Vagoun received an MS and BS in computer science from West Virginia University and a PhD in information systems from the University of Maryland.

References

[1] Chopra A, Schmidt H. "Federal cybersecurity R&D strategic plan released." *The White House Blog*. 2011 Dec 06. Available at: http://www.whitehouse.gov/blog/2011/12/06/federal-cybersecurity-rd-strategic-plan-released

[2] National Security Council. "The comprehensive national cybersecurity initiative." Available at: http://www.whitehouse.gov/cybersecurity/comprehensive-national-cybersecurity-initiative

[3] Responses to the CSIA Senior Steering Group's requests for input available at: http://www.nitrd.gov/fileupload/files/NCLY_submissions_public.pdf

[4] Chong F, Lee R, Acquisti A, Horne W, Palmer C, Ghosh A, Pendarakis D, Sanders W, Fleischman E, Teufel III H, Tsudik G, Dasgupta D, Hofmeyr S, Weinberger L. "National cyber leap year summit 2009: Co-chairs' report." 2009 Sep 16. Available at: http://www.nitrd.gov/fileupload/files/National_Cyber_Leap_Year_Summit_2009_CoChairs_Report.pdf

[5] "National cyber leap year summit 2009 participants' ideas report: Exploring paths to new cyber security paradigms." 2009 Sep 16. Available at: http://www.nitrd.gov/fileupload/files/National_Cyber_Leap_Year_Summit_2009_Participants_Ideas_Report.pdf

[6] Wing J, Landwehr C, Muoio P, Maughan D. "Toward a federal cybersecurity research agenda: Three game-changing themes." 2010 May 19. Available at: http://www.nitrd.gov/fileupload/files/NITRDCybersecurityR&DThemes20100519.ppt

[7] Orszag P, Holdren J. Memorandum for the heads of executive departments and agencies: Science and technology priorities for the FY 2012 budget. 2010 Jul 21. Available at: http://www.whitehouse.gov/sites/default/files/microsites/ostp/fy12-budget-guidance-memo.pdf

[8] National Coordination Office for Networking and Information Technology Research and Development. "The Networking and Information Technology Research and Development program: Supplement to the President's Budget FY 2013." 2012 Feb. Available at: http://www.nitrd.gov/PUBS/2013supplement/FY13NITRDSupplement.pdf

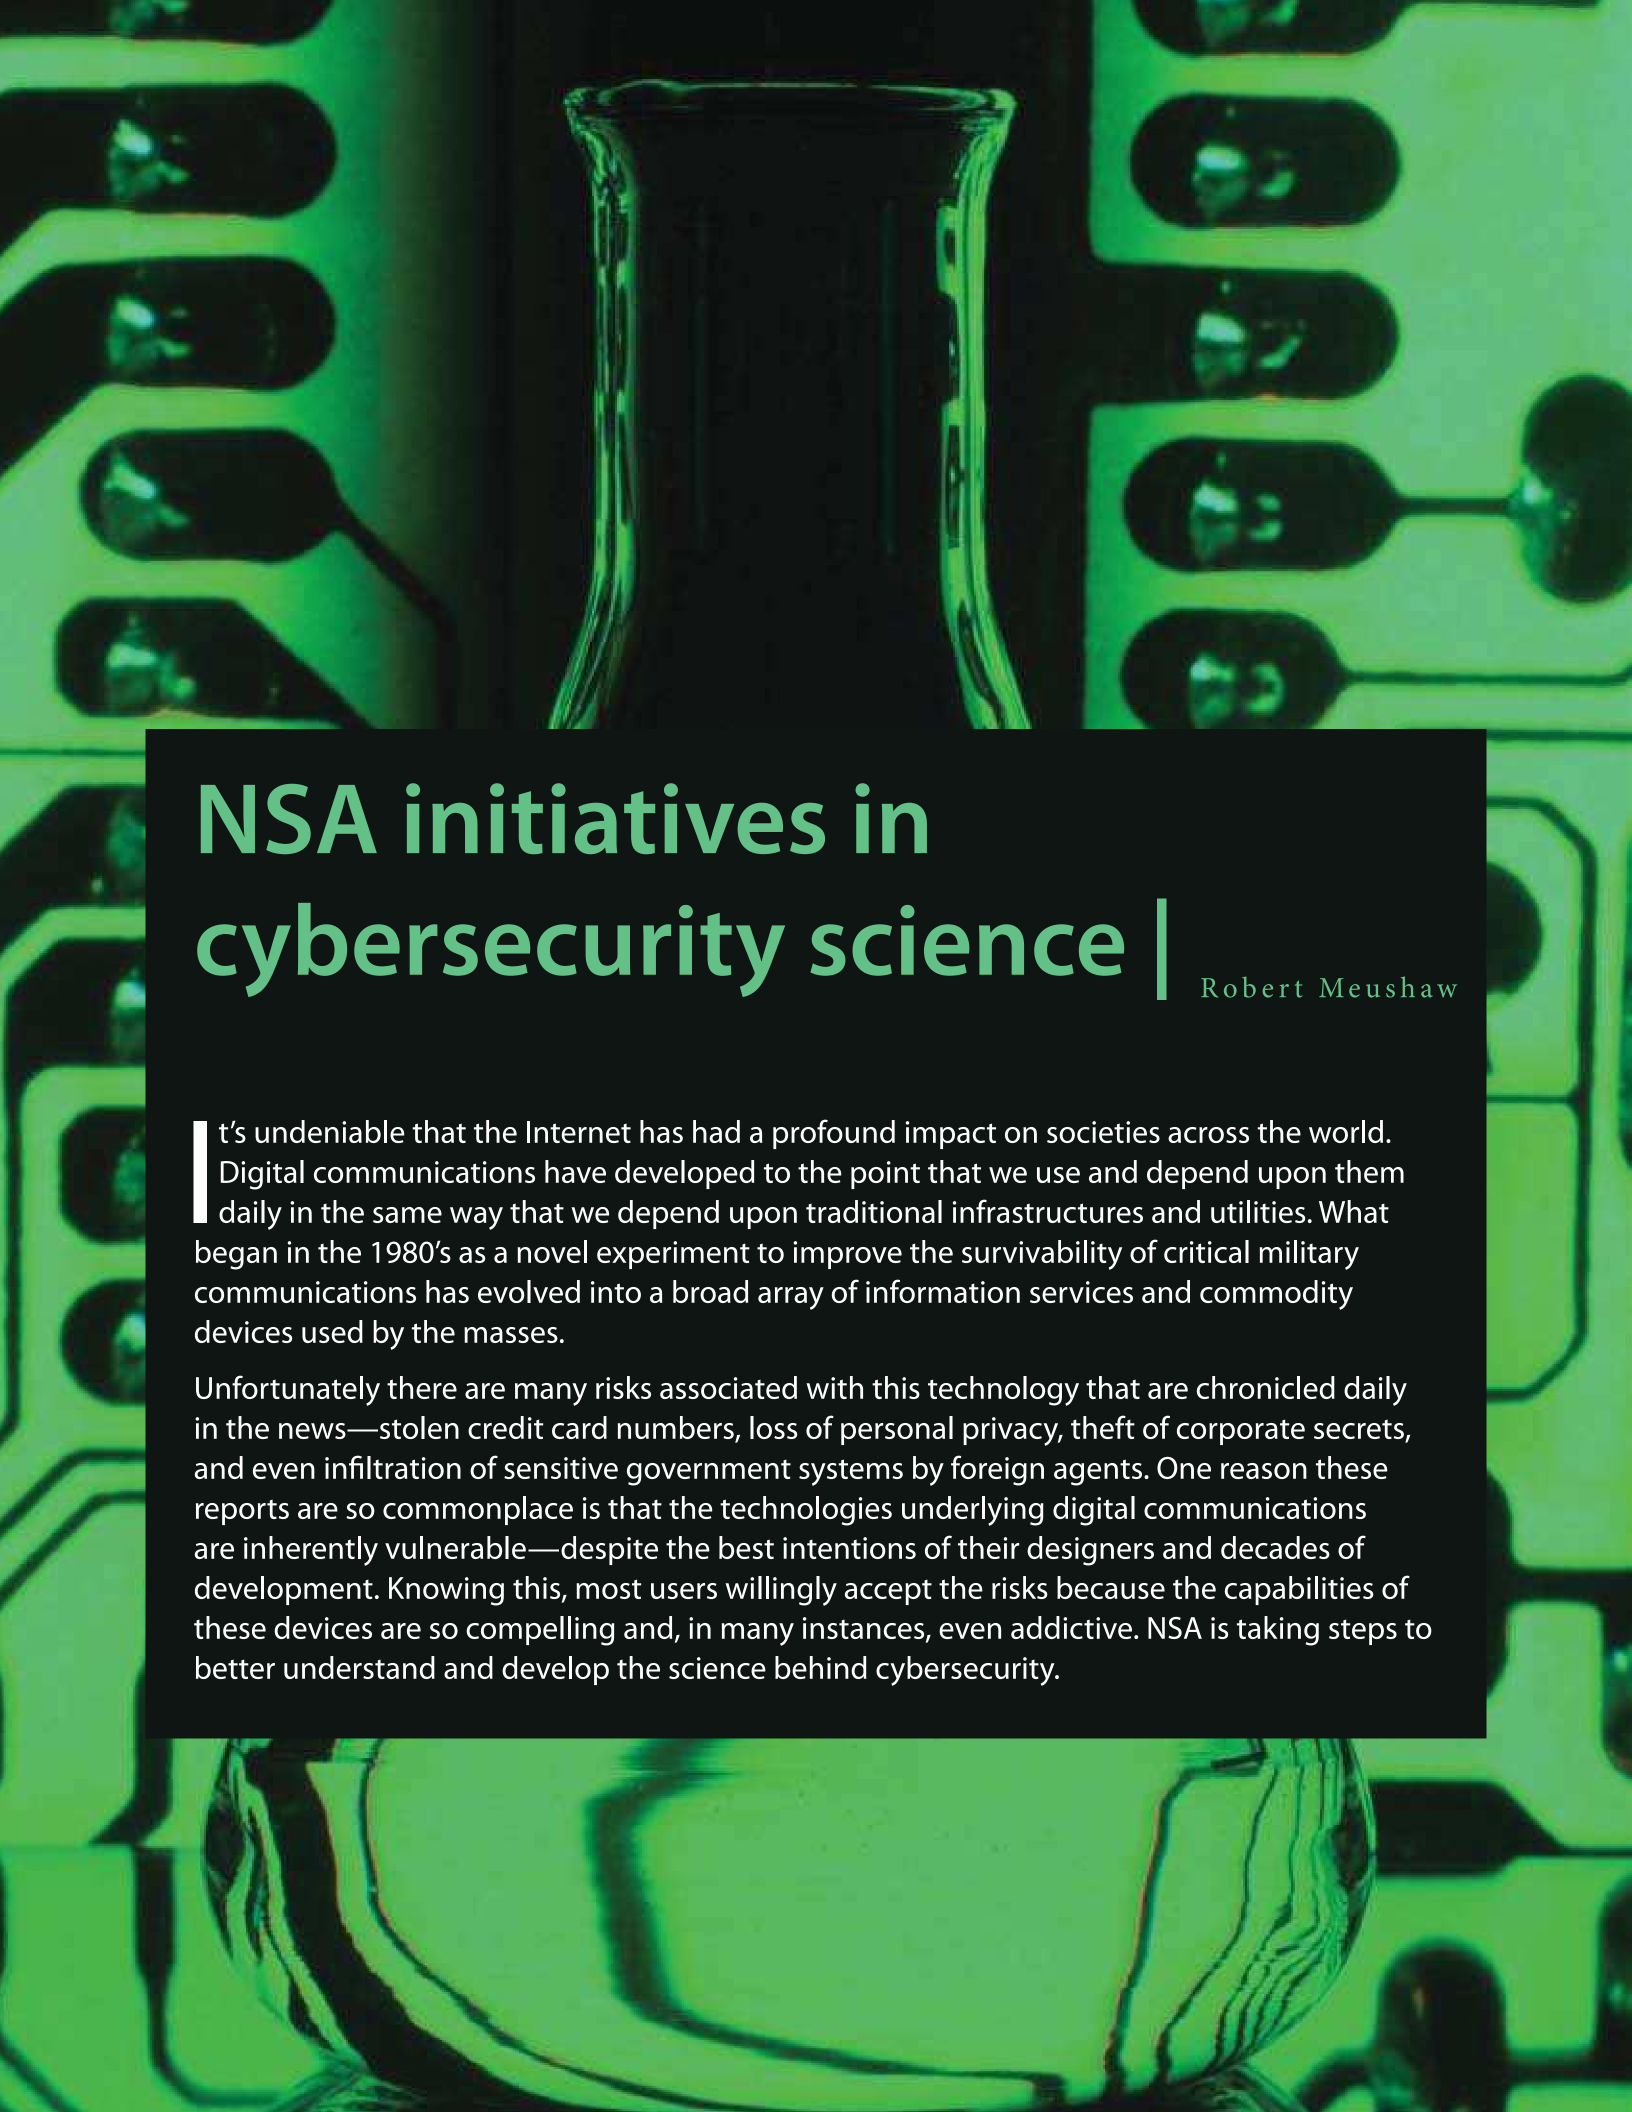

NSA initiatives in cybersecurity science |

Robert Meushaw

It's undeniable that the Internet has had a profound impact on societies across the world. Digital communications have developed to the point that we use and depend upon them daily in the same way that we depend upon traditional infrastructures and utilities. What began in the 1980's as a novel experiment to improve the survivability of critical military communications has evolved into a broad array of information services and commodity devices used by the masses.

Unfortunately there are many risks associated with this technology that are chronicled daily in the news—stolen credit card numbers, loss of personal privacy, theft of corporate secrets, and even infiltration of sensitive government systems by foreign agents. One reason these reports are so commonplace is that the technologies underlying digital communications are inherently vulnerable—despite the best intentions of their designers and decades of development. Knowing this, most users willingly accept the risks because the capabilities of these devices are so compelling and, in many instances, even addictive. NSA is taking steps to better understand and develop the science behind cybersecurity.

Realizing the need for cybersecurity science

NSA has played an active role in system security for over six decades—originally in the area of cryptography for classified communications and later in the development of a wide range of technologies to protect modern computing systems. To maintain its edge, NSA has a tradition of using expert panels for advice and guidance in critical technical areas. In 2008, the Information Security Panel initiated a discussion concerning the scientific underpinning for computer security engineering. Their concern stemmed from the growing use of commercial off-the-shelf technology in critical government systems, and they questioned whether the frequency of high profile security failures could be attributed to a lack of scientific rigor in security engineering. In contrast, they noted that the science and engineering associated with cryptographic systems, while still imperfect, seemed to result in far fewer catastrophic failures. The panel concluded that NSA's Information Assurance (IA) Research Group should review the state of cybersecurity science and consider establishing an initiative to put cybersecurity engineering on par with other established engineering disciplines.

The panel's concerns and challenge were welcomed as corporate-level acknowledgement of what security researchers at NSA and throughout the community had come to believe—that a new, strategic initiative was needed to advance security from the current patchwork of point solutions and ad hoc approaches and that resources should be shifted to focus on the development of a cohesive and organized body of knowledge as a foundation for the field of cybersecurity. The IA research group was convinced that the Agency's experience developing strong foundations for cryptography provided the model for what might be done in cybersecurity science and that the evolution of NSA's IA mission into the cyber domain provided more than enough motivation for it to take on a leadership role.

Assessing the state of cybersecurity science

Gauging the state of cybersecurity science, or any science, requires some method of determining what work truly qualifies as science. While there are myriad definitions of science that relate to testable hypotheses—for example, the ability to make predictions and the use of methodical procedures—a simplistic definition adopted by the IA research group was "any work that describes the *limits of what is possible.*" A good example of science consistent with this definition is Claude Shannon's seminal work on channel capacity, which established upper bounds on the rate of information transfer through a communications circuit. Shannon's results have provided the foundation upon which much of modern communications engineering is based.

Our simple litmus test provided us with a simple and straightforward way to distinguish scientific results in our review of security research. We began with a high-level review of research papers presented at prominent security conferences and then surveyed the security curricula of leading academic institutions. We concluded that most security work meeting our definition of science was concentrated in the areas of cryptography, cryptographic protocols, program correctness, fault tolerance, and formal methods. Much of the other research in security has been concerned with models of security (e.g., Bell-Lapadula and Biba), heuristic design principles, attack strategies, design/assessment of security components (e.g., firewalls, filters, and virtual private networks), risk assessment, intrusion analysis, etc. Although this body of research has contributed to the development of more trustworthy systems, it does not contribute to our understanding of the science of cybersecurity.

Overall, we concluded that the results of our review were consistent with the advisory panel's view of cybersecurity science. But an equally important conclusion we reached was that making significant strides in cybersecurity science would require an effort much larger than NSA alone could support. Unlike NSA's authority in the field of cryptography, no single government organization is charged with responsibility for cybersecurity technology and its scientific foundations. We felt that developing a body of science to support our nation's interests in cyberspace would require a large, long-term effort supported by the combined resources of government, industry, and academia. NSA's mission and experience in information assurance, and its six decades of investment in the science of cryptography, place it in a unique position to provide a leadership role for advancing the science of cybersecurity.

A holistic approach to cybersecurity science

To socialize the idea of a broad program focused specifically on science, we consulted with the other government organizations that have traditionally sponsored security research. Those discussions resulted in a decision to sponsor a workshop to explore the topic of cybersecurity science in depth with a broad group of representatives from government, academia, and industry. In November 2008, the Workshop on the Science of Security (i.e., science of cybersecurity) sponsored by the National Science Foundation (NSF), the Intelligence Advanced Research Projects Agency (IARPA), and NSA was held in Berkeley, California. Attendees included experts from traditional information security fields as well as others from a variety of nontraditional fields including biology, economics, and sociology. The range of topics discussed was equally broad and included such questions as:

- Is a science of cybersecurity possible?
- What might a science of cybersecurity look like?
- How can we reason about problems that seem impossibly hard?
- Is it possible to have scientific security metrics?
- What lessons can we learn from other disciplines?

Several days of discussions generated a broad and divergent set of ideas concerning the possibility of developing a science of cybersecurity. But there was general agreement on several areas where advances were sorely needed. The first concerned the need to account for human behavior in models of system security. While the difficulty of modeling intelligent adversarial behavior has long been recognized as a shortcoming in security models, it has also become increasingly apparent that a science of cybersecurity should account for human behavior associated with the overall operation and defense of cyber systems. In either case, however, the addition of a human dimension was acknowledged to add enormous complexity to the task of analyzing and designing secure systems.

There was also agreement that the ability to produce systems that are secure in the real world requires accounting for important factors beyond just the technical aspects of the security mechanisms used. The poor adoption rate and ineffective use of available security technology over the past several decades were viewed as evidence of this. Beyond the role of human behavior, the impact of financial and business constraints on the effectiveness of system security were highlighted.

While no specific plan of action emerged from the workshop, the collection of ideas generated significantly influenced the research programs of numerous funding groups, NSA's in particular. In a significant departure from past NSA research programs, our new cybersecurity science portfolio will seek to include a much more diverse set of disciplines than previously considered, including human perception, psychology, physiology, economics, data analytics, and game theory.

Strategies for advancing science

Recognizing the need to improve the scientific foundations of security was a useful first step, but it didn't provide insight regarding what strategy might best accomplish this goal. One seemingly obvious and straightforward approach was simply to increase funding for security research that specifically targeted science. It was clear that even sizable increases in current budgets—which weren't likely—would fall far short of producing the advances needed. But before proceeding with any specific strategy, it seemed prudent to investigate why more science hadn't already been produced. Some who have reviewed the broader ecosystem in which research is conducted believe that current incentives associated with security research weren't well suited to producing science. (See Tom Longstaff's article on page 14 for more on this subject.) This suggested that we should consider a strategy aimed at reshaping the incentive system. In the end, since it was not clear if either of these approaches would produce the desired results, we decided to adopt a mixed strategy—one that provides direct support for specific science research projects while, at the same time, seeking improvements in the overall conditions for producing science.

Experiments in funding science

For decades, government organizations including the Defense Advanced Research Projects Agency (DARPA), NSF, the Air Force Research Laboratory (AFRL), and the Army Research Office (ARO), as well as NSA have used direct funding for research targeted at specific security topics; so it seemed straightforward to apply the same approach for cybersecurity science. NSA's cybersecurity science initiative is exploring a number of variations of this strategy to assess their effectiveness. One approach, used shortly after the conclusion of the Berkeley workshop, provides supplemental funding to an ongoing security research program (i.e., NSF's Team for Research in Ubiquitous Secure Technology Science and Technology [TRUST] Center) specifically to encourage work in science. A second approach was adapted from industry: it involves funding specific work in science at a small number of academic research groups—referred to as lablets—at highly qualified institutions. The first three lablets, established at Carnegie Mellon University, University of Illinois, and North Carolina State University, were beneficiaries of funding provided to NSA that was specifically earmarked for cybersecurity science. (See page 46 for more information.) While the initial choice of lablets was limited by timing constraints placed on the funding, the number of institutions participating in the program increased through the inclusion of an outreach requirement for each lablet. The last funding approach included in our portfolio provides support to specific, high-impact problem areas identified through research reviews conducted across the security community. Composition is one cybersecurity science topic that is currently being supported with the goal of understanding how the security properties of a system can be derived from the properties of its component parts.

After several rounds of modest NSA funding supplements to NSF's TRUST Center, increased attention is being devoted to science and beginning to influence other work and researchers. NSA's lablet initiative, formally established in 2012, recently kicked off several dozen projects to explore how effective a multiuniversity, multidisciplinary team approach can be at advancing science and involving nontraditional partners. Early work has focused on identifying core hard problems in science that must be understood in order to deal with the security issues that plague the nation. We have long recognized that security research does not always lead to scientific understanding, and through collaboration with our lablet partners, we are maturing our joint understanding of how to shape research to maximize its contribution to science. Our work funding specific projects in science has just begun, but the quality of the investigators and their previous contributions to science make us confident that these efforts will provide a showcase for cybersecurity science research.

On applying strong inference to cybersecurity science

Carl E. Landwehr

In 1964, biophysicist John R. Platt observed that some scientific fields, such as molecular biology and high energy physics, seem to advance more quickly than others, and he argued that the use of a method he dubbed "strong inference" was responsible [1]. In strong inference, a tree of alternative hypotheses is developed and pruned in response to the results of critical experiments. Platt's paper created quite a stir at the time and has continued to inspire responses over the years. (See [2, 3] for two examples.)

Could this approach speed the development of a science of cybersecurity? To investigate this question, NSA sponsored a panel at the 2012 Institute of Electrical and Electronics Engineers Symposium on Security and Privacy. Five cybersecurity researchers active in economics, human behavior, systems, formal methods, and cryptography were asked to assess the suitability and actual use of strong inference in their respective fields. As organizer of the panel and moderator of the discussion, which included lively exchanges with the audience, my personal conclusions are that strong inference is not widely used in the field at present and that its potential benefit is strongest in those domains where natural phenomena, including human behavior, must be modeled. Its benefits are less clear in areas like cryptography and formal methods, where mathematics and logic predominate. Nevertheless, in any field, the intellectual rigor required to formulate a proposed research project as a hypothesis-testing exercise can only help.

Broadening research participation

A funding strategy that targets specific research projects unavoidably limits participation to a small group of researchers. To significantly broaden participation in cybersecurity science we are investigating ways to reshape the overall research environment to be more conducive to producing science. One goal is to increase the perceived value of research that advances science, even incrementally, rather than of work that tracks the latest security trends. If successful, we believe we can accelerate the creation of a cybersecurity science by leveraging a much larger community of researchers. The downside of such an indirect approach is that specific research outcomes are much less certain and the overall effectiveness of the investment is difficult to assess. While influencing the research environment seems simple notionally, developing a practical strategy to do this is challenging. Some of the approaches we are investigating include challenge problems, competitions, awards for scientific papers, and recognition of researchers' achievements. The strategy we adopt, as in other cases, will include a variety of these techniques.

Building community

Our report to NSA's advisory board observed that the scope of the effort needed to develop a science of cybersecurity was well beyond what NSA could accomplish on its own. But we also noted that NSA was in a unique position to lead a community activity to make this happen. One of the key aspects of our science initiative has been enlisting the support of NSA's many research partners including the Air Force Office of Scientific Research, the Department of Homeland Security, NSF, DARPA, IARPA, the federal laboratories, and other groups across the DoD and intelligence community. We have also sought the involvement of our foreign partners, particularly the UK and Canada. Although a government-wide cybersecurity science initiative does not yet exist, we have attempted to coordinate the collection of research projects to provide cohesion and balance.

In the past several years there has been a groundswell of interest in creating more robust scientific foundations for cybersecurity. Today, there are numerous cybersecurity science activities underway, with more being planned, and keeping track of them is becoming increasingly difficult. To deal with this problem and to encourage the development of a community surrounding work on cybersecurity science, NSA has taken a lead role in developing a web-based Science of Security Virtual Organization (SoS VO). This work leverages the Virtual Organization collaboration infrastructure developed by NSF to support its Cyber-Physical Systems (CPS) program. (Visit the CPS Virtual Organization at cps-vo.org.) The goal for the SoS VO is to provide "one stop shopping" for anything related to cybersecurity science. The website will provide information on conference events, research sponsors, current research programs, notices of future initiatives, research tools and data, etc. The research produced by these activities will be made available for review and distribution, and a future goal is to provide video streams of research reviews for wide viewing. The site is also intended to encourage and support collaboration by providing a variety of social networking features including discussion forums, chat, researcher blogs, and lists of challenge problems. (See article on page 20 for more information about the SoS VO.)

Transitioning findings to practice

New security systems continue to be developed despite limitations in existing science, so developers must make do with whatever practices are available, however imperfect. Because of this, an important consideration in our initiative is the rapid transition of emerging scientific results into the practice of security engineering. In our cybersecurity science lablet program, for example, we are seeking opportunities to develop courses that capture new science and to augment existing courses with improved scientific foundations. As new material is developed, we intend to leverage relationships with the National Institute of Standards and Technology and NSA's own Centers of Academic Excellence program in order to influence the design of new systems and future generations of developers. (For more information about NSA's Centers of Academic Excellence, visit http://www.nsa.gov/ia/academic_outreach/nat_cae.)

Measuring progress

Although the resources currently invested in cybersecurity science are relatively modest compared with other research areas, responsible program managers will still need to track the return on their investment. So, how can progress in cybersecurity science be measured? While breakthrough discoveries and near-term impact are always hoped for, scientific advances are often incremental and produced over periods measured in decades. Therefore, expectations for significant results need to be circumspect and mindful of the many ways in which scientific advance is observed. Types of scientific progress include:

- **Finding the new—**discovering scientific breakthroughs;

- **Taking a fresh look—**developing useful new ways to look at a given set of data;
- **Finding patterns—**discovering and explaining patterns in phenomena across time;
- **Finding connections—**linking theories and explanations across multiple fields of research; and
- **Influencing others—**stimulating further research, including research outside the field, and collaboration across different fields.

In addition, scientific progress may be seen in measures that show rising interest and excitement about a new field, including [4]:

- Established scientists begin to work in a new field;
- Highly promising junior scientists choose to pursue new concepts, methods, or lines of inquiry;
- Students increasingly enroll in courses and programs in a new field;
- The rate of publications in the field increases;
- Citations to publications in the field increase in both number and range across other scientific fields;
- Publications in the new field appear in prominent journals;
- New journals or societies appear; and
- Ideas from the field are adopted in other fields.

Conclusion

NSA's long-standing investment in cryptographic science and engineering has yielded the most robust encryption technology in the world. But the protection of our nation's cyber systems demands security design and analysis techniques that encompass much more than cryptography, yet are comparably grounded in science. While we do not expect that a science of cybersecurity can guarantee complete protection against cybersecurity threats any more than safety science can guarantee risk-free transportation, it should provide us with greater certainty about the capabilities and limitations of our security mechanisms, allowing us to make well-informed risk decisions. NSA's cybersecurity science initiative is the first step in a long-term endeavor to develop the broad understanding of security that we need to protect our national interests in cyberspace.

About the authors

Robert Meushaw is the former technical director of NSA's Information Assurance (IA) Research Laboratory. His current work focuses on developing new strategies and programs for the advancement of a science of cybersecurity. He retired from NSA in 2005 after 33 years of service, including over a decade of work in IA research. Meushaw's career at NSA also included significant stints in both the Production Development Group and the Security Evaluation Group of the IA Directorate. In addition to his technical responsibilities, he served for six years as a technical editor of NSA's *Tech Trend Notes* and *The Next Wave* publications. Meushaw holds degrees in electrical engineering from Princeton University and the Johns Hopkins University.

Carl E. Landwehr is an independent consultant in cybersecurity research. Until recently, he was a senior research scientist for the Institute for Systems Research at the University of Maryland, College Park. He received his BS in engineering and applied science from Yale University and his PhD in computer and communication sciences from the University of Michigan. Following a 23-year research career at the Naval Research Laboratory, he has for the past decade developed and managed research programs at the National Science Foundation and the Advanced Research Development Activity/Defense Technology Office/Intelligence Advanced Research Projects Activity. He is interested in all aspects of trustworthy computing. In December 2010, he completed a four-year term as editor in chief of *IEEE Security & Privacy Magazine.*

References

[1] Platt JR. "Strong Inference: Certain systematic methods of scientific thinking may produce much more rapid progress than others." *Science.* 1964;146(3642):347–353. DOI: 10.1126/science.146.3642.347

[2] O'Donohue W, Buchanan JA. "The weakness of strong inference." *Behavior and Philosophy.* 2001;29:1–20. Available at: http://mres.gmu.edu/pmwiki/uploads/Main/ODonohue2001.pdf

[3] Davis RH. "Strong inference: Rationale or inspiration?" *Perspectives in Biology and Medicine.* 2006;49(2):238–49. DOI: 10.1353/pbm.2006.0022

[4] Feller I, Stern PC, editors. *A Strategy for Assessing Science: Behavioral and Social Research on Aging.* Washington (DC): National Academies Press; 2007. Chapter 4, "Progress in science"; p. 67–94.

Barriers to achieving a science of cybersecurity

Tom Longstaff

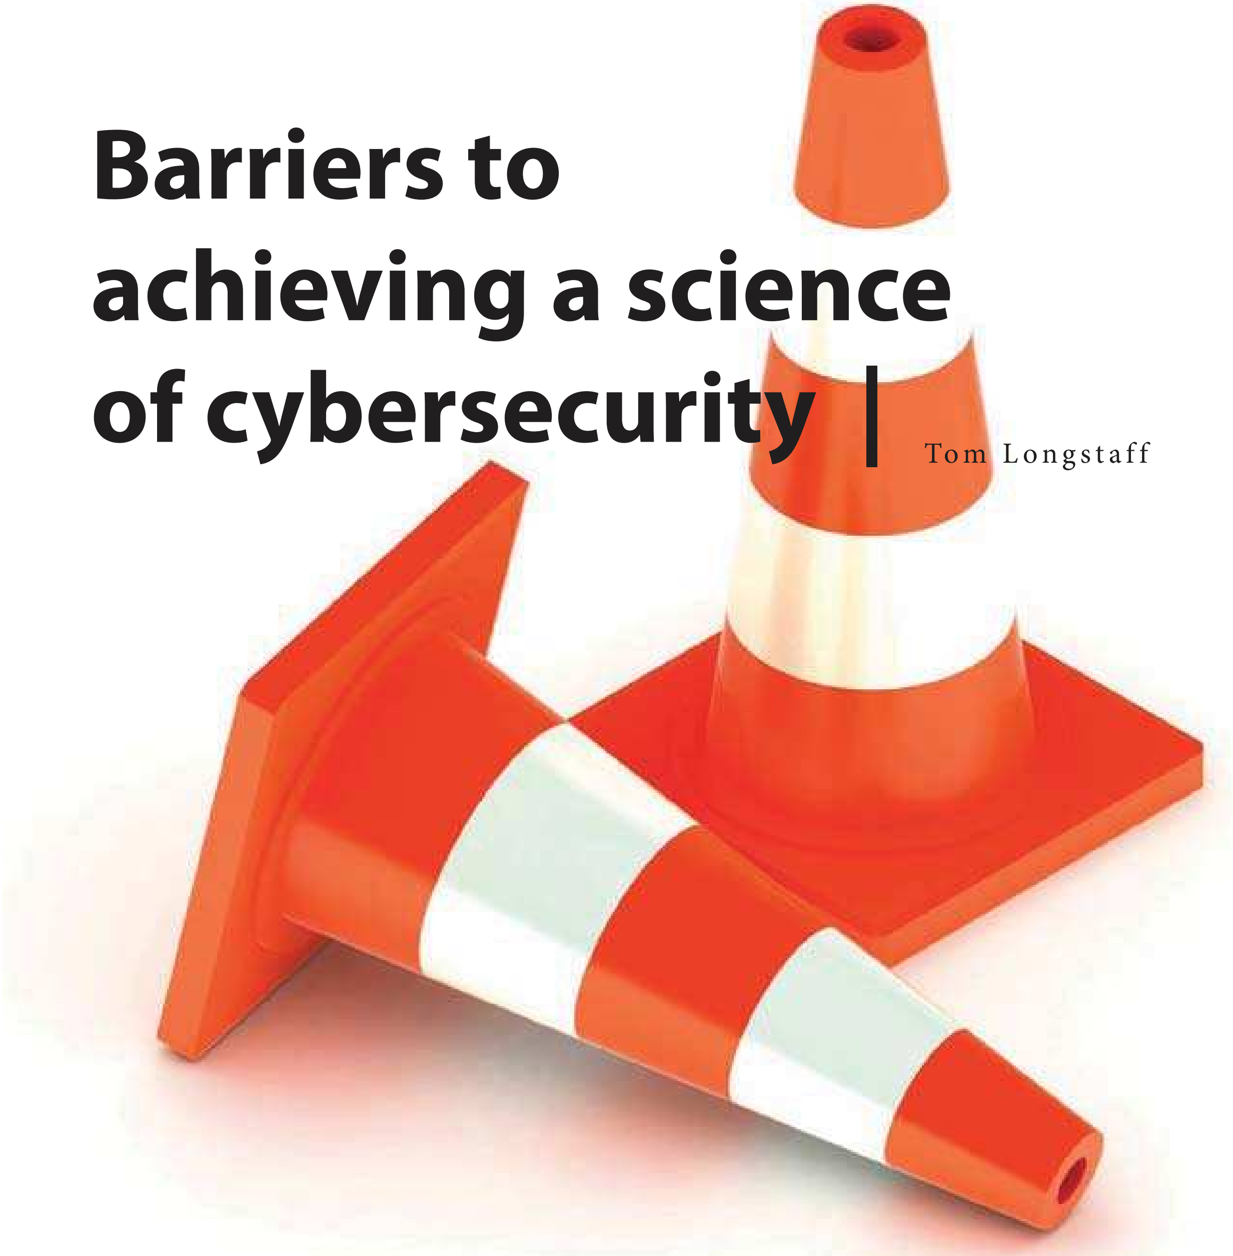

Several recent reports, such as the JASON "Science of cyber-security" report [1], point to examples and approaches for achieving success in applying science to cybersecurity. Audiences everywhere enthusiastically agree and thrash themselves for bypassing science all along, bemoaning the fact that we could be "so much further along" if we only did science. Of course, after the presentation is over, everyone goes back to the methods that have been used throughout our generation to create prototypes and tools with no regard for the scientific principles involved. Why?

During the winter of 2009, an informal group of three cybersecurity researchers—Roy Maxion from Carnegie Mellon University, Tom Longstaff from Johns Hopkins Applied Physics Laboratory, and John McHugh from the University of North Carolina—pondered this question based on their collective experience. The results of their discussion generated a presentation at the 2010 Annual Computer Security Applications Conference and a National Science Foundation (NSF) Washington Area Trustworthy Computing Hour (WATCH) lecture on March 15, 2012. (A transcript of the lecture can be found here, http://www.nsf.gov/events/event_summ.jsp?cntn_id=123376&org=NSF.)

At the NSF WATCH lecture, Tom Longstaff discussed some barriers to achieving a science of cybersecurity within the cybersecurity culture—barriers that seem to prevent well-meaning researchers from taking a more scientific approach to cybersecurity projects. Three of these barriers are described below.

1 Research begins after a conference is announced.

The informal group recognized that the publication cycle for cybersecurity papers is very short in comparison to other scientific fields, such as physics, chemistry, or psychology. The group noted that in other fields research is completed far in advance of a call for papers. In cybersecurity, however, common practice is to begin the research after a particular conference or venue is identified, often within six months of the submission deadline.

2 Program committees lack scientists.

The members of the informal group had been on many program committees before. They recognized that such committees were often made up of nonscientists who did not recognize or value the material in a scientific cybersecurity paper. Thus, papers accepted by these committees often did not include a methodology section, nor were authors encouraged to provide enough information to make their results repeatable or reproducible.

3 Publications favor articles about novelties in the field.

Finally, cybersecurity publications typically prefer articles or papers that indicate entirely new directions in cybersecurity, rather than incremental approaches that better describe the causal relationships found in cybersecurity. Being aware of this preference, authors do not spend time executing careful scientific experiments that lead to incremental approaches, but instead speculate or quickly produce a novel prototype.

While there are many incentives that could be added to address these three barriers, several were called out specifically in the WATCH lecture as likely to have a good long-term impact on the field of cybersecurity. They are to:

- Encourage the publication of longer-duration research in cybersecurity through preferential acceptance of such research in conferences and journals,
- Leverage the knowledge of traditional physical scientists in structuring scientific publications by encouraging coauthorship and collaboration with cybersecurity researchers,
- Train computer science students to use the scientific method through the development of new courses in experimental research and publication,
- Sponsor conferences and journals that promote the scientific method as a main acceptance criterion,
- Require authors of papers to use scientific rigor in their construction for sponsored conferences and journals,
- Create a publicly available body of knowledge consisting of a scientific publication in cybersecurity, and
- Create an explicit separation between scientific contributions and technological contributions (and reward scientific contributions).

Cybersecurity culture is rooted in performing rapid prototyping and programming ad hoc solutions to engineering problems. Changing this culture and overcoming the barriers described above will be difficult, but the benefits of encouraging science in cybersecurity will be well worth the effort.

About the author

Tom Longstaff is the technical director for the System Behavior Office within the NSA Research Directorate. He has spent the last 25 years leading research in Internet security, incident detection and response, and cyber resilience.

Reference

[1] JASON Program Office. "Science of cyber-security." 2010. McLean (VA): The Mitre Corporation. Report No.: JSR-10-102. Available at: http://www.fas.org/irp/agency/dod/jason/cyber.pdf

Funding research for a science of cybersecurity: The Air Force makes it a mission |

Dr. Robert Herklotz

The Air Force Office of Scientific Research (AFOSR) plans, coordinates, and executes the Air Force Research Laboratory's basic research program. AFOSR's technical experts identify and fund long-range technology options at Air Force, university, and industry research laboratories. This support ensures the timely transition of research results that lead to revolutionary scientific breakthroughs, enabling the Air Force and US industry to produce world-class, militarily significant, and commercially valuable products. Such research is inherently risky, sometimes outside of the mainstream, and often requires an extended period of support. This article describes several AFOSR initiatives that focus on the science of [cyber]security (SoS). The initiatives include a Multidisciplinary University Research Initiative (MURI), a Young Investigator Program (YIP) grant, and a Basic Research Initiative (BRI).

Multidisciplinary University Research Initiative (MURI)

In 2010, the deputy director for cybersecurity in the Information Systems and Cyber Security Directorate of the Assistant Secretary of Defense for Research and Engineering (ASD(R&E)) requested the AFOSR to fund a MURI focused specifically on the science of [cyber]security (SoS). The MURI program is DoD-wide and complements other DoD programs that support university research through the single-investigator awards. The MURI supports the research of teams of investigators whose backgrounds intersect multiple traditional science and engineering disciplines in order to accelerate research progress. The government team for this effort was led by Dr. Robert Herklotz, AFOSR, and included support from a number of research funding organizations including the Air Force Research Laboratory/Information Directorate; the Army Research Office; the Office of Naval Research; the National Science Foundation; the National Security Agency; the National Institute of Standards and Technology; and the Office of the Director, Defense Research and Engineering (now the ASD(R&E)).

The SoS MURI was prompted by the widely held belief in the security community that cybersecurity has been pursued largely as a reactive effort, with an endless cycle of new attacks and defensive responses. Many security experts have come to believe this cycle cannot be broken because today's information technology systems are too complex to ever be modeled with formally defined and verified security properties.

In fact, no formal definition of cybersecurity described in terms of system properties has yet been produced, let alone metrics capable of measuring those properties.

The objectives of the SoS MURI, as presented in the proposal solicitation, are to begin the development of an architecture or first principle foundation to define cybersecurity for such systems, to discover and define basic system properties that comprise system security and other useful attributes, and to identify system properties that can be verified and validated through theoretical proof and/or experimentation. A primary goal is to answer the following questions through the discovery and analysis of basic system properties:

- Can the system enforce the desired security policies in each system component?
- Can the system enforce the desired security policies across all system components simultaneously? If so, what are the security properties of the whole system?
- Can system capability, as defined in the first two bullets above, defend against each class of attack, once classes of cybersecurity attacks are defined?
- How can we formally define cybersecurity policies and mechanisms (including defense, monitoring, response, etc.) and assess their effectiveness against classes of attacks?
- Can an adversarial process model be formally defined that is capable of generating known classes of attacks?
- Can we define metrics for basic system properties and for the ability of a system to enforce a security policy that defends against a class of attacks?
- Can we define system properties and metrics dealing with system characteristics, such as scalability, adaptability, ease of use, etc., in order to compare alternative system designs?

The development of theoretical underpinnings (i.e., system properties and relationship to policies) and the theories and metrics (i.e., relationships between attacks, defenses, and policies) will allow us to create system engineering methodologies that can perform rigorous design trade-offs among cybersecurity properties, as well as other properties, in the development of complex systems. In addition, this research will:

- Enable the creation of new technologies and supporting tools grounded on sound principles,
- Establish a baseline for comparing technology capabilities among vendors,
- Encourage the creation of a new industry for security software engineering technologies, and
- Reduce development costs by providing scientifically supported evidence of security properties rather than applying exhaustive testing to look for evidence of insecurity.

The winning MURI proposal

The winning proposal, announced April 22, 2011, is entitled "Science of cybersecurity: Modeling, composition, and measurement." The work is to be performed by a multiuniversity team of researchers led by Professor John C. Mitchell of Stanford University.

Professor Mitchell's team proposed research to advance a science base for trustworthiness by developing concepts, relationships, and laws with predictive value. Their work will focus on problem areas amenable to rigorous treatment and generalizable solutions and is organized around the following three thrust areas:

1. **Security modeling.** A uniform approach to security modeling will allow systematic approaches to be developed and applied to a broad range of richly connected systems, supporting analysis of resilience against graduated classes of clearly defined threat models.
2. **Secure composition.** Principles of secure composition will be developed, analyzed, and evaluated for systematic and modular construction of trustworthy systems, relative to security properties that can be verified and validated through theoretical proof and/or experimentation.
3. **Security measurement.** New security measurement concepts will be devised and used to

determine relative strengths of defense mechanisms, whether security improves from one version of a system to another, and when additional security mechanisms are warranted, given incentives associated with system attackers and defenders.

Together, the advances anticipated for these three complementary thrusts will support a science base for future systems that proactively resist attacks through secure design, development, and implementation based on principled foundations.

Young Investigator Research Program

On January 11, 2012, the AFOSR announced it would award approximately $18 million in grants to 48 scientists and engineers who submitted research proposals through the Air Force's Young Investigator Research Program (YIP).

The YIP is open to scientists and engineers at research institutions across the US who received a PhD or an equivalent degree in the last five years and show exceptional ability and promise for conducting basic research. The objective of this program is to foster creative basic research in science and engineering, enhance early career development of outstanding young investigators, and increase opportunities for the young investigators.

Among the 2012 winners was Michael Clarkson, assistant professor in the Department of Computer Science at the George Washington University. His YIP proposal, "Making cybersecurity quantifiable," is focused on further development of his PhD thesis on hyperproperties, a very promising tool for security science.

Basic Research Initiative on cyber trust and suspicion

On March 27, 2012, the AFOSR announced a Basic Research Initiative (BRI) to build the foundational understanding of human trust and suspicion in the cyberspace domain. Cyberspace operations rely heavily on the degree to which users trust, or are suspicious of, their information technology systems. To date, there has been little or no work in providing any unified/comprehensive treatment of the impacts of social, cultural, economic, political, and emotional factors (to name a few) underlying trust and suspicion, especially in complex systems.

The winning proposal, "A social, cultural, and emotional basis for trust and suspicion," led by Dr. Eunice E. Santos of the Institute of Defense and Security at the University of Texas, El Paso (UTEP), was funded on September 14, 2012. Her team, which includes UTEP, Syracuse University, the University of Tulsa, the University of Houston, and Assured Information Security, Inc., proposed research to develop a model of system users and managers and insider behavior that accounts for and explains the social, cultural, and emotional basis for trust and suspicion.

Among the questions their research will address are:

1. How can different people be swayed (or sway others) based on trust or suspicion?
2. How and why do group member sociocultural characteristics, group size, information sharing patterns, and events affect group cohesion?
3. Is it possible to detect significant drops in situational awareness or when the level of trust is inappropriate in a given context?
4. What are the critical interrelationships between information, emotional responses, situational awareness, influences on decision making, and associated changes in task performance?
5. How do complex multiscale and multilevel factors affect insider threat detection?
6. Lastly, and most importantly, can this research be unified into a single overarching framework of social, cultural, and emotional factors underlying trust and suspicion?

The end product of their project is a methodology that can be used to better understand system users and managers and the insider threat by providing the social, cultural, and emotional basis of human behavior in the cyber domain and the impacts of trust and suspicion on cyberspace operations.

A legacy of research

The AFOSR was born out of the need to address a long-standing shortfall in military basic research. This deficiency became obvious during World War II, when massive civilian-led research and development

efforts were required to create the technology needed for our nation to dominate warfare in a physical battle space. Today the AFSOR continues its original mission by investing in the development of basic research to support domination of the emerging battle space in the cyber domain. Just as a well understood scientific foundation is necessary for secure and safe physical systems, a science of cybersecurity is needed for safety and security in the cyber world. To learn more about the AFOSR basic research program funding opportunities, download the broad agency announcement (i.e., BAA-AFOSR-2012-0001) from https://www.fbo.gov/spg/USAF/AFMC/AFOSR/BAA-AFOSR-2012-0001/listing.html.

About the author

Dr. Robert L. Herklotz is currently the program manager for the Information Operations and Security basic research program at the Air Force Office of Scientific Research in Arlington, Virginia. He invests in science to develop secure information systems for our warfighters and to deny the enemy such systems. His specific subareas of research include the science of cybersecurity, secure humans, secure networks, secure hardware, covert channels, secure execution on insecure systems, secure data, and secure systems-security policy. From 2000 to 2006, he managed the Air Force's basic research investment in three programs: software and systems, artificial intelligence, and external aerodynamics and hypersonics. Prior to that, he was a career Air Force officer, retiring in 1999.

Dr. Herklotz holds a PhD from the Massachusetts Institute of Technology, an MS from Purdue University, and a BS from the US Air Force Academy. His awards include two Silver Stars, four Distinguished Flying Crosses, eight Air Medals, and the Association for Computing Machinery 2012 Special Interest Group on Security, Audit and Control Outstanding Contributions Award.

Advancing the science of cybersecurity with a virtual organization

Frankie King, Heather Lucas, and Robert Meushaw

Origins

The National Science Foundation (NSF)'s Cyber-Physical Systems (CPS) program is a research initiative to support the development of systems that combine physical, computing, and communications components at very large scale and high complexity. Cyber-physical systems are not the traditional desktop computers, embedded/real-time systems, and sensor nets with which we are familiar today. They are characterized by cyber capabilities in all physical components, networking at multiple and extreme scales, high degrees of automation, dynamic reconfiguration and reorganization, and extreme requirements for dependability and reliability. Although cyber-physical systems are currently being planned and developed to support applications in numerous areas (e.g., the smart power grid, smart healthcare, and smart transportation), the scientific understanding and engineering tools needed to realize such systems with high-confidence reliability and dependability are lacking.

The CPS Virtual Organization (CPS VO), an offshoot of the CPS program, was envisioned as a tool to promote and support a broad spectrum of collaborative interactions among researchers to assist in solving complex, crosscutting problems requiring expertise from multiple domains. The CPS VO provides a web-based gathering place and clearinghouse for knowledge relevant to cyber-physical systems and to advance the theory, engineering, and operation of cyber-physical systems. A primary objective of the CPS VO is to overcome some of the major impediments to progress in complex systems science, such as the lack of integration and cross-fertilization of numerous traditionally isolated disciplines. The NSF intended the CPS VO to enable electronic community building and to provide a vehicle for sharing information among otherwise disparate researchers, students, educators, and industry practitioners within the growing cross-disciplinary field of cyber-physical systems.

Vanderbilt University was selected by NSF to develop and manage the CPS VO. It was built using DRUPAL, a widely used, free, and open-source content management system that provides the back end for at least two percent of all websites worldwide, including whitehouse.gov. The system is flexible and highly customizable, providing a rich set of capabilities for the CPS VO user community. The CPS VO was initially used to advertise the activities of the CPS program and to establish electronic forums for many of the common interest groups (e.g., medical, automotive, aviation, education, and architectures) within the national High Confidence Software and Systems Coordinating Group. The High Confidence Software and Systems Coordinating Group (HCSS CG) is part of the national Networking and Information Technology Research and Development (NITRD) Program. (For more information on NITRD, see www.nitrd.gov.)

Establishing a virtual organization for cybersecurity science

At a high level, NSF's CPS program and the federal cyber-physical systems research portfolio can be seen as a broad research initiative intended to develop the scientific foundations for designing complex systems. Many of the activities associated with cyber-physical systems have focused on identifying the technical challenges associated with various types of complex systems. In late 2010, NITRD agencies, led by NSA and NSF, launched one such activity related to the science of dependable and secure cyber-physical systems. This effort culminated in the Workshop on Foundations of Dependable and Secure Cyber-Physical Systems, held as part of CPS Week 2011 in Chicago, Illinois. (For more information, see https://www.truststc.org/conferences/11/CPSWeek/program.htm).

The workshop focused on topics that addressed fundamental challenges of making cyber-physical systems secure, dependable, and trustworthy. Particular emphasis was placed on the control and verification challenges arising from the complex interdependencies among networked systems. Such systems are in widespread use today, controlling the operation of critical infrastructures such as power transmission, water distribution, transportation, healthcare, building automation, and process control. The combination of various factors—including the widespread use of commodity components, Internet connectivity, and the malicious intents of hackers and cybercriminals—have made these types of systems extremely vulnerable. Despite attempts to apply security-oriented design guidelines and policies, much remains to be done to achieve a scientifically grounded and principled design approach to security, trustworthiness, and dependability in these systems.

The 2011 workshop was a first formal attempt to foster collaboration among researchers from a variety of fields including control *and* systems theory, embedded systems, game theory, software verification and formal methods, and computer security. One important outcome of the workshop was the recognition that the science of cybersecurity was critical to the overall success of the CPS program *and* of the cyber-physical systems field. This recognition aligned with the vision that had been previously put forward by the NITRD HCSS CG cochairs in a white paper to the Office of Science and Technology Policy (OSTP) titled "Winning the future with science and technology for 21st century smart systems." Workshop recommendations went even further, recommending that a virtual organization dedicated to cybersecurity science be established within the CPS VO—the Science of Security Virtual Organization (SoS VO).

Growing interest in cybersecurity science

At the same time as the CPS program moved toward creating a distinct cybersecurity science group, a number of governmental initiatives in cybersecurity science began appearing from organizations across the broader cybersecurity community, including several outside of the US. Unfortunately, without the benefit of any centralized resource to help coordinate

their efforts, these activities developed in isolation. As information about these efforts became more widely available, it became clear that the SoS VO could serve an even more valuable role if it provided a focal point for all things cybersecurity science related. Together, through NSA leadership and sponsorship, Vanderbilt's design goal for the CPS VO was augmented to provide a portal with a rich set of collaboration and sharing capabilities, leveraging and extending NSF's investment to support an enhanced data repository and content management system. This coordinated effort served well the interest of both the CPS VO and SoS VO communities. While this approach was significantly more ambitious, it offered better opportunities for advancing work in both cyber-physical systems and cybersecurity science much more quickly and efficiently. The integrated approach and the resulting extended capabilities will benefit other cyber-physical systems special interest groups as they begin building their online communities.

Content is king, search is queen

From its inception, the CPS VO was intended to grow into an established research resource by offering a storehouse of information with a robust search capability to mine it efficiently. Achieving this goal meant that the virtual organization needed to attract a large user population and provide services that were valuable, engaging, and easy to use. These objectives were adopted as the guiding principles for all decisions made in augmenting support for the SoS VO. The target audience was expanded to include researchers, program managers, educators, funding agents, system designers, and students—almost anyone having an interest in cybersecurity science. Attracting such a broad group meant the SoS VO had to provide an extensive and useful assortment of information, accessible intuitively and efficiently—a very tall order. If the SoS VO is able to create an enduring engagement center for cybersecurity science, user-contributed content should generate value and further help to build a cybersecurity science community.

Evolving an SoS VO capability

After a careful assessment of the needs identified for the SoS VO, a plan was developed to roll out new capabilities in three basic areas. The first set of capabilities was geared toward establishing the SoS VO as a focal point for information about ongoing activities related to cybersecurity science and as a repository for significant research results. The second phase of development would place emphasis on community development, information sharing, and interaction among researchers in the field. The last, and most ambitious, set of capabilities envisioned for the SoS VO would help to establish and support true collaboration in advancing cybersecurity science. (See figure 1 for a screenshot of the SoS VO home page.)

SoS VO capability phases

- **Phase 1. Build a resource center.**

  Creating a centralized information resource on cybersecurity science activity is the first step planned for the SoS VO and is key to helping establish a community. An important goal of this phase involves identifying and collecting information about the disparate cybersecurity science work currently being performed. Providing descriptions and contact information for the organizations conducting and supporting cybersecurity science work is a priority, as well as advertising new program funding opportunities. For organizations currently producing reports related to cybersecurity science, the SoS VO intends to provide a centralized library for cataloging, analyzing, searching, and distributing information. A calendar of events related to cybersecurity science is a core capability of the SoS VO and will appear early with the ability to sync to users' individual calendars.

- **Phase 2. Cultivate collaboration with virtual tools.**

  The second phase of planned SoS VO capabilities is intended to expand the reach of cybersecurity science information to a much broader community of users. One of the exciting features being developed will allow videos of research reviews to be viewed online in both real-time streaming and archived formats. This capability should permit users to become involved much more easily in reviews without the time and budget constraints of long distance travel. Discussion forums, blogs, content subscriptions, chat, wikis, and user profiles are being created to permit increased interaction among users and to promote simple forms of collaboration.

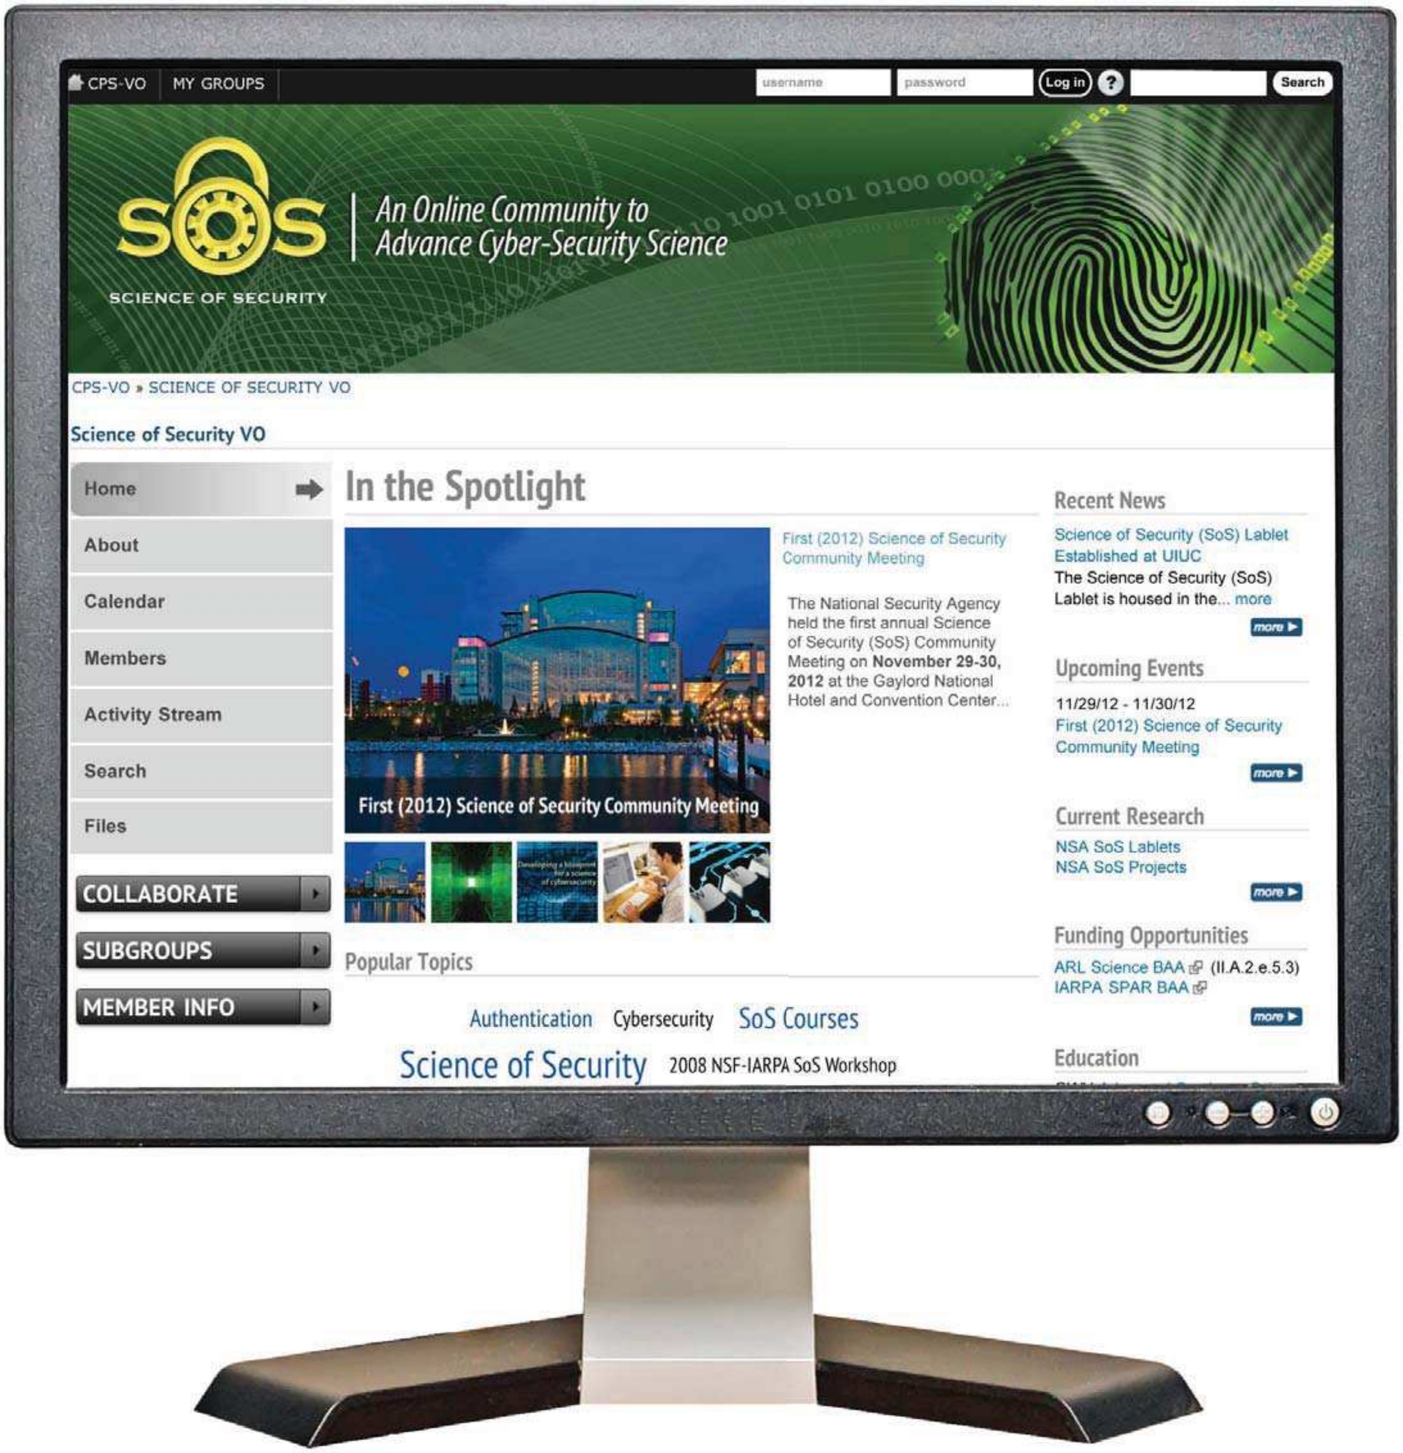

In the Spotlight

First (2012) Science of Security Community Meeting

The National Security Agency held the first annual Science of Security (SoS) Community Meeting on **November 29-30, 2012** at the Gaylord National Hotel and Convention Center...

Recent News

Science of Security (SoS) Lablet Established at UIUC

The Science of Security (SoS) Lablet is housed in the... more

Upcoming Events

11/29/12 - 11/30/12

First (2012) Science of Security Community Meeting

Current Research

NSA SoS Lablets

NSA SoS Projects

Funding Opportunities

ARL Science BAA (II.A.2.e.5.3)

IARPA SPAR BAA

Education

FIGURE 1. The Science of Security Virtual Organization (SoS VO) enables those interested in cybersecurity science to survey current research; stay current on news in the field; find out about events related to cybersecurity science; collaborate with others using chat, video conferencing, and forums; share work by uploading documents and creating wikis; and access educational resources contributed by members. Visit cps-vo.org/group/SoS to learn more.

- **Phase 3. Strengthen collaboration with social networking.**

  Ultimately, the capabilities delivered by the SoS VO, as well as the CPS VO, were conceived to promote community collaboration in order to advance science. The features deployed in the first phases of the SoS VO should help to create a broad community of users and establish a focal point for their interactions. But it is the last group of capabilities offered by the SoS VO that should enable the type of robust collaboration desired by blending elements of social networking with a rich set of communication and research tools. Some of the features currently being planned in this phase include:

  - Research toolsets and datasets;
  - On-demand video conferencing;
  - Desktop sharing;
  - Individual user space, dashboard, etc.;
  - Interface personalization;
  - Subscription services;
  - Cybersecurity science-related newsfeeds;
  - A multimedia library; and
  - Open research support.

SoS VO rollout

The establishment of the SoS VO is founded on the beliefs that open collaboration can play a key role in advancing cybersecurity science and that the availability of a platform where researchers can share, collaborate, and learn is vital to building community. The structure and features of the SoS VO attempt to leverage popular features provided by social networking technology with rich domain-specific content to create a focal point for cybersecurity science research. The pilot version of the SoS VO has evolved dramatically in form and content since its inception in 2011, and it will continue to evolve as user feedback is received when it becomes operational and as the cybersecurity science community matures.

About the authors

Frankie D. King is the assistant director of the Annapolis Technical Coordination Project Office at Vanderbilt University's Institute for Software Integrated Systems (VU-ISIS), where she is responsible for managing the coordination of collaborative research and development activities on the Cyber-Physical Systems Virtual Organization that are sponsored by federal agencies belonging to the Networking and Information Technology Research and Development (NITRD) Program. Before joining VU-ISIS, King served as the technical coordinator for the High Confidence Software and Systems Program Component Area at the National Coordination Office for the NITRD Program for nearly seven years. King has over 28 years of program development and management experience in domestic and international policy affairs where she has served in high-level capacities in the executive and legislative branches of the US government and the private sector. King received an MA from the University of Notre Dame and a BA from Fisk University, where she graduated summa cum laude.

Heather Lucas is a program director within the Trusted Systems Research Group at NSA and is the current program lead for NSA's Science of Security Virtual Organization effort.

Robert Meushaw is the former technical director of NSA's Information Assurance (IA) Research Laboratory. His current work focuses on developing new strategies and programs for the advancement of a science of cybersecurity. He retired from NSA in 2005 after 33 years of service, including over a decade of work in IA research. Meushaw's career at NSA also included significant stints in both the Product Development Group and the Security Evaluation Group of the IA Directorate. In addition to his technical responsibilities, he served for six years as technical editor of NSA's *Tech Trend Notes* and *The Next Wave* publications. Meushaw holds degrees in electrical engineering from Princeton University and the Johns Hopkins University.

UK’s new Research Institute investigates the science of cybersecurity

Government Communications Headquarters (GCHQ)

How do we know when we are “secure enough”? How do we decide how best to spend our precious security budget? How do we reduce our reliance on individual expert judgement and make better, more objective security decisions? It is always challenging to bring scientific rigor to bear on a complex, real world problem, and this challenge applies in spades to the relatively young discipline of cybersecurity. Practitioners must work hard to stay on top of ever changing technologies and a rapidly evolving threat environment, and simply keeping abreast of “best practice” is challenging. Yet we must—if we want to ever get ahead of the curve—develop a more systematic, rigorous approach based on foundational scientific knowledge and understanding.

The UK government recently announced the formation of a virtual Research Institute to improve understanding of the science behind the growing cybersecurity threat. The Institute, which is funded by a £3.8 million grant ($6.14 million US), is part of a cross-government commitment toward increasing the nation’s academic capability in all fields of cybersecurity.

Established by the Government Communications Headquarters (GCHQ), in partnership with the UK Research Councils (RCUK) and the Department for Business, Innovation and Skills (BIS), the Research Institute is a virtual organization involving seven universities. It will allow leading academics in the field of cybersecurity, including social scientists, mathematicians, and computer scientists from across the UK, to work together. It will also connect them with the collective expertise of industry security experts and international researchers in the field—with a particularly close relationship expected with the US. The Research Institute opened for business on October 1, 2012, and is funded for a period of three and a half years.

Universities were selected following a tough competitive process in which they had to devise new research programs to address one of two key challenges:

- How secure is my organization?
- How do we make better security decisions?

Addressing these very practical challenges requires a blended approach from researchers, drawing from both technological and behavioral disciplines. Four teams were successful:

- University College London, working with University of Aberdeen;
- Imperial College, working with Queen Mary College and Royal Holloway, University of London;
- Royal Holloway, University of London; and
- Newcastle University, working with Northumbria University.

University College London (UCL) was selected to host the Research Institute, with Professor Angela Sasse taking the role of director of research. At the press launch, Sasse acknowledged the strong multidisciplinary nature of the research portfolio, saying, "I am delighted to be leading the new Research Institute. This is an opportunity to work closely with colleagues from different scientific disciplines to tackle the technical, social, and psychological challenges that effective cybersecurity presents."

As well as being cross-disciplinary, the research portfolio is an exciting blend of theoretical work and experimentation in "the field"—with "the field" meaning real organizations, operational information technology (IT) systems, and real, live users. The work is unusual in being focused firmly on improving security within organizations rather than for individual citizens. It is equally applicable to governmental or commercial organizations. The collaborative approach between academia, industry, and government will ensure that research is relevant and inspired by real world, cutting edge security issues.

The winning projects

UCL's project is entitled "Productive security: Improving security compliance and productivity through measurement" and will focus on the behavior of users within the workplace. This work builds on a growing body of evidence that security policies and control are not fully effective because employees either cannot or will not comply with them [1, 2]. A key reason for noncompliance is the combination of employee workload and the complexity of security controls chosen. Yet many security decision makers do not factor the impact on employees, their tasks, and the company's business processes into their decision about which security controls to put in place. Current attempts to educate employees about the need for security are of questionable effectiveness because they simply push more information on people who are already overworked. Even in organizations with high security awareness, noncompliance can be observed because the security policy causes excessive friction or is not agile enough to meet the needs of the business [3, 4].

The project will work with at least two major companies to collect data on employees' workload, risk perception, and the resulting security behaviors. It will use that data to develop a decision support model to allow security professionals to balance the impact of security controls on employees and business processes against the risk mitigation the controls can achieve.

The lead researchers are Professor Angela Sasse of UCL and Professor David Pym of University of Aberdeen.

In contrast to UCL, the three-party team led by Imperial College will work on the Research Institute's most heavily theoretical program. The project, "Games and abstraction: The science of cybersecurity," will develop new approaches to decision support based on mathematical game theory. The project is academically ambitious in attempting to combine three major disciplines: game theory, machine learning, and abstract interpretation. For example, no connection has been established so far between abstract interpretation and these other areas.

Game theory, the theory developed for the mathematical analysis of multiperson strategic decision making [6], has been increasingly applied in the last decade in cybersecurity. Examples of applications can be found in the fields of intrusion detection systems, anonymity and privacy, economics of network security, and cryptography. A state of the art survey of these applications is given in Alpacan and Basar's *Network Security: A Decision and Game Theoretic Approach* [7]. This new work will build on the game theoretical model developed by Lye and Wing [5]. A limitation of this work is that the attacker model is based on a set of known strategies; part of the proposed research is to extend the approach to deal with previously unseen attacks (e.g., zero days) and emerging behaviors. The research objectives are to:

FIGURE 1. The University College London will host the Research Insitute, a virtual organization that will bring together cybersecurity experts from around the world.

- Model complex scenarios by developing mathematical abstraction techniques for stochastic games, using techniques originating in probabilistic abstract interpretation and machine learning;
- Provide a precise way to analyze how results of optimal behavior in the abstract models relate to the optimal or near-optimal behaviors in complex real scenarios; and
- Demonstrate the results by proof-of-concept implementations and test on realistic data provided through empirical studies.

The lead researchers are Professor Chris Hankin of Imperial College; Professor Dusko Pavlovic of Royal Holloway, University of London; and Dr. Pasquale Malacaria of Queen Mary College.

Royal Holloway, University of London's project is entitled "Cybersecurity cartographies." Its goal is to develop ways of visualizing the different means in which both people and technology protect important data. The project brings together the disciplines of art and design, network security, and organizational security in order to develop a range of visualization techniques that better inform security managers about the strength of data protection across their cyber estate.

Security managers use a combination of organizational, physical, and technical controls to provide robust information asset protection. Control lists, such as those in Annex A of ISO 27001 (i.e., an information security management system standard), have long acknowledged the need for the three types of control, but no methods are available to systematically combine them. In addition, risk management techniques do not include visualization methods that can present a combined picture. To address these gaps, the project will further develop existing research on the influence of cultural and organizational techniques on policy compliance [8]. It will also develop techniques to combine interpretive cartography with informational cartography using a visualization framework [9]. In addressing these gaps, the work will help security managers to develop well informed trade-offs between security and other business drivers, while supporting

FIGURE 2. The Research Insitute's director of research is Professor Angela Sasse of University College London.

their existing skills and expertise.

The lead researcher is Dr. Lizzie Coles-Kemp of Royal Holloway, University of London.

Finally, Newcastle University is working on the project "Choice architecture for information security." Newcastle's research hypothesis is that there exists a rigorous choice architecture which will nudge decision makers to make demonstrably better information security decisions. Newcastle's approach takes inspiration from the work on nudging from the behavioral economics community [10]. Nudging provides a framework to influence decision makers in a subtle way. The theory will be applied to scenarios relating to consumerization [11] (i.e., the use of personal devices in the workplace) and will also be relevant to the broader issue of work-life integration (i.e., the blurring of the boundaries between work and home life).

In addition, part of the novelty of the approach will be the ability to integrate rigorous security assessment with psychological ownership models adapted from the occupational psychology literature [12, 13].

The research objectives are to:

- Understand the psychological phenomena that dictate security behavior relevant for data protection in consumerization scenarios, from the various perspectives of the chief information security officer, IT administrators, and employees;
- Develop a choice architecture for these scenarios;
- Implement a toolset to implement the choice architecture—steering the decision maker to "better" decisions; and
- Experimentally evaluate the improvements delivered.

The lead researchers are Dr. Aad van Moorsel of Newcastle University and Professor Pamela Briggs of Northumbria University.

Conclusion

In mid-2012, GCHQ, BIS, and RCUK awarded the Academic Center of Excellence (ACE) in Cybersecurity Research to eight UK universities [14]. This initiative, the first part of a broad, joint response to the UK government's national cybersecurity strategy [15], will enhance the UK's cyber knowledge through original research.

The establishment of the Research Institute is another part of the broad response to the UK government's national cybersecurity strategy [15]. The strategy describes how the government is working with academia and industry to make the UK more resilient to cyberattacks. Both the ACE and the Research Institute initiatives are harnessing the vital role that academia has to play in supporting and developing the UK's capability in cybersecurity.

About GCHQ

Government Communications Headquarters (GCHQ) is one of three UK intelligence agencies. GCHQ provides intelligence, protects information, and informs relevant UK policy to keep our society safe and successful in the Internet age.

References

[1] Sasse MA, Brostoff S, Weirich D. “Transforming the ‘weakest link’—a human-computer interaction approach to usable and effective security.” *BT Technology Journal.* 2001;19(3):122–131. DOI: 10.1023/A:1011902718709

[2] Beautement A, Sasse MA, Wonham M. “The compliance budget: Managing security behaviour in organizations.” In: *Proceedings of the 2008 New Security Paradigms Workshop;* Sep 2008; Lake Tahoe, CA; p. 47–58. DOI: 10.1145/1595676.1595684

[3] Pallas F. “Information security inside organizations—a positive model and some normative arguments based on new institutional economics” [PhD thesis]. [Berlin (Germany)]: Technical University of Berlin; 2009.

[4] Albrechtsen E, Hovden J. “The information security digital divide between information security managers and users.” *Computers and Security.* 2009;28(6):476–490. DOI: 10.1016/j.cose.2009.01.003

[5] Lye KW, Wing J (2005). “Game strategies in network security.” *International Journal of Information Security.* 2005;4(1–2):71–86. DOI: 10.1007/s10207-004-0060-x

[6] Neumann J, Morgenstern O. *Theory of Games and Economic Behavior.* Princeton (NJ): Princeton University Press; 1944. ISBN-13: 978-0-691-13061-3

[7] Alpacan T, Basar T. *Network Security: A Decision and Game Theoretic Approach.* Cambridge (MA): Cambridge University Press; 2011. ISBN-13: 978-0-521-11932-0

[8] Pieters W, Coles-Kemp L. “Reducing normative conflicts in information security.” In: *Proceedings of the 2011 New Security Paradigms Workshop;* Sep 2011, Marin County, CA: p. 11–24. DOI: 10.1145/2073276.2073279

[9] Hall P. “Bubbles, lines and string: How information visualization shapes society.” In: Blauvelt A, Lupton E, editors. *Graphic Design Now in Production.* Minneapolis (MN): Walker Art Center; 2011. p. 170–185.

[10] Sunstein C, Thaler R. *Nudge: Improving Decisions about Health, Wealth, and Happiness.* New Haven (CT): Yale University Press; 2008. ISBN-13: 978-0-300-12223-7

[11] Microsoft Corporation (2011). “Strategies for embracing consumerization.” 2011. Available at: http://download.microsoft.com/download/E/F/5/EF5F8B95-5E27-4CDB-860F-F982E5B714B0/Strategies%20for%20Embracing%20Consumerization.pdf

[12] Ifinedo P. “Understanding information systems security policy compliance: An integration of the theory of planned behavior and the protection motivation theory.” *Computers and Security.* 2012;31(1):83–95. DOI: 10.1016/j.cose.2011.10.007

[13] Aurigemma S, Panko R. “A composite framework for behavioral compliance with information security policies.” In: *Proceedings of the 45th Annual Hawaii International Conference on System Sciences;* Jan 2012, Maui, HI: p. 3248–3257. DOI: 10.1109/HICSS.2012.49

[14] GCHQ. “UK universities awarded Academic Centre of Excellence status in Cyber Security Research” [Press release]. Available at: http://www.gchq.gov.uk/Press/Pages/Cyber-Security-Research-Centres-of-Excellence.aspx

[15] UK government cabinet office. “The UK Cyber Security Strategy: Protecting and promoting the UK in a digital world,” 25 Nov 2011. Available at: http://www.cabinetoffice.gov.uk/sites/default/files/resources/uk-cyber-security-strategy-final.pdf

Securing America's digital infrastructure through education

Bill Newhouse

On May 29, 2009, in the East Room of the White House, President Barack Obama announced that his administration will pursue a new comprehensive approach to securing America's digital infrastructure. During the speech on "Securing our nation's cyber infrastructure," [1] he noted the following:

> *. . . we will begin a national campaign to promote cybersecurity awareness and digital literacy from our boardrooms to our classrooms, and to build a digital workforce for the 21st century. And that's why we're making a new commitment to education in math and science, and historic investments in science and research and development. Because it's not enough for our children and students to master today's technologies—social networking and emailing and texting and blogging—we need them to pioneer the technologies that will allow us to work effectively through these new media and allow us to prosper in the future.*

"Building capacity for a digital nation," part II of the president's cyberspace policy review [2], included recommendations around the idea that the general public needs to be well informed to use technology safely, that the US needs a technologically advanced workforce to remain competitive in the twenty-first century economy, and that math and science must be a priority in schools. The review suggested that the US should initiate a K–12 cybersecurity education program for digital safety, ethics, and security; expand university curricula; and set the conditions to create a competent workforce for the digital age. To help achieve these goals, the review stated that the nation should:

- Promote cybersecurity risk awareness for all citizens;
- Build an education system that will enhance understanding of cybersecurity and allow the US to retain and expand upon its scientific, engineering, and market leadership in information technology;
- Expand and train the workforce to protect the nation's competitive advantage; and
- Help organizations and individuals make smart choices as they manage risk.

In response to the president's cyberspace policy review, the National Security Staff (NSS)'s Cybersecurity Directorate and the Office of the Director of National Intelligence (ODNI)'s Joint Interagency Cyber Task Force formed an interagency working group to expand the Comprehensive National Cybersecurity Initiative (CNCI)'s initiative #8—Expand Cyber Education—to encompass a national, rather than federal, focus. The goal of the working group was to formulate a recommendation for the Information and Communications Infrastructure Interagency Policy Committee (ICI-IPC) on a way forward for a national program to improve cybersecurity awareness, education, workforce structure, and training.

The working group consisted of representatives from the NSS Cybersecurity Directorate staff; ODNI; the Departments of Commerce, Defense (DoD), Education, Homeland Security (DHS), Justice (DoJ), Labor (DoL), State, and Treasury; NSA; the Office of Personnel Management (OPM); the Office of Management and Budget; and the Office of Science and Technology Policy. The group worked for several months to finalize a recommendation to the ICI-IPC on the governance model for a national cybersecurity education program. The recommendation resulted in the March 2010 creation of an interagency structure and governance model for the National Cybersecurity Education Initiative, renaming it the National Initiative for Cybersecurity Education (NICE) [3].

National Initiative for Cybersecurity Education (NICE)

With NICE, the federal government aims to enhance the overall cybersecurity posture of the US by accelerating the availability of educational and training resources designed to improve the cyber behavior, skills, and knowledge of every segment of the population. This will enable a safer cyberspace for all. The initiative has established three underlying goals:

- Raise national awareness about risks in cyberspace,
- Broaden the pool of individuals prepared to enter the cybersecurity workforce, and
- Cultivate a globally competitive cybersecurity workforce.

The recommendation identified the National Institute of Standards and Technology (NIST) as the overall lead with four components (shown in figure 1).

Interagency structure

NICE will be represented by the following four components.

1. **National cybersecurity awareness campaign.** The goal of this component, led by DHS, is to improve the cybersecurity behavior of the American public. DHS is doing this by delivering a national public awareness campaign—Stop.Think.Connect. [4]—aimed at increasing the understanding of cyber threats and empowering the American public to be safer and more secure online. A core strategy of the campaign is a National Cyber Awareness Coalition [5], which comprises federal agency partners as well as state and local governments. The Coalition offers a mechanism for message and materials dissemination. Making effective use of the communications channels and outreach capabilities of the Coalition members is key to extending

the campaign's reach. Projects within this component include:

- Planning and executing Cyber Tours [6] nationwide to directly engage communities in promoting awareness and initiating a dialogue about the dangers individuals face online;
- Launching and expanding the National Network, a spin-off of the National Cyber Awareness Coalition, which will mirror the Coalition but be open for membership from any national nonprofit organization;
- Improving the Stop.Think.Connect. resources, such as the Toolkit [7];
- Finding new outreach opportunities and mechanisms to spread the campaign's message; and
- Increasing coordination of the campaign and National Cyber Security Awareness Month (NCSAM), including incorporating Stop.Think.Connect. language in the state proclamations and conducting a Cyber Tour during NCSAM.

2. **Formal cybersecurity education.** The goal of this component, led by the Education Department and National Science Foundation (NSF), is to broaden the pool of skilled workers for a cybersecure nation. It is responsible for supporting formal education to increase both the number of people with cybersecurity knowledge, skills, and abilities and the quality of the cybersecurity capabilities held by those people. Projects within this component include:

   - Making the connection between cybersecurity and science, technology, engineering, and mathematics (STEM);
   - Disseminating common evidence standards in pre-K–12 education;
   - Promoting the growth of effective cybersecurity competitions in high schools and higher education;
   - Facilitating the development of curricular recommendations in high schools and higher education; and
   - Coordinating a learning network of virtual national cybersecurity laboratories.

3. **Cybersecurity workforce structure.** The goal of this component, led by DHS and supported by OPM, is to define cybersecurity jobs, attraction, recruitment, retention, and career path strategies. This component contains the following subcomponent areas: the federal workforce (led by OPM), the government (nonfederal) workforce (led by DHS), and the private sector workforce (led by the Small Business Administration, DoL, and NIST).

   This component focuses on talent management of cybersecurity professionals. It aims to evaluate the professionalization of the workforce, recommend best practices for forecasting future cybersecurity needs, and define national strategies for recruitment and retention. Projects within this component include:

   - **Professionalization—**establishing a methodology for identifying cybersecurity areas to be professionalized [8] and providing a central national resource for cybersecurity professionalization.
   - **Workforce planning—**delivering a methodology for accurately forecasting cybersecurity workforces across government, industry, and academia.

- **Recruitment and retention—**providing, disseminating, and maintaining a strategy and set of materials for recruiting and retaining cybersecurity professionals at the national level.

4. **Cybersecurity workforce training and development.** The goal of this component, led by DHS, DoD, and ODNI, is to develop and maintain an unrivaled cybersecurity workforce. It contains the following functional areas: general IT use (led by DHS and the Department of the Navy); information technology infrastructure, operations, maintenance, and information assurance (led by DoD and DHS); domestic enforcement and counterintelligence (led by the Defense Cyber Crime Center, the Office of the National Counterintelligence Executive, DoJ, and the US Secret Service); and specialized cybersecurity operations (led by NSA).

   This component is responsible for defining the cybersecurity workforce and identifying the training and professional development required for the nation's cybersecurity workforce. Projects within this component include:

   - **National Cybersecurity Workforce Framework [9]—**providing a common language to define cybersecurity work. The Framework defines specialty areas; knowledge, skills, and abilities (KSAs); and competencies.
   - **Training catalog/National Institute for Cybersecurity Studies portal—**serving as a national online resource for information about cybersecurity awareness, education, careers, and professional development. It provides an online web resource that has a robust and representative collection of training opportunities mapped to the National Cybersecurity Workforce Framework.
   - **Workforce inventory—**collecting data to baseline and identify the current state of the IT workforce and assess current cybersecurity capabilities.
   - **Training gap analysis—**ensuring that available training is appropriate in terms of quality, need, and content.

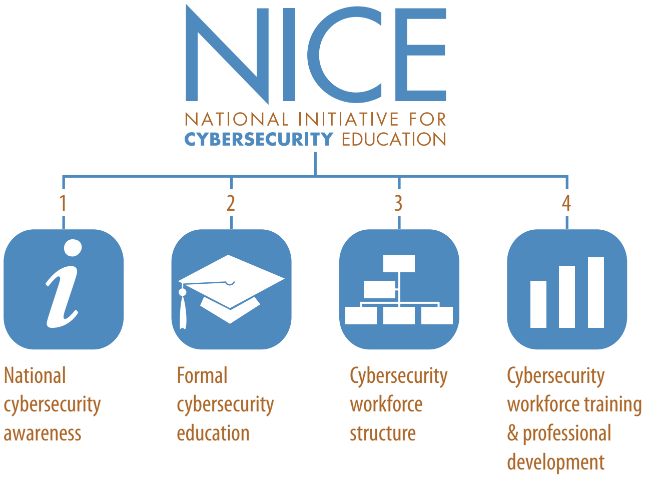
NICE
NATIONAL INITIATIVE FOR
CYBERSECURITY EDUCATION
1
2
3
4
National cybersecurity awareness
Formal cybersecurity education
Cybersecurity workforce structure
Cybersecurity workforce training & professional development

FIGURE 1. The National Initiative for Cybersecurity Education (NICE) is broken into four components aimed at enhancing the overall cybersecurity posture of the US.

   - **Professional development road maps—**developing resources which depict career progression from entry to expert within each specialty area.

Relationship to the cybersecurity R&D science of security thrust

In December 2011, the White House released "Trustworthy cyberspace: Strategic plan for the federal cybersecurity research and development program" [10] that included a thrust on developing scientific foundations. This thrust challenges the research and development (R&D) community to organize the knowledge in the field of cybersecurity and to investigate universal concepts that are predictive and transcend specific systems, attacks, and defenses resulting in a cohesive understanding of underlying principles of cybersecurity. This thrust will enable investigations that affect large-scale systems and will promote the development of hypotheses subject to experimental validation; it will also support high-risk explorations needed to establish a scientific basis and to form public-private partnerships of government agencies, universities, and industry.

NICE seeks to organize the knowledge in the field of cybersecurity education by supporting the development of cybersecurity awareness and educational content appropriate for different audiences and students.

NICE also seeks to identify and develop consensus on universal concepts that support increased cybersecurity awareness, expand cybersecurity education, and nurture a cybersecurity workforce that is prepared to support our nation's future.

NICE will continue to form public-private partnerships to achieve its goals. Leadership from the private and academic sectors is critical to the success of the NICE strategy to help organize disparate areas of knowledge. The R&D strategy noted that developing a strong, rigorous scientific foundation to cybersecurity helps the field by providing structure and organization to a broad-based body of knowledge in the form of testable models and predictions. This is true for NICE as well, but rather than testable models and predictions, NICE needs to develop common core state standards [11] for cybersecurity that will enable cybersecurity to be incorporated into K–12 education. The formation of cybersecurity education and awareness into a common core standard like the one already designed for mathematics [12] will help define what students should understand and be able to demonstrate in their study of cybersecurity.

Increased exposure to cybersecurity concepts, including computational thinking [13] in K–12 education, and an overall STEM emphasis in K–12 education will produce more students with the skills necessary to perform cybersecurity R&D as they matriculate through universities, academies, colleges, and institutes of technology. NICE believes that the innovative skills gained while performing R&D in an academic environment will translate into more people capable of performing and leading cybersecurity R&D activities within both the federal government and the nation's high-tech industries. NICE also recognizes the need to keep up with the innovations developed by the R&D community as the initiative continues its pursuit of its strategic goals.

The science of cybersecurity workforce

The National Cybersecurity Workforce Framework provides a common set of definitions for the cybersecurity workforce. The Framework brings consistency to how cybersecurity work is defined and described. It provides a common language to discuss and understand the work requirements of cybersecurity professionals, empowering our nation's agencies and industries to:

- Baseline capabilities,
- Identify skill gaps,
- Develop cybersecurity talent in the workforce, and
- Prepare the pipeline of future talent.

The Framework organizes the cybersecurity workforce into seven broad categories, then into thirty-one specialty areas. These specialty areas are further broken down into work roles and then KSAs. Some organizations may mix roles or specialty areas; this is a major strength for the Framework in that it can be customized to fit the needs of an organization and still maintain its integrity. The Framework was developed in collaboration with subject matter experts from government, nonprofits, academia, and the private sector.

The Framework concept began before the establishment of NICE and grew out of the recognition that the cybersecurity workforce (federal and private industry) could not be measured and that the roles needed to support our nation's cybersecurity were undefined. To combat this challenge, the federal Chief Information Officers (CIO) Council [14] began a Cybersecurity Workforce Development Matrix effort in 2008, when the organization was tasked to provide a standard framework to understand the cybersecurity roles within the federal government. In 2008, the CIO Council's Information Technology Workforce Committee (ITWC) conducted an environmental scan and produced a research report that referenced where other information technology professional development efforts were also underway, including the "Essential Body of Knowledge (EBK) report" and "The Committee of National Security Systems (CNSS) standards." Specific roles were identified as needed by agencies to conduct cybersecurity work.

In November 2011, thirteen roles were identified and four cybersecurity development matrices were published by the federal CIO Council along with the "Cybersecurity workforce development matrix resource guide" [15] to instruct managers on how to use the matrices. The roles and initial matrices were created based on input from focus groups consisting of subject matter experts from many federal agencies. The federal CIO Council's Information Security and Identity Management Committee (ISIMC) and ITWC advised on the project. Plans are underway to link the matrices to the Framework by providing sample illustrations of how the specialty areas within the Framework can be mapped to create various cybersecurity roles.

The Framework is comprehensive and inherently flexible, allowing organizations to adapt its content to their human capital and workforce planning needs. The work conducted in the federal CIO Council's Cybersecurity Workforce Development Matrix project will be leveraged to provide government organizations with sample applications of how they can adjust the Framework to suit their own workforce needs. These sample applications provide an option for each department or agency to customize their template through the Framework model. Over time, these examples will be expanded to include the education, experience, credentials, and training needed by an individual for each role.

The Framework [9], published in August 2012, enabled the issuance of cybersecurity functional codes by OPM on October 1, 2012, in their "Guide to data standards" [16]. Use of these cybersecurity function codes will enable OPM and federal agencies to identify the cybersecurity workforce; determine baseline capabilities; examine hiring trends; identify skill gaps; and more effectively recruit, hire, train, develop and retain a valuable cybersecurity workforce.

An increased focus on the science of security at our nation's institutions of higher learning based on the R&D strategic plan's thrust of developing scientific foundations will produce graduates ready to enter the cybersecurity workforce with the skills to organize disparate areas of knowledge, leverage the universal laws to be discovered, and apply scientific method to their work. The National Cybersecurity Workforce Framework developers recognize that it will be vital for the workforce and science and technology communities to work together to acknowledge and communicate the importance of these skills and other newly discovered KSA's needed within our nation's workforce.

NICE end-state vision

Looking to the future, NICE envisions a developed workforce that is prepared to ensure an organized and unified response to cyber incidents. NICE envisions a nation that is prepared to work together to secure America's information and communications networks. Public-private partnerships, established to meet the NICE goals, will continue to collaborate to meet the demands of new threats and to utilize cutting-edge R&D which is delivering the innovation and discovery that the nation needs to meet the challenges of our time. NICE envisions increased cybersecurity awareness from our boardrooms to our classrooms and a strong cybersecurity workforce for the twenty-first century.

About the author

Bill Newhouse is a cybersecurity program lead in the Computer Security Division, one of six divisions in the Information Technology Laboratory at National Institute of Standards and Technology (NIST). Newhouse represents NIST in several collaborative efforts including (1) the National Initiative for Cybersecurity Education, (2) a partnership with the Department of Homeland Security and the financial sector to develop and test innovative cybersecurity technologies and processes, and (3) as a member of federal interagency cybersecurity R&D committees.

Before coming to NIST in 2010, Newhouse spent five years in the Office of the Secretary of Defense, where he focused on the cybersecurity and information assurance R&D portfolio, first with the assistant secretary of defense for research and engineering and then with the assistant secretary of defense for Networks and Information Integration (NII). While in NII, he championed Defense Venture Catalyst Initiative workshops to focus on cybersecurity solutions from innovative companies. He is an electrical engineering graduate of both the Georgia Institute of Technology and George Washington University and has been with the federal government for over 25 years, beginning as a cooperative education student at NSA in 1986.

References

[1] The White House Office of the Press Secretary. "Remarks by the President on securing our nation's cyber infrastructure." 2009 May 29. Available at: http://www.whitehouse.gov/the-press-office/remarks-president-securing-our-nations-cyber-infrastructure

[2] The White House. "Cyberspace policy review: Assuring a trusted and resilient information and communications infrastructure." Available at: http://www.whitehouse.gov/assets/documents/Cyberspace_Policy_Review_final.pdf

[3] To learn more about the National Initiative for Cybersecurity Education, visit http://csrc.nist.gov/nice/.

[4] To learn more about Stop.Think.Connect., visit http://stopthinkconnect.org/.

[5] To learn more about the National Cyber Awareness Coalition, visit http://www.dhs.gov/stopthinkconnect-coalition.

[6] To learn more about the Cyber Tours program, visit http://stopthinkconnect.org/get-involved/homeland-security-campaign/cyber-tours-program.

[7] To learn more about the Stop.Think.Connect. Toolkit, visit http://www.dhs.gov/files/events/stop-think-connect-campaign-materials.shtm#1.

[8] To learn more about the project to professionalize the nation's workforce, visit http://sites.nationalacademies.org/CSTB/CurrentProjects/CSTB_070783.

[9] To learn more about the National Cybersecurity Workforce Framework, visit http://csrc.nist.gov/nice/framework/

[10] Executive Office of the President National Science and Technology Council. "Trustworthy cyberspace: Strategic plan for the federal cybersecurity research and development program." 2011 Dec. Available at: http://www.whitehouse.gov/sites/default/files/microsites/ostp/fed_cybersecurity_rd_strategic_plan_2011.pdf

[11] To learn more about common core state standards, visit http://www.corestandards.org/.

[12] National Governor's Association and Council of Chief State School Officers. "Common core state standards for mathematics." Available at: http://www.corestandards.org/assets/CCSSI_Math%20Standards.pdf

[13] To learn more about computational thinking, visit http://www.cs.cmu.edu/~CompThink/.

[14] To learn more about the Chief Information Officers Council, visit http://www.cio.gov/.

[15] Chief Information Officers Council. "Cybersecurity workforce development matrix resource guide." 2011 Oct. Available at: http://www.cio.gov/documents/Cybersecurity_Workforce_Development_Matrix_Resource_Guide_Oct_2011.pdf

[16] Office of Personnel Management. "Guide to data standards." 2012 Oct. 01. Page A-107. Available at: http://www.opm.gov/feddata/GDS/GDS_A09.pdf

Toward a secure and trustworthy cyberspace

Nina Amla, Vijayalakshmi Atluri, Jeremy Epstein, Sol Greenspan, Peter Muhlberger, Victor P. Piotrowski, Andrew Pollington, Kevin Thompson, Zhi Tian, and Sam Weber

Cyberspace, a global "virtual" village enabled by hyperconnected digital infrastructures, has transformed the daily lives of people for the better. Regardless of distance and location, families and friends can see and talk with one another as if in the same room. Cyber economies create new opportunities. Every sector of the society, every discipline, has been transformed by cyberspace. It is no surprise that today cyberspace is critical to our national priorities in commerce, education, energy, financial services, healthcare, manufacturing, and defense.

The rush to adopt cyberspace, however, has exposed its fragility. The risks of hyperconnectedness have become painfully obvious. The privacy of personally identifiable information is often violated on a massive scale by persons unknown. Competitive advantage is eroded by the exfiltration of significant intellectual property. Law enforcement is hobbled by the difficulty of attribution, by national boundaries, and by uncertain legal and ethical frameworks. All these concerns now affect the public's trust of cyberspace and the ability of institutions to fulfill their missions.

Cybersecurity is arguably the most important challenge confronting society in the information age. No one—whether government, business, or individual—is exempt from the ravages of malicious cyber acts upon information technologies. The intelligent cyber adversary, whether human or software, learns and evolves to exploit, disrupt, and overpower cyber defenses, even as they are improved and strengthened. But posing cyber conflict solely in terms of classic attackers and defenders shortchanges the diversity and subtlety of the motivations, incentives, ethics, asymmetries, and strategies of the constituent actors and players in cyberspace. Addressing the challenge of securing cyberspace requires a coordinated multidisciplinary approach including computer scientists, mathematicians and statisticians, economists, behavioral scientists and sociologists, education experts, and engineers from many areas, all contributing to the body of knowledge on cybersecurity. Ultimately, the goal of such a multidisciplinary effort is the development of a science of cybersecurity, leading to practical, usable, and deployable technologies.

As a step toward creating such a science of cybersecurity, the National Science and Technology Council (NSTC) with the cooperation of the National Science Foundation (NSF) put forth a 2011 report, "Trustworthy cyberspace: Strategic plan for the federal cybersecurity research and development program" [1]. The plan identifies a broad, coordinated research agenda to make cyberspace secure and trustworthy. Research in cybersecurity must "change the game," check the misuses of cyber technology, bolster education and training in cybersecurity, establish a science of cybersecurity, and transition promising cybersecurity research into practice. The objective is to make cyberspace worthy of the public's trust.

NSF's Secure and Trustworthy Cyberspace (SaTC) program

NSF's new program for secure and trustworthy cyberspace (SaTC) supports the NSTC strategic plan for a trustworthy cyberspace. It recognizes that cyberspace will continue to grow and evolve and that advances in the sciences and technologies will create new leap-ahead opportunities expanding cyberspace. It recognizes that cybersecurity must also grow and coevolve along with cyberspace and that a secure and trustworthy cyberspace will ensure continued economic growth and future technological innovation.

The SaTC program is seeking research proposals that address cybersecurity from three distinct perspectives:

- Trustworthy computing systems;
- Social, behavioral, and economic sciences; and
- Transition to practice.

In addition, the SaTC program is seeking research proposals that integrate research addressing two or more of these perspectives, as well as proposals focusing entirely on cybersecurity education.

The following sections of this article describe the SaTC cybersecurity research perspectives. Each section outlines the projects and proposals that are of interest to the SaTC program within the relevant research perspective.

Trustworthy computing systems perspective

The trustworthy computing systems perspective aims to provide the basis for designing, building, and operating a cyber infrastructure with improved resistance and improved resilience to attack that can be tailored to meet a wide range of technical and policy requirements, including both privacy and accountability. The broad scope of this work supports all research approaches from theoretical to experimental, including participation by human subjects. Theories, models, cryptography, algorithms, methods, architectures, languages, software, tools, systems, and evaluation frameworks are all of interest as potential research projects.

Of particular interest is research that addresses how better to design desired security and privacy properties into components and systems. Methods for raising attacker costs by incorporating diversity and change into systems, while preserving system manageability, are also relevant.

The SaTC program welcomes studies of the trade-offs among trustworthy computing properties (e.g., security and usability, or accountability and privacy) as well as work that examines the tension between security and human values, such as openness and transparency. Also, methods to assess, reason about, and predict system trustworthiness, including observable metrics, analytical methods, simulation, experimental deployment—especially deployment on live test beds for experimentation at scale—will be considered. Statistical, mathematical, and computational methods in the area of cryptographic methods, new algorithms, risk assessments, and statistical methods in cybersecurity are also of interest to the program.

Social, behavioral, and economic sciences perspective

Research addressing the social, behavioral, and economic sciences (SBE) perspective of cybersecurity may focus on the individual, group, organizational, market, and societal levels, identifying cybersecurity risks and exploring the feasibility of potential solutions. All research approaches, including (but not limited to) theoretical, experimental, observational, statistical, survey, and simulation-based are of interest.

A variety of methods can be used in research from the SBE perspective, including field data, laboratory experiments, observational studies, simulations, and theoretical development.

Not all work that examines aspects involving people falls within the SBE perspective. If such aspects are not the primary focus of the proposal, or if the aspects involving people merely apply the social, behavioral, or economic sciences instead of contributing to them, the proposal might fit under the trustworthy computing systems perspective as human factors research.

Research with the SBE perspective as its primary perspective must have the social, behavioral, or economic sciences as its main focus and must involve theoretical or methodological contributions to those sciences. Contributions to the social, behavioral, or economic sciences may include identifying generalizable theories and regularities and should push the boundaries of the current understanding of social, behavioral, or economic phenomena in cybersecurity. The SaTC program seeks research that holds the promise of constructing new social, behavioral, or economic science theories that would apply to a variety of domains, or new generalizations of existing theory which clarify the conditions under which such generalizations hold (i.e., scope conditions).

More inductive or interpretative approaches may contribute to the social, behavioral, or economic sciences as well, especially if they set the groundwork for generalizable research or reveal broad connections that advance understanding in those sciences. The SBE perspective proposals should clearly state and elaborate how the proposed research will contribute to the social, behavioral, or economic sciences. Research proposals that involve the SBE perspective, but not as their primary perspective, must include at least an application of the social, behavioral, or economic sciences but need not involve a theoretical or methodological contribution.

All SBE perspective proposals must, like all SaTC proposals, also contribute toward the goal of creating a secure and trustworthy cyberspace. The social, behavioral, or economic sciences contribution of any SBE perspective proposal must be related to bringing about that goal.

The strongest research proposals should demonstrate the capabilities of the research team to bring to bear state-of-the-art research in the human sciences. These proposals should seek to understand, predict, and explain prevention, attack, and/or defense behaviors and should contribute to developing strategies for remediation. Proposals that contribute to the design of incentives, markets, or institutions to reduce either the likelihood of cyberattack or the negative consequences of cyberattack are especially welcome, as are proposals that examine incentives and motivations of individuals.

Research proposals submitted with an SBE perspective will be evaluated with careful attention to their:

- Mutual application of, and contribution to, basic social, behavioral, or economic science research;
- Generalizability to multiple cybersecurity settings;
- Ultimate contribution to the construction of institutions that induce optimal behavior; and
- Value toward creating a secure and trustworthy cyberspace.

Given the nascent state of social, behavioral, and economic science research in cybersecurity, work that proposes workshops and other opportunities for intellectual engagements is welcomed. Such proposals, however, must clarify how the efforts are likely to enable future contributions to the SBE perspective, preferably from a range of social, behavioral, and economic sciences. For research proposals that are infrastructure-oriented, those that contribute directly to research and go beyond merely providing a resource for other researchers are of special interest.

Transition-to-practice perspective

Research proposals with the transition-to-practice perspective should address the challenge of moving from research to capability. These proposals will typically leverage successful results from previous and current basic research and focus on later stage activities in the research and development life cycle (e.g., applied research, development, prototyping, testing, and experimental deployment). Strong preference will be given to projects whose outcomes result in fielded capabilities and innovations of direct benefit to networks, systems, and environments supporting NSF science and engineering research and education. Any software that is developed in this program area will be required to be released under an open source license listed by the Open Source Initiative [2]. Industry partnerships and collaborations are strongly encouraged.

Research proposals that are submitted with a transition-to-practice perspective will be evaluated with careful attention to:

- The expected impact on the deployed environment described in the proposal;
- The extent to which the value of the proposed cybersecurity research and development is described in the context of a needed capability required by science and engineering and potential impact across a broader segment of the NSF community;
- The feasibility, utility, and interoperability of the capability in its proposed operational role;
- A project plan that addresses in its goals and milestones the demonstration and evaluation of a working system in the target environment; and
- Tangible metrics described to evaluate the success of the capabilities developed and the steps necessary to take the system from prototype status to production use.

Cybersecurity education perspective

The results of SaTC funded research may lead to widespread changes in our understanding of the fundamentals of cybersecurity that can, in turn, lead to fundamentally new ways to motivate and educate students about cybersecurity. Proposals submitted with this perspective should leverage successful results from previous and current basic research in cybersecurity and research on student learning, both in terms of intellectual merit and broader impact, to address the challenge of expanding existing educational opportunities and resources in cybersecurity. This might include, but is not limited to, the following efforts:

- Defining a cybersecurity body of knowledge and establishing curricular recommendations for new courses (both traditional and online), degree programs, and educational pathways leading to wide adoption nationally;
- Evaluating the effects of these curricula on student learning;
- Encouraging the participation of a broad and diverse student population in cybersecurity education;
- Developing virtual laboratories to promote collaboration and resource sharing in cybersecurity education;
- Developing partnerships between centers of research in cybersecurity and institutions of higher education that lead to improved models for the integration of research experiences into cybersecurity degree programs; and
- Developing and evaluating the effectiveness of cybersecurity competitions, games, and other outreach and retention activities.

Additional information on NSF's SaTC program solicitation NSF 12-596 is available at http://www.nsf.gov/pubs/2012/nsf12596/nsf12596.htm.

About the authors

Nina Amla, Vijayalakshmi Atluri, Jeremy Epstein, Sol Greenspan, and **Samuel Weber** are program officers for the National Science Foundation (NSF)'s Directorate for Computer and Information Science and Engineering. The Directorate for Social, Behavioral, and Economic Sciences is represented by program officer **Peter Muhlberger** and the Directorate for Mathematical and Physical Sciences by **Andrew Pollington**. **Kevin Thompson** is a program officer in the NSF Office of Cyberinfrastructure, while **Victor P. Piotrowski** and **Zhi Tian** are program officers in the Directorate of Education and Human Resources and the Directorate of Engineering, respectively.

References

[1] Executive Office of the President National Science and Technology Council. "Trustworthy cyberspace: Strategic plan for the federal cybersecurity research and development program." 2011 Dec. Available at: http://www.whitehouse.gov/sites/default/files/microsites/ostp/fed_cybersecurity_rd_strategic_plan_2011.pdf

[2] To learn more about the Open Source Initiative, visit http://www.opensource.org/.

GLOBE AT A GLANCE

NSF programs in Secure and Trustworthy Cyberspace

Cybersecurity is arguably the most important challenge confronting society in the information age. Addressing this challenge requires a coordinated multidisciplinary approach, contributing to the body of knowledge on cybersecurity in the respective disciplines and leading to practical usable deployable technologies. The National Science Foundation's Secure and Trustworthy Cyberspace (SaTC) department is responding to this challenge by funding programs across the nation. This map shows the top 20 universities with the most, active SaTC programs as of December 2012. For more information about SaTC programs, see page 37.

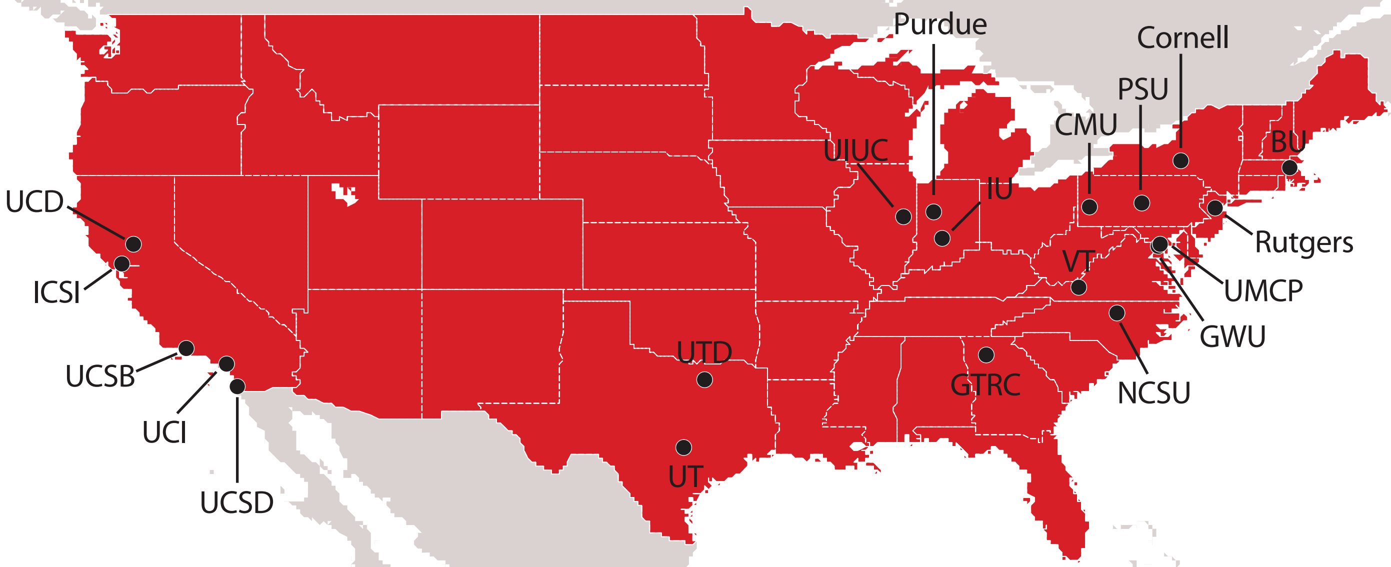

UNIVERSITIES WITH THE MOST NSF SaTC PROGRAMS

Abbreviation	University	No. of Programs
CMU	Carnegie Mellon University	11
UCSD	University of California, San Diego	11
Cornell	Cornell University	7
IU	Indiana University	7
PSU	Pennsylvania State University, University Park	7
UIUC	University of Illinois at Urbana-Champaigne	7
Purdue	Purdue University	6
UT	University of Texas at Austin	6
GTRC	Georgia Tech Research Corporation	5
ICSI	International Computer Science Institute	5
Rutgers	Rutgers University–New Brunswick	5
BU	Trustees of Boston University	5
UCD	University of California, Davis	5
UCI	University of California, Irvine	5
UCSB	University of California, Santa Barbara	5
UMCP	University of Maryland, College Park	5
VT	Virginia Polytechnic Institute and State University	5
GWU	George Washington University	4
NCSU	North Carolina State University	4
UTD	University of Texas at Dallas	4

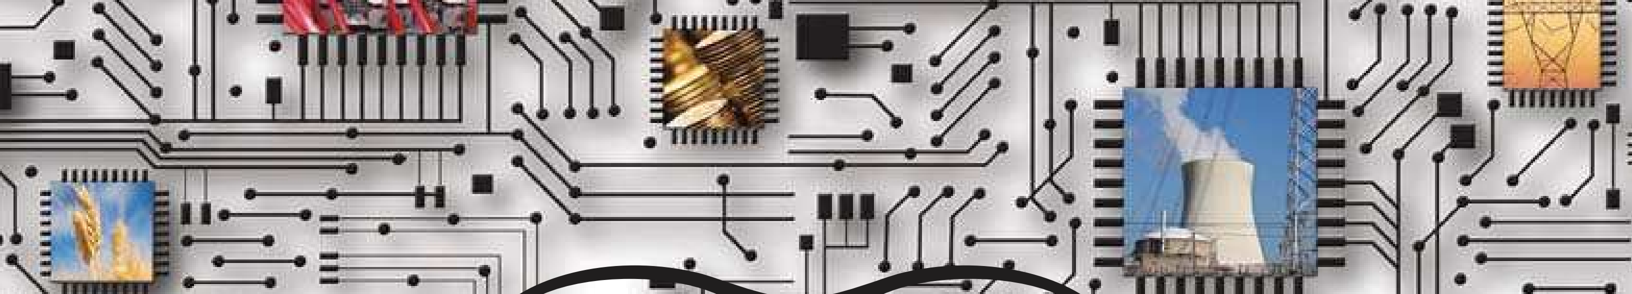

ACCORDING TO THE EXPERTS

Cyber threats to US infrastructure on the rise

The Department of Homeland Security (DHS) Control Systems Security Program manages and operates the Industrial Control Systems Cyber Emergency Response Team (ICS-CERT) to provide focused operational capabilities for defense of control system environments against emerging cyber threats. ICS-CERT responds to cyber threats that affect organizations that own and operate control systems associated with critical infrastructure and key resources including agriculture and food, banking and finance, chemical, commercial facilities, critical manufacturing, dams, defense industrial base, drinking water and water treatment systems, emergency services, energy, government facilities, information technology, national monuments and icons, nuclear reactors and materials and waste, postal and shipping, public health and healthcare, telecommunications, and transportation systems.

To accomplish this mission, ICS-CERT

- Responds to and analyzes control systems related incidents,
- Conducts vulnerability and malware analysis,
- Provides on-site support for incident response and forensic analysis,
- Provides situational awareness in the form of actionable intelligence,
- Coordinates the responsible disclosure of vulnerabilities/mitigations, and
- Shares and coordinates vulnerability information and threat analysis through information products and alerts.

Companies report cybersecurity incidents to ICS-CERT and request analysis support to help determine the extent of the compromise and gather information about cyber attacks, including the adversary's techniques and tactics. This information helps asset owners evaluate their security posture and take measures to strengthen their control systems and network security. Typical incident response support consists of analysis performed in ICS-CERT's Advanced Analytics Lab (AAL) on digital media, malware, log files, and other artifacts.

Figure 1 illustrates the number of incident report tickets and incident report on-site deployments between 2010 and 2011.

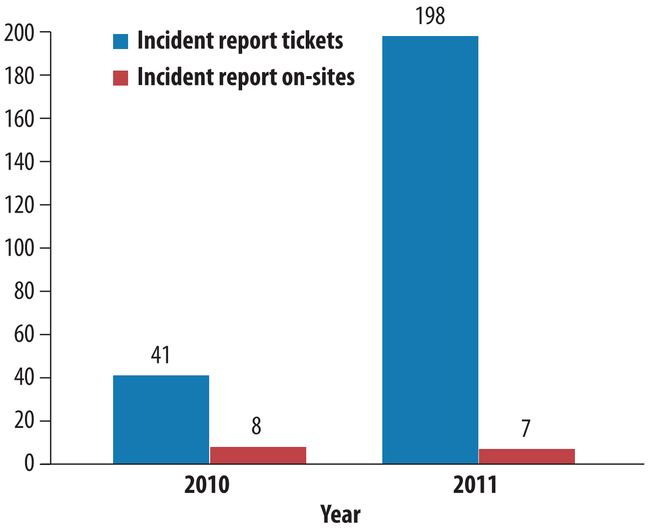

FIGURE 1. The number of cyber incident report tickets and on-site deployments for 2010 and 2011.

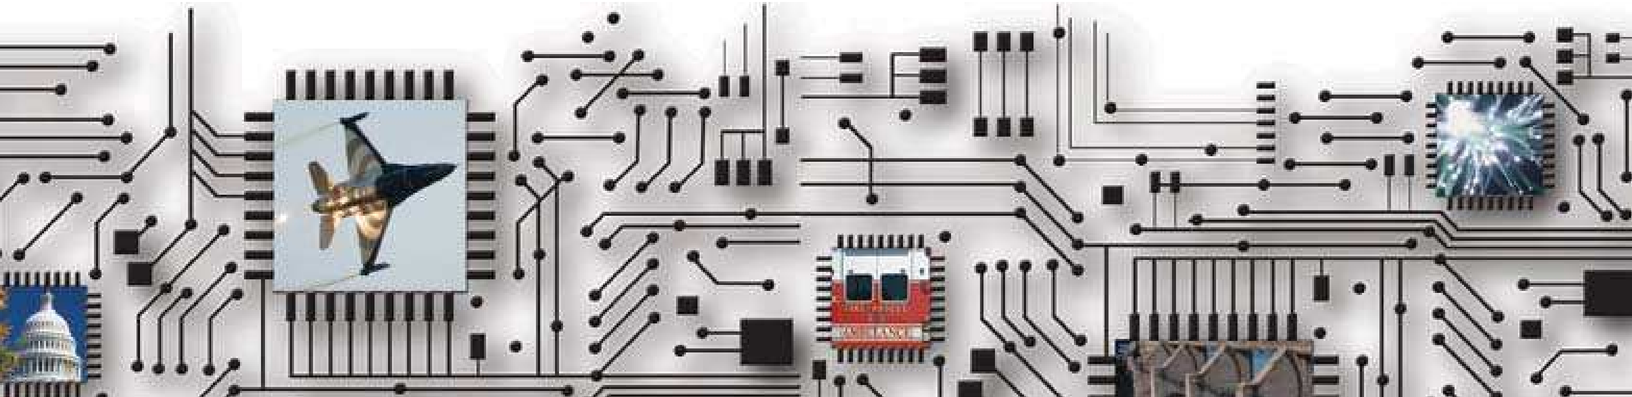

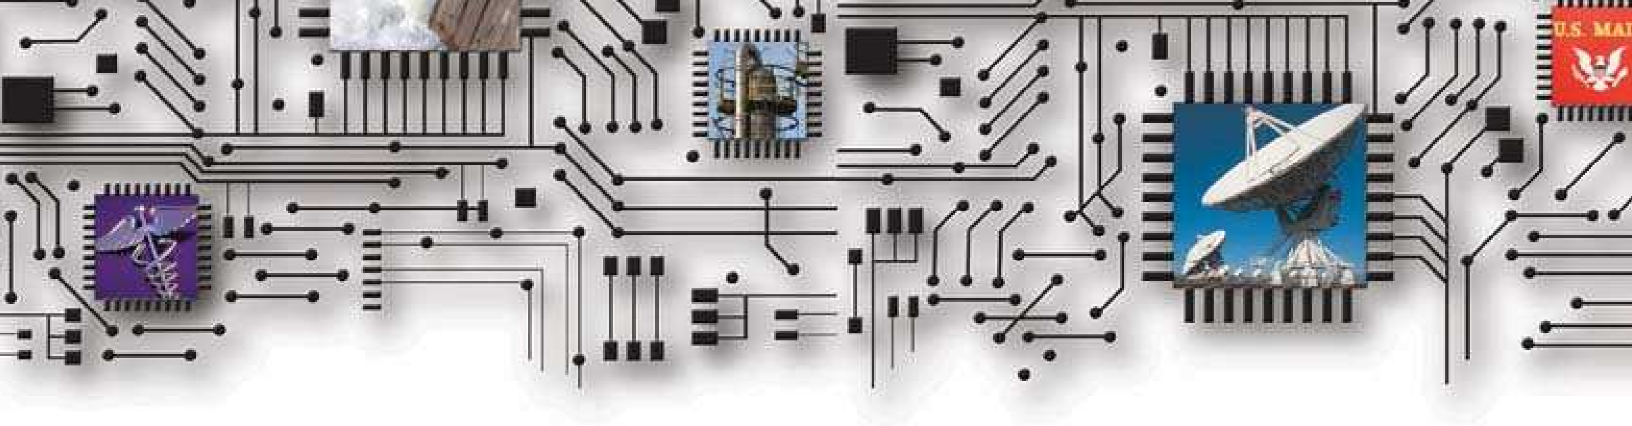

In 2010, 41 incident reports were received. Of the 41, eight resulted in the deployment of on-site response teams. An additional seven incidents involved remote analysis by the AAL. Figure 2 illustrates the breakout of incidents by sector.

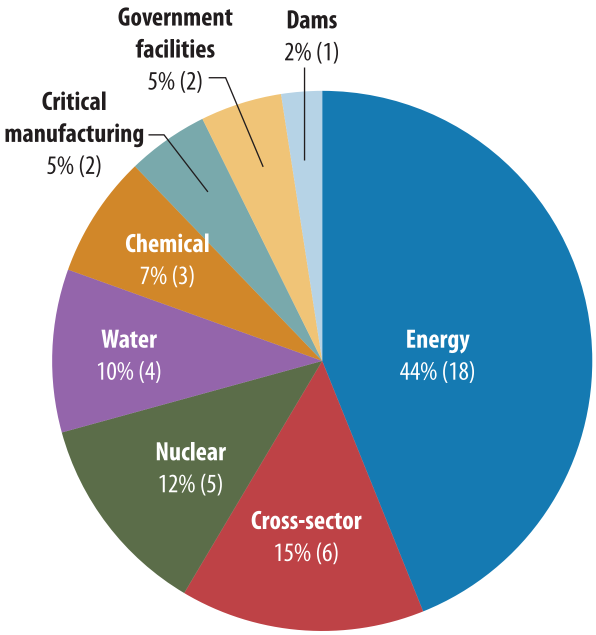

FIGURE 2. The number of cyber incident reports by sector in 2010.

In 2011, ICS-CERT received 198 reports of incidents. Of those 198, seven resulted in the deployment of on-site incident response teams. An additional 21 incidents involved analysis efforts by the AAL to identify malware and techniques used by the threat actors. Figure 3 displays the sector distribution for all incidents reported in 2011. Incidents specific to the water sector, when added to those that impacted multiple sectors, accounted for over half of the incidents due to a large number of Internet facing control system devices reported by independent researchers.

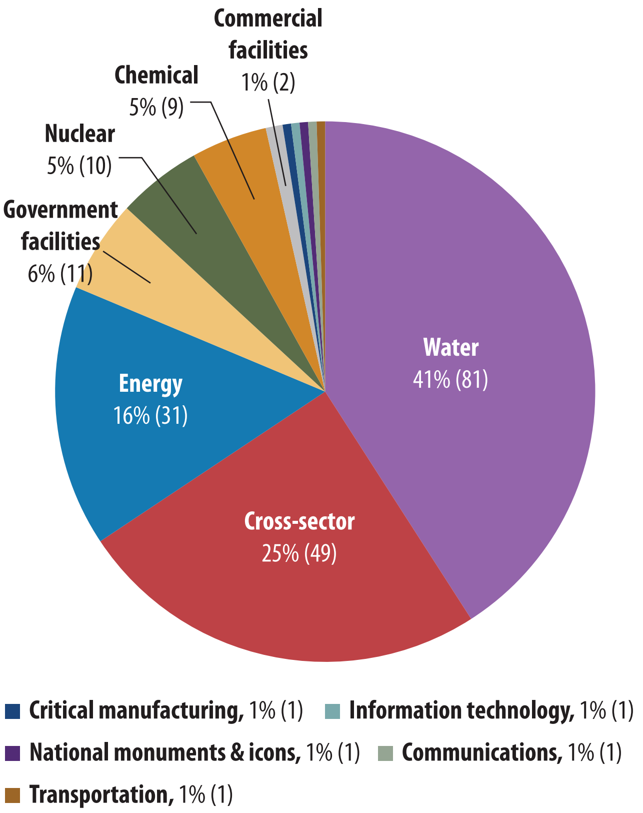

FIGURE 3. The number of cyber incident reports by sector in 2011.

For more information about ICS-CERT, or to report a cybersecurity incident, visit http://www.us-cert.gov/control_systems/ics-cert/.

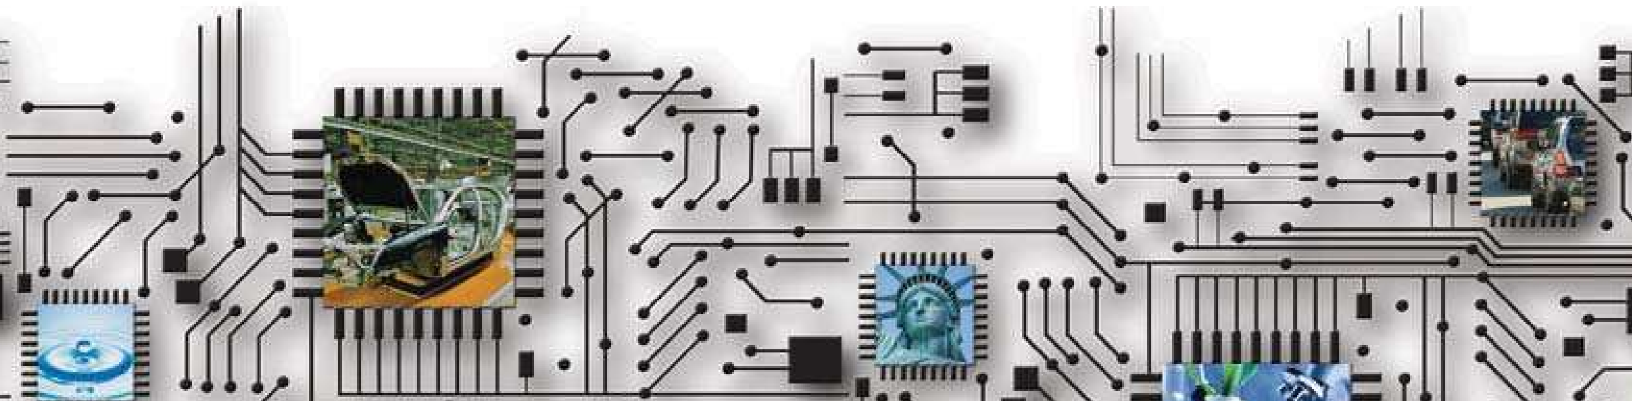

POINTERS

NSA sponsors science of cybersecurity lablets

NSA granted $2.5 million to Carnegie Mellon University, the University of Illinois at Urbana-Champaign, and North Carolina State University to fund research lablets devoted to developing a more scientific basis for the design and analysis of trusted cyber systems—a science of [cyber]security (SoS). NSA approved the schools' first research proposals for the lablets in December of 2011.

NSA's goal with these lablets is to create a unified body of knowledge in addition to analytics methods and tools that can serve as the basis of a trust engineering discipline, curriculum, and rigorous design methodologies. The results of SoS lablet research are to be extensively documented and widely distributed through the use of a new, network-based collaboration environment—the SoS virtual organization. The intention is for that environment to be the primary resource for learning about ongoing work in cybersecurity science and to be a place to participate with others in advancing the state of the art. (For more information about the SoS virtual organization, see page 20.)

The lablets' work will draw on several fundamental areas of computing research. Some ideas from fault-tolerant computing can be adapted to the context of security. Strategies from control theory will be extended to account for the high variation and uncertainty that may be present in systems when they are under attack. Game theory and decision theory principles will be used to explore the interplay between attack and defense. Formal methods will be applied to develop formal notions of resiliency. End-to-end system analysis will be employed to investigate resiliency of large systems against cyber attack. The lablets' work will draw upon ideas from other areas of mathematics and engineering as well.

Carnegie Mellon University SoS lablet

The broad goal of the Carnegie Mellon University (CMU) SoS lablet is to identify scientific principles that can lead to approaches to the development, evaluation, and evolution of secure systems at scale. The focus on scalability derives from a recognition that modern software-intensive systems have more components and a greater diversity of suppliers. The theme of scalability includes two principal areas of focus, which are composability and usability. Projects within the lablet may address diverse and possibly conflicting technical approaches in order to most effectively address the overall thematic goals.

Contributing technical areas include safe programming languages, binary and source code analysis, data-intensive systems analysis, self-healing and resilient architecture, assured API (application programming interface) and framework compliance, sociotechnical ecosystems, development environments, trusted computing, specification and verification, concurrent and distributed systems, requirements and policy, usable security and privacy, intrusion and malware detection, dynamic network analysis, model checking, secure coding practice, secure process separation, verification of cyber-physical systems, and others.

The lead principal investigator of the CMU SoS lablet is William Scherlis, professor in the School of Computer Science at CMU. He is the founding director of CMU's PhD Program in Software Engineering and director of CMU's Institute for Software Research in the School of Computer Science. His research relates to software assurance, software analysis, and assured safe concurrency.

The lablet's projects include:

- A language and framework for development of secure mobile applications,
- Architecture based self-securing systems,
- Improving the usability of security requirements by software developers through empirical studies and analysis,
- Learned resiliency: Secure multilevel systems,
- Secure composition of systems and policies,
- Security reasoning for distributed systems with uncertainties,
- Systematic testing of distributed and multithreaded systems at scale, and
- Validating productivity benefits of type-like behavioral specifications.

University of Illinois at Urbana-Champaign SoS lablet

The University of Illinois at Urbana-Champaign SoS lablet, which will be housed in the Information Trust Institute at Illinois, will leverage Illinois' expertise in resiliency, which in this context means a system's demonstrable ability to maintain security properties even during ongoing cyber attacks.

David M. Nicol, the lablet's principal investigator, explains, "The complexity of software systems guarantees that there will almost always be errors that can be exploited by attackers. We have a critical need for foundational design principles that anticipate penetrations, contain them, and limit their effects, even if the penetration isn't detected."

Nicol is a professor of electrical and computer engineering at Illinois and the director of the Information Trust Institute. The lablet's leadership is shared with coprincipal investigators William H. Sanders, who is an ECE professor and director of the Coordinated Science Laboratory at Illinois, and José Meseguer, a professor of computer science.

The lablet's projects include:

- Classification of cyber-physical system adversaries,
- End-to-end analysis of side channels,
- Enhancing cybersecurity through networks resilient to targeted attacks,
- From measurements to security science: Data-driven approach,
- Protocol verification: Beyond reachability properties,

- Quantitative assessment of access control in complex distributed systems,
- Quantitative security metrics for cyber-human systems,
- Scalable methods for security against distributed attacks,
- Secure platforms via stochastic computing,
- The science of summarizing systems: Generating security properties using data mining and formal analysis,
- Theoretical foundations of threat assessment by inverse optimal control,
- Toward a theory of resilience in systems: A game-theoretic approach,
- Towards a science of securing network forwarding, and
- Trust from explicit evidence: Integrating digital signatures and formal proofs.

North Carolina State University SoS lablet

The North Carolina State University (NC State) SoS lablet, which will be housed in the Institute for Next Generation IT Systems, will leverage NC State's expertise and experience in analytics, including the extensive expertise available in the NC State Institute of Advanced Analytics.

The coprincipal investigators for the NC State SoS lablet are Dr. Laurie Williams, professor of computer science, and Dr. Michael Rappa, director of the Institute of Advanced Analytics and professor of computer science.

"The security fortification technique of data encryption has a sound mathematical basis, providing a predictable and quantifiable level of security based upon the strength of the encryption algorithm," Williams says. "Conversely, the science behind other security techniques that provide vulnerability prevention, detection, and fortification is either rudimentary or does not exist. As a result, the principles of designing trustworthy systems often are not rooted in science. The three SoS lablets established by the NSA will research techniques to provide this scientific basis."

The lablet's projects include:

Full proposals

- An investigation of scientific principles involved in software security engineering,
- Attaining least privilege through automatic partitioning of hybrid programs,
- Argumentation as a basis for reasoning about security,
- Developing a user profile to predict phishing susceptibility and security technology acceptance,
- Empirical privacy and empirical utility of anonymized data,
- Improving the usability of security requirements by software developers through empirical studies and analysis,

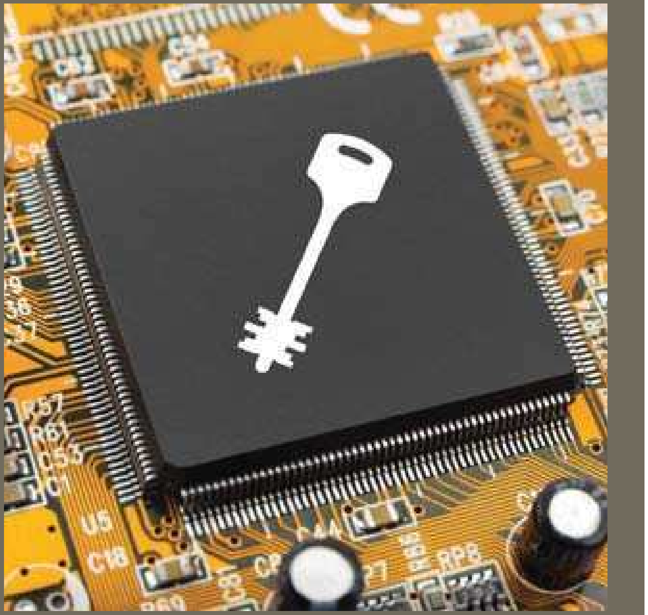

- Security metrics, and
- Towards a scientific basis for user-centric security design.

Seedlings

- A science of timing channels in modern cloud environments,
- An adoption theory of secure software development tools,
- Multitarget visualizations for visual analytics,
- Normative trust toward a principled basis for enabling trustworthy decision making,
- Quantifying underpinnings for network analytics as components of composable security,
- Quantifying mobile malware threats,
- Spatiotemporal security analytics and human cognition, and
- Studying latency and stability of closed-loop sensing-based security systems.

SPINOUTS

News from the Technology Transfer Program

Shared technology, shared defense: Spinning out the Vulnerability Tool Suit

One of NSA's critical missions is creating tools and techniques to provide information assurance and computer network defense for systems and networks throughout the US government. One such product is the Vulnerability Tool Suite (VTS).

The VTS is a collection of software and hardware computer network defense tools that has been developed to support the warfighter and critical national security communications systems. Typical components include methods to detect unauthorized hardware and software installations as well as tools to monitor system baseline configurations. NSA shares this toolset with military and civilian government organizations using a mechanism called a technology transfer sharing agreement (TTSA) administered by NSA's Technology Transfer Program (TTP).

Unlike patent license agreements, TTSAs are effectively no-cost licenses allowing other government agencies and partners to obtain proprietary NSA technology through interagency agreements. After entering into a TTSA with NSA, recipient agencies and partners are provided access to specific technologies, periodic updates and upgrades, and in some cases, training. All TTSAs contain standard legal references regarding intellectual property rights and each party's responsibilities. TTSAs typically are in place for three years.

In the case of the VTS, the TTP and the Information Assurance Directorate (IAD) are the primary interfaces between NSA and potential recipients. The IAD sends the VTS referrals to the TTP on a nearly daily basis and the IAD and TTP work collaboratively to execute the agreements. The TTP and IAD also showcase the VTS at various workshops and conferences throughout the year. The TTP and IAD meet periodically to update the VTS toolset contents and protection plan parameters.

As a result of the collaboration between the IAD and TTP, the VTS makes up almost 40% of all TTSAs executed by the Agency. Since mid-2007, NSA's TTP has executed 123 TTSAs for the VTS.

The VTS TTSA is just one example of how NSA is providing collaborative network assurance and cyber defense to all agencies of the US government.

FSC
logo

iN055686

National Security Agency

Central Security Service

Defending Our Nation. Securing The Future