COMPUTING edge

SECURITY

Also in this issue:

- Multimedia Hashing and Networking
- Reflecting on Quality

NOVEMBER 2016

www.computer.org

IEEE computer society

CELEBRATING 70 YEARS


HISCOX
business insurance®
“I don’t worry about getting sued. I plan for it.”
Even if you haven’t made a mistake, defending a lawsuit can be a big cost for your business. Customized insurance from Hiscox can help you keep moving forward with confidence.
Get a fast, free quote at Hiscox.com/planonit or call our licensed insurance agents at 866-930-1054 Mon–Fri, 8:00am–10:00pm ET. Your policy could start as low as $22.50/mo.
#encouragecourage.
© 2016 Hiscox Inc. All rights reserved.

NOVEMBER 2016 • VOLUME 2, NUMBER 11

COMPUTING edge

26

Addressing Pressing Cybersecurity Issues through Collaboration

Departments

Subscribe to ***ComputingEdge*** for free at **www.computer.org/computingedge.**

COMPUTING edge

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

Printed with inks containing soy and/or vegetable oils

IEEE computer society
CELEBRATING 70 YEARS

STAFF

Editor
Lee Garber

Contributing Staff
Christine Anthony, Lori Cameron, Carrie Clark, Chris Nelson, Meghan O'Dell, Dennis Taylor, Bonnie Wylie

Production & Design
Carmen Flores-Garvey, Monette Velasco, Jennie Zhu-Mai, Mark Bartosik

Senior Manager, Editorial Services
Robin Baldwin

Director, Products and Services
Evan Butterfield

Senior Advertising Coordinator
Debbie Sims

BPA WORLDWIDE

Circulation: ComputingEdge (ISSN 2469-7087) 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 address changes to ComputingEdge-IEEE Membership Processing Dept., 445 Hoes Lane, Piscataway, NJ 08855. Periodicals Postage Paid at New York, New York, and at additional mailing offices. Printed in USA.

Editorial: Unless otherwise stated, bylined articles, as well as product and service descriptions, reflect the author's or firm's opinion. Inclusion in ComputingEdge does not necessarily constitute endorsement by the IEEE or the Computer Society. All submissions are subject to editing for style, clarity, and space.

Reuse Rights and Reprint Permissions: Educational or personal use of this material is permitted without fee, provided such use: 1) is not made for profit; 2) includes this notice and a full citation to the original work on the first page of the copy; and 3) does not imply IEEE endorsement of any third-party products or services. Authors and their companies are permitted to post the accepted version of IEEE-copyrighted material on their own Web servers without permission, provided that the IEEE copyright notice and a full citation to the original work appear on the first scree n of the posted copy. An accepted manuscript is a version which has been revised by the author to incorporate review suggestions, but not the published version with copy-editing, proofreading, and formatting added by IEEE. For more information, please go to: http://www.ieee .org/publications_standards/publications/rights/paperversionpolicy.html. Permission to reprint/republish this material for commercial, advertising, or promotional purposes or for creating new collective works for resale or redistribution must be obtained from IEEE by writing to the IEEE Intellectual Property Rights Office, 445 Hoes Lane, Piscataway, NJ 08854-4141 or pubs-permissions@ieee .org. Copyright © 2016 IEEE. All rights reserved.

Abstracting and Library Use: Abstracting is permitted with credit to the source. Libraries are permitted to photocopy for private use of patrons, provided the per-copy fee indicated in the code at the bottom of the first page is paid through the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923.

Unsubscribe: If you no longer wish to receive this ComputingEdge mailing, please email IEEE Computer Society Customer Service at help@computer.org and type "unsubscribe ComputingEdge" in your subject line.

IEEE prohibits discrimination, harassment, and bullying. For more information, visit www.ieee.org/web/aboutus/whatis/policies/p9-26.html.

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
Ahmad-Reza Sadeghi, *Technical University of Darmstadt*

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*

IEEE computer society

IEEE

2017 B. Ramakrishna Rau Award

Call for Nominations

Honoring contributions to the computer microarchitecture field

New Deadline: 1 May 2017

Established in memory of Dr. B. (Bob) Ramakrishna Rau, the award recognizes his distinguished career in promoting and expanding the use of innovative computer microarchitecture techniques, including his innovation in complier technology, his leadership in academic and industrial computer architecture, and his extremely high personal and ethical standards.

WHO IS ELIGIBLE? The candidate will have made an outstanding innovative contribution or contributions to microarchitecture, use of novel microarchitectural techniques or compiler/architecture interfacing. It is hoped, but not required, that the winner will have also contributed to the computer microarchitecture community through teaching, mentoring, or community service.

AWARD: Certificate and a $2,000 honorarium.

PRESENTATION: Annually presented at the ACM/IEEE International Symposium on Microarchitecture

NOMINATION SUBMISSION: This award requires 3 endorsements. Nominations are being accepted electronically: www.computer.org/web/awards/rau

CONTACT US: Send any award-related questions to awards@computer.org

www.computer.org/awards

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

Advances in medicine and clinical care are increasingly tied to computing technologies. *Computer*'s October 2016 special issue explores emerging trends in **smart health** and the benefits they bring to individual patients and to society as a whole.

IEEE Software

Developers of **systems of systems** (SoSs) face numerous challenges. The authors of "Monitoring Requirements in Systems of Systems," from *IEEE Software*'s September/October 2016 issue, discuss their ReMinds tool, designed to meet these challenges by instrumenting SoSs to extract events and data at runtime. ReMinds then defines requirements as constraints to check for expected behavior and properties.

IEEE Internet Computing

Measuring spam's cost for users and network operators and identifying who pays for it is difficult. In "Measuring, Characterizing, and Avoiding Spam Traffic Costs," from *IEEE Internet Computing*'s July/August 2016 issue, the authors provide a way to quantify those costs. They show that stub networks incur high spam-traffic costs but that some networks actually profit from spam traffic and might not want to filter it. They then present an algorithm to identify networks that would benefit from cooperating to filter such traffic.

Computing in Science & Engineering

As scientific-data volumes continue to grow, researchers increasingly need a flexible computational infrastructure that can support the entire data-science lifecycle. The authors of "A Case for Data Commons: Toward Data Science as a Service," from *CiSE*'s September/October 2016 issue, explain their development of an **interoperable data commons infrastructure** that collocates data, storage, and computing resources with common

Published by the IEEE Computer Society
2469-7087/16/$33.00 © 2016 IEEE

analysis tools. Challenges remain, but development of such an infrastructure brings us one step closer to data science as a service for the scientific research community.

I*EEE Security & Privacy*

The gaming industry collects participants' data to generate marketing-related revenue and improve the playing experience. However, the need to protect player privacy complicates this process. "**Incorporating Privacy into Digital Game Platform Design**: The What, Why, and How," from *IEEE S&P*'s July/August 2016 issue, details an iterative approach that includes privacy-by-design principles in game development.

IEEE Cloud Computing

IEEE Cloud Computing's July/August 2016 special issue includes articles on topics such as the economics and strategy of **manufacturing and the cloud**; and cloud manufacturing's security, privacy, and forensic concerns.

IEEE Computer Graphics and Applications

"**Designing for Insight**: A Case Study from Tennis Player Analysis," which appears in *CG&A*'s July/August 2016 issue, describes a combinatory design process that uses incremental addition to generate increasingly complex data arrangements and thus create new ways to see the information and discover new insights about topics being analyzed.

IEEE Intelligent Systems

"Design of a **Multiagent System for Real-Time Traffic Control**," from *IEEE Intelligent Systems*' July/August 2016 issue, examines the various steps involved in analyzing and designing such a system for use at isolated street intersections. In the authors' model, the many agents designed for isolated intersections create, manage, and evolve their own traffic-signal plans.

IEEE MultiMedia

State-of-the-art hashing techniques are widely used in high-efficiency multimedia storage, indexing, and retrieval. The authors of "**Multimedia Hashing and Networking**," from *IEEE MultiMedia*'s July–September 2016 issue, summarize shallow-learning-based and deep-learning-based hashing, and introduce multimedia information networks as a way to incorporate both visual and textual information to make deep learning practical in multimedia applications.

IEEE Annals of the History of Computing

"Two **Early Interactive Computer Network Experiments**," from *IEEE Annals*' July–September 2016 issue, looks at experiments that joined a System Development Corp. time-sharing computer with a system at the Stanford Research Institute in 1963 and with one at MIT Lincoln Laboratory in both 1966 and 1967.

IEEE Pervasive Computing

The authors of "Displays as a Material: A Route to **Making Displays More Pervasive**," from *IEEE Pervasive Computing*'s July–September 2016 issue, advocate using an architecture that relies on autonomous pixels that independently sense input and convert it to a visual output. They also discuss their two display prototypes.

IT Professional

Rich visual information is becoming increasingly important in today's web. In "**Visual Information Retrieval**: The State of the Art," from *IT Pro*'s July/August 2016 issue, the author examines the process of searching and retrieving images using a visual query, also called content-based image retrieval.

IEEE Micro

"Ten Open Questions for Techno-Optimists," from *IEEE Micro*'s July/August 2016 issue, discusses some of the open questions regarding **productivity growth and economic gains** resulting from innovative IT.

Computing Now

The Computing Now website (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.

Hardware-Enforced Privacy

This installment highlighting the work published in IEEE Computer Society journals comes from IEEE Computer Architecture Letters.

Simha Sethumadhavan, Columbia University

Privacy is important to modern, civilized life: it's been described as a fundamental human right, a key to democratic processes, and—as evidenced by the rise of data brokers and markets—a driver of economic growth. Yet support for privacy isn't considered a requirement in the design of hardware systems. In our paper "Hardware Enforced Statistical Privacy," Matthew Maycock and I describe the motivation for and benefits of enforcing privacy in hardware (*IEEE Computer Architecture Letters*, vol. 15, no. 1, 2016, pp. 21–24).

The backdrop for our work is the current Internet of Things era. Today's electronic devices have both computational power and the means to communicate not only with one other but also with remote third parties. Such devices' increasing ubiquity has resulted in swaths of disparate user data that's used to construct fine-grained behavioral profiles. These profiles can serve both noble causes, such as improving healthcare, and nefarious activities, such as illegal surveillance and targeting. As such, it's vital that users have the tools necessary to control their privacy. Complicating privacy control, however, is the fact that users might want to give different-quality data to different parties, for example, doctors versus insurance companies or advertisers.

We argue that relegating such complex privacy needs to purely software solutions isn't desirable. Software is mutable, generally buggier than hardware, might have coverage holes due to heterogeneity and layering, and might implement incorrect privacy notions. Hardware, however, is immutable and can sit between data sources (sensors) and data consumers (software accessing the data), guaranteeing coverage and a universal, minimum notion of privacy. Additionally, hardware combined with other trusted hardware mechanisms can provide proof of privacy—which isn't easily achieved with software implementations alone.

We've proposed a hardware privacy protection unit (PPU): a standalone chip or an integrated hardware module that sits between the sensors and the remaining system (see Figure 1). The PPU can mediate data access, access sensor data, fuzz data, log overall transactions, and raise hardware exceptions for any issue at any step. The PPU can also be programmed with different policies. The type of fuzzing we studied was ε-differential privacy, a form of statistical privacy that involves noising the sensor outputs to provide strong mathematical guarantees.

Hardware-enforced privacy offers a cleaner way to engineer privacy into computing platforms. If successful, our approach could lead to a firmer foundation for reaping the rewards of the Internet of Things era.

SIMHA SETHUMADHAVAN is an associate professor of computer science at Columbia University. Contact him at simha@columbia.edu.

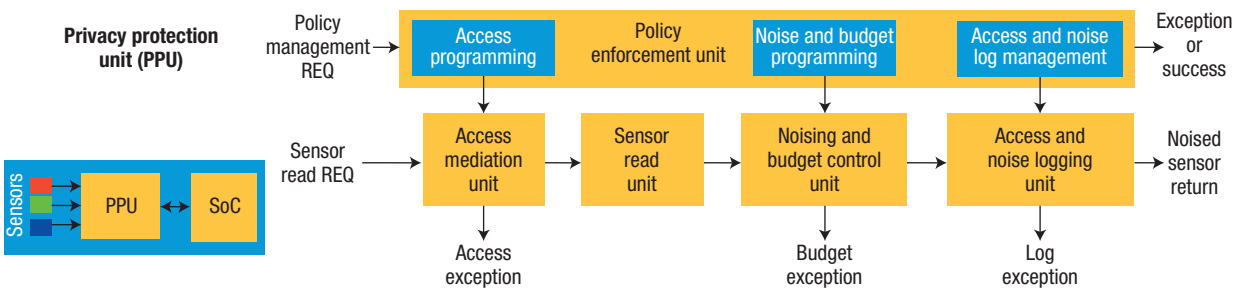

Figure 1. Hardware privacy protection unit. SoC: system on chip.

 Published by the IEEE Computer Society 2469-7087/16/$33.00 © 2016 IEEE

EDITOR'S NOTE

Securing Today's Systems

Cybersecurity is a field that touches every aspect of computing and communications: hardware, software, networking, big data, the Internet of Things, wireless technology, and so on. This *ComputingEdge* issue looks at various aspects of this critical topic.

No system is immune from attack, so cybersecurity advances should take a backseat to improving human awareness of security risks, according to *Computer*'s " 'Good Enough' Security: The Best We'll Ever Have."

Adversarial samples are slightly modified inputs crafted to mislead and corrupt systems whose security is based on computationally learned models. The authors of *IEEE Security & Privacy*'s "Machine Learning in Adversarial Settings" examine adversarial samples and potential future countermeasures.

Existing bot-detection approaches balance proof of human identity with workflow unobtrusiveness. According to "Natural Interaction for Bot Detection" from *IEEE Internet Computing*, cognitive modeling and natural interaction might provide stronger security and even less intrusiveness.

The US National Institute of Standards and Technology's National Cybersecurity Center of Excellence (NCCoE) focuses on building reference designs and example solutions that can improve security. The author of *IT Professional*'s "Addressing Pressing Cybersecurity Issues through Collaboration" discusses how NCCoE works with the private sector, academia, and other government organizations to identify pressing security problems.

Building a secure system requires rigorous threat analysis followed by systematic transformation of identified threats into security-related requirements, which can be tracked throughout the development lifecycle, explains "Keeping Ahead of Our Adversaries," from *IEEE Software*.

In *IEEE Security & Privacy*'s "Silver Bullet Talks with Jacob West," West—chief architect for security products at NetSuite, a vendor of cloud-based enterprise-software—discusses secure design, the critical difference between bugs and flaws, and wearable device security.

ComputingEdge articles on topics other than cybersecurity include the following:

- The authors of *IEEE MultiMedia*'s "Multimedia Hashing and Networking" summarize shallow- and deep-learning-based hashing, and introduce multimedia information networks as a way to incorporate both visual and textual information to make deep learning practical in multimedia applications.
- *Computing in Science & Engineering*'s "The Future of NSF Advanced Computing Infrastructure Revisited" examines a US government report on possible future directions for the National Science Foundation's study of an advanced computing infrastructure that would support domestic science and engineering.
- "Reflecting on Quality," from *IEEE Software*, looks at tools and techniques that provide the data-driven quality management necessary for effective software production.

2469-7087/16/$33.00 © 2016 IEEE Published by the IEEE Computer Society November 2016

"Good Enough" Security: The Best We'll Ever Have

George Hurlburt, STEMCorp

Given that no system is completely immune from attack, cybersecurity advances, while important, should take a backseat to improving human awareness of security risks. An objective, consensus-based rating system is one means to achieve this end.

"Good morning. Your mission, should you choose to accept it, is to design a totally secure computer, impenetrable by anyone or anything. This message will self-destruct in five seconds." Cue the *Mission Impossible* theme!

What's a totally secure computer? It's likely not connected to the Internet, making it immune to network attacks. For that matter, it likely has no external ports at all, preventing any form of malicious manipulation from the outside world. Power would likely be supplied by internal fuel cells, eliminating yet another subtle path for external exploitation. Most importantly, it would have no user interfaces, much less any human accessibility, thus bypassing the prevalent insider threat. Ideally, it would be encased in reinforced concrete and buried deep in the ground to harden it against the ravages of man or nature. Finally, what little software is allowed to run on it would be pure assembly language, or better yet, firmware—thoroughly scrubbed of any malicious content. In short, a totally secure computer would serve absolutely no useful purpose whatsoever.

Given that total security is impossible, what security is good enough? "Good enough" clearly constitutes a risk-management approach in which risk is assessed as a consequence of compromise. Typically, compromise occurs in the form of disruptive loss of functionality, a catastrophically costly data breach, or both, suggesting, at best, multiple probability-based assessment metrics. Unfortunately, such metrics are purely theoretical, as humans are usually engaged at all levels of cybersecurity. Far too often, the human element becomes the spoiler.

THE HUMAN FACTOR

Significant security breaches often result more from faulty policies or lax procedures than poor system design. Prevalent 1992–1993 security taxonomies focused on so-called deliberate and accidental events: intentional perpetrators of fraud, espionage, and vandalism committed deliberate events; anything else was an accident. By 2001, further research introduced the notions of human error through gaps or faults in the human skill, rule, or

 Published by the IEEE Computer Society 2469-7087/16/$33.00 © 2016 IEEE

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

knowledge base. These faults helped enrich the accidental element of the security lexicon.[1] Soon thereafter, best security practices began including policies, awareness, and education to address human error.[2]

In its *2014 Cyber Security Intelligence Index*, IBM noted that "over 95 percent of all [security] incidents investigated recognize 'human error' as a contributing factor. The most commonly recorded form of human errors include system misconfiguration, poor patch management, use of default user names and passwords or easy-to-guess passwords, lost laptops or mobile devices, and disclosure of regulated information via use of an incorrect email address. The most prevalent contributing human error? 'Double clicking' on an infected attachment or unsafe URL."[3]

Two recent well-publicized incidents illustrate humans' central role in security failures.

In late 2014, a hacker group called the Guardians of Peace leaked a whopping 100 Tbytes of sensitive data stolen from Sony Pictures Entertainment, including yet-unreleased movies, confidential emails, and personally identifiable information about employees such as salaries and Social Security numbers, to the Internet. Possibly sponsored by North Korea, the hackers had penetrated the company's network and been extracting the data for more than a year. The data breach was followed by a sophisticated and destructive wiper malware attack distributed via Sony's email system that erased server data and made many machines unusable. A SANS Institute report noted that Sony had failed to implement basic critical security controls—for example, most of the stolen data wasn't encrypted or even secured by passwords.[4] The lax security was all the more perplexing given that Sony had suffered multiple costly and embarrassing data breaches in 2011.

A far more damaging cyberattack, this time against a US government agency, occurred at about the same time. In late 2013, suspected Chinese hackers breached Office of Personnel Management (OPM) databases and in separate but related intrusions in 2014–2015 stole background investigation data and personnel records of tens of millions of current, former, and prospective federal employees and contractors. While the exact means of exploitation remains unstated, OPM IT security personnel had received repeated warnings about malicious activity and failed to implement adequate safeguards. The breach had major national security implications and led to resignations of senior OPM IT officials and a massive overhaul of agency security procedures.[5]

UPFRONT DESIGN

Even if most security failures are attributable to human error, the importance of upfront design can't be discounted. In fact, security by design is thriving in industry, as evidenced by ongoing controversies surrounding encryption technology.

In the wake of Edward Snowden's revelations about government surveillance, many tech companies have publicly committed to protecting user privacy by hardening or expanding their encryption capabilities. For example, social media messaging systems like WhatsApp and Snapchat employ end-to-end encryption schemes for all user traffic as well as options such as message self-destruction. These schemes are difficult to penetrate and are increasing frustration among federal and local law-enforcement agencies pursuing criminals and terrorists.

However, in the endless cat-and-mouse game between security designers and hackers, no system is totally immune from compromise given sufficient will and resources. Consider the recent standoff between Apple and the Federal Bureau of Investigation over iPhone encryption. In most iPhone models, 10 failed attempts to correctly enter a four-digit passcode thoroughly scrambles the phone's data. This posed a dilemma after the 2 December 2015 mass shooting in San Bernardino, California, when the FBI sought encrypted data on an iPhone 5C belonging to one of the terrorists. After Apple resisted pressure to create software that would unlock the device, the FBI hinted that it paid a third party about $1.3 million to successfully hack into the device.[6]

SECURITY VERSUS PRIVACY

Upfront design can certainly contribute to "good enough" security, but what if the security is too good? Should an encryption scheme prevent authorities from obtaining critical information that could thwart a future terrorist attack, or allow undetectable nefarious communication traffic to occur with relative impunity?

Protecting individuals' privacy and protecting society against clear and present dangers are both legitimate goals. In recent years, the debate between privacy and security advocates has intensified. Much of the legislation addressing this issue is outdated, and doesn't reflect current technological capabilities. For example, the Electronic Communications and Privacy Act, which regulates what digital information can be collected on citizens, was passed in 1986. In its dispute with Apple over encryption, the US government invoked the All Writs Act of 1789. Ultimately, Congress and the courts will have to draw the line between security and privacy.

In the meantime, media reports suggest that cyberattacks against corporations and government agencies are increasing in both frequency and severity. Relatively simple exploits, often involving spoofing, phishing,

and other malicious activities propagated through email, can seriously disrupt organizational operations, exfiltrate sensitive information, and, as evidenced by the recent rash of ransomware attacks on hospitals,[7] block access to computers and data.

As the world becomes ever more connected and the Internet of Things grows, the potential risks of large, devastating cyberattacks will raise the stakes in the privacy–security dispute. For example, attacks on supervisory control and data acquisition networks, such as the 2010 Stuxnet attack on Iranian nuclear-enrichment facilities[8] and the attack on the Ukrainian power grid in 2015,[9] threaten to cripple entire critical infrastructures.

When does security trump privacy, and what jurisdictional laws apply to international attacks in cyberspace, where there are no physical borders?

TOWARD A NEW PARADIGM

Given that no computing system is completely immune from attack, and that accommodations for privacy protection must be made and could vary considerably depending on the system, what constitutes "good enough" security?

In an earlier era, we might have argued that security is "good enough" when it discourages a would-be adversary from expending the necessary resources to mount a successful attack. This logic is consistent with the metaphor of a neighborhood in which homes with unlocked doors are far more vulnerable to crime than those with defenses—a burglar faced with deadbolts, alarms, or vicious dogs is more likely to move on to a less well-fortified home.

However, such a simplistic model is inadequate for dealing with today's realities, necessitating a paradigm shift. Many modern cyberattacks originate from multiple hosts, involve multiple sessions, and aim at multiple targets simultaneously. Such multipoint attacks defy traditional approaches to systems security such as standalone malware or intrusion detection. Rather, the presence of distributed attacks must be derived from behavioral patterns involving evidence from many hosts.[10] Likewise, patterns hold valuable behavioral information.

Consider credit scores. A credit score combines many behavioral characteristics—for example, too many credit inquiries, failure to regularly pay bills, or maintaining a significant balance on a revolving credit account regardless of payment record will lower the overall score. High scores suggest acceptable credit risk, while low scores indicate elevated risk.

Credit scores make it possible to assess individuals buying a car, applying for a credit card, buying a home, and so on across a common rating scale. Such a scale incentivizes low-rated individuals to raise their scores to improve their lot in life while establishing high-performance levels in which "reasonable" behavior statistically reduces relative risk.

Other industries have similar types of ratings—for example, automotive crashworthiness and safety evaluations by vehicle make and model.

Such rating systems are nongovernmental; they're used for reporting, not regulatory, purposes. In addition, they encourage participants to work toward a common good by establishing accepted objective standards.

A METRIC-BASED APPROACH

A similar metric-based approach would enable different organizations to gauge the relative security risks of all types of cyberactivities.

The makings of such a structure already exist in the systems security realm. Standards such as the Common Criteria for Information Technology Security Evaluation (ISO/IEC 15408), Systems Security Engineering—Capability Maturity Model (ISO/IEC 21827), and Security Requirements for Cryptographic Modules (FIPS-140), and guidelines such as the National Institute of Standards and Technology *Performance Measurement Guide for Information Security* (SPP 800-55) seek to establish useful metrics against a number of security-related factors including but not limited to people, processes, networks, hardware, and software.[11]

Likewise, the Open Web Application Security Project (OWASP; www.owasp.org), the SANS Institute (www.sans.org), the Software Engineering Institute's Computer Emergency Response Team (CERT; www.cert.org), and many other organizations maintain extensive databases of prior security incidents that could help form the basis of useful metrics.

Using such metrics, an industry-supported private-sector organization could evaluate and report on the relative risk of various products, corporations, government agencies, and nongovernmental organizations against a dynamic, objective, and common scale. The ratings would let stakeholders determine whether others' security is "good enough" to do business with them. Low-rated organizations would be motivated to improve their security posture, while those partnering with them could take extra precautions. By assessing security risk on an actuarial basis, the scale could also lead to the creation of insurance underwriters for systems security.

Similar stakeholder models already exist in segments of the test, reliability, and safety spaces. For example, the Radio Technical Commission for Aeronautics administers the Software Considerations in Airborne Systems and Equipment Certification (DO-178C), a private-sector standard that guides the Federal Aviation Administration and other aviation agencies around the world in certifying airworthy software. Likewise, the Food and Drug Administration relies on consensus standards[12] in certifying medical devices for commercial use under the 510(k) rapid certification program.[13] The International Electrotechnical Commission sponsors the Functional Safety of Electrical/Electronic/Programmable Electronic Safety-Related Systems (IEC 61508), which represents many other like standards in the nuclear safety realm.[14]

This article originally appeared in Computer, *vol. 49, no. 7, 2016.*

Although none of these standards specify what is "good enough" in their particular space, each establishes a working threshold that defines it by default. As a first step toward defining acceptable risk as a function of testing, reliability, and safety, the security community must follow other communities' lead in creating consensus-based standards.

So long as people are involved, systems security will never be better than "good enough." Even with robust upfront design, any system can be penetrated by adversaries with sufficient desire, time, and resources. Sometimes all it takes to compromise a system is duping a single user into opening a door; other times the door is left wide open by faulty policies or lax procedures. Government pressure to facilitate surveillance—for example, through built-in back doors—could inadvertently make systems more vulnerable to sophisticated hackers.

Assuming that a system is completely secure is foolhardy and arrogant—the potential damage from a breach can be costly at best and life-threatening at worst. Consequently, technological and design advances in cybersecurity, while important, should take a backseat to improving human awareness of security risks. Toward this end, an objective, consensus-based rating system offers a fair means of ranking the relative risks of different products and organizations and raises the bar to achieve ever-higher levels of protection against cyberattacks.

Unfortunately, this message won't self-destruct in five seconds.

REFERENCES

1. G.P. Im and R.L. Baskerville, "Longitudinal Study of Information System Threat Categories: The Enduring Problem of Human Error," *ACM SIGMIS Database*, vol. 36, no. 4, 2005, pp. 68–72.
2. M.E. Whitman, "In Defense of the Realm: Understanding Threats to Information Security," *Int'l J. Information Management*, vol. 24, no. 1, 2004, pp. 43–57.
3. IBM Global Technology Services, *IBM Security Services 2014 Cyber Security Intelligence Index: Analysis of Cyber Attack and Incident Data from IBM's Worldwide Security Operations*, research report, IBM Corp., June 2014; www.slideshare.net/ibmsecurity/2014-cyber-security-intelligence-index.
4. G. Sanchez, *Case Study: Critical Controls That Sony Should Have Implemented*, white paper, SANS Inst., 1 June 2015; www.sans.org/reading-room/whitepapers/casestudies/case-study-critical-controls-sony-implemented-36022.
5. K. Finklea et al., *Cyber Intrusion into U.S. Office of Personnel Management: In Brief*, CRS report R4411, Congressional Research Service, 17 July 2015; www.fas.org/sgp/crs/natsec/R44111.pdf.
6. E. Lichtblau and K. Benner, "F.B.I. Director Suggests Bill for iPhone Hacking Topped $1.3 Million," *The New York Times*, 21 Apr. 2016; www.nytimes.com/2016/04/22/us/politics/fbi-director-suggests-bill-for-iphone-hacking-was-1-3-million.html?_r=0.
7. B. Siwicki, "Cybersecurity Special Report: Ransomware Will Get Worse, Hackers Targeting Whales, Medical Devices and IoT Trigger New Vulnerabilities," *Healthcare IT News*, 21 May 2016; www.healthcareitnews.com/news/cybersecurity-special-report-ransomware-will-get-worse-hackers-targeting-whales-medical-devices.
8. K. Zetter, *Countdown to Zero Day: Stuxnet and the Launch of the World's First Digital Weapon*, Crown, 2014.
9. K. Zetter, "Inside the Cunning, Unprecedented Hack of Ukraine's Power Grid," *Wired*, 3 Mar. 2016; www.wired.com/2016/03/inside-cunning-unprecedented-hack-ukraines-power-grid.
10. R.K. Cunningham et al., "Detecting Computer Attackers: Recognizing Patterns of Malicious, Stealthy Behavior," presentation, Center for Education and Research in Information Assurance and Security, Purdue Univ., 29 Nov. 2000; www.cerias.purdue.edu/news_and_events/events/security_seminar/presentations/11-29-2000.pdf.
11. Z. Abbadi, "Security Metrics? What Can We Measure?," presentation, Open Web Application Security Project (Northern VA chapter), 19 Apr. 2007; www.owasp.org/images/b/b2/Security_Metics-_What_can_we_measure-_Zed_Abbadi.pdf.
12. *Guidance for Industry and FDA Staff: Recognition and Use of Consensus Standards*, Center for Devices and Radiological Health Standards, Food and Drug Administration, US Dept. of Health and Human Services, 17 Sept. 2007; www.fda.gov/downloads/MedicalDevices/DeviceRegulationandGuidance/GuidanceDocuments/ucm077295.pdf.
13. *The New 510(k) Paradigm: Alternate Approaches to Demonstrating Substantial Equivalence in Premarket Notifications*, Center for Devices and Radiological Health Standards, Food and Drug Administration, US Dept. of Health and Human Services, 20 Mar. 1998; www.fda.gov/MedicalDevices/DeviceRegulationandGuidance/GuidanceDocuments/ucm080187.htm.
14. P. Baufreton et al., "Multi-Domain Comparison of Safety Standards," *Proc. 5th Int'l Conf. Embedded Real Time Software and Systems* (ERTS 10), 2010; http://web1.see.asso.fr/erts2010/Site/0ANDGY78/Fichier/PAPIERS%20ERTS%202010/Corrected-Blanquart.PDF.

GEORGE HURLBURT is the chief scientist at STEMCorp, a nonprofit corporation that works in the public sector to further economic development via adoption of network science to advance autonomous technologies as useful tools for human use. Contact him at ghurlburt@change-index.com.

Editors: Sean W. Smith, sws@cs.dartmouth.edu | William Enck, whenck@ncsu.edu

Machine Learning in Adversarial Settings

Patrick McDaniel, Nicolas Papernot, and Z. Berkay Celik | Pennsylvania State University

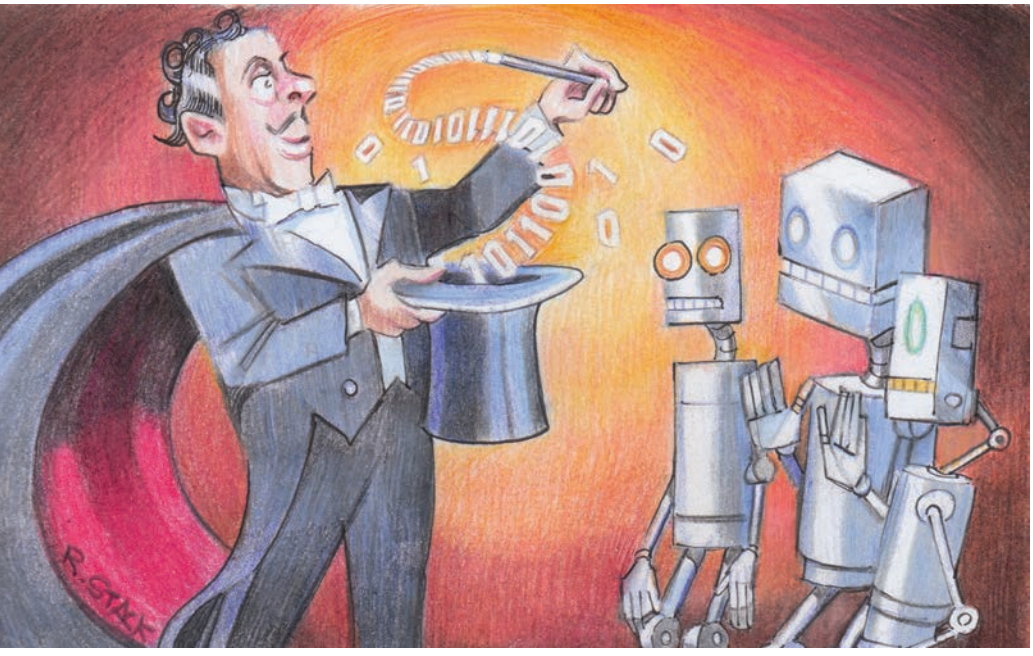

Advances in machine learning have led to transformational new fields of technology and introduced capabilities not previously possible. Emerging applications in self-driving cars, data analytics on massive datasets, adaptive and interactive entertainment, and Web search and sentiment analysis are but a few of the technologies that will impact society in the decades to come.

Perhaps no technology field has relied on or benefited from advances in machine learning more than systems and computer security. Machine learning is the basis for almost all nonsignature-based detection, whether identifying malware, network intrusions, spam, rogue processes, fraudulent transactions, or other malicious activity. Indeed, machine learning has become so intertwined with security that the technical community's ability to apply machine learning securely will likely be crucial to future environments.

In this article, we consider whether today's use of machine learning in security-sensitive applications is vulnerable to nonobvious and potentially dangerous manipulation. Here, we examine sensitivity not only in the context of computer security but also in any application whose misuse might lead to harm—for instance, crashing an autonomous vehicle or bypassing a content filter. We explore the use of machine learning in this area particularly in light of recent advances in the computationally efficient creation of adversarial samples targeted at widely used classes of machine-learning approaches.

Machine Learning in Practice

Consider a generalized use of machine learning as a classifier and an example system identifying spam email. (For brevity, we restrict ourselves to machine-learning classifiers—identifying a sample as being from some output class, among a predefined finite set of [potentially] many classes—trained on labeled data. Many other kinds of machine-learning systems and training techniques exist; our arguments apply almost universally.) A classifier is a system that takes an input sample and identifies it as one of several output classes (or none, if the sample can't be identified confidently). In this example, the system determines whether the item is in the "spam" or "not spam" class.

In machine learning, each sample is input into the classification process as a vector of features that describe the sample. For email, typical features might be keywords, sender and recipient domain names, existence of embedded content, or number of emails of a particular type. The

 Published by the IEEE Computer Society 2469-7087/16/$33.00 © 2016 IEEE

system determination is based on how that set of input features is interpreted by the model for the classification process—in this case, a model of how email input features indicate spam or not.

Conceptually, a model encodes semantic information about how certain features or sets of features relate to the output class. For example, certain keywords or keyword combinations could be strong indicators of an email being spam. In practice, models will encode many different such relationships, each weighted on the basis of the association's strength. An aggregate calculation over the feature associations with respect to the input features results in an output classification and/or confidence score.

To date, the key assessment metric for these systems has been accuracy: How often does the model pick the correct class for a sample? Several accuracy measures exist, including precision, sensitivity, and specificity. These quality assessments directly relate to assumptions about the expected distribution of the classification system input and don't account for adversarial behavior, which often falls outside of this expected input distribution. In other words, accuracy can be viewed as a measure of the system's average performance, whereas the security evaluation is interested in worst-case performance.

Adversarial Samples

One of the limitations of machine learning in practice is that it's subject to adversarial samples. Adversarial samples are carefully modified inputs crafted to dictate a selected output. In the context of classification, adversarial samples are crafted to force a target model to classify them in a class different from their legitimate class—for instance, spam emails that bypass a spam filter. The modifications, called *perturbations*, are introduced to yield a specific adversary-selected misclassification. In general, adversaries want to perturb the sample as little as possible so that to a human observer, for example, it remains indistinguishable from the original unaltered sample.

In the context of classification, adversarial samples are crafted to force a target model to classify them in a class different from their legitimate class.

Over the past few years, several algorithms used to automate adversarial-sample generation have emerged for multiclass classifiers built, for example, with deep neural networks. In late 2013, Christian Szegedy and his colleagues were the first to reveal the vulnerability of trained deep neural networks to slight perturbations of their inputs when they cast sample generation as an approximate optimization.[1] Ian Goodfellow and his colleagues followed with a fast gradient sign method, which linearly approximates the cost function in the neighborhood of legitimate samples to allow faster crafting of adversarial samples.[2] Finally, Nicolas Papernot and his colleagues proposed an iterative crafting algorithm that uses the model's Jacobian to select perturbations yielding the adversary's desired classification. They showed that adversaries can reliably achieve chosen adversarial target classes for any legitimate source class.[3] Their iterative approach also allows greater control over the introduced perturbations, thus reducing their magnitude. These more recent works expand on the classical adversarial machine-learning efforts described by Pavel Laskov and Richard Lippmann.[4] For example, related past work explored the formalization of worst-case errors against learned binary classifiers,[5] reverse engineering of binary linear classifiers to identify inputs they misclassify,[6] and contamination of training data jeopardizing binary classifiers' integrity and availability.[7]

Consider the following real-world scenario in which an autonomous vehicle uses a camera to identify and recognize roadside signs (see Figure 1). Once a sign has been identified, its image is fed to a neural network for classification in one of the predefined sign classes. Here, the neural network identifies the sign as a stop sign. Now, consider adversaries capable of altering the input of this neural network. They can force the model to output a wrong class upon processing a slightly perturbed variant of the stop sign's image. If adversaries can transfer perturbations to the neural network's image input, the autonomous system can be misled into misclassifying signs—reading stop signs as yield signs, for instance—potentially resulting in vehicles crashing into one another.

Again, to humans, adversarial samples are often indistinguishable from original samples. Humans would classify both images in Figure 2 as stop signs. In real-world tests using the Papernot algorithm, a trained deep-learning neural network classifies Figure 2a as a stop sign and Figure 2b as a yield sign. In actuality, the image on the left is an ordinary image of a stop sign, whereas the image on the right is an adversarial sample crafted by solving the earlier optimization problem.

Learning Models from Training Data

To understand why adversarial samples exist, it's important to explore how learning models are

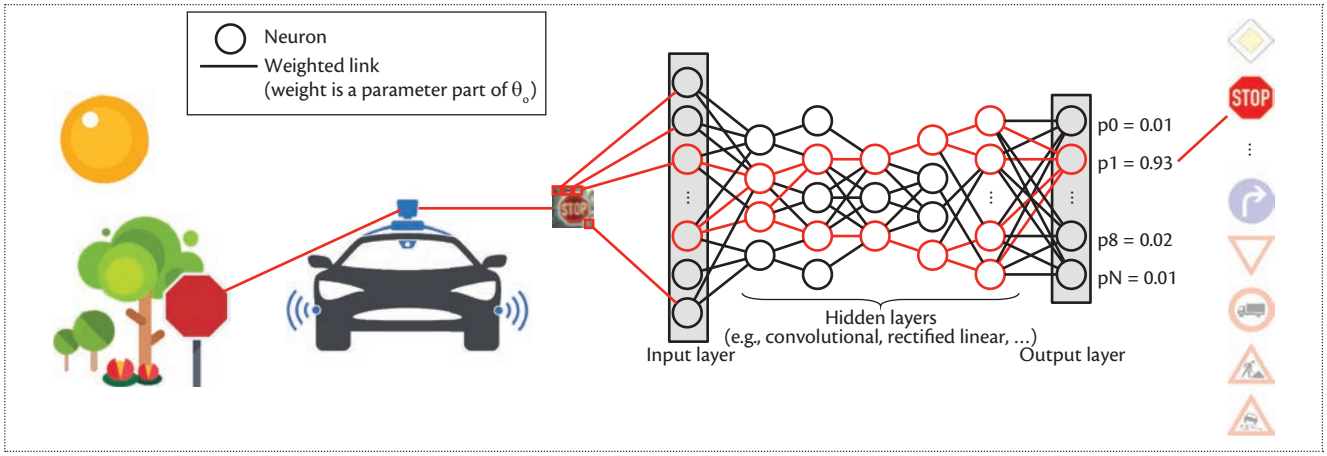

Figure 1. An autonomous vehicle uses a camera to identify and recognize roadside signs. Once a sign has been identified, its image is fed to a neural network for classification in one of the predefined sign classes. Here, the neural network identifies the sign as a stop sign.

built. Although there are other approaches, the models we discuss here are trained in a supervised fashion using labeled training data. This training data is a corpus of samples taken from the expected input distribution and labeled with their class. In the case of our spam system, this sample data would be a large number of emails that indicate whether or not they are spam. In the sign recognition system, the training data would include numerous signs and their type: stop, yield, and so on. These labels are taken as ground truth in constructing the models to be used at runtime.

Generally, model training begins with a null model representing no information. The training method iteratively processes each input sample in the training data and updates the model. This iterative refinement process strengthens or weakens the classification associations as supporting evidence is identified. Generally, the larger and more diverse the training data is, the more accurate the system becomes.

The refinement process of the input data's internal representations is specific to the kind of machine-learning technique employed: shallow and deep neural networks represent the model as a complex feed-forward network of mathematical neurons (parameterized elementary computing units), support vector machines use high-dimensional hyperplanes to separate classes, and random forests represent the model as a collection of learned decision trees. Some machine-learning techniques don't store a model—for instance, nearest neighbors—but simply use lazy evaluation to compare unseen samples to the training samples.

Regardless of technique, the model represents an approximation of the phenomena being modeled; unless the training data contains all possible input feature vectors, it can't fully capture a complete model of the target domain. In nonadversarial environments, this often isn't a problem. Data representative of the expected input distribution is sufficient for training. With enough input emails or images of signs to train on, input normally encountered at runtime will be sufficiently similar to allow the model to output a correct classification prediction by extrapolating from training samples.

Exploiting Natural Complexity in Decision Boundaries

A problem arises when adversaries exploit the system by providing input samples that aren't within the expected input domain. Here, they use information about the system to find where the model is inaccurate owing to items missing from the training set.

Consider an unsophisticated sample-generation algorithm in which adversaries simply test different input samples until they find a combination of input features that reliably achieves the desired classification. For example, spammers could simply modify email typography, vocabulary, addresses, and domains; test against the system; and see which are marked as legitimate (not spam). Indeed, this is common practice today; each new spam campaign contains carefully tested and selected email features that reliably bypass online spam filtering systems.

Figure 3 illustrates model training and use. In this figure, the plane represents all possible input feature vectors. For each sample, the input feature values uniquely identify its coordinates in the plane. Two classes A and B (that is, spam and not spam) are regions in a two-dimensional plane separated by the smooth curved line. All samples above the smooth curved line are in class A, and those below are in class B. This line is called the *real*

decision boundary. The model is trained using the input samples labeled X. On the basis of these samples, the training algorithm approximates the class separation as the linear dashed line—the *model decision boundary*. The distance between the real and model decision boundary is called the *model error* or space of adversarial samples (adversarial regions).

One might intuit that the model the algorithm learned by was faulty, but this isn't true.

This is a legitimate and highly accurate model for the training data: every sample in the input sample distribution is correctly classified. Indeed, one can't do better than this without more samples or information; a natural error is introduced by the fact that the training data can't, in almost all circumstances, cover the entire feature space or provide enough data to illuminate the real decision boundary with anything other than approximate accuracy. Making matters worse, the real decision boundary generally becomes more complex as the phenomenon becomes more nuanced and the feature and dimension space becomes larger.

It's this complexity that adversaries exploit. They simply take a sample and use trial and error (as in our earlier spam example) or information about the model error (as in the recently developed adversarial sample algorithms) to find a few perturbations that "move" the sample into the region of adversarial samples.

Herein lies the crux of these systems' vulnerability. Because adversaries can control the input sample features, they explicitly drive the malicious sample into the regions of the input space that are ambiguous with respect to the model. In short, they search for or calculate a sample that's in one class (for example, spam or stop sign) but, owing to the ambiguity resulting from incomplete training data, is classified as being in another class (for example, not spam or yield sign).

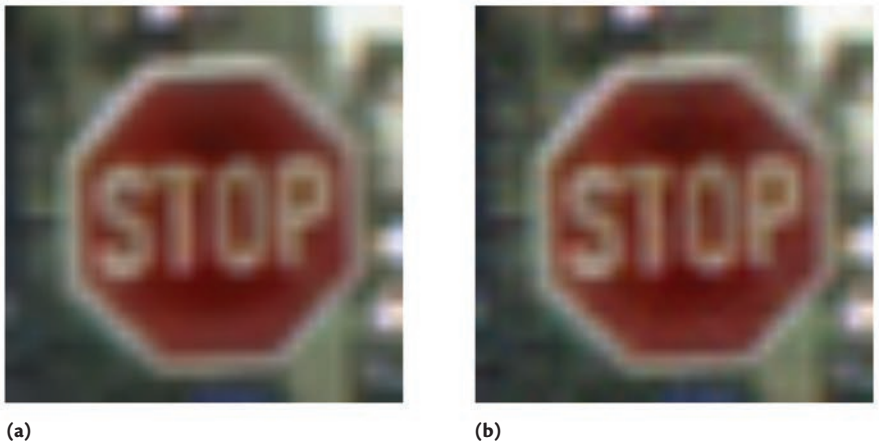

Figure 2. To humans, adversarial samples are indistinguishable from original samples. (a) An ordinary image of a stop sign. (b) An image crafted by an adversary.

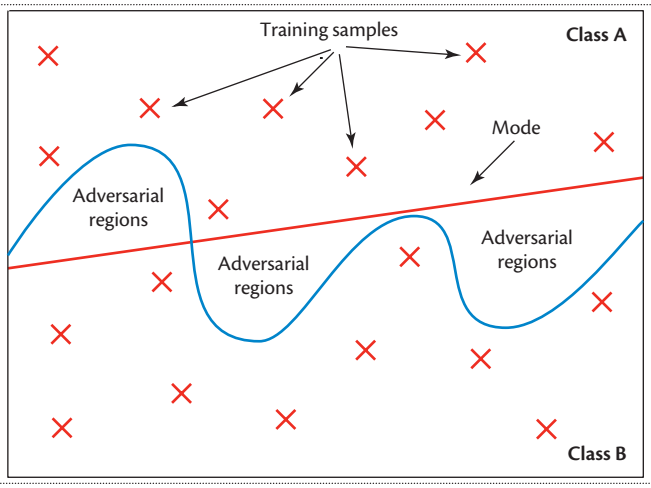

Figure 3. Model training and use. The plane represents all possible input feature vectors. For each sample, the input feature values uniquely identify its coordinates in the plane. Two classes A and B (that is, spam and not spam) are regions in a two-dimensional plane separated by the smooth curved line.

The Importance of Model Resilience

We argue that to address adversarial action, a new metric for machine-learning model quality is needed: *model resilience*.[8] Model resilience can be defined as robustness to perturbations of its input. Simply put, the more perturbation needed to move a sample from its legitimate class to an adversarial class, the more robust the model is to adversarial manipulations of its inputs.

In practice, we can achieve resilience in several ways. In the simplest approach, we can simply require higher confidence in outputs. This would move the decision boundaries further apart and thus leave fewer regions of ambiguity. This of course would affect model accuracy. Other approaches would be to refine the training process to smooth decision boundaries, or to measure each input's likelihood of

being an adversarial sample based on its characteristics, for example, closeness to the centroid of a nonselected class. Such approaches aren't well understood, but they're certain to be a necessary element to securing the future of machine learning in adversarial settings.

Machine learning is driving rapid innovation and providing new insights into how we can interpret and control complex data and environments. With these advances, adversaries will seek to circumvent their controls and drive systems for their malicious ends. In recognition of this reality, the machine-learning and security communities must endeavor to inoculate systems against such misuse. Thus, we must revisit our measures of quality for machine-learning techniques and weigh not only the results they produce but also their ability to resist samples carefully generated by adversaries. ■

References

1. C. Szegedy et al., "Intriguing Properties of Neural Networks," *Proc. Int'l Conf. Learning Representations* (ICLR 14), 2014; http://arxiv.org/abs/1312.6199.
2. I.J. Goodfellow, J. Shlens, and C. Szegedy, "Explaining and Harnessing Adversarial Examples," *Proc. Int'l Conf. Learning Representations* (ICLR 15), 2015; http://arxiv.org/abs/1412.6572.
3. N. Papernot et al., "The Limitations of Deep Learning in Adversarial Settings," *Proc. 1st IEEE European Symp. Security and Privacy* (EuroS&P 16), 2016; http://arxiv.org/abs/1511.07528.
4. P. Laskov and R. Lippmann, "Machine Learning in Adversarial Environments," *Machine Learning*, vol. 81, no. 2, 2010, pp. 115–119.
5. M. Kearns and M. Li, "Learning in the Presence of Malicious Errors," *J. Computing*, vol. 22, no. 4, 1993, pp. 807–837.
6. D. Lowd and C. Meek, "Adversarial Learning," *Proc. Knowledge Discovery in Data Mining* (SIGKDD 05), 2005, pp. 641–647.
7. B.A. Nelson, "Behavior of Machine Learning Algorithms in Adversarial Environments," PhD thesis, Dept. of Computer Science, Univ. California, Berkeley, 2010.
8. N. Papernot et al., "Distillation as a Defense to Adversarial Perturbations against Deep Neural Networks," to be published in *Proc. 37th IEEE Symp. Security and Privacy*, 2016.

Patrick McDaniel is a Distinguished Professor in the School of Electrical Engineering and Computer Science at the Pennsylvania State University. Contact him at mcdaniel@cse.psu.edu.

Nicolas Papernot is a graduate student at the Pennsylvania State University. Contact him at ngp5056@cse.psu.edu.

Z. Berkay Celik is a graduate student at the Pennsylvania State University. Contact him at zbc102@cse.psu.edu.

This article originally appeared in IEEE Security & Privacy, *vol. 14, no. 3, 2016.*

IEEE computer society | CELEBRATING 70 YEARS

PURPOSE: The IEEE Computer Society is the world's largest association of computing professionals and is the leading provider of technical information in the field.
MEMBERSHIP: Members receive the monthly magazine *Computer*, discounts, and opportunities to serve (all activities are led by volunteer members). Membership is open to all IEEE members, affiliate society members, and others interested in the computer field.
OMBUDSMAN: Email ombudsman@computer.org.
COMPUTER SOCIETY WEBSITE: www.computer.org

Next Board Meeting: 13–14 November 2016, New Brunswick, NJ, USA

EXECUTIVE COMMITTEE
President: Roger U. Fujii
President-Elect: Jean-Luc Gaudiot; **Past President:** Thomas M. Conte;
Secretary: Gregory T. Byrd; **Treasurer:** Forrest Shull; **VP, Professional and Educational Activities:** Andy T. Chen; **VP, Member & Geographic Activities:** Nita K. Patel;
VP, Publications: David S. Ebert; **VP, Standards Activities:** Mark Paulk;
VP, Technical & Conference Activities: Hausi A. Müller; **2016 IEEE Director & Delegate Division VIII:** John W. Walz; **2016 IEEE Director & Delegate Division V:** Harold Javid;
2017 IEEE Director-Elect & Delegate Division VIII: Dejan S. Milojičić

BOARD OF GOVERNORS
Term Expriring 2016: David A. Bader, Pierre Bourque, Dennis J. Frailey, Jill I. Gostin, Atsuhiro Goto, Rob Reilly, Christina M. Schober
Term Expiring 2017: David Lomet, Ming C. Lin, Gregory T. Byrd, Alfredo Benso, Forrest Shull, Fabrizio Lombardi, Hausi A. Müller
Term Expiring 2018: Ann DeMarle, Fred Douglis, Vladimir Getov, Bruce M. McMillin, Cecilia Metra, Kunio Uchiyama, Stefano Zanero

EXECUTIVE STAFF
Executive Director: Angela R. Burgess; **Director, Governance & Associate Executive Director:** Anne Marie Kelly; **Director, Finance & Accounting:** Sunny Hwang; **Director, Information Technology & Services:** Sumit Kacker; **Director, Membership Development:** Eric Berkowitz; **Director, Products & Services:** Evan M. Butterfield; **Director, Sales & Marketing:** Chris Jensen

COMPUTER SOCIETY OFFICES
Washington, D.C.: 2001 L St., Ste. 700, Washington, D.C. 20036-4928
Phone: +1 202 371 0101 • **Fax:** +1 202 728 9614 • Email: hq.ofc@computer.org
Los Alamitos: 10662 Los Vaqueros Circle, Los Alamitos, CA 90720
Phone: +1 714 821 8380 • **Email:** help@computer.org

MEMBERSHIP & PUBLICATION ORDERS
Phone: +1 800 272 6657 • **Fax:** +1 714 821 4641 • **Email:** help@computer.org
Asia/Pacific: Watanabe Building, 1-4-2 Minami-Aoyama, Minato-ku, Tokyo 107-0062, Japan • Phone: +81 3 3408 3118 • **Fax:** +81 3 3408 3553 • Email: tokyo.ofc@computer.org

IEEE BOARD OF DIRECTORS
President & CEO: Barry L. Shoop; **President-Elect:** Karen Bartleson; **Past President:** Howard E. Michel; **Secretary:** Parviz Famouri; **Treasurer:** Jerry L. Hudgins; **Director & President, IEEE-USA:** Peter Alan Eckstein; **Director & President, Standards Association:** Bruce P. Kraemer; **Director & VP, Educational Activities:** S.K. Ramesh; **Director & VP, Membership and Geographic Activities:** Wai-Choong (Lawrence) Wong; **Director & VP, Publication Services and Products:** Sheila Hemami; **Director & VP, Technical Activities:** Jose M.F. Moura; **Director & Delegate Division V:** Harold Javid; **Director & Delegate Division VIII:** John W. Walz

revised 10 June 2016

IEEE

Editor: Munindar P. Singh • m.singh@ieee.org

Natural Interaction for Bot Detection

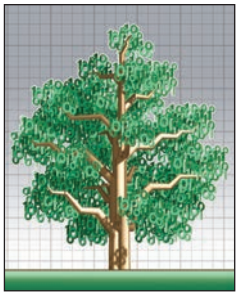

Robert St. Amant and David L. Roberts • *North Carolina State University*

Bot detection – identifying a software program that's using a computer system – is an increasingly necessary security task. Existing solutions balance proof of human identity with unobtrusiveness in users' workflows. Cognitive modeling and natural interaction might provide stronger security and less intrusiveness.

Bot detection has become an important topic in security. Bots are software programs that use a computer, typically a personal computer, for malicious use or at least use unintended by its owner. So-called "aiming bots" were once quite popular in online multiplayer first-person shooter games. These bots allowed players to bypass the game mechanics for targeting opponents, giving them perfect aim every time, and enabling them to artificially improve their standing in the game. More commonly these days, bots are often employed to register for free email accounts and send spam or phishing messages. These, among many others, are scenarios that bot detection techniques are designed to avoid.

Two families of techniques for bot detection are in common use today. One is represented by CAPTCHA (Completely Automated Public Turing test to tell Computers and Humans Apart)[1] technology. The premise behind CAPTCHA is to require a human to interactively solve a problem that's difficult (or more desirably, impossible) for a computer to solve. The now-ubiquitous CAPTCHA technology on the Internet involves having users look at distorted images of words or listen to distorted audio of words and type in the letters. Another approach, common in massively multiuser online games (MMOGs), involves monitoring a user's input to identify characteristic differences between human and bot-like behavior. Steven Gianvecchio and his colleagues, for example, show that differences in the distributions of keystroke durations and the efficiency of mouse movement can be used to distinguish humans from bots.[2] In online poker, systems can use this information to identify poker bots, along with other heuristic clues such as playing too many games continuously for too long a period of time. Yang-Wai Chow and his colleagues[3] propose that CAPTCHAs can be integrated into a MMOG as a mini-game: for example, making progress in a fantasy adventure game might require players to decode spells, presented visually as CAPTCHAs. The result is more natural interaction, integrated into the context of the game, with players potentially enjoying the challenge and getting better with practice.

Both of these approaches require users to "prove" that they're human; one requires explicit action on the user's part, while the other is passive. In other words, one is a human interactive proof (HIP), the other a human observational proof (HOP).[2] CAPTCHAs are a common examples of HIPs. An example of an observational proof is examining the spatial signature of mouse click locations, as influenced by an interface layout.[4] In this article, we discuss HIPs and HOPs, along with other recent developments (namely, human subtlety proofs, or HSPs), and their potential for natural interaction. Many of our examples will come from online games, in part because of the enormous growth in popularity over the past decade or so, and in part because games support a variety of interactions that can be considered natural in a specific context.

Natural Interaction

What would natural techniques for bot detection look like? As a term of art in human-computer interaction, natural user interfaces have a few key properties. Following the account of Daniel Wigdor

2469-7087/16/$33.00 © 2016 IEEE Published by the IEEE Computer Society November 2016

and Dennis Wixon,[5] users enjoy interacting with a natural user interface; users become more skilled with practice; and their interaction is appropriate to context. Wigdor and Wixon summarize by saying that a natural user interface is one that makes a user act and feel like a natural.

We take naturalness to be an important part of interactive security techniques. Security measures tend to be viewed by the average computer user as overhead — perhaps necessary, but still secondary in importance to carrying out other tasks.[6] In the early days of computer security, Jerome Saltzer and Michael Schroeder[7] identified psychological acceptability as an essential aspect of the human interface, making the correct use of protection mechanisms routine, but this goal has yet to be reached. To illustrate what we mean by "natural," let's discuss user authentication to establish some of the context for bot detection.

The use of usernames and passwords is a counterexample to natural interaction for user authentication. It's long been understood that remembering a username and password depends on access to semantic information in long-term memory. Although typing can become automatic and natural, the retrieval of specific information from memory, with precise syntactic properties, has less of the same naturalness. To overcome this artificiality, users tend to rely on tricks that lead to potential holes in security, such as choosing passwords based on semantically meaningful information (birthdates, family relationships, and so forth), not all of which can be kept private.

Other user authentication techniques, also in the challenge-response family, offer more naturalness. Biometric authentication, in which a user is identified by some physiological characteristic unique to that person, includes fingerprint analysis, face recognition, retina or iris scanning, speech and vocal sampling, gesture, hand geometry matching, and other techniques. A "natural" evolution of a technique such as fingerprint authentication is its integration into contexts in which users automatically place their fingertips in contact with the sensing device: a computer mouse, for example, or even the keys on a keyboard. Context plays an important role in the naturalness of the interaction. For example, Alexander Chan and his colleagues[8] describe the use of a Leap Motion to capture hand geometry and gestures, with greater than 99 percent accuracy in authentication; such a system would be useful, providing even continuous authentication, in a context in which gestures are an integral part of interaction with the system.

The more natural interface techniques share one thing in common: they're characteristically observational in nature.

A different avenue toward natural authentication is to rely not on fixed or slowly changing physiological characteristics of the user but on behaviors, which might be learned or practiced. These approaches tend to have a smaller footprint than those described above. The most common approach is to monitor a user's mouse movement or keyboard actions, matching observed patterns against the user's "signature." Techniques typically involve automatically extracting or generating a large number of features from the time series of mouse movements or key presses, inducing a model from some subset of the features, and using that model to classify new observed data.[4] Accuracy can be greater than 90 percent, with large enough samples of data, although acquiring enough data poses a tradeoff with respect to time.

The more natural interface techniques share one thing in common: they're characteristically observational in nature. These approaches work in the background, collecting information about users to compare against a model, and require no explicit, intentional interaction from users — and therefore have little or no cost to users.

HIPs, HOPs, and HSPs

HIPs and HOPs both have significant limitations, however. HOPs are susceptible to imitation attacks, in which bots carry out scripted actions designed to look like human behavior. HIPs, on the other hand, tend to be more secure because they require explicit action from a user to complete a dynamically generated test. Because these tests are dynamically generated, solutions to them can't (reasonably) be predicted, scripted, or generated by computer systems; however, because humans have to expend cognitive effort to pass HIPs, they can be disruptive or reduce productivity, violate the good design principles of natural interfaces, and even result in users seeking alternative systems to use.

With this in mind, some of the work happening at North Carolina State University is aimed at developing knowledge and techniques to enable human subtlety proofs. HSPs blend the stronger security characteristics of HIPs with the unobtrusiveness of HOPs. We're examining how subtle cognitive biases affect interaction with software in predictable and repeatable ways. Our goal is to leverage those biases to make small changes to interfaces that will subtly — not substantively — affect the interaction of either bots or humans. By making changes to interfaces strategically and looking for evidence of the subtle changes that

only humans would exhibit (and bots would find very difficult to fake), we expect HSPs to combine the strengths of both HIPs and HOPs.

One of our projects is based on touch interfaces on tablet computers. A sample interface presents a scattering of circular targets on the display; the user is to tap each target, upon which it disappears, to complete the task. Touch interfaces are associated with a much higher error rate in target selection than GUI interfaces used with a mouse or touchpad, and in practice a touch sometimes fails to register. Experimentally, with a tablet instrumented to collect touch and gaze data, we've identified different ways in which users respond to such errors.

One behavior involves a gaze fixation on a target, a tap, and then a pause to verify that the tap has been recognized. If the tap is successful the next target is handled, but if the tap fails then the target is tapped again. Another behavior is to tap targets without waiting for verification, returning to those that were missed. In either case, visual attention might remain on the target under consideration until a successful tap or move on to the next target. A different behavior, apparently derived from gaming experience, relies on peripheral vision to locate targets, with no obvious relationship between gaze fixations and tap locations. Yet another behavior involves a brief planning phase in which gaze moves between different targets before any one is tapped. These behaviors can be decomposed into what are referred to in the cognitive modeling literature as strategies (or micro-strategies) dimensions. Behaviors can be associated with different strategies by analysis of gaze fixations, tap locations, and the duration and ordering of events.

Critically, we can influence the target-selection error rate, either directly (by making targets larger or smaller) or indirectly (by simply ignoring taps, with some probability). The implication for HSPs is that if users are sensitive to the difference in error rates (we have evidence that this is the case), then we might be able to manipulate the interface to see if the user reacts in a way that we expect — for example, adopting a slower, more "careful" strategy for a higher error rate. Multiple target selection is a common enough task in touch-based interfaces that it could potentially act as the background for an HSP; this is part of our ongoing research.

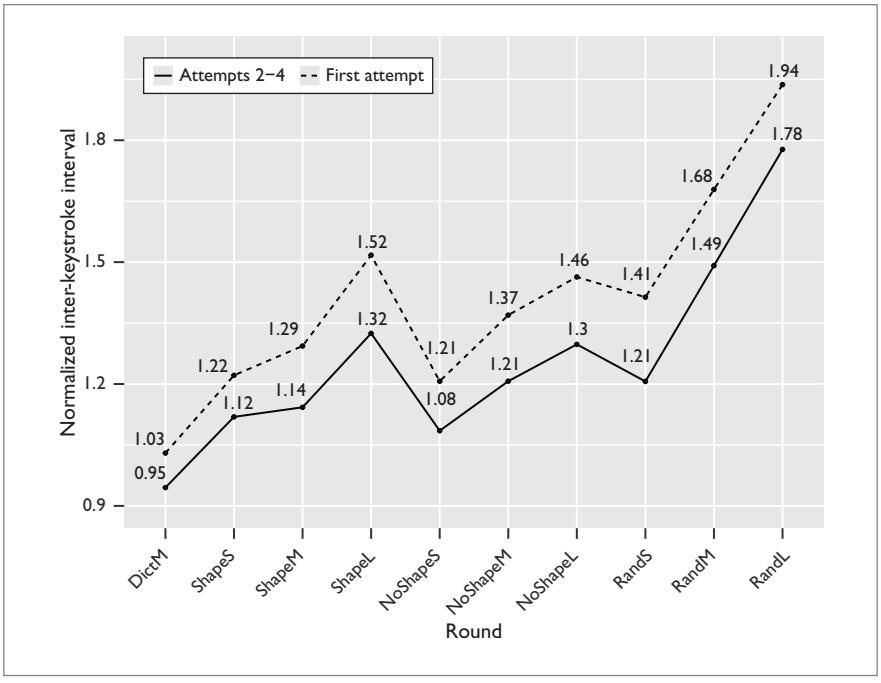

Figure 1. Mean inter-keystroke interval (IKI) for different word types, grouped by attempt. A dashed line indicates the first attempt, and a solid line represents the mean of subsequent attempts.

In another project, we used a typing game as an experimental platform. In the game, players typed words with differing characteristics and of varying length: dictionary words, dictionary words with transposed letters, and words composed of random letters. The game rewarded players for typing as quickly and accurately as possible; one factor that varied in the game was whether players could retry a given round consisting of a certain type of words, without a penalty, to improve their performance.

We found that typing speed improves with familiarity with words and with practice, but that these are independent of the number of mistakes that are made when typing. Specifically, the inter-keystroke interval (IKI), which is a measure of the time between key presses on a keyboard, was higher for misspelled words or random letter combinations. In all cases, the IKI decreased as players got practice typing those words or letter combinations. Figure 1 illustrates how practice reduces the IKI consistently for all word types, independent of players' familiarity with what they're typing. Figure 2 depicts the number of errors, which doesn't significantly change across different word types.

Our interpretation of these results was the players were sensitive to the speed/accuracy tradeoff, depending on the cost of the typing errors they made, which has implications for security. The number of incorrect attempts is a commonly used security measure, but our data suggest that speed has a more direct relationship to the nature of what's

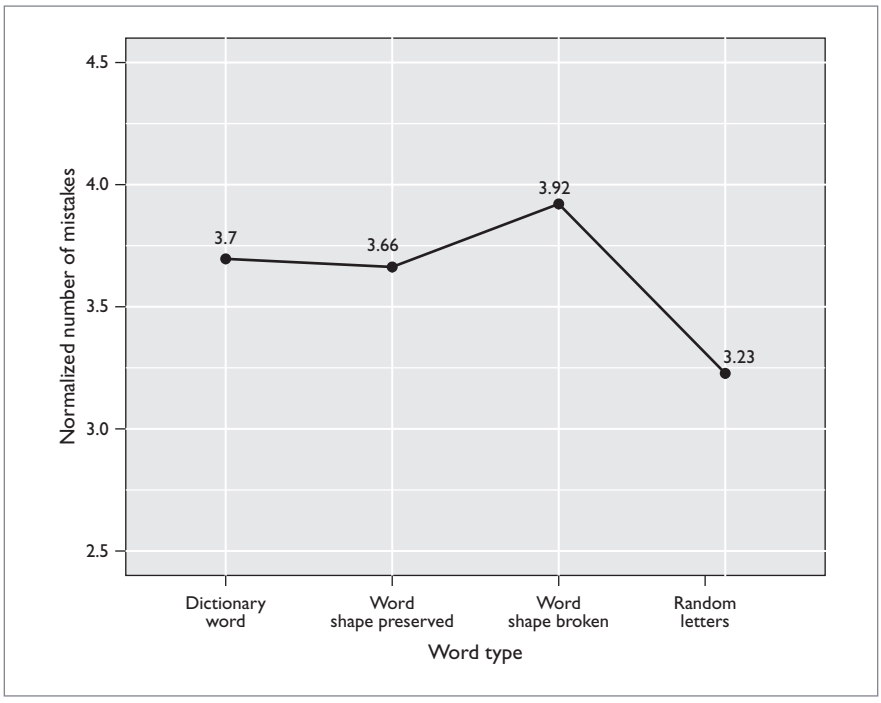

Figure 2. The mean number of mistakes made by players for different word types.

being typed than the number of mistakes that are made while typing. The implication of this finding is that by inspecting typing speed in conjunction with errors, a system can be more effective at detecting anomalies than looking at the number of incorrect attempts alone. Similarly, our results indicate that typing speed is correlated with the familiarity of the text being typed, which can be used, for example, to help systems ensure users create more secure passwords. It's a common IT policy to ask users to create new passwords, but it's impossible to know if a user's new password is entirely new or has been used on another system somewhere. By comparing to a known baseline, such as the user typing their previous passwords, the subtleties of the manifestation of cognitive function on typing can be brought to bear to identify familiarity. While challenging the user to type a specific word is the most obvious and direct application of this work (an HIP task), the underlying reasons for user performance might be applicable in an HSP.

HSPs are a new technology under early development, but show promise in producing new avenues for creating more secure systems. Human perception, cognition, and motor systems govern the ways in which users interact with input devices, and those systems are heavily influenced by the task being performed and the interface used for the task.

For example, many users have the taskbar configured to auto-hide when the mouse isn't in close proximity. When switching tasks while the taskbar is hidden, users are moving the mouse pointer to a target in their memory, rather than one identified by their visual perception. Under these different conditions, the motion of the mouse will differ in response to the brain's different perceptual and cognitive functions. Identifying these differences bears a resemblance to standard HOPs; however, by showing the taskbar at different times, a system can influence those perceptual and cognitive functions in predicable ways to produce expected changes in mouse movements. Making this change to the task environment gives this approach more in common with standard HIPs. Combined, however, the unobtrusive observational nature of monitoring mouse movement with the subtle, but still on-task, change to the taskbar's visibility will engage the subtlety of human cognition to produce an HSP.

Further down the line, similar approaches could yield techniques for detecting other classes of users, as well. Although bots and humans often have highly different interaction signatures, what about a user who's very distracted and likely to make a mistake compared to a focused user? How about an expert who has familiarity with the task and interface in comparison to a novice? These distinctions might represent new avenues for securing systems from authorized – but unintended – uses.

Acknowledgment

This work was funded by the US National Security Agency and US National Science Foundation through grant IIS-1451172.

References

1. L. von Ahn et al., "CAPTCHA: Using Hard AI Problems for Security," *Proc. Eurocrypt 2003: Int'l Conf. Theory and Applications of Cryptographic Techniques*, 2003, pp. 249–311.
2. S. Gianvecchio et al., "Battle of Botcraft: Fighting Bots in Online Games with Human Observational Proofs," *Proc. ACM Conf. Computer and Comm. Security*, 2009, pp. 256–268.
3. Y.W. Chow, W. Susilo, and H.Y. Zhou, "CAPTCHA Challenges for Massively Multiplayer Online Games: Mini-Game CAPTCHAs," *Proc. Int'l Conf. Cyberworlds*, 2010, pp. 254–261.
4. Z. Jorgensen and T. Yu, "On Mouse Dynamics as a Behavioral Biometric for Authentication," *Proc. 6th ACM Symp. Information, Computer and Comm. Security*, 2011, pp. 476–482.
5. D. Wigdor and D. Wixon, *Brave NUI World: Designing Natural User Interfaces for Touch and Gesture*, Elsevier, 2011.
6. D. Weirich and M.A. Sasse, "Pretty Good Persuasion: A First Step Towards Effective

Password Security in the Real World," *Proc. Workshop on New Security Paradigms*, 2001, pp. 137–143.
7. J. Saltzer and M. Schroeder, "The Protection of Information in Computer Systems," *Proc. IEEE*, 1975, pp. 1278–1308.
8. A. Chan, T. Halevi, and N. Memon, "Leap Motion Controller for Authentication via Hand Geometry and Gestures," *Human Aspects of Information Security, Privacy, and Trust*, LNCS 9190, Springer, 2015, pp. 13–22.

Robert St. Amant is an associate professor in the Computer Science Department at North Carolina State University. His research interests include human-computer interaction, intelligent user interfaces, and cognitive modeling. St. Amant has a PhD in computer science from the University of Massachusetts. Contact him at stamant@ncsu.edu.

David L. Roberts is an assistant professor of computer science at North Carolina State University, where he directs the Computational Intelligence and Interactive Games Research Lab. His research interests lie at the intersection of behavior, data, and computational modeling, with a particular emphasis on the role of computation in understanding and influencing behavior. Roberts has a PhD in interactive computing from the Georgia Institute of Technology. His work has been featured in national and international media outlets such as *The Wall Street Journal, PBS Newshour, IFLS*!, and the BBC. Contact him at robertsd@csc.ncsu.edu.

This article originally appeared in IEEE Internet Computing, *vol. 20, no. 4, 2016.*

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

PREFERRED PLUS

TRAINING & DEVELOPMENT

RESEARCH

BASIC

STUDENT

New Membership Options for a Better Fit

And a better match for your career goals. Now IEEE Computer Society lets you choose your membership — and the benefits it provides — to fit your specific career needs. With four professional membership categories and one student package, you can select the precise industry resources, offered exclusively through the Computer Society, that will help you achieve your goals.

IEEE computer society

Learn more at www.computer.org/membership.

IEEE Computer Society Is Where You Choose the Resources that Fit Your Career

Find the membership that fits you best. IEEE Computer Society lets you choose your membership — and the benefits it provides — to meet your specific career needs. With four professional membership categories and one student package, you can select the precise industry resources, offered exclusively through the Computer Society, that will help you achieve your goals.

Select your membership	Preferred Plus		Training & Development		Research		Basic		Student
	$60 *IEEE Member*	**$126** *Affiliate Member*	**$55** *IEEE Member*	**$115** *Affiliate Member*	**$55** *IEEE Member*	**$115** *Affiliate Member*	**$40** *IEEE Member*	**$99** *Affiliate Member*	**$8** *Does not include IEEE membership*
Computer magazine (12 digital issues)*	✓		✓		✓		✓		✓
ComputingEdge magazine (12 issues)	✓		✓		✓		✓		✓
Members-only discounts on conferences and events	✓		✓		✓		✓		✓
Members-only webinars	✓		✓		✓		✓		✓
Unlimited access to *Computing Now*, computer.org, and the new mobile-ready myCS	✓		✓		✓		✓		✓
Local chapter membership	✓		✓		✓		✓		✓
Safari Books Online (600 titles and 50 training videos)	✓		✓						✓
Skillsoft online solutions (courses, certifications, practice exams, videos, mentoring)	✓		✓						✓
Two complimentary Computer Society magazine subscriptions	✓				✓				
myComputer mobile app	*30 tokens*				*30 tokens*				*30 tokens*
Computer Society Digital Library	*12 FREE downloads*		*Member pricing*		*12 FREE downloads*		*Member pricing*		*Included*
Training webinars	*3 FREE webinars*		*3 FREE webinars*		*Member pricing*		*Member pricing*		*Member pricing*
Priority registration to Computer Society events	✓								
Right to vote and hold office	✓		✓		✓		✓		
One-time 20% Computer Society online store discount	✓								

* *Print publications are available for an additional fee. See catalog for details.*

IEEE computer society

www.computer.org/membership

EDITORS: Rick Kuhn, US National Institute of Standards and Technology, kuhn@nist.gov
Tim Weil, Scram Systems, tweil.ieee@gmail.com

Addressing Pressing Cybersecurity Issues through Collaboration

Bill Fisher,
National Cybersecurity Center of Excellence

Earlier this year, a hospital in California was forced to pay $17,000 to recover its files from a ransomware attack.[1] This is just one of many examples that highlight a concerning trend in malware attacks, which the Cyber Threat Alliance estimates resulted in $325 million in damages in 2015.[2] But looking only at the economic loss is insufficient. In the case of the California hospital, the organization had to revert to using pen and paper for medical records. The operational and reputational damage might be hard to quantify, but many will agree that it is as important as the pure economic loss.

To improve cybersecurity within the US economy, collaborations such as information-sharing organizations have been developed, including sector-specific centers such as the Financial Services Information Sharing and Analysis Center (FS-ISAC). These organizations focus on sharing cyber threat intelligence to help combat critical cyber threats facing businesses today.

The National Cybersecurity Center of Excellence (NCCoE) at the US National Institute of Standards and Technology (NIST) is another organization that relies on collaboration but focuses on building reference designs or example solutions that others can use to improve an organization's cybersecurity.

As the US's only national lab focused solely on cybersecurity, the NCCoE works with the private sector, academia, and other government organizations to identify pressing cybersecurity problems. In collaboration with these communities, NCCoE engineers generate detailed technical problem descriptions and identify the criteria for a reference design utilizing current standards and best practices. After seeking comments from the public to ensure the relevance of an example solution, the NCCoE works with technology vendors to use commercially available tools to develop a worked example of the reference design in the lab. The reference design and architectures are published as NIST Cybersecurity Practice Guides (Special Publication series 1800)—freely available documents that detail how to implement and adopt the reference designs.

In the case of the attack noted earlier, the NCCoE is working on a new project, Data Integrity, to help minimize the risks posed by ransomware and destructive malware. In fact, in 2016 alone, the NCCoE

 Published by the IEEE Computer Society 2469-7087/16/$33.00 © 2016 IEEE

is supporting 10 projects focused on improving the cybersecurity of businesses' infrastructure, including Domain Name System (DNS) Based Secured Email and Derived Personal Identity Verification (PIV) Credentials. We next look briefly at each project, and describe how you can add your expertise to shape the project and contribute to a reference design.

About the NCCoE

The National Cybersecurity Center of Excellence was created in February 2012 as a public-private partnership between the US Commerce Department's National Institute of Standards and Technology (NIST), the State of Maryland, and Montgomery County. As the only national laboratory focused on cybersecurity, the NCCoE collaborates with members of industry, government, and academia to identify and address businesses' most pressing cybersecurity challenges with practical, standards-based solutions using commercially available technologies.

Data Integrity

Businesses face a near-constant threat of destructive malware, ransomware, and malicious insider activities that can alter or destroy critical data. Customer data, transaction records, and correspondence can be rich targets for malicious actors looking to insert, modify, or delete information.

Typical attack vectors include phishing, drive-by website downloads, unmitigated vulnerabilities on external-facing resources, and malicious or infected attachments. Once malware gains a foothold in an organization, it can use multiple techniques to spread and corrupt data. Even honest mistakes can alter data in ways that could cause a significant loss to a company's reputation, business operations, and bottom line. These types of data integrity attacks, especially when they target an entire organization, can have a catastrophic impact on an enterprise's ability to operate. To reduce this risk, organizations need to be able to recover quickly from a data integrity attack and trust the accuracy and precision of the recovered data.

Multiple systems need to work together to prevent, detect, notify, and recover from events that corrupt data. The NCCoE Data Integrity project will explore methods to effectively recover operating systems, databases, user files, applications, and software and system configurations. It will also explore auditing and reporting issues to support recovery and investigations, including user activity monitoring, file system monitoring, database monitoring, scanning backups and snapshots for malware, and rapid-recovery solutions.

More detailed project information is available at https://nccoe.nist.gov/projects/building_blocks/data_integrity.

DNS-Based Secure Email

For most organizations, email is the foundation for conducting business and communicating in today's world. This also means it is a critical target for malicious attacks. The need to protect business plans and strategies; the integrity of transactions, and financial and other proprietary information; and the privacy of employees and clients are only some of the factors motivating organizations to secure their enterprise email.[3]

Transport Layer Security (TLS) cryptographic functions are the primary method organizations use to secure email transactions, whether the need is to authenticate the source of an email message, detect the alteration of messages by some unauthorized party, or keep message content confidential. Unfortunately, many current email security mechanisms are vulnerable to, and have been defeated by, attacks on the integrity of the cryptographic implementations on which they depend.

The most common malicious incidents include intrusion and man-in-the-middle attacks during cryptographic service negotiation, resulting in reading or modification of information by unauthorized third parties. An attacker can pose as one of the parties to an email exchange and send email that contains links to malware-ridden websites. If other content in a fraudulent message successfully motivates the user to click on the link, or if the user's system is configured to automatically follow some links or download content other than text, the malware will infect the user's system. Most confirmed data breaches involve the inclusion of links to malware.

Although tools that implement security standards—such as Domain Name System-Based Authentication of Named Entities (DANE) and other DNS Security Extensions (DNSSEC) applications—are available to combat these threats, there is a shortage of easily ported software libraries, and many current implementations degrade mail-delivery performance.

The NCCoE is addressing these challenges through collaborative efforts with email service providers and cybersecurity technology vendors to demonstrate tools that help organizations to automatically implement the following without degrading service:

- encrypt email traffic between servers across organizational boundaries;
- allow individual email users to digitally sign or encrypt email messages to other end users;

- allow individual email users to obtain other users' certificates to validate signed email or send encrypted email; and
- generate information that can be queried by email recipients to identify valid email senders for a domain and determine that a given message originated from one of these valid senders.

More detailed project information is available at https://nccoe.nist.gov/projects/building_blocks/secured_email.

Derived PIV Credentials

Organizations protect their information systems, in part by limiting access to the minimum set of users required to perform a function. The federal government, and increasingly the private sector, rely on multifactor authentication with smart cards or badges, and PINs to verify that users are who they say they are.

In the federal government, identity verification for physical and logical access requires smart-card-based credentials that conform to the Personal Identity Verification (PIV) of Federal Employees and Contractors standard (FIPS 201-2). Although many desktop and laptop computers have built-in card readers, enterprises today rely heavily on mobile devices (that is, smartphones and tablets) that do not easily accommodate card readers. Despite the movement toward mobile work, multifactor authentication continues to be an important security mechanism.

To help private and public sector organizations increase the deployment and use of strong authentication in mobile environments, the NCCoE will demonstrate, initially using PIV cards, how derived smart card credentials can be used with mobile devices for remote authentication to IT systems in operational environments.

More detailed project information is available at https://nccoe.nist.gov/projects/building_blocks/piv_credentials.

Share Your Expertise

The NCCoE encourages participation from the private, public, and academic sectors. Your expertise helps inform and guide a project and supports the broad outreach and engagement necessary to improve cybersecurity.

Join a Community of Interest

Those interested in supporting the NCCoE's projects are invited to join a *community of interest* (COI)—a group of professionals that share business insights, technical expertise, challenges, and perspectives to guide NCCoE projects. The NCCoE relies on this robust collaboration with experts and innovators to inform real-world cybersecurity challenges and our reference designs. COIs often include senior-level professionals and researchers from across the private, public, and academic sectors, and members typically meet monthly by teleconference.

To find out which projects are currently seeking COI members, visit https://nccoe.nist.gov/about_the_center/COI.

Technology Vendors

Our focus on using commercially available technologies means that we often work with vendors. We accept *letters of interest* (LOIs) from technology companies interested in collaborating on our projects. Partner companies help build a reference design that will result in a publicly available NIST Cybersecurity Practice Guide. Several projects are currently in need of technology vendors to participate in the development of a reference design. For more information on how to participate, please check the Federal Register at https://www.federalregister.gov and search for the following projects:

- Derived PIV Credentials,
- DNS-Based Secure Email,
- Mobile Device Security,
- Attribute Based Access Control,
- Access Rights Management for the Financial Services Sector, or
- Securing Wireless Medical Infusion Pumps.

Public Comments and Suggestions

The NCCoE is always seeking feedback and suggestions. We gladly welcome public involvement, from comments on our current projects to ideas about emerging cybersecurity issues to address. Whether as a COI member, technology vendor, or member of the general public, your input is key to fulfilling our vision of cultivating a secure cyber infrastructure that inspires technological innovation and fosters economic growth. Visit us at https://www.nccoe.nist.gov, or email us at nccoe@nist.gov to find out more about how you can participate. IT

References

1. R. Winton, "Hollywood Hospital Pays $17,000 in Bitcoin to Hackers; FBI Investigating," *Los Angeles Times*, 18 Feb. 2016; www.latimes.com/business/technology/la-me-ln-hollywood-hospital-bitcoin-20160217-story.html.
2. *Lucrative Ransomware Attacks: Analysis of the CryptoWall Version 3 Threat*, tech. report, Cyber Threat Alliance, Oct. 2015; http://cyberthreatalliance.org/cryptowall-report.pdf.
3. K.R. Harney, "Cybercriminals Sneak in to Realty Deals and Sneak Out with $100,000 or More," *Washington Post*, 30 Mar. 2016; https://www.washingtonpost.com/realestate/cybercriminals-sneak-in-to-realty-deals-and-sneak-out-with-100000-or-more/2016/03/29/d3f77a20-f4f9-11e5-9804-537-defcc3cf6_story.html.

Bill Fisher *is a security engineer at the National Cybersecurity Center of Excellence (NCCoE), where he leads a team of engineers that work collaboratively with industry partners to address cybersecurity business challenges facing the US. He leads the center's Attribute Based Access Control (ABAC) project and focuses on all topics related to identity and access management. Fisher's research interests include identity and access management for autonomous nonhuman entities, cloud security, context-driven access management, and mobile application single sign-on. Contact him at William.Fisher@nist.gov.*

This article originally appeared in IT Professional, *vol. 18, no. 4, 2016.*

Editor: **Jane Cleland-Huang**
DePaul University
jhuang@cs.depaul.edu

Keeping Ahead of Our Adversaries

Jane Cleland-Huang, Tamara Denning, Tadayoshi Kohno, Forrest Shull, and Samuel Weber

EVERY SOFTWARE SYSTEM is potentially vulnerable in ways that aren't always imagined during development. For example, pacemakers and implantable cardioverter-defibrillators (ICDs), which monitor and regulate cardiac rhythms, typically provide wireless access to healthcare providers so that they can modify settings and collect telemetry. However, a malicious user could transmit commands to ICDs to collect private data or change the device's therapy settings.[1,2] Recently, well-known hacker Barnaby Jack claimed to have developed software that let him shock patients within a 50-foot radius. Anticipating such potential threats, doctors proactively disabled the wireless features of former US Vice President Dick Cheney's pacemaker.

Malicious attacks' potential to cause real (and diverse) harm holds true for numerous other software systems. For example, University of Michigan researchers demonstrated how easy it was to take control of a traffic light system: a person could ensure that the lights were always green along his or her route or could seriously disrupt traffic by turning all the lights red.[3] Similarly, University of California, San Diego and University of Washington researchers used a car's telematics unit to remotely disable the brakes, turn off the headlights, and manipulate dashboard gauges.[4,5]

White-collar crime involving data breaches are rampant, and governments are investigating the potential for terrorist attacks on power grids, airplanes, and other public services. Technology is a double-edged blade: although computers let us pursue ever-more-impressive innovations, we're likewise subjected to growing possibilities for abuse.

So, how do we build secure products that are hardened against adversarial attacks? Let's take a look.

Thinking about Threats

Many steps to improve security can be taken at various stages of software development. However, an important place to begin is with a dedicated analysis of potential threats. Without a sound understanding of the possible threats against a given system, it's unlikely that developers will be able to adequately defend against them.[6] Surprisingly, this step is often performed hastily or skipped. One problem is that developers often assume they understand all common attack patterns and therefore fail to explore each system's specific vulnerabilities. Alternatively, they might assume they can patch in security later in the design by following accepted security policies and procedures.

To some extent, each system is unique. Even supposing that the system components are well understood and previously have been composed in the same way, the ways a deployed system is used, misused, or reappropriated can introduce unanticipated security vulnerabilities. Building a secure system requires proactive, rigorous analysis of the threats to which it might be exposed, followed by systematic transformation of those threats into security-related requirements. These requirements can then be tracked throughout the development life cycle.

Threat modeling aims to

- identify attackers' potential abilities and goals and
- catalog possible threats that the system must be designed to mitigate.

We consider threat modeling a requirements activity. The most benefit comes from understanding what security requirements are needed and

 Published by the IEEE Computer Society 2469-7087/16/$33.00 © 2016 IEEE

using those requirements to drive architecture decisions, develop test strategies, and engage in other software development activities.

However, in reality, threat-modeling techniques vary and can be applied to both existing and greenfield systems; different techniques are more suited to different software development activities and different development domains. Some techniques are like checklists, enumerating possible threats developers should consider in the context of their system. Others are less deterministic and try to inject more creativity to stimulate thinking about unusual attack vectors. All techniques encourage developers to think more critically about their system and about ways to subvert it; this contrasts with the more usual approaches that focus on functionality. As you might expect, developers find it exceptionally difficult to be complete and consistent and to truly put themselves in the shoes of an attacker.

Security Cards: A Threat Brainstorming Toolkit

To assist threat analysis, Tamara Denning, Batya Friedman, and Tadayoshi Kohno developed the Security Cards.[7] The Security Cards consist of 42 cards divided into four categories, or dimensions: Human Impact, Adversary's Motivations, Adversary's Resources, and Adversary's Methods.

Here, we illustrate how the cards might serve as starting points to explore potential threats to a technological system—in this case, an ICD. This thought exercise is only to explore what these software development processes would look like when applied to a system concept. We don't intend to make statements regarding current ICD security or what security considerations have been incorporated into the development process.

Human Impact

This dimension explores how security breaches could affect humans. The impacts range from personal-privacy violations to widespread societal impact. Threat-modeling sessions could start by ranking the Human Impact cards according to their relevance to the system under consideration. In this case, a highly relevant card is the Physical Well-being card (the first card in Figure 1). It asks us to think about how a misused or compromised ICD could impact people's physical well-being. However, we could also consider cards such as Emotional Wellbeing (for example, patients are aware of the threat to their health), Financial Wellbeing or Relationships (for example, the attack aims to discredit

FIGURE 1. Four Security Cards. Developers can use Security Cards to explore potential threats to a technological system—in this case, an implantable cardioverter-defibrillator (ICD).

the ICD company), or Personal Data (for example, the attacker wants to use the identifying personal data stored on the device).

Adversary's Motivations

This dimension explores why someone might want to attack a system. It helps provide a framework to explore a potential attack's scope and intended targets. For example, the Malice or Revenge card (the second card in Figure 1) might lead us to consider the situation in which an adversary attacks the ICD user owing to extreme emotion. Other motivation-related cards could include Self-Promotion (for example, the attacker wants to demonstrate technical prowess) or Diplomacy or Warfare (for example, the attacker aims to take down a political enemy who happens to have an ICD). Considering potential adversaries' motivations helps determine what resources they might have and helps us construct attacker profiles.

Adversary's Resources

This dimension explores assets an adversary might use to launch an attack. These include hardware and software tools, technical expertise, and various forms of influence. In this case, we select the Expertise card (the third card in Figure 1) and consider the hacker's potential technical skills. Another relevant card could be A Future World, which considers potential future attacks, given that interest exists in increasing the capabilities of remote checkups. We might also consider Impunity (for example, the attack might be difficult to pin on a particular person or to prosecute) or Inside Knowledge (for example, a former employee uses detailed, proprietary knowledge about the architecture).

Adversary's Methods

This dimension explores how an adversary might attack the system, including technology, coercion, and abusing logistical and bureaucratic processes. We might select the Technological Attack card (the fourth card in Figure 1), given that researchers have previously demonstrated such attacks. We also could consider cards such as Multi-Phase Attack (for example, the adversary tampers with software in the doctor's office responsible for sending commands to the ICD), Indirect Attack, or Attack Cover-up.

From Threats to Requirements

Exhaustively cataloging threats is of limited use if we don't use the information to improve the software we're developing. To illustrate moving from threats to requirements, suppose our threat model contains the following threat, written from a malicious user's perspective: "As an IT specialist with intent to physically harm an ICD patient, I'll launch an attack on the device that will change the intended effects on the patient's heart."

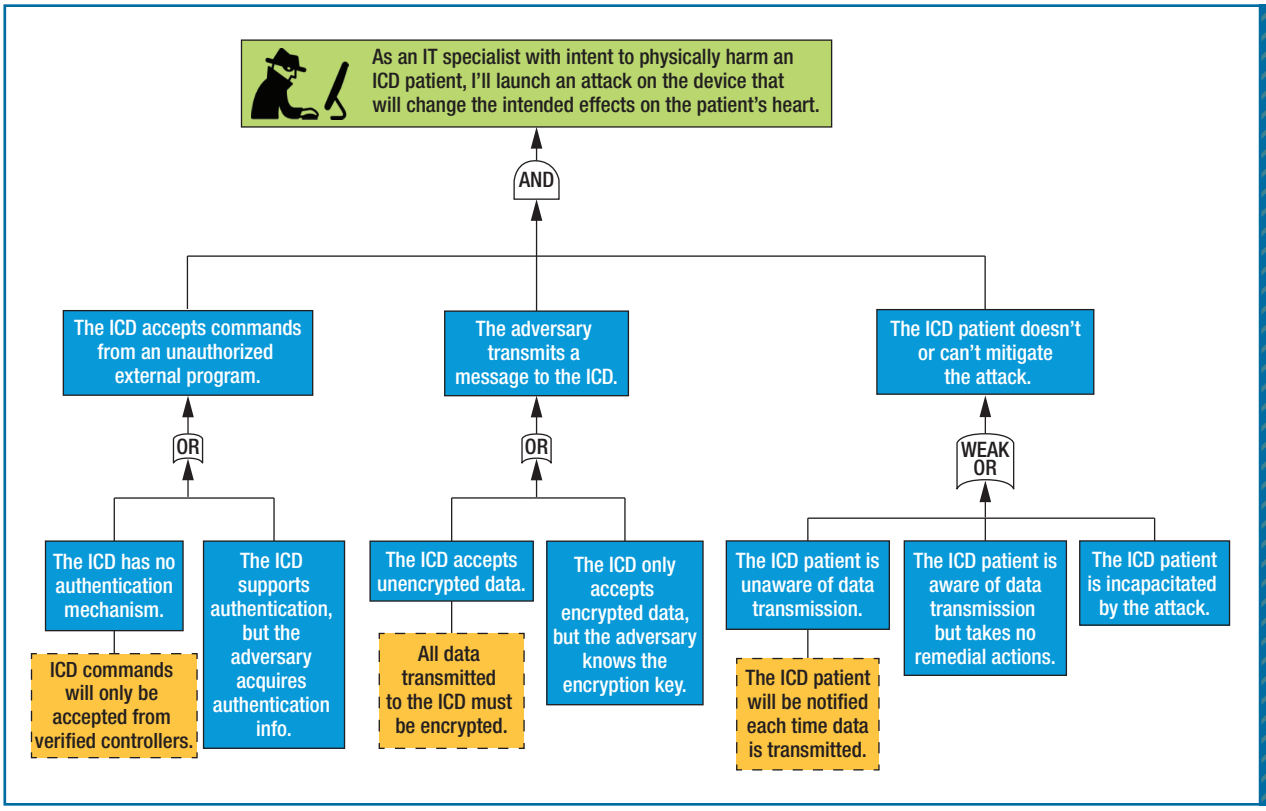

FIGURE 2. A threat tree models system vulnerabilities that potentially enable the threat—in this case, an attack on an ICD. Additional vulnerabilities could exist that the figure doesn't show.

We need to identify and specify requirements that prevent this adversary from achieving this goal. The first step is to identify vulnerabilities that enable each specific threat.

Figure 2 illustrates this through a partial analysis of vulnerabilities and issues that might facilitate an attack on an ICD. The ICD is vulnerable if it lacks an authentication mechanism or the adversary acquires authentication by stealing login information or eavesdropping. To execute the attack, the adversary must successfully transmit a valid command to the ICD. For the attack to succeed, the ICD patient should either be unaware of the attack and therefore unable to take remedial action (such as moving out of transmission range) or be immediately incapacitated.

We then analyze the associated vulnerabilities, evaluate possible mitigations, and specify them as candidate requirements. The scenario in which the ICD has no authentication mechanism is a potential problem for embedded medical devices, in which power consumption is crucial. To address this problem, we consider specifying the requirement, "ICD commands will be accepted only from verified controllers." To address the scenario in which the attacker acquires authentication, we consider specifying the requirement, "All data transmitted to the ICD must be encrypted." However, we must carefully examine both requirements and balance them against the need for additional processing, which would drain battery life. Furthermore, decreased accessibility could inhibit access to the ICD in an emergency.

So, we might consider an alternate requirement. In lieu of limiting access to verified controllers and encrypting data, a next-best option might be to provide an audible warning to ICD patients each time the device starts communicating with a controller. Such a requirement would clearly be a tradeoff. It's unlikely to provide sufficient security in the face of Barnaby Jack's shock attack, for example, but it might partially protect the user from privacy invasions or unauthorized reconfigurations.

Maintaining Traceability

Security requirements are driven by the threat-modeling process but are ultimately constrained by hardware and software tradeoffs. However, being able to maintain traceability from specific requirements back to the threats they address will likely be useful as tradeoffs among requirements are negotiated. It's important to ensure that some mitigation for important threats is maintained, even if the form of that mitigation needs to adapt and evolve.

Luckily, most of us don't have to build systems that must resist attacks by determined nations and whose failure would cause people to die. Attempting to provide perfect security for such systems not only isn't necessary but also almost certainly wouldn't be cost effective. Another benefit of threat modeling (and according to some people, the primary benefit) is documenting what threats won't be mitigated.

Without a documented, consistent understanding of what threats are out of scope, systems typically end up with extremely poor security. For example, many stories exist of systems whose password reset functionalities totally undermined all the other security features. After all, an attacker only needs to target the weakest point in a system's defenses. As another example, many organizations decide that it's not worth building technical solutions to counter insider attacks (employees deliberately doing malicious actions to their employer's systems). However, without explicitly documenting this decision, it's all too easy to overlook that the system must be built so that after employees have been fired, their knowledge of system passwords and procedures can't be used against their ex-employer.

In the end, we all agree that security requirements are needed. However, writing them without engaging in threat modeling will likely lead to cookie-cutter requirements that capture the same old problems. We will probably remember to include standard security functions, although we might not remember to specify them in the requirements document. It's less likely we'll think about specific threats that might be unique to our system. So, threat modeling is an essential activity that should form a natural prelude to the requirements process.

Acknowledgments

This material is based partly on work funded and supported by the US Department of Defense under contract FA8721-05-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center.

References

1. D. Halperin et al., "Pacemakers and Implantable Cardiac Defibrillators: Software Radio Attacks and Zero-Power Defenses," *Proc. 2008 IEEE Symp. Security and Privacy* (SP 08), 2008, pp. 129–142.
2. S. Gollakota et al., "They Can Hear Your Heartbeats: Non-invasive Security for Implanted Medical Devices," *Proc. ACM SIGCOMM 2011 Conf.* (SIGCOMM 11), 2011, pp. 2–13.
3. B. Ghena et al., "Green Lights Forever: Analyzing the Security of Traffic Infrastructure," *Proc. 8th USENIX Workshop Offensive Technologies* (WOOT 14), 2014; www.usenix.org/conference/woot14/workshop-program/presentation/ghena.
4. A. Czeskis et al., "Experimental Security Analysis of a Modern Automobile," *Proc. 2010 IEEE Symp. Security and Privacy* (SP 10), 2010, pp. 447–462.
5. S. Checkoway et al., "Comprehensive Experimental Analyses of Automotive Attack Surfaces," *Proc. 20th USENIX Conf. Security* (SEC 11), 2011, p. 6.
6. A. Shostack, *Threat Modeling: Designing for Security*, John Wiley & Sons, 2014.
7. T. Denning, B. Friedman, and T. Kohno, *The Security Cards: A Security Threat Brainstorming Toolkit*, Univ. Washington, 2013; http://securitycards.cs.washington.edu.

JANE CLELAND-HUANG is a professor of software engineering at DePaul University. Contact her at jhuang@cs.depaul.edu.

TAMARA DENNING is an assistant professor at the University of Utah's School of Computing. Contact her at tdenning@cs.utah.edu.

TADAYOSHI KOHNO is the Short-Dooley Professor in the Department of Computer Science & Engineering and an adjunct associate professor in the Information School at the University of Washington. Contact him at yoshi@cs.washington.edu.

FORREST SHULL is the assistant director for empirical research at Carnegie Mellon University's Software Engineering Institute. He's editor in chief emeritus of *IEEE Software*. Contact him at fjshull@sei.cmu.edu.

SAMUEL WEBER is a senior research staff member at Carnegie Mellon University's Software Engineering Institute, where he's a member of the Science of Cyber-security group. Contact him at samweber@cert.org.

This article originally appeared in IEEE Software, *vol. 33, no. 3, 2016.*

Editor: Gary McGraw, gem@cigital.com

Silver Bullet Talks with Jacob West

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.

Jacob West is the chief architect for security products at NetSuite, where he leads research and development for technology to identify and mitigate security threats. Prior to that, West was the chief technology officer for Enterprise Security Products at HP. In 2007, West coauthored *Secure Programming with Static Analysis* with Brian Chess, which was published in my Addison-Wesley Software Security Series. He is also a coauthor of the BSIMM (Building Security In Maturity Model) and a founding member of the IEEE Center for Secure Design (CSD).

Way back in episode 78 of this show, we discussed the arc of your career from intern to pundit, building static analysis tools, and the BSIMM. I think BSIMM4 had just come out, and now we're on BSIMM6. This time, I'm interested in discussing the CSD work you just published. What is the CSD, and why is it important?

We founded the CSD a couple years ago to shift the focus in security from what we think the main focus has been, which is finding and fixing bugs, to looking for and actively avoiding design flaws that lead to very serious security problems. Instead of just focusing on the implementation, we're thinking about the security implications of the design from a project's inception to its completion.

Can you give examples of a bug and a flaw so people can understand the difference?

An example of a bug is something like cross-site scripting. Let's say developers are trying to build a webpage and want to generate content for their users. What they're not thinking about is an attacker supplying some malicious value. Because they don't scrub that value, they don't validate it before they output it in the page—that's a bug, a mistake they made in their code. The attacker can deliver not only characters the way the programmer might have expected but also script that might run to the user's browser and execute an attack. No one intended for the website to have that feature, but the programmers made a mistake when building the site that allows that attack to succeed. That's a traditional security vulnerability or bug.

A flaw is a design decision—something the system intended to do but that probably isn't a good idea from a security standpoint. One example would be thinking about authentication mechanisms and how we authenticate a user. A system might be designed to allow a simple login and password, and it might be designed to allow any arbitrary password—even very insecure ones like a dictionary word. That wasn't a mistake a programmer made when building the site; it was part of the design. It's a requirement that was missing from the design, which is that the password scheme requires strong passwords.

You gave a great example of cross-site scripting as a bug, but if you think about the APIs that programmers are using to get input and the frameworks they're using, there might be ways to solve that entire class of bugs with a design tweak. In fact, Google has done that.

This is one of my favorite interview questions: asking people to decide

About Jacob West

Jacob West is the chief architect for security products at NetSuite, where he leads research and development for technology to identify and mitigate security threats. He has also served as chief technology officer for Enterprise Security Products at HP, where he founded and led HP Security Research. West coauthored the book *Secure Programming with Static Analysis* with Brian Chess in 2007, which is the only comprehensive guide to using static analysis to avoid the most prevalent and dangerous vulnerabilities in code. West is a founding member of the IEEE Center for Secure Design and the Application Security Advisory Council, and is a coauthor of the Building Security In Maturity Model. He is a frequent keynote speaker at industry events worldwide.

which would be easier to fix, SQL injection or cross-site scripting. For a long time, we've had a designed-in mechanism to make SQL injection very easy to avoid: parameterized queries. And we haven't had, until more recently, an equivalent on the cross-site scripting side.

How prevalent are flaws, and who needs to know about them and avoid them?

Experts like you have often cited a 50/50 split between flaws and bugs that leads to security problems. Some things can be addressed with a design decision; if they're not addressed in that way, they might lead to an implementation bug. So it's not necessarily a black-and-white division. Really, everyone involved in software development and production needs to be aware of secure design principles. Certainly, from an architect's standpoint, as you're imagining the initial system design, security is a huge consideration. But even as smaller design decisions are made, as the overall system coalesces and development proceeds, developers and product managers all need to be aware of the security implications of typically missing security requirements at the design phase and as that design matures throughout the project.

Tell me more about the CSD.

We pulled experts in software security from three main sources. The first was folks like me in the commercial world—many of us have built security products in the past or oversee security practices for our employers. Another group was from academia—folks who are responsible for teaching some of the security implications of the architecture and design principles that the CSD is so concerned about. The third was folks from government who probably have a different view of some of these systems. They might have different considerations from a process standpoint about how systems are designed, or maybe who their attackers are.

We think bringing together these three diverse groups provided a very wide-ranging set of perspectives. We actually asked people to bring real data to the first meeting. The first document we published listed the top 10 software security design flaws. We built that document out of the raw data and experience the group brought together during our first meeting.

That was published in 2014, and you recently published another report with the CSD called *WearFit: Security Design Analysis of a Wearable Fitness Tracker*, which is not a real product. But it brings the top 10 flaws to a real example so people can understand what a flaw might look like in an actual system.

In the initial top 10 flaws document, we give examples wherever we can of how the general guidance for avoiding the flaw applies to different kinds of systems. With this latest report, we inverted that equation and started with a real system design. As you mentioned, WearFit isn't an actual company; we weren't looking to pick on anyone in particular, but we looked at wearable fitness tracking devices across the industry and tried to understand how they're designed with hardware and protocol constraints, implementation decisions, and so on.

We used that information to design a fictitious system that closely resembles real-world systems. Once we had that design, we took the top 10 flaws from the original CSD document and applied them systematically, one by one, to complete a 10-step design review of the WearFit system we created. We talk through what would have happened in that design review, the discussions that would have occurred, and the key parts of the system design that would have been reviewed, and we try to add color about why certain system designs would have been made to achieve certain security properties in the final system.

Who were your coauthors?

I was lucky enough to work with really good friends: Yoshi Kohno from the University of Washington; David Lindsay from Synopsis; and Joe Sechman, who was a colleague of mine at HP Enterprise.

What was the most difficult flaw to work with when thinking about the design?

I think we spent the most time talking about the interplay between privacy and cryptography. This relates to a few different flaws about what's

important and how you protect it. I think this was particularly interesting in the fitness tracker scenario because it's significantly hardware constrained. You have to be able to encrypt data on a very small wearable device with not a lot of power or CPU.

The data's sensitivity also isn't exactly clear. If the data was credit card numbers, no one would argue that it wasn't sensitive information. If it was ambient temperature readings, you can get that information anywhere; it's public already. But what about steps and heart rate and the other information this device is collecting about a person? How sensitive is that? How should it be protected, and how should it be shared? These were topics we spent a lot of time discussing.

Let's talk about the importance of the process by which you find flaws like this. I would imagine you guys did it in an ad hoc manner, but systematizing an approach to architecture risk analysis and threat modeling has been a big challenge.
Documents about process can be off-putting to some people. In our report, we try to walk readers through a process without really talking about the process itself. Up front we talk about the system's design in pretty technical detail. Then we talk about the different attack categories and the kinds of threats we think the system might face. We group those into the high-level buckets—things like denial of service, compromising the device's integrity, stealing a user's health data—and we enumerate examples of those potential attacks.

With the combination of that design and a very ad hoc threat model in terms of what we were concerned about, we were then able to proceed through the top 10 design flaws, thinking about how the design and the threats would interact with a system that was eventually implemented. This worked really well.

Why is design review important, who should do it, how should they do it, and should it be a process?
Everyone should do design review, meaning every organization that's building software. In terms of who should conduct the design review, you have to know something about design, and you have to know something about security. Somebody with architecture chops, whether that's their title or not, is pretty important, but they have to understand the security side of it. Most likely, you need someone with experience in software security specifically.

We need to get to the point where security is treated as a fundamental property of software and, therefore, of computer science.

The vast majority of what we talk about in the report and the design decisions that we believe have security implications are related to nonsecurity functionality. It's not the crypto or the authentication mechanism necessarily; it's about how the system moves data around and services its users. It's not just about the security features; it's about the security implications of the way the rest of the features were designed.

Are there processes for doing design reviews that are more principled than "be really smart and have a lot of security experience"?
I think one of the best ways to learn is to work with someone who has experience. Whether it's someone you hire from a consulting firm or someone you know, or whether you build up that capability internally, the best way to learn is to go on a ride with somebody else and see how they do it.

You mentioned the possibility of figuring out how to do architecture analysis inside your own organization. Do you think that's possible, or do you really need to find somebody who's done this for a while and knows how to look at these problems? Have you had any experience trying to create this capability in an organization?
I have, but it has always involved very exceptional individuals. I don't think it's impossible for a firm to do this on its own; it needs to be able to get the right people, which is a challenge. It's hard to hire a good architect even independent of security. It's not something that I think is 100 percent reproducible, which is one of the reasons we'll continue to see outside firms provide a lot of help in this area.

One of my pet peeves, which is something I speak publicly about as often as I can, is the challenge we have in finding new security people. We have a huge talent shortage in security. Firms can't hire enough people to solve their problems. And that doesn't even account for tackling areas that aren't really being looked at today, like design reviews. At the same time, the top universities and best computer science programs in the country are still doing very little to instruct undergraduates in software security and secure coding.

The industry is going to have to do everything we can, kind of guerilla style, to create these skills. But I think we should also put pressure on universities so they can start to meet this demand as well.

What are your thoughts about the evolution of software security as a discipline since your time as a

student at the University of California, Berkeley?

I think it hasn't evolved nearly enough, frankly. I graduated from Berkeley in 2004 and, at the time, had really no exposure through coursework to any security topic—certainly nothing to do with robust, secure programming skills or secure design principles.

Unfortunately, more than a decade later, we haven't seen much change there. You can graduate from any one of the top universities in the US today with a degree in computer science and really never be exposed to software security topics. That's not to say those classes aren't available in many schools today, but you have to hunt them down.

We need to get to the point where security is treated as a fundamental property of software and, therefore, of computer science. It's taught as part of every discipline we teach today: OSs, networks, data structures, databases. All of these have security implications, and we need to teach these subjects in the right way.

The challenge is, of course, that real-world architecture often differs from academic project architecture. There might be security implications, but you don't get any real experience with walls of code like you might find in a financial institution, for example.

I think we're always going to have a delta between someone leaving a degree program and going into the workforce. There's always going to be a gap between academic scale and commercial scale. What we can do is start to inject more software security along the way, so when graduates get to that final step of scaling up to the real thing, they've at least seen the important building blocks of software security. Today, security is just missing from that equation.

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 via http://garymcgraw.com.

This article originally appeared in IEEE Security & Privacy, *vol. 14, no. 3, 2016.*

IEEE computer society

Read all your IEEE magazines and journals your**WAY** on

myCS

Introducing **myCS**, the digital magazine portal from IEEE Computer Society. Go beyond static, hard-to-read PDFs with an easily accessible, customizable, and adaptive experience.

There's No Additional Cost!

▶ LEARN MORE AT: mycs.computer.org

Rong Yan
Snapchat

Multimedia Hashing and Networking

Wei Liu
Tencent AI Lab

Tongtao Zhang
Rensselaer Polytechnic Institute

Many Internet companies frequently handle heterogeneous, multimedia data. Well-known social media websites, including Facebook, Twitter, and YouTube, along with mobile apps such as Instagram, Snapchat, and WeChat, all face the same problem—how can they efficiently and effectively store, index, search, manage, analyze, and understand multimedia data? Here, we attempt to address this problem by studying two popular topics in multimedia: hashing and networking.

Multimedia Hashing

We explore two different methodologies related to multimedia hashing—shallow-learning-based hashing and deep-learning-based hashing—demonstrating state-of-the-art techniques for enabling efficient multimedia storage, indexing, and retrieval.

Hashing by Shallow Learning

Hashing[1] has attracted considerable attention from researchers and practitioners in computer vision, machine learning, data mining, information retrieval, and other related areas. A variety of hashing techniques have been developed to encode documents, images, videos, or other types of data into a set of binary codes (used as hash keys), while preserving certain similarities among the original data. With such binary codes, similarity searches can be rapidly performed over massive datasets, thanks to the high efficiency of pairwise comparison using the Hamming distance.

Early endeavors in hashing concentrated on employing random permutations or projections to construct hash functions. Well-known representatives include Min-wise Hashing (MinHash)[2] and Locality-Sensitive Hashing (LSH).[3] MinHash estimates the Jaccard set similarity, while LSH accommodates various distance or similarity metrics—such as the ℓ_p distance for $p \in (0, \quad 2]$, cosine similarity, and kernel similarity. Due to randomized hashing, more bits per hash table are required to achieve high precision. This typically reduces recall, and multiple hash tables are thus required to achieve satisfactory accuracy of retrieved nearest neighbors. The overall number of hash bits used in one application can easily run into the thousands.

Beyond data-independent randomized hashing schemes, a recent trend in machine learning is to develop data-dependent hashing techniques that learn a set of compact hash codes based on a training dataset (a multimedia database, for example). Binary codes have been popular in this scenario because of their simplicity and efficiency in computation. The compact hashing scheme can accomplish almost a constant-time nearest neighbor search, after encoding the entire dataset into short binary codes and then aggregating them into a hash table. Additionally, compact hashing is particularly beneficial for storing massive-scale data. For example, saving one hundred million samples, each with 100 binary bits, costs less than 1.5 Gbytes, which can easily fit in memory.

To create effective compact hash codes, numerous methods have been presented, including unsupervised and supervised methods. The state-of-the-art unsupervised hashing method, Discrete Graph Hashing (DGH),[4] leverages the concept of "anchor graphs" to capture the neighborhood structure inherent in a given massive dataset, and then formulates a graph-based hashing model over the entire dataset. This model hinges on a novel discrete optimization procedure to achieve nearly balanced and uncorrelated hash bits, where the binary constraints are explicitly imposed and handled. The DGH technique has been demonstrated to outperform the conventional unsupervised hashing methods, such as Iterative Quantization, Spectral Hashing, and Anchor Graph Hashing,[1] which fail to sufficiently

2469-7087/16/$33.00 © 2016 IEEE Published by the IEEE Computer Society

capture local neighborhoods of raw data in the discrete code space.

The state-of-the-art supervised hashing method, Supervised Discrete Hashing (SDH),[5] incorporates supervised label information and formulates hashing in terms of linear classification, where the learned binary codes are expected to be optimal for classification. SDH applies a joint optimization procedure that jointly learns a binary embedding and a linear classifier. The SDH technique has also been demonstrated to outperform previous supervised hashing methods.[1]

There exist many other interesting hashing techniques, such as document hashing,[6] video hashing,[7] structured data hashing,[8] and intermedia hashing.[9] Note that all of the techniques we have mentioned depend on shallow-learning algorithms. Nonetheless, owing to the high speed of shallow-learning-based hashing, the state-of-the-art hashing techniques have been widely used in high-efficiency multimedia storage, indexing, and retrieval, especially in multimedia search applications on smartphone devices. Several well-known startups, such as Snapchat, Pinterest, SenseTime, and Face++, use proper hashing techniques to manage and search through millions or even billions of images.

Hashing by Deep Learning

Since 2006, *deep learning,*[10] also known as *deep neural networks,* has drawn enormous attention and research efforts in a variety of artificial intelligence areas, including speech recognition, computer vision, machine learning, and text mining. Deep learning aims to learn robust and powerful feature representations for complexly shaped data, so it's natural to leverage deep learning for pursuing compact hash codes, which can be regarded as binarized representations of data. Here, we briefly introduce two recently developed hashing techniques related to deep learning: Convolutional Neural Network Hashing (CNNHash) and Deep Neural Network Hashing (DNNHash).

Previous hashing techniques, relying on deep neural networks, took a vector of hand-crafted visual features extracted from an image as input. The quality of the generated hash codes thus heavily depended on the quality of the hand-crafted features. To remove this barrier, the CNNHash approach was recently developed to integrate image-feature learning and hash-code learning into a joint learning model.[11] This model consists of a stage of learning approximate hash codes given pairwise supervised information and a stage of training a deep Convolutional Neural Network (CNN).[12] Benefiting from the power of CNNs, the latter stage of the joint model can simultaneously learn image features and hash codes, directly working on raw image pixels. The deployed CNN comprises three convolution-pooling layers, a standard fully connected layer, and an output layer with softmax functions. The final hash codes are then produced by quantizing the softmax activations of the output layer.

While the CNNHash approach[11] requires separately learning approximate hash codes to guide the subsequent learning of image representation and finer hash codes, a more recent approach, DNNHash, goes further.[13] With DNNHash, image representation and hash codes are learned in one stage so that representation learning and hash learning are tightly coupled to benefit each other. The DNNHash approach incorporates listwise supervised information to train a deep CNN, leading to a currently deepest architecture for supervised hashing. The pipeline of the deep hashing architecture includes three building blocks:

- a triplet of images, which are fed to the CNN and upon which a triplet ranking loss is designed to characterize the listwise supervised information;
- a shared subnetwork, with a stack of eight convolution layers to generate the intermediate image features; and
- a divide-and-encode module to divide the intermediate image features into multiple channels, each of which is encoded into a single hash bit.

Within the divide-and-encode module, there is one fully connected layer and one hash layer. Eventually, the hash code of an image is yielded by thresholding the output of the hash layer. The DNNHash has been shown to outperform CNNHash and several shallow-learning-based supervised hashing approaches in terms of image search accuracy.[13]

However, for both CNNHash and DNNHash, note that researchers have not yet investigated or reported on the time required for hash code generation. In real-world search scenarios, the speed for generating hashes should be substantially fast. There might be concern about the

hashing speed of these deep-neural-network driven approaches, especially those involving image feature learning, because it might take longer to hash an image with deep learning compared to with shallow-learning-driven approaches.

Multimedia Networking

Here, we introduce the latest Multimedia Information Networks (MINets). As an example of leveraging MINets, we present the cross-media coreference, which incorporates both visual and textual information to reach a sensible event coreference resolution.

Multimedia Information Networks

Recent developments in Web technology—especially in fast connection and large-scale storage systems—have enabled social and news media to publish more in-depth content in a timely manner. However, such developments also raise some issues, such as overwhelming social media information and distracting news media content. In many emergent scenarios, such as encountering a natural disaster (for example, Hurricane Irene in 2011 or Hurricane Sandy in 2012), tweets and news are often repeatedly spread and forwarded in certain circles, so the corresponding content is overlapping. Browsing these messages and pages is unpleasant and inefficient, so an automatic summarization of tweets and news is desired, among which ranking is the most intuitive way to inform users of highly relevant content.

A passive (and common) solution is to prompt users to add more keywords when typing search queries. However, without prior knowledge, and given word limits, it's never trivial to establish a satisfying ranking list for the topics that attract the most public attention. Recent changes in the Google search engine have integrated the image search component and adopted some heterogeneous content analysis. Nevertheless, the connections between images and relevant keywords are still arbitrarily determined by users, so the current search quality is far from optimal.

Active solutions that attempt to summarize information only concentrate on single data modalities. Researchers have developed a context-sensitive topical PageRank method to extract topical key phrases from Twitter as a way to summarize twitter contents.[14] From a new perspective, the Latent Dirichlet Allocation (LDA)[15] model was employed to annotate images,[16] but this doesn't firmly integrate the information across different data modalities. Researchers have also developed a tweet ranking approach,[17] but it only focuses on a single data modality (text).

The connections between images and relevant keywords are still arbitrarily determined by users.

Other conventional solutions for analyzing the relationships or links between data instances include PageRank and VisualRank.[18] The former has been extensively used in heterogeneous networks (webpages and resources), but it mainly concerns linkage. VisualRank, which extends PageRank to the image domain, is a content-based linkage method, but it's confined to homogeneous networks.

A novel MINets[19] representation was recently proposed to create a basic ontology of a powerful ranking system, which aims to integrate cross-media inference and create the linkage among the multimodal information extracted from heterogeneous data. Beyond traditional ranking approaches, designed for homogeneous networks or simple heterogeneous networks, many researchers are developing a series of novel ranking approaches to exploit the properties of MINets, leading to startups such as Toutiao and Tumblr.

Cross-Media Coreference

However, such information networks, where each node represents one event, can suffer from redundant events and low efficiency due to the repeated nodes, because the same stories are often reported by multiple newscast agents. Moreover, to strengthen the impact on audiences and readers, the same stories and events are reported multiple times, especially on TV and in radio broadcasts.

These properties call for automatic methods that can cluster information and remove redundancy. A method has been proposed[20] that not only deals with information from both visual (video contents) and textual (enclosed captions) channels but also analyzes event coreferences.

Thus this method can fully exploit TV news (or newscasts) containing audio and videos.

A good starting point for cross-media coreference is the processing of closed captions (CCs) that accompany videos in a newscast. Such CCs are either generated by automatic speech recognition (ASR) systems or transcribed by a human stenotype operator who inputs phonetics, which are instantly and automatically translated into texts from which events can be extracted. Different from written news, a newscast is often limited in time due to fixed TV program schedules, so anchors and journalists are trained and expected to organize reports that are comprehensively informative with complementary visual and CC descriptions within a short time. These two descriptions have minimal overlap, even though they're interdependent. For example, anchors and reporters introduce background stories that aren't presented in the videos, so the events extracted from the CCs often lack key information about participants.

Another challenge comes from the mistakes that reside in CCs, caused by errors made by human operators or ASR systems. For example, in two similar newscasts, where the death of Jordanian pilot was reported, the closed caption in one cast was mistakenly printed as "It's not clear when it was killed," where "it" should have been "he," referring to the Jordanian pilot.[20] The other newscast had another flawed CC: "Jordan just executed two ISIS prisoners, direct retaliation for the capture of the killing Jordanian pilot" (instead of "capture of the Jordanian pilot"). It's impossible for any existing text-based coreference resolution approach to cluster the two Life.Die event mentions into the same event, because in most natural language processing systems, "it" must not be linked to "Jordanian pilot." Fortunately, videos often illustrate brief descriptions with vivid visual content, and both newscasts adopted the video frames demonstrating the capture of the Jordanian pilot, so these two event mentions can be considered as the same one.

In fact, diverse anchors, reporters, and TV channels tend to use similar or even identical video content to describe the same story, even though they usually use different words and phrases. Therefore, the challenges in coreference resolution methods relying on text information can be addressed by incorporating visual similarity.

Similar work has explored methods for linking visual cues with texts.[21–23] However, these methods mainly focus on connecting image concepts with entities in text mentions, and some didn't clearly distinguish `entity` from `event` in the documents, because the definitions of visual concepts often require both. In addition, the work[21–23] is mostly dedicated to improving visual content recognition by introducing textual features, while the more recent work[20] takes the opposite route by leveraging visual information to improve event coreference resolution.

In the future, we expect to apply the techniques discussed here to make deep learning practical in realistic multimedia applications. For example, we plan to develop deep neural network compressing techniques to endow deep-learning-driven hashing methods with the real-time hashing speed. We also plan to introduce end-to-end memory networks to understand visual and textual information more thoroughly, leading to stronger cross-media event coreference methods. **MM**

References

1. J. Wang et al., "Learning to Hash for Indexing Big Data—A Survey," *Proc. IEEE*, vol. 104, no. 1, 2015, pp. 34–57.
2. A.Z. Broder et al., "Min-Wise Independent Permutations," *J. Computer and System Sciences*, vol. 60, no. 3, 2000, pp. 630–659.
3. A. Andoni and P. Indyk, "Near-Optimal Hashing Algorithms for Approximate Nearest Neighbor in High Dimensions," *Comm. ACM*, vol. 51, no. 1, 2008, pp. 117–122.
4. W. Liu et al., "Discrete Graph Hashing," *Advances in Neural Information Processing Systems*, vol. 27, 2014, pp. 3419–3427.
5. F. Shen et al., "Supervised Discrete Hashing," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 2015, pp. 37–45.
6. H. Li, W. Liu, and H. Ji, "Two-Stage Hashing for Fast Document Retrieval," *Proc. Ann. Meeting of the Assoc. Computational Linguistics*, 2014, pp. 495–500.
7. G. Ye et al., "Large-Scale Video Hashing via Structure Learning," *Proc. IEEE Int'l Conf. Computer Vision*, 2013, pp. 2272–2279.
8. W. Liu et al., "Compact Hyperplane Hashing with Bilinear Functions," *Proc. Int'l Conf. Machine Learning* (ICML), 2012, pp. 17–24.
9. J. Song et al., "Inter-Media Hashing for Large-Scale Retrieval from Heterogeneous Data Sources," *Proc. ACM SIGMOD Int'l Conf. Management of Data*, 2013, pp. 785–796.

10. Y. LeCun, Y. Bengio, and G. Hinton, "Deep Learning," *Nature*, vol. 521, 2015, pp. 436–444.
11. R. Xia et al., "Supervised Hashing for Image Retrieval via Image Representation Learning," *Proc. AAAI Conf. Artificial Intelligence*, 2014, pp. 2156–2162.
12. A. Krizhevsky, I. Sutskever, and G. Hinton, "Imagenet Classification with Deep Convolutional Neural Networks," *Advances in Neural Information Processing Systems*, vol. 25, 2012, pp. 1106–1114.
13. H. Lai et al., "Simultaneous Feature Learning and Hash Coding with Deep Neural Networks," *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 2015, pp. 3270–3278.
14. W. X. Zhao et al., "Topical Keyphrase Extraction from Twitter," *Proc. 49th Annual Meeting of the Assoc. Computational Linguistics: Human Language Technologies—Volume* 1, 2011, pp. 379–388.
15. D. Blei, A. Ng, and M. Jordan, "Latent Dirichlet Allocation," *J. Machine Learning Research*, vol. 3, 2003, pp. 993–1022.
16. Y. Feng and M. Lapata, "Topic Models for Image Annotation and Text Illustration," *Proc. Ann. Conf. North Am. Chapter of the Assoc. Computational Linguistics*, 2010, pp. 831–839.
17. H. Huang et al., "Tweet Ranking Based on Heterogeneous Networks," *Proc. Int'l Committee on Computational Linguistics and the Assoc. Computational Linguistics* (COLING), 2012, pp. 1239–1256.
18. Y. Jing and S. Baluja, "Visualrank: Applying Pagerank to Large-Scale Image Search.," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 30, no. 11, 2008, pp. 1877–1890.
19. T. Zhang et al., "Cross-Media Cross-Genre Information Ranking Multi-Media Information Networks," *V&L Net*, 2014, p. 74.
20. T. Zhang et al., "Cross-Document Event Coreference Resolution Based On Cross-Media Features," *Proc. Conf. Empirical Methods in Natural Language Processing*, 2015, pp. 201–206.
21. V. Ramanathan et al., "Video Event Understanding Using Natural Language Descriptions," *Proc. Int'l Conf. Computer Vision*, 2013, pp. 905–912.
22. V. Ramanathan et al., "Linking People in Videos with 'Their' Names Using Coreference Resolution," *Proc. European Conf. Computer Vision*, 2014, pp. 95–110.
23. C. Kong et al., "What Are You Talking About? Text-to-Image Coreference," *Proc. Conf. Computer Vision and Pattern Recognition*, 2014, pp. 3558–3565.

Wei Liu is a technical leader and research manager at Tencent AI Lab. Contact him at wliu@ee.columbia.edu.

Tongtao Zhang is a PhD candidate in the Computer Science Department at Rensselaer Polytechnic Institute. Contact him at zhangt13@rpi.edu.

This article originally appeared in IEEE Multimedia, vol. 23, no. 2, 2016.

Keeping YOU at the Center of Technology

IEEE Computer Society Publications

Stay Informed

Access to Computer Society books, technical magazines and research journals arm you with Industry intelligence to keep you ahead of the learning curve.

- 3,000 technical books included with membership from books 24 x 7 and Safari Books Online
- 13 technical magazines
- 20 research journals

Learn something new. Check out Computer Society publications today!

Stay relevant with the IEEE Computer Society

More at www.computer.org/publications

IEEE computer society

Steven Gottlieb
Indiana University

The Future of NSF Advanced Computing Infrastructure Revisited

I am in Sunriver, Oregon, having just enjoyed three days at the annual Blue Waters Symposium for Petascale Science and Beyond. It was a perfect opportunity to catch up on all the wonderful science being done on Blue Waters, the National Science Foundation's flagship supercomputer, located at the University of Illinois's National Center for Supercomputing Applications (NCSA). To be honest, you can't really catch up on *all* the science: most of the presentations are in parallel sessions with four simultaneous talks. There were also very interesting tutorials to help attendees make the best use of Blue Waters.

But what I'm most interested in discussing here isn't the petascale science, but the "beyond" issue. *CiSE* readers might recall that in the March/April 2015 issue, I used this space for a column entitled "Whither the Future of NSF Advanced Computing Infrastructure?" (vol. 17, no. 2, 2015, pp. 4–6). One focus of that piece was the interim report of the Committee on Future Directions for NSF Advanced Computing Infrastructure to Support US Science in 2017–2020. This committee was appointed through the Computer Science and Telecommunications Board of the National Research Council (NRC) and was expected to issue a final report in mid-2015 (in fact, it was announced nearly a year later, in a 4 May 2016 NSF press release). I had a chance to sit down with Bill Gropp (University of Illinois Urbana-Champaign), who cochaired the committee with Robert Harrison (Stony Brook) and gave a very well-received after-dinner talk at the symposium about the report.

Over the years, there has been a growing gap between requests for computer time through NSF's XSEDE (Extreme Science and Engineering Discovery Environment) program and the availability of such time. Making matters worse, Blue Waters is scheduled to shut down in 2018. At the symposium, William Kramer announced that the NCSA had requested a zero-cost extension to continue operations of Blue Waters until sometime in 2019. Extension of Blue Waters operations would be a very positive development. Unfortunately, the NSF hasn't announced a plan to replace Blue Waters with a more powerful computer, even in light of the NSF's role in the National Strategic Computer Initiative announced by President Obama on 29 July 2015. There could be a very serious shortage of computer time in the next few years that would broadly impact science and engineering research in the US.

My previous article mentioned that the Division of Advanced Cyberinfrastructure (ACI) is now part of the NSF's Directorate of Computer & Information Science & Engineering (CISE). Previously, the Office of Cyberinfrastructure reported directly to the NSF director. The NSF has asked for comments on the impact of this change, but the deadline is 30 June, well before you'll see this column. The NSF's request for comments was a major topic of conversation in an open meeting at the symposium held by NCSA Director Ed Seidel. I plan to let the NSF know that I think it's essential to go back to the previous arrangement: scientific computing isn't part of computer science, and it's very important that the people at the NSF planning for supercomputing be at the same level as the science directorates in order to get direct input on each directorate's computing needs.

The committee report I mentioned earlier has seven recommendations, most of which contain subpoints (see the "Committee Recommendations" sidebar for more information). The recommendations are organized into four main issues: maintaining US leadership in science and engineering, ensuring that resources meet community needs, helping computational scientists deal with the rapid changes in high-end computers, and sustaining the

 Published by the IEEE Computer Society 2469-7087/16/$33.00 © 2016 IEEE

Committee Recommendations

The full report is at http://tinyurl.com/advcomp17-20; the text here is a verbatim, unedited excerpt, reprinted with permission from "Future Directions for NSF Advanced Computing Infrastructure to Support US Science and Engineering in 2017-2020," Nat'l Academy of Sciences, 2015 (doi:10.17226/21886).

A: Position US for continued leadership in science and engineering

Recommendation 1. NSF should sustain and seek to grow its investments in advanced computing—to include hardware and services, software and algorithms, and expertise—to ensure that the nation's researchers can continue to work at frontiers of science and engineering.

Recommendation 1.1. NSF should ensure that adequate advanced computing resources are focused on systems and services that support scientific research. In the future, these requirements will be captured in its road maps.

Recommendation 1.2. Within today's limited budget envelope, this will mean, first and foremost, ensuring that a predominant share of advanced computing investments be focused on production capabilities and that this focus not be diluted by undertaking too many experimental or research activities as part of NSF's advanced computing program.

Recommendation 1.3. NSF should explore partnerships, both strategic and financial, with federal agencies that also provide advanced computing capabilities as well as federal agencies that rely on NSF facilities to provide computing support for their grantees.

Recommendation 2. As it supports the full range of science requirements for advanced computing in the 2017-2020 timeframe, NSF should pay particular attention to providing support for the revolution in data driven science along with simulation. It should ensure that it can provide unique capabilities to support large-scale simulations and/or data analytics that would otherwise be unavailable to researchers and continue to monitor the cost-effectiveness of commercial cloud services.

Recommendation 2.1. NSF should integrate support for the revolution in data-driven science into NSF's strategy for advanced computing by (a) requiring most future systems and services and all those that are intended to be general purpose to be more data-capable in both hardware and software and (b) expanding the portfolio of facilities and services optimized for data-intensive as well as numerically-intensive computing, and (c) carefully evaluating inclusion of facilities and services optimized for data-intensive computing in its portfolio of advanced computing services.

Recommendation 2.2. NSF should (a) provide one or more systems for applications that require a single, large, tightly coupled parallel computer and (b) broaden the accessibility and utility of these large-scale platforms by allocating high-throughput as well as high-performance work flows to them.

Recommendation 2.3. NSF should (a) eliminate barriers to cost-effective academic use of the commercial cloud and (b) carefully evaluate the full cost and other attributes (e.g., productivity and match to science work flows) of all services and infrastructure models to determine whether such services can supply resources that meet the science needs of segments of the community in the most effective ways.

B. Ensure resources meet community needs

Recommendation 3. To inform decisions about capabilities planned for 2020 and beyond, NSF should collect community requirements and construct and publish roadmaps to allow NSF to set priorities better and make more strategic decisions about advanced computing.

Recommendation 3.1. NSF should inform its strategy and decisions about investment trade-offs using a requirements analysis that draws on community input, information on requirements contained in research proposals, allocation requests, and foundation-wide information gathering.

Recommendation 3.2. NSF should construct and periodically update roadmaps for advanced computing that reflect these requirements and anticipated technology trends to help NSF set priorities and make more strategic decisions about science and engineering and to enable the researchers that use advanced computing to make plans and set priorities.

Recommendation 3.3. NSF should document and publish on a regular basis the amount and types of advanced computing capabilities that are needed to respond to science and engineering research opportunities.

Recommendation 3.4. NSF should employ this requirements analysis and resulting roadmaps to explore whether there are more opportunities to use shared advanced computing facilities to support individual science programs such as Major Research Equipment and Facilities Construction projects.

Recommendation 4. NSF should adopt approaches that allow investments in advanced computing hardware acquisition, computing services, data services, expertise, algorithms, and software to be considered in an integrated manner.

Recommendation 4.1. NSF should consider requiring that all proposals contain an estimate of the advanced computing resources required to carry out the proposed work and creating a standardized template for collection of the information as one step

of potentially many toward more efficient individual and collective use of these finite, expensive, shared resources. (This information would also inform the requirements process.)

Recommendation 4.2. NSF should inform users and program managers of the cost of advanced computing allocation requests in dollars to illuminate the total cost and value of proposed research activities.

C. Aid the scientific community in keeping up with the revolution in computing

Recommendation 5. NSF should support the development and maintenance of expertise, scientific software, and software tools that are needed to make efficient use of its advanced computing resources.

Recommendation 5.1. NSF should continue to develop, sustain, and leverage expertise in all programs that supply or use advanced computing to help researchers use today's advanced computing more effectively and prepare for future machine architectures.

Recommendation 5.2. NSF should explore ways to provision expertise in more effective and scalable ways to enable researchers to make their software more efficient; for instance, by making more pervasive the XSEDE (Extreme Science and Engineering Discovery Environment) practice that permits researchers to request an allocation of staff time along with computer time.

Recommendation 5.3. NSF should continue to invest in and support scientific software and update the software to support new systems and incorporate new algorithms, recognizing that this work is not primarily a research activity but rather is support of software infrastructure.

Recommendation 6. NSF should also invest modestly to explore next-generation hardware and software technologies to explore new ideas for delivering capabilities that can be used effectively for scientific research, tested, and transitioned into production where successful. Not all communities will be ready to adopt radically new technologies quickly, and NSF should provision advanced computing resources accordingly.

D. Sustain the infrastructure for advanced computing

Recommendation 7. NSF should manage advanced computing investments in a more predictable and sustainable way.

Recommendation 7.1. NSF should consider funding models for advanced computing facilities that emphasize continuity of support.

Recommendation 7.2. NSF should explore and possibly pilot the use of a special account (such as that used for Major Research Equipment and Facilities Construction) to support large-scale advanced computing facilities.

Recommendation 7.3. NSF should consider longer-term commitments to center-like entities that can provide advanced computing resources and the expertise to use them effectively in the scientific community.

Recommendation 7.4. NSF should establish regular processes for rigorous review of these center-like entities and not just their individual procurements.

infrastructure for advanced computing. When I asked Gropp about the report's main message, he told me that "the community needs to get involved for the NSF to implement the recommendations." That's because we'll need to do a better job of describing our needs and our scientific plans. Gropp emphasized that it's important to distinguish between our wants and our needs. For example, Recommendation 3 calls on the NSF to collect information on the needs of the scientific community for advanced computing—one possibility is that all grant applications will need to supply information about their computing needs in a standard form (see recommendation 4.1).

The report also emphasizes that data-driven science needs to be supported along with simulation. The latter has often driven machine design, but there are many interesting scientific problems for which access to large amounts of data is the bottleneck, and there are also now many simulations that produce large volumes of data that must be read, stored, and visualized. It will be best to purchase computers that can support both requirements well.

"For many years, we have been blessed with rapid growth in computing power," Gropp stated, but in referring to stagnant clock speeds, he noted, "that period is over." New supercomputers are going to employ new technologies that will require new programming techniques to deal with the massive parallelism and deep memory hierarchies. Gropp quoted Ken Kennedy as saying that software transformations can take

10 years to reach maturity. I note that my own community is eight years into GPU code development and three to four years into development for Intel Xeon Phi. The effort is continuing in anticipation of the next generation of supercomputers. The report strongly emphasizes that the NSF must help users to adapt their codes (Recommendation 5 and its subpoints).

Before my conversation with Gropp ended, I asked him about the delay from the original mid-2015 target date for the report's release. He mentioned the "grueling review process" and the need to respond to every comment. However, he said there were many thoughtful, useful comments and that responding to them made the report much better. Finally, Gropp left me with the thought that "Writing the report is not the end, it is the beginning." I certainly hope that my fellow *CiSE* readers will take that to heart and get involved with helping the NSF plan for our needs for advanced computing. You can find the entire report at http://tinyurl.com/advcomp17-20.

This article originally appeared in Computing in Science & Engineering, *vol. 18, no. 5, 2016.*

Steven Gottlieb is a distinguished professor of physics at Indiana University, where he directs the PhD minor in scientific computing. He's also an associate editor in chief of *CiSE*. Gottlieb's research is in lattice quantum chromodynamics, and he has a PhD in physics from Princeton University. Contact him at sg@indiana.edu.

Are Enemy Hackers Slipping through Your Team's Defenses?

Protect Your Organization from Hackers by Thinking Like Them

Take Our E-Learning Courses in the Art of Hacking

You and your staff can take these courses where you are and at your own pace, getting hands-on, real-world training that you can put to work immediately.

www.computer.org/artofhacking

IEEE NOT SO SECURE IEEE computer society

Editor in Chief: **Diomidis Spinellis**
Athens University of Economics and Business, dds@computer.org

Reflecting on Quality

Diomidis Spinellis

MONITORING ROADWORK quality differs completely from actually building roads. When the road is being built, you take samples of materials and test them in a lab. When it's ready, you use specialized equipment to look for slippery or uneven pavements. And when the road is opened to traffic, you set up cameras and other sensors to see how it's used. In software engineering, we have it much easier because we can monitor how we build software (the process), the software we build (the product), and the product's actual use, simply with yet more software.

There's a deep reason why software systems are reflective—why software is used to build and analyze them. Large software systems form the most complex artifacts designed and built by humans. Managing this complexity requires tools of matching capabilities, so these are necessarily also software. Examples include compilers, runtime libraries, version control systems, issue trackers, application servers, and OSs. Without these the modern software industry would grind to a halt.

The Process, the Product, and the Product's Use

The tools that manage the development process provide ample opportunities to monitor its quality. Every code commit is a heartbeat that can trigger static analysis (for instance, in the form of style checks and code smell detectors); unit, integration, and regression tests; and, inevitably, test coverage analysis. The test results identify possible areas of concern in three dimensions: product functionality, software modules, and developer teams. Feature requests and bug reports on the project's issue tracker let us assess daily progress and, again, pinpoint problem areas.

IDEs provide finer-granularity data on how software is developed before a change matures for an eventual commit. This data can describe crashes, automated refactorings, and newly created entities. Logs of online code reviews reveal the details of caught (and missed) snafus. We can even apply sentiment analysis and other natural-language-processing techniques on the project's email lists, forums, and chat logs to improve our understanding of the developers' performance.

When the software runs, we can either have it instrumented to blab about its quality or apply other tools to it to make it talk involuntarily, as it were. An important element of internal instrumentation is assertion statements: logic fuses that blow when our assumptions about the program state no longer hold. Logging instrumentation, typically implemented through purpose-built libraries and frameworks, can provide extensive details about what's happening in a system, thus letting us reason more deeply about possible problems.

We can also apply software tools that probe or slightly modify the software's internal workings to give us data regarding its functioning. With CPU-profiling tools, we can find where the code spends most of its time, memory profiling lets us see where memory is allocated and leaks, and tracing tools show us library and OS interactions. More intrusive tools help us locate out-of-bounds memory accesses, parallelism bugs, or security vulnerabilities. And when things go south, we can collect and process the details of crash reports, such as the stack trace and the software's log up until the crash.

Finally, we can monitor how our customers actually use the software. We can easily do this externally—for example, by looking at web-server logs or key presses. However, we can obtain much better results if we instrument the software to log its use: invocations, input and output data, button clicks, command exe-

IEEE Software Mission Statement

To be the best source of reliable, useful, peer-reviewed information for leading software practitioners—the developers and managers who want to keep up with rapid technology change.

 Published by the IEEE Computer Society 2469-7087/16/$33.00 © 2016 IEEE

This article originally appeared in IEEE Software, *vol. 33, no. 4, 2016.*

cutions, latency, and throughput. If our software provides a cloud-based service, all we need is some additional logging. Otherwise, our software must ship the corresponding data back to the mothership over the Internet. Increasingly, software also gives its users an explicit say by letting them vote on feature requests or prompting them to fill out satisfaction surveys.

Explicitly designing our development process, our product, and its use to generate precisely the data we need reduces the collection effort and improves the data's quality. This can involve trivial adjustments, such as configuring the format and retention of log files, or larger-scale software instrumentation initiatives. Invariably, well-designed software and processes are also easier to monitor. For example, in one case I could obtain precise usage data from a desktop application because its hundreds of diverse commands were all uniquely identified with a mnemonic string and were dispatched from a single central point.

Exploiting the Data

With so much data reflecting the software's quality readily available, flying blind is inexcusable. When managing a software business, we must ensure that the types of data I outlined are generated, collected, and, more important, used. At a minimum, their widespread availability throughout an organization (subject to appropriate confidentiality safeguards) can help all stakeholders generate the intelligence they require. Other organizations might deliberately institute detailed monitoring procedures for collecting data, triggers that get pulled when something goes wrong, and corrective actions to fix problems. Dashboards, alarms, and periodic reports help us access the data when needed. If all this sounds like a tall order, there are also companies that collect and analyze software data as a service.

Data-driven quality management enables the efficient allocation of finite (and perennially constrained) resources. Recently, while going over a software product's crash logs, I found that just two easily fixed crashes caused more than 20 percent of 1,200 collected crash reports. Other areas in which we can utilize the collected data include feature selection, software performance optimization, team allocation, development process tuning, bug triaging, software evolution planning, hardware allocation, and marketing-channel selection. Significantly, we can expect that our actions' results will later show up in the data we collect, thus giving us feedback on whether we're going in the correct direction.

For extra points, we can look for opportunities arising from integrating process, product, and usage data. For example, consider driving profile-based optimization from actual usage data rather than synthetic benchmarks. Or, we can investigate how software crashes map back to the static analysis or code reviews of the corresponding change.

In the future, when the world is an Internet of Things running on software, road builders and other engineers will have at their disposal the wealth of data we developers take for granted. At that point, we'll have to tell our engineering colleagues how over time we learned to use data to build the quality software our society deserves.

got flaws?

IEEE CYBER SECURITY

Find out more and get involved:

cybersecurity.ieee.org

IEEE CENTER FOR SECURE DESIGN

IEEE computer society

CELEBRATING 70 YEARS

Take the CS Library wherever you go!

IEEE Computer Society magazines and Transactions are available to subscribers in the portable ePub format.

Just download the articles from the IEEE Computer Society Digital Library, and you can read them on any device that supports ePub, including:

- Adobe Digital Editions (PC, MAC)
- iBooks (iPad, iPhone, iPod touch)
- Nook (Nook, PC, MAC, Android, iPad, iPhone, iPod, other devices)
- EPUBReader (FireFox Add-on)
- Stanza (iPad, iPhone, iPod touch)
- ibis Reader (Online)
- Sony Reader Library (Sony Reader devices, PC, Mac)
- Aldiko (Android)
- Bluefire Reader (iPad, iPhone, iPod touch)
- Calibre (PC, MAC, Linux)
  (Can convert EPUB to MOBI format for Kindle)

www.computer.org/epub

IEEE IEEE computer society

Finding the Cybersecurity Job You Want

The rapidly increasing popularity of cloud computing, smartphone technology, big data, and the Internet of Things has made cybersecurity more important than ever. There are thus many opportunities for finding jobs in the field. That is the focus of this month's article, which features an interview with Hewlett Packard Enterprise research engineer Massimo Felici, who co-authored the article "What's New in the Economics of Cybersecurity" in *IEEE Security & Privacy*'s May/June 2016 issue. His research interests include support for information and communication technology deployments, validation of engineering methodologies, and technological complexities throughout application-development lifecycles.

ComputingEdge: What careers in computer technology and cybersecurity will see the most growth in the next several years?

Felici: Among today's main technological trends are cloud computing and containerization, big data and analytics, networking and virtualization, and the Internet of Things. These trends are transforming various industries and have cybersecurity as a key concern. There will thus be an increasing demand for skills—such as automating digital infrastructures, drawing insights from big data analytics, and managing edge points in hybrid ecosystems—that enable cybersecurity within emerging technologies.

ComputingEdge: What advice would you give college students to provide them with an advantage over the competition?

Felici: Industry forecasts say that organizations will need to use their limited resources to protect an increasing amount of data and deal with technological complexity. This will require problem-solving skills to develop innovative, well-grounded solutions addressing cybersecurity and other problems. Students will have to acquire skills combining foundational subjects (such as algorithms, data structures, and probability theory), methodologies (such as machine learning and statistical analysis), engineering practices (such as testing

2469-7087/16/$33.00 © 2016 IEEE Published by the IEEE Computer Society November 2016

and group work), and technologies (such as cloud computing, programming languages, and analytical frameworks).

ComputingEdge: What advice would you give people changing careers midstream?

Felici: Many people already have at least some of the skills necessary to change careers to one in cybersecurity. The field of cybersecurity demands a wide range of capabilities, including some that can be transferred from experience in other industries such as economics, sociology, behavioral sciences, law, and human-computer interaction.

ComputingEdge: What do you consider to be the best strategies for professional networking?

Felici: Engage with professional communities to identify the kinds of problems that concern specific groups in your field. This helps you understand how your skills and expertise relate to professional environments and practices. Don't be afraid to present your work and ideas to others. Any feedback helps you improve your work and position yourself within professional networks. Also, you should work collaboratively with others on projects, which helps you build professional relationships.

ComputingEdge: What should applicants keep in mind when applying for computer and cybersecurity jobs?

Felici: Developing a career in any industry or domain is a bumpy journey. Always be prepared to acquire new skills, and never stop learning.

ComputingEdge's Lori Cameron interviewed Felici for this article. Contact her at l.cameron@computer.org if you would like to contribute to a future ComputingEdge article on computing careers. Contact Felici at massimo.felici@hpe.com.

IEEE Security & Privacy **magazine provides articles with both a practical and research bent by the top thinkers in the field.**

- stay current on the latest security tools and theories and gain invaluable practical and research knowledge,
- learn more about the latest techniques and cutting-edge technology, and
- discover case studies, tutorials, columns, and in-depth interviews and podcasts for the information security industry.

computer.org/security

ENGINEER – SOFTWARE: ERICSSON INC. has an opening for the position of **ENGINEER - SOFTWARE** in **EL SEGUNDO, CA** to work with product management & syst management team to design & implement new software features for mission critical & highly scalable applications. Up to 20% domestic and/or international travel required. To apply mail resume to Ericsson Inc. 6300 Legacy Dr, R1-C12, Plano, TX 75024 & refer to 16-CA-3732.

PROJECT MANAGER: ERICSSON INC. has an opening for the position of **PROJECT MANAGER** in I**RVINE, CA** to establish project plan baseline: define project scope, secure the necessary resources & plans & monitor all activities. Requires up to 10% domestic travel. To apply mail resume to Ericsson Inc. 6300 Legacy Dr, R1-C12, Plano, TX 75024 & refer to 16-CA-2049.

ENGINEER-SERVICES SOFTWARE: ERICSSON INC. has an opening for the position of **ENGINEER- SERVICES SOFTWARE** in **BELLEVUE, WA** to perform customer solution requirements analysis & translate into detailed software requirements. To apply mail resume to Ericsson Inc. 6300 Legacy Dr, R1-C12 Plano, TX 75024 & refer to 16-WA-3661.

TECHNICAL SUPPORT ENGINEER: ERICSSON INC. has an opening for the position of **TECHNICAL SUPPORT ENGINEER** in **ATLANTA, GA** for support of solution deployment, operation readiness, client-acceptance, maintenance, system troubleshooting including incident managements, impact analysis & root cause analysis. To apply mail resume to Ericsson Inc. 6300 Legacy Dr, R1-C12 Plano, TX 75024 & refer to 16-GA-3768.

ENGINEER-SERVICES SOFTWARE: ERICSSON INC. has an opening for the position of **ENGINEER – SERVICES SOFTWARE** in **RESTON, VA** to provide support for Network Development Lab Integration & Certification. To apply mail resume to Ericsson Inc. 6300 Legacy Dr, R1-C12, Plano, TX 75024 & refer to 16-VA-1589.

SOLUTIONS ARCHITECT: ERICSSON INC. has an opening for the position of **SOLUTIONS ARCHITECT** in **PLANO, TEXAS** to analyze customer business plans; propose technical & competence development solutions in new areas & domains to enhance customer's competitive position. Position requires frequent domestic and/or int'l travel. To apply mail resume to Ericsson Inc. 6300 Legacy Dr, R1-C12 Plano, TX 75024 & refer to 16-TX-3838.

Multiple Tenure-Track Faculty Positions in Computer Science

The Department of Computer Science at the National University of Singapore (NUS) invites applications for several tenure-track as well as tenured faculty positions. We have positions dedicated to cyber security, machine learning, robotics (particularly, robot learning), computer vision, computer systems (in particular, embedded systems), data analytics, particularly in statistical and algorithmic foundations, and visual analytics. We also welcome strong applications in other areas of computing. NUS' Department of Computer Science is highly ranked internationally. It enjoys ample research funding, moderate teaching loads, 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 www.comp.nus.edu.sg/careers

NUS offers highly competitive salaries and is situated in Singapore, an English-speaking cosmopolitan city that is a melting pot of many cultures, both the east and the west. Singapore offers high-quality education and healthcare at all levels, high levels of personal freedom and security, as well as very low tax rates.

While we are primarily looking for candidates for Assistant Professor positions, we also welcome applications from exceptional candidates for Associate and full Professor positions. Candidates for Assistant Professor positions should demonstrate excellent research potential and a strong commitment to teaching. Truly outstanding Assistant Professor applicants will also be considered for the prestigious Sung Kah Kay Assistant Professorship. Candidates at more senior levels should have an established record of outstanding and recognized research achievements.

Application Details:

- Submit the following documents (in a single PDF) online via: https://faces.comp.nus.edu.sg
  - A cover letter that indicates the position applied for and the main research interests
  - Curriculum Vitae
  - A teaching statement
  - A research statement
- Provide the contact information of 3 referees when submitting your online application, or, arrange for at least 3 references to be sent directly to csrec@comp.nus.edu.sg
- Application reviews will commence on 1 October 2016 and continue until positions are filled
- Please submit your application by 15 December 2016
- If you have further enquiries, please contact the Search Committee Chair, Weng-Fai Wong, at csrec@comp.nus.edu.sg

Oracle America, Inc.

has openings for

PRODUCT MANAGER, SOLUTIONS MANAGEMENT

positions in **Columbia, MD**.

Job duties include: Plan, initiate, and manage information technology (IT) projects. Travel to various unanticipated sites throughout the United States required. May telecommute from home.

Apply by e-mailing resume to angela.andresen@oracle.com, referencing 385.18285.

Oracle supports workforce diversity.

JUNIPER NETWORKS

Juniper Networks is recruiting for our Sunnyvale, CA office:

ASIC Engineer #26058: Develop test plans and test strategies for module in ASIC chip. Develop test bench using UVM methodology. Write tests and debug the RTL.

Software Engineer #29062: Design, develop, troubleshoot and debug software enhancements for the Company's JUNOS operating system's manageability subsystems.

QA Engineer #11980: Review design specifications and functional specifications for Juniper firewall products. Work with IPv6, IPSec, VPN, NAT, Routing and Switching, and PKI.

Juniper Networks is recruiting for our Herndon, VA office:

Technical Support Engineer #37649: Provide high-level expertise on company specific products. Deliver in-depth diagnostics & root-cause analysis for network impacting issues on Juniper's routing products to large internet service providers & enterprise customers.

Juniper Networks is recruiting for our Westford, MA office:

Sales Demonstration Engineer #33809: Responsible for developing, testing and deploying the equipment and automation to enable the functionality of Juniper Cloud Labs (JCL).

Mail single-sided resume with job code # to
Juniper Networks
Attn: MS A.4.435
1133 Innovation Way
Sunnyvale, CA 94089

CUSTOMER PROJECT MANAGER: ERICSSON INC. has an opening for the position of **CUSTOMER PROJECT MANAGER** in **PLANO, TEXAS** to identify, design, develop, test, & pilot all new processes & tools that can be automated to improve efficiencies in service industry pre-sales & customer project delivery. Up to 50% of domestic travel required. To apply mail resume to Ericsson Inc. 6300 Legacy Dr., R1-C12, Plano, TX 75024 & refer to 16-TX-3805.

ENGINEER-SERVICES SOFTWARE: ERICSSON INC. has an opening for the position of **ENGINEER – SERVICES SOFTWARE** in **PLANO, TEXAS** to provide client support with testing & trials with on live network determined by project need. Up to 20% domestic travel required. To apply, mail resume to Ericsson Inc. 6300 Legacy Dr., R1-C12, Plano, TX 75024 & refer to 16-TX-2584.

COMPUTING edge
WWW.COMPUTER.ORG
/COMPUTINGEDGE

THE HONG KONG POLYTECHNIC UNIVERSITY
香港理工大學

The Hong Kong Polytechnic University (PolyU) is a government-funded tertiary institution in Hong Kong. It offers programmes at various levels including Doctorate, Master's and Bachelor's degrees. It has a full-time academic staff strength of around 1,200. The total consolidated expenditure budget of the University is about HK$6.6 billion (US$1 = HK$7.8 approximately) per year. Committed to academic excellence in a professional context, PolyU aspires to become a world-class university with an emphasis on the application value of its programmes and research. Its vision is to become a leading university that excels in professional education, applied research and partnership for the betterment of Hong Kong, the nation and the world.

The University is now inviting applications or nominations for the following post:

Head of Department of Computing

The successful candidate will be appointed as Chair Professor/Professor, commensurate with his/her qualifications and experience, and hold a concurrent headship appointment. The headship appointment is normally for an aggregate period of six years in two three-year terms of office. Post specification can be obtained from http://www.polyu.edu.hk/hro/job/en/external_adv/deans-heads.php. Other suitable candidate(s), if deemed appropriate by the University, may be appointed as Chair Professor/Professor.

Remuneration and Conditions of Service

Terms of appointment and remuneration package are negotiable and highly competitive. For general information on terms and conditions for appointment of academic staff in the University, please visit the website at http://www.polyu.edu.hk/hro/TC.htm.

Application

Applicants are invited to send detailed curriculum vitae with the names and addresses of three referees and direct any enquiries to **Human Resources Office, 13/F, Li Ka Shing Tower, The Hong Kong Polytechnic University, Hung Hom, Kowloon, Hong Kong [Fax: (852) 2764 3374; E-mail: hrsccomp@polyu.edu.hk], quoting the position being applied for and the reference number. Recruitment will continue until the position is filled. Initial consideration of applications will commence at the end of November 2016.** Candidature may be obtained by nomination. The University reserves the right to make an appointment by invitation or not to fill the position. General information about the University and the Department of Computing is available on the University's Homepage http://www.polyu.edu.hk and http://www.comp.polyu.edu.hk respectively, or from the Human Resources Office [Tel: (852) 3400 3420]. The University Personal Information Collection Statement for recruitment can be found at http://www.polyu.edu.hk/hro/job/en/guide_forms/pics.php.

www.polyu.edu.hk

Opening Minds • Shaping the Future

FACULTY CLUSTER HIRE IN CYBERSECURITY

The University of Texas at San Antonio (UTSA) has embarked on a focused cluster hiring plan under the Gold Star Initiative, to recruit top-tier researchers over a four year period. The plan will focus on strategic areas of research excellence, to include cybersecurity, cloud computing and data analytics. UTSA is currently looking for candidates to fill six faculty positions to foster collaborative research, education and outreach and to create interdisciplinary areas of knowledge that will advance the field of cybersecurity.

UTSA is a recognized leader in the field of infrastructure assurance and security by the National Security Agency and the Department of Homeland Security and is a designated Center of Academic Excellence in Information Assurance Education (CAE). In spring 2014, UTSA was ranked #1 nationally for Cyber Security Programs according to a national survey of certified information technology security professionals conducted for Hewlett-Packard. UTSA is home to the Institute for Cyber Security (ICS), which conducts basic and applied cybersecurity research in partnership with academia, government and industry. The Center for Infrastructure Assurance and Security (CIAS), also located at UTSA, has developed the world's foremost center for multidisciplinary education and development of operational capabilities in the areas of infrastructure assurance and security. In complement to the ICS and CIAS, the Center for Education and Research in Information and Infrastructure Security (CERI2S) conducts high impact research, as well as educates the cybersecurity workforce within the San Antonio area and beyond. In partnership with Rackspace, UTSA houses the largest open cloud infrastructure in academia. The Open Cloud Institute is an initiative to develop degree programs in cloud computing and foster collaboration with industry, positioning UTSA and San Antonio as world leaders in open cloud technology.

Required Qualifications and Responsibilities:

The successful applicant will have a Doctoral degree (Ph.D.) and publications commensurate with appointment levels in the department of interest. Successful candidates will be expected to develop and maintain externally funded research programs, engage in both undergraduate and graduate education, and contribute their leadership and innovative thinking towards global prominence in cybersecurity. Teaching opportunities will vary by department and teaching qualifications will be a consideration for fit within their respective department

Application Process:

Applicants must submit their full application package via the respective link to each position. For more information about this cluster and to access the position links, please visit http://research.utsa.edu/research-news/cyber/. For general questions or additional information on the Gold Star Initiative, please contact: Bernard Arulanandam, Interim Vice President for Research at Bernard.arulanandam@utsa.edu or 210-458-8176.

Enterprise Security

Tenure-track Assistant Professor in the Department of Communication. This position is targeted towards faculty with expertise and interest in the areas of situational awareness and decision-making, cyber data analysis, attack and response, human-machine interactions, organizational communication, information networks, and cybersecurity training.

Cyber Analytics

Tenure-track Assistant, Tenured/tenure-track Associate Professor in the Department of Information Systems and Cyber Security. This position is targeted towards faculty with expertise and interest in conducting transformative research and developing tangible "big data" solutions to cyber analytics challenges with interests in the following domains: embedded system security, cloud security, enterprise security situational awareness and decision making; and/or digital forensics.

Cyber Decision Support

Tenure-track Assistant, Tenured/tenure-track Associate Professor in the Department of Information Systems and Cyber Security. This position is targeted towards faculty with expertise and interest in conducting transformative research that enables organizations to make cyber related decisions quickly, effectively, and accurately. We are particularly interested in the following areas: cyber resiliency; enterprise security situational awareness and decision making; and/or risk assessment and management.

Cloud Computing Security

Tenure-track Assistant Professor in the Department of Electrical and Computer Engineering. This position is targeted towards faculty with expertise and interest in security and privacy issues in cloud computing. Specific topics of interest include secure and privacy-aware data analytics in cloud, data analytics techniques to enhance cloud security, secure software defined networking and network function virtualization, cloud monitoring, dependability issues (availability, assurance and recover) in cloud, secure multi-tenancy, hardware architectures to improve cloud security, etc.

Embedded Systems Security

Tenure-track Assistant Professor in the Department of Electrical and Computer Engineering. This position is targeted towards faculty with expertise and interest in embedded systems security. Areas of particular interest are: security of embedded systems with applications to cyber physical systems such as the Internet of Things, energy, transportation, building design, automation, healthcare and manufacturing.

Privacy and Data Protection

Tenure-track Assistant Professor in the Department of Computer Sciences. This position is targeted towards faculty with expertise and interest in privacy protection and security. Applicants with expertise in software engineering, programming languages and compilers, or big data analytics are particularly encouraged to apply.

As an Equal Employment Opportunity and Affirmative Action employer, it is the policy of The University of Texas at San Antonio to promote and ensure equal employment opportunity for all individuals without regard to race, color, religion, sex, national origin, age, sexual orientation, gender identity, disability, or veteran status. The University is committed to the Affirmative Action Program in compliance with all government requirements to ensure nondiscrimination. The UTSA campus is accessible to persons with disabilities.

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, *Computing Edge* 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, *Computing Edge* 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." *Computing Edge* reserves the right to append to any advertisement without specific notice to the advertiser. Experience ranges are suggested minimum requirements, not maximums. *Computing Edge* 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. *Computing Edge* 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.

CUSTOMER SOLUTIONS SALES MANAGER: ERICSSON INC. has an opening for the position of **CUSTOMER SOLUTIONS SALES MANAGER** in **PLANO, TEXAS** to be accountable for customer centric offerings & maintain responsibility for sales of offerings within respective practice. To apply mail resume to Ericsson Inc. 6300 Legacy Drive, R1-C12, Plano, TX 75024 and refer to 16-TX-148.

ENGINEER-SERVICES SOFTWARE: ERICSSON INC. has an opening for the position of **ENGINEER – SERVICES SOFTWARE** in **PLANO, TEXAS** to develop adaptive solutions for customers including add-on & supplement systs, applications, tools, & scripts on top of the standard products. Up to 50% domestic &/or international travel is required. To apply mail resume to Ericsson Inc. 6300 Legacy Dr., R1-C12, Plano, TX 75024 & refer to 16-TX-3782.

ENGINEER- SERVICES RF: ERICSSON INC. has an opening for the position of **ENGINEER-SERVICES RF** in **PLANO, TX** to perform radio network design, RF tuning & optimization for high capacity wireless networks using GSM, CDMA, LTE and WCDMA technologies. To apply mail resume to Ericsson Inc. 6300 Legacy Dr., R1-C12 Plano, TX 75024 & refer to 16-TX-1913.

ENTERPRISE SAP BASIS MANAGER. Roswell, GA, General Motors. Lead &administer installation, upgrades, &configuration of SAP soft. Design &implement SAP based solutions for company global SAP landscape, &that align with corporate IT strategy in areas of security, availability, reliability, performance, scalability, systems integration, systems monitoring, minimized operating costs, &user satisfaction. Supervise 8 SAP Basis Administrators. Lead investigation into alternative approaches, architecting solutions, documenting &communicating architecture to stakeholders. Analyze SAP Early Watch Alerts to preempt stability &performance issues. Partner with app architects to create/maintain stable SAP solution footprint. Administer SAP ECC, EP, BW/BI, NW, SCM, CRM &Solution Mngr systems. Patch SAP Kernel, upgrade support package, implement OSS notes, install add-on/plug-ins in qlty assurance, training &production systems. Monitor servers, users &trace activities, check sys log for warnings &errors, analyze &resolve the ABAP runtime errors, check lock entries, terminated updates &logs of the regular maintenance jobs. Manage SAP sys refresh, database refresh, on line &off line database &archive backup, restore, &recover, EWA recommendation anlys &implementation. Master, Computer Science. 12 mos exp as Delivery App Architect Manager, SAP Architect or Senior Basis Engineer administering SAP ECC, EP, BW/BI, NW, SCM, CRM &Solution Mngr systems; patching SAP Kernel; implementing OSS notes, installing add-on/plug-ins in qlty assurance, training &production systems; monitoring servers, users &trace their activities, analyzing &resolving ABAP runtime errors, checking lock entries, terminated updates &logs of regular maintenance jobs; &managing SAP sys refresh, database refresh, on line &off line database &archive backup, restore, &recover, with EWA recommendation anlys &implementation. Mail resume to Alicia Scott-Wears, GM Global Mobility, 300 Renaissance Center, Mail Code 482-C32-D44, Detroit, MI 48265, Ref#2215.

PENN STATE | ONLINE

Earn a Master's Degree in Engineering—Entirely Online

Eric Lasway
Systems Engineering Graduate

- Software Engineering
- Systems Engineering
- Engineering Management

- Five- or seven-week courses over six semesters
- GREs not required
- Finish in as little as two years

Achieve your career goals—apply today!

BEST ONLINE PROGRAMS
U.S.News & WORLD REPORT
GRAD ENGINEERING
2016

PENNSTATE
1855
World Campus

worldcampus.psu.edu/psueng

U.Ed.OUT 17-0107/17-WC-0442sms/bjm

IT PROFESSIONALS ESTABLISHED IT COMPANY HAS OPENINGS FOR THE FOLLOWING POSITIONS: BUSINESS ANALYST: will ensure the business & technical needs of clients predominantly in the property & casualty insurance industry are analyzed, documented, & translated into detailed specifications & delivered with quality. (MS deg. or the equiv. in IT, Soft. Engg. or related field & 36 mos. of exp. as a Soft. Engr., Sr. Rating Content Anlst., Bus. &/or Systems Anlst. or related in the commercial property & casualty insurance soft. industry. Will consider candidates with a bachelor's deg. in CS, IT, Soft. Engg., or a related field & 5 yrs. of progressively responsible exp. as a Soft. Engr., Sr. Rating Content Anlst, Bus. &/or Systems Anlst. or related in the commercial property & casualty insurance soft. industry). **SENIOR SYSTEMS ANALYST:** will serve as hands-on developer to deliver one or more modules for projects that provide the functionality & customizations of our core proprietary software products which include complete policy administration software for our customers. (MS deg. or the equiv. in CS, Comp. Info. Systems, Comp. Engg., or related field & 36 mos. of exp. as a Java Dev., Sys. Programmer, or closely related position. Must have 3 yrs. of exp. working with J2EE, Java, XML & HTML. Will consider candidates with a bachelor's deg. in CS, Comp. Info. Systems, Comp. Engg., or related field & 5 yrs. of progressively responsible exp. as a Java Dev., Sys. Programmer, or closely related position). Positions based out of company's US headquarters in Morristown, NJ & subject to relocation to various office & client sites throughout the US. Send resumes to Karen Fernandes, HR Dept., Cover-All, 412 Mt. Kemble Ave., Ste. 110C, Morristown, NJ 07960.

SR. ENGG SVC ARCHTCT. (NY, NY) bld, test & release Svc Vrtlztn s'wre. Prvde on-site & rmte asstnce to spprt & deploy beta prgms, recreate cust issues in CA lab envrmnts & gnrte references. Prvde tech asstnce re: prod dvlpmnt, mktg & quality of solutns. Interact w/ sr mgmt, prod mgmt & mktg. REQ: Bach deg or for equiv in Comp Sci, Math, Engg (any) or rel + 5 yrs of prog exp in job offered &/or a rel occup. Must have exp w/ Svc Vrtlztn using CA LISA; Dsgn & archtng Svc Vrtlztn implmntns; Using Eclipse to crte cust extns; MQ & CICS Transport Protocols; Copybook, CTG, REST, SOAP, & JSON data protocols; Triaging Svc Vrtlztn archtctres & upgrading older Svc Vrtlztn installatns; Prsntg & asstng cust in their adptn & expnsn of Svc Vrtlztn prods; Undrstndg & prsnting DevOps & Cont Dlvry; Bus Trvl Req approx. 25%. Wrk fr home anywhere in the US. Send resume to: Althea Wilson, CA Technologies, 201 North Franklin Street, Suite 2200, Tampa, FL, 33602, Refer to Requisition #139621.

TEST CONSULTANT. Multiple positions available in Enfield, CT. Create testing strategies and identify appropriate levels of testing risk tolerance, including regression, integration, and acceptance testing. Develop and execute detailed testing scenarios and test plans in accordance with testing strategy and risk based test methodology using Quality Center/QTP Communicate progress, defects, issues and risks as they are encountered using the test component of the project communication strategy. Document test results, tracks, and coordinate fixes. Apply: L. Sawtelle, Massachusetts Mutual Life Insurance Company, 1295 State St, Springfield, MA 01111. Please reference Job ID: 708202400.

Three Tenure-Track Assistant Professor Positions, and One Lecturer Position, in Computer Science and Engineering at the University of Nevada, Reno.

The Department of Computer Science and Engineering at the University of Nevada, Reno, invites applications for three tenure-track faculty positions and one Lecturer position. Three of the positions are at the Assistant Professor level with expertise in the areas of (i) high performance computing, (ii) cybersecurity, and (iii) data science and engineering with emphasis on bioinformatics. The fourth position is at the Lecturer level.

Applicants for the Tenure Track Faculty positions must have a Ph.D. in Computer Science or Computer Engineering by July 1, 2017 and must be strongly committed to excellence in research and teaching and should demonstrate potential for developing robust externally funded research programs. Candidates for the Lecturer position, must be strongly committed to excellence in teaching and must have a M.S. or Ph.D. in Computer Science or Computer Engineering by July 1, 2017. In the last five years, the College of Engineering has witnessed an unprecedented growth in student enrollment and number of faculty positions. The College is positioned to further enhance its growth of its students, faculty, staff, facilities as well as its research productivity and its graduate and undergraduate programs.

The University of Nevada, Reno recognizes that diversity promotes excellence in education and research. We are an inclusive and engaged community and recognize the added value that students, faculty, and staff from different backgrounds bring to the educational experience.

Interested candidates must apply online

For the high performance computing position, apply to
www.unrsearch.com/postings/22243
For the cybersecurity position, apply to
www.unrsearch.com/postings/22183
For the data science and engineering position, apply to
www.unrsearch.com/postings/22239
For the lecturer position, apply to
www.unrsearch.com/postings/22187

Application process includes: a detailed letter of application which also indicates how you would contribute to the diversity and excellence of the academic community through your research, curriculum vitae, statement of teaching philosophy, statement of research and plans, (research and plans does not apply the lecturer position) and contact information for three professional references. Review of applications will begin on January 5, 2017 and will continue until the search closes on February 15, 2017. Inquiries should be directed to Ms. Lisa Cody, lacody@unr.edu.

The University of Nevada serves over 21,000 students. The university is ranked as a Tier 1 institution by "U.S. News and World Report" and offers an array of degree programs at all levels. Reno is located in the foothills of the Sierra Nevada, a 30-minute drive from Lake Tahoe. Reno/Tahoe is recognized as a world-class outdoor recreation area. Additional nearby areas of interest include the Black Rock Desert, Sacramento, Yosemite National Park, Napa/Sonoma, the San Francisco Bay Area.

EEO/AA Women and under-represented groups, individuals with disabilities, and veterans are encouraged to apply.

University at Buffalo, The State University of New York

Department of Computer Science and Engineering

Multiple Tenured, Tenure-Track Faculty Positions and Non-Tenure Track Lecturer Positions

The Department of Computer Science and Engineering, University at Buffalo invites candidates to apply for **multiple tenured and tenure-track faculty positions beginning** in the 2017-2018 academic year. Candidates at all ranks from **all areas of computer science and engineering**, including but not limited to areas covered by existing faculty strength such as Algorithms, Big Data, Cyber Security, Cyber Physical Systems (or Internet of Things), Databases, Distributed Systems, Embedded Systems, Machine Learning, Mobile Computing, Multimedia, Pattern Recognition, Robotics, and Theory. Applicants must have a Ph.D. in computer science or a related area by August 2017 and demonstrate potential for excellence in research, teaching, service and mentoring. Applicants from underrepresented groups, especially women and minorities, are strongly encouraged. 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 until **January 15, 2017.** Applicants must submit their application electronically via www.ubjobs.buffalo.edu. Posting number 1600687. The University at Buffalo is an Equal Opportunity Employer.

The Department also invites candidates to apply for non-tenure track lecturer positions beginning in fall 2017. We invite applications from candidates from all areas of Computer Science and Computer Engineering who have a passion for teaching. We are particularly looking for candidates who can operate effectively in a diverse community of students and faculty and share our vision of helping all constituents reach their potential. Applicants from underrepresented groups, especially women and minorities, are strongly encouraged. 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.

Lecturer's 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.

Minimum Qualifications (Position): Ideally, applicants should have a PhD degree in Computer Science, Computer Engineering, or a related field by August 2017. Exceptional applicants with a MS 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.

Applications will be accepted until **January 15, 2017.** Applicants must submit their application electronically via www.ubjobs.buffalo.edu. Posting number: 1600718.

Computer Science and Engineering Department

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

The department currently has 38 tenured/tenure-track faculty, 7 teaching faculty, and approximately 900 undergraduate majors, 450 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 eight 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 is about $5 Million dollars.

University at Buffalo (UB)

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

City and Region

The city of Buffalo is the second largest city in New York state, and was recently voted as one of the top ten best places to live and raise a family by Forbes magazine. 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, and has a packed year-round calendar of cultural events and sporting activities, coupled with relatively low house prices and great schools. The economic renaissance of the region is underlined by a revitalized downtown waterfront and an energetic tech and start-up community. In an extraordinary recognition of Western New York's potential, Governor Andrew M. Cuomo has committed an historic $1 billion investment in the Buffalo area economy to create thousands of jobs and spur billions in new investment and economic activity over the next several years.

SPLUNK INC.. seeks Senior Software Engineer in San Francisco, CA: Design & dev high performance, scalable & stable sw that functions across ops sys & within multiple envir. Refer to Req#9ZZRSQ & 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.

SOFTWARE

Oracle America, Inc.

has openings for

SOFTWARE DEVELOPER

positions in **Reston, VA**.

Job duties include: Design, develop, troubleshoot and/or test/QA software.

Apply by e-mailing resume to shakeeb.rahman@oracle.com, referencing 385.16138.

Oracle supports workforce diversity.

BAYLOR UNIVERSITY

Baylor University is a private Christian university and a nationally ranked research institution, consistently listed with highest honors among *The Chronicle of Higher Education*'s "Great Colleges to Work For." Chartered in 1845 by the Republic of Texas through the efforts of Baptist pioneers, Baylor is the oldest continuously operating university in Texas. The university provides a vibrant campus community for over 15,000 students from all 50 states and more than 80 countries by blending interdisciplinary research with an international reputation for educational excellence and a faculty commitment to teaching and scholarship. Baylor is actively recruiting new faculty with a strong commitment to the classroom and an equally strong commitment to discovering new knowledge as we pursue our bold vision, Pro Futuris. (www.baylor.edu/profuturis/).

FACULTY POSITION

Assistant, Associate or Full Professor of Computer Science in the area of Software Engineering

The Department of Computer Science at Baylor University seeks a productive scholar and dedicated teacher for a tenured or tenure-track position beginning August 2017. Viable candidates must have a PhD in Computer Science or a closely related field, demonstrate scholarly capability and an established and active independent research agenda in software engineering and related areas. The ideal candidate will also have leadership experience, a commitment to undergraduate and graduate education, effective communication and organization skills and a strong research record that includes significant external funding. Qualities of a successful candidate for a senior position include leadership experience and a strong record of an independently funded agenda. All candidates are expected to exhibit a passion for teaching and mentoring at the graduate and undergraduate level.

The Department offers a B.S. in Computer Science degree, a B.A. degree with a major in Computer Science, a B.S. in Informatics degree with a major in Bioinformatics, a B.S. in Computer Science with a major in Computer Science Fellows and a M.S. degree in Computer Science and a Ph.D. degree in Computer Science, and is rapidly expanding with its new Ph.D. program.

Applications will be accepted until the position is filled. To ensure full consideration, complete applications must be submitted by **12/01/2016**.

APPLICATION PROCEDURE: Please submit 1) a letter of application, which includes the applicant's anticipated rank, 2) a current curriculum vitae, 3) a statement of teaching interests, 4) a statement of research plans related to Baylor's programs, 5) transcripts, 6) the names, addresses, and phone numbers of three individuals from whom you have requested letters of recommendation to: Dr. Eunjee Song (Search Committee Chair), Baylor University, One Bear Place #97141, Waco, Texas 76798-7141.

In addition, as an openly Christian institution, Baylor seeks to establish an intentionally Christian environment, where students and faculty are encouraged to pursue both faith and reason. In order to achieve this goal, we ask that applicants submit a brief statement addressing their Christian commitment. This may include church or religious affiliation, active membership in Christian organizations, or a personal statement.

Materials may be submitted electronically to Dr. Eunjee Song at eunjee_song@baylor.edu. Please combine all submitted material into a single pdf file.

To learn more about the above position, the Department of Computer Science, the School of Engineering and Computer Science, and Baylor University, please visit the appropriate URL:

https://jobs.baylor.edu/postings/1276;
http://www.ecs.baylor.edu/computerscience;
http://www.ecs.baylor.edu;
or http://www.baylor.edu.

Baylor University is a private not-for-profit university affiliated with the Baptist General Convention of Texas. As an Affirmative Action/Equal Opportunity employer, Baylor is committed to compliance with all applicable anti-discrimination laws, including those regarding age, race, color, sex, national origin, marital status, pregnancy status, military service, genetic information, and disability. As a religious educational institution, Baylor is lawfully permitted to consider an applicant's religion as a selection criterion. Baylor encourages women, minorities, veterans and individuals with disabilities to apply.

SENIOR SOFTWARE ENGINEER F/T. (Poughkeepsie, NY). Position involves travel to various unanticipated worksites up to 100% of the time anywhere in the United States. Must have Master deg or the foreign equiv in Comp Sci, Comp Sci & Engg, Engg, or related w/1 yr of exp performing application analysis, requirement gathering, design, development, implementation and testing of software applications through full product development life cycle and release process. Create high-level and detailed design documents, technical specifications, flow diagrams, and class diagrams, etc. Responsible for timely delivery of code in accordance with coding standards and best practices. Leads full application development life cycle and assists other team members for any technical challenges. Perform code reviews, release notes walk-throughs, mentor the junior team members, and troubleshoot production issues and provide root cause analysis. Develop Web and Enterprise applications using following tools/technologies: Java/J2EE, Spring, Web Services (SOAP and RESTFul), Hibernate, JDBC, HTML, CSS, AJAX, JavaScript, jQuery, Bootstrap, JSON, XML, JSP, Servlets, Struts, EJB, WebLogic, Tomcat, SQL, Oracle, MySQL, UNIX, Shell, ANT, MAVEN, JUnit. Send resume: Indotronix Int.l Corp., Recruiting (MS), 331 Main St, Poughkeepsie, NY 12601.

TECHNICAL LEAD F/T. (Poughkeepsie, NY). Position involves travel to various unanticipated worksites up to 100% of the time anywhere in the United States. Must have Bach deg or the foreign equiv in Comp Infor Sys, Comp Appls, or related w/5 yrs of progressive exp in the position offered or Master deg or the foreign equiv in Comp Info Sys, Comp Appls, or related w/1 yr of exp leading technical analyst with development expertise on following tools/languages: Oracle 9i, Oracle 11g, SQL Server, SQL Server Query Analyzer, PL/SQL, Optimization of Databases, SAP Data migration using Cransoft and SAP R/3, HP PPM V7.0 and V8.0. Deep understanding of data migration/ transformation. Analyze/develop Business Requirements with Client Stakeholders/Executives and determine test strategy, test deliverables (Plan, scripts, Requirement Traceability Matrix and defect management) and test solution in order to determine project approach for waterfall and Agile Methodology using tools like HP QC, Version 1 and JIRA. Accountable for overall delivery of projects, Operational Performance, Release Management and Testing/QA Services for Client in US. Send resume: Indotronix Int.l Corp., Recruiting (NJ), 331 Main St, Poughkeepsie, NY 12601.

SOFTWARE DEVELOPER. Participate in the full Software Development Lifecycle ("SDLC") involving design, development, testing & implementation in n-tier development environment by utilizing knowledge of Microsoft .NET, SQL Server, JavaScript, HTML. Responsible to maximize customer satisfaction by maintaining, understanding, adding functionality & meeting end-user acceptance criteria through the design & implementation of software test strategies for new features. Analyze, participate, review the requirements, create & assist in performing various types of tests & code review. Mail resumes to Code Ace Solutions Inc. 50 Cragwood Rd. Ste 217, South Plainfield, NJ 07080.

SOFTWARE
Oracle America, Inc.
has openings for
SOFTWARE DEVELOPER
positions in **Morrisville, NC**.

Job duties include: Design, develop, troubleshoot and/or test/QA software. As a member of the software engineering division, apply knowledge of software architecture to perform tasks associated with developing, debugging, or designing software applications or operating systems according to provided design specifications.

Apply by e-mailing resume to david.kesselring@oracle.com, referencing 385.18604.
Oracle supports workforce diversity.

COMPUTER
Oracle America, Inc.
has openings for
PROGRAMMER ANALYST - IT
positions in **Frisco, TX**.

Job duties include: Analyze complex programs and formulate logic for complex internal systems. May telecommute from home.

Apply by e-mailing resume to spencer.chappell@oracle.com, referencing 385.19417.

Oracle supports workforce diversity.

TECHNICAL
Oracle America, Inc.
has openings for
TECHNICAL ANALYST
positions in **Orlando, FL.**

Job duties include: Analyze user requirements to develop, implement, and/or support Oracle's global infrastructure. As a member of the IT organization, assist with the design, development, modifications, debugging, and evaluation of programs for use in internal systems within a specific function area.

Apply by e-mailing resume to sharan.alva@oracle.com, referencing 385.18553.
Oracle supports workforce diversity.

ASSOCIATE OR FULL PROFESSOR OF COMPUTER SCIENCE. The Dartmouth College Department of Computer Science invites applications for a tenured faculty position at the level of associate or full professor. We seek candidates who will be excellent researchers and teachers in the broad range of areas related to cyber-security. This position is the first of three hires that the College anticipates making in the area of cyber-security. We particularly seek candidates who will help lead, initiate, and participate in collaborative research projects within Computer Science and beyond, including Dartmouth researchers from other Arts & Sciences departments, Geisel School of Medicine, Thayer School of Engineering, and Tuck School of Business. The Computer Science department is home to 21 tenured and tenure-track faculty members and two research faculty members. Research areas of the department encompass the areas of security, computational biology, machine learning, robotics, systems, algorithms, theory, digital arts, vision, and graphics. The Computer Science department has strong Ph.D. and M.S. programs and outstanding undergraduate majors. The department's security faculty are affiliated with Dartmouth's Institute for Security, Technology, and Society (ISTS), which also involves faculty from Engineering, Sociology, and Business. Dartmouth College, a member of the Ivy League, is located in Hanover, New Hampshire (on the Vermont border). Dartmouth has a beautiful, historic campus, located in a scenic area on the Connecticut River. Recreational opportunities abound in all four seasons. We seek candidates who have a demonstrated ability to contribute to Dartmouth's undergraduate diversity initiatives in STEM research, such as the Women in Science Program, E. E. Just STEM Scholars Program, and Academic Summer Undergraduate Research Experience (ASURE). We are especially interested in applicants with a demonstrated track record of successful teaching and mentoring of students from all backgrounds (including first-generation college students, low-income students, racial and ethnic minorities, women, LGBTQ, etc.). Applicants are invited to submit a cover letter and CV via Interfolio at http://apply.interfolio.com/36691. Email David.F.Kotz@Dartmouth.edu with any questions. Dartmouth College is an equal opportunity/affirmative action employer with a strong commitment to diversity and inclusion. We prohibit discrimination on the basis of race, color, religion, sex, age, national origin, sexual orientation, gender identity or expression, disability, veteran status, marital status, or any other legally protected status. Applications by members of all underrepresented groups are encouraged. Application review will begin November 1, 2016, and continue until the position is filled.

CALYPSO TECH. seeks Quality Assurance Analyst in SF, CA to guarantee quality & perstnce of SW by testing corrctns of bugs & enhancmnts. Ref Job ID: 9DPSZM & mail res. to Calypso, Attn: HR, 595 Market St, Ste. 1800, SF, CA 94105.

TECHNICAL

Oracle America, Inc.

has openings for

TECHNICAL ANALYST

positions in **Orlando, FL.**

Job duties include: Analyze user requirements to develop, implement, and/or support Oracle's global infrastructure.

Apply by e-mailing resume to andre.luis.priosti@oracle.com, referencing 385.19557.

Oracle supports workforce diversity.

SOFTWARE

Oracle America, Inc.

has openings for

SOFTWARE DEVELOPER

positions in **Broomfield, CO**.

Job duties include: Design, develop, troubleshoot and/or test/QA software.

Apply by e-mailing resume to david.palmer@oracle.com, referencing 385.16927.

Oracle supports workforce diversity.

TECHNICAL

Oracle America, Inc.

has openings for

TECHNICAL ANALYST

positions in **Colorado Springs, CO**.

Job duties include: Deliver solutions to the Oracle customer base while serving as an advocate for customer needs. Offer strategic technical support to assure the highest level of customer satisfaction.

Apply by e-mailing resume to brad.ericksen@oracle.com, referencing 385.20005.

Oracle supports workforce diversity.

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

ALPHARETTA, GA: Software Engineer (Ref.# ALP1): Responsible for the definition, design, development, test, debugging, release, enhancement or maintenance of networking software.

AUSTIN, TX: Database Administrator (Ref.# AUS22): Provide database design and management function for business and/or engineering computer databases. Telecommuting permitted.

BELLEVUE, WA: Network Consulting Engineer (Ref.# BEL4): 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. **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 loctations throughout the United States. **Network Consulting Engineer (Ref.# BEL7)**: Responsible for the support and delivery of Advanced Services to company's major accounts. Telecommuting permitted.

BLOOMINGTON, MN: Network Consulting Engineer (Ref.#BLO1): Responsible for the support and delivery of Advanced Services to company's major accounts. Telecommuting permitted and travel may be required to various unanticipated locations throughout the United States.

BOXBOROUGH, MA: Technical Marketing Engineer (Ref.# BOX18): Responsible for enlarging company's market and increasing revenue by marketing, supporting, and promoting company's technology to customers.

COLUMBUS, IN: Customer Support Engineer (Ref.#: COL1): Responsible for providing technical support regarding the company's proprietary systems and software. Telecommuting permitted.

LINDON, UT: Software Engineer (Ref.#: LIN1): Responsible for the definition, design, development, test, debugging, release, enhancement or maintenance of networking software.

NEW YORK, NY: Network Consulting Engineer (Ref.# NY12): Responsible for the support and delivery of Advanced Services to company's major accounts. Telecommuting permitted.

OVERLAND PARK, KS: Network Consulting Engineer (Ref.# OVE1): Responsible for the support and delivery of Advanced Services to company's major accounts.

RESEARCH TRIANGLE PARK, NC: Test Engineer (Ref.# RTP17): Build test equipment and test diagnostics for new products based on manufacturing designs.

SAN FRANCISCO, CA: CNG Member of Technical Staff (Ref.# SF9): Design, implement, and test software for a web application used by our customers for IT management.

SAN JOSE/MILPITAS/SANTA CLARA, CA: Data Scientist (Ref.# SJ585): Run simulations and statistical models to identify optimal Repair Life Cycle (RLC) conditions and enable exception driven management for Service Supply Chain.

SEATTLE, WA: Network Consulting Engineer (Ref.# SEA9): Responsible for the support and delivery of Advanced Services to company's major accounts. Telecommuting permitted and travel may be required to various unanticipated locations throughout the United States.

PLEASE MAIL RESUMES WITH REFERENCE NUMBER TO CISCO SYSTEMS, INC., ATTN: V51B, 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

BAYLOR UNIVERSITY

Faculty Position

The Electrical and Computer Engineering Department of Baylor University seeks faculty applicants for a Lecturer or a Senior Lecturer Position. Any area of expertise in ECE will be considered but applicants in computer engineering will be given special consideration. Applicants must demonstrate potential for excellent teaching; applicants for Senior Lecturer must additionally present evidence of achievement in 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, 2017 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 interests,
4) the names and contact information for at least four professional references.

Additional information is available at www.ecs.baylor.edu. Should you have any questions on the position, feel free to contact the search chair, Dr. Keith Schubert at keith_schubert@baylor.edu. Please submit materials to https://apply.interfolio.com/38070.

Baylor University is a private not-for-profit university affiliated with the Baptist General Convention of Texas. As an Affirmative Action/Equal Opportunity employer, Baylor is committed to compliance with all applicable anti-discrimination laws, including those regarding age, race, color, sex, national origin, marital status, pregnancy status, military service, genetic information, and disability. As a religious educational institution, Baylor is lawfully permitted to consider an applicant's religion as a selection criterion. Baylor encourages women, minorities, veterans and individuals with disabilities to apply.

SENIOR APPLICATION ENGINEER: Peterson Technology Partners Inc. seeks qualified sr. application engineer for its headquarters located in Rolling Meadows, IL & various & unanticipated work locations throughout the U.S. Resp. for designing & implementing complex & scalable enterprise applications w/ optimum mobile & web interfaces for clients. Master's degree in IT Engg, Info System Technology, IT Mgmt, or a closely related field of study (Will accept Bachelor's degree in above fields plus 5 yrs related progressive exp in lieu of Master's degree) each alternative degree requirements w/ at least 1 yr exp in: (i) database design & devel, algorithms, computer networks, operating systems, service oriented architectures, multimedia networks, software engg, object oriented system analysis; & (ii) developing software solutions using Java, HTML C#, Shell, jQuery, Angular JS, jQuery Mobile, JSON, & Perl. An EOE. Respond by mail to Peterson Technology Partners, 1600 Golf Rd, Ste 1206, Rolling Meadows, IL 60008. Refer to ad code: PTP-0916.

QA ANALYST

Oracle America, Inc.

has openings for

QA ANALYST-ProdDev

positions in **West Conshohocken, PA**.

Job duties include: Responsible for developing, applying, and maintaining quality standards for company products with adherence to both internal and external standards. Develop and execute software test plans. Some positions may allow for telecommuting.

Apply by e-mailing resume to melissa.blandon@oracle.com, referencing 385.16882.
Oracle supports workforce diversity.

FIU | FLORIDA INTERNATIONAL UNIVERSITY

Florida International University is classified by Carnegie as a R1: Doctoral Universities - Highest Research Activity and recognized as a Carnegie engaged university. It is a public research university with colleges and schools that offers 196 bachelor's, master's and doctoral programs in fields such as engineering, computer science, international relations, architecture, law and medicine. As one of South Florida's anchor institutions, FIU contributes almost $9 billion each year to the local economy. FIU is Worlds Ahead in finding solutions to the most challenging problems of our time. FIU emphasizes research as a major component of its mission. FIU has awarded more than 220,000 degrees and enrolls more than 54,000 students in two campuses and three centers including FIU Downtown on Brickell, FIU@I-75, and the Miami Beach Urban Studios. FIU's Medina Aquarius Program houses the Aquarius Reef Base, a unique underwater research facility in the Florida Keys. FIU also supports artistic and cultural engagement through its three museums: Patricia & Phillip Frost Art Museum, the Wolfsonian-FIU, and the Jewish Museum of Florida-FIU. FIU is a member of Conference USA and has more than 400 student-athletes participating in 18 sports. For more information about FIU, visit http://www.fiu.edu/.

FIU's School of Computing and Information Sciences (SCIS) is a rapidly growing program of excellence at Florida International University (FIU). The School has 29 tenure-track faculty members and over 2,000 students, including over 90 Ph.D. students. The School is engaged in on-going and exciting new and expanding programs for research, education and outreach. The School offers B.S., M.S., and Ph.D. degrees in Computer Science, and M.S. degrees in Telecommunications and Networking, Cyber-security, and Information Technology as well as B.S./B.A degrees in Information Technology. NSF ranks FIU 43rd nationwide in externally-funded research expenditures. SCIS has six research centers/clusters with first-class computing and support infrastructure, and enjoys broad and dynamic industry and international partnerships.

We invite applications from exceptionally qualified faculty at all levels with particular emphasis on networking, cyber-security, computer systems or data sciences, and other related areas. Ideal candidates for junior positions should have a record of exceptional research in their early careers and a demonstrated ability to pursue and lead a research program. Candidates for senior positions must have an active and sustainable record of excellence in funded research, publications and professional service as well as demonstrated leadership in collaborative or interdisciplinary research. 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. Applications are encouraged from candidates with highly transformative research programs and seminal ideas that extend the frontiers of computing and networking across other disciplines. A Ph.D. in Computer Science or related disciplines is required.

HOW TO APPLY:

Qualified candidates are encouraged to apply to Job Opening ID **(Job Opening ID # 512441)** at ***facultycareers.fiu.edu*** and attach a cover letter, curriculum vitae, statement of teaching philosophy, research statement, etc as *individual attachments*. Candidates will be requested 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 31, 2016. Review will continue until position is filled.

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.

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

Software Engineer Applications (Req# 9UQT3T) Build Apple's next-gen Employee Sys pltfrm & suite of prods. Travel req'd 25%.

Systems Design Engineer (Req# A3Q3UP) Evaluate baseband & baseband related components.

Technical Support Engineer (Req# 9E3PEV) Provide tech & customer sup, status info, issue resolution, and other forms of assistance re: escalations from Apple's contact centers WW.

Systems Design Engineer (Req# 9FZW4M) Analyze Radio Frequency (RF) system and antenna performance for various wireless technologies. Travel Required: 30%.

Firmware Engineer (Req# 9SSUJX) Des & dev Firmware/SW for embedded accessories.

Software Development Engineer (Req# 9XE24T) Diagnose protocol issues in lab and field testing.

Software Development Engineer (Req# 9M5U4E) Dev & implem user interface & interaction SW on Apple pltfrms for health and fitness apps.

Hardware Development Engineer (Req# 9CYM8L) Des, implmnt, integrate & qual complex comp sys.

Software Engineer Applications (Req# A845LD) Dsgn, dvlp, & build scalable server-side recommendations servcs & cloud based features.

Software Development Engineer (Req# A3Z39P) Prfrm cellular certification testing of iOS devices against 3GPP protocol specifictions.

Development Engineer (Req# AAGU5A) Dtect and rspnd to information security threats.

Software Engineer Applications (Req# A2Q26K) Dsgn, dvlp, test & deploy high volume, high scalable, & high fault tolerant email infrastructure.

ASIC Design Engineer (Req# 9T67DA) Implmnt complex, hi-perf & low power CPU units using gate-level logic design, P&R, & HDL synth.

Product Design Engineer (Req# 9TDQV3) Dev & practice anlytcl methds based on multphyscs modlng & simulatn.

Software Development Engineer (Req# 9F4V73) Dsgn & dvlp Assisted GPS and GLONASS testing for Emergency Location Services.

Software Development Engineer (Req# A362CE) Dsgn and dvlp core security & security automation solutions to create web application firewall systms using machine learning techniques.

Software Engineer Applications (Req# 9EGPRS) Dsgn, rev, dev, dbg & supp SW for bldng, dplymnt & tst of dist sys.

Software Development Engineer (Req# 9YX3Z2) Dsgn & dvlp SW for telecom systms.

System Design Engineer (Req# 9TU6KZ) Dsgn & dvlp instrumentation for calibration & testing of sensors. Travel req. 25%.

Software Engineer Applications (Req# 9R6THF) Dsgn & dvlp SW to spprt building TVMLKit based aplctns for Apple tvOS pltfrm.

Product Design Engineer (Req# 9UQTBB) Dev materials tst methods, conduct tst & analyze tst data for Apple consumer prods incl metallic, polymeric and composite materials. Travel req'd 20%.

Software Development Engineer (Req# 9XQU7B) Tst new sm & lrg scale feats & apps on iOS & OS X across QA teams.

Software Engineer Applications (Req# A2TU5F) Des, dev & mntain dist data storage plats using Apache Hadoop Distributed File System (HDFS Internals).

Software Engineer Applications (Req# A85W4M) Rspnsble for driving the enterprise MFi (Made for iPhone/iPad/iPod) systms to the nxt levl.

Software Engineer Applications (Req# 9X44SJ) Design and develop point of sale (POS) applications for Apple Retail Platform.

Software Development Engineer (Req# A2UTX6) Dev SW tools & perform analyses for undrstndng lrge scale battery life trnds & idntfy areas for imprvmnt.

Software Engineer Applications (Req# A735W8) Provide IT operational supprt & performance monitoring services in the ISO.

Software Engineer Systems (Req# A593FH) Dvlp SW systems to support existing & new iCloud features.

Software Engineer Systems (Req# A9FLQ5) Rsrch, dsgn, implmnt, & tst storage subsys in OS that spprt various storage-rltd techs, including kernel dvce drvrs & user lvl frmwrks. Travel req 15%.

Software Development Engineer (Req# 9XL2U9) Bld a pltfrm to flly utlze probe & 3rd prty data to imprv the Maps exp.

Software Engineer Applications (Req# 9WYVNK) Dsgn, build & tst srvr-drvn mbile e-comrce app. SW on the Apple iOS platfrm.

Software Development Engineer (Req# A57PZ8) Build hi qualty search sys for prdcts incldng Maps search, Siri local search, & othr featres focsd on undstadng wht's intrstng in the wrld arond a givn loction.

Software Development Engineer (Req# 9XD3S6) Build high perf & scal apps. Des, code & test major features.

Software Engineer Applications (Req# 9VG2P3) Arch & dev efficnt, secure, highly avail, scalable, supportable nxt gen SW sys.

Software Engineer Applications (Req# 9KLPYJ) Dsgn, dvlp and deploy data warehouse & analytics solns for multple biz groups at Apple

Software Engineer Applications (Req# 9KHUBT) Des & dev solns for app & data mngmt of intrnl defect, proj, prgrm, & test srvcs.

Mechanical Design Engineer (Req# 9XJVXC) Selct & imprve manuf HW & dev key processes for prodctn, w/ spcialty on thin film optcal coatngs. Travel req: 35%.

Software Development Engineer (Req# 9TWR2F) Dev features for Apple Maps.

Software Development Engineer (Req# A3H2J9) Rspnsbl for validation & spprt of file sys, storage sys drivers, & NAND spprt on iOS & OSX.

Software Engineer Applications (Req# 9ZK3K2) Build SW & sys to manage infrstrctre & apps through automation.

Software Quality Assurance Engineer (Req# 9YU29P) Asst in bring up & valid of iPhone, iPad & Apple Watch, incl Apple made USB accessories.

Software Engineer Applications (Req# A3D437) Architect, prototype & collaborate to deliver scalable/performant runtime systms.

Localization Producer (Software QA Engineer) (Req# 9BUTV9) Work on localization proj for iTunes Store, iCloud, & other Apple entities. Fluency in traditional & simplified Chinese language req'd.

Software Development Engineer (Req# 9SV6WT) Dvlp software tools for test automation.

Software Development Engineer (Req# A8C3F9) Monitor, maintain, and sup infrstrctr for app rev.

Software Quality Assurance Engineer (Req# A362DH) Des & dev fully automated test sols for sys lvl validation of HW and SW.

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

Systems Programmer (Req# 9F4VWV) Support, maintain & operate wrkld & procss automation tools & sys.

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

ASIC Design Engineer (Req# 9PBPSW) Rspnsbl for all aspects of timng incldng wrking w designers for timng changes, help construct/modify flows, timng analysis & timng closure.

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

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

The Computer Science and Engineering Department at **THE OHIO STATE UNIVERSITY :** seeks to fill multiple tenure-track positions at the assistant professor level. We are particularly interested in recruiting in the following areas: cybersecurity, machine learning, distributed systems & cloud computing, and data management. The department is committed to enhancing faculty diversity; women, minorities, and individuals with disabilities are especially encouraged to apply. Some of these positions are partially funded by the university-wide Discovery Themes Initiative, a significant investment in key thematic areas, including the Data Analytics Collaborative which will establish a singular presence in data analytics at Ohio State. The university is also responsive to dual-career families and strongly promotes work-life balance through a suite of institutionalized policies. Applicants should hold or be completing a PhD in computer science & engineering or a closely related field, have a commitment to and demonstrated record of excellence in research, and a commitment to excellence in teaching. To apply, please submit your application via the online database. The link can be found at: https://web.cse.ohio-state.edu/cgi-bin/portal/fsearch/apply.cgi Review of applications will begin in December and will continue until the positions are filled. The Ohio State University is an Equal Opportunity/Affirmative Action Employer.

CLOUDERA, INC. is recruiting for our Palo Alto, CA office: Software Engineer: design & implement large distributed systems that scale well – to petabytes of data & 10s of 1000s of nodes. Mail resume w/job code #37023 to: Cloudera, Attn.: HR, 1001 Page Mill Rd., Bldg. 2, Palo Alto, CA 94304.

SENIOR PROJECT MANAGERS : in Reston, VA sought by project mgmt firm. Qualified candidates will have Master's deg in IT or rltd field & 60 mos as S/ware Eng./Dev. or rltd position. Exp w/ Open System Interconnection (OSI) reference model, TCP/IP Model, Nessus, Wireshark, Firewalls, MS-Access, SQL Server, SQL, Distributed Database System & Reports, architecture diagram & process dsgn using MS Visio, C, C++, PHP, JAVA, XML, PERL, IIS, Windows Server Suite, Linux, VMWare Platform ESX/ESXi, BlueCoat Certified Professional SE, & Sourcefire Certified Professional SE-3D System are req'd. Qualified applicants submit resume: Singhal & Co., Inc., 1952 Isaac Newton Sq. W, Reston, VA 20190.

COMPUTER PROGRAMMER : Create, modify, and test the code, forms, & script that allow computer applications to run. Write computer programs to store, locate, & retrieve specific documents, data, & information.Update & maintain computer programs & software packages. Utilize OOPS concepts involving Java/.NET, PL/SQL relational database applications, Oracle, XML,Webservices, HTML, CSS, Ajax, Design Patterns, Javascript, & Unix, Windows Will work in unanticipated locations. Req. 2 yrs exp. Mail resume to HR - Lorhan Corp. Inc, 400 South Avenue , Suite 9, Middlesex, NJ - 08846.

SR. ADVSR, : Bus Unit Ops (Framingham, MA) Undrstnd cust & partner issues, assgn proper CA Mainframe Solutn. Work w/ mainframe team to map CA tech to customers' bus drivers. Undrstnd & articulate CA's tech vision & strategy. Communicate CA's mainframe capablties & vision for prods to mainframe partnrs. REQS: 5 yrs exp in job &/or a rel occup. Must have exp w/ provdng strategic dirctn & acting as a thought leader for dvlpmnt & exectn of innovative mainframe partnr strategies; Collab w/sales teams to crte, maintn & execute partner bus plan for mainframe prods based on opportunity cust need & strategic direction; Adptng goals fr the bus func based on strategy changes & ensurng changes are planned & exectd in alignment w/ desired bus outcomes; Defing & bldng Route-to-Market (RTM) bus plns & perf msurmnt frmwrks, geo partner prgrms & through-partner mktng strategies; Reprsntg mainframe in industry grps & events & creating intrnl & extrnl awareness; Provdng pipeline forecasts & bus health checks to executive sales mgmt & partners; Bldg & mntng relshps w/ partners at all levels & dvlpng/imprvng tools, processes & procedures to incr partner efficiency; Nat'l & internat'l travel req approx 35% of the time.40 hours/week; M-F; 8:30 am-5:30 pm; Send resume to: Althea Wilson, CA Technologies, 201 N Franklin Street, Suite 2200, Tampa FL US 33602 Refer to Requisition # 139862.

COMPUTER PROGRAMMER: Create, modify, and test the code, forms, & script that allow computer applications to run. Write computer programs to store, locate, & retrieve specific documents, data, & information. Update & maintain computer programs & software packages. Utilize C#.NET, WCF, MVC, HTML, ASP.NET, VB.Net, SQL Server, T-SQL, PL/SQL, Oracle, AJAX, Java Script, Web Services, SSAS, SSIS, SSRS, Dev Express. Will work in unanticipated locations. Req. 2 years experience, of which 2 years experience required in C#.NET, WCF, MVC, HTML, ASP.NET, VB.Net, SQL Server, T-SQL, PL/SQL, Oracle, AJAX, Java Script, Web Services, SSAS, SSIS, SSRS, Dev Express. Send resume to Quest IT Solutions Inc. 1449 Hwy 6 S, Suite 380, Sugar Land, TX 77478.

VLOCITY : seeks Sr. UI Engineer in San Francisco, CA to dsgn & develop industry specific apps. Ref Job ID: 9WRQ6K & send res to T. Dilley at hiring@vlocity.com.

TECHNOLOGY

Intuit Inc.

has openings for the following positions in **Mountain View, California:**

Senior Technical Data Analysts (Job code: I-2788): Create new and forward-thinking automated Small Medium Business (SMB) credit risk model, leveraging QuickBooks data and relevant third party data.

Openings in **San Diego, California**:

Software Engineers (Job code: I-972): Apply software development practices to design, implement, and support individual software projects. **Staff Systems Engineers (Job code: I-1360):** Design & develop new software systems, services, features & enhancements, & maintain existing software products.

Openings in **Plano, Texas**:

Staff Software Engineers (Job code: I-459): Apply master level software engineering and industry best practices to design, implement, and support software products and services.

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.

COMPUTER PROFESSIONALS. Central NJ IT Consulting Company requires candidates for following position at their primary Princeton, NJ location, Java Developer: to Design, develop & implement business software apps using Java & J2EE, MQ, struts, Jira and CVS,GIT. Candidates must have 1-2 yrs of mandatory exp in related -field. All positions require: MS in CS/Engineering/Comp Applications/Business or related. BS degree + 5yrs exp can be substituted for the MS requirement any combination of foreign edu + related exp equivalent to a US Masters, or any combination of foreign edu +related exp equivalent to a BS Degree will be accepted. Travel to several unanticipated locations all over US & might involve relocation consistent w/client needs & State & Local needs. Mail your resume to: Kellton Tech Inc., Attn: HR, 3 Independence Way, STE 209, Princeton NJ 08540.

SAP PROJECT MANAGER. Princeton, NJ based IT Consulting company requires SAP Project Manager: Oversee the successful planning & implementation of multiple projects/releases of SAP, using standard project management tools & procedures. Will be responsible to initiate, plan, execute, monitor & control and close, and the supporting processes. Position requires at least 12 months of applying Accounting principles to SAP Enterprise Resource Planning (ERP)+MS in Computer Science/Accounting Any combination of foreign edu equivalent to a US Masters will be accepted. Position involves travel to several unanticipated locations all over US & might involve relocation. Mail your resume to: Global HR, KelltonTech Inc., 3 Independence way, STE 209, Princeton NJ 08540.

SOLUTIONS ARCHITECT. Position available in Enfield, CT. Construct technology solutions for business needs and manage technical risk of delivery using Sparx Systems Enterprise Architect and Oracle databases. Assist vin the development of the technology architecture plan and the integration of corporate and business area architectures. Establish standards, best practices, and design/implementation patterns for new and existing technologies. Willingness to travel to company headquarters in Springfield, MA as needed required. Apply: L. Sawtelle, Massachusetts Mutual Life Insurance Company, 1295 State Street, Springfield, MA 01111; Please Reference Job ID: 708203100.

ATHOC, INC. has the following job opportunities in San Mateo, CA: **Sr.Systems Eng. (Req#AH2)** Perform test management and defect tracking. **Principal Build & Release Eng. (Req#AH1)** Desgn & build software products & sols. **Principal Sys. Eng. (Req#AH3)** Develop new APIs, web services & sw components. Refer to Req# & mail resume to BlackBerry Corp., P.O. Box 141394, Irving, TX 75014.

SENIOR SOFTWARE DEVELOPER. Valassis Communications, Inc. has an opening for the position of Senior Software Developer in Windsor, CT to analyze, write, test, & deploy software applications & research, design, document & modify software specifications. To apply mail resume to Valassis, Attn: Patty [CT-11307.7], 19975 Victor Parkway, Livonia, MI 48152.

SOFTWARE

CarRentals.com Inc.

currently has openings for

SOFTWARE ENGINEERS

(JOB ID: 728.1615)

in our **San Francisco, CA** office

(various/levels/types)

To design, implement, and debug software for computers including algorithms and data structures.

Send your resume to:
CarRentals.com/Expedia Recruiting,
333 108th Avenue NE, Bellevue, WA
98004. Must reference JOB ID# .

Oracle America, Inc.

has openings for

HARDWARE DEVELOPER

positions in **Burlington, MA**.

Job duties include: Evaluate reliability of materials, properties and techniques used in production; plan, design and develop electronic parts, components, integrated circuitry, mechanical systems, equipment and packaging, optical systems and/or DSP systems.

Apply by e-mailing resume to
anthony.lucca@oracle.com ,
referencing 385.18662.
Oracle supports workforce diversity.

SOFTWARE

Oracle America, Inc.

has openings for

SOFTWARE DEVELOPER

positions in **Washington, D.C**.

Job duties include: Design, develop, troubleshoot and/or test/QA software, specifically language processing modules for Arabic-script languages and other foreign languages. Fluency in spoken and written Arabic required. May telecommute from home.

Apply by e-mailing resume to
george.krupka@oracle.com,
referencing 385.17394.
Oracle supports workforce diversity.

SOUTH UNIVERSITY OF SCIENCE & TECHNOLOGY OF CHINA. Professor/Associate Professor/Assistant Professorship in Computer Science and Engineering. The University Established in 2012, the Southern University of Science and Technology (SUSTech) is a public institution funded by the municipal of Shenzhen, a special economic zone city in China. Shenzhen is a major city located in Southern China, situated immediately north of Hong Kong Special Administrative Region. As one of China's major gateways to the world, Shenzhen has been the country's fast-growing city in the past two decades. The city is the high-tech and manufacturing hub of southern China, home to the world's third-busiest container port, and the fourth-busiest airport on the Chinese mainland. A picturesque coastal city, Shenzhen is also a popular tourist destination and was named one of the world's 31 must-see tourist destinations in 2010 by The New York Times. By the end of 2013, there were over 10 million permanent residents in the city. The Southern University of Science and Technology (SUSTech) is a pioneer in higher education reform in China. The mission of the University is to become a globally recognized institution which emphasizes academic excellence and promotes innovation, creativity and entrepreneurship. The teaching language at SUSTech is bilingual, English and Mandarin Chinese. Set on five hundred acres of wooded landscape in the picturesque Nanshan (South Mountain) area, the new campus offers an ideal environment suitable for learning and research. The new campus occupies more than 1,940,000 square meters, and the total construction area will be approximately 630,000 square meters. Call for Applications The Southern University of Science and Technology (SUSTech) invites applications for faculty positions in the Department of Computer Science and Engineering. It is seeking to appoint up to eight teaching-track positions at all ranks. Candidates with teaching interests in all core fields of Computer Science and Engineering will be considered. These positions are full-time and tenure track posts although they are teaching focused. SUSTech adopts the tenure track system, which offers the recruited faculty members a clearly defined career path. Candidates should have demonstrated excellence in and commitment to teaching and research. A doctoral degree in Computer Science or a closely related discipline is required at the time of appointment. Candidates for senior positions must have an established track record of outstanding teaching at universities and previous research experience. Main Teaching Needs We are particularly interested in candidates who have expertise in teaching some of the following courses: Introduction to Computer Science, Foundation of Computer Programming, Java Programming, Object-Oriented Design and Programming, Databases, Data Structures, Computer Networks, Operating Systems, Compiler Design & Principles, Software Engineering, Computer Systems and Architectures, Computer Security, Intelligent Robotics, Artificial Intelligence, Machine Learning, Computer Graphics. Candidates may be requested to teach other courses than those listed above. Terms & Applications SUSTech offers internationally competitive salaries and fringe benefits including medical insurance, retirement and housing subsidies, which are among the best in China and competitive internationally. The salary and rank will commensurate with qualifications and experiences. More information can be found at http://talent.sustc.edu.cn/en. All candidates must be able to teach and communicate well in English. Up to eight positions are available. This call for applications is valid until all positions are filled. To apply, please provide a cover letter identifying the primary areas of teaching, previous teaching experiences (including subjects, number of students, student feedback, etc), curriculum vitae, teaching statements, previous research achievements, and arrange for at least three recommendation letters, all forward to cshire@sustc.edu.cn.

University of Illinois at Urbana-Champaign
Positions in Computing

The Department of Electrical and Computer Engineering (ECE) at the University of Illinois at Urbana-Champaign invites applications for faculty positions at all areas and levels in computing, broadly defined, with particular emphasis on reliable and secure computing; networked and distributed computing; high-performance, energy-efficient, and scientific computing; data center and storage systems; data science, machine learning and its applications; complex data analysis and decision science; bio-inspired computing; computational genomics; and health informatics, among other areas. Applications are encouraged from candidates whose research programs specialize in core as well as interdisciplinary areas of electrical and computer engineering. From the transistor and the first computer implementation based on von Neumann's architecture to the Blue Waters petascale computer—the fastest computer on any university campus, ECE Illinois faculty have always been at the forefront of computing research and innovation. The department is engaged in exciting new and expanding programs for research, education, and professional development, with strong ties to industry. The ECE Department has recently settled into its new 235,000 sq. ft. net-zero energy design building, which is a major campus addition with maximum space and minimal carbon footprint.

Qualified senior candidates may also be considered for tenured full Professor positions as part of the Grainger Engineering Breakthroughs Initiative (http://graingerinitiative.engineering.illinois.edu), which is backed by a $100-million gift from the Grainger Foundation.

Please visit http://jobs.illinois.edu to view the complete position announcement and application instructions. Full consideration will be given to applications received by December 15, 2017, but applications will continue to be accepted until all positions are filled.

Illinois is an EEO Employer/Vet/Disabled www.inclusiveillinois.illinois.edu.

The University of Illinois conducts criminal background checks on all job candidates upon acceptance of a contingent offer.

Join **MICHIGAN STATE UNIVERSITY'S** Global Impact Initiative, designed to address the grand challenges through the creation of over 100 new faculty positions in some of the most promising and exciting fields of research. We welcome applicants from diverse backgrounds. MSU offers an inclusive and collaborative work environment. To learn more visit research.msu.edu/global-impact. **Computer Science and Engineering Faculty Position in Biometrics.** The Department of Computer Science and Engineering (CSE) at Michigan State University (MSU) invites applications for a faculty position in the area of biometrics. While the position is primarily for a junior faculty, candidates at other ranks may be considered. The successful candidate will be expected to develop an externally-funded interdisciplinary research program of international prominence that includes fundamental research, publications in high-impact journals and conferences, and training graduate students. Multidisciplinary research is strongly encouraged and is being actively pursued by the faculty members at MSU. Leadership is expected in the development of innovative educational programs that provide state-of-the-art knowledge to both undergraduate and graduate students. Candidates should have a Ph.D. in Computer Science or a closely related field, with demonstrated evidence of research accomplishments, teaching skills, and ability to work effectively with other researchers within the Department and colleagues on campus. Appointments will start in August 2017. MSU enjoys a park-like campus with outlying research facilities and natural areas. The campus is in the city of East Lansing and adjacent to the capital city of Lansing. The Lansing metropolitan area has a diverse population of approximately 450,000 residents. Local communities have excellent school systems and place a high value on education. Michigan State University is pro-active in exploring opportunities for employment for dual career couples, both inside and outside the University: http://miwin.msu.edu/. Information about work and life at MSU and the College of Engineering can be found at http://www.egr.msu.edu/WE. Applicants should submit a cover letter, curriculum vitae, the names of at least three references, and statements of their research and teaching interests through http://jobs.msu.edu and refer to posting #4074. Applications will be reviewed on a continuing basis until the position is filled . Review of applications will begin on January 2, 2017. For questions about this position, contact the search committee chair at search@cse.msu.edu. Additional information about the university, college and department is available at: CSE Department - http://www.cse.msu.edu. College of Engineering - http://www.egr.msu.edu/ MSU – http://www.msu.edu/ Michigan State University has been advancing the common good with uncommon will for more than 160 years. A member of the Association of American Universities, MSU is a research-intensive institution with 17 degree-granting colleges. MSU is an affirmative-action, equal opportunity employer. MSU is committed to achieving excellence through a diverse workforce and inclusive culture that encourages all people to reach their full potential. The University actively encourages applications and/or nominations of women, persons of color, veterans and persons with disabilities.

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 multiple faculty positions at the assistant or associate professor levels, with research interests in all areas of computer science. Candidates must have a Ph.D. in computer science or related field at the time of appointment and a rank-appropriate record of scholarship and collaboration in computing research, broadly defined. Candidates should give evidence of sensitivity to issues of diversity and inclusion, and will be expected to teach graduate and undergraduate courses, mentor graduate students, and develop a high quality research program.

The department has 39 tenured/tenure-track faculty, over 700 majors (17% women), 68 MS students, and 181 PhD students. Departmental research expenditures last year were $14 million. Candidates will have the opportunity to work with a wide range of productive research groups including computational biology and bioinformatics, computational science, computer science education, cybersecurity, data analytics and machine learning, human-computer interaction, software engineering and computer systems. Multidisciplinary research centers led by CS faculty include Center for Human-Computer Interaction (**www.hci.vt.edu**), Discovery Analytics Center (**dac.cs.vt.edu**), stack@cs (**stack.cs.vt.edu**), Synergistic Environments for Experimental Computing (**research.cs.vt.edu/seec**), Biocomplexity Institute (**www.bi.vt.edu**), and Institute for Creativity, Arts, and Technology (**icat.vt.edu**). The department also plays a central role in several university-wide research and teaching initiatives (**provost.vt.edu/destination-areas**), including Data Analytics and Decision Sciences, Integrated Security, Intelligent Infrastructure for Human-Centered Communities, and Creative Technologies and Experiences. The department is in the College of Engineering, whose undergraduate program ranks 15th and graduate program ranks 21st among U.S. engineering schools (*USN&WR,* 2015).

These faculty positions are located at the main campus in Blacksburg, VA, in a region that is consistently ranked among the country's best places to live. The positions require occasional travel to professional meetings. Virginia Tech is committed to building a culturally diverse community and strongly encourages women and minorities to apply. Candidates must pass a criminal background check prior to employment.

Submit applications to **jobs.vt.edu**, posting **#TR0160116**. Screening begins November 20, 2016 and continues until positions are filled. Direct inquiries to Dr. Ali Butt, Search Committee Chair, **butta@cs.vt.edu**.

Virginia Tech is an Equal Opportunity/Affirmative Action Employer

Oracle America, Inc.

has openings for

CONSULTING PROJECT PRINCIPAL CONSULTANT

positions in **Redwood Shores, CA**.

Job duties include: Responsible for ensuring that a quality, integrated software solution is delivered in a timely manner, at budget, and to client satisfaction. Travel to work on projects at various, unanticipated sites throughout the United States required. May telecommute from home.

Apply by e-mailing resume to mark.griffith@oracle.com, referencing 385.13995.

Oracle supports workforce diversity.

CLOUDERA, INC. is recruiting for our Palo Alto, CA office: Software Engineer: design & implement engineering systems that scale well – to petabytes of data & 1000s of nodes. Mail resume w/job code #37801 to: Cloudera, Attn.: HR, 1001 Page Mill Rd., Bldg. 2, Palo Alto, CA 94304.

PROGRAMMER ANALYST, United Shore Financial Services, LLC, Troy, MI. Desg &dvlp apps using Asp.net, C# &Object Oriented Programming. Dvlp &implement improvements to tools to facilitate delivery of loan products to brokers, correspondents, financial institutions, &mortgagors, using Asp.Net, .NET framework, C#, HTML, Scheme, CSS, Ajax, &JavaScript. Perform Web dvlpmt in .Net using MS Visual Studio &SQL Server. Dvlp &code features to eliminate inefficient manual processes &optimize internal IT tools, such as loan origination systems, online broker portals, &mortgage payoff calculators. Use SVN &TFS Source controls. Interact with customer to understand reqmts. Use MVC Model &expertise in soft such as .NET, C#.Net, Vb.net, Visual basic, Jquery &Scheme. Use TFS extensively for version control of source code along with maintenance of builds &relevant documents. Bachelor, Computer Science, Engineering or related. 12 Months, Programmer Analyst dvlpg &implementing improvements to tools to facilitate delivery of loan products to brokers, correspondents, financial institutions, &mortgagors, using Asp.Net, .NET framework, C#, HTML, Scheme, CSS, Ajax, &JavaScript. Mail resume to Michelle Salvatore, 1414 E Maple Rd, Troy, MI 48083, ref#36925.

IEEE Intelligent Systems

THE #1 ARTIFICIAL INTELLIGENCE MAGAZINE!

IEEE Intelligent Systems delivers the latest peer-reviewed research on all aspects of artificial intelligence, focusing on practical, fielded applications. Contributors include leading experts in

• Intelligent Agents • The Semantic Web
• Natural Language Processing
• Robotics • Machine Learning

Visit us on the Web at
www.computer.org/intelligent

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

Facebook, Inc.

currently has the following openings in **Menlo Park, CA (multiple openings/various levels)**:

Software Engineer (SWEB1016J) Create web &/or mobile applications that reach over one billion people, & build high volume servers to support our content. Bachelor's degree required. Exp. may be required depending on level/type. **Software Engineer (SWEM1016J)** Create web &/or mobile applications that reach over one billion people, & build high-volume servers to support our content, utilizing graduate level knowledge. Master's degree required. Exp. may be required depending on level/type. **Network Engineering Manager (2220J)** Manage engineers working with our Datacenter, Backbone, Hardware and Software networking teams to build, scale, deploy, and support our global network infrastructure. **Lead Partner Engineer, LatAm (3062J)** Responsible for all integration projects and Partner Engineering activities in Latin America region, including: planning, execution, and quality of the integration projects. Position requires occasional travel.

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

Focus on Your Job Search

IEEE Computer Society Jobs helps you easily find a new job in IT, software development, computer engineering, research, programming, architecture, cloud computing, consulting, databases, and many other computer-related areas.

New feature: Find jobs recommending or requiring the IEEE CS CSDA or CSDP certifications!

Visit www.computer.org/jobs to search technical job openings, plus internships, from employers worldwide.

http://www.computer.org/jobs

IEEE computer society | JOBS

The IEEE Computer Society is a partner in the AIP Career Network, a collection of online job sites for scientists, engineers, and computing professionals. Other partners include Physics Today, the American Association of Physicists in Medicine (AAPM), American Association of Physics Teachers (AAPT), American Physical Society (APS), AVS Science and Technology, and the Society of Physics Students (SPS) and Sigma Pi Sigma.

IEEE Software offers pioneering ideas, expert analyses, and thoughtful insights for software professionals who need to keep up with rapid technology change. It's the authority on translating software theory into practice.

www.computer.org/software/**subscribe**

SUBSCRIBE TODAY

TECHNOLOGY

LinkedIn Corp.

has openings in our **Sunnyvale, CA** location for:

Software Engineer (All Levels/Types) (SWE1016SV) Design, develop & integrate cutting-edge software technologies; **Software Engineering Manager (6597.1696)** Architect, design, develop, & support the most visible Internet-scale products & infrastructure at LinkedIn; **Test Engineer (6597.1405)** Design, develop & integrate cutting-edge software technologies; **Senior Test Engineer (6597.1436)** Design & develop advanced test suites & necessary automation frameworks using object-oriented methodologies (OOM); **Database Engineer (6597.1422)** Review & deploy code changes, monitor Back-Ups, troubleshoot performance issues, & support the Development team.

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

Software Engineer (All Levels/Types) (SWE1016SF) Design, develop & integrate cutting-edge software technologies.

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

Move Your Career Forward

IEEE Computer Society Membership

Explore These Security Resources

Advance Your Career

Security Certificate of Achievement

The Computer Society now offers a Certificate of Achievement in security. Take advantage of this professional development opportunity to expand your expertise in this growing field and further advance your career by successfully completing these four courses:

- Secure Software Coding
- Secure Software Design
- Managing Secure Software Development
- Foundations of Software Security

Build Your Knowledge

IEEE Security & Privacy

The information security industry turns to *S&P* for both practical and research viewpoints from experts in the field through case studies, tutorials, columns, and in-depth interviews and podcasts.

IEEE Cybersecurity Initiative

Launched in 2014, the initiative helps students, educators, and practitioners improve their understanding of cybersecurity through IEEE's ongoing involvement in and elevation of the cybersecurity field. In 2015, the Initiative launched the IEEE Center for Secure Design (CSD), which shifted focus from finding bugs to identifying common design flaws in the hope that software architects could learn from others' mistakes. Currently, the initiative tackles security-related challenges through expanding computer security education, continuing development of a security-focused building code for critical software, recruiting underrepresented groups in the field, and more.

FOR DIRECT LINKS TO THESE RESOURCES, VISIT

www.computer.org/edge-resources

IEEE computer society

CELEBRATING 70 YEARS

IEEE computer society

For Early Pricing, Register Now!

TechIgnite

A ROCK STAR TECHNOLOGY EVENT

The Truth Behind Technology

March 21–22, 2017 | Burlingame, CA

FUTURE TECHNOLOGIES, TRENDS, TECH GURUS

Learn the latest trends, best practices and hear case studies from thirty-three of today's top technology gurus as they dispel the myths about disruptive technologies and demonstrate actionable problem solving techniques you can apply today.

Featuring

Steve Wozniak & Grady Booch

Also, Google's Head of Quantum Computing, CTO Homeland Security, and Uber's Machine Learning & AI Guru

www.computer.org/**TechIgnite**