COMPUTING edge

# Cybersecurity

Also in this issue:

> Code Evasion

> Days of Endless Time

OCTOBER 2015

www.computer.org

IEEE

IEEE computer society

# ENDLESS POSSIBILITIES ...

## AND COUNTLESS RISKS

**ENHANCED PROGRAM OPTIONS**

### Your technology career gives you both. Be prepared for them with the NEW and ENHANCED IEEE Member Professional Liability Insurance Program*.

From computing and sustainable energy systems, to aerospace, communications, robotics, cybersecurity, biomedical, and more, your career in the technology field offers you endless possibilities. But at the same time, it poses countless risks. No matter how well you design it or how accurate your advice is, you can still be sued. Whether the claim is founded or not, protecting your career, reputation and assets could be costly. That's why IEEE sponsors a Professional Liability Insurance Program. And now it's been ENHANCED to offer you more benefits, such as:

- NEW CHOICE PLATFORM gives you more coverage choices from leading IEEE-approved insurers
- Fill out one application to receive multiple quote options**
- Computer exposures and technology coverage
- Various deductible options
- Ideal protection for firms or self-employed individuals
- Exclusive member pricing

**Protect all your career possibilities from liability risk exposures today. Learn how this enhanced program can help you:**

**1-800-375-0775**
**IEEEINSURANCE.COM/NEWPL**

*The IEEE Member Professional Liability Insurance Program with the Choice Platform is available to active IEEE members who reside in the U.S. IEEE members in Canada (excluding Quebec) have access to the IEEE Member Professional Liability Insurance Plan through Marsh Canada Limited. Please visit www.ieeeinsurance.com/canadapl for more information. Marsh Canada Limited acts as the insurance broker with respect to residents of Canada.

**Coverage options may vary or may not be available in all states. Not all plan features will be available under all carriers or plan options. This program is administered by Mercer Consumer, a service of Mercer Health & Benefits Administration LLC.

COMPUTING edge

IEEE COMPUTER SOCIETY http://computer.org • +1 714 821 8380

Printed with inks containing soy and/or vegetable oils

IEEE

IEEE computer society

## STAFF



BPA WORLDWIDE

**Circulation:** ComputingEdge is published monthly by the IEEE Computer Society. IEEE Headquarters, Three Park Avenue, 17th Floor, New York, NY 10016-5997; IEEE Computer Society Publications Office, 10662 Los Vaqueros Circle, Los Alamitos, CA 90720; voice +1 714 821 8380; fax +1 714 821 4010; IEEE Computer Society Headquarters, 2001 L Street NW, Suite 700, Washington, DC 20036.

**Postmaster:** Send undelivered copies and address changes to IEEE Membership Processing Dept., 445 Hoes Lane, Piscataway, NJ 08855. Application to Mail at Periodicals Postage Prices is pending at New York, New York, and at additional mailing offices. Canadian GST #125634188. Canada Post Corporation (Canadian distribution) publications mail agreement number 40013885. Return undeliverable Canadian addresses to PO Box 122, Niagara Falls, ON L2E 6S8 Canada. Printed in USA.



## IEEE Computer Society Magazine Editors in Chief

***Computer***
Sumi Helal, *University of Florida*

***IEEE Software***
Diomidis Spinellis, *Athens University of Economics and Business*

***IEEE Internet Computing***
M. Brian Blake, *University of Miami*

***IT Professional***
San Murugesan, *BRITE Professional Services*

***IEEE Security & Privacy***
Shari Lawrence Pfleeger, *Dartmouth College*

***IEEE Micro***
Lieven Eeckhout, *Ghent University*

***IEEE Computer Graphics and Applications***
L. Miguel Encarnação, *ACT, Inc.*

***IEEE Pervasive Computing***
Maria Ebling, *IBM T.J. Watson Research Center*

***Computing in Science & Engineering***
George K. Thiruvathukal, *Loyola University Chicago*

***IEEE Intelligent Systems***
Daniel Zeng, *University of Arizona*

***IEEE MultiMedia***
Yong Rui, *Microsoft Research*

***IEEE Annals of the History of Computing***
Nathan Ensmenger, *Indiana University Bloomington*

***IEEE Cloud Computing***
Mazin Yousif, *T-Systems International*

OCTOBER 2015 • VOLUME 1, NUMBER 10

# COMPUTING edge

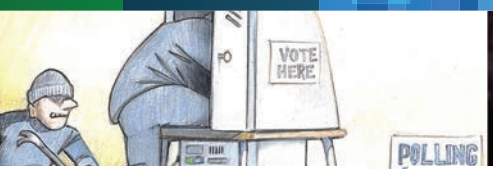


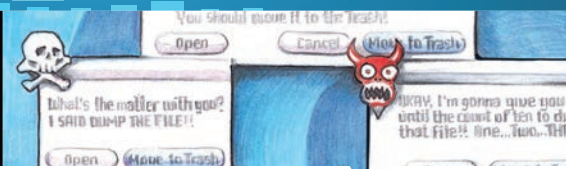

48

Murmurations: Drawing Together Art, Visualization, and Physical Phenomena

**Departments**



# An Efficient Algorithm for Exact Recovery of Vertex Variables from Edge Measurements

**Afonso S. Bandeira,** MIT

*This installment highlighting the work published in IEEE Computer Society journals comes from* IEEE Transactions on Network Science and Engineering.

What do community detection in complex networks, structure from motion in computer vision, and multiple-image registration have in common? They can each be formulated as an inverse problem on a graph. More specifically, we can associate each data unit—such as a node label, photo, or rotated image—to a graph node, and each pairwise measurement to an edge connecting the two nodes. In doing so, the problem becomes one of estimating for each node an unknown variable, such as community membership or rotation, from relative information between pairs of such variables. It's useful to think of node variables as taking values in a group $\mathcal{G}$ and relative measurements as revealing information about group ratios $g_i (g_j)^{-1} \in \mathcal{G}$.

Although much literature proposes algorithmic approaches to this problem, we still don't understand it from an information-theoretical or computational average-case-complexity viewpoint. In "Decoding Binary Node Labels from Censored Edge Measurements: Phase Transition and Efficient Recovery" (*IEEE Trans. Network Science and Eng.*, vol. 1, no. 1, 2014, pp. 10–22), my colleagues and I moved toward a better understanding by treating a simple, yet crucial, instance of this issue—the setting in which node labels take only two different values (that is, $\mathcal{G}$ is the group of two elements).[1]

In this setting, it's best to view the problem as community detection on an observation graph $G = (V, E)$. The unknown vertex labels $x^V$ represent community memberships and the edge measurements $y^E$ are noisy indications of whether node pairs belong to the same cluster. More precisely, given a noise level $\varepsilon < 1/2$, for each edge $(i, j)$, $y^E_{i,j}$ indicates whether $i$ and $j$ belong to the same community or not, making an error with probability $\in$ for each edge, and independently to errors in other edges. The goal is then to determine for which graphs $G$ and values of $\varepsilon$ we can recover $x^V$ given $y^E$, and whether this inverse problem can be solved efficiently. Note that $x^V$ can only be recovered up to a global flip, corresponding to a cluster relabeling. Indeed, because only relative information is available, the measurements are invariant to relabeling the clusters.

This model is closely related to the popular stochastic block model for two communities. In that model, a random graph is drawn on vertices belonging to two or more communities from a distribution with independent edges and probability $p$ if between two vertices of the same community and $q$ if otherwise. The objective is to recover the community memberships. The main difference between the models is that in the stochastic block model, every node pair provides information (the nonexistence of an edge is itself information), whereas in our model node pairs not connected in $G$ don't provide information.

Despite the models' differences, we've successfully adapted the techniques discussed in our paper[1] to a stochastic block model setting.[2]

Our paper's main contribution was demonstrating that, if we consider the observation graph an Erdös-Rényi graph (a random graph in which each node is connected by an edge, independently and with probability $p$), with average degree $(n–1)p$, exact recovery of $x^V$ is possible with high probability if and only if

$$\alpha = np / \log(n) > 2 / (1 - 2\varepsilon)^2 + o(2 / (1 - 2\varepsilon)^2).$$

If $\alpha < 1$, then we know with high probability that the observation graph contains isolated nodes. This renders recovery impossible even in a noiseless

 Published by the IEEE Computer Society 

setting, because there would be nodes with no available information.

On the algorithmic side, we propose the use of a semidefinite programming relaxation-based algorithm. Duality and estimates on spectral norms of certain random matrices show that this efficient algorithmic approach exactly recovers $x^V$ with high probability at regimes very close to the information-theoretical limit.[3] Remarkably, with the use of sharper estimates on the spectrum of random matrices, these results have since improved.[4,5]

The techniques in our paper have since been adapted to community detection with more than two communities.[5–8] Furthermore, although our work focused on exactly recovering $x^V$, more recent studies have solved one of the open problems we posed in our paper: how to partially recover $x^V$.[9,10] C

## REFERENCES


1. E. Abbe et al., "Decoding Binary Node Labels from Censored Edge Measurements: Phase Transition and Efficient Recovery," *IEEE Trans. Network Science and Eng.*, vol. 1, no. 1, 2014, pp. 10–22.
2. E. Abbe, A.S. Bandeira, and G. Hall, "Exact Recovery in the Stochastic Block Model," 2014; http://arxiv.org/abs/1405.3267.
3. J.A. Tropp, "User-Friendly Tail Bounds for Sums of Random Matrices," *Foundations of Computational Mathematics*, vol. 12, no. 4, 2012, pp. 389–434.
4. A.S. Bandeira, "Random Laplacian Matrices and Convex Relaxations," 2015; http://arxiv.org/abs/1504.03987.
5. B. Hajek, Y. Wu, and J. Xu, "Achieving Exact Cluster Recovery Threshold via Semidefinite Programming: Extensions," 2015; http://arxiv.org/abs/1502.07738.
6. N. Agarwal et al., "Multisection in the Stochastic Block Model Using Semidefinite Programming," 2015; http://arxiv.org/abs/1507.02323.
7. W. Perry and A.S. Wein, "A Semidefinite Program for Unbalanced Multisection in the Stochastic Block Model," 2015; http://arxiv.org/abs/1507.05605.
8. E. Abbe and C. Sandon, "Community Detection in General Stochastic Block Models: Fundamental Limits and Efficient Recovery Algorithms," 2015; http://arxiv.org/abs/1503.00609.
9. P. Chin, A. Rao, and V. Vu, "Stochastic Block Model and Community Detection in the Sparse Graphs: A Spectral Algorithm with Optimal Rate of Recovery," 2015; http://arxiv.org/abs/1501.05021.
10. A. Saade et al., "Spectral Detection in the Censored Block Model," 2015; http://arxiv.org/abs/1502.00163.



**AFONSO S. BANDEIRA** is an instructor and postdoctoral fellow in MIT's Department of Mathematics. Contact him at bandeira@mit.edu.

# Magazine Roundup

The IEEE Computer Society's lineup of 13 peer-reviewed technical magazines covers cutting-edge topics ranging from software design and computer graphics to Internet computing and security, from scientific applications and machine intelligence to cloud migration and microchip manufacturing. Here are highlights from recent issues.

## *Computer*

The Internet of Things entails connecting vast numbers of devices to the Internet, in essence turning them into smart networked devices. *Computer*'s September 2015 special issue focuses on **activating the Internet of Things**.

## *IEEE Software*

Properly designing, creating, and documenting a **software architecture** are critical to an application's success. *IEEE Software*'s September/October 2015 special issue explores this topic from different perspectives.

## *IEEE Internet Computing*

Before we can trust the **Internet of Things' infrastructure**, we must feel that our personal data is protected. The authors of "Governance Challenges for the Internet of Things," from *IEEE Internet Computing*'s July/August 2015 issue, discuss how legal data-protection frameworks must change for this to occur.

## *Computing in Science & Engineering*

**Open simulation laboratories** (OSLs)—the focus of *CiSE*'s September/October 2015 special issue—will become more common as their potential is better understood and as they begin providing larger segments of the scientific community with access to valuable datasets. New analysis tools and new ways to perform scientific research will develop as a result.

## *IEEE Security & Privacy*

The Internet of Things (IoT) is growing, which means household appliances are increasingly facing security threats that come with Internet connectivity. Focusing on one particularly security-sensitive product, "Garage Door Openers: An Internet of Things Case Study," from *IEEE S&P*'s July/August 2015 issue, investigates **how IoT devices might be designed with security in mind**.


Published by the IEEE Computer Society

*IEEE Cloud Computing*

"**Socioeconomics of Cloud Standards**," from *IEEE Cloud Computing*'s May/June 2015 issue, discusses factors that could lead to the emergence of standards. The article explores features that could limit their initial emergence and those that could accelerate their subsequent uptake. The main limiting factor so far has been the technology's rapidly changing nature, which doesn't allow for the technological simplification necessary for the development and adoption of standards.

*IEEE Computer Graphics and Applications*

**DreamWorks Animation's Premo** offers a state-of-the-art animator experience with fully deforming characters, complex environments, and real-time lighting that's indicative of final rendered images. Premo unlocks novel workflows, is efficient, and offers a natural interface for 3D animation, in contrast to many other professional animation tools' highly technical interfaces. In *CG&A*'s July/August 2015 issue, the author of "Premo: DreamWorks Animation's New Approach to Animation" analyzes the system.

*IEEE Intelligent Systems*

In *IEEE Intelligent Systems*' July/August 2015 issue, the authors of "**Saving Rhinos with Predictive Analytics**" look at the problem of animal poaching. They describe their Anti-Poaching Engine, which builds on behavior models of rhinos and poachers to protect as many animals as possible.

*IEEE MultiMedia*

Social media's proliferation has generated a huge amount of multimedia data. Because this information is unstructured and multimodal, **multimedia big data computing** has not only created unprecedented opportunities but has also created fundamental challenges in storage, processing, and analysis. In *IEEE MultiMedia*'s July–September 2015 issue, the authors of "Multimedia Big Data Computing" examine these matters, present various methodologies, and speculate on research opportunities.

*IEEE Annals of the History of Computing*

Before World War II, electric utilities adopted analog computing for various power-system calculations and controls. After the war, many made a slow transition to digital computing, as described in "**Transitions from Analog to Digital Computing in Electric Power Systems**" from *IEEE Annals*' July–September 2015 issue.

*IEEE Pervasive Computing*

Using technologies such as HTML and JavaScript to build Web apps promises to create portable applications for all platforms. Browsers would thus become application containers, requiring increased performance from browser engines. In *IEEE Pervasive Computing*'s July–September 2015 issue, the authors of "Concurrency in Mobile Browser Engines" discuss **advances in browser technologies** that exploit multicore processing and concurrency.

*IT Professional*

"**Cognitive Computing, Analytics, and Personalization**," from *IT Pro*'s July/August 2015 issue, addresses the increasing sophistication of search technology and its role in cognitive computing, as well as how machine learning can contribute to search personalization. For example, the article looks at how search engines predict which information users are likely to need and how the software retrieves and presents the data.

*IEEE Micro*

Initial general-purpose microprocessor designs were homogeneous, which had advantages in design, testing, and programmability. However, with the advent of multicore processors, **heterogeneous computing**—in general-purpose and specialized cores—became available to provide more diverse capabilities, as explored in *IEEE Micro*'s July/August 2015 special issue.

*Computing Now*

The Computing Now website (http://computingnow.computer.org) features **up-to-the-minute computing news** and blogs, along with articles ranging from peer-reviewed research to opinion pieces by industry leaders.

# The Challenge of Cybersecurity

Cybersecurity is frequently referred to as an "arms race." Security experts devise ways to stop and prevent known attacks, and hackers create new techniques and look for new penetration vectors. This cycle has ramifications for many aspects of computer technology, including hardware, software, and networks.

This month's *ComputingEdge* explores a variety of cybersecurity-related issues. For example, in *IEEE Internet Computing*'s "The Research Horizon: Four Nearly Practical Concepts," Hilarie Orman examines four promising approaches that could be among the next big things in security.

Researchers have been trying to develop secure electronic-voting systems for years. "End-to-End Verifiability in Voting Systems, from Theory to Practice," from *IEEE Security & Privacy*, suggests a new paradigm.

In recent years, the US government has introduced several security-related policy measures, but their implementation could create many concerns, according to *Computer*'s "Recent US Cybersecurity Policy Initiatives: Challenges and Implications."

In *IEEE Security & Privacy*'s "Scaring and Bullying People into Security Won't Work," Angela Sasse argues that users want reliable and credible risk information, as well as better threat detection and security tools, rather than fear mongering.

*IEEE Internet Computing*'s "Ancillary Impacts of Multipath TCP on Current and Future Network Security" looks closely at a potentially important security issue.

"Silver Bullet Talks with Katie Moussouris," from *IEEE Security & Privacy*, highlights an interview with security vendor HackerOne's chief policy officer about bug bounties and vulnerability disclosures.

*ComputingEdge* articles on other subjects include:

- *IEEE Software*'s "Code Evasion" says the challenge in writing reliable code is finding ways to remove code from an application by simplifying and generalizing, rather than continuing the trend of adding more.
- In *IEEE Computer Graphics and Application*'s "Murmurations: Drawing Together Art, Visualization, and Physical Phenomena," the authors interview Rhode Island School of Design professor Dennis Hylnsky about his work and how it could help bridge the gap between art and science practice.
- *Computing in Science & Engineering*'s "Days of Endless Time" by columnist Charles Day discusses how video installations are catching up to video games in the use of computer-generated imagery.


Published by the IEEE Computer Society 2376-113X/15/$31.00 © 2015 IEEE

Editors: Mohamed Kaâniche, mohamed.kaaniche@laas.fr | Aad Van Moorsel, aad.vanmoorsel@ncl.ac.uk

# End-to-End Verifiability in Voting Systems, from Theory to Practice

**Peter Y.A. Ryan** | University of Luxembourg
**Steve Schneider** | University of Surrey
**Vanessa Teague** | University of Melbourne

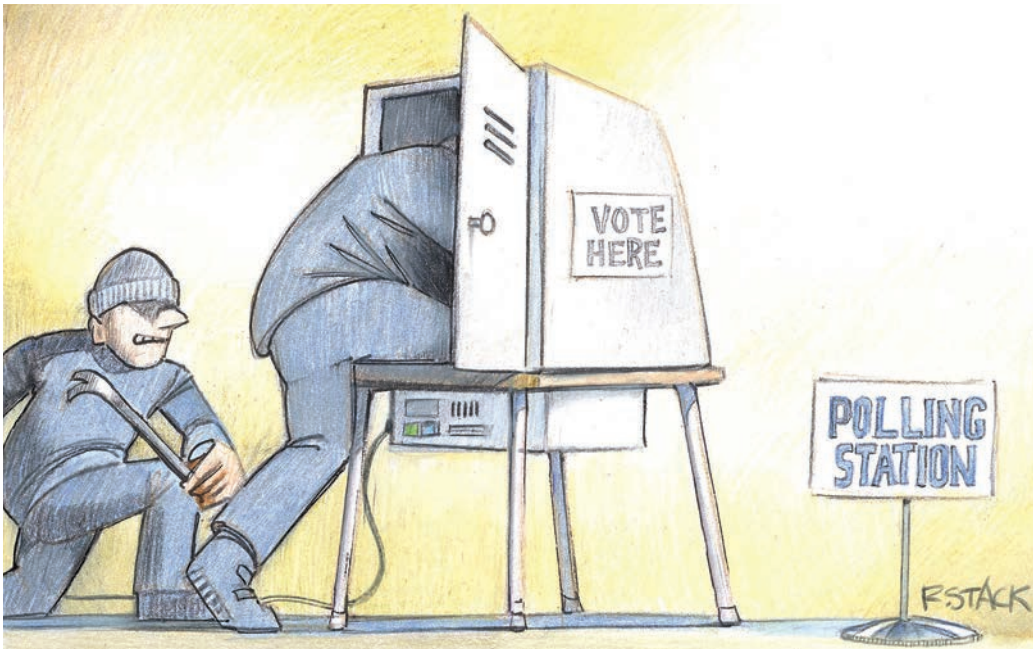


Since the dawn of democracy, societies have been experimenting with technological means to tackle corruption and avoid the need to trust officials. Excavations of Ancient Greece have revealed mechanisms that were clearly designed to ensure allotment: the randomness of the selection of people for office. In response to a rash of corrupted elections in the US in the late 19th century, countless devices were created that promised to provide incorruptible vote recording and counting. Thomas Edison even patented an electronic vote-recording device, and monstrous Metropolis-style lever machines persisted in some US states until very recently.

Throughout the history of democracy, there's been a battle between those trying to ensure the integrity of elections and those seeking to undermine them. The human ingenuity that has been poured into this war is truly impressive; see Andrew Gumbel's *Steal this Vote: Dirty Elections and the Rotten History of Democracy in America* for a highly entertaining—and somewhat terrifying—account.[1] The combat continues unabated, but now with new technology available to both sides. Cryptographers and those in information security have attempted to address the problem since the turn of the 21st century. Modern cryptography opens up a realm of new possibilities, but like all technology, cryptography and digital innovations are double-edged swords, opening up new threats.

Some argue that voting is a human activity that should remain in the traditional, even ceremonial realm: casting paper votes into ballot boxes and counting the resulting pile of ballots by hand. Others worry that any move to digital voting technology will enable systematic corruption. This position does hold some merit: it's true that any hasty, ill-thought-out innovation could result in disaster. Indeed, this has been demonstrated many times, such as with the California Top-to-Bottom Review of voting (https://www.sos.ca.gov/voting-systems/oversight/top-to-bottom-review.htm), where the team analyzing commercial voting systems in California declared that "virtually every important software security mechanism is vulnerable to circumvention." It's clear, then, that innovations must be developed with extreme care. But the argument that moving away from the traditional voting system will be disastrous is misguided.

## End-to-End Verifiability

The promise of *end-to-end verifiability* (E2EV) gives us hope that digital technologies can provide benefits in terms of security, and not just in terms of convenience and usability. E2EV uses some of the novel

2376-113X/15/$31.00 © 2015 IEEE Published by the IEEE Computer Society 

properties of modern cryptography to offer something completely new and quite remarkable: the means for voters to confirm that their vote is accurately included in the tally while preventing any third party from determining how they voted, even with their cooperation. In essence, voters can privately create an encryption of their vote. All encrypted votes are posted to a public website, where voters can confirm that their vote is correctly recorded. The batch of encrypted votes is anonymized and decrypted in a universally verifiable fashion and can then be tabulated.

The fundamental challenge in public voting is how to reconcile the conflict between demonstrable integrity and ballot privacy. The E2EV solution is the classic computer science way of introducing an indirection: the encryption and decryption of votes. A short, gentle introduction to E2EV can be found at http://arxiv.org/abs/1504.03778.

Although E2EV sounds simple, it's really quite complex. The implementation of E2EV has to be sufficiently simple and usable for voters, election officials, and candidates to feel comfortable. A particularly delicate step is encrypting the ballot in such a way so that voters are confident that their vote has been correctly encoded without involving a third party. The most common approach to achieving ballot assurance is the *Benaloh challenge*: voters tell the device how they wish to vote, and this commits to an encryption. Voters can now challenge this—requiring that the encryption be opened—or cast their ballot. Voters are free to repeat this as many times as they wish until they feel confident that the device is behaving correctly. Of course, it's essential that the device not know in advance how voters will choose.

In recent years, we've seen such systems start to move from academic articles into the real world. In 2009, the Scantegrity II system, which uses the E2EV approach, was successfully used in municipal elections in Takoma Park, Maryland.[2]

## vVote

Last November in Victoria, Australia, a system called vVote, based on the Prêt à Voter approach,[3] was successfully used by a section of the electorate. The system allowed for E2EV electronic voting in supervised polling places—the first time this was done in a politically binding statewide election—for voters with disabilities, such as vision impairment, and for Australian citizens voting remotely from London, England. Votes were cast privately in a voting booth and then transferred electronically to a central count. Because the electronic system ran in parallel with the traditional paper voting system, the final step in which the electronic votes were merged with the physical ones could be observed only by poll watchers who were present. Apart from that, all other steps could be verified by voters.

The key idea behind the Prêt à Voter approach, which vVote inherits, is to encode votes using a randomized candidate list, which ensures the secrecy of each vote and removes any bias. Once a ballot is marked by a voter, the candidate list is detached and destroyed. An encryption of the candidate order is preserved and used to extract the vote during tabulation.

This gives voters four steps of verification:

1. Before casting a vote, voters can confirm that the printed ballot with the randomized candidate list is properly constructed. When given a ballot, voters can choose to challenge it by demanding cryptographic proof of its correctness, which they can take home and verify. Voters can challenge as many ballots as they like before accepting one.
2. When the voting computer prints out their marked ballot, voters can check that the marks align properly with the randomized candidate list.
3. Once the candidate list is destroyed, voters leave the polling place with a receipt that includes their printed ballot and the encrypted candidate order. Voters can see that their ballot appears on a public list of accepted votes without revealing how they voted.
4. Anyone can verify that all the votes on the public list are properly shuffled and decrypted.

All of these steps—aside from the second—can be performed by or with the help of proxies of the voters' choice. Every aspect of the system is available for scrutiny: every check that voters perform with a computer can be independently recompiled, reimplemented, or performed by a completely independent party.

The source code for vVote is available at https://bitbucket.org/vvote. A nontechnical guide is available at http://electionwatch.edu.au/victoria-2014/click-here-democracy-e-vote-explained, and the complete system description and security analysis can be found in Chris Culnane and his colleagues' "vVote: A Verifiable Voting System."[4]

The vVote system was designed to handle up to hundreds of thousands of votes, though for this particular election, access to the system in the State of Victoria was restricted to 24 early voting centers and to voters with disabilities. In addition, voters in London, England, were able to use the system to cast their vote in a supervised polling place at the Australian High Commission. For these groups, 1,121 votes were cast using the system, more than the number of remote electronic votes cast in

2010, and with a quarter of the number of polling places available. A survey of the voters in London found that more than 75 percent agreed or strongly agreed with the statement that the system was easy to use.

## Issues and Challenges

Although voter feedback seems to be fairly positive, there are some issues regarding existing E2EV techniques. The very concept of being able to verify a vote rather than blindly trusting a system is novel for voters and requires an effort by the authorities to educate and motivate the electorate. Usability remains a challenge for E2EV systems, as discussed in Fatih Karayumak and his colleague's "User Study of the Improved Helios Voting System Interfaces."[5] Verification needs to be simple enough so voters can understand its purpose and feel motivated to perform the checks in significant numbers. It's not sufficient for voters to simply follow the system's instructions—without performing any checks—as attackers could manipulate the code issuing the instructions.

Another challenge is that a system can't simply be verifiable—it's essential that the system is actually verified randomly many times to ensure confidence in the result. In the case of the November 2014 election in Victoria, observation of the remote voters in London suggested that the majority did perform some check of the printed receipt against the candidate list, and around 13 percent of those using vVote checked receipts on the public website.[6]

There are a number of alternative commercial systems that claim to be verifiable but don't actually allow voters to perform their own checks. Of course, this can result in a more appealing "vote and go" user interface. With the iVote system, used in the 2015 state elections in Victoria's neighboring state of New South Wales, only a small number of chosen auditors could verify the system's output. Voters can check their own votes only by querying a database, instead of seeing the evidence themselves and checking it with their own machine as they can with E2EV voting.

One of the authors of this article co-discovered a serious security vulnerability in the 2015 New South Wales election. It was easily patched, but only after 66,000 votes had been cast.[7] Given that iVote's "verification" mechanism is unavailable for external review, there's a risk that it contains errors or security holes. This is important because trust in a small number of computers represents a potential avenue for undetectable, large-scale electoral manipulation if attackers can compromise that small set.

**The fundamental challenge in public voting is how to reconcile the conflict between demonstrable integrity and ballot privacy.**

## System Verification versus E2EV

It's important to note that the philosophy behind E2EV systems is quite different from what's usually meant by "system verification." In the latter, the idea is to perform a detailed analysis of a system's design and implementation against a set of required properties. Thus, as long as the verified code is running at execution time and the verification is complete and correct, the system should uphold the required properties. In practice, it's extremely difficult to achieve all this, especially due to the rather open, distributed nature of voting systems.

By contrast, E2EV seeks to ensure that the system execution is fully auditable. This idea is nicely captured in Josh Benaloh's maxim: "Verify the election, not the system." A related concept is Ronald L. Rivest and John P. Wack's notion of "software independence," which says that any error in the code that could result in a change in the outcome must be detectable at execution time (http://people.csail.mit.edu/rivest/RivestWack-OnTheNotionOfSoftwareIndependenceInVotingSystems.pdf). Of course, this doesn't mean that verification of the design and code should be neglected—it just means that the integrity of the outcome should not be dependent on assumptions about the correctness of the running code.

Another project is the End-to-End Verifiable Internet Voting Project (www.overseasvotefoundation.org/E2E-Verifiable-Internet-Voting-Project/News), which is examining E2EV in an attempt to define the real requirements of verifiability, so vendor systems that are not truly E2EV—but claim to be—can be differentiated from systems that are.

End-to-end verifiability represents a paradigm shift in electronic voting, providing a way to verify the integrity of elections by allowing voters to audit the information published by the system, rather than trusting that the system has behaved correctly. Recent deployments of E2EV systems in real elections demonstrate its practical applicability, and we hope to one day see E2EV as the normal expectation for electronic voting systems. ■

### References


1. A. Gumbel, *Steal this Vote: Dirty Elections and the Rotten History of Democracy in America*, Nation Books, 2005.

2. R. Carback et al., "Scantegrity II Municipal Election at Takoma Park: The First E2E Binding Governmental Election with Ballot Privacy," *Proc. 19th USENIX Conf. Security* (USENIX Security 10), 2010, p. 19.
3. P.Y.A. Ryan et al., "Prêt à Voter: A Voter-Verifiable Voting System," *IEEE Trans. Information Forensics and Security*, vol. 4, no. 4, 2009, pp. 662–673.
4. C. Culnane et al., "vVote: A Verifiable Voting System," arXiv:1404.6822, 2014; http://arxiv.org/abs/1404.6822.
5. F. Karayumak et al., "User Study of the Improved Helios Voting System Interfaces," *1st Workshop Socio-Technical Aspects in Security and Trust* (STAST 11), 2011, pp. 37–44.
6. C. Burton, C. Culnane, and S. Schneider, "Secure and Verifiable Electronic Voting in Practice: The Use of vVote in the Victorian State Election," arXiv:1504.07098, 2015; http://arxiv.org/abs/1504.07098.
7. V. Teague and A. Halderman, "Security Flaw in New South Wales Puts Thousands of Online Votes at Risk," *Freedom to Tinker*, 22 Mar. 2015, https://freedom-to-tinker.com/blog/teaguehalderman/ivote-vulnerability.

**Peter Y.A. Ryan** is a professor of applied security at the University of Luxembourg. Contact him at peter.ryan@uni.lu.

**Steve Schneider** is a professor of computing and Director of the Surrey Centre for Cyber Security at the University of Surrey. Contact him at s.schneider@surrey.ac.uk.

**Vanessa Teague** is a research fellow in the Department of Computing and Information Systems at the University of Melbourne. Contact her at vjteague@unimelb.edu.au.

*This article originally appeared in* IEEE Security & Privacy, *vol. 13, no. 3, 2015.*

Editor: Gary McGraw, gem@cigital.com

# Silver Bullet Talks with Katie Moussouris

**Gary McGraw** | Cigital

Hear the full podcast at www.computer.org/silverbullet. Show links, notes, and an online discussion can be found at www.cigital.com/silverbullet.


Katie Moussouris is the chief policy officer for HackerOne, a San Francisco–based startup that developed a vulnerability management and bug bounty platform. Moussouris oversees the company's philosophy and approach to vulnerability disclosure. For seven years before that, she worked at Microsoft, ending up as a senior security strategist. She also headed the company's bug bounty program, served as content chair for its BlueHat internal hacker conference, and helped with its secure development life cycle process.

**You started on a Commodore 64 programming in BASIC. What's the coolest program you wrote as a kid?**
It was probably my first one. I was really into Choose Your Own Adventure books, and I wrote a text-based adventure game. That was really fun for me.

**Did you use and adapt some code from other people?**
I think I took some examples from a book on case statements. I definitely borrowed some code. I'm sure it had bugs in it.

**Then you decided to study biology and made a quick switch back to the computer field.**
I wanted to cure cancer and AIDS. I got to work on the Human Genome Project, which was great. But I found myself going back to the computing world. I made a couple different career shifts, working as a systems administrator for a while. I learned the pain of having to defend your own network from attackers and polished my penetration-testing skills doing that when I worked at MIT.

I became a Linux developer when I first moved to California. I was mostly working with the different systems administration tools, fixing a lot of bugs. I realized we didn't have a formal computer security response program there, so I started one. That's where I began my career in security response and vulnerability handling and coordination.

**Tell us in very basic terms how a bug bounty program works.**
A bug bounty program is a cash reward offered by a vendor in exchange for vulnerability information. It's typically offered on a per-bug basis: one bug, one bounty. But a lot of the incentive programs I worked on, for example, at Microsoft, were looking at broader things like learning defensive techniques and new attack or mitigation bypass techniques. That's essentially a new way to exploit and bypass all the safety measures on a platform's latest version. These were more structured incentive programs that looked for a higher-order outcome than you look for with basic bug bounties.

**Is this what you do at HackerOne, or is it more bug bounty related?**
HackerOne is a platform and toolset that enables response teams to receive vulnerability reports. You can do that on our platform with or without a bug bounty. We're really about vulnerability coordination. The platform is free if you want to use it for vulnerability coordination. We charge a fee only if there's a bounty involved. We also advise our

2376-113X/15/$31.00 © 2015 IEEE Published by the IEEE Computer Society

## About Katie Moussouris

Katie Moussouris is the chief policy officer for HackerOne, a platform provider for coordinated vulnerability response and structured bounty programs. She oversees the company's philosophy on vulnerability disclosure, advises customers and researchers, and works to legitimize and promote security research to help make the Internet safer for everyone. After working at MIT, Moussouris joined Symantec and later Microsoft, where her work encompassed industry-leading initiatives such as Microsoft's bounty programs and Microsoft Vulnerability Research. She is also a subject matter expert for the US National Body of the International Organization for Standardization (ISO) on vulnerability disclosure and handling processes, secure development, and penetration testing.

customers on the best way to create an incentive program if that's what they're ready for. You really have to do a lot more security homework for a bounty.

**Bug bounties are something we saw emerge in the BSIMM [Building Security In Maturity Model] community (http://bsimm.com). It makes sense to have a bug bounty or vulnerability coordination system as a software vendor, but it's iffier when it comes to financial services and healthcare firms. What do you recommend for those sorts of firms?**
Financial services and healthcare firms—really any organizations or verticals at high risk of cyberattack—already hire penetration-testing companies. So they're very comfortable with the BSIMM, including the idea of bug bounties. Instead of looking at it as, "I'm only comfortable hearing about vulnerabilities from a specific small set of individuals that I hire," we want to open up that mindset to, "If anybody could tell me about a vulnerability in my product, I'm interested in hearing about it. And I'd rather hear about it first than wait for an attack to happen."

**I was going to ask whether it's better to use a HackerOne system or hire an internal penetration test team. But I know you're going to say both.**
It's absolutely both. There's a time and a place to do different kinds of testing as you're developing the software. You and I are both big advocates of secure software development, and there are appropriate times and methodologies and tools to use when you're testing along the entire life cycle. Using external folks of any kind, whether they're penetration testers under contract or the unwashed masses of the Internet, you really just want to know if there's any way to break into your system.

**Do some of the researchers or hackers that you coordinate with work on all 68 bug bounty programs? Do you have a stable of people who do that?**
Our platform has a built-in reputation system for hackers. And we're very transparent about how reputation is affected by different outcomes of a hacker reporting a vulnerability. You can imagine that their reputation improves when they find an issue that's serious enough to award a bounty. But to answer to your question, there are folks who are very prolific. They came in looking at one piece of software—maybe they came from Russia because they were looking at a specific piece of Russian software that was hosted on our platform for vulnerability reporting. And then they ended up spreading their knowledge and finding a lot of vulnerabilities in many other people's programs.

**Let's talk about breaking and building in security. Many conferences focus on breaking things and getting bounties. Is it possible to turn the people who do the breaking into builders who can develop secure systems? Do we even want to do that?**
I think it's possible. As much as we talk about getting people to think more like attackers, getting attackers to think like defenders is as hard or harder in a lot of ways. Generally speaking, they don't have to worry about the piece of software functioning in exactly the way the user or creator envisioned. But that functionality is going to be a huge part of securing software—you have to secure it so that it still works.

An example of incentivizing this switch is the bounties that I set up at Microsoft. There was a mitigation bypass bounty and a defense bounty. If you came up with a way to defeat all the mitigations but also came up with a way to defend against your new attack technique, you got more money.

This actually was paid out for the first time to researchers who work at ZDI [Zero Day Initiative]. Three bounties were launched in 2013, including one defensive bounty. Just a few months ago, ZDI's security researchers not only got [US]$100,000 from Microsoft for a new mitigation bypass but also claimed the very first defensive award of an additional $25,000. I thought it was cool that somebody actually created that shift in thinking from being a breaker to thinking, "Here are some viable ways that the platform could introduce measures to defend against this."

**So, Microsoft gets something for $25,000 that might cost a million bucks in the open market, if it were a real technology transfer.**

Who else except for Microsoft would buy a defensive technique? Looking for those gaps in markets and market motivations was definitely one of the things that went into my thinking when creating Microsoft's bounties. Create incentives for what you want and pay attention to those gaps in the market where you can develop an incentive for something that's useful only to you.

**I've always wondered whether you should teach developers to think like a bad guy. I used to think that you should, but then Microsoft's Steven Lipner changed my mind and convinced me that you can't teach all developers how to think like a bad guy. What do you think?**

You're right—you can't teach them all. I think one of the biggest problems I had as a penetration tester was convincing a developer that somebody would think like a bad guy.

**They look at you like you killed their puppy.**

Exactly. Who would kill my puppy? It was made of software. It was beautiful. So, convincing them of the likelihood that somebody would [think like a bad guy] is often the biggest challenge.

**What's the most effective way you found to do that?**

Unfortunately, you actually do have to show them something exploitable most of the time. I can't tell you how many times I've actually had to deliver an exploit to somebody when we were trying to tell them about a vulnerability. Then you have the danger of them thinking that exploit is the only vector.

**Tell me about some of your mentors. What did they help you with?**

My mentors have really run the gamut. Bob Bruen was my boss for a brief period and my friend for a lot longer when I worked at MIT. He would tell me if some crazy idea I had was actually viable and give me different strategies for making it happen. When I was working at MIT on network planning, I basically outlined how we should do it, but the folks were kind of politely and not so politely dismissing my advice. So I went into Bob's office and said, "I need to borrow your gray beard for a minute. Will you just tell them this, that, and the other thing?" And he said, "Look, I can tell them, and you're right that they'll probably listen to me. But eventually, you're going to have to grow your own gray beard."

**One of the biggest problems I had as a penetration tester was convincing a developer that somebody would think like a bad guy.**

He basically said that at some point, you're going to have to be able to communicate and command authority in areas where you have expertise. It's not always going to be easy, and you're right that I'm going to have an easier time walking into a room and having my expertise assumed, whereas you're going to have to work a little harder to establish it. But you can do it. I believe in you. And that was that.

Ever since then, I've been growing this gray beard. I'm kidding. But I've tried to make sure that I lead with my knowledge and try to be helpful in a situation first. And usually that ends up working. If you look young, you're going to have to lead with your authoritative foot.

I've had a lot of experiences with peer mentors. It ends up that your whole network is a mentor, in a way. You decide where you want to steer your ship, and then you go to your network of peers or colleagues and say, "I'm interested in moving into this area. Can you make some introductions?"

This past year at HackerOne, I've been fortunate in being able to take my expertise from vulnerability disclosure and coordination and parlay that into speaking to more and more policymakers and lawmakers about things that might affect security research, like the recent proposals to expand the US Computer Fraud and Abuse Act. This year has been about me leveraging my peer mentor network and getting introduced to people who are willing to hear about what this might mean to security.

The Silver Bullet Podcast with Gary McGraw is cosponsored by Cigital and this magazine and is syndicated by SearchSecurity.

**Gary McGraw** is Cigital's chief technology officer. He's the author of *Software Security: Building Security In* (Addison-Wesley 2006) and eight other books. McGraw received a BA in philosophy from the University of Virginia and a dual PhD in computer science and cognitive science from Indiana University. Contact him at gem@cigital.com.



*This article originally appeared in* IEEE Security & Privacy, *vol. 13, no. 4, 2015.*

EDITORS: Rick Kuhn, US National Institute of Standards and Technology, kuhn@nist.gov
Tim Weil, US Department of Interior, trweil@ieee.org

# Government-Funded R&D to Drive Cybersecurity Technologies

**Douglas Maughan,** *US Department of Homeland Security Science and Technology Directorate*
**David Balenson, Ulf Lindqvist, and Zachary Tudor,** *SRI International*

New and innovative cybersecurity technologies are essential to ensuring the security and resilience of our information systems and critical infrastructure. Such technologies must also meet the needs and requirements of IT professionals in the public and private sectors and be available as products via channels that are acceptable to those users. In this article, we present the US Department of Homeland Security (DHS) Science and Technology Directorate's (S&T) cybersecurity R&D program. This program funds top researchers in academia, industry, and government in developing new cybersecurity technologies across key areas, including trustworthy cyber infrastructure, cybersecurity research infrastructure, network and system security and investigations, cyber-physical systems, and transition and outreach.

Successful transition of cybersecurity technology from research to operational use is necessary to address rapidly evolving threats. The DHS S&T program works to significantly increase the transition rate of technologies out of R&D labs and into users' hands, and to increase the value those technologies bring to users. To illustrate the program's effectiveness, we describe examples of research technologies that are widely deployed and used. These examples show that a comprehensive cybersecurity R&D program can transition research results into operational use, with the potential to significantly improve cybersecurity for users in industry and government.

## Cybersecurity R&D

The cybersecurity problem is bigger than ever, with government and industry being victims of severe attacks, including successful ones against companies that specialize in security technology. For the past several years, we've seen rampant theft of sensitive information and intellectual property. We're also starting to see destructive attacks, with some targeting critical infrastructure. New and innovative solutions are desperately needed to get the problem under control, and these solutions must be widely deployed in operational settings to make a difference.

Two key actions are needed: we must increase R&D efforts, and we must get better at taking the best results of R&D all the way to deployable solutions. If we fail to accomplish either of these actions, our society will be set up for a disaster some years from now, when security solutions fail to match the challenges of the rapidly evolving IT world. It's therefore in the best interests of society to ensure that technologies are transitioned from research to users' hands.

## Mission and Strategy

The DHS S&T Cyber Security Division (CSD) executes the directorate's cybersecurity R&D program. Research requirements come from numerous sources across government and industry, including the White House and its National Security Council staff; DHS and its operational components; interagency collaborations across federal, state, and local governments; the critical infrastructure sectors, which are largely in the hands of private industry; and international partners. The sidebar identifies several key sources of cybersecurity R&D requirements from

 Published by the IEEE Computer Society 

the White House, DHS, and energy sector.

Through its core research program and activities, CSD plays a critical role in government-funded innovation to drive future cybersecurity technology. First and foremost, CSD funds research to develop and deliver new technologies, tools, and techniques to defend and secure current and future systems and networks. We describe this program next. As part of its R&D projects, CSD conducts and supports technology transition efforts with key stakeholders in government, IT security companies, venture capital, open source efforts, and international partners. Finally, CSD provides global R&D leadership and coordination within the government, academia, the private sector, and the international cybersecurity community.

The CSD research program and activities are largely funded through open competitive solicitations, including cybersecurity R&D Broad Agency Announcements (BAAs), long-range BAAs, and Small Business and Innovative Research (SBIR) solicitations that solicit R&D services to generate potential solutions in technical topic areas. Current solicitations can be found via the DHS BAA Program Portal (https://baa2.st.dhs.gov/portal/BAA/) or FedBizOps (www.fbo.gov). For example, the DHS S&T Cyber Security Division 5-Year BAA[1] seeks new solutions in the areas of distributed-denial-of-service defense (DDoSD), data privacy, mobile technology security, cyber-physical system security, and additional forthcoming areas.

## Key Sources of Cybersecurity R&D Requirements

- Comprehensive National Cybersecurity Initiative (CNCI): www.whitehouse.gov/sites/default/files/cybersecurity.pdf
- Strategic Plan for the Federal Cybersecurity Research and Development Program: www.whitehouse.gov/sites/default/files/microsites/ostp/fed_cybersecurity_rd_strategic_plan_2011.pdf
- The 2014 Quadrennial Homeland Security Review (QHSR): www.dhs.gov/quadrennial-homeland-security-review-qhsr
- Blueprint for a Secure Cyber Future: The Cybersecurity Strategy for the Homeland Security Enterprise: www.dhs.gov/blueprint-secure-cyber-future
- A Roadmap for Cybersecurity Research: www.dhs.gov/sites/default/files/publications/CSD-DHS-Cybersecurity-Roadmap.pdf
- Roadmap to Achieve Energy Delivery Systems Cybersecurity: energy.gov/oe/downloads/roadmap-achieve-energy-delivery-systems-cybersecurity-2011

## R&D Program

The CSD research program cuts across many key areas and facets of cybersecurity. The objective of the *trustworthy cyber infrastructure* area is to develop standards, policies, processes, and technologies to enable more secure and robust global cyber infrastructure. It also aims to identify components in need of protection, apply analysis capabilities to predict and respond to cyberattack effects, and provide situational understanding to providers. Projects in this area include secure protocols such as secure routing and DNS, Internet measurement and attack modeling (IMAM), and DDoSD. An article on Clique, a situational awareness technology arising from this area, appears in this issue's Smart Systems department.

The *network and system security and investigations* area is developing new and innovative methods, services, and capabilities to secure future networks and systems, ensuring that they are usable and that their security properties can be measured, and to provide the tools and techniques needed for combating cybercrime. This area includes security for cloud-based systems, mobile device security, identity management, data privacy, software quality assurance, usable security and security metrics, and investigation capabilities for law enforcement.

The *cyber-physical systems/process control systems* area seeks to ensure that necessary security enhancements are added to the design and implementation of ubiquitous cyber-physical systems and process control systems, with an emphasis on transportation, emergency response, energy, and oil and gas systems. Key initiatives include Cyber-Physical Systems Security (CPSSEC), Trustworthy Cyber Infrastructure for the Power Grid (TCIPG), and Linking the Oil and Gas Industry to Improve Cybersecurity (LOGIIC).

A critical supporting element is *research infrastructure* such as test facilities, realistic datasets, tools, and methodologies to enable global cybersecurity R&D community researchers to perform at-scale experimentation on their emerging technologies with respect to system performance goals. Projects in this area are an experimental research testbed (the DETER Project), a research data repository (the Protected Repository for the Defense of Infrastructure against Cyber Threats, or PREDICT), and the Software Assurance Marketplace (SWAMP).

Finally, to ensure that results affect operational systems and users, the *transition and outreach* area seeks to accelerate the transition of mature, federally funded cybersecurity R&D technology

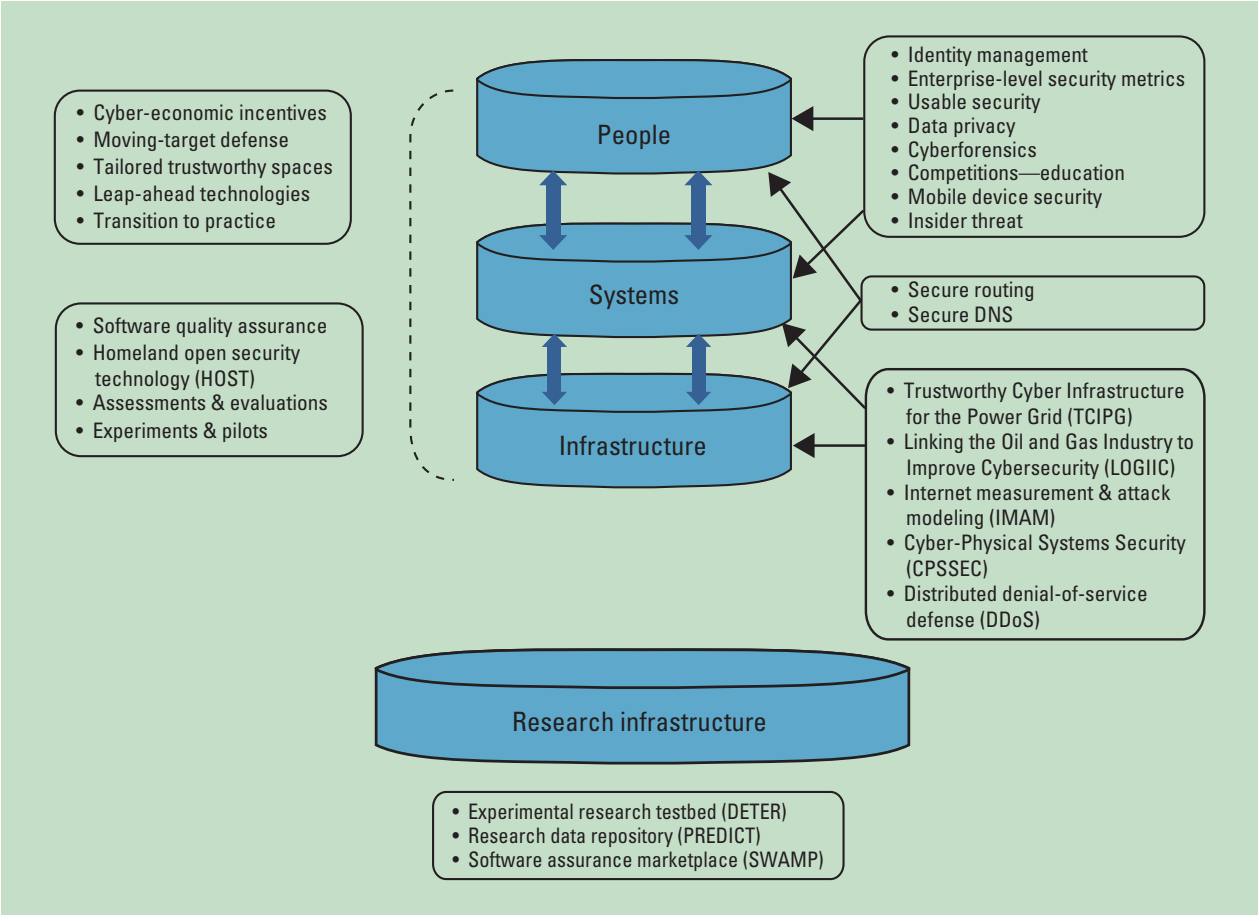


**Figure 1. Cybersecurity R&D projects and relationships. The Cyber Security Division research program develops new cybersecurity technologies that affect people, systems, and infrastructure; the program also develops key infrastructure that supports cybersecurity research.**

into widespread operational deployment, as well as educate and train the cybersecurity workforce of today and tomorrow through multiple methods, models, and activities. Projects in this area include Transition to Practice (TTP), cybersecurity competitions, and the National Initiative for Cybersecurity Education (NICE).

Figure 1 depicts the cybersecurity R&D projects and their relationships to people, systems, and infrastructure, and also shows the supporting research infrastructure.

## Transition of Research Results

Transition of research results into current, emerging, and future systems is clearly and explicitly stated to be integral to the DHS S&T cybersecurity R&D program (www.dhs.gov/cyber-research). The successful transition of cybersecurity technology from research to operational use is absolutely necessary to address rapidly evolving threats, but it's also a difficult endeavor with many challenges. We describe these challenges and the means to overcome them, including key elements of the DHS S&T cybersecurity R&D program, in further detail in prior work.[2] With technology transition built in as an integral component of the full research, development, test, evaluation, and transition life cycle, the program works to transition technologies from government-funded projects to the commercial marketplace through spin-offs, acquisitions, and commercial products, including open source software.

## Example Transition Successes

The DHS S&T cybersecurity R&D program has transitioned more than 30 products since 2004, including 18 commercial, six open source, and two government off-the-shelf products, along with research infrastructure, an open source standard, and multiple knowledge products. Table 1 lists some examples of technologies developed under or supported by the program that

*Table 1. Cybersecurity R&D technology transition successes.*

| | |
|---|---|
| Komoku Rootkit Detection | Advanced rootkit security detection technology; formed startup in 2004 and acquired by Microsoft in 2008 |
| Endeavor Systems Malware Analysis | Commercial botnet detection and mitigation tool; acquired by McAfee in 2009 |
| Open Information Security Foundation (OISF) Suricata | High-performance intrusion-detection system, intrusion-prevention system, and network security monitoring engine; active open source product first released in 2009 |
| IronKey Secure USB Device | Secure data storage and protection technology; acquired by Imation in 2011; renamed as Marble Cloud |
| HBGary Forensics Tools | Memory and malware analysis tools; acquired by ManTech International in 2012 |

have been successfully transitioned to widespread use.

Cybersecurity research is a key area of innovation to support our global economic and national security futures. The DHS S&T cybersecurity R&D program continues with an aggressive research agenda to solve the cybersecurity problems of current and future infrastructure and systems. The program strongly emphasizes technology transition and will affect cyber education, training, and awareness of our current and future cybersecurity workforce.

Industry and government bring distinctly different resources and expertise to the table, and partnerships between the two are essential to success in this important endeavor. Those who are funding and conducting R&D need to partner with the eventual technology users to ensure that the solutions provide high value to users and solve the most important problems.

As IT professionals, we encourage you to get involved by learning about current and planned future research and sharing your future needs and requirements. We also invite you to support pilots and experiments that help ensure that future cybersecurity technologies meet your needs. The ultimate success of cybersecurity research efforts depends on the IT community working with the research community to help find and deploy the best ideas and solutions to real-world problems. IT

## Acknowledgments

*The views and conclusions contained herein are the authors' and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the US Department of Homeland Security (DHS) or the US government. The work by SRI International was funded by the DHS Science and Technology Directorate under contract no. HSHQDC-10-C-00144. The SRI authors thank DHS S&T program managers Greg Wigton and Mike Pozmantier for their support.*

## References

1. "DHS S&T Cyber Security Division 5-Year Broad Agency Announcement," solicitation no. HSHQDC-14-R-B0005, 2014; www.fbo.gov/spg/DHS/OCPO/DHS-OCPO/HSHQDC-14-R-B0005/listing.html.
2. D. Maughan et al., "Crossing the 'Valley of Death': Transitioning Cybersecurity Research into Practice," *IEEE Security & Privacy*, vol. 11, no. 2, 2013, pp. 14–23; doi.ieeecomputersociety.org/10.1109/MSP.2013.31.

***Douglas Maughan*** *is director of the Cyber Security Division in the US Department of Homeland Security (DHS) Science and Technology (S&T) Directorate, Homeland Security Advanced Research Projects Agency, where he oversees cybersecurity R&D activities. Contact him at douglas.maughan@dhs.gov.*

***David Balenson*** *is a senior computer scientist in the Computer Science Laboratory at SRI International. He's interested in technology transition and provides technical support, subject matter expertise, and project management for the DHS S&T cybersecurity R&D program. He's a member of IEEE. Contact him at david.balenson@sri.com.*

***Ulf Lindqvist*** *is a program director in the Computer Science Laboratory at SRI International. He manages research in critical infrastructure security and leads SRI's support for the DHS S&T cybersecurity R&D program. He's a member of the IEEE Computer Society. Contact him at ulf.lindqvist@sri.com.*

***Zachary Tudor*** *is a program director in the Computer Science Laboratory at SRI International. He supports operational and R&D projects in critical infrastructure security and provides technical support, subject matter expertise, and project management for the DHS S&T cybersecurity R&D program. He's a member of the IEEE Computer Society. Contact him at zachary.tudor@sri.com.*

*This article originally appeared in* IT Professional, *vol. 17, no. 4, 2015.*

# Recent US Cybersecurity Policy Initiatives: Challenges and Implications

**Nir Kshetri,** University of North Carolina at Greensboro

*In recent years, the US government has introduced several policy measures aimed at tackling the growing cyberthreats facing the country, but many challenges and concerns could arise as a result of their implementation.*

The Obama administration recently introduced a range of initiatives to strengthen US cybersecurity (CS) policy. These initiatives, as emphasized in the January 2015 State of the Union address, aim to secure networks and trade secrets, protect privacy, and ensure that government agencies share intelligence to combat cyberthreats. On 13 February 2015, President Obama also signed Executive Order (EO) 13691, "Promoting Private Sector Cybersecurity Information Sharing," which lays out a strategy for expanded collaboration between private companies and the federal government (www.whitehouse.gov/the-press-office/2015/02/13/executive-order-promoting-private-sector-cybersecurity-information-shari).

These efforts to achieve a more secure cyberspace complement other CS-related policies and programs adopted in recent years, including 2011's National Strategy on Trusted Identities in Cyberspace (NSTIC),[1] which aims to create an "identity ecosystem" to increase individuals' and organizations' confidence in engaging in online transactions; the National Initiative on Cybersecurity Education (NICE) program (http://csrc.nist.gov/nice), which seeks to address the shortage of CS-related human capital; and the International Strategy for Cyberspace (www.whitehouse.gov/sites/default/files/rss_viewer/international_strategy_for_cyberspace.pdf), which establishes "norms of responsible behavior" for nations' cyberspace actions.

 Published by the IEEE Computer Society 2376-113X/15/$31.00 © 2015 IEEE

EDITOR **JEFFREY VOAS**
National Institute of Standards and Technology; j.voas@ieee.org

**TABLE 1.** Recent US cybersecurity (CS) policy initiatives.

| Policy initiative | Features and contribution to CS goals | Key challenges and concerns |
|---|---|---|
| Establish federal breach notification legislation to notify employees and customers of a data breach | Companies experiencing a data breach must notify affected consumers within 30 days | Weaker than current data breach laws in some states (for example, California) |
| | | Concerns regarding the appropriateness of the 30-day reporting timeline |
| Facilitate greater information sharing between the federal government and the private sector | Understanding past hacking activities will help prevent or combat future cyberattacks | Unclear added value over what's already being shared among companies |
| | Gives private sector "targeted" liability protection to share information, including various cyberthreat indicators | Fear of liability is only part of the problem |
| | Government will disclose more classified threat information to the private sector | Liability protection might discourage companies to strengthen CS practices |
| | Creates the Cyber Threat Intelligence Integration Center (CTIIC) | Proposal relies on privacy guidelines that haven't yet been written |
| | | Anonymization might offer false reassurance |
| | | Shared information might include confidential and proprietary information about, for example, a company's security system |
| Amend the Racketeer Influenced and Corrupt Organizations Act (RICO) and the Computer Fraud and Abuse Act (CFAA) | Is expected to modernize law enforcement agencies' tools to fight cybercrime | Proposed RICO changes invite potential abuse by law enforcement agencies |
| | | Proposed CFAA revisions contain vague language |

Here, I examine how the most recent policy initiatives can help achieve national CS objectives, and outline the challenges and concerns that might arise during their implementation.

## RECENT CS POLICY INITIATIVES

Table 1 summarizes three major CS policy initiatives recently announced by the Obama administration.

### Federal breach notification legislation

On 12 January 2015, President Obama proposed legislation requiring companies that experience a data breach to notify affected customers within 30 days of the breach discovery (www.whitehouse.gov/the-press-office/2015/01/12/remarks-president-federal-trade-commission). Currently, 47 states have different laws regarding how people should be notified when breaches involve personally identifiable information (PII).[2] The proposal unifies the complex patchwork of inconsistent state laws and regulations, and is expected to reduce compliance costs for businesses.

A similar requirement already exists for federal departments and agencies under the 2014 Federal Information Security Modernization Act (FISMA). FISMA requires the director of the Office of Management and Budget to periodically update federal agency data breach notification policies and guidelines, and to notify various congressional committees no later than 30 days after a data breach is discovered. FISMA also mandates federal agencies to notify those affected "as expeditiously as practicable and without unreasonable delay" after discovery of a data breach.[3]

### Information sharing between government agencies and the private sector

EO 13691 lays out a framework for US companies to share cyberthreat information with one another and with government agencies. This EO and the federal breach notification legislative proposal complement each other (www.whitehouse.gov/the-press-office/2015/02/12/fact-sheet-executive-order-promoting-private-sector-cybersecurity-inform) and are related to EO 13636, "Improving Critical Infrastructure Cybersecurity" (www.whitehouse.gov/issues/foreign-policy/cybersecurity/eo-13636), and Presidential Policy Directive 21 (PPD-21), "Critical Infrastructure Security and Resilience" (www.whitehouse.gov/the-press-office/2013/02/12/presidential-policy-directive-critical-infrastructure-security-and-resil)—both signed by President Obama on 12 February, 2013—in that they all emphasize the roles of the private sector and information sharing between business and government. EO 13636 mandated that the US government work with "owners and operators of critical infrastructure" to share cyberthreat information and create a framework for protecting critical infrastructure. It also sought to implement common CS standards.

A key provision of EO 13691 is the establishment of information sharing and analysis organizations (ISAOs), which will comply with voluntary standards envisioned by the EO. The Department of Homeland Security (DHS) and the newly created National Cybersecurity and Communications Integration Center (NCCIC) are given the authority to share data with ISAOs, so organizations will be able to access classified CS data.[4] Similarly, the Cyber Threat Intelligence Integration Center (CTIIC) was also created in February 2015 to carry out "coordinated cyberthreat assessments" based on information received from various sources. The CTIIC aims to provide "all-source analysis" of cyberthreats to policymakers and assist relevant agencies in dealing with those threats (www.whitehouse.gov/the-press-office/2015/02/25/fact-sheet-cyber-threat-intelligence-integration-center).

The rationale behind these proposals is that timely information sharing would facilitate a better understanding of past cyberattacks in order to prevent future ones, as perpetrators often use the same malware to infiltrate multiple targets. Although there are some industry-specific initiatives to share cyberthreat intelligence, many cybercrimes impact numerous industries. The proposals would make it easier for companies to share intelligence with the NCCIC, including various cyberthreat indicators such as attempts to access restricted files, the way in which a website runs, and the ways in which a company utilizes user data.

Targeted liability protection will be granted to share data.[5] To qualify for liability protection, companies are required to take reasonable measures to ensure that irrelevant PII is removed before sharing information. They're also required to comply with additional privacy guidelines created by the Director of National Intelligence, the Attorney General, and the DHS (www.defense.gov/news/newsarticle.aspx?id=123966; www.whitehouse.gov/the-press-office/2015/01/13/securing-cyberspace-president-obama-announces-new-cybersecurity-legislat).

### RICO and CFAA

In his 2015 State of the Union address, President Obama proposed including cybercrimes in the Racketeer Influenced and Corrupt Organizations Act (RICO), originally passed in 1970, to modernize law enforcement agencies' tools to fight cybercrime. Under the proposed legislation, the maximum penalty for serious cybercrimes such as running illegal marketplaces to sell drugs and stolen identify information will be 10 to 20 years in prison. However, the proposal aims to ensure that "insignificant conduct" (such as sharing passwords for online services such as Netflix and Hulu) will not fall within the statute's scope.

Likewise, President Obama proposed amending the Computer Fraud and Abuse Act (CFAA) of 1986 to expand the definition of "unauthorized access." Strictly applied, the proposed law makes it a crime to use a computer "for a purpose that the accesser knows is not authorized by the computer owner."

## KEY CHALLENGES AND CONCERNS

A number of challenges and concerns could arise if the recent CS policy initiatives are implemented (see Table 1).

### Federal breach notification legislation

The proposed federal breach notification legislation has been criticized on the grounds that it's weaker than some states' current data breach laws. For instance, California requires businesses to provide notice of a breach "without unreasonable delay" unless law enforcement determines that such notification might impede investigation. Companies are also required to notify the State Attorney General if the breach involves more than 500 users' information.[6]

The proposed legislation also doesn't make it clear when a security breach is viewed as having been discovered—for instance, upon suspicion or confirmation (www.coxsmithbanking.com/proposed-federal-data-security-breach-notification-law). Some investigations can take several weeks or even months. Moreover, initial awareness of a breach often doesn't reveal enough details to determine the best way to report it.

Some argue that adding a 30-day reporting timeline would intensify the challenges organizations face because assessing and diagnosing the impacts and origins of a cyberattack is a time-consuming process. Opponents of this view argue that 30 days is too long. For example, in the Target data breach of 2013, buying and selling of stolen credit cards started in

underground markets only a few days after the breach was discovered.[7]

### Information sharing between government agencies and the private sector

One criticism of the information sharing proposal pertains to the unclear added value of sharing information between the government and the private sector over what's already being shared among many companies. For instance, the Retail Cyber Intelligence Sharing Center (www.r-cisc.org) was established in 2014 by more than 50 retailers to share cyberthreat information. Likewise, the energy sector established the Oil and Natural Gas Information Sharing and Analysis Center (http://ongisac.org) for a similar purpose. The Financial Services Information Sharing and Analysis Center (FS-ISAC; www.fsisac.com) was launched in 1999 to promote sharing cyberthreat information among financial services firms. In 2013, FS-ISAC extended its charter to include financial services firms worldwide. Finally, CS vendors including Palo Alto Networks, Fortinet, and Symantec formed the Cyber Threat Alliance (http://cyberthreatalliance.org) in 2014 to share intelligence.

Regarding the role of liability protection as an incentive to share information, some critics point out that fear of liability is only part of the problem. A chief concern among businesses is that the government lacks the resources and experience to successfully prosecute cybercriminals. The government's poor track record supports this view.[8] Others maintain that liability protection might discourage companies from strengthening CS practices and could even stimulate widespread distribution of personal data.[9]

According to the American Civil Liberties Union's policy advisor, the proposal for information sharing doesn't sufficiently ensure that all PII will be stripped before sharing. Privacy advocates are concerned that even if the privacy guidelines are well developed, it's almost impossible to know whether the guidelines have been followed and enforced properly.[6]

Some critics have argued that the only positive aspect of the proposed Cyber Intelligence Sharing and Protection Act (CISPA) is the provision requiring "a process to anonymize and safeguard information."[10] CISPA in its original form was passed in the House in 2012 and again in 2013 but not by the Senate; an updated version of the bill was introduced in the House in 2015 but hasn't yet come to a vote. Various interest groups have argued that CISPA, as well as a similar law proposed in the Senate in 2014, the Cybersecurity Information Sharing Act (CISA), contain too few limits on the government's monitoring of PII. Researchers have found that it's possible to use a data aggregation process to convert semi-anonymous or certain personally nonidentifiable information into non-anonymous information or PII (www.isaca.org/Knowledge-Center/Research/ResearchDeliverables/Pages/Generating-Value-From-Big-Data-Analytics.aspx). Thus, anonymization might offer only false reassurance.

### RICO and CFAA

A main criticism of the proposed amendment to RICO is that due to its broad nature, the revised law would be left open to potential abuse by law enforcement agencies. Hackers, computer scientists, and curious users trying to find security holes could be prosecuted and face felony charges. For this reason, some argue that the proposed legislation could actually make cyberspace less secure.

Another challenge involves the novelty of cybercrime. RICO was enacted 45 years ago, so law enforcement agencies have long prosecuted organized crimes under the act. Because cybercrime is relatively new and often difficult to explain to judges and juries, critics have emphasized the importance of clear "red lines" to apply the law in modern times.[10]

> In addition to details about threats such as viruses, malware, spyware, and Trojan horses, shared information should also include perpetrators' modus operandi.

Critics also worry that the language of the proposed revisions to CFAA is too vague to translate into effective legislation. They say the legislation could encourage some prosecutors to take advantage of this vagueness to aggressively pursue computer scientists or curious users for hacking offenses. Likewise, an individual could be guilty of violating the law for engaging in innocent behavior such as sharing a Netflix password with family members or inadvertently clicking on a link that leads to unauthorized content.[11]

## IMPLICATIONS AND TAKEAWAYS FOR BUSINESSES

Cybercriminals are increasingly modifying their approaches to suit different purposes. Analysts have noted that techniques once found in state-sponsored cyberwarfare are being deployed against corporate targets. Likewise, industrial espionage is being expanded to control physical assets via hacking, which used to be deployed only to capture commercial secrets and intellectual property.[12] Information sharing, then, must extend beyond the current narrow industrial focus to include a broader national interest. A positive aspect of the proposed initiatives is that they aim to achieve this by expanding information sharing. In addition to details about threats such as viruses, malware, spyware,

and Trojan horses, shared information should also include perpetrators' modus operandi.

A large proportion of cybercriminals targeting US operations have jurisdictionally shielded themselves by operating from countries that lack strict law enforcement or have little or no cooperation with the US regarding cybercrime. In this regard, the proposals exhibit a low degree of outward orientation. The US–China Business Council, which represents about 230 US companies with operations in China such as Boeing, Caterpillar, Citigroup, and JPMorgan Chase, have asked the US and Chinese governments to work together to address the growing problem of cyberattacks.[13]

The recent Obama administration CS policy initiatives don't directly address how US organizations can better protect themselves against state-sponsored hackers such as those in North Korea. There has been limited progress in the development of international norms for cyberspace engagement as envisioned by the International Strategy for Cyberspace.[14] If implemented, the proposed legislation might have heterogeneous effects across firms. For example, if the main threats facing an organization are inside attackers using their credentials to attain illegitimate goals, a more severe punishment is likely to deter such criminality. Businesses experiencing attacks by mostly foreign hackers, on the other hand, might not necessarily be any safer.

As reflected in the NICE program, the development of a cyber-savvy workforce has been a key priority for the US (www.whitehouse.gov/issues/foreign-policy/cybersecurity/summit#section-five-things-to-know). One of NICE's goals is to increase qualified CS professionals by 20 percent by 2015.[15] However, there's currently a significant shortage of CS manpower. According to the National Institute of Standards and Technology, more than 700,000 CS professionals will be needed by 2015,[16] but there were more than 30,000 open CS positions in federal agencies in 2014.[17] Moreover, many CS specialists with practical computer expertise are often self-taught.[18] This shortage of CS experts underscores the importance of organizational initiatives to provide employees with CS-related training.

Due to the various shortcomings and imperfections of the current regulatory framework, businesses need to take more CS measures than those required by law. Organizations could implement effective self-regulatory strategies instead of waiting for CS laws to be enacted. It's critical to have a well-developed plan for post-breach resilience so businesses can quickly return to normal operations.[19] For instance, in addition to weak cyberdefense mechanisms, Sony Pictures Entertainment was criticized for its lack of disaster recovery provisions in the wake of the 2014 hack: current and former employees complained that they didn't get information about identifying protection measures or registering for free credit monitoring.[20] In this regard, the proposed initiatives put pressure on businesses to be better prepared to deal with data breaches and to make recovery easier and faster.

Interstate harmonization of data breach notification legislation is likely to result in lower compliance costs regarding data breaches. For businesses that operate in one or a few states, however, the costs related to reporting a data breach could increase or decrease. For instance, the proposed legislation is likely to have a favorable effect on businesses operating only in California, which already has strict reporting requirements. For a business operating in a state with looser reporting requirements, on the other hand, the proposed legislation could lead to an increase in related costs.

Interstate harmonization of data breach notification legislation is likely to result in lower compliance costs regarding data breaches.

Despite some privacy concerns that need to be addressed, greater information sharing between the federal government and the private sector will increase our understanding of cybercriminals' modus operandi and allow us to take defensive and precautionary measures to reduce the risk of becoming a victim. The severity of punishment under the proposed amendments of RICO and CFAA are likely to deter cybercrime, especially if the certainty of punishment is increased with stronger law enforcement measures against such crimes.

## REFERENCES


1. H.A. Schmidt, "The National Strategy for Trusted Identities in Cyberspace," blog, 25 June 2010; www.whitehouse.gov/blog/2010/06/25/national-strategy-trusted-identities-cyberspace.
2. R. King, "New EU Cyber Security Directive to Impact US Companies," blog, 7 Feb. 2013; http://blogs.wsj.com/cio/2013/02/07/new-eu-cyber-security-directive-to-impact-u-s-companies.
3. S.B. Hoar, "Congress Passes the Federal Information Security Modernization Act of 2014: Bringing Federal Agency Information Security into the New Millennium," blog, 18 Dec. 2014; www.privsecblog.com/2014/12/articles/cyber-national-security/congress-passes-the-federal-information-security-modernization-act-of-2014-bringing-federal-agency

-information-security-into-the-new-millennium.
4. T. Wolverton, "Silicon Valley: Obama Calls on Corporations to Work with Government to Prevent Cyber-attacks," *San Jose Mercury News*, 13 Feb. 2015; www.mercurynews.com/business/ci_27520838/obama-issues-cybersecurity-order-at-open-summit.
5. J.H. Davis, "Obama Calls for New Laws to Bolster Cybersecurity," *The New York Times*, 13 Jan. 2015; www.nytimes.com/2015/01/14/us/obama-to-announce-new-cyberattack-protections.html.
6. M. Jaycox and L. Tien, "Obama's Computer Security Solution Is a Mishmash of Old, Outdated Policy Solutions," Electronic Frontier Foundation, 16 Jan. 2015; www.eff.org/deeplinks/2015/01/obamas-computer-security-solution-mish-mash-old-outdated-policy-solutions.
7. R. King and C. Boulton, "CIOs Eye Obama Cybersecurity Push with 'High Level of Interest,'" blog, 20 Jan. 2015; http://blogs.wsj.com/cio/2015/01/20/cios-eye-obama-cybersecurity-push-with-high-level-of-interest.
8. D.M. Upton, "The Flaws in Obama's Cybersecurity Initiative," *Harvard Business Rev.*, 20 Jan. 2015; https://hbr.org/2015/01/the-flaws-in-obamas-cybersecurity-initiative.
9. D. Froomkin, "Obama's Cyber Proposals Sound Good, But Erode Information Security," *The Intercept*, 20 Jan. 2015; https://firstlook.org/theintercept/2015/01/20/obamas-cyber-proposals-sound-good-totally-clueless.
10. P. Tucker, Why Obama's Cybersecurity Plan May Not Make Americans Safer," *The Atlantic*, 22 Jan. 2015; www.theatlantic.com/technology/archive/2015/01/why-obamas-cybersecurity-plan-may-not-make-average-americans-safer/384733.
11. D. Storm, "Obama's Cybersecurity Plan: Share a Password, Click a Link, Go to Prison as a Hacker," *Computerworld*, 21 Jan. 2015; www.computerworld.com/article/2872368/obamas-cybersecurity-plan-share-a-password-click-a-link-go-to-prison-as-a-hacker.html.
12. C. Binham, "The Hacker Hunters," *FT Mag.*, 21 Nov. 2013; www.ft.com/intl/cms/s/2/bccc8f3c-523c-11e3-8c42-00144feabdc0.html.
13. D. Palmer, "Trade Group Wants U.S.-China Action on Cyber Security Threats," Reuters, 4 Feb. 2013; www.reuters.com/article/2013/02/04/us-usa-china-trade-idUSBRE9130Y220130204.
14. K. Eichensehr, "The US Needs a New International Strategy for Cyberspace," blog, 24 Nov. 2014; http://justsecurity.org/17729/time-u-s-international-strategy-cyberspace.
15. W.H. Tipton, "Gotta' Hand it to NICE: A Strategy with the Big Picture in Mind," blog, 2 Sep. 2011; http://breakinggov.com/2011/09/02/gotta-hand-it-to-nice-a-strategy-with-the-big-picture-in-mind
16. "Government Preps Next Generation of Cybersecurity Employees," Homeland Security News Wire, 8 Dec. 2011; www.homelandsecuritynewswire.com/bull20111208-government-preps-next-generation-of-cybersecurity-employees.
17. D.J. Summers, "For Uncle Sam, Trouble Raising a Cyber Army," *Fortune*, 3 Oct. 2014; http://fortune.com/2014/10/03/government-cyber-security-shortage.
18. T. Risen, "Companies Unprepared as Hacking Increases," *US News & World Report*, 28 May 2014; www.usnews.com/news/articles/2014/05/28/companies-unprepared-as-hacking-increases.
19. P.M. Barrett, "The Cybersecurity Myths That Small Companies Still Believe, *Bloomberg Business*, 24 Nov. 2014; www.bloomberg.com/bw/articles/2014-11-24/the-cyber-security-myths-that-small-companies-still-believe.
20. B. Fritz, "Victims of Sony Breach Left Fuming," *The Wall Street J.*, 8 Dec. 2014; www.wsj.com/articles/victims-of-sony-breach-left-fuming-1418082738.

*This article originally appeared in* Computer, *vol. 48, no. 7, 2015.*

**NIR KSHETRI** is a professor at the University of North Carolina at Greensboro and a research fellow at the Research Institute for Economics and Business Administration at Kobe University, Japan. Contact him at nbkshetr@uncg.edu.

Editor: Gustavo Rossi • gustavo@lifia.info.unlp.edu.ar

# Ancillary Impacts of Multipath TCP on Current and Future Network Security

**Catherine Pearce** • *Neohapsis*

**Sherali Zeadally** • *University of Kentucky*

Multipath TCP (MPTCP) is an experimental TCP extension designed to add functionality to TCP while remaining backwards-compatible with most networks and devices. MPTCP changes TCP's behavior from how it's commonly understood in ways that go beyond the security of MPTCP itself, with ancillary implications challenging how network security is practiced and implemented. Here, the authors investigate the implications for network security — both in the transitional state, where MPTCP is partially supported, and in a future where every device supports MPTCP. They find that while MPTCP isn't widely supported, increasing support will stimulate changes to common network security rationales and paradigms.

TCP is ubiquitous in modern networking, with much of the world's infrastructure depending upon it, but its single-path nature has significant shortcomings in many modern applications. TCP falls short where there's a need to combine the bandwidth of multiple network connections, or to roam across different network locations while maintaining a connection. Although many proposed solutions have been revolutionary (comprising complete replacements) – notably Stream Control Transmission Control Protocol (SCTP) – recent proposals have been more evolutionary extensions of TCP itself.

Given the importance of TCP to networked devices, any changes or extensions will have far-reaching implications. Multipath TCP (MPTCP) is one such extension. MPTCP adds the ability to split TCP flows over multiple network interfaces or paths, and to change them on-the-fly without dropping the connection. It has moved toward official IETF standardization[1] and runs transparently on most existing infrastructures. However, MPTCP fundamentally changes how TCP communication actually works. MPTCP appears the same as TCP on a network, and it's transparent to most software, but differs from TCP in various critical aspects. It changes how data flows are addressed, what's meant by a "connection," and how data flows between connected systems.

While securing the MPTCP protocol itself against attack has been the primary research focus thus far,[2,3] this approach isn't sufficient to ensure network security, as no protocol is used in isolation. Security assessments must take the environment into account; a protocol may *itself* be secure from attack, but interact with other systems in a way that introduces problems. MPTCP breaks common assumptions made by network devices that have been valid until now, with significant implications for security – both good and bad. It's this additional aspect – the security implications *around* MPTCP, that aren't directly related to the security of MPTCP in itself – which our work focuses on. These implications can be grouped into two key categories: those which apply during the transitional period when

 Published by the IEEE Computer Society 

only some infrastructure is aware of MPTCP, and those which apply when MPTCP is fully deployed and all infrastructures are aware of MPTCP.

Thus, here we introduce and briefly discuss several key issues surrounding an area that seldom has been discussed in the literature previously: the impact of MPTCP on security-related network administration. In particular, we discuss the impact on key items of intrusion-detection systems (IDS), traffic monitoring, and traffic filtering. Although we touch upon them, the focus of this article is neither on the security of MPTCP itself, nor on attacks against it. Additionally, although we speculate on ways MPTCP may be used to enhance privacy, the specifics of this are beyond this work's scope.

## Motivations and Background of MPTCP

TCP ensures reliable, bidirectional, in-order data delivery. TCP's reliability and versatility has made it the primary protocol for most non-multimedia applications, such as HTTP Web browsing and email.

### TCP Shortcomings and MPTCP Motivators

TCP defines each connection as between two singular fixed endpoints (in a 1–1 cardinality), each defined by an IP address and a TCP "port number." This static 1–1 mapping is the key shortcoming of TCP; as a result, TCP can't support increasingly more common scenarios where devices need to shift their communication address or interface regularly (mobile devices), use multiple redundant links to their full capacity (data centers and mesh networks), or otherwise utilize multiple different network connections in the same TCP connection. These shortcomings have motivated the development of Multipath TCP.[4]

### Multipath Networking

Multipath networking isn't a new idea, and multipath communications are a core concept of IP-based networks, sending packets along the best route at any given time and place. The idea of a multipath TCP itself dates back to at least 1995.[5] The SCTP protocol was developed to achieve many of the same goals, but didn't result in widespread adoption, primarily due to the requirement that software and hardware be rebuilt or modified to support it.[6] Most research into the feasibility, design, standardization, and implementation of what has become MPTCP has occurred in the last decade.[7] Regarding the actual mechanism for widely used multipath communications, the feasibility of extending TCP was investigated,[8–11] leading to the design of MPTCP as a TCP extension.[12]

### Multipath TCP (MPTCP)

The IETF Multipath TCP working group was set up in 2010 and was responsible for the development of what has become the MPTCP standard. The first MPTCP draft of what would become the core MPTCP specification is RFC 6824,[1] which was introduced in mid-2010. Protocol design and implementation have progressed significantly since then. MPTCP was designed to run on top of today's Internet Protocol stack by utilizing a standards-compliant TCP with additional details in the TCP options field of the TCP header. Furthermore, it's designed with a worst-case fallback to standard TCP if errors occur, to make continuing MPTCP communications impossible. Further details on the rationales behind MPTCP design are given elsewhere.[4]

An MPTCP connection is made up of one or more subflows, with each subflow being a TCP flow with a TCP option communicating MPTCP information. There's no such thing as an MPTCP packet — MPTCP packets are simply standard compliant TCP packets with additional TCP options to handle MPTCP details. As a result, MPTCP is effectively a subtype of TCP traffic, with additional capabilities added for multipath communications. More in-depth discussions of Multipath TCP's background are given elsewhere.[4,13–15]

MPTCP operates transparently at the transport layer; therefore, most existing software supports it without modification. Connections are set up along paths the same way as in TCP, but additional paths are handled behind the scenes by the MPTCP implementation. This provides backward compatibility, but some multipath situations simply don't fit the transparent model well — such as listening on multiple interfaces or a connection remaining open after the associated network interface is down. In the long term, even if operating system APIs are modified, the software will need to have a more complex understanding of addressing if it wishes to utilize multipath connectivity to the best effect.

### MPTCP Traffic Flow Control and Data/Subflow-Mapping

MPTCP uses an initial subflow to set up connection details, including keyed-hash message authentication code (HMAC) keys used for additional subflow authentication and MPTCP data signing. Beyond the initial connection setup, there's no differentiation between subflows; each party has full control of how it chooses to send its own traffic, as well as how it divides it among subflows. Most current MPTCP implementations use a combination of well-understood congestion algorithms and multipath adapted ones,[4] but there's no requirement that they do so. If a client wishes to send alternate bytes across alternate links, it's free to do so, however inefficient that might prove to be.

This work doesn't discuss standard path-selection and congestion-control algorithms, as the security issues arise from how traffic *may* be divided, rather than how it usually is. We aren't concerned with the way things are normally done within the

*Table 1. The results of our surveys of MPTCP support in Alexa's[21] top 1 million websites.**

| Response type | Meaning | December 2013 tests (%) | April 2014 tests (%) |
|---|---|---|---|
| ICMP with MPTCP | TCP is rejected, but MPTCP is able to travel the whole path | 257 (0.03%) | 422 (0.04%) |
| MPTCP | Multipath-capable TCP response | 825 (0.11%) | 1106 (0.11%) |
| NONE | No response | 17,047 (2.3%) | 32958 (3.4%) |
| TCP | TCP connection response | 725,727 (97.6%) | 934425 (96.4%) |
| **Total** | — | **743,856 (100%)** | **968,911 (100%)** |

* ICMP = Internet Control Message Protocol; MPTCP = multipath TCP.

protocol; we're highlighting risks that arise from how things *might* be done using this protocol.

### MPTCP Adoption

Although MPTCP isn't widely deployed and long-term widespread adoption is by no means certain,[16] several open source implementations exist, and limited support has been added into some major open source network tools. Operating systems with implementations available include (but aren't limited to) Linux,[4] Berkeley Software Distribution (BSD),[17] and various Android ports. In addition, some MPTCP-aware commercial products exist, including vendors that disclosed the existence of support[18] and vendors that didn't initially disclose MPTCP support, but still offered it in their products. The most notable of the initially non-disclosing vendors is Apple, which introduced MPTCP support into the voice-command function of their iOS devices with iOS 7.[19] Currently, no support exists for Microsoft Windows platforms. Although the two stated use cases (datacenters and mobility) motivating MPTCP largely use platforms other than Windows, Microsoft's support is likely to be crucial for long-term adoption of MPTCP.

### Scanning for MPTCP Support in the Wild

Gregory Detal and his colleagues[20] developed techniques for determining whether MPTCP is actively used or otherwise supported, and they also undertook preliminary scans of key websites. As part of our research, we built our own scanner using these techniques to scan a larger number of sites, and our survey with Alexa's[21] 1 million top websites confirms that adoption of MPTCP, as of early 2014, was quite low. Changes over time are well within the bounds of likely error.

Note that our approach is intended to be indicative only, rather than an authoritative assessment. We didn't test actual MPTCP communications, but verified a rudimentary level of support in a two-stage approach.

In the first step, we simply attempted to initiate an MPTCP connection and checked the response to the handshake for multipath-capable (`MP_CAPABLE`) header options. This was exactly as undertaken by Detal and his colleagues.[20] Also, as they encountered, we had a number of likely false positives, where systems repeated TCP options back to the client blindly. Table 1 shows the results.

In the second step, we verified MPTCP support in the 1,106 hosts in the MPTCP-capable set from April 2014, using a technique that we developed previously.[22] This method involves sending an `MP_JOIN` request with an invalid connection ID. If a host understands MPTCP, then the specification states that it must respond with a TCP reset. Therefore, if a host replies with a `TCP SYNACK`, then we know it doesn't actually understand MPTCP. Only three of the 1,106 hosts actually understood MPTCP. This indicates that actual support is probably incredibly low (well below any conceivable margin for error), at least in the HTTP space.

A related technique that we used is to check if the other party replies with a different MPTCP key in the response. If the key is different, then the response probably is generated at the other end's network stack rather than a repeat of what was sent from the client.

## Key Security-Relevant Differences between MPTCP and TCP

To avoid confusion and emphasize the goals of this work, we note that we aren't discussing situations where these implications are largely irrelevant and the situation is remarkably similar to TCP. Two key situations that we omit discussing are

- the degenerate case of MPTCP, where only a single and unchanging subflow is used throughout the life of a connection; and
- situations where an adversary or administrator controls a common point among all paths (such as an end host or a common router).

However, we do present some of the network security implications associated with MPTCP. Although in isolation these may seem trivial, the combined implications lead to network traffic flows that are quite different from well-understood TCP communications.

### Network Address Decoupling

MPTCP provides connections in a way that decouples them from the ubiquitous tuple of network address and port. MPTCP connections exist independent of the network tuples their subflows use. This is problematic when network monitoring devices are unable to reassemble MPTCP traffic, but still need to correlate communications and identify the related end hosts. This implication doesn't present a significant challenge on its own, because it's mitigated when network monitoring systems can process MPTCP. However, the decoupling of the network address becomes more challenging when combined with the other items in this list.

### Network Address Nonmonotonicity

MPTCP provides the ability for endpoints to add, remove, and otherwise change network addresses mid-connection — without dropping and reestablishing a connection, as TCP does. As a result, network flows must be correlated by their connection identifiers and not by their network addresses at any single point during the connection. Similar to the aforementioned network address decoupling, this property can cause problems for network-monitoring devices that are unable to process MPTCP correctly, but related issues are also partially mitigated when network-monitoring systems can process MPTCP and use MPTCP metadata for connections to specific end systems.

### Traffic Fragmentation

MPTCP leads to situations where a network-monitoring device can only see part of traffic. This doesn't greatly differ from traditional multi-homed network configurations, except that MPTCP lets single logical communications streams be fragmented across these different network paths. This has two main scenarios: TCP port fragmentation (where traffic is broken across different TCP connections with different endpoint port combinations) and TCP address fragmentation (where it's sent from devices with different network addresses and potentially over separate networks). The correct behavior of network-monitoring devices when presented with partial content appears largely unknown, and isn't likely to be easily mitigated — even if they're able to process MPTCP.

### Connection Resilience

MPTCP provides inbuilt redundancy of communication; if a problem is detected with a path (by MPTCP-level checksums or a path that consistently terminates subflows), then that path may be avoided without dropping the overall MPTCP connection. Some network intrusion-prevention systems (NIPS) work by resetting malicious or otherwise unauthorized TCP streams. However, existing network intrusion-prevention devices aren't designed to deal with MPTCP. Unless the IPS is able to terminate the overall MPTCP connection, killing an MPTCP connection requires that every single subflow be stopped simultaneously. If any valid MPTCP subflow survives, then the communication continues, and additional subflows can be established, beginning the process anew.

## Transitional State: Current Impact of MPTCP on Network Security

MPTCP changes networking heuristics in several ways, some of which are an issue only for a transitional period (where not all devices support MPTCP), but some of which are more fundamental and require a complete rethinking of how network security is approached. MPTCP runs on existing IP-based networks, and brings risks and benefits already when few endpoints understand it; this *transitional* state is the current state on most networks. The transition presents some unique security challenges, which will remain until the vast majority of infrastructure either supports or blocks MPTCP. In particular, MPTCP brings significant risks where network operators don't have MPTCP capabilities when devices on their networks do.

### Transitional State Network Security Risks with MPTCP for Network Operators

When a network provider isn't running an MTPCP-aware infrastructure, common network security approaches that rely on full traffic visibility and traffic inspection become significantly harder in the MPTCP transitional world. Endpoints using MPTCP can evade (or unwittingly bypass) security monitoring and analysis tools, as the following instances demonstrate.

**Broken correlation.** If a monitoring or intrusion-detection device can't reassemble MPTCP correctly, then it will miss events. This enables MPTCP to be used to undertake cross-path fragmentation attacks to evade detection by network inspection devices, even if they have full visibility into the subflow's traffic (see Figure 1).

**Moving targets.** MPTCP devices can shift network addresses and ports used as often as they like in an MPTCP connection. A network security device must not only be able to reassemble traffic from different subflows, it must be able to follow related traffic across disparate networks as they add and remove subflows and network addresses during a connection.

**Active control avoidance.** MPTCP adds another level of checking for traffic modifications (MPTCP checksums) to detect alterations to communications by network devices that don't understand MPTCP. Where modification is detected along a path, that path is avoided in the future. Some security devices, such as certain types of transparent proxies and application-aware

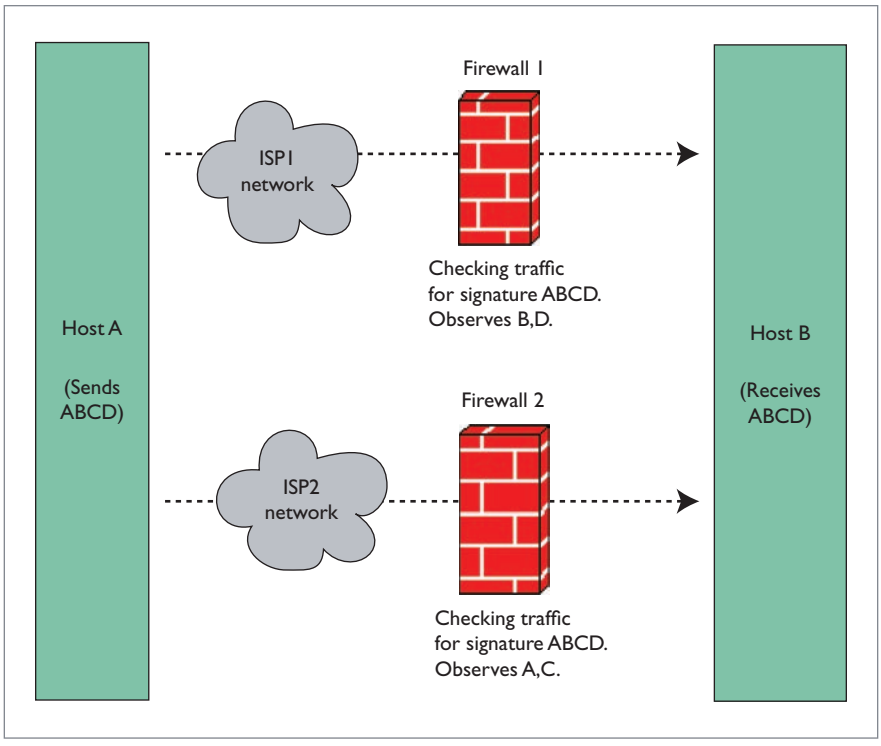


*Figure 1. MPTCP splits traffic over physical paths. If the individual stateful firewalls don't cross-correlate traffic, they won't be able to detect the content-inspection signature ABCD — even if they're MPTCP-aware. (ISP = Internet service provider.)*

intrusion-prevention systems (such as Web application firewalls), modify application data. If these middleboxes don't ensure that MPTCP metadata are correct, then they'll be bypassed trivially when the MPTCP stack detects their actions and switches to use alternative paths.

**Difficulty terminating network connections.** As we discussed previously, MPTCP's ability to route around failure can also lead to situations where MPTCP network connections are more difficult to terminate by network operators and security devices than a TCP connection is at present. Without correct MPTCP metadata, the MPTCP connection will stay open until every subflow is killed. Furthermore, backup subflows are easily overlooked by monitoring devices because these flows might not be actively creating traffic — and if missed, they'll continue the MPTCP connection seamlessly.

**Misunderstood connection direction.** MPTCP breaks the previously reliable heuristic that the sender of a TCP SYN packet is the client in a connection. Under MPTCP, a `TCP SYN` packet doesn't necessarily indicate a new connection; it simply could be a new subflow on an MPTCP connection initially established in the other direction. Although we haven't found reverse connections implemented yet, the specification allows for them. If this is implemented, then many existing firewall rules that distinguish "inbound" and "outbound" connections will make incorrect decisions, potentially resulting in legacy firewalls both accepting inbound traffic that's mistaken for outbound and outbound traffic mistaken for inbound.

## Future State: Long-Term Impact of MTPCP on Network Security

If MPTCP is widely deployed, there are some changes to network security approaches that still will be needed, even if all networking equipment is MPTCP-aware. In an IPv6 world, with a vastly increased number of addresses and networks, the scale of the available address space makes keeping track of MPTCP traffic far more difficult. How do we deal with traffic when a host may use hundreds of addresses in a single day?

At this point, we haven't identified any provable solutions to the following challenges, or addressed whether they exist at all. Nevertheless, it appears that network security must adapt to account for the challenges presented by multipath technology.

**Split traffic paths.** MPTCP causes traffic to fragment and it will often be fragmented across networks owned by different and competing organizations. Unless network security tools are able to terminate the session, they will need to make decisions on traffic with only partial content visibility. This will require the emergence of network security tools that work heuristically on malformed or fragmented content.

Network security and monitoring nodes can no longer determine traffic content by analyzing the traffic they observe locally. MPTCP gives monitoring systems a lot more to keep track of, and may require previously unnecessary (or even impossible) fine-grained coordination between network devices on different network segments, or even different organizations.

**Moving targets.** Hosts changing network addresses as they roam isn't an MPTCP-specific issue, but the ability to do it during the middle of a communication is. This address agility means that the addresses that started a communication need not be those in use at the end, or at any point along that way. This not only means that every path must be monitored, but also means that every *possible* path must be monitored. Tracking

hosts over moves is likely to be very resource-intensive.

**MPTCP and denial-of-service attacks.** MPTCP is designed to provide the ability to aggregate across all network interfaces, and this provides potentially much greater bandwidth available to each connection. DDOS prevention approaches will need to adapt, as novel multipath attacks may allow far greater traffic to flow from hosts than ever before, and in ways that are harder to respond to. Due to MPTCP's much more complicated state management and associated processing overhead, a client can potentially cause heavy load at the recipient with a lot of requests,[2] a flood of `TCP SYN` requests, or with the addition of a very large number of addresses to a connection). MPTCP also has the potential for bounce attacks, where a host is told new addresses it should connect to, such that the resulting traffic is difficult for the victim to trace back to the original attacker. This attack isn't new, and MPTCP has some mitigation against it, requiring a link be validated before sending data.

## MPTCP and Modification Detection through Timing Analysis

The integrity of user traffic against potential censorship and tampering is desirable, but the reality of the current Internet is that strong tamper detection and cryptography aren't universal. As an area for future research, we argue that it's possible to detect the modification of unprotected traffic via cross-path timing analysis. Specifically, the existence of an active attacker will be detectable because their actions will introduce an additional time delay. If an attacker is performing active interception on both paths, then it's necessary for that attacker to synchronize data between both paths before sending it on. If an attacker has control of two points that are along the target's communication path, but those two points are 100 ms apart, they can't alter the combined content for retransmission without first combining the data and sending it across that 100-ms link, thereby increasing the total time it takes to travel and introducing a potentially detectable delay.

## MPTCP and Privacy

If, and when, censorship and privacy tools (such as Tor) begin to take advantage of MPTCP,[23] governments and network providers will need to drastically change their models for mass surveillance or traffic manipulation to monitor multiple paths. Although it's possible for attackers to intercept all paths, this interception will be a more complicated undertaking.

Multipath TCP also has the potential to offer improved privacy against attackers who are able to observe or interfere with subflow traffic along a subset of paths. If a single secure path can be attained, then it might be possible to use it to make paths over other routes more secure. The spreading of traffic over multiple paths makes it less likely that attackers will get access to all of the data.[24] However, they might be able to infer some details about traffic on a path they aren't observing, such as how much data are flowing across it or what delay it introduces, with the use of protocol knowledge or statistical analyses.

Further research must be done to quantify the specifics, however; techniques appear possible to secure the connection such that if an attacker can't observe the data on *every* path, then they're unable to observe or modify the traffic on *any* path. This can happen through techniques such as sending cryptographic signing details using multiple paths, with possible approaches resembling Dirk Balfanz and his colleagues' Interactive Guy-Fawkes protocol[25] (but using different paths instead of out-of-band "location-limited" channels), or techniques such as multipath key reinforcement[26] (where keys are negotiated through communications spread over several paths). We also believe that privacy can be improved through the application of cryptographic chaining alternately across paths — for example, by using chained cipher modes such as cipher block chaining (CBC) or output feedback (OFB) — such that a person can only decrypt each block of data if he or she has access to the preceding blocks that traveled over different paths.

While the additional benefits of split-path techniques are debatable when contrasted with strong cryptography, the spreading of traffic over multiple paths (to reduce the implicit trust on any single path) is an area of research worth investigating further in the future. We argue that the existence of auxiliary MPTCP channels will enable different and better-trusted architectures for communications. Currently, both the information to be protected and the information needed to protect it (such as an encryption key or root certificates) are negotiated over the same network path. If MPTCP is used, then these are spread over multiple paths — making many existing attacks (such as intercepting HTTPS interception and replacing the certificate with an attacker one) infeasible.

MPTCP, an extension to the TCP protocol, has undergone substantial development in the past few years and promises improvement to the Internet. It enables features that align with recent technological changes (in particular, with mobility and cloud computing). Although few systems support MPTCP, most current IP networks can unwittingly let MPTCP traffic travel over them — and none of the common network intrusion-detection systems we examined are capable of inspecting it. This combination of backwards compatibility and naive

network security devices results in security threats against networks.

Although it's not possible to control a common point across all paths, such as a chokepoint or one of the endpoints, it's critical that network intrusion-detection systems and other network security systems be updated to support MPTCP reassembly. However, this isn't sufficient to maintain the level of network security commonly expected. MPTCP spreads traffic over different network paths, and existing node-based network monitors can't continue to make decisions based on only the part of the traffic data they happen to observe.

It's also worth noting that most of MPTCP's security implications are dependent upon the use of multiple diverse network paths. If an adversary, or administrator, is able to cause *all* related traffic to traverse a common point they control (such as a network bottleneck or compromised host), and to correlate it together effectively, then the differences are removed and the traffic security is essentially equivalent to the traditional, single-path TCP.

Where it's neither possible to control and monitor the endpoint's network traffic, nor bottleneck all traffic through a common point, existing content inspection approaches will be ineffective. This will require the development of new networks security systems, which either coordinate between nodes in reassembling session contents, or act probabilistically on the fragments of traffic they do happen to see.

### Acknowledgments

We express our gratitude to Michael Rabinovich, the former Editor-in-Chief of *IEEE Internet Computing,* for his pertinent comments and advice on early drafts of this article. We also thank Patrick Thomas and Gene Meltser from Neohapsis Labs for their constant support and valuable feedback and ideas throughout the preparation of this article, along with collaboration in related work. We also thank the anonymous reviewers, the Spotlight editor, and Lucy Stewart for their useful comments and suggestions.

### References

1. A. Ford, O.B.C. Raiciu, and M. Handley, *TCP Extensions for Multipath Operation with Multiple Addresses*, IETF RFC 6824, Jan. 2013; https://tools.ietf.org/html/rfc6824.
2. M. Bagnulo, *Threat Analysis for TCP Extensions for Multipath Operation with Multiple Addresses*, IETF RFC 6181, 2011, pp. 1–17; http://tools.ietf.org/html/rfc6181.
3. M. Bagnulo et al., *Analysis of MPTCP Residual Threats and Possible Fixes draft-ietf-mptcp-attacks-02*, IETF informational draft, 2014; http://tools.ietf.org/wg/mptcp/draft-ietf-mptcp-attacks/draft-ietf-mptcp-attacks-03-from-02.diff.html.
4. C. Raiciu et al., "How Hard Can It Be? Designing and Implementing a Deployable Multipath TCP," *Proc. 9th USENIX Symp. Networked Systems Design and Implementation*, 2012; www.usenix.org/conference/nsdi12/technical-sessions/presentation/raiciu.
5. C. Huitema, *Multi-Homed TCP*, IETF draft, 1995; http://tools.ietf.org/html/draft-huitema-multi-homed-01:IETF.
6. S. Barré, C. Paasch, and O. Bonaventure, "Multipath TCP: From Theory to Practice," *Proc. 10th Int'l IFIP TC6 Conf. Networking*, vol. I, 2011, pp. 444–457.
7. D. Wischik et al., "The Resource Pooling Principle," *Proc. ACM Sigcomm Computer Comm. Rev.*, vol. 38, no. 5, 2008. pp. 47–52.
8. M. Honda et al., "Is It Still Possible to Extend TCP?" *Proc. 2011 ACM Sigcomm Conf. Internet Measurement*, 2011, pp. 181–294.
9. A. Langley, *Probing the Viability of TCP Extensions*, tech. report, 2008, pp. 1–3; www.imperialviolet.org/binary/ecntest.pdf.
10. F. Kelly and T. Voice, "Stability of End-to-End Algorithms for Joint Routing and Rate Control," *Proc. ACM Sigcomm Computer Comm. Rev.*, vol. 35, no. 2, 2005, pp. 5–12.
11. P. Key, L. Massoulié, and D. Towsley, "Combining Multipath Routing and Congestion Control for Robustness," *Proc. 40th Ann. Conf. Information Sciences and Systems*, 2006, pp. 345–350.
12. H. Han et al., "Multi-Path TCP: A Joint Congestion Control and Routing Scheme to Exploit Path Diversity in the Internet," *IEEE/ACM Trans. Networking*, vol. 14, no. 6, 2006, pp. 1260–1271.
13. C. Paasch and O. Bonaventure, "Multipath TCP," *Comm. ACM*, vol. 57, no. 4, 2014, pp. 51–57.
14. B. Radunović et al., "Horizon: Balancing TCP over Multiple Paths in Wireless Mesh Network," *Proc. 14th ACM Int'l Conf. Mobile Computing and Networking*, 2008, pp. 247–258.
15. C. Raiciu et al., "Improving Datacenter Performance and Robustness with Multipath TCP," *Proc. ACM Sigcomm Computer Comm. Rev.*, vol. 41, no. 4, 2011, pp. 266–277.
16. A. Kostopoulos et al., "Towards Multipath TCP Adoption: Challenges and Opportunities. Next Generation Internet (NGI)," *Proc. 2010 6th EURO-NF Conf.*, 2010, pp.1–8.
17. N. Williams, L. Stewart, and G. Armitage, *Design Overview of Multipath TCP Version 0.3 for FreeBSD-10*, CAIA tech. report 130424A 2013; http://caia.swin.edu.au/reports/130424A/CAIA-TR-130424A.pdf.
18. J. Young and D. Wing, *MPTCP and Product Support Overview*, tech. report, Cisco, 2013; www.cisco.com/c/en/us/support/docs/ip/transmission-control-protocol-tcp/116519-technote-mptcp-00.pdf.
19. I. Apple, *iOS: Multipath TCP Support in iOS 7*, tech. report, Apple, 2014; http://support.apple.com/kb/HT5977.
20. G. Detal et al., "Revealing Middlebox Interference with Tracebox," *Proc. 2013 Conf. Internet Measurement*, 2013, pp. 1–8; http://dx.doi.org/10.1145/2504730.2504757.
21. Alexa, *Top 1-Million Websites,* 2013; http://s3.amazonaws.com/alexa-static/top-1m.csv.zip.
22. C. Pearce and P. Thomas, "Multipath TCP: Breaking Today's Networks with Tomorrow's Protocols," *Proc. Blackhat Briefings*, 2014; www.blackhat.com/docs/us-14/materials/us-14-Pearce-Multipath-TCP-Breaking-Todays-Networks-With-Tomorrows-Protocols-WP.pdf.
23. H.T. Karaoglu et al., "Multi Path Considerations for Anonymized Routing: Challenges and Opportunities," *Proc. 5th Int'l Conf. New Technologies, Mobility and Security*, 2012, pp. 1–5.

24. J. Yang and S. Papavassiliou, "Improving Network Security by Multipath Traffic Dispersion," *Proc. IEEE Military Comm. Conf. for Network Centric Operations: Creating the Information Force*, 2001, pp. 34–38.
25. D. Balfanz et al., "Talking to Strangers: Authentication in Ad-Hoc Wireless Networks," *Proc. 9th Ann. Symp. Network and Distributed System Security*, 2002; www.isoc.org/isoc/conferences/ndss/02/papers/balfan.pdf.
26. H. Chan, A. Perrig, and D. Song, "Random Key Predistribution Schemes for Sensor Networks," *Proc. IEEE Symp. Security and Privacy*, 2003, pp. 197–213.

**Catherine Pearce** is a senior security consultant at Neohapsis (now a part of Cisco). Her technical and research interests include application security and unexpected security-related interactions between technologies. Pearce has an MS in computer science from the University of Canterbury. Contact her at katpearc@cisco.com.

**Sherali Zeadally** is an associate professor in the College of Communication and Information, University of Kentucky. Zeadally has a PhD in computer science from the University of Buckingham, England. He's a Fellow of the British Computer Society and the Institution of Engineering Technology, England. Contact him at szeadally@uky.edu.

*This article originally appeared in* IEEE Internet Computing, *vol. 19, no. 5, 2015.*

# The Research Horizon: Four Nearly Practical Concepts

**Hilarie Orman** • *Purple Streak*

Does the world of academic research have anything to contribute to the world of practical security? That was the question I asked myself recently while attending the annual IEEE Symposium on Security and Privacy. This year there were 55 papers, a record number. At least one of them was considered practical, because it won the "Distinguished Practical Paper" award. I found at least three other papers that seemed both practical and interesting: ways of using a hybrid capability architecture for a variety of process protection mechanisms, a near-term solution to privacy of DNS queries, and a methodology that turned up several exploitable errors in Transport Layer Security (TLS) state machine implementations. That last paper was one of two selected for the overall "Distinguished Paper" award.

In some sense, almost all security research papers are practical. They must surmount a high bar by covering new ground. Solutions that haven't been tried before, attacks that haven't been analyzed, and new ways to analyze existing systems are all aspects of the collection of papers for a good security research conference.

Today's research papers usually address a real-world problem in some way. The researchers will apply analysis to gain a deeper understanding of the problem, or they'll investigate a proposed solution's effectiveness. Some solutions aren't practical due to computational overhead or the amount of infrastructure that would have to be changed, but generally the paper is about something that might be used in the foreseeable future. Nonetheless, some things are more practical than others, and some things are more important than others. The four papers I have selected to discuss in this issue's "Practical Security" column occupy different places on the axes of practicality, novelty, and importance.

## Practically Competing for Attention

The first paper is "Ad Injection at Scale: Assessing Deceptive Advertisement Modifications," and it's an investigation into how some ads get placed on webpages.[1] Webpages are a digital battleground when it comes to ads. There's a complex ecosystem feeding that war, and part of the system relies on some shady characters who use browser extensions to turn the ad selection system upside down.

When visiting a retailer's webpage, a user expects to see ads for products sold through the site. Similarly, when visiting a search results page, the user understands that some of the entries are paid for by advertisers. The retailers might think they're paying for an ideal kind of tailored shopping experience — a thoughtful selection of products suited to the shopper's preferences and budget and current best deals offered by the retailer. This is unlikely to be the experience for users who have installed browser extensions that promise an enhanced shopping experience.

A user might visit Walmart's website looking for a TV and see ads for Crazy Ed's Electronics, as a hypothetical example. Those ads weren't there when the Walmart server composed the page, but by the time the page is displayed in an afflicted user's browser, the entire process of selecting and inserting ads has been repeated by code that was installed as a browser extension. The extensions often are touted as advantageous for the user. In theory, they automatically do product searches and price comparisons. In practice, they're a nuisance because they slow down all page loads and don't deliver the promised value. On top of that, they can be almost impossible to remove.

The researchers were able to learn a great deal about the prevalence of ad injectors in the wild by utilizing their association with Google. Their methodology inserted monitoring code

 Published by the IEEE Computer Society 

into webpages delivered by Google servers, but the code did nothing to change what was displayed. Any ads that were inserted after the page was delivered were undisturbed. The monitoring software waited for the "unload" event generated when the user navigated away from the page, and at that point the page's contents were compared to what was originally delivered. In this way, they learned that a few percent of all webpage deliveries are affected by ad injectors. The injectors affect many kinds of browsers, notably Chrome, Firefox, and Internet Explorer.

The research project then moved on to investigating the variety of software libraries that implement ad injection, the software providers, and the revenue stream that supports the practice. Unsurprisingly, only a few entities dominate ad replacement technology: Superfish, JollyWallet, VisAdd, and Nav-Links. Superfish is notorious because of its short-lived deal with Lenovo to preinstall its ad software on Windows machines, thus opening a security hole.[2] The researchers found that the advertisers who work with the ad injection companies form a tangled web, sometimes looping back to the unknowing retailer, who can end up paying for his own ads to be replaced with different ads of his own!

An interesting side note on this subject is in a different paper from the IEEE Symposium on Security and Privacy, one that detects and analyzes the function of embedded scripts in webpages.[3] I wonder if the research behind it might have come across the monitoring scripts used in the ad injection research.

While you have to admire the depth of the investigation and interesting information about the danger of browser extensions, the work leaves me wondering, Where's the security, where's the practicality? Web browsers are applications with unlimited extensibility and no security model, and it seems obvious that they're vectors for unpleasant surprises. Is ad replacement a security problem, or is it really a tug-of-war between competing ad companies? Is the user harmed by the ad replacements? Could there be ad replacement services that are mutually beneficial to the user and the website owner? It's likely that the war for user attention is only beginning, and this article won't be the last word on the subject.

## Practical Privacy for DNS Queries

Web browsing is critically dependent on DNS lookups, a little piece of magic that turns a name like Walmart.com into an Internet address. Although the authenticity of the DNS response can be assured through cryptography, and although the webpage that you ultimately access may be delivered with privacy-preserving encryption, the DNS query itself is in plain view of all network-observing eyes. Can this privacy gap be closed?

Researchers at Verisign Labs and USC's Information Sciences Institute teamed up to validate the assertion that it's eminently feasible and practical to use TLS to gain this privacy.[4] Their work shows that the fears of server overload can be allayed through careful connection management, a few implementation changes, and a protocol extension. Their conclusions are based on analysis of a huge sample of DNS queries to three different, heavily used, servers: a proxy resolver, an ISP DNS server, and a root nameserver.

TLS runs over TCP, but DNS has historically used User Datagram Protocol (UDP). The connectionless nature of UDP originally seemed suited to the simplicity of the query/reply nature of DNS. But, there are downsides to this simplicity, one of them being an "amplification attack" in which malicious senders use their victim's IP address to trick DNS servers into sending large replies that overwhelm the victim's network capacity. There's evidence that most denial-of-service attacks utilize DNS amplification. TCP and TLS together thwart a good number of these simple attacks, giving yet another reason to switch away from UDP.

The TLS approach is an interesting contrast to an older proposal for DNS query privacy[5] that works over UDP, but has been slow to catch on. That method seems roughly comparable to this new proposal in terms of performance.

There are two nonobvious implementation changes that the researchers used to argue that TCP is a good choice for a transport protocol. Traditional query/response systems process requests in order, but this blocks the connection until the response is ready. If instead the server delivers replies when they're ready, even if they're out of order, the connection can be kept open without blocking. Because DNS already has provisions for matching responses to queries, this change isn't onerous for clients. The clients can improve things further by sending the queries without waiting for responses; the typical webpage has at least four unique domain names, and all those queries can be sent in one batch. Finally, clients could help servers avoid the expense of tearing down connections and restarting them by advising the server that it would be a good idea to extend the timeout period.

Based on the data available, the researchers feel that two other modifications would help with the rather expensive process of starting up a new TLS connection for DNS. They propose adding a new bit to the extensions vector to indicate that the client wants to use TLS immediately. They also recommend adopting the strategy of letting the server bundle the TLS session state into an opaque object that can be stored by the client while the session is closed. The client could request session resumption

by sending the opaque object back to the server, and both could then restore their cryptographic contexts securely.

How practical is all of this? The paper is based on traffic analyses, not actual implementations, but it suggests that the time to make this a practical reality is at hand. There's even a working group for that: the DNS PRIVate Exchange (dprive).[6]

## A Practical Approach to Finding TLS State Machine Errors

The TLS security protocol is widely used for authenticating and protecting connections to webservers. Briefly, it's the lock icon shown next to the address bar in a browser. The protocol has also been the subject of many detailed analyses and the source of many bugs.[7] Problems don't come from the protocol's formal definition, which has been analyzed extensively, but from the implementations that are burdened with legacy functionality, cipher suite variability, and complex error handling. Practically speaking, how can we vet a TLS implementation for correct handling of everything thrown at it?

The problem is especially difficult because there's not just one sequence of messages between a client and server that sets up the cryptographic context for a TLS session. If there were, it would be easy to test for correctness. At each state, a testing system would try sending one of the 10 or so messages associated with a different, out-of-order state transition. But TLS isn't that simple, and there are multiple valid pathways that weave through its state transition graph. Not all implementers coded this correctly, which led to a number of exploitable vulnerabilities.

Can all the errors in TLS implementations be found and corrected? Perhaps not all of them, but Benjamin Beurdouche and his colleagues[8] have an innovative testing methodology that automatically detects problems in real-life implementations of TLS. They probe at an implementation (either a client or a server) by sending messages out of order and looking for responses that fail to "alert" the bad message. This was a surprisingly fruitful investigation. For example, they found that one implementation incorrectly allowed some cryptographic setup messages to be skipped. In one case, this omission disabled subsequent data encryption.

One of the discoveries was an exploitable downgrade attack that has been named "FREAK." The problem comes about because some client implementations incorrectly accept a key exchange message, even though they already have the server's RSA key from its certificate. The key exchange message can have a very weak RSA key (from the days when US export controls required weakness in exported cryptographic products). As a result, a man-in-the-middle (MITM) can factor the key and read the traffic. In fact, for many implementations, the MITM can read the traffic for several days, because the key is changed infrequently.

A second part of the research produced a verified implementation of the complete TLS state machine. Their implementation requires annotations on critical memory areas to assure that they're well-defined and not overlapping. This verified implementation could be used within another implementation as a check that each protocol message left the implementation in a valid state.

As practical as this seems to be, as a way of locking down the TLS state machine, the authors note that the whole protocol, including all the cryptography, is unlikely to be verified anytime soon. Still, we can only applaud any method that finds bugs for the good guys to fix before the bad guys find them.

## Capabilities Made (Almost) Practical

Finally, wouldn't it be nice if the computer hardware protected processes from all unauthorized data accesses, rendered buffer overflow bugs innocuous, and generally kept us all safe? Of course it would, and that idea is behind the venerable capability machine which has long been the Holy Grail of the secure processing community.

Capability machines use special hardware to mediate interactions between different "domains" (roughly speaking, processes and/or libraries). A *capability architecture* can securely implement the principle of "least privilege" for trusted computing bases. A capability is an immutable system object that we can use for resource access. For example, we could control read access to a section of memory from another process through a capability rather than a bare memory pointer or kernel-mediated interprocess communication. Capabilities are also useful for controlling access to resources such as files or network connections.

The problem that has beset capability machines from their outset is that they're slow. Compared to traditional processors of similar instruction sets, they seem like sack race contestants trying to run Olympic sprints. That slowness comes from the number of independent address space descriptors needed to represent a capability-based process. Each one requires a memory-page mapping. The translation from address to page map index is done via the fast associative memory cache structure called the *translation lookaside buffer* or TLB (which is generally as well understood as a spleen). If there are more active entries than can fit in the TLB, memory lookups go through a much slower lookup mechanism that makes a modern computer seem to be running in low gear.

Undissuaded by this history, a large research team has embarked on the Capability Hardware-Enhanced RISC

*This article originally appeared in* IEEE Internet Computing, *vol. 19, no. 5, 2015.*

Instructions (CHERI) project to "do capabilities right" on their Business Environment Risk Intelligence (BERI) processor, a RISC field-programmable gate array (FPGA) softcore processor (open sourced). Their recent work[9] builds on a capability architecture, but sidesteps some of its associative memory limitations by way of a hybrid approach that selectively integrates a C compiler with software capabilities on a FreeBSD (BSD stands for Berkeley Software Distribution) operating system. If that sounds familiar, it's because the capabilities of FreeBSD were the subject of work published in 2010.[10] That paper investigated the "capabilitization" of common utilities like `tcpdump` and `zlib`.

The work described at the symposium is an interesting architecture that lets a developer choose security/performance tradeoffs in a capability architecture. At one end of the spectrum, everything could be a capability: all library accesses, all memory pointers, all interprocess communication. That would be prohibitively expensive, as noted in the past. That's why CHERI supports intermediate solutions that can set the protection boundaries in several different ways. Capabilities can be used to wrap an untrusted library operating within an otherwise unmodified application, for example. Or, a capability-based process can access "compartmentalized" legacy C code. The compartment limits the damage from buffer overflows or improper addressing to the resources of the software and data within the compartment's address space.

The performance results are greatly favorable to the capabilities when used correctly, as opposed to sandboxing. Interprocess memory copies, for example, are much speedier with a capability than with sandboxed processes. The `tcpdump` utility is an interesting case study, because it has known bugs when handling some kinds of malformed packets. CHERI is flexible enough to support isolation based on packet types or addresses, packet processing time, or other criteria. An error in a compartment is limited to damage within its own scope. These techniques mitigated all but two known vulnerabilities in the utility.

The CHERI work isn't ready for immediate use, and their FPGA processor isn't going to be fabricated for use in mobile devices anytime soon. Yet the work does show that stronger protections for real-world legacy code could be a reality in a future world that values secure processing. As a practical matter, security-enhanced processors still seem a distant hope, unless there's some overwhelmingly lucrative use for them. This seems a bit odd, because the estimated losses from cybercrime are astounding.

There's no crystal ball that tells us when research will turn into practice. Although computer security research is very much an applied science (I'm reminded of Edsger Dykstra's cutting remark that computer science is no more a science than surgery is "knife science"), and even though most researchers would be delighted to have their work used commercially, the pathway from the research lab to the sales-in-the-billions product is unpredictable. Economics, physics, and psychology must all align perfectly. The four examples presented here are novel approaches expanding computer security, and each of them represent ongoing work that could be in your hands within a few years. They're only a small sample of what comes out of academic and industrial research labs every year, so keep your eye out for the gems that might be the next big influence on security.

## References


1. K. Thomas et al., "Ad Injection at Scale: Assessing Deceptive Advertisement Modifications," *Proc. IEEE Symp. Security and Privacy*, 2015, pp. 151–167.
2. S. Rosenblatt, "Lenovo's Superfish Security Snafu Blows Up in Its Face," *CNET*, 20 Feb. 2015; www.cnet.com/news//superfish-torments-lenovo-owners-with-more-than-adware.
3. Y. Zhou and D. Evans, "Understanding and Monitoring Embedded Web Scripts," *Proc. Symp. IEEE Security and Privacy*, 2015, pp. 850–865.
4. L. Zhu et al., "Connection-Oriented DNS to Improve Privacy and Security," *Proc. IEEE Symp. Security and Privacy*, 2015, pp. 171–186.
5. M. Dempsky, *DNSCurve: Link-Level Security for the Domain Name System*, IETF draft, 26 Feb. 2010; http://tools.ietf.org/html/draft-dempsky-dnscurve-01.
6. DNS PRIVate Exchange (dprive), *Charter for Working Group*, 2015; https://datatracker.ietf.org/wg/dprive/charter.
7. C. Meyer and J. Schwenk, *Lessons Learned from Previous SSL/TLS Attacks – A Brief Chronology of Attacks and Weaknesses*, Int'l Assoc. for Cryptologic Research (IACR) eprint archive, 2013; http://eprint.iacr.org/2013/049.
8. B. Beurdouche et al., "A Messy State of the Union: Taming the Composite State Machines of TLS," *Proc. Symp. IEEE Security and Privacy*, 2015, pp. 535–552.
9. R.N.M. Watson et al., "CHERI: A Hybrid Capability-System Architecture for Scalable Software Compartmentalization," *Proc. IEEE Security and Privacy Symp.*, 2015, pp. 20–37.
10. R.N.M. Watson et al., "Capsicum: Practical Capabilities for Unix," *Proc. 19th Usenix Security Symp.*, 2010; www.usenix.org/legacy/event/sec10/tech/full_papers/Watson.pdf.


**Hilarie Orman** is a security consultant and president of Purple Streak. Her research interests include applied cryptography, secure operating systems, malware identification, security through semantic computing, and personal data mining. Orman has a BS in mathematics from the Massachusetts Institute of Technology. She's a former chair of the IEEE Computer Society's Technical Committee on Security and Privacy. Contact her at hilarie@purplestreak.com.

Editors: Michael Lesk, lesk@acm.org | Jeffrey MacKie-Mason, jmm@umich.edu

# Scaring and Bullying People into Security Won't Work

**Angela Sasse** | University College London

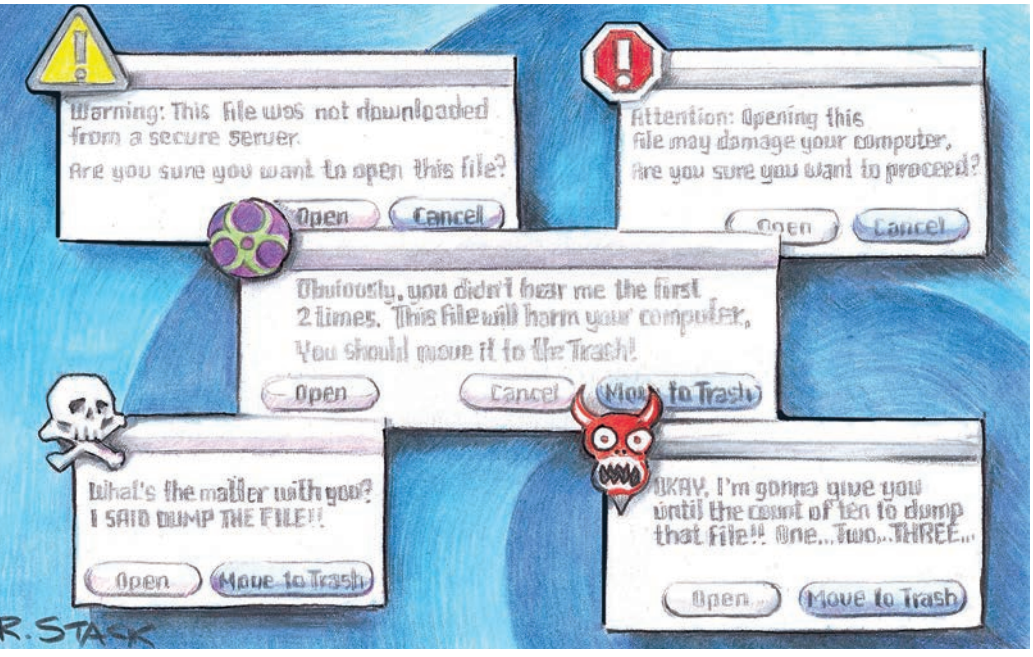


Usable security and privacy research began more than 15 years ago. In 1999, Alma Whitten and J.D. Tygar explained "Why Johnny Can't Encrypt,"[1] and Anne Adams and I pleaded that, even though they don't always comply with security policies, "Users Are Not the Enemy."[2] Today, there are several specialist conferences and workshops: publications on usable security and privacy are featured in top usability conferences, such as the ACM SIGCHI Conference on Human Factors in Computing Systems (CHI), and top security conferences, such as the IEEE Symposium on Security and Privacy.

An ongoing topic in usable security research is security warnings. Security experts despair that the vast majority of users ignore warnings—they just "swat" them away, as they do with most dialog boxes. Over the past six years, continuous efforts have focused on changing this behavior and getting users to pay more attention. SSL certificate warnings are a key example: all browser providers have evolved their warnings in an attempt to get users to take them more seriously. For instance, Mozilla Firefox increased the number of dialog boxes and clicks users must wade through to proceed with the connection, even though it might not be secure. However, this has made little difference to the many users who decide to ignore the warnings and proceed. Creating more elaborate warnings to guide users toward secure behavior is not necessarily the best course of action, as it doesn't align with the principles of user-centered design.

## Refining Warnings

At ACM CHI 2015, two studies reported on efforts to make more users heed warnings. Adrienne Porter Felt and her colleagues at Google designed a new SSL warning for Google Chrome, applying recommendations from current usable security research: keep warnings brief, use simple language to describe the specific risk, and illustrate the potential consequences of proceeding.[3] The authors hypothesized that if users understand the risks associated with a warning, they will heed rather than ignore it.

They tested these improved warnings in a series of mini surveys and found a modest but significant (12 percent) improvement in the number of participants who correctly identified the potential risks of proceeding, but no significant improvement in the number of participants who correctly identified the data at risk. In addition, compared to existing browser SSL warnings, there was no improvement in the number of participants who thought the warning was likely to be a false positive.


Published by the IEEE Computer Society

Felt and her colleagues reasoned that if they couldn't improve users' understanding, they might still be able to guide users toward secure choices. They applied what they called *opinionated design* to make it harder for participants to circumvent warnings, and visual design techniques to make the secure course of action look more attractive. In a field study, this technique led to a 30 percent increase in the number of participants who didn't proceed upon seeing the warning. The authors concluded that it's difficult to improve user comprehension of online risks with simple, brief, nontechnical, and specific warnings, yet they urge fellow researchers to keep trying to develop such warnings. In the meantime, they advise designers to use opinionated design to deter users from proceeding in the face of warnings by making them harder to circumvent and emphasizing the risks associated with doing so.

In the second paper, Bonnie Anderson and her colleagues examined 25 participants' brain responses to warnings using a functional magnetic resonance imaging (fMRI) scanner.[4] Previous studies using eye tracking showed that users habituate: the first time around, a warning catches their attention, but after repeated showings, it does not. Anderson and her colleagues found that the brain mirrors this habituation: when encountering a warning for the first time, participants' visual processing center in the superior parietal lobes showed elevated activation levels, but these disappeared with repeated showings of the warning.

The authors hypothesized that varying a warning's appearance, such as its size, color, and text ordering, should prevent habituation and keep participants paying attention. They found that participants indeed showed sustained activation levels when encountering these polymorphic warnings; participants' attention decreased on average only after the 13th variation of the same warning. They concluded that users can't help but habituate, and designers should combat this by creating warnings that force users to pay attention.

## Usability: When Does "Guiding" Become "Bullying"?

Both teams' work was motivated by an honorable intention—to help users choose the secure option. But as a security researcher with a usability background and many years of studying user behavior in the lab as well as in real-world settings, I am concerned by the suggestion that we should use design techniques to force users to keep paying attention and push them toward what we deem the secure—and hence better—option. It is a paternalistic, technology-centered perspective that assumes the security experts' solution is the correct way to manage a specific threat.

In the case of SSL, the authors recommended counteracting people's habituation response and keeping their attention focused on security. However, habituation is an evolved response that increases human efficiency in day-to-day interactions with the environment: we stop paying attention to signals we've deemed irrelevant. Crying wolf too often leads to alarm or alert fatigue; this has been demonstrated over many decades in industries such as construction and mining and, most recently, with the rapid increase of monitoring equipment in hospitals.

In 2013, the US Joint Commission issued an alert about the widespread phenomenon of alarm fatigue.[5] The main problem was desensitization to alarms, which led to staff missing critical events. An increase in workload and decrease in patient satisfaction were also noted.

Eminent software engineer and usability expert Alan Cooper identified the use of warnings in software as a problem more than a decade ago.[6] He pointed out that warnings should be reserved for genuine exceptions—events software developers couldn't reasonably anticipate and make provisions for. Perhaps on their legal advisors' suggestion, most developers have ignored Cooper's recommendation, and the increasing need for security has led to a marked increase in the number of dialog boxes or warnings that users have to "swat" away today.

Strategies such as opinionated design and forcibly attracting users' attention do not align with usability. As Cooper pointed out, usability's overall guiding principle is to support users in reaching their primary goals as efficiently as possible. Security that routinely diverts the attention and disrupts the activities of users in pursuit of these goals is thus the antithesis of a user-centered approach.

And where, in practical terms, would this approach lead us? A colleague with whom I discussed the studies commented: "Even with this polymorphic approach, users stop paying attention after 13 warning messages. I suppose the next step is to administer significant electrical shocks to users as they receive the warning messages, so that they are literally jolted into paying attention." (The colleague kindly allowed me to use the quote, but wishes to remain anonymous.) Scaring, tricking, and bullying users into secure behaviors is not usable security.

## Cost versus Benefit

In 2009, Turing award and von Neumann medal winner Butler Lampson pointed out that[7]

> [t]hings are so bad for usable security that we need to give up on perfection and focus on essentials. The root cause of the

> problem is economics: we don't know the costs either of getting security or of not having it, so users quite rationally don't care much about it. … To fix this we need to measure the cost of security, and especially the time users spend on it.

Lampson's observations haven't been heeded. User time and effort are rarely at the forefront of security studies; the focus is on whether users choose the behavior that researchers claim to be desirable because it's more secure. Even if users' interaction time with specific security mechanisms, such as a longer password, is measured, the cumulative longer-term effect of draining time from individual and organizational productivity isn't considered.

Over the past few years, researchers have declared the task of recalling and entering 15- or 20-character complex passwords "usable" because participants in Mechanical Turk studies were able to do so. But being able to do something a couple of times in the artificial constraints of such studies doesn't mean the vast majority of users could—or would want to—do so regularly in pursuit of their everyday goals.

Factors such as fatigue as well as habituation affect performance. In real-world environments, authentication fatigue isn't hard to detect: users reorganize their primary tasks to minimize exposure to secondary security tasks, stop using devices and services with onerous security, and don't pursue innovative ideas because they can't face any more "battles with security" that they anticipate on the path to realizing those ideas.[8] It's been disheartening to see that, in many organizations, users who circumvent security measures to remain productive are still seen as the root of the problem—"the enemy"[2]—and that the answer is to educate or threaten them into behavior security experts demand—rather than considering the possibility that security needs to be redesigned.

A good example is the currently popular notion that sending phishing messages to a company's employees, and directing them to pages about the dangers of clicking links, is a good way to get their attention and make them less likely to click in the future. Telling employees not to click on links can work in businesses in which there's no need to click embedded links. But if legitimate business tasks contain embedded links, employees can't examine and ponder every time they encounter a link without compromising productivity.

In addition, being tricked by a company's own security staff is a negative, adversarial experience that undermines the trust relationship between the organization and employees. Security experts who aim to make security work by "fixing" human shortcomings are ignoring key lessons from human factors and economics.

In modern, busy work environments, users will continue to circumvent security tasks that have a high workload and disrupt primary activities because they substantially decrease productivity. No amount of security education—a further distraction from primary tasks—will change that. Rather, any security measure should pass a cost–benefit test: Is it easy and quick to do, and does it offer a good level of protection?

Cormac Herley calculated that the economic cost of the time users spend on standard security measures such as passwords, antiphishing tools, and certificate warnings is billions of dollars in the US alone—and this when the security benefits of complying with the security advice are dubious.[9] SSL warnings have an overwhelming false-positive rate—close to 100 percent for many years[9]—so users developed alarm fatigue and learned to ignore them. In addition, longer (12- to 15-character) passwords, which are associated with a very real cost in recall and entry time and increased failure rates—especially on the now widely used touchscreens—offer no improvement in security.[10]

## Fitting the Task to the Human

The security-centered view assumes that users want to avoid risk and harm altogether. However, many users choose to accept some risks in pursuit of goals that are important to them. Security experts assume that users who don't choose the secure option are making a mistake, and thus preventing mistakes and educating users is the way forward.

However, a combination of usability and economics insights leads to a different way of thinking about usable security:

- Usable security starts by recognizing users' security goals rather than by imposing security experts' views on users.
- Usable security acknowledges that users are focused on their primary goals—for example, banking, shopping, or social networking. Rather than disrupting these primary tasks and creating a huge workload for users, security tasks should cause minimum friction.
- Security experts must acknowledge and support human capabilities and limitations. Rather than trying to "fix the human," experts should design technology and security mechanisms that don't burden and disrupt users.

Techniques from the human factors field can maximize performance while ensuring safety and security. A key principle is designing technology that fits users' physical and mental abilities—fitting the

task to the human. Rarely should we fit the human to the task, because this requires significant organizational investment in terms of behavior change through education and training. Security education and training are only worthwhile if the behavior fits with primary tasks. An organization could train its employees to become memory artists, enabling them to juggle a large number of changing PINs and passwords. But then employees would need time for routines and exercises that reinforce memory and recall.

Changing security policies and implementing mechanisms that enable employees to cope without training are more efficient. For instance, Michelle Steves and Mary Theofanos recommend a shift from explicit to implicit authentication[8]; in most environments, there are other ways to recognize legitimate users, including device and location information or behavioral biometrics, without disrupting users' workflow. They also point out that infrequent authentication requires different mechanisms that complement the workings of human memory—something Adams and I recommended after our first study 15 years ago[2]—but this rarely occurs in practice.

Users will pay attention to reliable and credible indicators of risks they want to avoid. Security mechanisms with a high false-positive rate undermine the credibility of security and train users to ignore them. We need more accurate detection and better security tools if we are to regain users' attention and respect, rather than scare, trick, and bully them into complying with security measures that obstruct human endeavors. ■

## References


1. A. Whitten and D. Tygar, "Why Johnny Can't Encrypt: A Usability Evaluation of PGP 5.0," *Proc. 8th Conf. USENIX Security Symp.*, vol. 9, 1999, p. 14.
2. A. Adams and M.A. Sasse, "Users Are Not the Enemy," *Comm. ACM*, vol. 42, no. 12, 1999, pp. 40–46.
3. A. Porter Felt et al., "Improving SSL Warnings: Comprehension and Adherence," *Proc. Conf. Human Factors and Computing Systems*, 2015; https://adrifelt.github.io/ssl interstitial-chi.pdf.
4. B.B. Anderson et al., "How Polymorphic Warnings Reduce Habituation in the Brain—Insights from an fMRI Study," *Proc. Conf. Human Factors and Computing Systems*, 2015; http://neurosecurity.byu.edu/media/Anderson_et_al._CHI_2015.pdf.
5. "Medical Device Alarm Safety in Hospitals," *Sentinel Event Alert*, no. 50, 8 Apr. 2013; www.pwrnewmedia.com/2013/joint_commission/medical_alarm_safety/downloads/SEA_50_alarms.pdf.
6. A. Cooper, *The Inmates Are Running the Asylum: Why High-Tech Products Drive Us Crazy and How to Restore the Sanity*, Sams–Pearson, 2004.
7. B. Lampson, "Usable Security: How to Get It," *Comm. ACM*, vol. 52, no. 11, 2009, pp. 25–27.
8. M.P. Steves and M.F. Theofanos, *Report: Authentication Diary Study*, tech. report NISTIR 7983, Nat'l Inst. Standards and Technology, 2014.
9. C. Herley, "So Long, and No Thanks for the Externalities: The Rational Rejection of Security Advice by Users," *Proc. 2009 Workshop New Security Paradigms*, 2009, pp. 133–144.
10. D. Florencio, C. Herley, and P.C. van Oorschot, "An Administrator's Guide to Internet Password Research," *Proc. USENIX Conf. Large Installation System Administration*, 2014, pp. 35–52.



**Angela Sasse** is a professor of human-centered technology at University College London. Contact her at a.sasse@cs.ucl.ac.uk.



*This article originally appeared in* IEEE Security & Privacy, *vol. 13, no. 3, 2015.*

by Charles Day
American Institute of Physics
cday@aip.org

# Days of Endless Time

*This article originally appeared in Computing in Science & Engineering, vol. 17, no. 4, 2015.*

One cold and bright Sunday in early spring, I left my house in the Capitol Hill neighborhood of Washington, DC, and walked to the National Mall. My destination was *Days of Endless Time*, an exhibition of video installations at the Hirshhorn Museum. The 14 works on display—to quote the program—"emphasize slower, more meditative forms of perception."

The first work that visitors encounter is Su-Mei Tse's *L'Echo* (2003). A wall-sized projection shows a woman, Tse herself, dressed in red and playing a cello with her back to the viewer. She's dwarfed by the sheer dark gray face of a distant cliff and by the expanse of bright green grass on which her stool rests. For the five minutes of the video, we hear the slow contemplative music that she plays echoed by the cliff that dominates the frame.

Despite its power to evoke both awe and tranquility, *L'Echo* was surpassed in impact by two works that relied heavily on computational image processing. *Afterimage* (2013) by Clemens von Wedemeyer consists of a large semicircular screen displaying what seems at first to be a video tour of a warehouse crammed with movie props. As the six-minute video plays out, the camera pans over the closely packed statues of Roman emperors, busts of medieval popes, and so on. But the camera also appears to pass through walls that dissolve as a new room is entered. Some of the props move.

At 12 minutes, David Claerbout's *Travel* (1996–2013) was the longest work in the exhibition. According to the caption,

> the artist spent years animating archetypes from nature, both sublime and quotidian, to produce a slow, solitary journey through an imaginary landscape. *Travel* incorporates painterly precision and a synthesis of "therapeutic music," and was inspired by vistas in France, Luxembourg, and Belgium. Evoking everywhere and nowhere, the work presents a wealth of visual detail as it conforms to well-worn imagery of the forest, jungle, and suburbs.

The effect of *Travel* is both mesmerizing and eerie. The foliage is rendered with near photographic accuracy, yet even the lightest leaves remain utterly still. And the camera—or, rather, the virtual camera—moves among the trees at the height of a toddler and with perfectly smooth tracking. No insects, birds, or other animals are present—just plants.

Video games, whose computer-animated worlds are increasingly intricate and convincing, might seem the diametric opposite of the works in *Days of Endless Time*, both in character and purpose. But not all video games are fast-paced, violent romps. Claerbout's imaginary landscapes reminded me of the scenery in *The Vanishing of Ethan Carter*, a mystery game for PCs and PlayStation released last year to much critical acclaim.

If a computer-generated video game can be slow and tranquil, then computer-generated art can be fast and exciting. Much as I enjoyed *Days of Endless Time*, I would love to see an exhibition whose works matched the intensity of three other acclaimed video games of 2014: *Titanfall*, *Dragon Age: Inquisition*, and *Civilization: Beyond Earth*. ■

**Charles Day** is the online editor at *Physics Today*.

 Published by the IEEE Computer Society 2376-113X/15/$31.00 

Editor: Gerard J. Holzmann
NASA/JPL
gholzmann@acm.org

# Code Evasion

Gerard J. Holzmann

**MOST OF US** have trouble reading other people's code, or even our own code if we haven't looked at it for a while. The art of writing more or less self-documenting code that any competent programmer can understand effortlessly is all too rare. Why?

One reason programs tend to lose their structure and clarity as they move from concept to product is the addition of error handling. Often, more than half of a code base ends up dedicated to various types of error detection and recovery, obscuring the nominal control flow that defines the basic structure.

```
if ((buf = malloc(N)) == NULL) {
    fprintf(stderr, "out of memory\n");
    exit(1);
}
if ((fd = fopen(fnm, "r")) == NULL) {
    fprintf(stderr, "cannot open %s\n", fnm);
    exit(2);
}
if (read(fd, buf, N) != N) {
   fprintf(stderr, "read error\n");
   exit(3);
}
```

(a)

```
buf = malloc(N);
fd = fopen(fnm, "r");
read(fd, buf, N);
```

(b)

**FIGURE 1.** A fragment of C code that allocates memory to a buffer, opens a file for reading, and reads data from that file into the buffer. (a) Typical code obfuscation. (b) The nominal path that is hiding in Figure 1a. The second version is a lot easier to read than the first, although it wouldn't be usable in this form.

## Obscuring the Flow

Consider a typical fragment of C code that allocates memory to a buffer, opens a file for reading, and reads data from that file into the buffer. Taking care of the possibly failing operations, that code would look something like Figure 1a. Hiding in these 14 lines of code is a small nominal execution path of just three statements, which looks like Figure 1b.

Clearly, the second version is a lot easier to read than the first, although it wouldn't be usable in this form. Figure 1a still keeps things simple, with any error triggering a straight exit from the program. In larger programs, the error-handling code could take up a lot more space if instead we try to recover from each type of error and continue executing. By doing so, we could end up obscuring the nominal program flow even more.

## Restoring Flow

If we adopt a single uniform policy for handling all errors, we can rewrite the code to clarify the nominal flow. We can do so by defining little interface functions for each potentially failing function call. For instance, for the call to the `malloc` routine, we can define an interface function that handles the out-of-memory condition internally, without bothering the caller (see Figure 2a). This then lets us switch to a simple call to `e_malloc` without having to disrupt the flow of the main code when an error occurs, as the final version in Figure 2b shows.

This approach no longer works if we want to divert the execution to alternate execution paths, depending on the type of error detected, when we're trying to implement more complex strategies for error recovery.

For the example in Figure 1a, four execution paths are possible, depending on which (if any) operation fails. Yet, the control-flow graph appears to have eight paths; visually, that's also the number your eyes start tracing through the code when you first look at it. This apparent flow determines the code's apparent complexity.

The alternative version in Figure 2b still has four possible executions, but on first inspection, your eyes see just the single nominal path. We could do that in this case not by ignoring errors but by standardizing the response to them. This turns out to be an important strategy in safety-critical-system design.

## Avoiding Failure

You'd be tempted to think that to make a system reliable, we must study every conceivable source of error and define detailed strategies to

2376-113X/15/$31.00 © 2015 IEEE Published by the IEEE Computer Society

```
void *
e_malloc(size_t n)
{ void *ptr = malloc(n);

  if (!ptr) {
    fprintf(stderr, "out of memory\n");
    exit(1);
  }
  return ptr;
}
(a)

buf = e_malloc(N);
fd = e_fopen(fnm, "r");
e_read(fd, buf, N);
(b)
```

**FIGURE 2.** Hiding error handling. (a) A simple interface function that handles the out-of-memory condition internally. (b) The cleaned-up code fragment. The simple call to **e_malloc** doesn't disrupt the flow of the main code when an error occurs.

handle each one, so that nominal execution can be restored in all cases. This is only partly true. It's true that before we build a system, we must have a good understanding of how it might fail under a broad range of conditions. As civil engineer and author Henry Petroski recently said, "The short definition of engineering is the avoidance of failure."[1]

But this doesn't mean that reliable code is doomed to be inscrutable. As I showed before, it's not the handling of errors in itself that necessarily complicates code structure, it's how the errors are handled. The safest method isn't always to devise elaborate mitigation strategies that can return the system to nominal execution under all imaginable circumstances. Sometimes the best strategy is to punt and merely move the system into a known state, so that an external user or operator can diagnose the problem with the benefit of a broader perspective. Doing so can reduce a system's complexity, which can make the system safer.

Unmanned spacecraft contain algorithms that continuously check the system's health, hunting for any type of anomalous condition. When an anomaly is detected, the simplest of these algorithms make no attempt to fix the problem. Instead, they place the spacecraft into a known *safe mode*, so that operators on the ground can diagnose the problem and devise a solution.

Safe mode is a system state with only the minimal functionality needed for a spacecraft to remain commandable, and therefore repairable. Stripping away the code needed to self-diagnose and repair all foreseen and unforeseen problems autonomously can significantly reduce the overall complexity (and improve the predictability) of system execution.

This approach to system safety has deep roots. When the software for the first moon landings was prepared at MIT in the mid '60s, uncontrolled complexity seemed to be putting the Apollo missions' time schedule at risk. NASA manager Bill Tindall reported in detail on the problems and how they were being addressed in his "Tindallgrams"—short updates he sent with some frequency to NASA mission planners in Houston. In his first Tindallgram from 31 May 1966, he wrote,

> *I am still very concerned about unnecessary sophistication in the program and the effect of this "frosting on the cake" on schedule and storage. It is our intention to go through the entire program, eliminating as much of this sort of thing as possible, I am talking about complete routines, such as "Computer Self-checks" ….*[2]

It seems counterintuitive that error-handling code would end up being the culprit of so much system complexity. But note that this code generally aims to defend against the highly unpredictable types of errors and failures that real life can throw at us. If that's almost doomed to failure, the safest strategy might be to practice extreme frugality.

## Hey, Reboot!

The original design of Unix also avoided elaborate error-handling code to reduce complexity, which might be partly why it so quickly became popular with developers. Early versions of the Unix kernel source code are legendary among programmers for their clarity and readability, even inspiring the well-known line-by-line exegesis by John Lions.[3]

A developer on another project, Tom Van Vleck, once described a discussion he had with Dennis Ritchie about the difficulty of error handling.[4] He said, "I remarked to Dennis that easily half the code I was writing in Multics was error recovery code." Dennis answered, "We left all that stuff out. If there's an error, we have this routine called panic, and when it is called, the machine crashes, and you holler down the hall, 'Hey, reboot it.'"

The simple strategy to reboot a misbehaving system is familiar to anyone who has used a computer, no matter what OS it runs, or any device that uses software, for that mat-

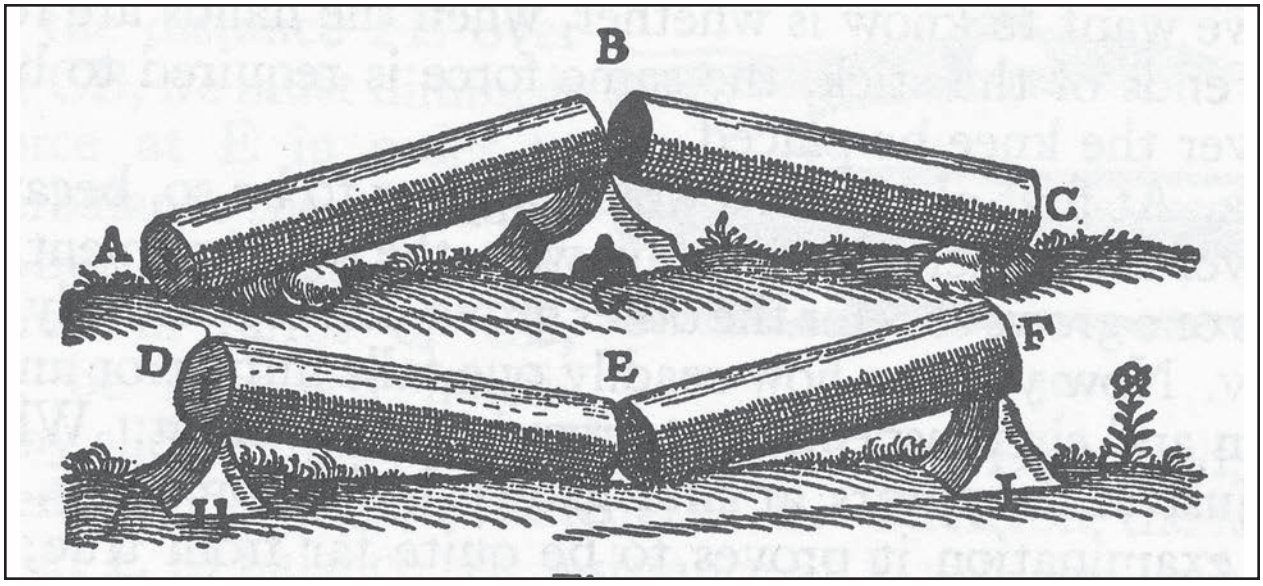


**FIGURE 3.** Fault protection can introduce new failure modes.[6] Instead of protecting the column, adding a third support could add another reason for the column to break in two.

ter. The strategy is fine and works most of the time, unless of course the error is hiding in the reboot code itself. This scenario might sound far-fetched, but it does occur. It indirectly caused the infamous Sol-18 problem that struck the first of two Mars Exploration Rovers shortly after it successfully landed on Mars in January 2004. Correctly diagnosing and fixing a problem such as this can be a true challenge even for humans back on Earth, as Glenn Reeves and Tracy Neilson documented in nail-biting detail.[5]

## Defect Rates

It's worth considering yet another point that can make us extra cautious in writing elaborate error-handling code. Error-handling code by its very nature executes infrequently, not just in system operation but also in system testing. Naturally, code that's not tested thoroughly tends to have a larger fraction of residual defects. This means that if we write large amounts of code to recover from obscure error conditions, that error-handling code will likely contain more defects than the code it's trying to protect. Hitting one of those new defects while recovering from an earlier error can aggravate the situation and put a system in a state that's even harder to diagnose and repair.

## Faulty Fault Protection

I noted earlier that fault protection systems can increase complexity and introduce entirely new failure modes into a system. An easy noncomputational example of this phenomenon is the familiar "Always Alert, Nobody Hurt" sign you can trip over in the dark.

An even better example dates back about two thousand years. In Roman times, marble columns were key components of most monumental buildings. The best columns were cut from a single piece of stone and transported to the worksite, where they had to be stored (horizontally) until they could be used.

For storage, the columns would sometimes be placed on small supports, which could prevent discoloration from prolonged contact with the ground. This, however, could cause long columns to break in the middle, perhaps just by the sheer force of gravity, or helped by playing Roman children who would no doubt jump up on the columns.

To prevent the breakage, a clever builder thought of placing an extra support at each column's midpoint, where breaks typically occurred. Ironically, this protective measure could introduce a new failure mode that was much more likely to strike. The added support could cause the column to break in the middle, this time in the opposite direction. Because a straight line goes through two points and not three, the fault protection increased the odds of breakage instead of reducing it. Figure 3, from a book by Galileo that first appeared in 1638, illustrates this problem.

The lesson in all this is that unless the cause of an error is well understood, well-intended remedial actions could actually make matters worse. In those cases, doing nothing might well turn out to be the better strategy. Doug McIlroy, the head of the department at Bell Labs where Unix, C, and C++ were born, once phrased it as, "The real hero of programming is the one who writes negative code" (www.azquotes.com/quote/819506).

The challenge in writing reliable code is similarly to find ways to remove code from an application by simplifying and generalizing, rather than continuing to add more. After all, the only code that can't fail is the code that isn't there to begin with.

## References


1. R. Latanision and C. Fletcher, "An Interview with … Henry Petroski," *The Bridge*, vol. 45, no. 1, 2015, pp. 49–55.
2. H.W. Tindall Jr., "Tindallgrams," D.T. Criag, ed.; http://web.mit.edu/digitalapollo/Documents/Chapter7/tindallgrams.pdf.
3. J. Lions, *Lion's Commentary on Unix 6th Edition with Source Code*, Peer-to-Peer Communications, 1996.
4. T. Van Vleck, "Unix and Multics," 1995; www.multicians.org/unix.html.
5. G. Reeves and T. Neilson, "The Mars Rover Spirit FLASH Anomaly," *Proc. 2005 IEEE Aerospace Conf.*, 2005, pp. 4186–4199.
6. G. Galilei, *Discourses and Mathematical Demonstrations Relating to Two New Sciences*, 1638.


**GERARD J. HOLZMANN** works at the Jet Propulsion Laboratory on developing stronger methods for software analysis, code review, and testing. Contact him at gholzmann@acm.org.



Selected CS articles and columns are also available for free at http://ComputingNow.computer.org.

*This article originally appeared in* IEEE Software, *vol. 32, no. 5, 2015.*

# IEEE COMPUTER SOCIETY: Be at the Center of It All

IEEE Computer Society membership puts you at the heart of the technology profession—and helps you grow with it.

**Here are 10 reasons why you need to belong.**

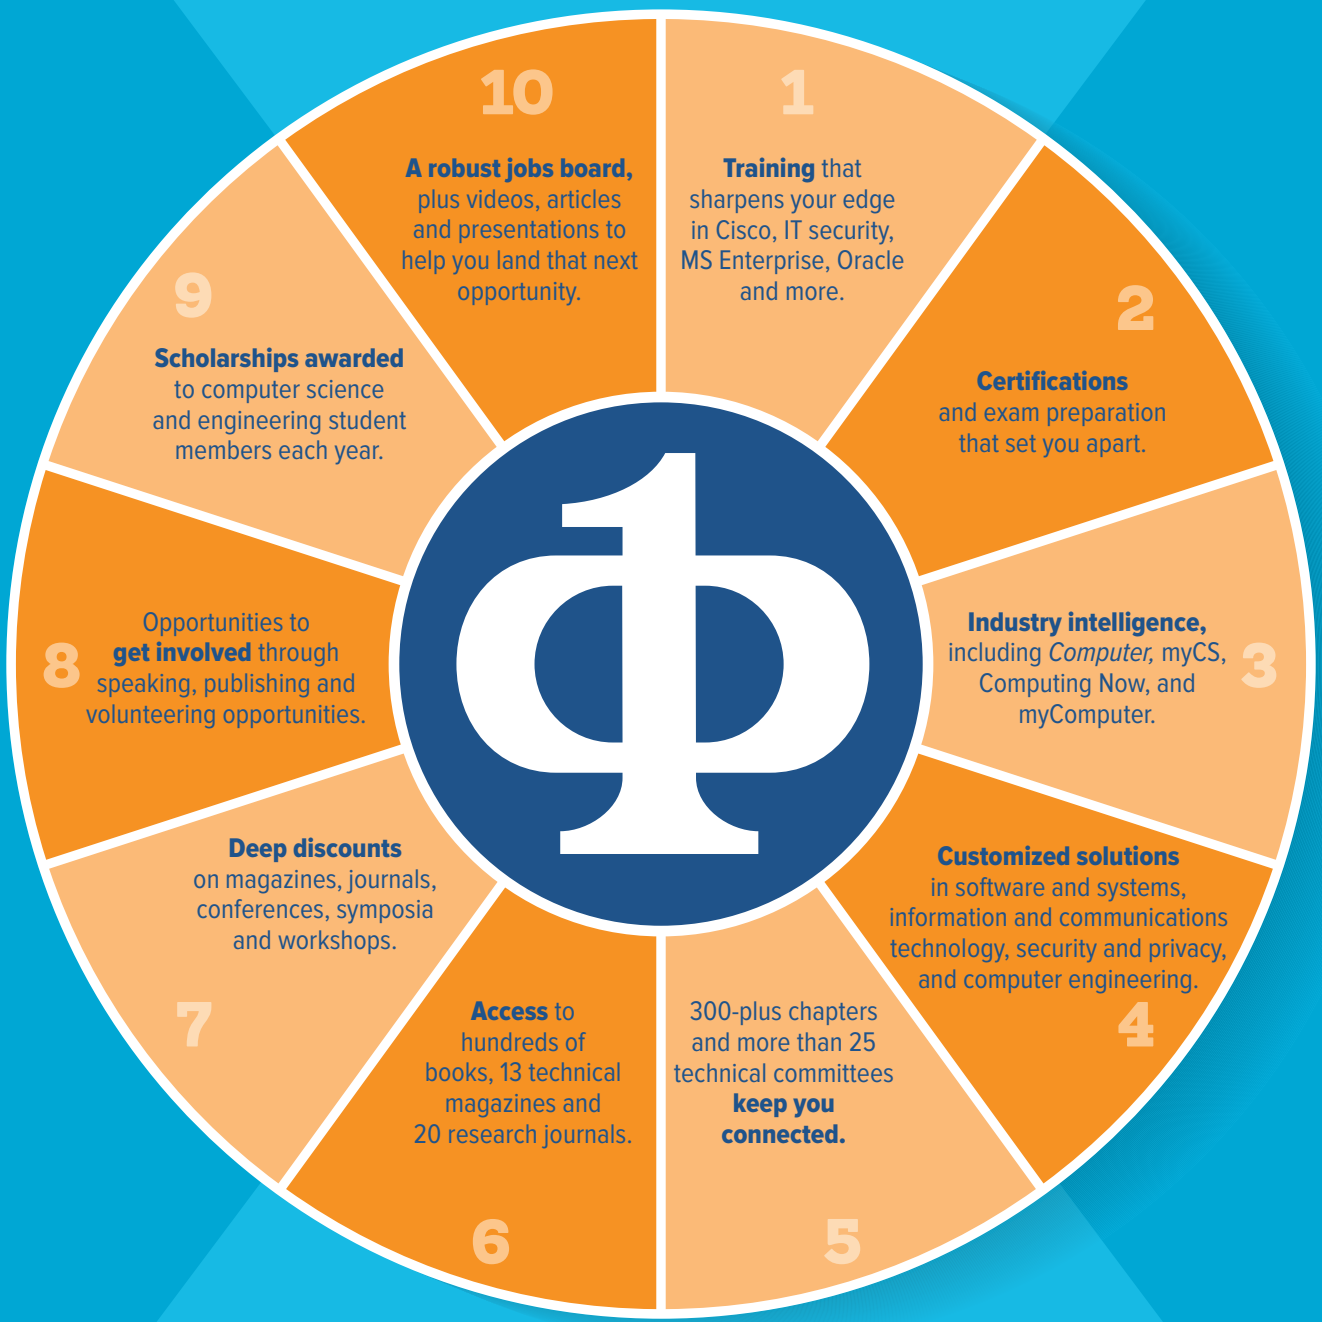


IEEE Computer Society—keeping you ahead of the game. Get involved today.

**www.computer.org/membership**

IEEE computer society

Editors: Bruce Campbell
and Francesca Samsel

# Murmurations: Drawing Together Art, Visualization, and Physical Phenomena

**Bruce D. Campbell**
*Rhode Island School of Design*

**Francesca Samsel**
*University of Texas–Austin*

## Editors' Note

This issue's Art on Graphics department is in an interview format. We are experimenting with a range of formats, seeking the best means of bridging the disciplines to facilitate dialogue. We thank Dennis Hlynsky for his willingness to join our experiment and provide feedback.

Dennis Hlynsky is a professor at the Rhode Island School of Design and head of the Film/Animation/Video Department. He is a practicing artist and early adopter of electronic media. His initial work with experimental video was followed by a decade of documentary filmmaking. As digital technologies evolved, they provided a refreshing outlet and reignited his interest in promoting potential experiments with shared video. At RISD he became a principal researcher for a RISD/IBM project to investigate artist workstations, and he began to study 3D modeling and animation. One of his current projects focuses on the movement of small animals and what animal visualizations can tell us.

***Francesca:*** The sparse poetic visual style of your work clearly conveys the complexity of movement and behaviors of the flocks of birds you document. A question arose between Bruce and me: "Might Dennis' sparse visual vocabulary inform new means of conveying multivariate time-varying data?" It is clear that your process as well as the work itself has many correlations with visualization practices. As discussions with you unfolded, surface connections led to deeper, mutually thought-provoking explorations. To bring our readers into the dialogue, I'd like to start with where some of the concrete areas of interest overlap and then move into the more abstract connections with your thoughts.

One of the key challenges in visualization is how to depict complex time-varying data in a clear, concise, and comprehensible fashion. Think of the birds as physical time-varying data points. As they participate in murmuration, they change direction, speed, and volume. Despite the complexity, the movements, directions, and speeds are instantly and intuitively understood, as shown in "Starlings Mapping Heaven" (see Figure 1). Could the shapes, which represent the birds, be abstracted, analyzed, and applied to visualizing other physical phenomena? Could we break down the shapes and their transformations and use the forms to convey other types of physical transformations?

***Dennis:*** I've been observing other group phenomena and noting similarities that video can capture for the visualization community to investigate. I think those investigations will continue to be useful to pursue. One of the ongoing technical discussions in the film community concerns frame rate. The standard rates of 24 and 30 frames per second (fps) are often used. In my work, I find it easier to reconsider the frame rate of video recording as the sample rate and the time the sample stays on screen as the refresh rate. My images employ a very slow refresh rate to evaluate a path taken.

***Francesca:*** You consistently mention your interest as a participant in observation science. Your datasets are time-dependent locations of subjects—group behavior of birds, insects, and small animals. Could you talk about observation in your work?

 Published by the IEEE Computer Society 



**Figure 1. Still from Hlynsky's "Starlings Mapping Heaven." Viewed abstractly, this provides a wealth of information on intuitional means of depicting location, direction, and speed, all challenges within the visualization community.**

***Dennis:*** Although science is thought of as a hypothesis-based process, it often starts as observation. As an aware human, I can't imagine pursuing a study of group behaviors without spending the time to observe the phenomena.

At the same time, I feel video recording is pithy. So much of the long-term inquiry feedback loop I experience is lost when capturing the experience as a video artifact. Using image processing, I hope to better convey my observations to someone who was absent at the time of the recording. It's the core problem of documentation. Deciding not to manipulate the recording for the sake of honesty disregards the inherent bias of the medium. But adding bogus narratives to drive home a point is selfish and misleading. It's tough to walk that line.

***Francesca:*** You use sampling techniques, choosing to keep a limited information set based on a variety of factors that are data-specific. Your work is similar to visualization in that you are processing the video in order to direct attention and make key phenomena easier to absorb. Can you describe the techniques you employ?

***Dennis:*** Contemporary digital film cameras snack on gigabytes. Once the camera is purchased, it's cheap to make data-intensive observations of the mundane world. This means one can choose a subject, in my case small animals en masse, and easily make a lengthy comparative study. My initial challenge was to focus on the individual paths the creatures took while in swarm in order to understand the swarm as a whole. I believe understanding the problem and then matching a solution to it always leads to the elegant solution.

Because I intended to study the mass, I needed to look at the data as broadly as possible. I began to question the standard of making recordings representing the moment-to-moment way we see. The standard video recording seemed focused on isolating distinct moments. Visually following such large numbers of moving creatures in "real time" was problematic because I just couldn't keep track of everything. Slowing the recordings down was a dead end. Tracking felt impossible. Time lapse was useless. Focusing on individual creatures and individual moments seemed wrong if I was motivated by group behavior.

I started imagining the recorded data en masse as well. So, I simply extended the amount of time the location data of the bird stayed on the screen. You can see the results in "Flight of a Small Brown Cloudyspot" (see Figure 2). In this way, it was easier to cognitively understand the movement as a whole. I developed a technique of combining pixel values from many video samples (frames) to draw the path but not alter the time base of the recording.

***Francesca:*** This is potentially useful to scientists in far afield research areas. Art has been drawing on science throughout history, and visualization in particular is an area where science can draw from art to further its purposes and processes. Can you provide any information that might assist the



**Figure 2. "Flight of a Small Brown Cloudyspot." Here Hlynsky uses a sampling of time steps but retains each image on the screen for a longer duration. This enables an easily understood documentation of the subject's movement and pathways over time.**

greater visualization community in applying your methods to visualization?

***Dennis:*** I am honored to be a part of your efforts to create a dialogue about art and science practice. My community has begun to speak of the data we leave behind as exhaust. Data exhaust and our ability to process these trails have brought much art making into the realm of research. As the communities of artists and scientists communicate more with each other, both the schisms and overlaps in the communities become clear. For me, something seems to be lost in processes when complex natural phenomena are reduced into usable bits to support existing aims of science. To a seasoned artist, good science is wonderfully informative, but it disregards many of the human observations we use to obtain evidence.

***Bruce:*** Spending time with you filming starling murmurations let me experience the depth that focused observation time affords the thought process. Are there other tools that can help with observing natural phenomena?

***Dennis:*** Sure. Consider those starling murmurations as a case study. Andrea Cavagna at the Italian National Institute for Condensed Matter Physics developed a stereo camera system that makes it possible to gather data from larger flocks of birds where previously scientists could only work out the positions of tens of birds at any one time.[1] Andrea has been dogged in creating technical solutions to the problem of occlusion in data gathering of large flocks.[2] His work creates marvelous tools. The flocking organization, once converted into a 3D data space, can be explored from any angle. Being able to import these datasets into Video FX graphic software may not improve the science but might offer better human processing.

For my work, I am not certain what a visualization of flying within the flock in virtual space will tell us. But it might offer a view of how a bird sees the flock from the inside. Although small, birdbrains and bird bodies are a complex system in and of themselves. While it's wonderful to describe the movements in mathematical terms, pure mathematical observation of the flock ignores issues that an artist might focus on. If we begin to look at social behavior and give the animals a little more credit than being dumb, instinctual, and completely predictable creatures, then contextual decision making adds a multifaceted and social element.

***Bruce:*** Is that why you film starlings from October to March, to observe behaviors?

***Dennis:*** Watching starlings is just an itch I have to scratch. Art, like science, is often inefficient and doesn't need to have specific reasons for being other than understanding the mystery of the universe. I find metaphorical parallels between the invasiveness of the starling species and our own human conquests. Social flocking through mass sharing of media objects is fascinating. I find the transition between the catatonic and epileptic flights of flocks captivating. Starlings are surprising because they are such good mimics. They seem to learn as a group over a season. I like to watch them evolve group behavior from autumn equinox to spring. What's there not to like?

I would never classify my fascinations and inquiries as science. I try to be honest and truthful to my subject. I hope my visual processing doesn't add spurious narratives. But as an artist I am free to make sense of nonsense. I enjoy it.

***Bruce:*** Are you considering art and science in the work that results in these videos?

***Dennis:*** At times I feel the people who practice art and science get their instruction sets from different DNA. The danger of being isolated in the laboratory or studio often blinds one to the big picture. I do believe each discipline can offer a useful reminder of why we do what we do. We are both concerned with quality of life issues and need to be reminded to evaluate if the research we do is contributing to our overall sustainability and happiness or if we are simply lost in our strategies. The people who practice art or science need each other when investigating certain types of phenomena. In those situations, progress is strongest when content and practices cross boundaries and are shared by multiple communities. I appreciate the moments when I find a phenomenon that appeals to both communities. The starling videos do this based on the response I've gotten from a wide audience. These recordings present factual data to analyze and ephemeral automatic skywriting to interpret (see Figure 3).

***Francesca:*** What response have you received from scientists?

***Dennis:*** As one example, Rusty Lansford came to me as a researcher looking at the cellular level, not at the organism level. Rusty immediately saw relevancy in a visualization technique. He called me to talk about the process I used to visualize what



**Figure 3. "Wings." This is a close-up look at skywriting, a phenomenon drawing interest from both the scientific and artistic communities.**

I call "small migrations." Because the memory of motion is kept, the flight paths become glyphs that naturally retain temporal and spatial information. Rusty had been looking at cell migrations in the formation of quail embryos and was running into a nagging issue of visualizing complex systems—a problem of data overload.[3] Being able to visualize paths would be useful to understanding just what was happening in his microcellular world. During the exchange, Rusty introduced me to the idea of autonomous collaboration, which speaks to the complex group movement shown in "Individual Starling Position Adjustments" (see Figure 4). We both came away with something.

***Francesca:*** I was immediately drawn to your work as an artist for the reasons you describe regarding



**Figure 4. Microcellular researcher Randy Lansford was drawn to "Individual Starling Position Adjustments." This piece provided inspiration for retaining temporal and spatial information within Lansford's research, the memory of motion. He introduced Hlynsky to the idea of automomous collaboration, of interest to Hlynsky because it speaks to flocking behaviors and how the birds learn as a group through the season.**

*This article originally appeared in* IEEE Computer Graphics and Applications, *vol. 35, no. 4, 2015.*

inquiry into why and how we capture phenomena of interest to science. Your film techniques immediately inspire further consideration. We started with pursuing some concrete means in which art and science can inform one another and yet our expansive discussions moved toward more conceptual connections that have provided much food for thought. Care to highlight any particular connections you've made by discussing your work with us?

***Dennis:*** In New England, 2015 has been an abnormally cold and wicked winter. I enjoyed the times Bruce and I met to film the starlings, standing above a cold parking lot on 30-foot mounds of snow conveying to each other the brand of fixations we practice. If only I could melt the frozen words carried on my breath.

If I were to brief my thoughts over the past few months, I would say that my observation of small animals en masse has provided a model to understanding complex systems. The question is, "How should I make use of this model?"

My tenure as a professor has given me a unique view. Students have washed in and out of my classrooms over the years and allowed me to witness generational modes of behavior (instruction sets). I believe the current generation of young adults is evidence of a sea change. The change deemphasizes leadership and individualism in favor of crowdsourcing social directives. Autonomous collaboration is a means of massing, where each member of the group acts without a leader but shares an instruction set with the group (DNA). The politics are neither anarchist nor socialist but rather found with abundance in nature. I hope the images I have made embody this narrative.

***Bruce:*** Now that you've been incorporating thoughts from sharing interest with scientists and artists, what's next?

***Dennis:*** The video recordings I have made continue to thrive in the blogosphere (see https://vimeo.com/user491023 and Figure 4). They have become a connecting tissue, a sort of interstitial substance that joins me to like-minded people. I receive daily email contact from people who associate work they are doing with these recordings. The investigations are varied and invoke a curiosity in me.

Recently I was contacted and asked to describe my process for an image I made for the Howard Hughes Medical Institute (HHMI) BioInteractive image of the week webpage.[4] HHMI provides an image bank for scientists who produce images relevant to their research. I believe what I do was misclassified because I am not a scientist. This will continue to provide an ongoing challenge to inhabit a privileged view between what are now two very different professional practices.

Currently I am intrigued by the starling visual system. What interests me is the notion that the lateral eye placement isn't ideal for stereovision, but the birds seem to land elegantly on wires in space. The beautiful arching flight paths seen in my Video FX tampering may be a means of determining self-location in space. I have also become interested in the evolution of behavior over the season. With starling life being an average of 2.5 years, a good portion of the flock is new to murmuration each fall. The process of learning from each other en masse presents a point of inquiry. Right now I am limited by my teaching job because I am only able to record sporadically throughout the season. I would like to install a series of Web cameras to observe the daily evolution of the flock from fall to spring in order to have access to a consistent observation presence.

## References

1. C. Barras, "Birdwatching in Stereo Captures Flocks in 3D," *New Scientist*, 7 May 2008; www.newscientist.com/article/dn13853-birdwatching-in-stereo-captures-flocks-in-3d.html.
2. A. Cavagna et al., "The STARFLAG Handbook on Collective Animal Behaviour: 2. Three-Dimensional Analysis," *Animal Behaviour*, vol. 76, no. 1, 2008, pp. 237–248.
3. Y. Sato et al., "Dynamic Analysis of Vascular Morphogenesis Using Transgenic Quail Embryos," *PLOS One*, vol. 5, no. 9, 2010; http://journals.plos.org/plosone/article?id=10.1371/journal.pone.0012674.
4. D. Hlynsky, "Following Tracks in the Sky," BioInteractive, Howard Hughes Medical Inst., 2015; www.hhmi.org/biointeractive/following-tracks-sky.

*Contact Dennis Hlynsky at dhlynsky@risd.edu.*

***Bruce Campbell*** *is an adjunct faculty member at the Rhode Island School of Design. Contact him at bcampbel01@risd.edu.*

***Francesca Samsel*** *is a research associate in the Center for Agile Technology at the University of Texas–Austin and an artist-in-residence at the Los Alamos National Laboratory. Contact her at figs@cat.utexas.edu.*

**IT ANALYST** wanted in Detroit, Michigan for the design and development of software for control of production flow and inventory. Send resume to Courtney Hodare, Director, Talent Management & Culture, Little Caesar Enterprises Inc., 2211 Woodward Avenue, Detroit, MI 48201.

**SENIOR SYSTEM ENGINEER-BI.** (Mult Open) sought by Bara Consulting, Inc in Edison, NJ. Bach's deg Info Tech or rltd w/5 yrs exp. Work w/client's bus. mgrs, affiliates, & corp IT app support team. Implmt client reqs in adherence to client's SDLC process. Dsgn data model for data warehouse. Dvlp end-to-end execution strategy of ETL. Perform analysis/dvlpmt of bus. models, ETL dsgn, & detail technical spec dsgn. Dvlp Informatica mappings; perform end-to-end testing; & coord w/reporting team & bus. for user testing. Dvlp tech objects in Informatica, Teradata, Oracle, Linux, & DB2 platforms. Migration to production envrmt & shadow execution run. Performance tuning of loading process to increase reporting efficiencies. Setup pilot solutions & prototypes. Respond to client needs & resolve issues w/in client's SLA framework. Mail resumes to Bara, 860 Route 1, Ste 101, Edison, NJ 08817, Attn: HR.

**SR. ENGG SVC ARCHTCT.** (NY, NY & client sites nat'l & int'l). Archt, config, dsgn, dvlp & dply CA's Mobility Solutions. Assess & cnfrm config settings. Dvlp go-to-mkt plans. Gather diagnostic data for product mngmt. Drve prod launches & gthr/anlyze prod perf, op metrics, & win/loss analyses. Prvde pre-sales, sales & engg svcs supp. REQ: Bach Deg or for equiv in Comp Sci, Math, Engg (any) or rel + 5 yrs prog exp in job &/or a rel occup. Will accept Masters in Comp Sci, Math, Engg (any) or rel + 2 yrs of exp. Must have exp: Internal archtre & dsgn for mobile dev mgmt techs; Archtctng & dplyng highly available & scalable pltfrms that use mobile dev mgmt & supp infrstrctre tech & interoperate w/ multi servers & data centers; Assisting prod mgmt & engg teams w/ archtre, dsgn, trblshtng, testing, & validation; Gathering diagnostic data & idntfyng relevant envrmntl info to relay to product mgmt. & engg teams to expedite prob resolution or dplymnt of new prod releases w/in a client envrmnt; Mngng escalated support cases to resolution;

BAYLOR UNIVERSITY

1845

## Faculty Position

**The Electrical and Computer Engineering Department of Baylor University seeks faculty applicants for a tenured/tenure-track Faculty Position at any level. Any area of expertise will be considered but applicants in computer engineering will be given special consideration. Applicants for assistant professor must demonstrate potential for sustained, funded scholarship and excellent teaching; applicants for associate or full professor must present evidence of achievement in research and teaching commensurate with the desired rank. The ECE department offers B.S., M.S., M.E. and Ph.D. degrees and is rapidly expanding its faculty size. Facilities include the *Baylor Research and Innovation Collaborative* (BRIC), a newly-established research park minutes from the main campus.**

**Chartered in 1845 by the Republic of Texas, Baylor University is the oldest university in Texas. Baylor has an enrollment of over 15,000 students and is a member of the Big XII Conference. Baylor's mission is to educate men and women for worldwide leadership and service by integrating academic excellence and Christian commitment within a caring community. The department seeks to hire faculty with an active Christian faith; applicants are encouraged to read about Baylor's vision for the integration of faith and learning at www.baylor.edu/profuturis/.**

**Applications will be considered on a rolling basis until the January 1, 2015 deadline. Applications must include:**

1. **a letter of interest that identifies the applicant's anticipated rank,**
2. **a complete CV,**
3. **a concise statement of teaching and research interests,**
4. **the names and contact information for at least four professional references.**

**Additional information is available at www.ecs.baylor.edu. Send materials via email to Dr. Keith Schubert at keith_schubert@baylor.edu. Please combine all submitted material into a single pdf file.**

**Baylor University is affiliated with the Baptist General Convention of Texas. As an Affirmative Action/Equal Employment Opportunity employer, Baylor encourages candidates of the Christian faith who are minorities, women, veterans, and persons with disabilities to apply.**

UNIVERSITY AT ALBANY
State University of New York

## Computer Science Department

## Assistant Professor

The Computer Science Department in the College of Engineering and Applied Sciences at the University at Albany SUNY is searching for tenure-track faculty beginning either January 2016 or Fall 2016. While the search is primarily at the Assistant Professor level, we will consider senior applicants with appropriate credentials. We seek candidates with research expertise in cyber-security and privacy, an area that the University has identified as a priority for development.

Applicants must have a Ph.D. in Computer Science, Computer Engineering, Informatics, or a closely related discipline.

For a complete job description and application procedures, visit:
**https://albany.interviewexchange.com/jobofferdetails.jsp?JOBID=53468**

Questions regarding the position may be addressed to **CSCyber@albany.edu** .

The Computer Science Department offers Bachelor's, Master's and Ph.D. degrees. For additional information, please visit **http://www.cs.albany.edu/**.

*The University at Albany is an EO/AA/IRCA/ADA*

 2376-113X/15/$31.00 © 2015 IEEE

Supporting sales proposals & cust prod evaluations fr transactions in excess of $250,000 new license revenue; Mobile dev tech, device mgmt, secrty, config options parameters, & tools; Securing mobile email & in-house dvlpd mobile apps; Freq travel to unanticip nat'l & interat'l client sites; Work fr home anywhere in US. Send resume to: Althea Wilson, CA Technologies, One CA Plaza, Islandia, NY 11749, Refer to Requisition #109661.

**SENIOR SOFTWARE ENGINEER.** (mult. openings) sought by Q1W Holdings, LLC in Orlando, FL w/a BS in Comp Eng. or rltd & 5 yrs exp. Modify existing s/ware to correct errors, adapt to new h/ware, improve its performance. Dvlp & direct s/ware systm testing & validation procedures, prgmg & documentation. Confer w/systms analysts/engineers/programmers to dsgn systm & to obtain info on project limitations & capabilities, performance reqmts & interfaces. Analyze user needs & s/ware reqmts to determine feasibility of dsgn w/in time & cost restraints. Dsgn/dvlp/modify s/ware systms using scientific analysis & mathematical models to predict/measure outcome. Store/retrieve/manipulate data for analysis of systm capabilities & reqmts. Consult w/customers about s/ware systm dsgn/maintenance. Coord s/ware systm installation & monitor eqpmt functioning to ensure specifications are met. Obtain & eval info. on factors such as reporting formats req'd, costs & security needs to determine h/ware configuration. Resumes to: 1500-A Tradeport Dr., Orlando, FL 32824, Attn: HR.

NUS National University of Singapore

**Multiple Tenure-Track or Tenured Faculty Positions in Computer Science**

The Department of Computer Science at National University of Singapore (NUS) invites applications for several tenure-track or tenured faculty positions. We have positions dedicated to cyber security or big data analytics as well as positions open to all areas of computer science. While our main focus is on the assistant professor level, we also welcome exceptional candidates at the associate and full professor levels. For applications at the assistant professor level, candidates should demonstrate excellent research potential and a strong commitment to teaching. Candidates at more senior levels should have an established record of outstanding research achievements.

The Department of Computer Science at NUS is highly ranked internationally. It enjoys ample research funding, moderate teaching load, excellent facilities, and extensive international collaborations. The department covers all major research areas in computer science and boasts a thriving PhD program that attracts the brightest students from the region and beyond. More information is available at http://www.comp.nus.edu.sg/.

NUS offers highly competitive salaries and is situated in Singapore, an English-speaking cosmopolitan city and a meeting point of many cultures, both the east and the west. Singapore offers high-quality education, healthcare, and extremely low tax rates.

Interested candidates are invited to send, via electronic submission, the following materials to the Chair of the CS Faculty Search Committee, Prof. David Hsu, at csrec@comp.nus.edu.sg:

- A cover letter that clearly indicates main research interests
- Curriculum Vitae
- A teaching statement
- A research statement

Please arrange for

- at least 3 references

to be sent directly to the same e-mail address or provide the contact information.

Application review will commence on October 1, 2015 and continue until the positions are filled. To ensure maximal consideration, please submit your application by December 15, 2015.

## ADVERTISER INFORMATION • OCTOBER 2015

**Advertising Personnel**

**Debbie Sims: Advertising Coordinator**
**Email: dsims@computer.org**
**Phone: +1 714 816 2138 | Fax: +1 714 821 4010**

**Chris Ruoff: Senior Sales Manager**
**Email: cruoff@computer.org**
**Phone: +1 714 816 2168 | Fax: +1 714 821 4010**

**Advertising Sales Representatives (display)**

**Central, Northwest, Far East:**
**Eric Kincaid; Email: e.kincaid@computer.org**
**Phone: +1 214 673 3742; Fax: +1 888 886 8599**

**Northeast, Midwest, Europe, Middle East:**
**David Schissler; Email: d.schissler@computer.org**
**Phone: +1 508 394 4026; Fax: +1 508 394 1707**

**Southwest, California:**
**Mike Hughes; Email: mikehughes@computer.org**
**Phone: +1 805 529 6790**

**Southeast:**
**Heather Buonadies; Email: h.buonadies@computer.org**
**Phone: +1 201 887 1703**

**Advertising Sales Representatives (Classifieds & Jobs Board)**

**Heather Buonadies; Email: h.buonadies@computer.org**
**Phone: +1 201 887 1703**

**HCL AMERICA, INC.** has the following openings in multiple locations. Travel & work @ various unanticipated client sites as assigned: **Business Analysts:** Transform business requirements into functional specifications, focusing on workflow analysis & design, business process reengineering, user interface design, & process flow modeling. **Cary, NC:** Reqs BS*+0 (HCL459); MS+1/BS+5 (HCL461). **Frisco, TX:** Reqs BS*+0 (HCL462); MS+1/BS+5 (HCL464). **Redmond, WA:** Reqs BS*+0 (HCL465); MS+1/BS+5 (HCL467). **Sunnyvale, CA:** Reqs BS*+0 (HCL468); MS+1/BS+5 (HCL470). **Jersey City, NJ:** Reqs BS*+0 (HCL471); MS+1/BS+5 (HCL473). **Database Administrators:** Install, configure, maintain, build, & back-up databases. **Cary, NC:** Reqs BS*+0 (HCL474); MS+1/BS+5 (HCL476). **Frisco, TX:** Reqs BS*+0 (HCL477); MS+1/BS+5 (HCL479). **Redmond, WA:** Reqs BS*+0 (HCL480); MS+1/BS+5 (HCL482). **Sunnyvale, CA:** Reqs BS*+0 (HCL483); MS+1/BS+5 (HCL485). **Jersey City, NJ:** Reqs BS*+0 (HCL486); MS+1/BS+5 (HCL488). **Network Administrators:** Analyze, design, troubleshoot, implement, maintain, & manage network solutions. **Cary, NC:** Reqs BS*+0 (HCL489); MS+1/BS+5 (HCL491). **Frisco, TX:** Reqs BS*+0 (HCL492); MS+1/BS+5 (HCL494). **Redmond, WA:** Reqs BS*+0 (HCL495); MS+1/BS+5 (HCL497). **Sunnyvale, CA:** Reqs BS*+0 (HCL498); MS+1/BS+5 (HCL500). **Jersey City, NJ:** Reqs BS*+0 (HCL501); MS+1/BS+5 (HCL503). **Programmer Analysts:** Define, develop, code, & test programs & applications. **Cary, NC:** Reqs BS*+0 (HCL504); MS+1/BS+5 (HCL506). **Frisco, TX:** Reqs BS*+0 (HCL507); MS+1/BS+5 (HCL509). **Redmond, WA:** Reqs BS*+0 (HCL510); MS+1/BS+5 (HCL512). **Sunnyvale, CA:** Reqs BS*+0 (HCL513); MS+1/BS+5 (HCL515). **Jersey City, NJ:** Reqs BS*+0 (HCL516); MS+1/BS+5 (HCL518). **Software Engineers:** Involved with software implementation, design, testing, & coding. **Cary, NC:** Reqs BS*+0 (HCL519); MS+1/BS+5 (HCL521). **Frisco, TX:** Reqs BS*+0 (HCL522); MS+1/BS+5 (HCL524). **Redmond, WA:** Reqs BS*+0 (HCL525); MS+1/BS+5 (HCL527). **Sunnyvale, CA:** Reqs BS*+0 (HCL528); MS+1/BS+5 (HCL530). **Jersey City, NJ:** Reqs BS*+0 (HCL531); MS+1/BS+5 (HCL533). **Systems Analysts:** Define systems strategy & develop systems requirements. **Cary, NC:** Reqs BS*+0 (HCL534); MS+1/BS+5 (HCL536). **Frisco, TX:** Reqs BS*+0 (HCL537); MS+1/BS+5 (HCL539). **Redmond, WA:** Reqs BS*+0 (HCL540); MS+1/BS+5 (HCL542). **Sunnyvale, CA:** Reqs BS*+0 (HCL543); MS+1/BS+5 (HCL545). **Jersey City, NJ:** Reqs BS*+0 (HCL546); MS+1/BS+5 (HCL548). **Project Managers:** Responsible for managing, planning, coordinating, supervising & directing IT professionals. **Cary, NC:** Reqs MS+1/BS+5 (HCL559). **Frisco, TX:** Reqs MS+1/BS+5 (HCL560). **Redmond, WA:** Reqs MS+1/BS+5 (HCL561). **Sunnyvale, CA:** Reqs MS+1/BS+5 (HCL562). **Jersey City, NJ:** Reqs MS+1/BS+5 (HCL563). *Employer will accept a combination of edu. & exp. as determined by a qualified evaluation service as equivalent to a Bachelor's degree. Multiple job openings available. How to apply: Mail resume, referencing HCL job code, including job history, to: HCL America, Inc., Attn: ISG Department, 330 Potrero Ave, Sunnyvale, CA 94085. HCL is an Equal Opportunity Employer.

**HCL AMERICA, INC.** has the following openings in multiple locations. Travel

## Cisco Systems, Inc. is accepting resumes for the following positions:

**BELLEVUE, WA: Technical Marketing Engineer** (Ref.# BEL8): Responsible for enlarging company's market and increasing revenue by marketing, supporting, and promoting company's technology to customers. Travel may be required to various unanticipated locations throughout the United States.

**BOXBOROUGH, MA: Project Manager** (Ref.#: BOX17): Coordinate small, medium, large/complex and multiple projects throughout the project lifecycle (initiate, plan, execute, control, close) or a portion of a larger, more complex project. Telecommuting permitted and travel may be required to various unanticipated locations throughout the United States. **Network Consulting Engineer** (Ref.# BOX11): Responsible for the support and delivery of Advanced Services to company's major accounts.

**CENTENNIAL, CO: Systems Engineer** (Ref.# CENT102): Provide business-level guidance to the account team or operation on technology trends and competitive threats, both at a technical and business level. Telecommuting permitted and Travel may be required to various unanticipated locations throughout the United States.

**COLUMBIA, MD: Technical Support Engineer** (Ref.# COLU3): Process telephone and email requests from enterprise customers. **Software Development Manager** (Ref.# COLU4): Lead a team in the design and development of company's hardware or software products.

**IRVINE, CA: Solutions Consultant** (Ref.# IRV14): Responsible for planning, designing, implementing, operating and optimizing (PDIOO) Safety and Security solutions utilizing multiple technologies, and the company's and partner's products. Telecommuting permitted and travel may be required to various unanticipated locations throughout the United States.

**ISELIN/EDISON, NJ: Network Consulting Engineer** (Ref#: ED7): Responsible for the support and delivery of Advanced Services to company's major accounts. Travel may be required to various unanticipated locations throughout the United States. **Network Consulting Engineer** (Ref#: ED9): Responsible for the support and delivery of Advanced Services to company's major accounts. Telecommuting permitted.

**JACKSONVILLE, FL: Solutions Architect** (Ref.# JAC1): Responsible for IT advisory and technical consulting services development and delivery.

**RESEARCH TRIANGLE PARK, NC: Network Engineer** (Ref.# RTP541): Responsible for the operational support of internal network systems. **Product Manager** (Ref.# RTP621): Create high level marketing strategies and concepts for company solutions for markets and segments worldwide. **Network Consulting Engineer** (Ref.# RTP2): Responsible for the support and delivery of Advanced Services to company's major accounts. **IT Manager** (Ref.# RTP301): Design, Architect and Implement Data Center infrastructure.

**RICHARDSON, TX: Customer Support Engineer** (Ref.# RIC919): Responsible for providing technical support regarding the company's proprietary systems and software. Travel may be required to various unanticipated locations throughout the United States. **Manager, Technical Services** (Ref.# RIC18): Responsible for leading a team in the delivery of world-class customer support on a line of products or for a targeted group of customers. Telecommuting permitted. **Solutions Specialist** (Ref.# RIC341): Increase sales of Advanced Services for Advanced Technologies by gathering customer requirements in order to design an Advanced Technologies Services Solution. Telecommuting permitted. **Customer Support Engineer** (Ref.# RIC1): Responsible for providing technical support regarding the company's proprietary systems. **Technical Marketing Engineer** (Ref.# RIC13): Responsible for enlarging company's market and increasing revenue by marketing, supporting, and promoting company's technology to customers.

**RICHMOND, VA: Network Consulting Engineer** (Ref.# RIV21): Responsible for the support and delivery of Advanced Services to company's major accounts. Telecommuting permitted.

**SAN FRANCISCO, CA: User Experience Engineer** (Ref.# SF22): Identify user interaction requirements and develop user experience interface specifications and guidelines. **CNG Member of Technical Staff** (Ref.# SF25): Design, implement, and test software for cloud-managed security appliance devices.

**SAN JOSE/MILPITAS/SANTA CLARA, CA: Project Manager** (Ref.# SJ18): Coordinate small, medium, large/complex and multiple projects throughout the project lifecycle (initiate, plan, execute, control, close) or a portion of a larger, more complex project. **Customer Solutions Architect** (Ref.# SJ983): Responsible for IT advisory and technical consulting services development and delivery. Telecommuting permitted and travel may be required to various unanticipated locations throughout the United States. **Video Solutions Architect** (Ref.# SJ637): Help establish a customers roadmap for pervasive video with specific product milestones and commitments and key outcome benefits. Telecommuting permitted and travel may be required to various unanticipated locations throughout the United States. **Optical Design Engineer** (Ref.# SJ337): Simulation, design and characterization of physical structures suitable for maintaining good optical coupling. **Network Consulting Engineer** (Ref.# SJ9): Responsible for the support and delivery of Advanced Services to company's major accounts.

**SEATTLE, WA: Network Consulting Engineer** (Ref.# SEA7): Responsible for the support and delivery of Advanced Services to company's major accounts. Telecommuting permitted.

**PLEASE MAIL RESUMES WITH REFERENCE NUMBER TO CISCO SYSTEMS, INC., ATTN: M51H, 170 W. Tasman Drive, Mail Stop: SJC 5/1/4, San Jose, CA 95134. No phone calls please. Must be legally authorized to work in the U.S. without sponsorship. EOE.**

**www.cisco.com**

& work @ various unanticipated client sites as assigned: **Sales Engineers:** Responsible for selling various IT services to clients & pre-sales engineering support & guidance to customers. **Cary, NC:** Reqs BS*+0 (HCL564); MS+1/BS+5 (HCL566). **Frisco, TX:** Reqs BS*+0 (HCL567); MS+1/BS+5 (HCL569). **Redmond, WA:** Reqs BS*+0 (HCL570); MS+1/BS+5 (HCL572). **Sunnyvale, CA:** Reqs BS*+0 (HCL573); MS+1/BS+5 (HCL575). **Jersey City, NJ:** Reqs BS*+0 (HCL576); MS+1/BS+5 (HCL578). **Business Development Managers:** Responsible for business development, discussions, planning, & product development, & working towards connecting & building relationships with key decision makers in customer organizations. **Cary, NC:** Reqs MS+1/BS+5 (HCL579). **Frisco, TX:** Reqs MS+1/BS+5 (HCL580). **Redmond, WA:** Reqs MS+1/BS+5 (HCL581). **Sunnyvale, CA:** Reqs MS+1/BS+5 (HCL582). **Jersey City, NJ:** Reqs MS+1/BS+5 (HCL583). **General & Operations Managers:** Coordinate business initiatives & integrate people processes across the company, & spearhead development, communication, & implementation of effective growth strategies & processes. **Cary, NC:** Reqs MS+1/BS+5 (HCL584). **Frisco, TX:** Reqs MS+1/BS+5 (HCL585). **Redmond, WA:** Reqs MS+1/BS+5 (HCL586). **Sunnyvale, CA:** Reqs MS+1/BS+5 (HCL587). **Jersey City, NJ:** Reqs MS+1/BS+5 (HCL588). **Sales Managers:** Responsible for sales operations in the area of infrastructure management, managing & directing sales operations, & carrying out business development. **Cary, NC:** Reqs MS+1/BS+5 (HCL589). **Frisco, TX:** Reqs MS+1/BS+5 (HCL590). **Redmond, WA:** Reqs MS+1/BS+5 (HCL591). **Sunnyvale, CA:** Reqs MS+1/BS+5 (HCL592). **Jersey City, NJ:** Reqs MS+1/BS+5 (HCL593). Multiple job openings are available. How to apply: Mail resume, referencing HCL job code, including job history, to: HCL America, Inc., Attn: ISG Department, 330 Potrero Ave, Sunnyvale, CA 94085. HCL is an Equal Opportunity Employer. **Senior Market Research Analysts:** Conduct market research & price comparisons, & analyze market data to identify the competition, target groups, & market opportunities. **Frisco, TX:** Reqs MS+1/BS+5 (HCL594). *Employer will accept a combination of education and experience as determined by a qualified evaluation service as equivalent to a Bachelor's degree. Multiple job openings are available. How to apply: Mail resume, referencing HCL job code, including job history, to: HCL America, Inc., Attn: ISG Department, 330 Potrero Ave, Sunnyvale, CA 94085. HCL is an Equal Opportunity Employer.


**THE DEPARTMENT OF ELECTRICAL ENGINEERING AND COMPUTER SCIENCE (EECS) AT VANDERBILT UNIVERSITY** is seeking candidates for two tenured/tenure-track (T/TK) faculty positions in Computer Science. Appointments at all ranks will be considered; we prefer to fill at least one position at the assistant professor level. The areas of primary emphasis are Big Data and Cyber-physical Systems. The Vanderbilt CS program provides a unique, collaborative, and interdisciplinary research environment that spans a wide range of topics in Computer Science. New trans-institutional programs are creating opportunities for research on issues of broad significance that create and extend collaborations across multiple fields. We seek an active, culturally and academically diverse faculty of the highest caliber, skilled in both scholarship and teaching. Average T/TK faculty funding is nearly ~$1M per year from DARPA, DoD, NASA, NIH, NSF, and many industry sponsors. All junior faculty members hired during the past decade have received prestigious young investigator awards, such as NSF CAREER and DARPA CSSG. Successful candidates are expected to (1) teach at the undergraduate and graduate levels, (2) develop and grow vigorous programs of externally funded research, and (3) strengthen the CS program's research areas and enhance the School of Engineering's strategic directions in health care and medicine, security, and energy and natural resources.

Vanderbilt University is a private, internationally renowned research university located in vibrant Nashville, Tennessee. Its 10 schools share a single cohesive campus that nurtures interdisciplinary activities. The university has a student body of over 12,000 undergraduate, graduate, and professional students, including 25.3% minority students and 1,169 international students from 84 countries. The School of Engineering currently comprises 90 tenured and tenure-track faculty, operates with an annual budget of $100 million including over $70 million of externally funded research, and teaches 1,400 undergraduate and nearly 500 graduate students. The School of Engineering over the past decade has been on a strong upward trajectory in national and international stature and prominence, and is in the process of building infrastructure to support a significant expansion in faculty size over the next five years. In the 2015 rankings of graduate engineering programs by U.S. News & World Report, the School ranks third among programs with fewer than 100 faculty members. With a metro population of approximately 1.5 million people, Nashville has been named one of the 15 best U.S. cities for work and family by Fortune magazine, was ranked as the #1 most popular U.S. city for corporate relocations by Expansion Management magazine, and was named by Forbes magazine as one of the 25 cities most likely to have the country's highest job growth over the coming five years. Major industries include tourism, printing and publishing, manufacturing technology, music production, higher education, finance, insurance, automobile production and health care management.

Vanderbilt University is an equal-opportunity, affirmative-action employer that aspires to become a leader among peer institutions in making meaningful and lasting progress in responding to the needs and concerns of women and members of under-represented minority groups. Applications should be submitted on-line at: https://academicjobsonline.org/ajo/jobs/6129. For more information, please visit our web site: http://engineering.vanderbilt.edu/eecs/. Applications will be reviewed on a rolling basis, with an initial deadline of December 1, 2015, but will be accepted until the positions are filled.

UNIVERSITY of WASHINGTON | BOTHELL
SCHOOL OF SCIENCE, TECHNOLOGY, ENGINEERING AND MATHEMATICS

## University of Washington, Bothell

### Division of Computing and Software Systems, School of STEM — Multiple Lecturers or Senior Lecturers in Computer Science/Software Engineering (AA13785)

The Computing and Software Systems (CSS) Division of the School of Science, Technology, Engineering and Mathematics (STEM) at the University of Washington Bothell (UWB) invites applications for multiple Lecturer or Senior Lecturer positions on a full-time, nine month, 1–3 year renewable basis beginning in Autumn 2016. Faculty duties include teaching and mentoring undergraduate and graduate students, including capstone projects/cooperative education. We are interested in candidates who can contribute to the core of our computer science and software engineering curriculum, including introductory programming, data structures, algorithms, software engineering, technical communications, web development, project management, design, databases, operating systems, architecture, embedded systems, and networking, as well as work with faculty and students in other STEM divisions. An interest and expertise in on-line teaching and/or demonstrated ability to teach large lecture classes would be pluses.

Required qualifications for the position include:

- a master's degree or earned doctorate, or foreign equivalent, in computer science, computer engineering, software engineering, or another relevant technical field,
- a body of work that warrants UWB appointment at the rank of Lecturer or Senior Lecturer,
- demonstrated experience in and commitment to excellence in undergraduate and graduate education, and
- experience with or commitment to working with and enhancing learning for diverse student and community populations.

CSS is among the largest and fastest growing computer science departments in the Pacific Northwest. We currently offer six degrees: a Bachelor of Science in Computer Engineering, a Bachelor of Science in Computer Science and Software Engineering, a Bachelor of Arts in Applied Computing, a Bachelor of Arts in Interactive Media Design, a Master of Science in Computer Science and Software Engineering, and a Master of Science in Cyber Security Engineering. All of our curricula are broadly-based in computer science and software engineering.

The 19 full-time CSS faculty members are excellent interdisciplinary teachers and scholars, actively conducting research in Computational Biology, Computer Graphics, Computer Science Education, Computer Vision, Cybersecurity, Digital Humanities, Embedded Systems, Human-Computer Interaction, Mobile Computing, Multimedia Database Systems, Parallel and Distributed Computing, Scientific Computing, Social Computing, Software Engineering, and Wireless Networks. All University of Washington faculty engage in teaching, research, and service.

The School of Science, Technology, Engineering and Mathematics combines all of the STEM fields in one academic area, allowing for cross-disciplinary training and project work. The School envisions being a leader in providing accessible, innovative, and effective education and research that promotes responsible engagement with our world and society. Our mission is to support and promote excellence in STEM research, scholarship, and education through commitment to our core values of collaboration, opportunity, rigor, and engagement. The School offers twelve undergraduate degrees and three graduate degrees within its four Divisions of Biological Sciences, Computing and Software Systems, Engineering and Mathematics, and Physical Sciences. The Bothell campus of the University of Washington was founded as an innovative, interdisciplinary campus within the University of Washington system — one of the premier institutions of higher education in the US. Faculty members have full access to the resources of a major research university, with the culture and close relationships with students of a small college. Situated just 20 miles from downtown Seattle in the midst of the Pacific Northwest's technology corridor — one of the largest and most dynamic software and technology industry concentrations in the world — UWB offers unmatched opportunities for collaborative work with industry and an excellent environment for developing creative approaches to teaching, research, and community collaborations.

UW Bothell has one of the most diverse student populations in Washington State: 64% of our incoming students are underrepresented minorities or first generation college attendees; 35% are Pell Grant eligible; 8% are international. We value engaged scholarship and experiential learning relevant to the diverse student populations and communities we serve.

Opportunity is a core value of the School of STEM. We believe that all students, regardless of background, should have the opportunity to succeed and become effective critical thinkers. Catalyzing the power of diversity enriches all of us by exposing us to a range of ways to understand and engage with the world, identify challenges, and to discover, design and deliver solutions. The School of STEM prepares professionals to work in an increasingly diverse and global society by promoting equity and justice for all individuals. We actively work to eliminate barriers and obstacles created by institutional discrimination. In your application, please describe your experiences with diversity in your professional work or educational experience and/ or your potential to enhance diversity in the School of STEM, and also discuss your potential to mentor and educate students who will serve diverse populations.

**How to apply:** Only complete applications will be considered. Please submit a single electronic file to **uwbcsss@uw.edu** with the subject line "AA13785 Lecturer Sr Lecturer CS and SE IEEE." The file should contain the following:: (1) a cover letter, (2) a curriculum vitae, (3) a list of a minimum of three professional references including contact information, (4) a research plan, (5) a statement of teaching philosophy, (6) evidence of teaching effectiveness, and (7) a description of your experiences with diversity in your professional work or educational experience and/ or your potential to enhance diversity in the School of STEM, and also a discussion of your potential to mentor and educate students who will serve diverse populations. Review of applications will begin on November 1, 2015; the position will remain open until filled.

For additional information, please see our website at **http://www.uwb.edu/CSS/.**

*University of Washington is an affirmative action and equal opportunity employer. All qualified applicants will receive consideration for employment without regard to race, color, religion, sex, sexual orientation, gender identity, national origin, age, protected veteran or disabled status, or genetic information.*

If you have a question about the details of this search/position please contact the hiring unit directly. For questions regarding potential disability accommodations, please contact Annie Brelsford at **abrelsford@uwb.edu**. Thank you for your interest in this position at the University of Washington.

**The Department of Computer Science and Engineering at the University of Notre Dame** invites applications for Assistant, Associate, or Full Professor. Excellent candidates in all areas will be considered.

The Department offers the Ph.D. degree and undergraduate Computer Science and Computer Engineering degrees. Faculty are expected to excel in classroom teaching and to lead highly-visible research projects that attract substantial external funding.

Applicants should submit a CV, statement of teaching and research interests, and contact information for three professional references at http://apply.interfolio.com/30991.

### Equal Opportunity Employment Statement

The University of Notre Dame seeks to attract, develop, and retain the highest quality faculty, staff and administration. The University is an Equal Opportunity Employer, and is committed to building a culturally diverse workplace. We strongly encourage applications from female and minority candidates and those candidates attracted to a university with a Catholic identity. Moreover, Notre Dame prohibits discrimination against veterans or disabled qualified individuals, and requires affirmative action by covered contractors to employ and advance veterans and qualified individuals with disabilities in compliance with 41 CFR 60-741.5(a) and 41 CFR 60-300.5(a).

## UNIVERSITY OF ALABAMA

### Tenured/Tenure-Track Faculty Position, Computer Science

### All Areas of Computer Science

The Department of Computer Science at the University of Alabama invites applications for a tenure-track faculty position at either the Assistant or Associate level, to begin August 2016. Outstanding candidates in all areas of computer science will be considered.

Located in Tuscaloosa, AL, the University of Alabama is the capstone of higher education in the State. The student-centered research institution is also the State's largest, with roughly 36,000 students. Housed in the College of Engineering, the Computer Science Department has 22 faculty members (14 tenured/tenure track faculty), and approximately 500 undergraduate and 50 graduate students.

For additional details, visit http://cs.ua.edu/ or contact Dr. Susan Vrbsky (faculty.search@cs.ua.edu). To apply, visit: http://facultyjobs.ua.edu/postings/37597. Review of applications will begin November 30 and will continue until the position is filled. The University of Alabama is an equal opportunity/affirmative action employee. Women and minority applicants are particularly encouraged to apply.

**SVC ARCHT.** (NY, NY & unanticip client sites US) Rsrch, elicit, analyze, validate & docmnt bus reqs. Undrstnd & evaluate cust bus needs fr IT infrastrctre monitoring. Provide feedback to prod dvlpmnt teams. Prtcpte in pre-sales tech discussions fr CA monitoring tools. Def customers' long term svc monitoring strategy & dvlp roadmaps. REQS: Bach deg or for equiv in Comp Sci, Math, Engg (any), Bus Admin or rel + 5 yrs prog exp in job &/or rel occup. Will also accept Master's deg or for eqiuv in Comp Sci, Math, Engg (any) Bus Admin or rel + 3 yrs exp in job off &/or rel occup. Must have exp w/ prvdng IT cnsltng services fr various bus clients incl medical, finance, commrcl & gov't; Dvlpng & dsgng architectural diagrams & implmntn plans; Dvlpng custom code in a client svc envrmnt to address interoperability issues; Provdng specific training & thought leadership to clients during implmntn proj; Dvlpng cust spec docs fr product implmntn; Freq travel to unanticip client sites thrght the US; Wrk fr home anywhere in the US when not at client sites.Send resume to: Althea Wilson, CA Technologies, One CA Plaza, Islandia, NY 11749, Refer to Requisition #110742.

**ASSISTANT VICE PRESIDENT-DATA SCIENCE.** Position available in Boston, MA. Design and analyze variably sourced data sets using quantitative methodologies, computational frameworks, and systems. Develop machine learning algorithms and probabilistic models, create prototype systems, visualizations, and web applications and use SQL and NoSQL systems. Design experiments and analyze data using mathematical modelling package R. Disseminate findings to non-technical audiences through various media. Apply: B. O'Brien, Massachusetts Mutual Life Insurance Company, 1295 State Street, Springfield, MA 01111; Please Reference Job ID: 708201800.

**ASSISTANT / ASSOCIATE / FULL RESEARCH PROFESSOR IN CYBER SECURITY.** Qatar University invites applications for research faculty positions at all levels. Candidates will cultivate and lead research projects at the KINDI Center for Computing Research in the area of Cyber Security. Qatar University offers competitive benefits package including a 3-year renewable contract, tax free salary, free furnished accommodation, and more. Apply by posting your application on the QU online recruitment system at careers.qu.edu.qa under "College of Engineering".

UNIVERSITY of WASHINGTON | BOTHELL
SCHOOL OF SCIENCE, TECHNOLOGY, ENGINEERING AND MATHEMATICS

## University of Washington, Bothell

### Division of Computing and Software Systems, School of STEM — Assistant Professor in Computational Science (AA13788)

The Computing and Software Systems (CSS) Division of the School of Science, Technology, Engineering and Mathematics (STEM) at the University of Washington Bothell (UWB) is seeking candidates for a tenure track Assistant Professor on a full-time, nine-month academic year basis beginning Autumn 2016. We are particularly interested in candidates with research and teaching interests in computational science (such as but not limited to big data, data visualization, or cloud computing). Successful candidates are expected to develop externally sponsored research programs, supervise graduate students, and teach and provide academic advising to students at all levels.

CSS is among the largest and fastest growing computer science departments in the Pacific Northwest. We currently offer six degrees: a Bachelor of Science in Computer Engineering, a Bachelor of Science in Computer Science and Software Engineering, a Bachelor of Arts in Applied Computing, a Bachelor of Arts in Interactive Media Design, a Master of Science in Computer Science and Software Engineering, and a Master of Science in Cyber Security Engineering. All of our curricula are broadly-based in computer science and software engineering.

The 19 full-time CSS faculty members are excellent interdisciplinary teachers and scholars, actively conducting research in Computational Biology, Computer Graphics, Computer Science Education, Computer Vision, Cybersecurity, Digital Humanities, Embedded Systems, Human-Computer Interaction, Mobile Computing, Multimedia Database Systems, Parallel and Distributed Computing, Scientific Computing, Social Computing, Software Engineering, and Wireless Networks. All University of Washington faculty engage in teaching, research and service.

The School of Science, Technology, Engineering and Mathematics combines all of the STEM fields in one academic area, allowing for cross-disciplinary training and project work. The School envisions being a leader in providing accessible, innovative, and effective education and research that promotes responsible engagement with our world and society. Our mission is to support and promote excellence in STEM research, scholarship, and education through commitment to our core values of collaboration, opportunity, rigor, and engagement. The School offers twelve undergraduate degrees and three graduate degrees within its four Divisions of Biological Sciences, Computing and Software Systems, Engineering and Mathematics, and Physical Sciences.

This Assistant Professor position in computational science is expected to strengthen the educational and research opportunities in computational sciences throughout the School of STEM including assisting with the development of a computational science graduate program and an analytics minor. The opportunity to collaborate with faculty in computational biology, computational physics, and mathematics/statistics/analytics is expected and encouraged.

The Bothell campus of the University of Washington was founded as an innovative, interdisciplinary campus within the University of Washington system — one of the premier institutions of higher education in the US. Faculty members have full access to the resources of a major research university, with the culture and close relationships with students of a small college. Situated just 20 miles from downtown Seattle in the midst of the Pacific Northwest's technology corridor — one of the largest and most dynamic software and technology industry concentrations in the world — UWB offers unmatched opportunities for collaborative work with industry and an excellent environment for developing creative approaches to teaching, research, and community collaborations.

UW Bothell has one of the most diverse student populations in Washington State: 64% of our incoming students are underrepresented minorities or first generation college attendees; 35% are Pell Grant eligible; 8% are international. We value engaged scholarship and experiential learning relevant to the diverse student populations and communities we serve.

Opportunity is a core value of the School of STEM. We believe that all students, regardless of background, should have the opportunity to succeed and become effective critical thinkers. Catalyzing the power of diversity enriches all of us by exposing us to a range of ways to understand and engage with the world, identify challenges, and to discover, design and deliver solutions. The School of STEM prepares professionals to work in an increasingly diverse and global society by promoting equity and justice for all individuals. We actively work to eliminate barriers and obstacles created by institutional discrimination. In your application, please describe your experiences with diversity in your professional work or educational experience and/ or your potential to enhance diversity in the School of STEM, and also discuss your potential to mentor and educate students who will serve diverse populations.

Required qualifications for the position include:

- an earned doctorate, or foreign equivalent, in computer science or another relevant technical field,
- a body of teaching and scholarship, or demonstrated promise for future work, that warrants UWB appointment as an Assistant Professor,
- demonstrated commitment to excellence in undergraduate and graduate education, and
- experience with or commitment to working with and enhancing learning for diverse student and community populations.

**To apply:** Only complete applications will be considered. Please submit a single electronic file to **uwbcsss@uw.edu** with the subject line "AA13788 Assistant Professor Computational Science IEEE" containing: (1) a cover letter, (2) a curriculum vitae, (3) a list of a minimum of three professional references including contact information, (4) a research plan, (5) a statement of teaching philosophy, (6) evidence of teaching effectiveness, and (7) a description of your experiences with diversity in your professional work or educational experience and/ or your potential to enhance diversity in the School of STEM, and also a discussion of your potential to mentor and educate students who will serve diverse populations. Review of applications will begin on November 1, 2015; the position will remain open until filled.

For additional information, please see our website at http://www.uwb.edu/CSS/.

*University of Washington is an affirmative action and equal opportunity employer. All qualified applicants will receive consideration for employment without regard to race, color, religion, sex, sexual orientation, gender identity, national origin, age, protected veteran or disabled status, or genetic information.*

If you have a question about the details of this search/position please contact the hiring unit directly. For questions regarding potential disability accommodations, please contact Annie Brelsford at **abrelsford@uwb.edu**. Thank you for your interest in this position at the University of Washington.

**ENGG PROJ MGR.** (Ewing, NJ) Coord all aspects of projects. Idntfy mkt opps. Suprt mgmt of mnfrme s'ware tech dvlpmnt thr entire prod life cycle. Act as an agile prod owner & internal quality control check for projects. Asst in making buy vs build decisions. Collab w/ bus dvlpmnt to build relationships to move prod line forward. REQS: 5 yrs of exp in job &/or rel occup. Must have exp w/Coord multi-team/multi-loc agile projcts; Evltng projcts using agile metrics incl burn down, velocity & defect count; Epic & story creation & prioritization, release & sprint planning; Internal & customer sprint reviews; Collab w/Proj Mgr, scrum team & cust rel to prod dvlpmnt; Mnfrme tech in an enterprise data center envrmnt; Enterprise s'ware dvlpmnt; Adapting new & emerging cloud tech to ensure business req are aligned fr new & existing products; Contributing to mnfrme IT mgmt s'ware strategies; Meeting w/ clients to gather reqs, trnslte those reqs to provide dvlpmnt w/ roadmaps; Addressing cust issues by reviewing reqs; replicating issues & prov spprt; Leading cross-func orgs to deliver quality prods incl mgmt & engg; Send resume to: Althea Wilson, CA Technologies, One CA Plaza, Islandia, NY 11749, Refer to Requisition #110741.

**VLOCITY** is seeking a Software Engineer in San Francisco, CA to design & imp. products that are usable, scalable, extens., & main. on the Force.com platform. Ref Job ID: 9PBVNG & send res. To T. Dilley at hiring@vlocity.com.

**SWITCHFLY, INC.** in San Francisco, CA seeks Sr. SW QA Engineer: Eval new functionality, write test scenarios & execute test plans. Review bug reports & assist w/resolution. Reqs incl. MS or for equiv in Enginrng Mgmt, CS, Enginrng or rel +3yrs exp. Mail resume

to HR at 601 Montgomery St., 17th Fl. SF, CA 94111. Include job code 74514 in reply. EOE.

**ERICSSON INC.** has an opening for the position of Project Manager in Plano, TX to manage derived turnkey projects involving multiple 3rd party vendors and multiple stakeholders from the customer side. To apply please mail resume to Ericsson Inc. 6300 Legacy Drive, R1-C12 Plano, TX 75024 and indicate applying for 15-TX-1715.

**FUTURE CARE CONSULTANTS, LLC** in NYC seeks softw dvlpr to design, dvlp & modify softw sys's, using .NET Framework, scientific analysis & math models, maintain tech architectures & frameworks. Manage & provide direction for dvlpmt team in support of bus. ops. Collab. w/ dvlprs & bus. owners in testing of new softw programs & apps. Must have Bach degree in Computer, Eng'g or Math, plus 5 yrs exp & strong knowl of X12 & HL7 Healthcare Standards req'd. Mail resume: Future Care Consultants, LLC, 545 8th Ave, Ste 840, NY, NY 10018 (ATTN: Samuel Stern).

JUNIPER NETWORKS

**Juniper Networks is recruiting for our Sunnyvale, CA office:**

**Resident Engineer Staff #24951**: Define, design, implement, and test company networking product features and functionality. Analyze, propose and implement system enhancements and network configurations by assisting inventory tracking and management in improving product performance and reliability. May work at other undetermined locations throughout the U.S. Telecommuting allowed.

**Solution Consultant Senior Staff #21627**: Match company product knowledge and Professional Services with customer requirements, conduct research and planning, and provide the customers with network solutions based on company products. Test network products and features, designs, and configurations for suitability to customer environments and use cases. May work at other undetermined worksites throughout the U.S. Relocation required.

**Technical Support Engineer #17351**: Provide technical support for secured routing products, working directly with our customers and partners, field engineers, technicians, and product support personnel in troubleshooting, repairing and debugging complex systems. Work with a highly knowledgeable group of customers and as an escalation point for other Technical Assistance Center (TAC) groups within the organization.

**Software Engineer #30135**: Design networking, kernel, TCP and IP solutions. Debug OS kernel, networking infrastructure in kernel, and networking application solutions.

**Software Engineer #29389**: Design, develop, and prototype large scale, highly available, and cloud-ready software applications.

**ASIC Engineer #35029**: Perform physical design implementation of ASICs, including block level and fullchip floor planning, placement and routing, timing closure and physical verification. Handle physical design implementation of blocks and partitions.

**Mail single-sided resume with job code # to**
**Juniper Networks**
**Attn: MS A.8.429A**
**1133 Innovation Way**
**Sunnyvale, CA 94089**


## MASSACHUSETTS INSTITUTE OF TECHNOLOGY

### FACULTY POSITIONS

The Department of Electrical Engineering and Computer Science (EECS) seeks candidates for faculty positions starting in September 2016. Appointment will be at the assistant or untenured associate professor level. In special cases, a senior faculty appointment may be possible. Faculty duties include teaching at the undergraduate and graduate levels, research, and supervision of student research. Candidates should hold a Ph.D. in electrical engineering and computer science or a related field by the start of employment. We will consider candidates with research and teaching interests in any area of electrical engineering and computer science.

Candidates must register with the EECS search website at https://eecs-search.eecs.mit.edu, and must submit application materials electronically to this website. Candidate applications should include a description of professional interests and goals in both teaching and research. Each application should include a curriculum vitae and the names and addresses of three or more individuals who will provide letters of recommendation. Letter writers should submit their letters directly to MIT, preferably on the website or by mailing to the address below. Complete applications should be received by December 1, 2015. **Applications will be considered complete only when both the applicant materials and at least three letters of recommendation are received.**

**It is the responsibility of the candidate to arrange reference letters to be uploaded at https://eecs-search.eecs.mit.edu by December 1, 2015.**

Send all materials not submitted on the website to:

Professor Anantha Chandrakasan
Department Head, Electrical Engineering and Computer Science
Massachusetts Institute of Technology
Room 38-401
77 Massachusetts Avenue
Cambridge, MA 02139

*M.I.T. is an equal opportunity/affirmative action employer.*

# Apple Inc. has the following job opportunities in Cupertino, CA:

**Financial Analyst (REQ#9T8RZ5).** Resp for data for the Worldwide Credit & Worldwide Apple Financial Services teams.

**Hardware Development Engineer (REQ#9H5V3E)** Resp for Thin Film Transistor (TFT) dsgn & process dvlpmnt for HW for dsply. Travel req. 25%.

**Software Development Engineer (REQ#9GXVKW)** Dsgn & dvlp kernel SW, systems SW, and tools for SW debgging & prfrmnce analysis.

**ASIC Design Engineer (REQ#9E9U4F)** Prfrm dsgn verification of SOC's. Read specifications of indust std interfaces & write bus fnctnl models.

**IT Senior Software Developer (REQ#9R2NLX).** Des and dev SW for IT Inventory management.

**Software Development Manager (REQ#9F637R)** Respsbl for technical mngmnt & leadership of a team of sys SW engineers implementing camera & media processing frameworks.

**Software Engineer Applications (REQ#9D2VKE).** Deliver projects aimed at improving Apple's online store.

**ASIC Design Engineer (REQ#9FLQ9B)** Des verifications of complex SOCs (System On a Chip) for consumer prdcts.

**Software Engineer Applications (REQ#9NBVBH)** Create & maintain tools used by Engineering Computer Services (ECS) and customers of ECS.

**Software Engineer Systems (multiple positions open) (REQ#9DCNPS).** Respon for algorithm develop, tuning, & optimization of real-time algorithms for cam firmware & img signal process pipeline.

**Systems Design Engineer (REQ#9FF395)** Dvlp & optimize RF automation sys used on Apple's newest prdcts including iPhones, iPods, iPads & others. Travel req. 20%.

**Software Development Engineer (REQ#9A7T3Y)** Research, dsgn, dvlp, & debug statistical pattern recog. sys. applied to speech & txt input sys.

**Software Engineer Systems (REQ#9UGSWE).** Research, dsgn, dvlp, implmnt & debug embedded SW for Ethernet ntwrkng in nxt gen products.

**Systems Design Engineer (REQ#9FER53).** Eval the latest iPad, iPhone and iPod HW & SW systems. Travel req'd 30%.

**Mechanical Quality Engineer (REQ#9NRTCA).** Contrib to dsgn of products from quality side (dsgn for quality). Modernize & update quality standards. Travel required 30%.

**Software Development Engineer (REQ#9BWUDV)** Dsgn & implmnt iOS apps, incl arch app strcture.

**ASIC Design Engineer (REQ#9CKT4Q).** Respon for IC pkg design eng including implement of physical pkg design, coordination w/ cross-func team on pkg select, feasibility study & develop.

**Operations Engineering Program Manager (REQ#9HDTVE).** Respon for long-term build capacity all progs & OEM sites.

**Software Engineer Systems (REQ#9E62Y7).** Dsgn & dvlp SW for cloud computing systems servicing Apple's eng teams.

**Software Development Engineer (REQ#9LV5DB).** Resp for SW dvlpmnt, prfrmnce analysis & imprvmnt in the WebKit browser engine.

**Software Engineer Applications (REQ#9JH27A)** Dvlp & deliver highly scalable, high prfmce, highly available mission critical enterprise apps.

**Software Engineer, Applications (REQ#9NTSAN).** Use comp apps SW dsgn, dvlpmnt, impl & maintenance to supp large-scale retail POS system.

**Engineering Project Coordinator (REQ#9KXR7W)** Dsgn & dvlp SW apps for Apple's service mgmt. space to spprt Retail & Partner channel systms.

**Software QA Engineer (REQ#9MCTVL)** Perf integration QA & lead testing efforts for cross-functional retail projects.

**Software Quality Assurance Engineer (REQ#9JNQLE).** Dev, des, exec & analyze tests for valid of thrml ctrl SW on mbl dvcs.

**Software Development Engineer (REQ#9T66RY).** Brainstorm, prototype, code, debug, & polish features for the Safari web browsers on OS X and iOS.

**Systems Design Engineer (REQ#9DSQEY).** Respon for eval latest iPad, iPhone & iPod HW & SW systs. Travel req: 30%.

**Software Development Engineer (REQ#9JDLFB).** Conduct SW Qual Qual Test to ensure qual of grdbreaking tech for lg scale sys, spoken lang, big data, & artif intelligence w/ focus on Korean user exp. Lang req: Korean.

**Systems Design Engineer (REQ#9DE23G).** Dev & deliv tech presentations for delivery in Apple exec briefing centers for customers & Apple exec as well as edu & biz leaders. Travel rq'd: 25%.

**Software Development Engineer (REQ#9PRUAP).** Diag protocol issues in lab & fld test, screen thru & debug hundreds of issues reported from certif test, IOT & fld tests, driving for contin perform imprvnts, prov input to dev team on failure scen & potential fixes.

**Software Development Engineer (REQ#9FNPBE).** Res, des, and dev compiler optimizations w focus on vectorization tech to speed-up C/C++ and Objective-C/Swift applications on iOS/OS X

**Software Engineer (REQ#9UZQMX)** Dsgn, build, & support new critical infrastructural sys. & frameworks.

**Software Development Engineer (REQ#9FSPPM)** Dvlp automation for

stability testing & stress execution for wireless techs like WiFi/BT/ Cellular/NFC.

**Software Engineer Applications (REQ#9H4MYD).** Des and dev SW for media content systems.

**Mechanical Quality Engineer (REQ#9DZVHX).** Contrib to the des. of future Apple prod's w/ a focus on des. for quality. Travel req'd 30%.

**Software Engineer Applications (REQ#9E5VSX).** Archtct innov solns while plyng a hands-on des & dev role to dlvr prods in highly avail, scalable and integ envt.

**Software Engineer Applications (REQ#9KLPMY).** Dsgn, dvlp & deploy data warehouse sltns for multiple bus groups at Apple.

## Apple Inc. has the following job opportunities in Austin, TX:

**ASIC Design Engineer (Multiple Positions Open) (REQ# 9GHU5W).** Verify complex CPU (microprocessor).

**Refer to Req# & mail resume to Apple Inc., ATTN: L.M. 1 Infinite Loop 104-1GM Cupertino, CA 95014.**

**Apple is an EOE/AA m/f/ disability/vets.**

UNIVERSITY of WASHINGTON | BOTHELL

SCHOOL OF SCIENCE, TECHNOLOGY, ENGINEERING AND MATHEMATICS

## University of Washington, Bothell

### Division of Computing and Software Systems, School of STEM — Assistant Professor in Computer Science (AA13789)

The Computing and Software Systems (CSS) Division of the School of Science, Technology, Engineering and Mathematics (STEM) at the University of Washington Bothell (UWB) is seeking candidates for a tenure track Assistant Professor on a full-time, nine-month academic year basis beginning Autumn 2016. We are particularly interested in candidates with research and teaching expertise in but not limited to machine learning, artificial intelligence, geographic information systems, operating systems, mobile computing, or video games. Successful candidates are expected to develop externally sponsored research programs, supervise graduate students, and teach and provide academic advising to students at all levels.

CSS is among the largest and fastest growing computer science departments in the Pacific Northwest. We currently offer six degrees: a Bachelor of Science in Computer Engineering, a Bachelor of Science in Computer Science and Software Engineering, a Bachelor of Arts in Applied Computing, a Bachelor of Arts in Interactive Media Design, a Master of Science in Computer Science and Software Engineering, and a Master of Science in Cyber Security Engineering. All of our curricula are broadly-based in computer science and software engineering.

The 19 full-time CSS faculty members are excellent interdisciplinary teachers and scholars, actively conducting research in Computational Biology, Computer Graphics, Computer Science Education, Computer Vision, Cybersecurity, Digital Humanities, Embedded Systems, Human-Computer Interaction, Mobile Computing, Multimedia Database Systems, Parallel and Distributed Computing, Scientific Computing, Social Computing, Software Engineering, and Wireless Networks. All University of Washington faculty engage in teaching, research, and service.

The School of Science, Technology, Engineering and Mathematics combines all of the STEM fields in one academic area, allowing for cross-disciplinary training and project work. The School envisions being a leader in providing accessible, innovative, and effective education and research that promotes responsible engagement with our world and society. Our mission is to support and promote excellence in STEM research, scholarship, and education through commitment to our core values of collaboration, opportunity, rigor, and engagement. The School offers twelve undergraduate degrees and three graduate degrees within its four Divisions of Biological Sciences, Computing and Software Systems, Engineering and Mathematics, and Physical Sciences.

The Bothell campus of the University of Washington was founded as an innovative, interdisciplinary campus within the University of Washington system — one of the premier institutions of higher education in the US. Faculty members have full access to the resources of a major research university, with the culture and close relationships with students of a small college. Situated just 20 miles from downtown Seattle in the midst of the Pacific Northwest's technology corridor — one of the largest and most dynamic software and technology industry concentrations in the world — UWB offers unmatched opportunities for collaborative work with industry and an excellent environment for developing creative approaches to teaching, research, and community collaborations.

UW Bothell has one of the most diverse student populations in Washington State: 64% of our incoming students are underrepresented minorities or first generation college attendees; 35% are Pell Grant eligible; 8% are international. We value engaged scholarship and experiential learning relevant to the diverse student populations and communities we serve.

Opportunity is a core value of the School of STEM. We believe that all students, regardless of background, should have the opportunity to succeed and become effective critical thinkers. Catalyzing the power of diversity enriches all of us by exposing us to a range of ways to understand and engage with the world, identify challenges, and to discover, design and deliver solutions. The School of STEM prepares professionals to work in an increasingly diverse and global society by promoting equity and justice for all individuals. We actively work to eliminate barriers and obstacles created by institutional discrimination. In your application, please describe your experiences with diversity in your professional work or educational experience and/ or your potential to enhance diversity in the School of STEM, and also discuss your potential to mentor and educate students who will serve diverse populations.

Required qualifications for the position include:

- an earned doctorate, or foreign equivalent, in computer science or another relevant technical field,
- a body of teaching and scholarship, or demonstrated promise for future work, that warrants UWB appointment as an Assistant Professor,
- demonstrated commitment to excellence in undergraduate and graduate education, and
- experience with or commitment to working with and enhancing learning for diverse student and community populations.

**To apply:** Only complete applications will be considered. Please submit a single electronic file to **uwbcsss@uw.edu** with the subject line "AA13789 Assistant Professor Computer Science IEEE" containing: (1) a cover letter, (2) a curriculum vitae, (3) a list of a minimum of three professional references including contact information, (4) a research plan, (5) a statement of teaching philosophy, (6) evidence of teaching effectiveness, and (7) a description of your experiences with diversity in your professional work or educational experience and/ or your potential to enhance diversity in the School of STEM, and also a discussion of your potential to mentor and educate students who will serve diverse populations. Review of applications will begin on November 1, 2015; the position will remain open until filled.

For additional information, please see our website at **http://www.uwb.edu/CSS/**.

*University of Washington is an affirmative action and equal opportunity employer. All qualified applicants will receive consideration for employment without regard to race, color, religion, sex, sexual orientation, gender identity, national origin, age, protected veteran or disabled status, or genetic information.*

If you have a question about the details of this search/position please contact the hiring unit directly. For questions regarding potential disability accommodations, please contact Annie Brelsford at **abrelsford@uwb.edu**. Thank you for your interest in this position at the University of Washington.

# State University of New York at Buffalo

## Department of Computer Science and Engineering

The State University of New York at Buffalo Department of Computer Science and Engineering invites candidates to apply for **multiple tenured and tenure-track faculty positions** beginning in the 2016-2017 academic year. We invite candidates from all areas of computer science and engineering, especially Computer Vision and Pattern Recognition, Artificial Intelligence, all aspects of Big Data, including Cyber Security, Cyber Physical Systems (or Internet of Things) and Human Computer Interaction. Applicants must have a Ph.D. in computer science or a related area by August 2016 and demonstrate potential for excellence in research, teaching, service and mentoring. We are looking for candidates who can operate effectively in a diverse community of students and faculty and share our vision of keeping all constituents reach their potential.

Applications will be accepted from **October 15, 2015** to **January 15, 2016.** Applicants must submit their application electronically via www.ubjobs.buffalo.edu.Any questions can be directed to Search Committee Co-Chairs, Prof. Rohini Srihari and Chang-wen Chen at cse-recruit@buffalo.edu. The University at Buffalo is an Equal Opportunity Employer.

**Computer Science and Engineering Department**

Housed in a new $75M building, and as a part of School of Engineering and Applied Sciences, the Computer Science and Engineering department offers both BA and BS degrees in Computer Science and a BS in Computer Engineering (accredited by ABET), a combined 5-year BS/MS program, a minor in Computer Science, and two joint programs (BA/MBA and Computational Physics) as well as MS and PhD programs.

The department has 37 tenured/tenure-track faculty, 4 teaching faculty, and approximately 750 undergraduate majors, 470 masters students, and 160 PhD students. Eighteen faculty including 16 junior faculty have been hired since 2010, and we are continuing to expand. Two members of our faculty currently hold key university leadership positions and seven members of our faculty are IEEE and/or ACM Fellows. Our faculty members are actively involved in cutting-edge research and successful interdisciplinary programs and centers devoted to biometrics; bioinformatics; biomedical computing; computational and data science and engineering, document analysis and recognition; high performance computing; information assurance and cyber security; embedded, networked and distributed systems, and sustainable transportation. Our annual research expenditure exceeds 4.6 Millions on average over the last five years.

**The State University of New York at Buffalo (UB)**

The State University of New York at Buffalo (UB) is New York's largest and most comprehensive public university, with approximately 20,000 undergraduate students and 10,000 graduate students.

**City and Region**

Buffalo is the second largest city in New York state, and was rated the 10th best place to raise a family in America by Forbes magazine in 2010 due to its short commutes and affordability. Located in scenic Western New York, Buffalo is near the world-famous Niagara Falls, the Finger Lakes, and the Niagara Wine Trail. The city is renowned for its architecture and features excellent museums, dining, cultural attractions, and several professional sports teams, a revitalized downtown waterfront as well as a growing local tech and start-up community. Buffalo is home to Z80, a start-up incubator, and 43 North, the world's largest business plan competition.

# Department of Computer Science and Engineering

## LECTURER POSITION AVAILABLE

**The State University of New York at Buffalo Department of Computer Science and Engineering** invites candidates to apply for a non-tenure track lecturer position beginning in Fall 2016. We invite candidates from all areas of computer science and computer engineering who have a passion for teaching to apply. The department is dedicated to the goal of building a diverse and inclusive teaching, research, and working environment. We are looking for candidates who can operate effectively in a diverse community of students and faculty and share our vision of keeping all constituents reach their potential.

The department has a strong commitment to hiring and retaining a lecturer for this career-oriented position, renewable for an unlimited number of 3-year terms. Lecturers are eligible for the in-house titles of Teaching Assistant Professor, Teaching Associate Professor and Teaching Professor. The department encourages and supports lecturers to conduct research in addition to teaching and service. Current lecturers have been conducting successful research with both internal and external funding. The department is undertaking both research and education initiatives in several exciting areas related to Internet of Things (IoT) including Big Data, cloud computing, cyber security and cyber physical systems and embedded systems. This lecturer position also has opportunities to greatly contribute to the development of new certificate and degree problems at both the undergraduate and M.S. levels in the department.

Ideally, applicants should have a PhD degree in computer science, computer engineering, or a related field, by August 15, 2016. Exceptional applicants with M.S. degree will also be considered. The ability to teach at all levels of the undergraduate curriculum is essential, as is potential for excellence in teaching, mentoring, service and research. A background in computer science and computer engineering education, a commitment to K-12 outreach, and addressing the recruitment and retention of underrepresented students are definite assets.

Duties include teaching and development of undergraduate Computer Science and Computer Engineering courses (with an emphasis on lower-division), advising undergraduate students, as well as participation in department and university governance (service). Contribution to research is encouraged.

Review of applications will begin on January 15, 2016, but will continue until the position is filled. Applicants must submit their application electronically via www.ubjobs.buffalo.edu. Any questions can be directed to Search Committee Co-Chairs, Prof. Rohini Srihari and Chang-wen Chen at cse-recruit@buffalo.edu. The University at Buffalo is an Equal Opportunity Employer.

**The Department, School and University**

Housed in a new $75M building, and as a part of School of Engineering and Applied Sciences, the Computer Science and Engineering department offers both BA and BS degrees in Computer Science and a BS in Computer Engineering (accredited by ABET), a combined 5-year BS/MS program, a minor in Computer Science, and two joint programs (BA/MBA and Computational Physics) as well as MS and PhD programs.

The department has 37 tenured/tenure-track faculty, 4 teaching faculty, and approximately 750 undergraduate majors, 470 masters students, and 160 PhD students. Eighteen faculty including 16 junior faculty have been hired since 2010, and we are continuing to expand.

The State University of New York at Buffalo (UB) is New York's largest and most comprehensive public university, with approximately 20,000 undergraduate students and 10,000 graduate students.

**City and Region**

Buffalo is the second largest city in New York state, and was rated the 10th best place to raise a family in America by Forbes magazine in 2010 due to its short commutes and affordability. Located in scenic Western New York, Buffalo is near the world-famous Niagara Falls, the Finger Lakes, and the Niagara Wine Trail. The city is renowned for its architecture and features excellent museums, dining, cultural attractions, and several professional sports teams, a revitalized downtown waterfront as well as a growing local tech and start-up community. Buffalo is home to Z80, a start-up incubator, and 43 North, the world's largest business plan competition.

**CLASSIFIED LINE AD SUBMISSION DETAILS:** Rates are $425.00 per column inch ($640 minimum). Eight lines per column inch and average five typeset words per line. Send copy at least one month prior to publication date to: Debbie Sims, Classified Advertising, *Computer* Magazine, 10662 Los Vaqueros Circle, Los Alamitos, CA 90720; (714) 816-2138; fax (714) 821-4010. Email: dsims@computer.org.

In order to conform to the Age Discrimination in Employment Act and to discourage age discrimination, *Computer* may reject any advertisement containing any of these phrases or similar ones: "...recent college grads...," "...1–4 years maximum experience...," "...up to 5 years experience," or "...10 years maximum experience." *Computer* reserves the right to append to any advertisement without specific notice to the advertiser. Experience ranges are suggested minimum requirements, not maximums. *Computer* assumes that since advertisers have been notified of this policy in advance, they agree that any experience requirements, whether stated as ranges or otherwise, will be construed by the reader as minimum requirements only. *Computer* encourages employers to offer salaries that are competitive, but occasionally a salary may be offered that is significantly below currently acceptable levels. In such cases the reader may wish to inquire of the employer whether extenuating circumstances apply.

**WhatsApp, Inc.**
currently has the following opening in
**Mountain View, CA**
**(various levels/types):**

# PRODUCT DESIGNER

**(4139J)**

Design, prototype, and build new features for WhatsApp's mobile applications.

Mail resume to: WhatsApp, Inc. c/o Facebook Inc. Attn: SB-GMI, 1 Hacker Way, Menlo Park, CA 94025. Must reference job title and job# shown above, when applying.

FIU | FLORIDA INTERNATIONAL UNIVERSITY

Florida International University is a comprehensive university offering 340 majors in 188 degree programs in 23 colleges and schools, with innovative bachelor's, master's and doctoral programs across all disciplines including medicine, public health, law, journalism, hospitality, and architecture. FIU is Carnegie-designated as both a research university with high research activity and a community-engaged university. Located in the heart of the dynamic south Florida urban region, our multiple campuses serve over 55,000 students, placing FIU among the ten largest universities in the nation. Our annual research expenditures in excess of $100 million and our deep commitment to engagement have made FIU the go-to solutions center for issues ranging from local to global. FIU leads the nation in granting bachelor's degrees, including in the STEM fields, to minority students and is first in awarding STEM master's degrees to Hispanics. Our students, faculty, and staff reflect Miami's diverse population, earning FIU the designation of Hispanic-Serving Institution. At FIU, we are proud to be 'Worlds Ahead'! For more information about FIU, visit *fiu.edu*.

The School of Computing and Information Sciences (SCIS) seeks exceptionally qualified candidates for tenure-track and tenured faculty positions at all levels as well as non-tenure track faculty positions at the level of Instructor, including visiting instructor appointments. SCIS is a rapidly growing program of excellence at the University, with 30 tenure-track faculty members and over 2,000 students, including over 80 Ph.D. students. SCIS offers B.S., M.S., and Ph.D. degrees in Computer Science, an M.S. degree in Telecommunications and Networking, an M.S. degree in Cybersecurity, and B.S., B.A., and M.S. degrees in Information Technology. SCIS has received over $22M in the last four years in external research funding, has six research centers/clusters with first-class computing and support infrastructure, and enjoys broad and dynamic industry and international partnerships.

***Open-Rank Tenure Track/Tenured Positions (Job ID# 508676)***

SCIS seeks exceptionally qualified candidates for tenure-track and tenured faculty positions at all levels. We seek well-qualified candidates in all areas; researchers in the areas of computer systems, cybersecurity, cognitive computing, data science, health informatics, and networking are particularly encouraged to apply. Preference will be given to candidates who will enhance or complement our existing research strengths.

Ideal candidates for junior positions should have a record of exceptional research in their early careers. Candidates for senior positions must have an active and proven record of excellence in funded research, publications, and professional service, as well as a demonstrated ability to develop and lead collaborative research projects. In addition to developing or expanding a high-quality research program, all successful applicants must be committed to excellence in teaching at both the graduate and undergraduate levels. An earned Ph.D. in Computer Science or related disciplines is required.

***Non-tenure track instructor positions (Job Opening 507474)***

We seek well-qualified candidates in all areas of Computer Science and Information Technology. Ideal candidates must be committed to excellence in teaching a variety of courses at the undergraduate level. A graduate degree in Computer Science or related disciplines is required; significant prior teaching and industry experience and/or a Ph.D. in Computer Science is preferred.

**HOW TO APPLY:**

Qualified candidates for open- rank faculty positions are encouraged to apply to (Job Opening ID #508676); and candidates for instructor positions are encouraged to apply to (Job Opening ID# 507474). Submit applications at *facultycareers.fiu.edu* and attach cover letter, curriculum vitae, statement of teaching philosophy, research statement, etc as *individual attachments*. Candidates will be required to provide names and contact information for at least three references who will be contacted *as determined by the search committee*. To receive full consideration, applications and required materials should be received by December 31st, 2015. Review will continue until position is filled.

If you are interested in a visiting appointment please contact the department directly by emailing Dr. Mark Weiss at *Weiss@cis.fiu.edu*. All other applicants should apply by going to *facultycareers.fiu.edu*.

***FIU is a member of the State University System of Florida and an Equal Opportunity, Equal Access Affirmative Action Employer. All qualified applicants will receive consideration for employment without regard to race, color, religion, sex, national origin, disability status, protected veteran status, or any other characteristic protected by law.***

**TIBCO DEVELOPER:** design & devel. software integration using TIBCO; define & implem. tech & arch. reqs for current & next-gen sys. MS in CS, EE or related, + 3 yrs of relev. IT exp using TIBCO OR BS+5. Email Kelly.Reinhardt@laureate.net w/ Job# 8960BR in subj. line. Laureate Education, Inc. 650 S. Exeter St., Baltimore MD 21202. EOE.

**SPLUNK INC.** has the following job opportunities in San Francisco, CA: Senior Software Engineer [REQ#8ZZ2J4]. Dsgn & dev apps supported by Co.'s core platform. Software Engineer REQ#94NU83]. Archtct & impl distr fault-tolerant data store that runs on mult ops sys (HPUX, Linux, Solaris, windows, MacOS & AIX). Refer to Req# & mail resume to Splunk Inc., ATTN: J. Aldax, 250 Brannan Street, San Francisco CA 94107. Individuals seeking employment at Splunk are considered without regards to race, religion, color, national origin, ancestry, sex, gender, gender identity, gender expression, sexual orientation, marital status, age, physical or mental disability or medical condition (except where physical fitness is a valid occupational qualification), genetic information, veteran status, or any other consideration made unlawful by federal, state or local laws. To review US DOL's EEO is The Law notice please visit: https://careers.jobvite.com/Splunk/EEO_poster.pdf. To review Splunk's EEO Policy Statement please visit: http://careers.jobvite.com/Careers/Splunk/EEO-Policy-Statement.pdf. Pursuant to the San Francisco Fair Chance Ordinance, we will consider for employment qualified applicants with arrest and conviction records.

**SENIOR SOFTWARE ENGINEER.** (mult. openings) sought by Q1W Holdings, LLC in Orlando, FL w/a BS in Comp Eng. or rltd & 5 yrs exp. Modify existing s/ware to correct errors, adapt to new h/ware, improve its performance. Dvlp & direct s/ware systm testing & validation procedures, prgmg & documentation. Confer w/systms analysts/engineers/programmers to dsgn systm & to obtain info on project limitations & capabilities, performance reqmts & interfaces. Analyze user needs & s/ware reqmts to determine feasibility of dsgn w/in time & cost restraints. Dsgn/dvlp/modify s/ware systms using scientific analysis & mathematical models to predict/measure outcome. Store/retrieve/manipulate data for analysis of systm capabilities & reqmts. Consult w/customers about s/ware systm dsgn/maintenance. Coord s/ware systm installation & monitor eqpmt functioning to ensure specifications are met. Obtain & eval info. on factors such as reporting formats req'd, costs & security needs to determine h/ware configuration. Resumes to: 1500-A Tradeport Dr., Orlando, FL 32824, Attn: HR.

**ENG. PROJECT MGR.** (Bethesda, MD) needed w/ Oracle Cert. in Java & exp. using Oracle Database, PL/SQL, OOD, OOA, Waterfall, JavaScript, & JQuery. Resume to: Overture Technologies, Inc., Attn: R. Bhamidipati, 6900 Wisconsin Avenue, Suite 200, Bethesda, MD 20815.

**DIRECTOR,** Practice Svcs (NY, NY and locations throughout the US). Mnge team of Svc Arctcts & Cnsltnts charged w/ dsgn & dlvrng access control & DLP solutions. Mnge staffing efforts & annul targets & act as SME for tech & sales inquiries. REQ: Bach Deg or for equiv in CS or Bus or rel field + 5 yrs prog exp in job &/or rel occup. Will accept a Master's deg or for equiv in CS or Bus or rel field + 3 yrs exp in job &/or rel occup. Must have exp w/ CA DataMinder solution def, archtre, dsgn & implmntn; CA Privileged Identity Manager solution def, archtre, dsgn & implmntn. Engaging

# The University of North Texas

## Department of Computer Science and Engineering Assistant/Associate Professor and Lecturer

The University of North Texas (UNT) invited applications for the following faculty positions in the **Department of Computer Science and Engineering (CSE).**

**Tenure or tenure track positions:**

Position 1: Assistant or Associate Professor in **Computer Security**, including cryptography, cloud and network security, mobile security, intrusion detection, secure hardware and software systems, and vulnerability and threat analysis.

Position 2: Assistant or Associate Professor in **Computer Systems**, including computer architecture, real-time operating systems, runtime systems and virtualization, memory and storage systems, distributed systems, embedded systems, resilient, secure and survivable systems, and performance measurement and tuning.

The candidates are expected to teach CSE undergraduate and graduate courses, develop a strong research program funded by external sources, support and mentor graduate students, and provide service to the University and the profession. Minimum qualifications include an earned doctorate in computer science, computer engineering or a closely related field. For the Assistant Professor position, a strong publication record and the potential to succeed in securing research funding and mentoring graduate students are required. For the senior positions, a sustained record of providing mentoring to junior faculty, advising graduate students, providing service to the University and profession, and securing external funding for research activities with current research funding are also required. Post-doctoral research experience or industrial research experience is preferred.

**Non-tenure track position:**

Position 3: Lecturer, Senior Lecturer, or Principal Lecturer with a three-year appointment renewable annually, depending on performance and the availability of funding. The primary responsibility is teaching at the undergraduate and graduate level. Additional expectations include participating in departmental activities, assisting with ABET accreditation, and providing career guidance to undergraduate students.

*The Computer Science and Engineering department is home to 1,027 bachelors students, 151 masters students and 90 Ph. D. students. The UNT CSE department is the home department of the interdisciplinary Center for Information and Computer Security and the lead academic institution of the NSF Industry/University Cooperative Research Center for Net-centric and Cloud Software and Systems. Additional information about the department and the centers are available at the websites: www.cse.unt.edu, www.cics.unt.edu, and netcentric.cse.unt.edu, respectively.*

**Application Procedure:**

All applicants must apply online to: https://facultyjobs.unt.edu. Submit nominations and questions regarding the tenure track position in computer security (system identification number 6001152) to Dr. Ram Dantu (Ram.Dantu@unt.edu), the tenure track position in computer systems (system identification number 6001153) to Dr. Bill Buckles (Bill.Buckles@unt.edu), and lecturer position (system identification number 6001154) to Dr. Phil Sweany (Philip.Sweany@unt.edu).

**Application Deadline:**

The committee will begin its review of applications on November 1, 2015 and continue to accept and review applications until the positions are filled.

**The University:**

As the nation's 24th largest public university and the largest, most comprehensive in the Dallas-Fort Worth area, UNT is dedicated to providing an excellent educational experience to its 37,000 students while powering the North Texas region, state and nation through innovative education and research. UNT is strategically located in Denton, Texas, a vibrant city with a lively arts and music culture.

*The University of North Texas is an AA/ADA/EOE committed to diversity in its educational programs.*

in sales process invlvng access control & data loss prevention solutions; Mngng solution specific def, archtre, dsgn and implmntn of access control & data loss prevention solutions; Freq visits to unanticipated client sites; Wrk fr home anywhere in US. Send resume to: Althea Wilson, CA Technologies, One CA Plaza, Islandia, NY 11749, Refer to Requisition 113502.

**CLOUDERA, INC.** is recruiting for our Palo Alto, CA office: Dedicated Support Manager: ensure that critical customer issues are addressed quickly & effectively. Triage, diagnose & potentially escalate customer inquiries during their engineering & operations efforts. Mail resume w/job code #36490 to: Cloudera, Attn.: HR, 1001 Page Mill Rd., Bldg. 2, Palo Alto, CA 94304.

**BIG SWITCH NETWORKS** seeks Systems Engr for its Santa Clara, CA office to design network for customers & test routing designs, features & functionality. May work from home but must be based on the East Coast. Travel 30-50% of the time to meet with unanticipated customers on the East Coast & attend meetings/training at HQ in Santa Clara, CA. Send resume w/ad to 3965 Freedom Circle #300, Santa Clara, CA 95054. Attn: HR/SG.



VirginiaTech
*Invent the Future*

# Faculty Positions Department of Computer Science

The Department of Computer Science at Virginia Tech (**www.cs.vt.edu**) seeks applicants for tenure-track faculty positions in three areas: interactive computing, cyber security, and data analytics. Candidates should have a Ph.D. in Computer Science or related field at the time of appointment, a rank-appropriate record of scholarship and collaboration in computing and interdisciplinary areas, sensitivity to issues of diversity in the campus community, and will be required to teach at the undergraduate and/or graduate levels. The position requires occasional travel to professional conferences and meetings.

**Tenure-track Assistant Professor in Interactive Computing – Blacksburg, VA**
Strong candidates from any area related to interactive computing are encouraged to apply. Exceptional candidates at higher ranks will also be considered. Candidates who complement existing strengths in human-computer interaction, graphics, intelligent user interfaces, visualization, visual analytics, human interaction with big data, augmented reality, tangible interfaces, human-robot interaction, game design, creativity support, or computing in the arts and humanities are especially encouraged. Candidates have opportunities for collaboration in the interdisciplinary Center for Human-Computer Interaction (**hci.vt.edu**) which includes 30 faculty across campus, the Institute for Creativity, Arts, and Technology (**icat.vt.edu**) housed in the new Moss Center for the Arts, and the Discovery Analytics Center (**dac.cs.vt.edu**). Applications must be submitted online to **http://jobs.vt.edu** for posting **#TR0150108**. Applicant screening will begin on November 20, 2015 and continue until each position is filled. Inquiries should be directed to **Dr. Chris North, Search Committee Chair, north@vt.edu**.

**Tenure-track Assistant Professor in Cyber Security – Blacksburg, VA**
Candidates with expertise in cyber security, including technologies for and applications in information security, network security, and trustworthy computing are encouraged to apply. Candidates focusing on security issues of cyber-physical systems, embedded systems, sensor networks, robotics, Internet of Things (IoT), etc. are especially encouraged. The candidate will join the CS department and also participate in an interdisciplinary team of five faculty in Advanced Manufacturing and share common space and equipment, leveraging established labs and the Commonwealth Center for Advanced Manufacturing ( **www.ccam-va.com/** ), a public-private partnership in Virginia. There is an active group of cyber security faculty in CS and ECE departments collaborating in research as well as graduate and undergraduate education (see: **www.cyber.vt.edu/**). Applications must be submitted online to **http://jobs.vt.edu** for posting **#TR0150107**. Applicant screening will begin on November 20, 2015 and continue until the position is filled. Inquiries should be directed to **Dr. Ali Butt, Search Committee Chair, butta@cs.vt.edu.**

**Associate/Full Professor in Data Analytics – National Capital Region (NCR)** Candidates with research depth and breadth in data analytics, data mining, "big data", or data science are encouraged to apply. Candidates working at the intersection of data analytics and cyber-security and at the intersection of data analytics and urban computing are especially encouraged. Candidates should present a proven ability to initiate and sustain collaborations within computing as well as with application specialists. The department is home to the Discovery Analytics Center (**dac.cs.vt.edu**) that leads "big data" research on campus. The successful candidate will contribute to the research and graduate programs in the NCR and collaborate with faculty at Virginia Tech's campus in Blacksburg, VA. The NCR campus (**www.ncr.vt.edu**) is located near the Washington D.C./Falls Church area and houses the Virginia Tech Research Center (**www.ncr.vt.edu/arlington**) in Arlington, VA. Applications must be submitted online to **http://jobs.vt.edu** for posting **#TR0150106.** Applicant screening will begin on November 20, 2015 and continue until the position is filled. Inquiries should be directed to **Dr. Naren Ramakrishnan, Search Committee Chair, naren@cs.vt.edu**.

**The Department of Computer Science** has 40 research oriented tenure-track faculty and ~10 postdocs/research faculty. There are a total 12 NSF/DOE CAREER awardees in the department. Research expenditures for FY2015 were $412 thousand per tenure-track faculty member (i.e., a total of $15.5 million); total research funding at the beginning of FY2015 was $43.4 million. BS, MS, and PhD degrees are offered, with a growing enrollment of over 610 undergraduate majors (14% women) and over 270 PhD/MS students. In 2010, **CS@VT** was ranked 5th in the country in recruiting quality of CS undergrads by the *Wall Street Journal*. The department is in the College of Engineering, whose undergraduate program was ranked 8th and graduate program was ranked 12th among public engineering schools in 2014 by *U.S. News & World Report*.

Early applications are encouraged. We welcome applications from women or minorities. Salary for suitably qualified applicants is competitive and commensurate with experience. Selected candidates must pass a criminal background check prior to employment.

*Virginia Tech is an AA/EEO employer; applications from members of underrepresented groups are especially encouraged.*

# Samsung Research America, Inc.

has the following opportunities (various levels) available in **Mountain View, CA**:

**Sr. Industrial Designer (Ref# MTV15G01)**

**Systems Design/Architecture Engineer (Ref# MTV15G02)**

**Industrial Designer (Ref# MTV15G03)**

**Software Engineer (Ref# MTV15H01)**

**Software Engineer, Staff 1 (Ref# MTV15H02)**

**Software Engineer, Senior (Ref# MTV15H03)**

**Software Engineer, Sr. Staff 1 (Ref# MTV15H04)**

**Research Engineer (Ref# MTV15H05)**

**Sr. Research Engineer (Ref# MTV15H06)**

**Research Engineer, Staff 1 (MTV15H07)**

**Research Engineer, Staff 2 (MTV15H08)**

**Interaction Designer, Staff 1 (Ref# MTV15J01)**

**Interaction Designer, Staff 2 (Ref# MTV15J02)**

**Sr. Product Manager (Ref# MTV15E05)**

**Staff Engineer (Ref# MTV15J04)**

**Sr. Graphics Driver Engineer (Ref# MTV15J05)**

**Sr. UX Researcher (Ref# MTV15J06)**

**Hardware Engineer, Sr. Staff 1 (Ref# MTV15J07)**

Specific requirements apply. All of these positions will involve developing technologies for company's computer, digital television, mobile telephone, printer, or other electronic products. Mail your resume referencing job title and Ref# to farhat.k@samsung.com.

Help build the next generation of systems behind Facebook's products.

**Facebook, Inc.** currently has the following openings in **Menlo Park, CA (various levels/types)**:

**Partner Engineer (165J)** Work with strategic partners to help them build, grow, & monetize their products using the Facebook platform. Drive engineering effort, commuincate cross-functionality, act as subject matter expert & a solutions architect to parterns. **Manager, Data Engineering (6327J)** Proactively drive the vision for Business Intelligence & Data Warehousing across the company, & define & execute on a plan to achieve that vision. **Audience Insights Analyst (5373J)** Apply your expertise in quantitative analysis, data mining, & the presentation of data to uncover unique actionable insights about people, events & media. **Technical Program Engineer (2249J)** Develop & handle end-to-end IT project plans & ensure on-time delivery of technical solutions. **Application Engineer, Hyperion (5990J)** Design & develop Hyperion systems. Enhance Hyperion applications for budget, forecast & long range plan for financial planning & analysis (FP&A). **Data Engineer (DE915J)** Use data to influence decisions made about the development of Facebook products. **Front End Engineer (4296J)** Work with Product Designers to implement the next generation of Company's products. Build efficient & reusable front-end abstractions & systems. Identify & address performance bottlenecks. **Quality Assurance Lead (3775J)** Execute manual & automated tests, & identify actionable bugs quickly. Handle QA coverage of multiple mobile based projects. **Software Engineer (4533J)** Apply expertise in materials science & engineering to evaluate & create new 3D sensor technologies. **Community Operations Specialist (3015NJ)** Understand, identify, & investigate the trends underlying our operational metrics to drive optimizations.

Facebook, Inc. currently has the following openings in **New York, NY (various levels/types)**:

**Global Client Partner (5355J)** Earn a place as a trusted advisor to our most senior clients by partnering with them to understand their business objectives & help define where & how Facebook can play a role in achieving them. Domestic &/or international travel may be required.

Mail resume to: Facebook, Inc. Attn: SB-GIM, 1 Hacker Way, Menlo Park, CA 94025.
Must reference job title & job# shown above, when applying.

IEEE computer society

# ROCK STARS OF CYBER SECURITY

**Chris Calvert**
Global Director, HP
Enterprise Solutions Products

**Marcus H. Sachs**
Senior Vice President,
Chief Security Officer (NERC)

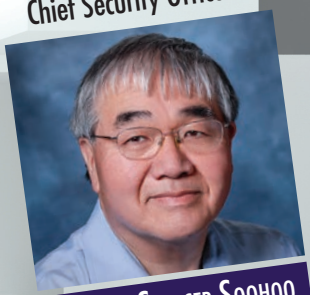

**Dr. Spencer Soohoo**
CSO/Director, Scientific Computing
Cedars-Sinai Medical Center

## Win the New Cybersecurity War with the New Rock Stars of Cybersecurity

Cybercrime is no longer a matter of credit card breaches. Cybercriminals are now trying to take down countries as well as top companies. Keep your organization safe. Come to the premier, one-day, high-level event designed to give real, actionable solutions to these cybersecurity threats.

**Learn from and collaborate with the experts—**

**27 October 2015**
**The Fourth Street Summit Center**
**San Jose, CA**

**REGISTER NOW**

**Early Discount Pricing Now Available!**

**computer.org/cyber2015**

IEEE

TECHNOLOGY

# LinkedIn Corp.

LinkedIn Corp. has openings in our **Mtn View, CA** location for:

**Software Engineer (All Levels/Types) (SWE915MV)** Design, develop & integrate cutting-edge software technologies; **Senior Software Engineer (6597.700)** Design, develop, test and integrate software technologies; **Test Engineer (6597.873)** Coordinate & execute the localization testing process for all languages except English; **Web Developer (6597.882)** Build rich, dynamic client-side interfaces using Javascript, leveraging new technologies like HTML5 & CSS3, to launch fast/scalable products for global users; **Software Engineering Manager (6597.37)** Lead a team of engineers in the design, development & integrating cutting edge software technologies; **Manager, Test Engineering (6597.96)** Propose & build investments in automation frameworks & tools, to drive improvements & efficiencies across the larger test engineering organization.

LinkedIn Corp. has openings in our **Sunnyvale, CA** location for:

**Software Engineer (All Levels/Types) (SWE915SV)** Design, develop & integrate cutting-edge software technologies; **Staff Software Engineer (6597.802)** Develop innovative new technologies, features, & products that help connect the world's professionals, increasing their productivity & success; **Test Engineer (6597.1171)** Write & build automated test suites, continuously design creative ways to break software, & identify potential bugs; **Senior Systems Engineer, Applications Operation (6597.1357)** Work closely with developers & application owners to ensure stability for LinkedIn's internal web & standalone applications. **Developer, Sales Systems (6597.688)** Build customized solutions in Salesforce.com & associated systems that support LinkedIn's business requirements & processes.

LinkedIn Corp. has openings in our **San Francisco, CA** location for:

**Technical Services Manager (6597.745)** Oversee the technical support of company customers.

LinkedIn Corp. has openings in our **Carpinteria, CA** location for:

**Lead Informatica Engineer (6597.1336)** Design, develop, maintain, monitor, & optimize PowerCenter & Informatica Cloud Services ETLs in support of company's website & data management systems.

Please email resume to: 6597@linkedin.com. Must ref. job code above when applying.

## TECHNOLOGY

**Intuit Inc.** currently has openings for the following positions in **Santa Clara County, including Mountain View, California** or any office within normal commuting distance:

**Staff Data Engineers (Job code: I-386):** Use technical expertise to develop code and unit test for software and/or analyze user needs and/or software requirements to determine required software improvements and/or modifications. Design, develop, and implement data movement and integration processes in preparation for analysis, data warehousing, and operational data stores, involving very large quantities of data. **Staff Software Engineers (Job code: I-177):** Use technical expertise to develop code and unit test for software and/or analyze user needs and/or software requirements to determine required software improvements and/or modifications. **Data Scientists (Job code: I-101):** Provides guidance and support leadership to Business leaders and stakeholders on how best to harness available data in support of critical business needs and goals. Leads the full cycle of iterative big data exploration, including hypothesis formulation, algorithm development, data cleansing, testing, insight generation/visualization, and action planning. **Senior Business Data Analysts (Job code: I-1837):** Collaborate with the data warehousing team, ensuring that data infrastructure supports the needs of Intuit's analytics team and validating data quality. **Staff Software Engineers (Job code: I-46):** Exercise senior level knowledge in selecting methods and techniques to design, implement, modify and support a variety of software products and services to meet user or system specifications. May telecommute from home 10% of the time. **Managers, Group Research and Analysis (Job code: I-288):** Define the roadmap to achieve strategies that will drive quality product experiences for customers and will accelerate business growth.

Positions located in **San Diego, California**:

**Senior Software Engineers in Quality (Job code: I-371):** Apply senior level software engineering practices and procedures to design, influence, and drive quality and testability of products and services. **Software Engineers (Job code: I-395):** Apply software development practices to design, implement, and support individual software projects. **Staff Software Engineers (Job code: I-330):** Partner with cross-functional leaders and team members to deliver Intuit products, with greater efficiency and speed. Test software, including creating test cases, test plans, test data, and defect write ups. **Software Engineers in Quality (Job code: I-468):** Apply best software engineering practices to ensure quality of products and services by designing and implementing test strategies, test automation, and quality tools and processes. **Senior Application Operations Engineers (Job code: I-339):** Responsible for ensuring all apps are ready for Scalability, Availability, Monitoring, Ops Support, Functionality (SAMOF).

To apply, submit resume to Intuit Inc., Attn: Olivia Sawyer, J203-6, 2800 E. Commerce Center Place, Tucson, AZ 85706. You must include the job code on your resume/cover letter. Intuit supports workforce diversity.

#CES2016


SEEING WHAT'S NEXT IS GREAT, BUT SEEING WHAT'S AFTER NEXT IS WHERE LIMITLESS OPPORTUNITY RESIDES. DISCOVER BOTH AND BEYOND AT CES. REGISTER NOW.

**CES® 2016 JAN. 6-9, 2016**

TECH EAST • TECH WEST • TECH SOUTH
LAS VEGAS, NV

REGISTER NOW at **CESweb.org**

CES®
THE GLOBAL STAGE FOR INNOVATION

PRODUCED BY CEA