06.15

Computer

WEARABLE COMPUTING

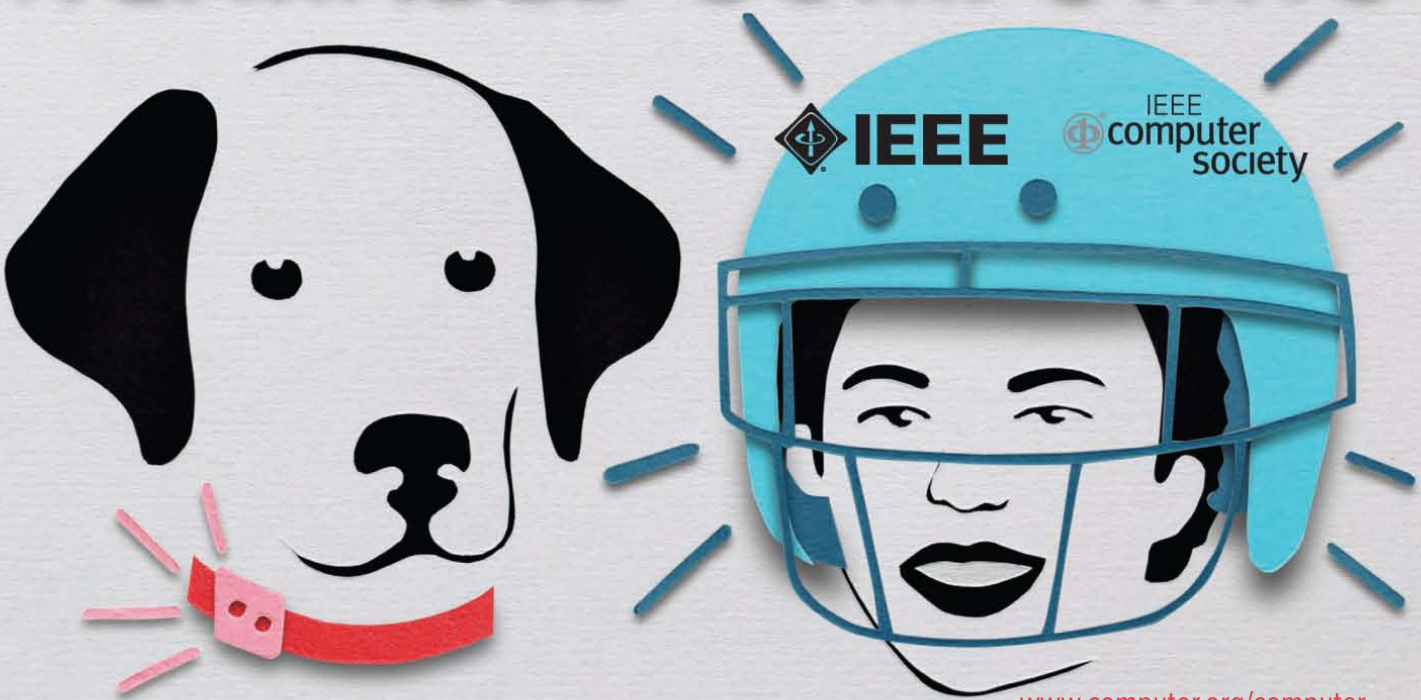

www.computer.org/computer

Computer

JUNE 2015

FEATURES

JUNE 2015

CONTENTS

ABOUT THIS ISSUE
WEARABLE COMPUTING

As computers come onto the body they compensate for what we lack—be it expanded, enhanced, or even just basic abilities.

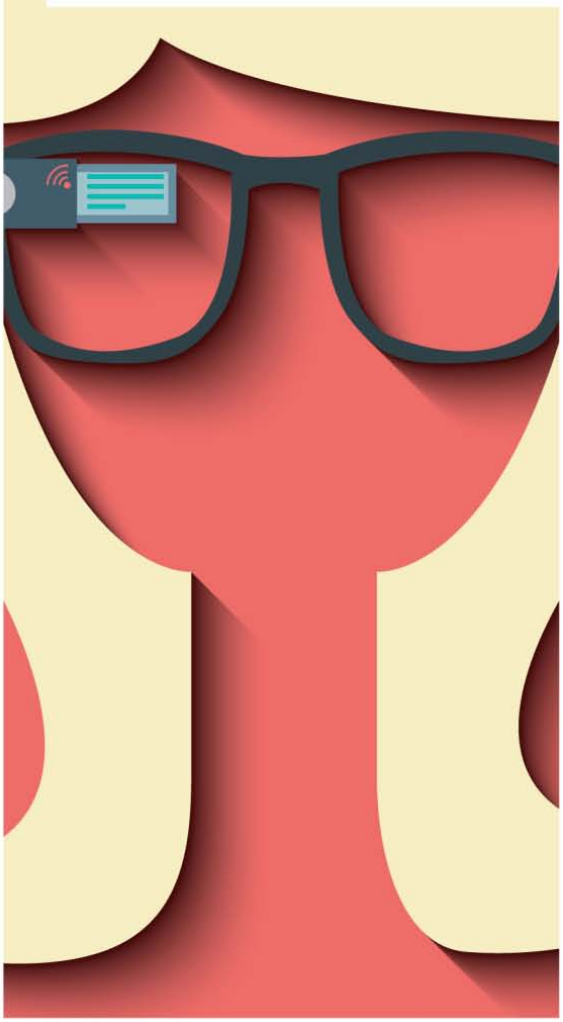

See **www.computer.org/computer-multimedia** for multimedia content related to the features in this issue

MORE FEATURES

RESEARCH FEATURE

COLUMNS

Membership News

Departments

For more information on computing topics, visit the Computer Society Digital Library at **www.computer.org/csdl.**

Computer

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

ELSEWHERE IN THE CS

EDITOR **LEE GARBER**
lgarber@computer.org

Computer Highlights Society Magazines

The IEEE Computer Society's lineup of 13 peer-reviewed technical magazines covers cutting-edge topics in computing, including scientific applications, Internet computing, machine intelligence, pervasive computing, security and privacy, digital graphics, cloud computing, and computer history. Here, we highlight recent issues of other Computer Society magazines.

IEEE Software

Software companies frequently find that they must create multiple versions of products for their different customers. This requires effective **variability management**. However, the number of changes necessary for each version and the dependencies between variation points can skyrocket and create significant problems. Facing these challenges is the topic of a special section in *IEEE Software*'s May/June 2015 issue titled "Trends in Systems and Software Variability."

IEEE Internet Computing

Physical-cyber-social (PCS) computing is the subject of *IEEE Internet Computing*'s May/June 2015 special section. PCS computing involves the integration of data and knowledge from the physical, cyber, and social worlds for analyses that produce contextually relevant abstractions to both humans and the applications that serve them. This approach is becoming increasingly important as mobile technology, social media, and the Internet of Things become more popular and as they connect physical and social elements with traditional cyber entities.

Computing in Science & Engineering

The US **National Energy Research Scientific Computing Center (NERSC)**, the state-of-the-art facility that serves government, industry, and academic users, celebrated its 40th anniversary in 2014. *CiSE*'s May/June 2015 special issue includes articles that document NERSC's history and discusses some of the facility's major contributions, including research on a new supercomputer architecture.

IEEE Security & Privacy

Policymakers are currently focusing on **informed consent** as a way to defend privacy. Many countries have laws requiring this. However, behavioral studies cast doubt on the approach's effectiveness, as many people tend to agree with almost any request they see on their screens. A related technique that both protects and empowers individuals would improve privacy. This is addressed in "Informed Consent: We Can Do Better to Defend Privacy," in *IEEE S&P*'s March/April 2015 issue.

IEEE Intelligent Systems

Researchers have tried to use footsteps as a basis for biometric systems. However, constructing an identity-verification system based on footsteps remains challenging. They are not only part of our physical makeup but can also reflect how we're feeling, the surface on which we're walking, and the shoes we're wearing. "Footstep-Identification System Based on Walking Interval," which appears in *IEEE Intelligent Systems*' March/April 2015 issue, describes a **novel footstep-identification system**.

IT Professional

In "Could What Happened to Sony Happen to Us?" from *IT Pro*'s March/April 2015 issue, the Open Web Application Security Project's Walter Houser says that until there are **cybersecurity standards of practice**, Internet commerce will experience problems like 2014's Sony exploit. We continue to put a growing percentage of our assets online, providing vandals, criminals, and terrorists with lucrative targets. As details of the Sony incident become public, we're seeing important lessons for online enterprises and notable implications for the discipline of security engineering.

0018-9162/15/$31.00 © 2015 IEEE

IEEE pervasive COMPUTING

Pervasive computing has yet to be widely exploited in physical spaces. According to the guest editors of the special issue on smart spaces in *IEEE Pervasive Computing*'s April–June 2015 issue, pervasive computing in physical spaces could provide many business and social opportunities that are currently being missed. They say the advent of the Internet of Things and the rise of consumer-based wearable computers have made it appropriate to look at this issue and revisit the notion of a smart space. The articles in this special issue consider new technologies and approaches for developing **pervasive smart spaces**.

IEEE micro

IEEE Micro's March/April 2015 special issue contains articles about a set of processors presented at 2014's **Hot Chips** 26 conference. Over the years, Hot Chips has become a leading venue for diverse companies—including industry leaders with a broad set of products and newer companies with a smaller portfolio—to release technical details about their latest chips. Hot Chips has consistently provided an early look at trends in the processor industry.

IEEE Computer Graphics and Applications

Visual analytics (VA) combines the strengths of human and machine intelligence to enable the discovery of interesting patterns in challenging datasets. Historically, most attention has been given to developing machine capabilities. But it's also essential to develop the **abilities of the visual analysts** themselves, especially at the beginning of their careers. For the past several years, the University of British Columbia—with support from Boeing—has experimented with ways to prepare undergraduates for VA careers. This has led to creation of a third-year undergraduate course. Details are provided in "Preparing Undergraduates for Visual Analytics," which appears in *IEEE CG&A*'s March/April 2015 issue.

IEEE MultiMedia

Traditionally, multimedia has been associated with video, audio, or other media content. However, in recent years, multimedia's definition has expanded to include media context such as video motion features or timestamps, and media connectivity features such as the degree of friendship between two social-media participants. This is the basis of *IEEE MultiMedia*'s April–June 2015 special issue titled "**Multimedia Goes beyond Content**."

IEEE Annals of the History of Computing

In "The Production and Interpretation of Arpanet Maps" from *IEEE Annals*' January–March 2015 issue, Bradley Fidler and Morgan Currie of the University of California, Los Angeles, explore a 20-year series of network topology maps produced in the 1960s, 1970s, and 1980s by the firm Bolt Beranek and Newman (now BBN Technologies). These maps were a way to represent **Arpanet, a precursor of today's Internet**, to engineers.

IEEE CLOUD COMPUTING

Recent studies by Cisco Systems and IBM show that we generate 2.5 quintillion bytes of data per day, and they predict that this rate will explode to 40 yottabytes daily by 2020. Much of this information is and will be generated from Internet of Things devices and sensors, and will be stored in cloud-accessible datacenters. The article "Processing Distributed Internet of Things Data in Clouds," from *IEEE Cloud Computing*'s January/February 2015 issue, discusses the capabilities and limitations of big data technologies for collecting and analyzing **distributed big datasets across multiple datacenters**.

Circulation: *Computer* (ISSN 0018-9162) 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.EEE Computer Society membership includes a subscription to Computer magazine.

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

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

SPOTLIGHT ON TRANSACTIONS

Distributed Optimal Channel Access in Cognitive Radio Networks

Prasant Mohapatra, University of California, Davis

This installment of Computer's *series highlighting the work published in IEEE Computer Society journals comes from* IEEE Transactions on Mobile Computing.

Cognitive radio networks (CRNs) let unlicensed or secondary users (SUs) access licensed spectra, improving wireless resource utilization and addressing the problem of spectrum shortage. SUs are allowed spectrum access only when they don't cause unacceptable levels of interference to licensed or primary users (PUs). Although throughput-optimal algorithms based on the well-known maximal-weight scheduling algorithm exist for CRNs, they require central processing of network-wide SU information. Developing a distributed implementation that can fully leverage the spectrum opportunities for SUs has so far remained elusive.

Because distributed algorithms eliminate the need for network-wide state information collection and avoid costly algorithm executions, researchers Shuang Li, Eylem Ekici, and Ness Shroff developed a fully distributed channel-access solution that uses only local queue-length information.

In their article "Throughput-Optimal Queue Length–Based CSMA/CA Algorithm for Cognitive Radio Networks" (*IEEE Trans. Mobile Computing*, vol. 14, no. 5, 2015, pp. 1098—1108), the authors identify a new class of *queue length–based carrier sense multiple access* (CSMA) algorithms as a promising approach to solving the distributed and optimal channel-access problems in CRNs. This new class of random-access algorithms combines channel-sensing results with locally available queue-length information to determine channel-access probabilities, achieving full capacity in ad hoc wireless networks in a distributed manner. However, the throughput optimality is realized for nonfading and always-available channel cases, a condition not met in CRNs due to PU activity.

Li, Ekici, and Shroff overcome this problem with a novel algorithm based on a new system-state representation that tracks the channel-availability information. This new definition of the system state, along with the differentiated treatment of SUs inside and outside the PU interference range, leads to a distributed algorithm design that achieves provable throughput optimality for SUs without disrupting PU operation. Significant capacity gains realized by the algorithm vis-à-vis other queue length–based CSMA algorithms highlight the efficacy of the authors' design approach.

The researchers' results show that it's possible to garner the full potential of CRNs without relying on complex centralized algorithms. This work is the first step toward highly efficient, decentralized deployment of CRNs and will lead to the development of other distributed optimal algorithms for networks with time-varying channel conditions.

PRASANT MOHAPATRA is a professor in the Department of Computer Science at the University of California, Davis, and the editor in chief of *IEEE Transactions on Mobile Computing*. Contact him at pmohapatra@ucdavis.edu.

0018-9162/15/$31.00 © 2015 IEEE

COMPUTING CONVERSATIONS

EDITOR CHARLES SEVERANCE
University of Michigan; csev@umich.edu

VIDEO

AUDIO

Roy T. Fielding: Understanding the REST Style

Charles Severance, University of Michigan

Roy T. Fielding reminisces about his PhD dissertation, which defined the Representational State Transfer architectural style.

Scratch the surface of most modern networked applications, and you'll see the various application parts exchanging data using an approach called Representational State Transfer (REST). Developers can learn about the origins of REST by reading Roy T. Fielding's PhD thesis for the University of California, Irvine (UCI), published in 2000 (www.ics.uci.edu/~fielding/pubs/dissertation/top.htm). This is required reading for those wanting to be REST experts. I spoke with Roy about how the REST concepts were developed and how they came to be described in his thesis. You can see the entire interview with Roy at www.computer.org/computingconversations.

See **www.computer.org/computer-multimedia** for multimedia content related to this article.

Roy was an early participant in Web protocol and software development as it emerged in the early 1990s. The first implementations of Web browsers and servers were created by Tim Berners-Lee at CERN using a NeXt computer, and they were quickly ported to many different computers and operating systems. For each new use of the Web, ideas about how its protocol could be improved to better support the new idea often followed. To change the protocol, you had to send an email to a list that included virtually all of the Web developers in the world:

> *HTTP and HTML were built informally, primarily using mailing lists as the coordination mechanism. [Developers from all over the world] would talk about a new feature. Frequently, we would come up with an idea in one time zone and someone would implement it in another time zone; by the next morning, you would know what worked and what didn't work with that feature. So it was very free-form and very fast. At the time we were all computer scientists or researchers who were playing around, doing research, or providing our own resources on the Web.*

COMPUTING CONVERSATIONS

By 1994, the impetus for Web improvement moved from computer scientists and Web developers who were exploring a new space to an increasingly critical infrastructure for growing Web-based businesses:

> *As companies became involved, they wanted to find ways to use the Web corporately as one of their platforms, so it needed to be more "businesslike." One of the ways to make things more suitable for business is to create common standards for everyone to adhere to, rather than to adopt things as you go along.*

Berners-Lee and others formed the World Wide Web Consortium (W3C) at MIT to develop the standards that guide the Web today. Roy was invited to work on the standards for URLs and HTML, and eventually became the lead editor for HTTP:

> *As one of the developers of the Web protocol library libwww-perl, I was asked to help work on the standards. Because I was a graduate student at UCI, I had all the freedom in the world. I had finished all my classwork and started working on my dissertation. That gave me the ability to write for the Web in addition to the programming I was still doing for my thesis.*

As the Web grew rapidly in the mid-1990s and powerful businesses started embracing the Web, there were even more suggestions from the new stakeholders as to how to improve the protocols:

> *I was faced with the dilemma of having many competing interests working toward making the Web what they thought would be a better place. I needed to differentiate between ideas that were actually better for the Web and those that went back to some older version of architecture that didn't make any sense on the Internet.*

Roy coined the term REST to capture the idea that the elements of networked applications were exchanging their states using protocols across the network.

Instead of debating each new idea separately, Roy developed a model that captured the principles of how distributed applications should interact with one another:

> *I called it the "HTTP Object Model." At the time, object models were the "in" thing, so I called it an object model even though it had nothing to do with objects. It was a model of how I expected Web applications to behave. If someone offered a feature or described something they thought was wrong with the Web, I would use the model to show how HTTP worked and how the new feature might hurt or help. The model gave me some intellectual leverage to affect how the HTTP standard evolved.*

Roy's PhD research on software architecture helped him better understand and explain his model:

> *It wasn't until after I had done the literature search for software architecture that I figured out the right way to describe the HTTP model I had developed. I read a paper by Dewayne Perry and Alexander Wolf ["Foundations for the Study of Software Architecture," ACM SIGSOFT Software Eng. Notes, vol. 17, no. 4, 1992, pp. 40–52]—it was the only software architecture paper I could find that described architecture in terms of not only the components and connectors of typical architecture diagrams, but also the data that's processed through the system.*

During the late 1990s, Roy was oscillating between the high-stakes world of Internet standards development and the more academic approach for his PhD thesis:

> *One of the great benefits I had was the freedom to pursue these different areas. I was working on a team that was doing research on global software engineering environments, so I was trying to use the Web as a platform for software engineering. My research project was essentially what GitHub is today. As part of that, I could do all this other work that was related to it. It was one of the nice things about general research funding.*

But Roy also needed to finish his thesis and graduate:

> *I come from an academic background. My father is a professor of geography and urban economics, so I always wanted to complete a PhD. It was never a question of running off and joining a startup, even though my startup friends were becoming millionaires left and right.*

As his PhD program extended into its sixth year, Roy felt pressure from many directions to finish his thesis and move on. But although he had explored many interesting questions and had played an important part in the foundation of the Web industry, he still wasn't sure what his thesis would be about:

> *I was talking to a colleague, Larry Masinter, about a related subject,*

and was telling him how I'd done all this work but I didn't know what to do for my dissertation. He said, "You're the only one who can describe HTTP's design rationale—the details of why it's there and what it's good for—why don't you just do that?" But at the time, I hadn't considered that to be my academic work; it was just my practical work for the Web. For me, going back to researching and finding the real knowledge framework for architectural styles was my way of fitting it all together. I learned an entirely new vocabulary in order to write my dissertation.

Roy wrote his dissertation as if he were writing a book on the architecture of HTTP. He coined the term REST to capture the idea that the elements of networked applications were exchanging their states using protocols across the network. He never expected that his thesis would be so widely read:

It's gratifying that people like to read my dissertation. It's an accessible piece of work. There's one equation, which is there just to have an equation. It's not actually necessary, but I felt I had to have one.

It's rare that a PhD student has the flexibility and support of his or her institution to participate in developing an entirely new industry while researching a thesis about that industry. Roy sees his participation in both the academic and industrial worlds as an advanced form of technology transfer:

This freedom gave me the ability to do technology transfer beyond my wildest imagination. What's hilarious from my standpoint is that I was just having fun. I was given resources so I could do my good deed for the universe, but it was really about working with enjoyable people and having wonderful conversations. And I learned an incredible amount.

They say that timing is everything. As Roy was starting his PhD, the Web was being invented. Roy finished his thesis in 2000, when the Web had forever changed communication, technology, and commerce worldwide. Finishing his PhD at the beginning of a new decade provided an opportunity to pause and reflect on the essential structure of Web technologies. His thesis continues to guide the architecture of today's modern distributed REST-style applications.

CHARLES SEVERANCE, Computing Conversations column editor and *Computer*'s multimedia editor, is a clinical associate professor and teaches in the School of Information at the University of Michigan. Follow him on Twitter @drchuck or contact him at csev@umich.edu.

32 & 16 YEARS AGO

EDITOR **NEVILLE HOLMES**
holmeswn@yahoo.com.au

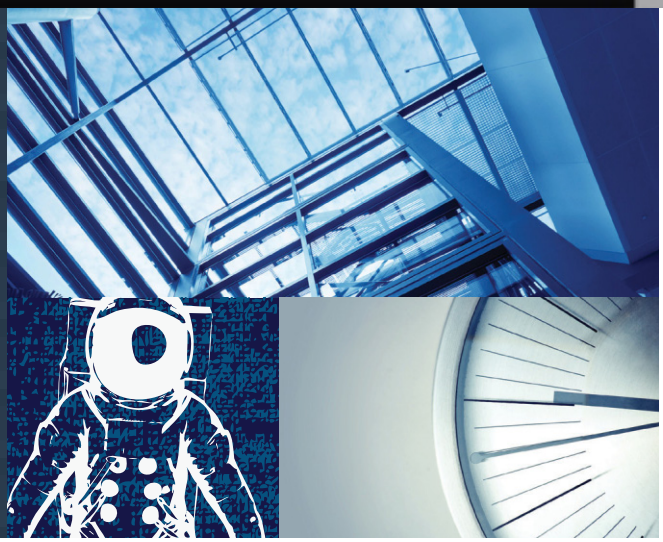

JUNE 1983

www.computer.org/csdl/mags/co/1983/06/index.html

Introduction (p. 13) "This month's issue is a sequel to *Computer*'s coverage of array processor architecture (September 1981). That issue presented articles on the internal architecture and system capabilities of these machines. The articles in this issue describe end-user applications, showing how the array processor fits into the system and how it can be used to achieve the necessary high-speed calculations. The main purposes are to show the utility and payoff of the array processor for specific applications and to explore its many significant uses."

Medical Imaging (p. 17) "Current imaging modalities include nuclear medical imaging, computed tomography scanners, ultrasound techniques, positron emission tomography, digital radiography, and nuclear magnetic resonance. This article focuses on just three: computed tomography, digital radiography, and nuclear magnetic resonance. These modalities represent the three techniques where the use of [array processors] is most easily justified."

Radar Processing (p. 32) "A number of signal processing techniques have been developed over the years to separate desired signals from undesired signals in radar raw video. Many of these techniques were implemented initially in the form of analog circuits, but as digital circuits became faster, a growing number of these functions were implemented in digital form. In addition, many new techniques have been developed that are possible only with digital processors."

Nuclear Fusion Simulation (p. 44) "Nuclear energy research at Lawrence Livermore National Laboratory indicates that field-reversed magnetic [FRM] mirror machines may offer greater efficiency and higher electrical yields than current reactor designs. Since theoretical analysis of the FRM is difficult, we are modeling the ion kinetics and electron physics of a fusion plasma to analyze the buildup and decay processes that might be characteristic of an actual machine. We have developed a simulation that accurately models the developing current distribution and use an array processor to obtain faster run times."

Nonlinear Filtering (p. 51) "Our application concerns developing code to build the best phase demodulator for use in areas such as deep-space and submarine communications. The current linear design, known as the phase-lock loop, is suboptimal. Our approach was to synthesize a nonlinear filter, addressing three specific problems."

Flight Simulation (p. 62) "Real-time pilot-in-the-loop flight simulation presents a number of unique computational problems: fast and efficient execution of differential equations, linear interpolation of many large tables, and a computer operating system that allows for timely interrupt handling without unnecessary interference. Special-purpose digital computers, such as array processors, can significantly enhance the performance of a simulation computer system by handling some of these tasks."

Computational Astrophysics (p. 73) "We have used the [array processor] to solve diverse problems in theoretical astrophysics and general relativity problems that encompass many different physical regimes and, naturally, require a variety of solutions. The problems examined thus far deal either with the dynamical structure and evolution of self-gravitating, large *N*-body systems or with the spherical collapse of stars to black holes."

Finite Element Analysis (p. 85) "Engineering firms are finding computer systems dedicated to engineering design activity to be cost-effective. However, the [finite element analysis] process has high computational requirements; a superminicomputer that can handle data preparation and interpretation of analysis might not have the capacity for the computationally intensive solution portion of analysis. A high-speed array processor—an attached processor—added to a superminicomputer can offload it for the more interactive portions of the design and analysis cycle."

Outlier Problems (p. 95) "Special problems confront managers of computer installations in developing countries. Interestingly enough, lack of funds may not be one of them, since the companies and institutions that install computer systems are usually financially strong. Other problems, such as turnaround time, quality of operating systems,

0018-9162/15/$31.00 © 2015 IEEE

and rapidly changing technology, are worldwide in nature and not peculiar to developing nations. The areas that do present special difficulties are staff recruitment, training, and retention; resource acquisition and maintenance; and infrastructural inadequacies."

JUNE 1999

www.computer.org/csdl/mags/co/1999/06/index.html

Letter 1 (p. 6) "[Kenneth L. Thompson] cannot be unaware of the extent to which large, viable commercial and noncommercial entities are deploying Linux these days precisely for the uses for which he deems Linux to be 'unreliable.' And his reference to Linux as less reliable than Microsoft is simply and utterly unsupportable."

A New Virus (p. 16) "There was concern because Melissa spread quickly via infected email attachments that, when opened, sent the virus to people in unsuspecting victims' address books. This meant that Melissa spread very fast throughout the world, overwhelming email servers, before it could be stopped."

Latin Y2K (p. 22) "Joyce Amenta, Y2K program coordinator for the World Bank's InfoDev program, which grants financial assistance to developing nations' IT investments, said that not one Latin American nation surveyed by InfoDev reported its Y2K correction program was on schedule or completed."

Emerging Standards (p. 42) "System on chip [SOC] opens the way for a multitude of electronics products not economically feasible before. SOC technology holds the key to increasingly complex applications by enabling high-performance, embedded processing solutions at a low, single-chip cost. Over the past two years, the Virtual Socket Interface Alliance development working groups have made great progress in defining the interface challenges, identifying the required functions, and recommending specific data formats and guidelines."

Chip Testing (p. 52) "The attributes that make system chips built with embedded IP cores an attractive methodology ... also make testing and debugging these chips a complex challenge. The authors review the various alternatives for testing embedded cores and describe solutions and proposed standards that are expected to play a key role in developing the core-based design paradigm."

SOC Products (p. 61) "To dispel a potential misperception that system-on-chip designs are a long way off, the authors describe several examples of new products that derive benefits from using SOCs. These designs, like those for embedded processors, emphasize the combination of a complementary set of functions into an economically viable package."

Data Indexing (p. 67) "Autonomous Citation Indexing [ACI] can help to evaluate the importance of individual contributions more accurately and quickly. Digital libraries incorporating ACI can help organize scientific literature and may significantly improve the efficiency of dissemination and feedback. ACI may also help speed the transition to scholarly electronic publishing."

Cache-Only Memory Architecture (p. 72) "The authors explain the functionality, architecture, performance, and complexity of [Cache-Only Memory Architecture (COMA)] systems. They also outline different COMA designs, compare COMA to traditional nonuniform memory access (NUMA) systems, and describe proposed improvements in NUMA systems that target the same performance obstacles as COMA."

Peripheral Component Interconnect (p. 80) "While the proliferation of peripheral component interconnect (PCI)-related standards attests to PCI's success, some confusion arises because it is not always clear why there are so many standards, what the differences are between them, and which standards apply to a specific application. To address these questions, the authors provide a detailed description of the PCI bus architecture. They also review the various standards that enable the use of PCI silicon, software, and firmware in various segments of the computer industry."

Computers in Space (p. 115) "Over the past few decades, advances in ground-based processing and space-to-Earth links have fallen further behind NASA's requirements for handling observation data. Thus NASA's ability to observe and capture Earth phenomena of theoretical and practical interest far outstrips its ability to transfer, process, or store such data."

Internet Shopping (p. 120) "US corporations spend more on IT than their European counterparts because they're using that money to attract and satisfy global customers. While European companies chafe against state and local regulations that keep that region's offline stores closed evenings and weekends, US corporations build 24/7 shopping malls in cyberspace. While Europeans wrestle with privacy concerns and promote the adoption of more stringent privacy laws, US consumers willingly exchange personal information for shopping convenience."

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

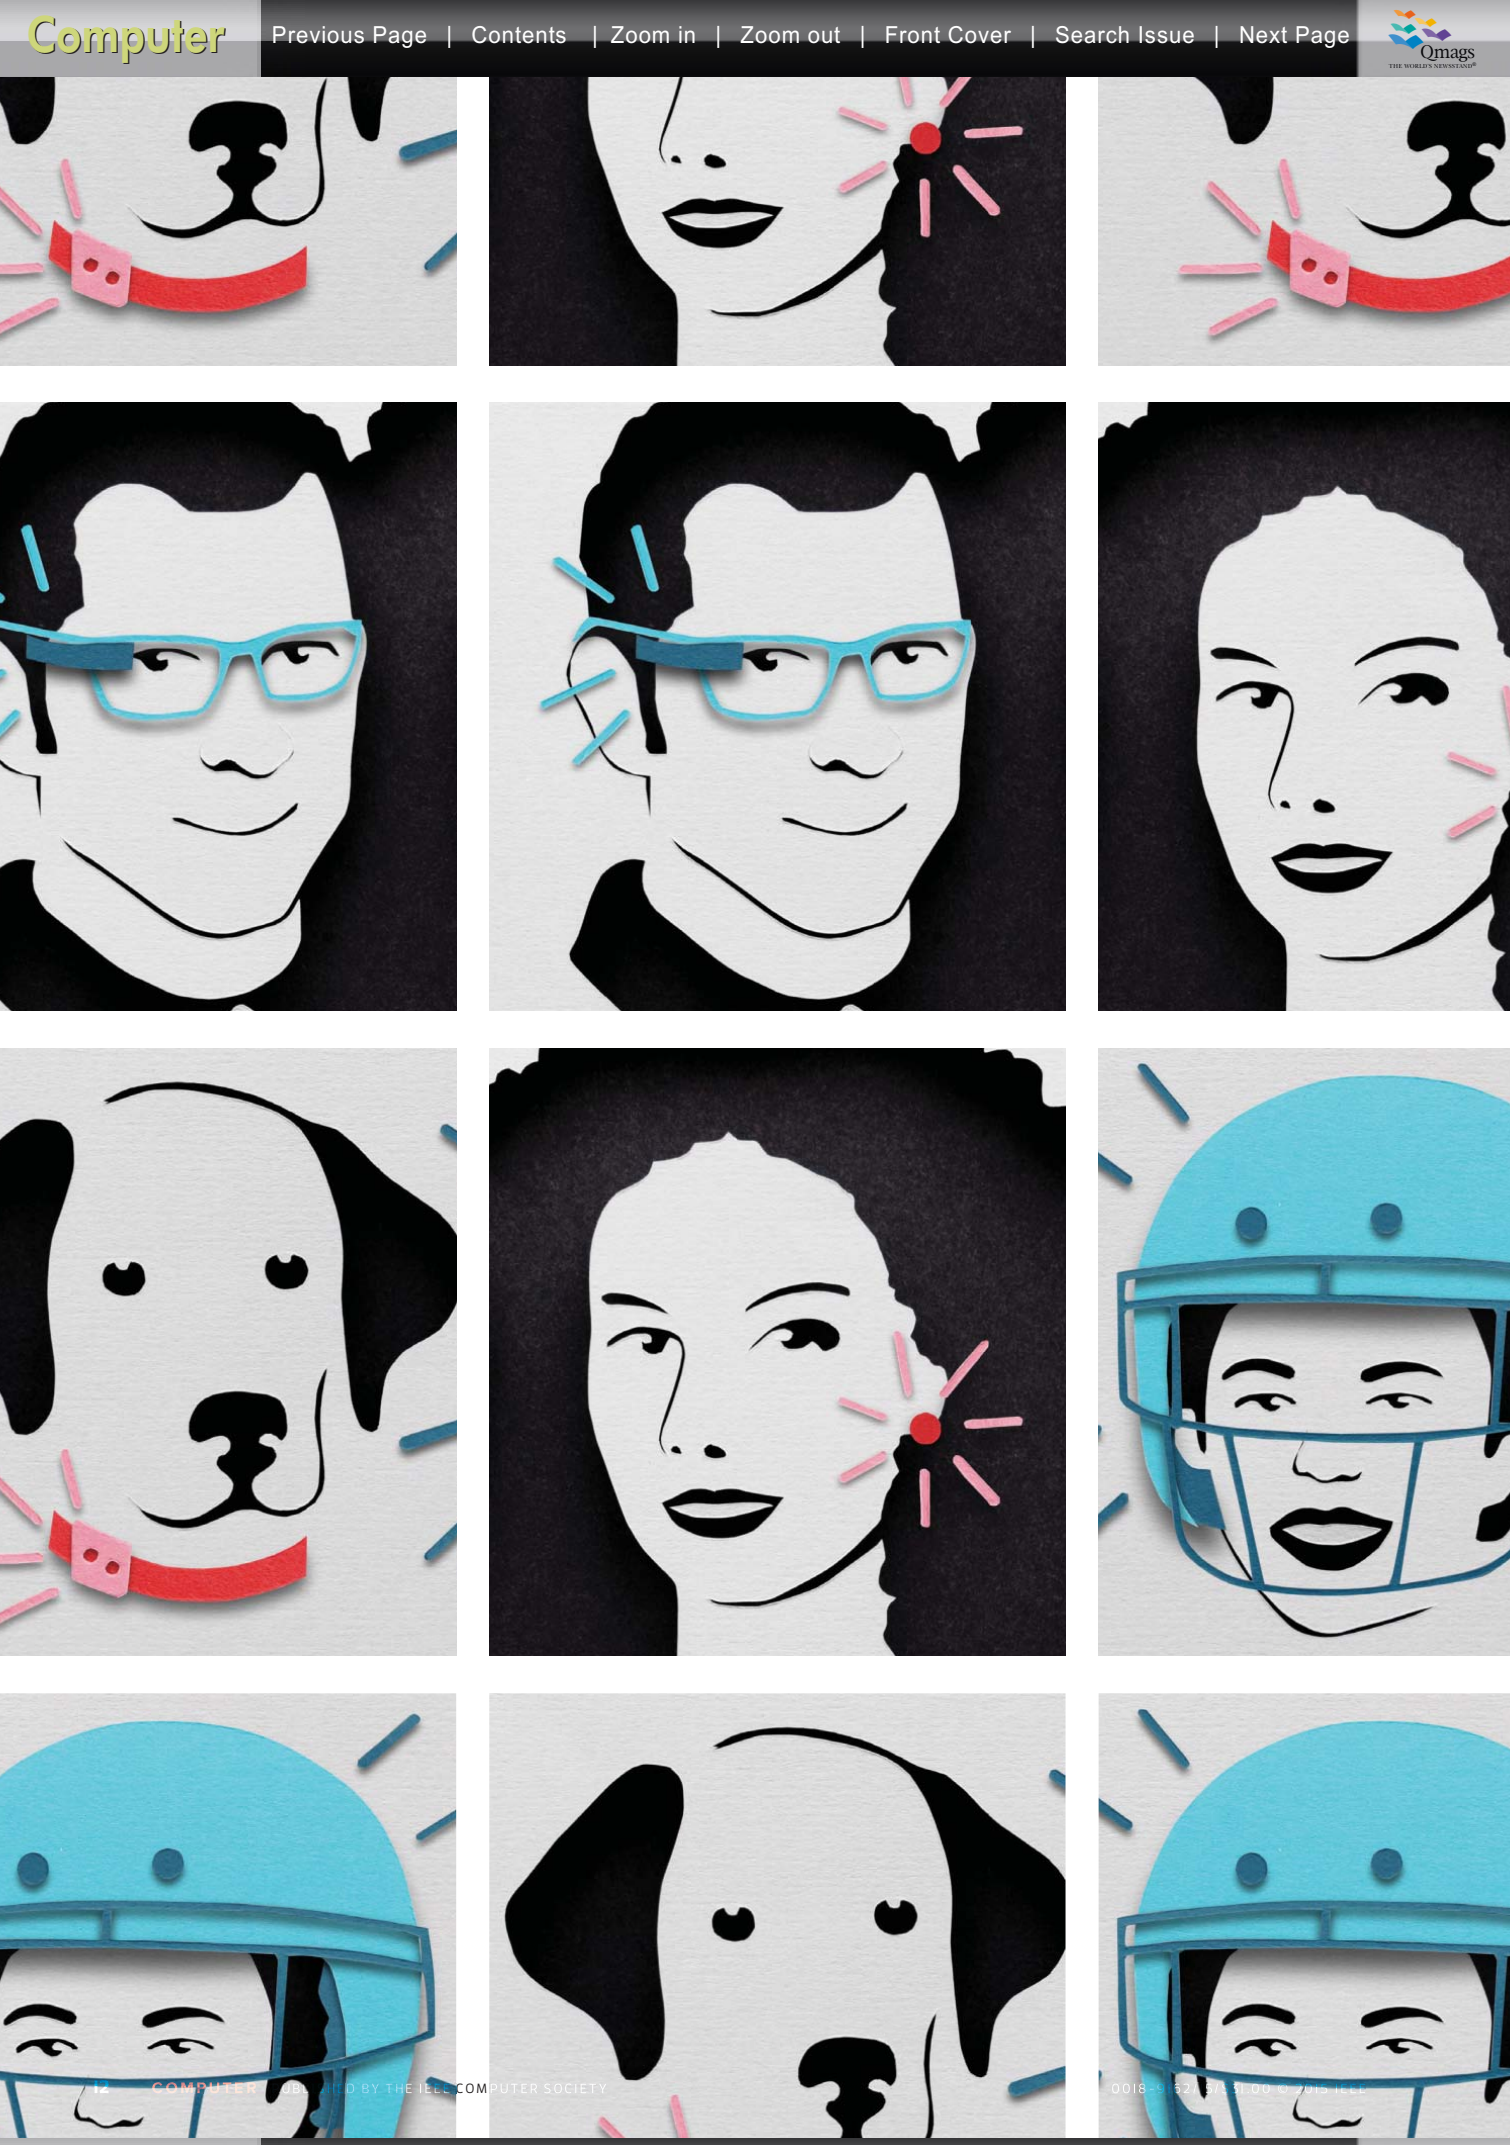


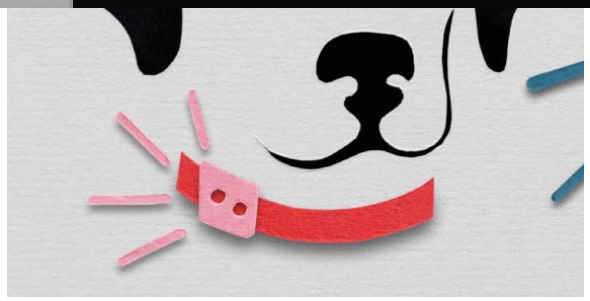


COVER FEATURE **GUEST EDITORS' INTRODUCTION**

WEARABLE COMPUTING: THE NEW DRESS CODE

Thad Starner, Georgia Tech
Tom Martin, Virginia Tech

Technology's presence on the body expands our capabilities and opportunities.




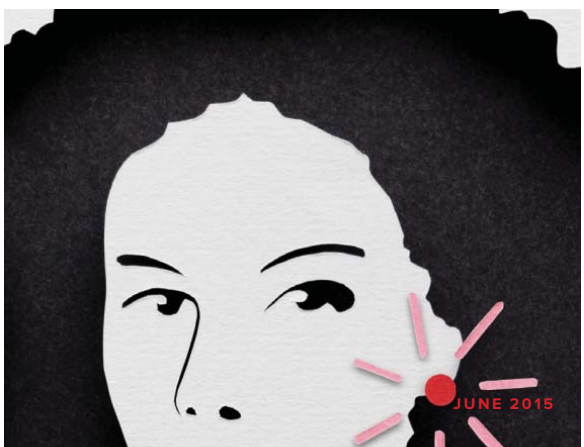


GUEST EDITORS' INTRODUCTION

Paradoxically, the closer technology gets to the body, the more it tends to disappear into our lives. Smartphones, MP3 players, Bluetooth earpieces, and fitness trackers are increasingly part of today's dress code. These devices empower us to communicate, relax, and work out more effectively. The new generation of "wearables" promises to go even further: new technologies will enable these on-body devices to sense and understand our environment and provide just-in-time information in ways that do not distract from the primary task, improving our efficiency and encouraging our creativity.

Technologists often think of augmenting the human body or its abilities when developing a new technology, thus the idea of wearable computers is almost as old as computing itself. In 1945, when Vannevar Bush articulated his vision for the computer of the future, he included concepts that we now recognize as the personal computer, the online encyclopedia, hypertext, the Internet, *and* a head-mounted camera.[1]

DAWN OF WEARABLE COMPUTERS

According to Bush, scientists and researchers of the future would work with these tools; the head-mounted camera, for example, was a method for capturing interesting data for incorporation into a memex library that could later be searched, cross-referenced, and augmented by others. Indeed, these technologies are used today in a wide variety of applications, from helping scientists to collect experimental data, as predicted by Bush,[2] to enabling parents to capture and share a child's first steps.

Similarly, in 1967, Hubert Upton devised a computer system for helping people with hearing loss to read lips; it used a head-up display (HUD) on a pair of eyeglasses so that the user could watch his conversational partner's lips and see acoustic information overlaid in real time. The user could remain focused on the primary task—the lip movements and the conversation—as opposed to glancing continuously at a handheld display and possibly missing vital information. Perhaps smartphone designers could take a lesson from this early wearable technology and make a device that augments face-to-face conversation instead of distracting from it!

Timekeeping devices also moved to the body. Though the first timepieces were bulky and designed to be seen from across the town square, they went through continuous refinement and miniaturization until they were small enough to be used in the home—placed in a hallway or on the mantel. With further improvement and miniaturization, the clock could finally be moved onto the body and be worn on the wrist—thus the wearer finally "owned" the time.

Checking the time became instantaneous—a task delegated to the wrist, where it could be checked with little effort. This advancement helped support other technical advancements; for example, because early pilots already had their hands full with the yoke, they could not afford the extra visual and manual attention required by a pocket watch, thus the wristwatch was essential. Similarly, World War I infantrymen had to coordinate the timing of their charges with great precision across long distances while holding their rifles and struggling out of trenches, thus wristwatches became standard issue for enlisted British soldiers. Of course, in 2015, several manufacturers are bringing their smartwatches to market, hoping for history to repeat itself. Will the smartphone eventually be relegated to pocket watch status?

Technologists' fascination with on-body technology is understandable—it promises great convenience and improved efficiency, and even an augmentation of our own human abilities. As wearable computers shrink and become further integrated into daily life, they become extensions of ourselves.

> AS WEARABLE COMPUTERS SHRINK AND BECOME FURTHER INTEGRATED INTO DAILY LIFE, THEY BECOME EXTENSIONS OF OURSELVES.

IN THIS ISSUE

In "Order Picking with Head-Up Displays," Anhong Guo, Xiaolong Wu, Zhengyang Shen, Thad Starner, Hannes Baumann, and Scott Gilliland show that a HUD provides just-in-time task guidance for improving speed and accuracy during order picking. These experiments show how wearable

ABOUT THE AUTHORS

THAD STARNER is a professor of computing at Georgia Tech and one of the founders of the field of wearable computing. His research projects and interests include Google Glass, brain–computer interfaces that use American Sign Language, wearable computers for service dogs, and pattern discovery to aid research in two-way communication with wild dolphins. Starner received a PhD in wearable computing from MIT. He is a member of IEEE and ACM. Contact him at thad@gatech.edu.

TOM MARTIN is a professor of computer engineering at Virginia Tech, with a courtesy appointment in the School of Architecture + Design, as well as a Senior Fellow at the Institute for Creativity, Arts, and Technology. His research interests include wearable computing, electronic textiles, and interdisciplinary design for pervasive computing. Martin received a PhD in electrical and computer engineering from Carnegie Mellon University. Contact him at tlmartin@vt.edu.

DISCLAIMER

Peer review and decision-making for "Order Picking with Head-Up Displays" by Anhong Guo, Xiaolong Wu, Zhengyang Shen, Thad Starner, Hannes Baumann, and Scott Gilliland were handled by an independent, nonconflicted *Computer* editorial board member; the guest editors had no role in determining the outcome of this article.

computers reduce workload, assist users with their tasks, and require less attention—all while minimizing inconvenience compared with current order-picking techniques.

As with the earlier Upton prototype, wearables can be of great assistance to people with disabilities. In "Can Disability Discrimination Law Expand the Availability of Wearable Computers?," Reuben Kirkham argues that the capacity to improve lives should override potential privacy and copyright concerns when it comes to these on-body devices. By making an analogy to service dogs—which are permissible in many venues where animals are not normally allowed—he makes the case that wearable devices could also open the door to better employment opportunities and access to public services. The wearable computing community should, Kirkham says, "fully take into account disability discrimination law in order to maximize the opportunities available for both disabled and non-disabled users," and he outlines several approaches to this effect.

In "Inferring Mobile User Status with Usage Cues," Jon C. Hammer and Tingxin Yan explore how variations in smartphone usage can help predict whether a user is busy, happy, alone, or stressed. Interestingly, one of the most telling cues is whether the user checks the phone without opening an application, presumably to check the time or some other status. The authors hypothesize that their findings can be adapted for wearable computers in general. The methods tested are more energy efficient than the addition of hardware sensors, which points to the practicality of this approach.

Finally, in "Beeping Socks and Chirping Arm Bands: Wearables that Foster Free Play," Andrea Rosales, Sergio Sayago, and Josep Blat show us the playful side of wearable computers. They describe how four qualities of wearables—individuality, natural interaction, ubiquity, and intimacy—can foster rich and diverse free play, which is key to learning in young children. The authors study three accessories worn on the feet, wrist, and waist, and explore how their findings can help support future learning development.

The diversity of topics in this special issue of *Computer* reflects the eclectic and highly multidisciplinary nature of the wearable computing field. We hope you enjoy these articles and will consider contributing to the field yourself through the Google+ Wearable Computing community or the annual International Symposium on Wearable Computers.

REFERENCES

1. V. Bush, "As We May Think," *Life*, 10 Sept. 1945, pp. 112–118.
2. P.M. Scholl and K. Van Laerhoven, "Wearable Digitization of Life Science Experiments," *Proc. ACM Int'l Joint Conf. Pervasive and Ubiquitous Computing* (UbiComp 14), 2014, pp. 1381–1388.

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

Order Picking with Head-Up Displays

VIDEO

Anhong Guo, Carnegie Mellon University
Xiaolong Wu, Zhengyang Shen, and Thad Starner, Georgia Tech
Hannes Baumann, Ubimax
Scott Gilliland, Georgia Tech

Experiments suggest that using head-up displays like Google Glass to support parts picking for distribution results in fewer errors than current processes. Making Glass opaque instead of transparent further improves selection efficiency.

Globally, roughly US$1 trillion in goods are distributed from nearly a million warehouse sites each year, and for many businesses such activity represents 20 percent of their logistics costs.[1] Order picking—the process of selecting items from inventory racks with *pick bins* and sorting them into *order bins* for distribution—accounts for about 60 percent of these warehouses' total operational costs.[2]

Current robotic systems lack the dexterity to handle the variety of parts on most pick lines, so the vast majority of western European warehouses still use manual picking,[1] which is costly and time consuming. Although manual methods vary, most warehouses still use paper lists that include each item's location, identifying number, and required amount. Such systems are error prone and can cause significant losses, as in automobile manufacturing, where the wrong part can halt an assembly line. In e-commerce, inventory errors can compromise order fulfillment, possibly increasing customer dissatisfaction.

Although technologies like parts-to-picker systems that bring parts bins to warehouse workers can facilitate various parts of the picking process,[1] such systems are expensive and relatively rare. Thus, parts bins are typically stationary, and the picker must rely on a paper list or expensive pick-by-light systems that use displays at each bin to indicate which parts to pick.[3]

Wearable computers are becoming more popular with manufacturers as a way to guide pickers, so we constructed an environment for two experiments to compare pick errors and speed. The first looked at paper, pick-by-light, head-up display (HUD), and cart-mounted display (CMD) systems, and the second evaluated a transparent versus opaque version of Google Glass. Our experiments show that HUDs outperform other technologies and that the opaque version of Google Glass might

0018-9162/15/$31.00 © 2015 IEEE

be more efficient for order picking than the default transparent version.

PICKING METHODS

Picking methods range from the basic paper list to task-guidance systems that are either worn or embedded in the environment. All methods aim to help the picker select the correct parts from the pick bin and deliver them to the correct order bin.

Paper

In the pick-by-paper method, the picker refers to a printout of items to be selected and their locations—one list for each task, which can have multiple orders. The advantages are simplicity and a relatively low implementation cost. On the downside, text-only lists can be difficult to read or interpret when product numbers are long. Often, the first optimization done on a pick line is shortening product numbers to the last three digits (if they are unambiguous).

The list-based picking process forces the picker to perform each step in two stages: read and interpret the list's fine detail and then move a part from one area to another. Often, pickers have the list in one hand while reaching for the parts, which limits the number of parts they can grasp at any one time.[2]

Light

The pick-by-light method is becoming more prevalent in manufacturing warehouses, even though system implementation cost can be as high as US$1,200 per meter. Plant managers report that the virtual elimination of pick errors and the increased picking speed are well worth the investment.[2]

As Figure 1 shows, with pick-by-light, warehouse bins often have small displays and push buttons. When a picker traverses the aisles, the displays illuminate on the pick bins, typically showing the quantity to be picked. Most pick-by-light systems require that pickers press a button to indicate they have picked from the correct bin. In more sophisticated systems, sensors try to detect the picker's reach into each bin. In such systems, the display on a pick bin goes out after each pick is confirmed, and another light illuminates the order bin where the picked parts are to be placed. Pickers tend to dislike the buttons and sensors on pick bins because the sensors do not always correctly detect the reach, which requires more button pushes to fix.

Proximity or weight sensors, or a button press by the picker, indicate when an item is placed in the order bins. Sensor activation or the button press triggers the display of the next task's picks. Although they dislike sensors on pick bins, pickers generally find the order bin sensors valuable.

Cart-mounted display

Pick-by-CMD displays a graphical representation of the picks on the order cart. In some cases, warehouses use a high-resolution LCD display to display picks for an entire shelving unit instead of the pick-by-light systems that indicate each pick bin. Instrumenting the order cart seems more effective and less expensive than instrumenting each shelving unit, but we have seen few such installations.

Head-up display

In this system, the picker wears a HUD that shows the pick charts needed for each shelving unit. Similar to the pick-by-light method, when the picker drops items into the order bin, the HUD displays the next pick chart. Figure 2 shows various technologies suitable for a HUD-based picking system, including an opaque version of Google Glass. Glass is lightweight and self-contained, requires no additional computer, and has sufficient processing power and network connectivity to run order-picking software.

Other methods

Commercial pick-by-voice systems use wearable computers to guide pickers

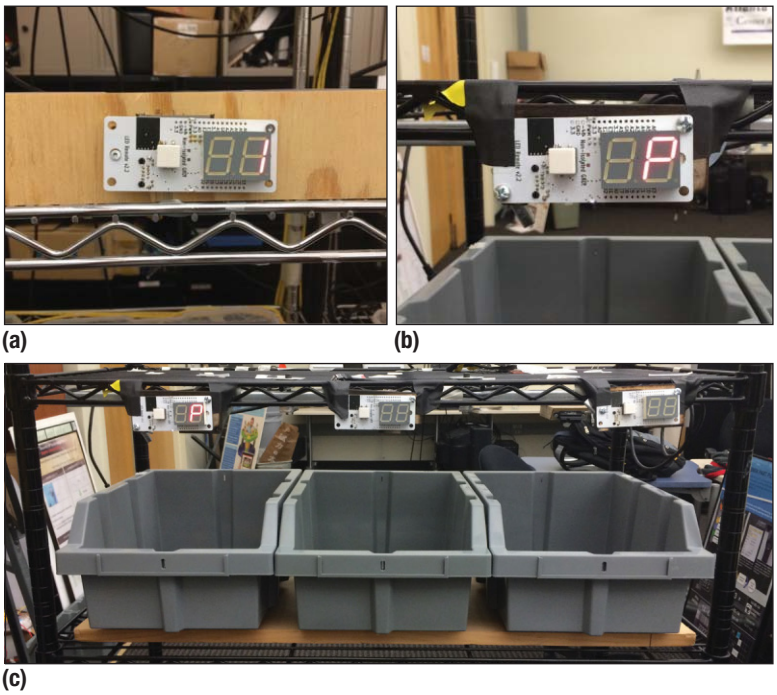

FIGURE 1. Pick-by-light method. (a) The pick bin displays the number of parts to pick (1), while (b) the order bin shows where to place (P) the order. (c) Typically, the order bins are mounted on a wheeled order cart so that the pickers can transport the order bins as they walk through the stacks of pick bins in the warehouse.

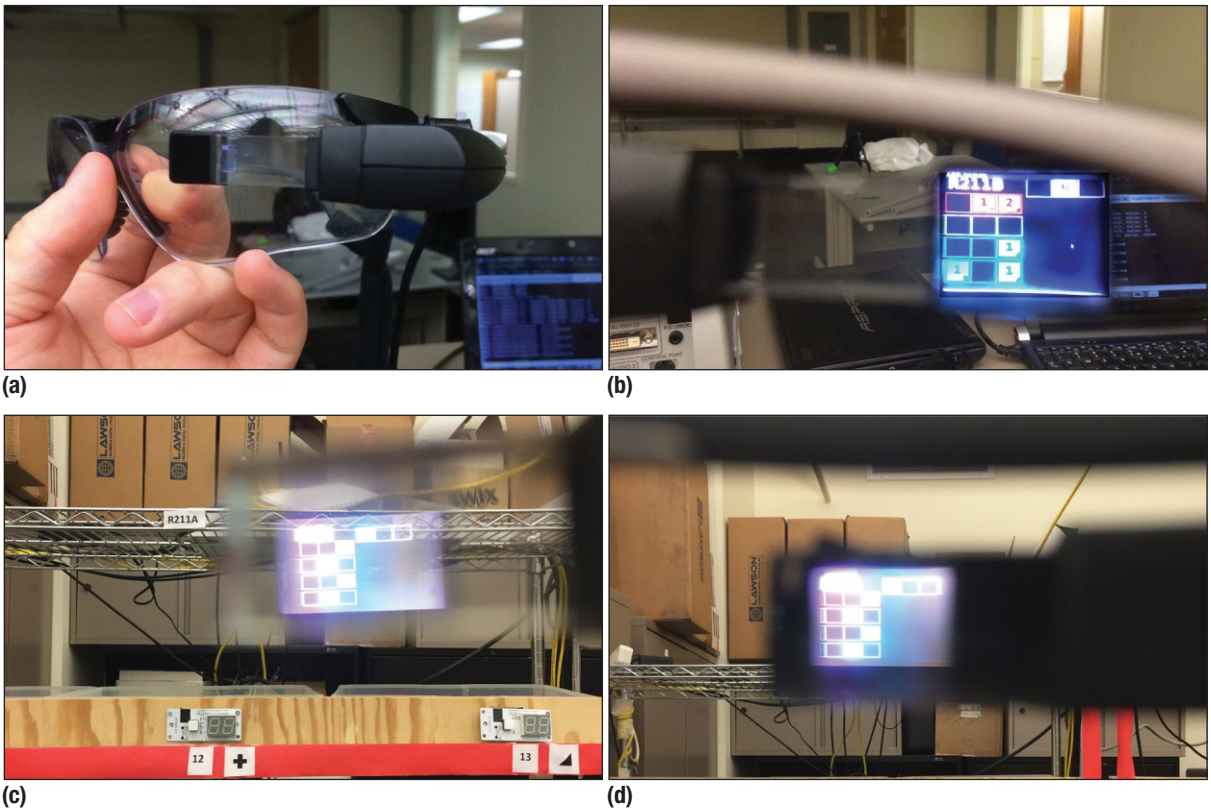

FIGURE 2. Displays used in pick-by-head-up display (HUD) systems. (a) MicroOptical SV-3 HUD mounted on safety goggles. (b) View through the SV-3, (c) view through Google Glass with a transparent display (default), and (d) view through an opaque display.

through their tasks using spoken audio prompts and simple speech recognition. In previous experiments, these systems eliminated many errors but were significantly slower than all other systems, including the standard pick list.[4,5] Similarly, some commercial systems require pickers to confirm each pick by using a handheld scanner to scan a barcode on each pick bin. These systems also eliminate many errors but seem even slower than paper or voice.

Pick-by-vision[6] uses a HUD and motion tracker to overlay 3D graphic tunnels onto the picker's visual field. The tunnels then guide the picker to the right position in the warehouse and highlight the correct pick and order bins with a surrounding frame. Augmented reality systems that use such registered graphics have not yet shown a significant improvement in either accuracy or speed over a paper pick list.

Recent experiments in Germany introduced a pick-by-projector method which guides pickers using video projectors that illuminate pick and order bins. One implementation[5] uses Microsoft Kinect sensors to sense the picker's motions and provide feedback. Preliminary results are promising relative to paper-, voice-, and vision-based picking methods and might eventually prove a viable, albeit more expensive, alternative to the pick-by-HUD system.

EXPERIMENTAL ENVIRONMENT

To evaluate order-picking technologies, we built a dense-picking environment in our research laboratory that uses intelligent storage locations and order batching.[2] Figure 3 shows our experimental environment. Figure 4 shows how the paper- and display-based prompts correspond to the shelves and bins.

Our warehouse environment consisted of 24 pick bins divided between two shelving units, A and B. Each shelving unit has four rows and three columns, and each pick bin contains 20 to 40 items. The order cart, shown in Figure 3b, has three order bins coded with a square, cross, or triangle. The cart's top row holds the paper tasks or the CMD.

We used our environment to conduct two experiments: the first compared the paper, light, HUD, and CMD methods, and the second compared a pick-by-HUD system using transparent and opaque versions of Google Glass. In these experiments, unlike our earlier work,[2,4] we increased task variety, attempting to induce more performance errors for comparison purposes.

We define key terms in our experiments as

- *pick*—one reach into a pick bin and the removal of one or more parts;
- *place*—putting all items currently being carried into an order bin;

FIGURE 3. Experimental dense-picking environment. (a) The environment consists of 24 pick bins arranged on two shelving units in four rows and three columns per unit. (b) It also contains an order cart with three order bins and a cart-mounted display (CMD) on the top left shelf, which shows a graphical pick chart corresponding to the colored labels in (a). (c) A room diagram shows two video cameras (bottom and top left) to capture the picker's actions, a still camera (bottom right) to capture the picked parts in the order bins for later error analysis, the picker (green figure), and two experimenters (black figures).

- *task*—a collection of up to six subtasks; and
- *subtask*—moving one to seven items from one shelving unit to one order bin.

A task consists of up to six subtasks corresponding to three orders: three sets of picks from shelving unit A, placed into three order bins respectively, followed by a set of three sets of picks from shelving unit B, also placed into the three respective order bins. For each subtask, we randomly assigned one to seven items to be picked from one to five pick bins on shelving unit A or B. These items are placed in a single order bin (1, 2, or 3). For example, in Figure 3a, the subtask involves shelving unit A, and 4 of the 12 possible pick bins are involved. The five picked parts are placed into order bin 1 (Figure 3b). Figure 3c shows the layout of the experimental venue.

For our testing, we have simplified the pick list and attempted to make it as efficient as possible. Figure 4a shows a sample. For the first subtask in the sample list, the participant would pick one item from row 1, column 3; one item from row 2, column 2; two items from row 3, column 3; and one item from row 4, column 2. The participant would place the five items into order bin 1.

Figure 4b shows a pick chart, a graphical representation of the pick list. We use these pick charts for our pick-by-HUD and pick-by-CMD methods. Each pick chart shows the arrangement of the items to be picked for each order bin from each shelving unit. The colors on the pick chart correspond to the colors of the shelf rows shown in Figure 3a. Each row has three bins, and each bin bears a square, cross, or triangle. The pick chart uses these symbols to help cue the picker. In previous research, we showed that such color coding and symbol cues improve pick accuracy.[2] Adjacent colors are arranged so as to be unambiguous to color-blind pickers. Pick charts are significantly more efficient and accurate than pick lists, even when used in paper form, as in Figure 4c.[4]

EXPERIMENT 1: COMPARING FOUR METHODS

In our first experiment, conducted in late 2013 to evaluate the paper, light, HUD, and CMD picking methods, we enlisted eight participants, ages 22 to 27. All the participants—five males and three females—were novices in order picking. Four were left-eye dominant.

We paid each participant US$20, and the study lasted approximately two hours. The participants were instructed to complete the tasks as quickly and accurately as possible. In both phases, we used a Latin square to counterbalance the order in which the participants performed each picking method.

For pick-by-light, CMD, and HUD systems, a participant could see instructions for only one subtask at a time. After the participant completed each subtask, an experimenter activated the next subtask, emulating the automatic process enabled by a weight or proximity sensor in the order bin. When a participant completed a task, the experimenter replaced the full order bins with empty ones and proceeded to assign the next task. For pick-by-light, we decided not to require the button-push confirmation for picks, since the pickers we interviewed did not like this requirement.

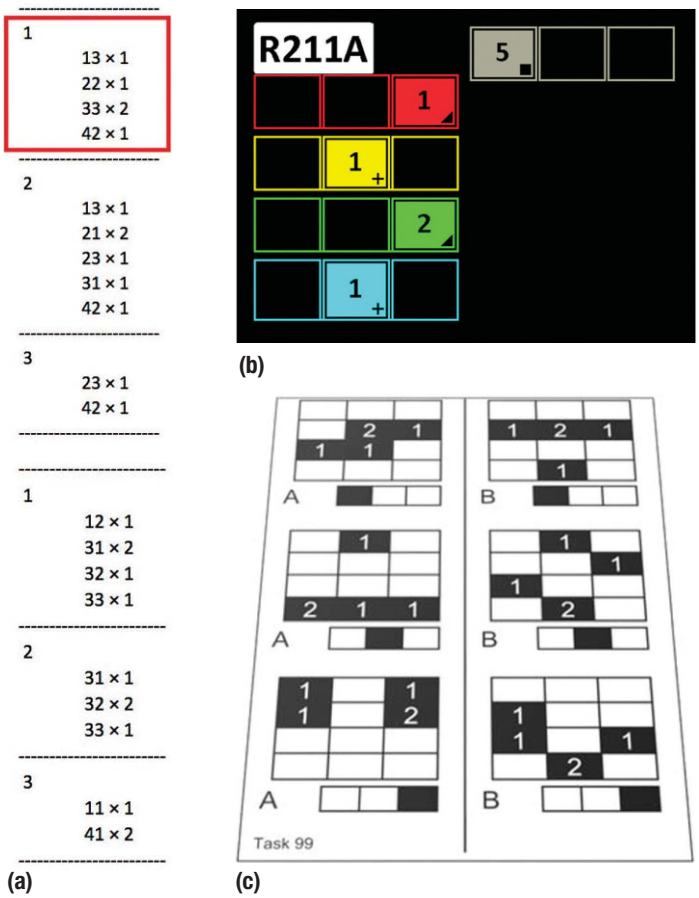

FIGURE 4. Task representations. (a) Sample paper pick list for task 001, subtask 1, on shelving unit A; (b) a graphical pick chart for the same task for the pick-by-CMD and pick-by-HUD methods (overlaid on the pick shelves in Figure 3a); and (c) a paper rendering of the graphical pick chart used in previous experiments.

For pick-by-CMD, we used an LCD-based laptop mounted on the order cart to display the pick chart so that the picker could refer to it with a simple turning of the head while picking. Our pick-by-CMD method was developed based on testing with order pickers at a major automobile manufacturer.[2]

For the HUD system, we used a MicroOptical SV-3 opaque HUD to display pick charts. The display was tethered to a laptop worn in a backpack (Figure 2a) Because Google Glass had not been released at that time, the SV-3 display served as a good proxy, as it has a similar field of view, head weight, and resolution.

For each of the four picking methods, participants performed five practice tasks as part of their training session. Afterward, participants performed 10 test tasks such that each experiment had a total of 20 practice and 40 test tasks.

After using each picking method, participants completed a NASA-TLX (Task Load Index)[7] survey. At the end of the testing phase, they ranked the four methods according to overall preference, ease of learning, comfort, speed, and accuracy.

Errors

To avoid learning-curve effects,[2] our study considered only the last eight tasks from each testing session. We hypothesized that the pick-by-HUD method would have a lower average error per pick and less average task time than the other three approaches. Our experiments were within-subject, and we used one-tailed, paired-sample *t*-tests, with a significance level of $\alpha = 0.05$.

Picks could have one of three types of errors: *item mistakes*, *wrong number*, and *wrong order bin*. *Item mistakes* had three subcategories:

- *substitution error*—when one part was swapped for another—which could be wrong row, column, or shelving unit;
- *missing-part error*—when the participant omitted a part; and
- *additional-part error*—when the participant placed an unrequested part in an order bin.

Wrong number errors have two subcategories—*too many* or *too few*—occurring when the participant selected too many or too few of the correct parts. *Wrong order-bin errors* occurred when the participant placed the items into the incorrect order bin. Errors of the types *too many* and *additional part* are not severe enough to stop an assembly line and might be discounted or ignored depending on the order picking domain.

Results

Figure 5a shows the methods' average task times. The HUD method's average task time was significantly shorter than using pick-by-light ($p = 0.002$)

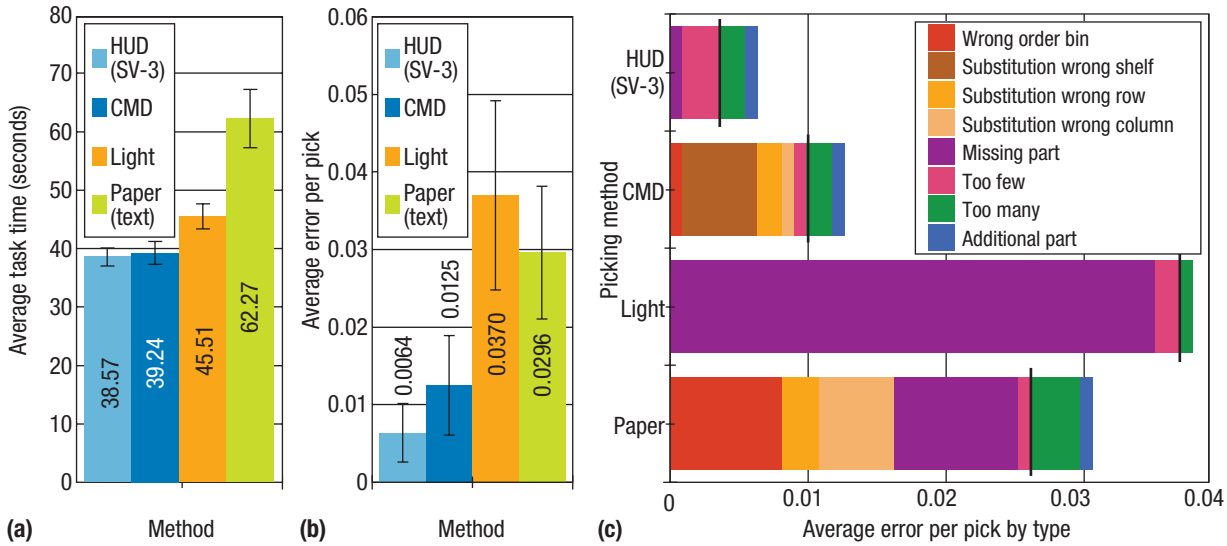

FIGURE 5. Experiment 1 results. (a) Average task time, (b) average error per pick, and (c) average error per pick by type. Errors in (c) are shown from most to least severe, with most severe to the left of the black vertical line.

and paper ($p = 0.0001$). Figure 5b shows the average error per pick. We counted one error for each error subcategory and used the total number of errors to calculate the method's average error per pick. Using the MicroOptical SV-3 HUD resulted in significantly fewer errors than pick-by-light ($p = 0.007$) and paper ($p = 0.018$).[8] Figure 5c shows the errors for all four methods divided by specific error type.

The results supported our hypothesis: the pick-by-HUD method had fewer errors and a shorter task time than either paper or light. Participants also rated the method more favorably, stating that it was much less work. The paper method's poor speed might stem from the need to use one hand to hold the list, and the many errors might be the result of the need to parse the text and remember it while scanning the shelves for the right bin.

HUD method. Our test results show that using pick-by-HUD can virtually eliminate errors and improve speed by approximately 30 percent over paper lists (Figure 5).[8] In addition, the average pick time with the HUD method is a statistically significant speed improvement over pick-by-light. The pick errors with pick-by-HUD, which were fewer than with any other method, also tend to be less severe. Study participants preferred pick-by-HUD over all other methods tested, and the method had the lowest workload as measured by the NASA TLX.

These results are surprisingly strong. Pick-by-HUD costs considerably less to implement than pick-by-light, making it a promising new tool for order picking. Products such as Ubimax's XPick (www.xpick.de) are bringing the technology to manufacturing environments, but much optimization work remains.

Light method. To our surprise, the pick-by-light method was slower and more error prone than pick-by-HUD. During testing, we saw that pickers were too close to the shelving units to see which other bins were lit, causing them to skip picks and not plan their motions as effectively as they could with a HUD, which offers a task overview.

The low performance is more understandable when considering the participants' view. Pick-by-light users often scan the shelving unit visually from top left to bottom right, which takes time. They also tend to step back frequently to see the entire shelving unit because they cannot keep the complete context in their heads. Even so, they tend to skip pick bins by accident, as evidenced by the large number of missing part errors in Figure 5c; bins on the shelving unit's periphery were particularly troublesome. Without making the picker press a button to confirm a pick, errors were worse than with a paper list. However, adding a button-push to each pick will slow the process even more.

CMD method. The high performance of the pick-by-CMD method was also a surprise. Speed, accuracy, workload, and preference results approach those of the pick-by-HUD method. However, the errors in the pick-by-CMD method are more critical. As Figure 5c shows, pickers picked from the wrong bin as opposed to not picking the right number of a given part (HUD method).

Perhaps the difference is procedural: pickers turn their heads often between the CMD and the pick bins. A picker who is focused on the cart's display might use peripheral vision to pick and place, resulting in more errors and workload. Regardless of the differences, our results suggest that both pick-by-HUD and pick-by-CMD methods are promising replacements for light and paper methods. Both seem

to reduce errors significantly, have lower setup costs than the pick-by-light method, and provide more flexibility.

EXPERIMENT 2: COMPARING HUD SYSTEMS

In our first experiment, our pick-by-HUD method used MicroOptical's SV-3 display, which performed well but required a tether to a backpack computer. The backpack is both inconvenient and uncomfortably hot.[2] For the second study (transparent versus opaque HUD), we used a self-contained Google Glass device without the backpack or tether, both alone and with a piece of electrical tape added to the Glass display to make it opaque.

Our second experiment used the same environment and most of the same procedure. The goal was to evaluate transparent Glass against an opaque version.

Transparent versus opaque

Our experiment has some similarity to Robert Laramee and Colin Ware's laboratory study evaluating the effects of transparent and opaque HUDs on a participant's ability to perform a table lookup task on a fixed screen.[9] Participants were stationary, and the task was designed for visual searching, which should be particularly sensitive to the visual interference possible with a transparent HUD. Indeed, their participants performed significantly better using the opaque screen.[9]

We wondered how much visual interference a transparent HUD might cause in order picking, because the user is glancing at the display but is focused more on the environment. We hypothesized that using an opaque display would result in lower average errors per pick and less average task time than a transparent display.

Procedure

We repeated the same testing procedure as in Experiment 1 except that we added 12 participants, 5 females and 7 males. All were ages 21 to 35, new to our studies, and novices at order picking. No one wore regular eyeglasses during the study, since glasses do not fit well with the Glass version we use. All participants who needed vision correction wore contact lenses.

For each of the two picking methods, the participants completed 5 training tasks and then proceeded to 10 test tasks. We again used a Latin square to counterbalance the order of the picking method chosen. In both the transparent and opaque versions, participants saw the same graphical chart (Figures 2c and 2d). To make Glass opaque, we covered the back of the display with black tape.

As in Experiment 1, we used only the last eight tasks from each testing session, randomizing task order for both sessions. We could then conduct a within-subject, paired-sample comparison between the methods for each task, which should be more sensitive than comparing average task performance. We accidentally skipped one task for one participant, which resulted in 95 (12 participants × 8 tasks – 1) task pairs for testing. This number was enough to estimate p values through Monte Carlo permutation methods (2^{95} combinations). Specifically, for each task pair we randomly selected one result for class A, the other for class B, and subtracted B from A. From these 95 differences, we calculated the average. We repeated this calculation process a million times to obtain a distribution of these averages, which should center on zero.

We calculated the actual experimental difference in average task times between the transparent and opaque conditions. Because our hypothesis was that the opaque display would outperform the transparent one, we determined the estimated p value by the percentage of the million trials that shows a difference greater than the experimental difference. We set the significance level at $\alpha = 0.05$ (meaning that only 5 percent of the random groupings show a greater difference). Out of curiosity, we also ran a standard one-tailed, paired t test on the data; the resulting p values were very similar.

> THROUGH REPEATED TESTING WITH ORDER PICKING, HUD OPTICAL PROPERTIES MIGHT BE OPTIMIZED FOR SIMILAR INDUSTRIAL TASKS.

Discussion

As Figure 6a shows, the average task time using the opaque display is less than when using the transparent one. The difference passes the significance test ($p = 0.0416$, 99 percent confidence interval [0.0410, 0.0421]). As shown in Figure 6b, the average error per pick for the two methods is close. Figure 6c

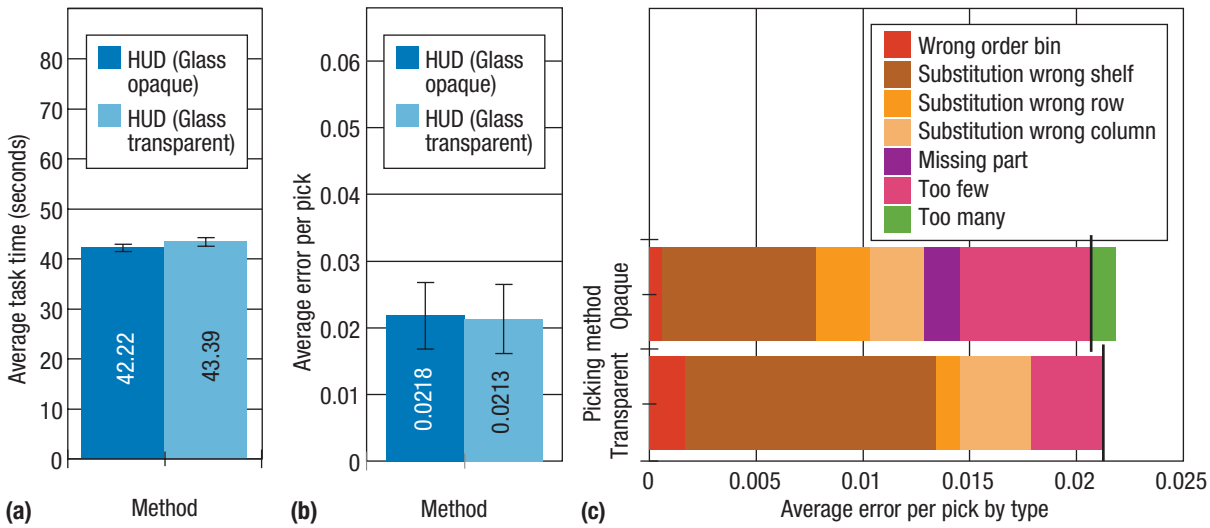

FIGURE 6. Results of comparing opaque and transparent Google Glass versions for order picking. (a) Average task time, (b) average error per pick, and (c) error per pick by type.

shows the errors divided into specific error types.

Although we were correct in our hypothesis that the opaque HUD results in faster picking than the transparent display, the difference is less than 3 percent—much lower than observed in Laramee and Ware's study with a stationary user.[9] Perhaps the difference is procedural; order picking alternates rapidly between virtual and physical worlds, while Laramee's stationary user dealt mostly with the virtual realm, which would maximize any effect from visual interference with the physical world.

A deeper investigation into the data shows a clear learning effect, which surprised us until we realized that we had given participants only 10 practice tasks instead of the 20 practice tasks for participants in the first experiment. This oversight might also explain the higher error rate: 2 percent with Glass overall relative to 0.6 percent with SV-3.

Our experimental design seems to have sufficient sensitivity to help select among picking system variations, enabling several lines for continuing work. Our picking environment is closely modeled after those we observed in the automobile industry, but our results should generalize to similar environments. Indeed, DHL (a major shipping firm), Ricoh, and Ubimax have recently seen a 25 percent improvement in picking performance with smart glasses. However, for pick environments that differ significantly from the one we tested, we could easily modify the experimental environment and retest.

Order picking might prove to be a good reference task for optimizing HUD characteristics for industrial environments, given that it requires constant movement, interaction with the physical world, and quick glances to a pick list or chart. Through repeated testing with order picking, the optical properties of HUDs (for example, brightness, contrast, field of view, color depth, focus, eye box, and bi-ocular versus monocular presentation) might be optimized for similar industrial tasks. A comparison of Glass and the SV-3 is a case in point. Although the two have a similar field of view, Glass is centered above the user's field of vision on the right eye, while the SV-3 is adjustable, and the view is slightly below the participant's line of sight. The SV-3 also has an adjustable focus and is often worn on the dominant eye, left or right. Is the higher error rate observed with Glass versus the SV-3 because of placement as opposed to the learning effect we hypothesized? A relatively quick study could determine the truth. Glass is becoming a popular experimental platform in industry, making it a logical choice for continued testing.

Safety is another important concern. Perhaps a HUD-based picking interface could in fact be safer than current paper lists because the HUD keeps pickers' hands free and is less distracting. If so, comparing the use of new HUD models to standard paper-based pick lists might provide a measure of safety assurance—much as tests in the automotive industry evaluate the safety of new dashboard interfaces by comparing the degree of distraction between using the interface and changing a radio channel.

The poor performance of the pick-by-light method also suggests an area for further study. How will including pick-confirming sensors or a button push affect accuracy and speed relative to the pick-by-HUD method? Adding scales under each order bin could reduce errors for all methods but might be particularly effective against missing parts or picking too few parts of a given type. If including both sensors and scales could virtually eliminate these errors, will pick-by-HUD

ABOUT THE AUTHORS

ANHONG GUO is a doctoral student at Carnegie Mellon University's Human–Computer Interaction (HCI) Institute. His research interests include wearable computing, crowdsourcing, and accessibility. Guo received an MS in HCI from Georgia Tech. Contact him at anhongg@cs.cmu.edu.

XIAOLONG WU is a master's student in HCI at Georgia Tech's School of Interactive Computing. His research interests include HCI, wearable computing, and user experience research. Wu received a BEng in industrial design from Shanghai Jiaotong University, China. Contact him at xwu86@gatech.edu.

ZHENGYANG SHEN is a master's student in HCI at Georgia Tech's School of Interactive Computing. His research interests include HCI, information visualization, and ubiquitous computing. Shen received a BEng in electronic information science and technology from Tsinghua University, China. Contact him at szy91@gatech.edu.

THAD STARNER is a professor of computing at Georgia Tech. His research interests include pattern discovery, wearable computing, and sign language. Starner received a PhD in wearable computing from MIT. He is a member of IEEE and ACM. Contact him at thad@gatech.edu.

HANNES BAUMANN is a project leader and software developer at Ubimax in Germany. While performing the work described in this article, he was a doctoral student in the Wearable Computing group of the Technologie-Zentrum Informatik und Informationstechnik, Universität Bremen, Germany. His research interests include HCI, wearable computing, and image processing. Baumann received a PhD in computer science from Universität Bremen. He is a member of the German Informatics Society (GI). Contact him at hannes.baumann@ubimax.de.

SCOTT GILLILAND is a research scientist at the Georgia Tech Interactive Media Technology Center. His research interests include rapid prototyping and embedded systems. Gilliland received an MS in computer science from Georgia Tech. Contact him at scott.gilliland@gatech.edu.

still be faster than pick-by-light and, if so, by how much? Answering these questions is just one of many directions for future research and could move industry more quickly toward assistance in performing a tedious and error-prone process.

ACKNOWLEDGMENTS

We thank Shashank Raghu, Xuwen Xie, Saad Ismail, Xiaohui Luo, and Joseph Simoneau, who helped produce the original pick-by-light study; Caleb Southern, who helped structure earlier versions of this article, the anonymous *Computer* reviewers, who provided valuable comments; and Nancy Talbert, who helped edit the final version of this article.

REFERENCES

1. R. De Koster, T. Le-Duc, and K.J. Roodbergen, "Design and Control of Warehouse Order Picking: A Literature Review," *European J. Operational Research*, vol. 182, no. 2, 2007, pp. 481–501.
2. H. Baumann, "Order Picking Supported by Mobile Computing," PhD dissertation, Univ. of Bremen, 2013.
3. M. De Koster, "Warehouse Assessment in a Single Tour," *Warehousing in the Global Supply Chain*, M. Lahmar, ed., Springer, 2012, pp. 457–473.
4. K.A. Weaver et al., "An Empirical Task Analysis of Warehouse Order Picking Using Head-Mounted Displays," *Proc. ACM Conf. Computer-Human Interface* (CHI 10), 2010, pp. 1695–1704.
5. J. Herter, "Augmented Reality Supported Order Picking Using Projected User Interfaces," bachelor's thesis 131, Univ. of Stuttgart, 2014.
6. B. Schwerdtfeger et al., "Pick-by-Vision: There Is Something to Pick at the End of the Augmented Tunnel," *Virtual Reality*, vol. 15, 2011, pp. 213–223.
7. S.G. Hart and L.E. Staveland, "Development of NASA-TLX (Task Load Index): Results of Empirical and Theoretical Research," *Advances in Psychology*, vol. 52, 1988, pp. 139–183.
8. A. Guo et al., "A Comparison of Order Picking Assisted by Head-Up Display (HUD), Cart-Mounted Display (CMD), Light, and Paper Pick List," *Proc. IEEE Int'l Symp. Wearable Computers* (ISWC 14), 2014, pp. 71–78.
9. R.S. Laramee and C. Ware, "Rivalry and Interference with a Head-Mounted Display," *ACM Trans. Computer-Human Interaction*, vol. 9, no. 3, 2002, pp. 238–251.

See **www.computer.org/computer-multimedia** for multimedia content related to this article.

COVER FEATURE **WEARABLE COMPUTING**

Can Disability Discrimination Law Expand the Availability of Wearable Computers?

Reuben Kirkham, Newcastle University

An exploration of current discrimination law reveals new opportunities for promoting the public use of wearable assistive technologies, which can be a powerful impetus for both wearable advocates and researchers exploring design alternatives.

Wearable computers have always presented novel legal challenges. For example, wearable devices to predict roulette spins became popular enough by 1985 that Nevada outlawed their use in casinos under a blanket statute banning any device for calculating probabilities[1] (revised in 2011 to "any computerized electronic, electrical or mechanical device which is designed, constructed, altered or programmed to obtain an advantage at playing any game."[2]) Fast forward to today when, in response to Google Glass and other forthcoming devices with head-mounted optical displays such as Microsoft HoloLens, legislation is pending in seven states to prohibit the use of such devices while driving (https://legisweb.state.wy.us/LSOResearch/2014/14FS004.pdf).

Google Glass and its successors are prompting debates on a range of legal issues—for example, the difficulty of proving that someone wearing a technology is also using it.[3] These debates are pushing wearable controversies into the public consciousness. Are those wearing such devices endangering others on roadways, pirating the movies they watch in theaters, or compromising the privacy of patrons in restaurants? Such questions have led to the preemptive ban of Glass and similar devices in various public places.

Many legal arguments are rooted in the notion of proportionality—namely, finding the right balance between the *collective* benefits and risks and the rights of *individuals* for a particular matter. However, legal debates about wearables have thus far centered only on their potential

harm, not on a user's right to wear the technology in public and the law's role in supporting that right.[3]

Disability discrimination law might be the basis for a legal route to expand and support wearables' use. The core principle is that people with disabilities have a right to reasonable accommodations that might ameliorate any disadvantages arising from their disability. Over the past two decades, this principle has been codified in many laws and regulations requiring organizations to take significant steps to assist people with disabilities, prohibiting direct discrimination against such persons, and permitting organizations to inconvenience those with disabilities only through a policy or practice that can be objectively justified in all circumstances.

As such, the application of disability discrimination law could greatly expand public acceptance and use of wearable assistive technologies (WATs). Just as disabled individuals have the right to take a guide dog into a shop or a taxi, they should also have the right to wear a WAT in public—they both achieve the same goal, which is to neutralize the disadvantages people face from their disabilities. Some WATs also assist those who prefer to conceal their disabilities from the general public, which helps make wearables permissible for those without disabilities as well. Moreover, auxiliary services such as Wi-Fi hotspots and indoor navigation systems might be provided automatically as a reasonable adjustment for WAT users, thereby leveraging the benefits of wearable technology in more scenarios.

Ultimately, disability discrimination law could induce changes in how wearables are created, enabling more design choices and socially oriented solutions to what many regard as purely technical challenges.

KEY EQUALITY PROVISIONS

Arguably, the advent of modern disability discrimination legislation began in 1990 with the Americans with Disabilities Act (ADA). This law requires both public and private institutions to make reasonable accommodations for a broadly defined range of disabilities, in a wide spectrum of circumstances, with special emphasis on public areas and workplaces (private homes are excluded). The ADA was followed by the UK's Disability Discrimination Act (DDA) of 1995 and similar legislation in other jurisdictions.[4] Global efforts to protect those with disabilities culminated in 2006 with the passage of the UN Convention on the Rights of Persons with Disabilities (UN CRPD), which 150 countries have since ratified.

"To promote equality and eliminate discrimination" against those with disabilities, UN CRPD signatories are obligated to "take all appropriate steps to ensure that reasonable accommodation is provided" (Article 5.1). The convention defines reasonable accommodation as "necessary and appropriate modification and adjustments not imposing a disproportionate or undue burden, where needed in a particular case, to ensure to persons with disabilities the enjoyment or exercise on an equal basis with others of all human rights and fundamental freedoms" (Article 2).

A key aspect of reasonable accommodation is accessibility:

> *To enable persons with disabilities to live independently and participate fully in all aspects of life, State Parties shall take appropriate measures to ensure to persons with disabilities access, on an equal basis with others, to the physical environment, to transportation, to information and communications, including information and communications technologies and systems, and to other facilities and services open or provided to the public, both in urban and in rural areas. These measures ... shall include the identification and elimination of obstacles and barriers to buildings, roads, transportation and other indoor and outdoor facilities, including schools, housing, medical facilities and workplaces, and to Information, communications and other services, including electronic services and emergency services (Article 9.1).*

This provision is ambitious in the rights it affords to those with disabilities, including rights related to technology use, and thus is a useful context for exploring the practical application of these rights. However, the convention is essentially a high-level framework, not domestic law, so understanding it requires reviewing its implementation in national legislation. As a UN CRPD–ratifying country, the UK has made strong inroads into implementing the convention into its domestic law, so its disability provisions are a sensible basis for exploring the UN CRPD.

UK DISABILITY PROVISIONS

The UK's Equality Act of 2010 (the DDA's successor) defines disability as "a physical or mental impairment [that] has a substantial and long-term adverse effect on [a person's] ability to carry out normal day-to-day activities" (Section 6). The UK government recognizes a wide range

TABLE 1. Key cases that address reasonable adjustment.

Case	Ruling	Reference
Royal Bank of Scotland Group Plc v. Allen (2009)	Court of Appeal (CoA) ordered a privately owned bank to spend £200,000 to install a platform lift in one of its branches, even though online services were available and the installation would involve the loss of an interview room.	www.bailii.org/ew/cases/EWCA/Civ/2009/1213.html
Leeds Teaching Hospital NHS Trust v. Foster (2011)	Employment Appeal Tribunal (EAT) awarded £50,000 to the appellant, who was suffering undue "stress" under a particular manager. The appellant claimed that stress was due in part to bullying and harassment. EAT ruled that the employer failed to redeploy the appellant even though it was highly unlikely that redeployment would successfully rehabilitate him in the workplace.	www.bailii.org/uk/cases/UKEAT/2011/0052_10_1406.html
Cordell v. Foreign and Commonwealth Office (2011)	EAT ultimately ruled that providing a team of lip speakers in response to the appellant's hearing impairment was not a reasonable adjustment, because the costs of more than £100,000 would be excessive.	www.bailii.org/uk/cases/UKEAT/2011/0016_11_0510.html
Croft Vets Ltd. and Others v. Butcher (2013)	EAT found that the employer was in breach of the reasonable adjustment duty because of its failure to pay for private medical treatment to treat the employee's depression—despite the availability of a free comprehensive public health service.	www.bailii.org/uk/cases/UKEAT/2013/0430_12_0210.html
FirstGroup Plc v. Paulley (2014)	CoA determined that a bus operator had no obligation to force—as opposed to strongly encourage—someone without a disability to leave a disabled space on a bus.	www.bailii.org/ew/cases/EWCA/Civ/2014/1573.html

of disabilities—mobility, visual, and hearing impairments as well as chronic medical conditions such as chronic fatigue syndrome and psychosocial disabilities such as autism or bipolar disorder (www.gov.uk/government/uploads/system/uploads/attachment_data/file/85010/disability-definition.pdf)—and rarely contests claims of disability. Most countries' legal systems have a similarly broad definition of disability.

Reasonable adjustments

The UK's equivalent terminology for reasonable accommodations is reasonable adjustments, which normally stem from a disabled person's specific request to change a criterion, policy, or practice to ameliorate one or more disadvantages arising from the disability or impairment. Reasonable adjustments include providing auxiliary aids and supports—for example, a sign-language interpreter for a deaf individual or an induction loop to support a hearing-aid user. They also give disabled employees the right to a wide range of assistive hardware to enable them to perform their occupational role and, to some extent, modify that role to accommodate their disability. Reasonable adjustments can include human support and training, or coaching for both the disabled person and the person's coworkers.

In line with the UN CRPD, the Equality Act relies on the proportionality principle to determine whether an organization is making reasonable adjustments on behalf of a person with disabilities. Table 1 lists examples of legal cases that deal with reasonable adjustments.

Proportionality in the individual case

Most UK jurisdictions adopt a proportionality test to determine if an individual's request for adjustment is reasonable. In *Leeds Teaching Hospital NHS Trust v. Foster*, the Employment Appeal Tribunal (EAT) awarded £50,000 to the appellant because the employer failed to make a reasonable adjustment even though the proposed action was not likely to improve the situation. In *Croft Vets Ltd. and Others v. Butcher*, the EAT also found that the employer was in breach of the reasonable-adjustment duty even when other appropriate support was available.

As these cases show, the reasonable-adjustment test is weighted heavily in support of the disabled person's right to be included in society; indeed, until *Cordell v. Foreign and Commonwealth Office* in 2011, the adjustment cost to large organizations appeared to be limitless. Nevertheless, with rare exceptions such as that evidenced by *Cordell*, the reasonable-adjustment duty has considerable strength, with individuals sometimes requiring dramatic and costly changes.

The anticipatory duty

The Equality Act also requires an organization to fulfill its anticipatory duty to those with disabilities by making reasonable adjustments without any consultation or advanced notice. The anticipatory duty is relevant when it is reasonably foreseeable that a disabled person is likely to suffer a disadvantage because of the disability and that adjustments in response to the

EUROPEAN CONVENTION ON HUMAN RIGHTS

The European Convention on Human Rights (ECHR), drafted in 1953 in the aftermath of World War II, expresses a wide range of fundamental human rights, including the right to life, a fair trial, and freedom of thought and conscience. The human rights more relevant to the use of wearable computing include the (qualified) right to privacy (Article 8) and the right to freedom of expression (Article 10), as well as the expectation of nondiscrimination in implementing the Convention (Article 14).

As an internationally binding Convention that covers all of Europe, including Russia, the ECHR is ultimately enforced by the European Court of Human Rights, which convenes in Strasbourg, France, and falls under the Council of Europe. However, the Convention is often implemented directly in the individual nations' domestic laws. For the UK, the implementing law is the Human Rights Act of 1998, and the Convention's emphasis on proportionality has effectively changed the way many judges interpret the Human Rights Act.[1]

The convention is also "a living instrument," in that its interpretation reflects changing times and circumstances.[2] In the technological domain, Strasbourg judges have responded to the introduction of technologies with rulings that focus on a subject's rights and personal data. As well as judgments relating to databases and biometrics,[3] rulings include *Google Spain v. AEPD and González*, which expanded the "right to be forgotten." In this 2014 case, Internet providers were expected to remove certain search results about individuals in order to make them more difficult to locate, resulting in much debate on balancing privacy and freedom of expression.

References

1. A. Baker, "Proportionality and Employment Discrimination in the UK," *Industrial Law J.*, vol. 37, no. 4, 2008; http://ssrn.com/abstract=1305007.
2. L. Wildhaber, "The European Court of Rights in Action," *Ritsumei Law Rev.*, vol. 21, 2004; www.ritsumei.ac.jp/acd/cg/law/lex/rlr21/wildhaber.pdf.
3. P. DeHert, "Balancing Security and Liberty within the European Human Rights Framework," *Utrech Law Rev.*, vol. 1, no. 1, pp. 68–96.

individual's specific request cannot mitigate the situation. For example, making a website or mobile application accessible to those with visual impairments would be an anticipatory reasonable adjustment, as would ensuring that a newly constructed building has an appropriately accessible physical or digital infrastructure.

In *Royal Bank of Scotland Group Plc v. Allen*, the reasoning was that a core part of the bank's service was face-to-face interaction and that the goal of disability discrimination law is to give those with disabilities the right to be included in society on an equal basis, rather than merely giving them access to the same overarching services as those without disabilities.

However, the anticipatory duty has some limits. In *FirstGroup Plc v. Paulley*, the Court of Appeal pointed out the potential impracticality of turning bus drivers into adjudicators and noted that the rights of a wheelchair user do not entirely trump those of people without disabilities. Even so, the court gave the appellant's needs as a disabled person considerable weight. Thus, in many respects, the anticipatory duty can be as strong as the obligation to respond to individual requests.

Other relevant provisions

The Equality Act of 2010 includes a range of other legal safeguards for persons with disabilities. These include protection from discrimination because of a disability (Section 15), the obligation of organizations to avoid indirect discrimination unless objectively justifiable (Section 19), and protection from victimization for anyone exercising or attempting to exercise their societal rights in good faith (Section 27).

An important example of these safeguards is the 2009 case *Dean v. Abercrombie & Fitch* (www.solicitorsjournal.com/news/employment/discrimination/abercrombie-failed-make-reasonable-adjustments-disabled-law-student), in which an employee objected to the employer's requirement that she cover her prosthetic arm with a cardigan. The court ruled that this action was unlawful harassment and awarded the claimant nearly £9,000. Extending this ruling to WATs, it might be considered unlawful for anyone to harass a WAT user, regardless of the assistive technology type.

The one clear proviso is that any WAT must comply with all other legislation. Notably, this includes the Data Protection Act of 1998, which means that users cannot use the device to

take photos in certain locations or store others' personal data. (The same point applies to WAT usage that might run afoul of comparable laws in other countries; for example, it is a felony in some US states to record conversations without permission.)

However, in the UK, the Human Rights Act of 1998, which implements the rights enumerated in the European Convention on Human Rights (ECHR) within the UK, has primacy over other domestic human-rights laws, including those pertaining to people with disabilities. Thus, any legislation impacting a disabled individual's human rights found to be incompatible with the Human Rights Act might have to be amended accordingly. As the sidebar "European Convention on Human Rights" explains in more depth, the ECHR rights rest at the pinnacle of a hierarchy of human-rights laws throughout most of Europe.

WEARABLES IN CONTENTIOUS APPLICATIONS

In keeping with the UK's reasonable-adjustments provision, guide dogs and even assistive miniature horses are allowed in public accommodations and on public transportation in a range of contexts. Wearable computers might benefit from this precedent, since in many applications they can help mitigate disabilities in ways that conventional adjustments cannot, as the sidebar "Wearable Assistive Technologies" describes.

If a disabled person chooses to publicly use a wearable in a potentially contentious way to support a specific disability, will the law support that use? If so, does this support apply only in the context of specifically developed WATs, or does it

WEARABLE ASSISTIVE TECHNOLOGIES

Various wearable assistive technologies (WATs) and applications are beginning to enter practical use. Perhaps the most prominent example is SenseCam, now commercially available as Autographer, which takes a photo at fixed intervals and stores the images for review as a memory aid by a user with, for example, amnesia or dementia.[1] Research shows public acceptance of SenseCam when used as an assistive technology,[2] although this attitude certainly could change if the device becomes more widespread.

Other prototype WATs are based on Google Glass, which has less acceptance in public places. Challenging ethical issues can arise in more extreme cases, including some legal concerns, such as the reappropriation of others' personal information as part of a social assistive system.[3] On the other hand, exploratory investigations with Parkinson's patients[4] did not encounter these issues; rather, researchers reported that this target group accepted Glass and felt it had genuine utility.

Still other prototype WATs are context aware and thus more privacy sensitive. These systems use inertial sensors to recognize discrete human activities. Examples include a system designed to track certain types of stereotypical behaviors in people with autism[5] and a system for the visually impaired[6] that uses a wristband for iPhone interaction instead of a touchscreen, enhancing the acceptability of assistive applications in social settings.

These emerging systems demonstrate the rich variety of research that will continue to develop as wearable computing gains popularity as an assistive technology.

References

1. S. Hodges et al., "SenseCam: A Retrospective Memory Aid," *Proc. ACM Conf. Ubiquitous Computing* (Ubicomp 06), LNCS 4206, 2006, pp. 177–193.
2. D.H. Nguyen et al., "Encountering SenseCam: Personal Recording Technologies in Everyday Life," *Proc. ACM Conf. Ubiquitous Computing* (Ubicomp 09), LNCS 4209, 2009, pp. 165–174.
3. R. Kirkham and C. Greenhalgh, "Social Access vs. Privacy in Wearable Computing," *IEEE Pervasive Computing*, vol. 14, no. 1, 2015, pp. 26–33.
4. R. McNaney et al., "Exploring the Acceptability of Google Glass as an Everyday Assistive Device for People with Parkinson's," *Proc. SIGCH Conf. Human Factors in Computing Systems* (CHI 14), 2014, pp. 2551–2554.
5. F. Albinali et al., "Detecting Stereotypical Motor Movements in the Classroom Using Accelerometry and Pattern Recognition Algorithms," *Pervasive and Mobile Computing*, vol. 8, no. 1, 2012, pp. 103–114.
6. H. Ye et al., "Current and Future Mobile and Wearable Device Use by People with Visual Impairments," *Proc. SIGCH Conf. Human Factors in Computing Systems* (CHI 14), 2014, pp. 3123–3132.

extend to more general wearables, such as Google Glass?

Privacy and copyright

Although some less-visible wearable technologies—cochlear implants, hearing aids, and Fitbits alike—face no real prospect of being banned from public use, more obvious wearables, such as Glass, have raised enough privacy and copyright concerns that many theaters, bars, and casinos have prohibited their use. To date, this proscription has gone largely unchallenged, but in light of reasonable-adjustment requirements, privacy and copyright violations might not be sustainable objections.

Privacy. Privacy is generally viewed as the most pressing wearables issue, with legal challenges arising as early as 2003.[5] Smartphones pose the same risk to the public of being recorded or photographed surreptitiously or without consent, so it seems unfair to prohibit disabled people from using wearables in the same environment.

The ECHR provides clear guidelines on how to balance privacy and disability, asserting that each citizen is entitled to "respect for his private and family life, his home and his correspondence" (Article 8.1) but also stating that privacy can be displaced to further the "protection of the rights and freedoms of others" (Article 8.2). Rights and freedoms within the ECHR's scope would include not discriminating against any ECHR-based implementation. Moreover, the ECHR's jurisprudence focuses primarily on privacy in homes and documents such as medical records, not public places. This focus is consistent with US law, which also tends to largely disregard individual privacy in public places.[6]

Modern data protection legislation, particularly in Europe, expects that personal data is processed lawfully and appropriately, and any system design would need to meet this expectation regardless of the information captured. However, contrary to common misconception, this legislation does not intend to make privacy a means for preventing the use of a properly configured wearable as an assistive technology.

Copyright. Copyright law is still largely based on the international Berne Convention. Although jurisdictions implement copyright law differently,[7] some aspects are handled consistently across jurisdictions; for example, both the US's fair-use doctrine and the UK's fair-dealing provisions allow the public to incidentally capture copyrighted material subject to identical concerns. Because copyright protection applies largely to the "substantial" portion of the material, only an extended recording would be a violation.

As the Court of Appeal noted in *Ashdown v. Telegraph Group Ltd.* (www.bailii.org/ew/cases/EWCA/Civ/2001/1142.html), the ECHR could transcend copyright legislation when appropriate. Thus, copyright law is much less likely than disability discrimination legislation to influence wearable permissibility. Moreover, the UK's Copyright and Rights in Performances (Disability) Regulations of 2014 is making copyright restrictions less problematic for people with disabilities by expanding the circumstances (including the range of disabilities; traditionally this exemption only applied to the visually impaired in very particular circumstances) in which someone has permission to copy a work without the copyright owner's explicit consent (www.legislation.gov.uk/ukdsi/2014/9780111112694).

Other copyright protections, such as digitally watermarking all content that a wearable device records or processes, would be sufficient and proportionate to deter the content's unlawful dissemination because the mark would identify the copier. Thus, a legal ban on an appropriately designed general wearable used as an assistive technology seems unlikely.

Use of generic WATs

Should the law protect WAT use if the device is generic, such as Glass, or only if it is tailored to support a specific disability? The main concern is that disability-specific devices with controlled use will create a serious cost barrier to WAT adoption and limit device customization. For example, before the iPhone, which now provides a unified environment-navigation medium in one physical device, the visually impaired had to use multiple bulky and expensive devices. Without permission to use generic wearables, users with other disabilities might have a similar problem.

Fortunately, there is little legal justification for requiring WATs to be disability-specific and for controlling their use on the grounds that more general technology might be used inappropriately. An exception might be the need to preserve confidentiality, in which case an employer might have the right to ban a personal wearable on its site. However, the employer would then need to fulfill its reasonable adjustment duty by providing (at its own expense) a separate but more restricted wearable for the employee's use on their premises.

> THE CASUAL OBSERVER SHOULD NOT BE ABLE TO DISCERN IF TWO USERS WITH THE SAME WEARABLE ARE DISABLED OR NOT.

Nonassistive use

Should wearable permissibility extend to those who might apply a WAT for reasons other than supporting people with disabilities? Any practical restriction on wearable use would require some proof of disability, much like disabled users mark guide dogs or carry special parking permits to identify their need.

Clearly marked assistive technologies can help identify a disabled individual in a way that enables others to make allowances for that disability,[8] and many disabled people deliberately identify themselves with markers so that they can readily access needed support. For example, a visually impaired individual might use a white cane to navigate a traffic crossing.

In other cases, disabled people might want to keep their disability confidential. For example, someone with a social disability might rely on a device to prompt them in social situations, or a person with Parkinson's disease might require a cuing system to help determine when to swallow. Making such devices noticeable negates one of the devices' core benefits and might embarrass users to a degree the law considers unusual and degrading.

Provisions in the UK's Equality Act strongly protect people whom others mistakenly treat as disabled individuals. For example, if a child was (falsely) perceived to have a learning disability and consequently was refused admission to a school or college, the child would have a claim even if she did not actually have that disability. By extension, a nondisabled person with a wearable that others could perceive as a WAT would likely have a case if someone attempted to obstruct its use. Thus, on occasions, the law effectively extends privileges designed for disabled persons to those without disabilities. In the US, Title II of the ADA deals with retaliation for any attempt to exploit the Act's provisions, so it also provides a potential route toward these extended privileges.

However, disability discrimination usually applies only to someone who appears to be using a WAT. It does not mean that nondisabled users have carte blanche to every wearable device feature in any environment. Rather, it means that the casual observer should not be able to discern if two users with the same wearable are disabled or nondisabled.

Eventually, the law might remove that concern by requiring all wearables to comply with rules imposed by the user's surroundings. The basis for this requirement might be a *virtual wall*[9]—a marker-based system that aims to enforce policies about a given technology's use in response to a specific environment. For example, the wearable, whether operated by a disabled or nondisabled user, would have to respect the virtual wall's direction to not store any video or audio data at that location because of the need to protect confidentiality or copyright.

The need to support a user's disabilities could override this mandate—for example, if a visually impaired user needed the system's camera in a theater to detect objects. For the most part, however, a virtual-wall breach would be considered unlawful and, in the theater example, not respecting that system would be a criminal violation of the UK's Copyright, Designs, and Patents Act of 1988.

The virtual wall would eliminate several concerns with nonassistive use: there would be little or no substantive legal objection to allowing wearable devices in more sensitive public places, and there would be a clear message that a wearable system is not a law unto itself, but rather a carefully controlled balance between both user needs and the needs of others within a given environment. Consequently, a virtual wall or other similar solution is likely to help system acceptance and result in better support for and inclusion of the disabled individual in society.

INFRASTRUCTURE FOR WEARABLES

Ramps, dropped curbs, and other building accessibility features are examples of the widespread accommodations (later known as reasonable adjustments) that began in 1963 as voluntary architectural concessions, largely in response to the popular *Designing for the Disabled*.[10] Such features are now prominently mentioned in disability discrimination law provisions.

A similar infrastructure is emerging in the digital realm, including the requirement for publicly available induction loops to assist hearing aid users. The ADA mandates assistive listening systems in a wide range of public places, and the UK's Equality Act states a similar expectation for service providers to provide induction loops in public places as appropriate. This expectation is emphasized in government guidance related to the Act (www.gov.uk/government/uploads/system/uploads/attachment_data/file/85011/disability.pdf).

Induction loop installation in public places is reasonable under UK disability discrimination law for several reasons:

- the cost is relatively small: only £300 to £500 for a simple

ABOUT THE AUTHOR

REUBEN KIRKHAM is a doctoral student in the Digital Interaction Group at Newcastle University. His research interests are the intersection of wearable and ubiquitous computing with disability legislation. Kirkham received an MSc in human–computer interaction from the University of Nottingham. He is a student member of ACM. Contact him at r.kirkham@newcastle.ac.uk.

installation, possibly with minimal additional maintenance costs;
- hearing-aid users, a significant subset of disabled people, gain a common and substantial benefit that helps address some of the disadvantages they face because of their disability;
- system implementation is standardized, so there is no need to install multiple different induction loop systems; and
- neither induction loops nor hearing aids impose a noticeable disruption or inconvenience to the establishment or its customers.

The point is simple. If developing a service that fulfills the reasonable-adjustment duty is both economically feasible and practical, and the service is not disruptive yet also uniquely benefits a sufficiently large group of disabled people, its installation is likely to be mandatory in most indoor public places. An explicit mandate for this provision could greatly assist the construction of an environment that supports all wearable systems.

Required technology. In a wearables infrastructure, the virtual wall is one essential technology. Two others are free quality Wi-Fi and indoor navigation systems.

Both have obvious rationales for disabled persons. With reliable Wi-Fi, for example, blind users can receive annotations of what they photographed,[11] or a person with learning disabilities can receive confidential prompts. The indoor navigation system would be useful to visually impaired users or perhaps to an individual with an anxiety disorder who needs to find familiar people in a crowded building.

If a WAT could demonstrate its support in cases like these, its service would probably meet the requirement for anticipatory reasonable adjustments, as long as it was consistent with public policy. Service provision will likely be incremental. The more frequent and varied the context is for device support, the stronger the obligation will be to provide that service and the more the undue burden test will apply to service providers. The ultimate goal would be to expand support systems for a wide range of wearables, benefiting both disabled and nondisabled users.

Accelerating availability. Technology firms such as Google are willing to amass huge costs to bring about or maintain laws and judicial policies suited to their particular business needs. In the recent case of *Hegglin v. Persons Unknown and Google Inc.* (www.bailii.org/ew/cases/EWHC/QB/2014/3793.html), in which the appellant effectively wanted to extend the right to be forgotten, Google presumably did not want to incur the (continuing) cost of removing the libelous search results the appellant complained of and spent nearly £1 million just in the pretrial stage. Companies with that much ready capital might fund the citywide trial deployment of a wearable support system and thus demonstrate the system's usefulness to those with disabilities, making its installation a reasonable adjustment. If the system and its deployment followed an open standard, then demonstrating the system's legal reasonableness in this way could accelerate its expanded use.

Thus, agreement on a sensible standard and the ability to demonstrate its reasonableness and compatibility with disability discrimination law might be the catalyst for rapid adoption of a standard for wearable systems. Together, the adopted standard and current development resources could produce globally available wearable support systems for disabled users far more rapidly than implementing narrowly tailored anticipatory adjustments, such as Web accessibility.

Disability discrimination law provides an opportunity to profoundly shape wearable computing and expand its use. The legal provisions outlined here are a route toward maximizing wearables' potential benefits for society in general, not just those with disabilities.

The need to address reasonable adjustments and to balance benefits, risks, and individual rights (proportionality) means that privacy, copyright, and confidentiality concerns are not likely to be strong enough to outweigh normal WAT use. Of course, a WAT user does not have unlimited rights; the system must provide lawful, not just effective, support. The limits likely lie where a technology might be used gratuitously on one hand, and where serious potential concerns arise from system use on the other. These limits will generate legal challenges, and the response will be the judiciary's determination of individual fringe cases or even the construction of new laws.

In designing the next generation of wearable computers, the research community should assume that wearable support systems, such as Wi-Fi and indoor navigation systems, will be available in a wide range of environments. Going forward, disability discrimination law would therefore provide a way to circumvent highly challenging technical problems.

Likewise, although a strong need remains to ensure that systems for people with disabilities are nonadversarial, it does not follow that a potentially socially offensive system should not be developed. Designers of wearable computers should be taking advantage of legal expertise to better inform system design and implementation, and avoid unduly deferring to privacy- and copyright-related concerns.

Finally, disability discrimination law remains the primary source for guidance on wearable use in public spaces. In the near future, evolving legislation could impede the availability and permissibility of a wide range of wearables. Developing and implementing wearables to fulfill the reasonable adjustment duty would require lawmakers to consider both aspects of proportionality when creating new legislation. Given the UN CRPD's forceful demands, it will be much harder for politicians to enact legal provisions that prevent or curtail the use of WATs that have been already deployed.

DISCLAIMER

This article is a high-level overview and thus not intended to be specific legal advice. Neither the author nor the IEEE accept any liability for use of the material herein.

REFERENCES

1. E.O. Thorp, "The Invention of the First Wearable Computer," *Proc. ACM Int'l Symp. Wearable Computers* (ISWC 98), 1998, pp. 4–8.
2. D.W. Schnell-Davis, "High-Tech Casino Advantage Play: Legislative Approaches to the Threat of Predictive Devices," *UNLV Gaming Law J.*, vol. 3, no. 2, 2012, pp. 299–346.
3. R. Kirkham et al., "Disability Law as a Driver for Innovation and Social Responsibility in Ubiquitous Computing," 2013; http://reubenkirkham.com/wp-content/uploads/2013/04/LegalDEVersion2013-CRversionrk-v2.pdf.
4. A. Lawson, *Disability and Equality Law in Britain: The Role of Reasonable Adjustment*, Hart Publishing, 2008.
5. S.S. Intille and A.M. Intille, *New Challenges for Privacy Law: Wearable Computers That Create Electronic Digital Diaries*, tech. report, Dept. of Architecture, MIT, 2003.
6. A. Kozinski, "The Dead Past," *Stanford Law Rev. Online*, vol. 64, no. 117, 2012; www.stanfordlawreview.org/online/privacy-paradox/dead-past.
7. R.M. Hilty and S. Nérisson, *Balancing Copyright: A Survey of National Approaches*, Springer, 2012.
8. K. Shinohara and J.O. Wobbrock, "In the Shadow of Misperception: Assistive Technology Use and Social Interactions," *Proc. SIGCHI Conf. Human Factors in Computing Systems* (CHI 11), 2011, pp. 705–714.
9. A. Kapadia et al., "Virtual Walls: Protecting Digital Privacy in Pervasive Environments," *Proc. 5th Int'l Conf. Pervasive Computing* (Pervasive 07), 2007, pp. 162–179.
10. S. Goldsmith, *Designing for the Disabled*, Royal Inst. of Architects, 1963; http://trove.nla.gov.au/work/11322065?selectedversion=NBD4435565.
11. J.P. Bigham et al., "VizWiz: Nearly Real-Time Answers to Visual Questions," *Proc. 23rd Ann. ACM Symp. User Interface Software and Technology* (UIST 10), 2010; http://dl.acm.org/citation.cfm?id=1866080.

Inferring Mobile User Status with Usage Cues

Jon C. Hammer and Tingxin Yan, University of Arkansas

An online inference engine monitors and extracts usage cues from wearables and identifies whether the user is busy, alone, happy, or stressed. Experiments show 85 percent identification accuracy with negligible energy cost.

Methods to understand and infer user context—a core research challenge for wearable computing—typically rely on sensory measurements to build inference models. Although sensors can effectively determine user context, such as physical activity, location, and vital signs, the energy cost to continuously collect and process sensory measurements is prohibitive.[1] In addition, sensors cannot reliably measure certain contextual elements, such as those with semantic meanings. For example, ambient sound sensing might accurately convey that a user is in a coffee shop, but it cannot always tell if the user is alone or actively engaged in a group.

With the increasing app centricity of mobile phones and wearables, researchers have an alternative to purely sensor-based context inference. As users spend more time with mobile devices and the technology continues to incorporate sophisticated interaction methods like voice control and notifications, usage cues—how people use and interact with wearables—become a more significant context source.

Usage cues complement sensory measurements, are inexpensive to retrieve and process, and uniquely correlate to semantic context. However, they can also be noisy, sporadically available, and highly user-dependent. Converting individual cues into a useful and reliable context source gives rise to three fundamental challenges:

- determining which usage cues are continuously available and useful for context inference,
- identifying the context patterns that can be reliably inferred from these cues, and
- measuring the cost of context inference based on usage cues.

In our study of usage cues as a basis for context inference, we determined that three usage cues are continuously available: application usage, application notifications, and connected networks. We also developed a generic learning framework that can infer a mobile user's status as busy, alone, happy, or stressed with more than 85 percent accuracy. Building on previous work,[2] we designed a learning framework that infers this information with an energy cost of less than 0.5 percent of total battery life when used continuously—a 90 percent reduction relative to the energy cost of

0018-9162/15/$31.00 © 2015 IEEE

TABLE 1. Three usage cues, or virtual sensors, and their features.

Virtual sensors	Features
Application usage	Session type (null or app), time between app sessions, trigger app and trigger type, session length (session time and number of unique apps), app category, and time spent in app
Application notifications	Arrival and response latency, response method, and meta information (owner app, arrival time, interarrival time, interdeparture time, response, response ratio, sender)
Connected network information	Cell ID, Basic Service Set Identifier (BSSID), connected or disconnected, enabled or disabled, time of day, day of week

hardware sensor–based approaches. These results show that usage cues can infer certain context elements with reasonable accuracy and reliability, serving as virtual sensors in combination with hardware sensors.

VIRTUAL SENSORS AND FEATURES

Current mobile operating systems record a rich set of usage information, including system events, services, and logs. Although our study focused on smartphones, the usage information recorded applies to wearable devices, such as smart watches and glasses, and correlates intuitively to user status for any mobile device. For example, videoconferencing implies a busy user, whereas gaming suggests that a user is alone. However, these usage cues are available only sporadically and are often inaccurate.

Our study goal was to find usage cues that have strong and consistent correlations to the logical status of mobile users. We considered application usage, notifications and user responses, and connected networks because these usage cues are universal across mobile devices and are extremely energy efficient to retrieve and process.

Table 1 summarizes the features within each usage cue, or virtual sensor, many of which are consistently available and extremely inexpensive to compute.

Application usage

Researchers have measured application usage as a way to infer a user's logical status, such as mood,[3,4] but a more difficult problem is how to extract a feature set from app usage patterns that is both reliable and useful in inferring user status. Specific apps, such as games, are not always used and therefore are not reliable features. We identified the features in Table 1 as more commonly available.

Session type. Consecutive apps used within a specific time period can be grouped into sessions.[5] We extended this idea by separating the time domain into slices that are based on screen-state transitions: screen-on and screen-off events. These events separate consecutive sessions and enable a natural grouping of mobile device use periods.

On average, the users we observed did not employ applications in 83 percent of the sessions, which is consistent with real-world scenarios, in which either the device is dormant or the user is checking widgets or time without opening any applications. These *null sessions* contrast with *app sessions*, the remaining 17 percent of the sessions, in which the user opens an app. Because a user is always in either an app or null session, session type is a consistently available feature, irrespective of current app use.

Time between app sessions. Although the screen is off most of the time, the time between consecutive app sessions is informative and consistently available. Small time-interval values indicate rapid use, common when responding to frequent texts or email. Large values indicate dormancy, perhaps when the user is at work or asleep.

Trigger app and trigger type. The *trigger app*—the first application in an app session—is an important feature.[5] We observed that an app session can start in one of two ways: either actively, when mobile users proactively use applications, or passively, when users are attracted by events, such as an incoming notification or call. Our study indicates that the usage patterns for active and passive sessions vary widely across users, which suggests different logical statuses. For example, only 15 percent of the sessions in our study were started passively, but a mobile user is at least 8 percent more likely to be busy in a passive session than in an active one.

Session length. We defined session length in terms of both session time and the number of unique applications used in one session. Session length could also be related to logical status. Intuitively, shorter sessions might correspond to quick glances at the phone, which indicate that the user is occupied with another task. Longer sessions might signal that the user is more actively engaged and less likely to be busy or with a group.

App category and time spent in app. In combination, these two features improve the chances of correctly inferring a user's status. For example, short bursts of time spent with productivity apps could indicate that the user is working and possibly stressed; long periods spent playing a game could indicate that the user is alone.

Application notifications

Users commonly interact with their devices through notifications, although other methods are emerging, such as voice and even gaze control. Our study focused on notifications, but we plan to extend our work to include these newer methods. Although many mobile apps frequently use notifications, they are not always available. We identified more consistently available features related to notifications as response time and latency, method of response, and meta information like sender type and number of accumulated notifications.

Arrival and response latency. With the notification's arrival and departure time, it is easy to calculate how long the user took to process the notification. If this value is very low, either the notification itself was significant, such as the receipt of an important email, or the user was unoccupied when the notification arrived. Larger values would imply the converse. We further partitioned response time into the *awareness point*—the time at which the user became aware of the notification—and the remaining time. *Arrival latency* is the time between the notification arrival and the awareness point, whereas *response latency* is the time from the awareness point to notification departure.

Intuitively, arrival latency corresponds with the responsiveness of mobile users to all notifications, irrespective of meta information such as the notifying app. In contrast, response latency corresponds with user responsiveness to individual notifications when meta information is known.

> THE CHALLENGE IS TO EXTRACT FEATURES FROM APPLICATION USAGE THAT ARE RELIABLE AND USEFUL IN INFERRING USER STATUS.

A long arrival latency might reflect a failure to recognize the device's notification alert, such as a text alarm, or that the user is deliberately ignoring the alert because of some preoccupation, such as a meeting. A short arrival latency might signal that the user is not preoccupied.

Response latency might be more helpful in refining predictions, because it could be viewed as the effective response time. Replacing total response time with effective response time could help reduce sensor noise, as there is no need to consider the time before the awareness point.

Response method. Users generally respond to a notification by tapping on and opening it or by swiping and dismissing it. Because no program can distinguish between a tap and a swipe, we wrote a program to monitor both notification events and applications in the mobile screen's foreground. If the monitor detected the same foreground application after a notification left the system, we reasoned that the user probably responded to the notification, and therefore a tap is more likely than a swipe. Similarly, if the foreground application is not the same as the notification, the user probably dismissed the notification with a swipe.

With this information, it is possible to record user response as either a tap or a swipe for all app notifications. We used either the raw swipe and tap counts or the ratio between them as the features.

Meta information. Certain other notification properties could be useful in inferring user status, such as the sender in messaging notifications. In general, several notifications from work colleagues might imply that the user is at work, whereas notifications from friends might indicate a social situation. The number of visible notifications is also useful, as an

accumulation of notifications might signal that the user is occupied.

Connected network information

The user's location is also a popular usage cue, but common methods to determine location, such as GPS and Wi-Fi, are energy intensive. Instead, we used connected network information, such as the Basic Service Set Identifier (BSSID) of Wi-Fi access points (APs) or cellular tower IDs, to indicate location. Without the need to scan the network, our connected network information incurred close to zero energy cost.

Recording changes in network status lets us know when a user has physically moved from a location. The lack of wireless connection is also a valid state, so connected network information is always available.

Combining connected network information and time of day enables the calculation of the user's logical location state as home, work, other, or null. Home is where the user is most likely to be in the middle of the night, and work is where the user is most likely to be during the morning and afternoon hours of a weekday. Other typically refers to an unknown or uncommon location, and null signifies no absolute knowledge of location, likely because there are no available Wi-Fi APs or cellular network information.

By observing which networks the user is connected to during a specific period, such as midnight or working hours, we can use the BSSID of the Wi-Fi APs to label the user's logical location. For example, if a device connects to an AP that it has frequently connected to during working hours, we label the device user's logical location as "at work." In addition, each radio on the device (cellular, Wi-Fi, or Bluetooth) has an enabled/disabled state that the user sets. Because an enabled radio can be either connected or disconnected, depending on network availability, the connected/disconnected state provides valuable insight into the user's context, which can help refine logical status predictions.

LEARNING FRAMEWORK

After we derived virtual sensors and their features, we created a framework to evaluate their usefulness in inferring the user's logical status. In designing the framework, we had two main challenges: how to preprocess raw features from continuous data streams and align them with potentially noisy training labels, and how to design a generic and low-cost online learning approach for logical status inference.

Preprocessing of virtual sensing data

Because our study emphasized logical user statuses, we needed input from mobile users to train our learning models. However, collecting a large amount of training data could be extremely intrusive, and the data could be erroneous because of missing values and entry mistakes.

Data from virtual sensors is typically heterogeneous and does not align well with time periods. Consequently, our first task was to separate sensing data into timeframes to smooth the transient noise and enable the use of learning algorithms that can handle discrete features. Our second task was to filter the training labels to remove errors.

Feature separation. The two logical separations of continuous mobile user data are sessions and moving windows. Sessions are intrinsically coupled to the screen state, while moving windows have no such limitation. Our empirical study indicated that a frame of 30 seconds provided the best results, which is the fixed framewidth our system adopts.

We aggregated all data collected within that grouping and used the aggregate as a single instance of the problem space. For example, we simply counted the unique applications within a given frame and recorded the most common state within that timeframe. Where relevant, we also used statistical aggregation types, such as minimum, maximum, and average.

Label filtering. To counteract the noise in training labels, we used the labels to train an empirical classifier and created a second set of labels by asking the classifier to predict the label for each training instance. We then gave both sets of labels associated confidence values. We assumed that confidence in the training labels (user-generated) would decrease over time and that confidence in the classifier's labels would remain static. We then compared the two labels—baseline (without label filtering) and filtered—for each instance and chose the filtered label as the one with the highest confidence.

Logical status inference

We used the processed data to infer the user's logical status by automatically selecting features and classifying them.

Automatic feature selection. When training the framework for a specific logical status, we did not know the optimal feature subset, so we included all possible features. Before making a prediction, the framework used algorithms

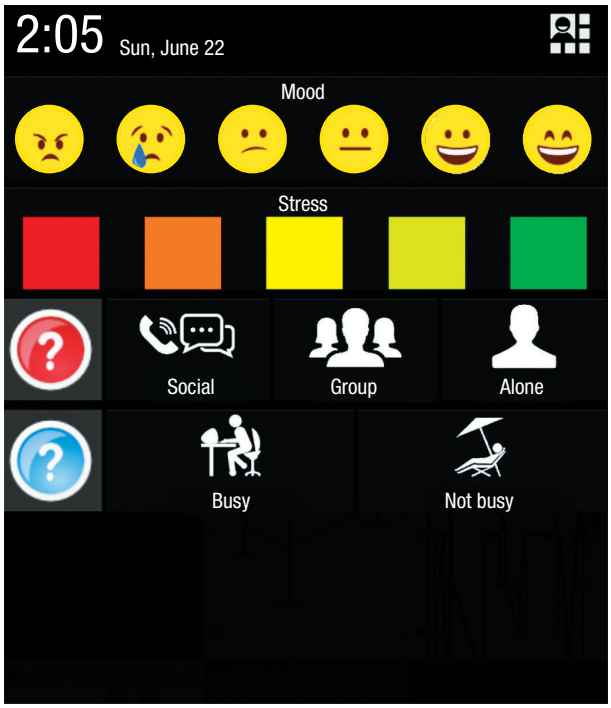

FIGURE 1. A user interface for logging training data. Note that the interface is notification based for minimum intrusiveness and the notifications are issued only when necessary.

based on principal component analysis (PCA) algorithms to reduce the feature candidates to only the most salient subset. Because the framework can calculate many logical statuses simultaneously, and each status has an optimal feature subset, it had to calculate only features that fell within the union of optimal feature subsets for all currently active logical statuses. If no logical status model used a feature, the framework did not calculate it, which saved both time and energy.

Feature classification. To perform the logical status inference, we needed to select a machine-learning classifier. We tested several learning classifier algorithms empirically by performing 10-fold cross-validation on our collected training data. These included nonsequential methods—random forests, K-NN, J48 decision trees, and support vector machines (SVMs)—as well as the hidden Markov model's sequential classifier. Of these, the J48 decision tree best balanced evaluation speed and predictive accuracy, so we used it for all our experiments.

Online learning

Feature extraction and learning algorithms can be computationally expensive, so a common practice is to offload work to a remote server, which incurs additional latency and raises privacy concerns. With some preparation, learning can be performed online on the device with an acceptable amount of overhead. Online learning is desirable because logical status changes frequently and models are highly user-specific.

We were able to use online learning because our framework met two criteria: feature extraction requires little computation, and the learning model needs updating only occasionally. We chose our features specifically because they do not require complex calculations (fast Fourier transform (FFT) or mel-frequency cepstral coefficient (MFCC) for microphone data), which reduce the work per instance. The need for only occasional model updates significantly reduces computational workload, although a single update might take longer. For many applications, updates are needed only a few times per day until enough data has been collected to make accurate predictions. The model can be trained online incrementally without significantly affecting either responsiveness or the device's battery life.

SYSTEM PROTOTYPING AND EVALUATION

We implemented our framework on the Android platform because we could record operating system information without jailbreaking the device. Our inference engine consists of an instrumenting service (IS) and a feature-extraction and inference module. The IS, a user app available from Google Play, collects both raw data from the virtual sensors and training labels from users. To make the labeling process less obtrusive, the app issues custom notifications that let the user log any logical status change. The feature-extraction and inference module classifies logical status.

To evaluate the framework's performance, we selected two critical usage features as microbenchmarks and then analyzed the framework's label filtering and inference performance overall as well as its energy consumption.

Participants and training data

Our evaluation is based on a trial study of nine participants over four weeks. All participants, one female and eight males, were college students and staff drawn primarily from the University of Arkansas Computer Science Department. Despite the small scale, we collected 65,000 operating system events and 2,500 user-generated labels to identify four logical statuses. Our aim was to show that our

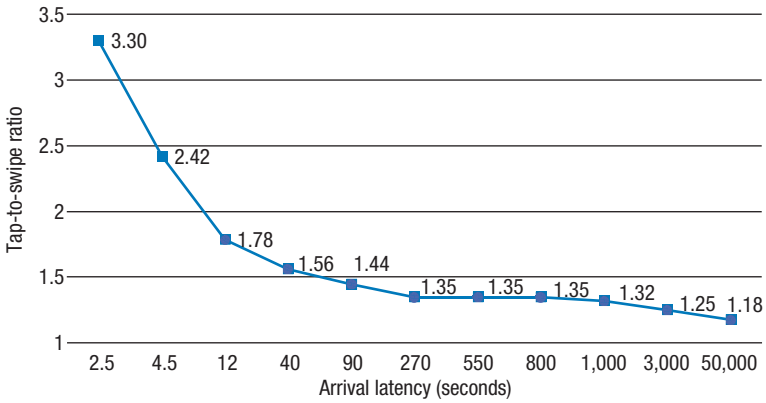

FIGURE 2. Correlation between the arrival latency of a notification and a user's response method as a tap-to-swipe ratio. A ratio of 1.0 implies no preference between taps and swipes. Because all values are greater than 1.0, users were generally more likely to tap on a notification than to swipe it away, regardless of arrival latency size.

method is feasible; we would need a larger-scale examination to demonstrate unbiased accuracy claims.

We used a custom logical-status survey application to collect raw usage data from the virtual sensors on the device and training labels from users. As Figure 1 shows, the application presents users with a choice of logical status in one of four categories, and the user simply taps on the appropriate icon for each one.

To ensure timely input from users, we designed the survey application as a notification that will be pushed to users periodically. To reduce the intrusiveness, the screen shows only pertinent notifications. The application records only the user's perception of a status; we cannot know whether a user who reported being happy is actually in that state.

Analysis of selected features

The two features we chose as microbenchmarks are session type (null or app) and application notification, including the relationship between arrival latency and response method.

Session type. We studied both null and app sessions and their correlation to a user's logical status. Because 83 percent of all sessions are null, using features only from app sessions will overlook the most usage information. We found that app sessions have a longer average duration than null sessions, which is consistent with the intuition that null sessions generally correspond to brief use, such as checking the time.

We then examined the correlation between logical status and percentage of null sessions. We found that 62 percent of null sessions were labeled as not busy, 76 percent as alone, 76 percent as happy, and 78 percent as not stressed. This observation indicates that null sessions are correlated to these logical states and should be integrated into the framework.

Application notification. Figure 2 shows the results of analyzing the relationship between arrival latency in seconds and the ratio of taps to swipes for notifications. The ratio decreases monotonically and exponentially with arrival latency. When the arrival latency is very short—2.5 seconds or less—the user is three times more likely to tap on a notification than to swipe it. With large arrival latencies, a ratio of 1.18 implies that it is much more difficult to predict a given notification's response. This correlation could inform the learning algorithm as it attempts to predict logical status.

Framework performance

To evaluate the framework's performance as a whole, we used a J48 decision tree to cross-validate our training data, assuming that all ground truth came directly from user-provided labels. Figure 3 shows our validation results. Each logical status has four bars, corresponding to three settings with various session types and one setting with fixed-width timeframes. Each bar is divided into a lower segment, which represents baseline accuracy (a classifier that always predicts the majority class), and an upper segment, which represents accuracy with our method.

As Figure 3 shows, our frame-based method outperforms the baseline approach by a large measure, improving accuracy by 10 percent overall. Information from null sessions shows a 3 to 10 percent increase in classification accuracy, which makes sense as most sessions are null. It also implies that information from null sessions is more useful in feature classification.

Figure 3 also shows that null and app sessions together result in slightly less accuracy than null sessions alone. The implication is that app session information might be detrimental to predictive accuracy and that app usage alone is a versatile feature that is not available most of the time. Finally, we found that fixed-width frames typically perform better than any session-based approaches because timeframe-based approaches align training labels and remove random noises by discretizing features to better align with timeframes.

Evaluation of energy efficiency

We ran our framework on Android devices with the screen turned off. Empirical results show that our framework consumed only 1 mW and 0.5 percent of battery life. This battery-life

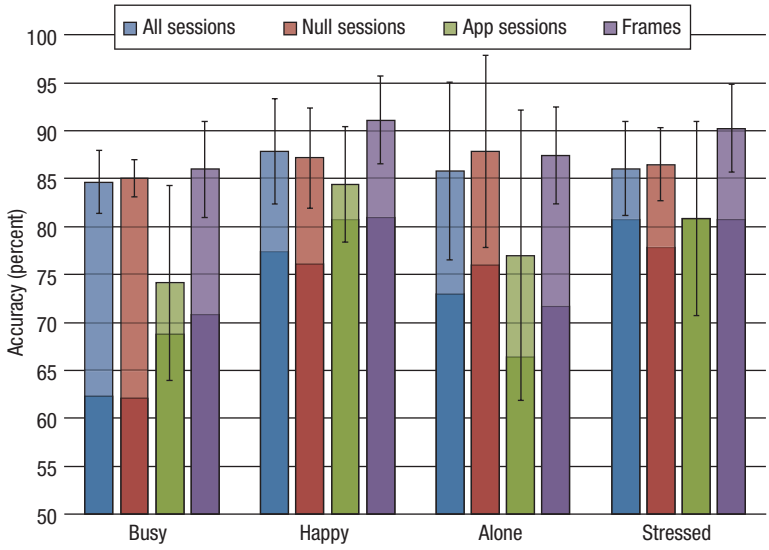

FIGURE 3. Predictive accuracy of each logical status in null, app, and all sessions and with fixed-width timeframes (frames). The vertical lines at the tops of the bars denote one standard deviation.

percentage was even lower when the screen was turned on, as the screen required much more power. In that sense, the 0.5 percent could be construed as a worst-case measurement.

This energy cost is several orders of magnitude less than what sensors incur, with accelerometers using 96 mW and GPS using 623 mW. Power usage this low means that our framework could run continuously without seriously degrading the user experience or drastically affecting the device's battery life.

Our logical status inference framework shows the feasibility of basing context inference on usage cues. With an energy consumption orders of magnitude smaller than that of hardware sensors, a hybrid sensing approach—usage cues and hardware sensing—could decrease overall energy cost and enable continuous context sensing. Fusing software and hardware sensing data into the same context inference framework could also boost the accuracy of context inference.

ACKNOWLEDGMENTS

This study is approved and monitored by the Information Research Board of the University of Arkansas under protocol 13-09-068.

REFERENCES

1. H. Lu et al., "The Jigsaw Continuous Sensing Engine for Mobile Phone Applications," *Proc. 8th ACM Conf. Embedded Networked Sensor Systems* (SenSys 10), 2010, pp. 71–84.
2. J.C. Hammer and T. Yan, "Exploiting Usage Statistics for Energy-Efficient Logical Status Inference on Mobile Phones," *Proc. IEEE Int'l Symp. Wearable Computers* (ISWC 14), 2014, pp. 35–42.
3. R. LiKamWa et al., "MoodScope: Building a Mood Sensor from Smartphone Usage Patterns," *Proc. 11th ACM Int'l Conf. Mobile Systems, Applications, and Services* (MobiSys 13), 2013, pp. 389–402.
4. K.K. Rachuri et al., "EmotionSense: A Mobile Phones Based Adaptive Platform for Experimental Social Psychology Research," *Proc. 12th ACM Int'l Conf. Ubiquitous Computing* (UbiComp 10), 2010, pp. 281–290.
5. T. Yan et al., "Fast App Launching for Mobile Devices Using Predictive User Context," *Proc. 10th ACM Int'l Conf. Mobile Systems, Applications, and Services* (MobiSys 12), 2012, pp. 113–126.

ABOUT THE AUTHORS

JON C. HAMMER is a doctoral student in computer science in the Department of Computer Science and Computer Engineering at the University of Arkansas. His research interests include energy-efficient context sensing and mobile computing. He is a member of IEEE. Contact him at jhammer@uark.edu.

TINGXIN YAN is an assistant professor in the Department of Computer Science and Computer Engineering at the University of Arkansas. His research interests include ubiquitous computing, energy-efficient mobile systems, and mobile crowdsourcing. Yan received a PhD in computer science from the University of Massachusetts. He is a member of IEEE. Contact him at tingxiny@uark.edu.

COVER FEATURE **WEARABLE COMPUTING**

Beeping Socks and Chirping Arm Bands: Wearables That Foster Free Play

VIDEO

Andrea Rosales, Open University of Catalonia
Sergio Sayago, University of Lleida
Josep Blat, Pompeu Fabra University

Three playful wearable accessories incorporate design principles that promote individuality, natural interaction, ubiquity, and intimacy to encourage spontaneous open-ended social interaction in school-age children.

Digitally augmented toys and play areas have the potential to encourage play that is spontaneous, collaborative, open-ended, physical, and social. Sensors can add a new dimension to traditional fantasy and active play as well as social games such as hopscotch. Experts agree that free play is essential in cognitive and emotional development, helping children learn how to evolve and enforce rules and improve their social interaction.[1]

Despite these benefits, free play is declining—in part because of the increased reliance on screen-based entertainment and extracurricular activities.[2] Even attempts to broaden structured games through the use of sensing and reacting technologies fall short of exploiting the full range of free-play opportunities. The sidebar "Technology in Free Play" describes some examples.

Incorporating these technologies into fixed areas or objects prevents children from engaging in free play incidentally in the doctor's waiting room or when shopping with parents, for example. Typically, children engage in structured and often sedentary activities at this time, such as drawing on a tablet.

In contrast, incorporating these technologies into wearables gives children more freedom to create and explore largely independently of their environment. For example, shaking the wrist in different ways might change a colored light or produce a unique sound. This freedom aligns with the idea that children see play opportunities in all kinds of things—a twig, fabric, a paper bag, and so on—which in turn immerses them in highly personalized physical and creative challenges. The potential of wearables is intuitive, but to date few

TECHNOLOGY IN FREE PLAY

Digitally augmented play areas and toys aim to encourage open-ended social play. For example, the Cardboard Box Garden, a sound-based collection of cardboard boxes, stimulates discovery, play, and adventure.[1] Toinggg consists of three interactive trampolines that provide different and exchangeable animal sound feedback.[2]

Wearables are relatively new to the digital toy market. Examples include TagURIt, which uses proximity sensors to detect when the chaser reaches a target and, through a Lumalive display on the chaser's shirt, makes a token appear on the chaser's chest.[3] Another wearable, Dancing in the Streets (DITS), is based on *Dance Dance Revolution*, a popular arcade game. DITS includes three-axis accelerometers, which players wear as anklets. DITS then sends dance movement information to a smartphone, where they become graphics as part of a game.[4]

Although both TagURIt and DITS encourage movement and present children with physical challenges that could enrich their free play, these wearables' benefits have not yet been studied. Consequently, it is not clear how they encourage free play or how their qualities relate to free-play opportunities.

References

1. K. Ferris and L. Bannon, "'...a Load of Ould Boxology!'" *Proc. Designing Interactive Systems* (DIS 02), 2002, pp. 41–49.
2. B. Van Hoeve, L. De Valk, and T. Bekker, "Toinggg: How Changes in Children's Activity Level Influence Creativity in Open-Ended Play," *Proc. Conf. Advances Computer Entertainment* (ACE13), 2013, pp. 642–645.
3. S. Cheng, "TagURIt: A Proximity-Based Game of Tag Using Lumalive e-Textile Displays," *Proc. ACM SIGCHI Conf. Human Factors in Computer Systems* (CHI 11), 2011, pp. 1147–1152.
4. J. Clawson, N. Patel, and T. Starner, "Dancing in the Streets: The Design and Evaluation of a Wearable Health Game," *Proc. Int'l Symp. Wearable Computers* (ISWC 10), 2010; doi: 10.1109/ISWC.2010.566584.

FIGURE 1. Children playing blind man's bluff with Statue. The sound feedback guides the blindfolded child, making the game different and challenging as the other children devise new strategies to keep their location hidden.

studies have examined to what degree they promote the characteristics of free play.

To better understand this relationship, we worked closely with school-age children to design three wearable playful accessories (PAs) in succession, using lessons from the previous PA design to inform the next one. The resulting PAs—Statue, FeetUp, and Wearable Sound (WS)—function primarily as clothing, but their sensing and reactive technologies are responsive to children's movements.

To test our PAs, we conducted a series of user studies at after-school centers and an electronic arts festival (WS only), and assigned observers at each session to take notes on how the children played with the three devices. We then analyzed and classified the observers' field notes to identify

- what wearable qualities foster positive free-play experiences;
- how Statue, FeetUp, and WS support these qualities; and
- how wearable mechanics and characteristics relate to free-play dynamics.

As we incorporated lessons from these studies and from our consecutive PA designs, four wearable qualities emerged as significantly encouraging free play—individuality, natural interaction, ubiquity, and intimacy—which can be the basis for designing products that foster rich and diverse play experiences.

WEARABLES IN PLAY

Our user studies consisted of multiple play sessions with varying numbers of children and at least two researchers and one monitor or teacher observing. After each session, all the researchers reviewed the observers' notes and individually categorized and classified them using an iterative thematic analysis[3] to

identify wearable qualities and free-play experiences.

Statue

Statue is a belt pouch with a Lilypad microprocessor, an accelerometer, LEDs, and a piezo speaker. When the child moves, Statue emits sounds and lights up. It uses the accelerometer's horizontal *x*-axis and diagonal *z*-axis to detect movement, but ignores the vertical *y*-axis. If wearers move slowly enough vertically, they can beat the system and avoid giving their location away with sound, which gives them an advantage in certain games. We designed Statue to augment games related to being a statue, such as freeze tag, but children were free to create their own games as well.[4]

Study. We conducted two sessions in two after-school centers with a total of 24 school-age children divided into four groups. One session had 16 children, while the other had 8, and both sessions had two researchers and two teachers observing.

Results. Observers reported that all children integrated Statue quickly into their free play, taking less than a minute to incorporate Statue into a game. Each group played two or three games, but no group chose the same games, which ranged from statue-related (blind man's bluff, freeze tag, red light/green light, musical statues, and serious faces) to adapted versions of everyday games with statue rules (hide-and-seek, cops-and-robbers, Simon says, and slow race). Figure 1 shows children using Statue in a blind man's bluff game.

Statue allows free-play scaffolding, as observed in the sessions; children moved from familiar structured games to creative alternatives based on social skills. For example, Statue helped children practice negotiating rules for group behavior when a group had to invent new ways to fool the blindfolded child in blind man's bluff. Statue also appeared to encourage fantasy games. For example, when they heard its computer-like sound, two children began to move like robots.

FIGURE 2. Two children playing with FeetUp. Although much free play was individual, children tended to share their experiences, in this case making sounds in a certain way by pushing their feet off the wall or the floor.

FeetUp

FeetUp is a pair of socks similar to slipper socks that emit sounds and light when both the wearer's feet are completely off the ground, as in jumping or lying down, with no pressure on either foot.[5] The Jeenode microprocessor's RF system works with pressure sensors in the sock's sole to detect the jump, and the LEDs and piezo speaker provide audiovisual feedback.

Study. As with Statue, two researchers and two teachers observed 24 school-age children, but this time we divided them into groups across two after-school centers. Most children were new to the study, although a few from the Statue study participated. One of the centers was a new participant.

Results. Observers confirmed that FeetUp encouraged the use of gross

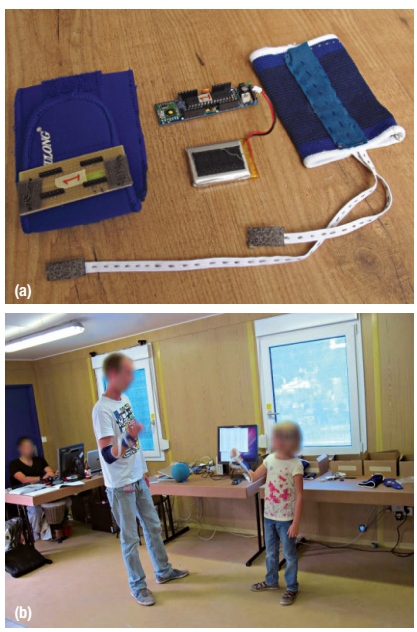

FIGURE 3. Wearable Sound (WS) components. (a) WS includes an elastic band with a bend sensor that can be attached to the elbow, the wrist, or the knee, as well as a Velcro band that attaches to the elastic band through conductive elastic and Velcro and with a computer through an RF-equipped microcontroller. The Velcro band includes the Jeenode microcontroller, also with RF. Users press buttons on the Velcro band to select one of 30 sounds, such as chirping birds. (b) The computer plays the sound selected by each user according to his or her movements.

FIGURE 4. Playing with WS. (a) In kinetic play, a dad and his daughter invent a game to create sounds by tapping on each other's heads, while (b) in fantasy play, two girls imitate octopus movements to the sound of bubbles.

motor skills. Most children started with individual play, such as jumping while dancing or doing gymnastics. Three children taught their groups routines—including capoeira (a traditional Brazilian dance), breakdancing, and ninja moves. Although these routines were not as popular as individual play, the others in their groups stopped playing their games, first to watch and then to practice the new activity. The other three groups played by imitating each other, or exploring odd movements and sounds as in Figure 2.

Some children used FeetUp's feedback to prove their jumping skills, counting the number of beeps while jumping. Observers noted that all children integrated FeetUp in free-play activities, and even though much play was individual, the children shared their jumping activities and created games that involved others.

WS

As Figure 3 shows, WS is a wide elastic armband that responds to bends with up to 30 familiar sounds, such as engine noise or chirping birds. WS's fabric-based bend sensor and Jeenode microprocessor interconnect through elastics and Velcro, both of which are conductive. Unlike Statue and FeetUp, WS uses a PC as well as the wearable. Users select the desired sound from the band, the microprocessor sends the choice as an RF signal to the PC, and the PC uses the WS software patch[6] to play the sound each time the user bends an arm or leg (depending on band location).

Study. Rather than observe WS in an after-school center, we made it available at the 2012 Ars Electronica Festival, an annual event in Linz, Austria, to view cutting-edge computing applications in a variety of disciplines. More than 250 people, including 100 children, played with WS. Different groups, typically children and their parents, played for various durations and in a range of ways. Three researchers and one monitor took notes—including participants' age, gender, choice of band position (arm or leg), sound preferences, and time spent in play—and videotaped the groups as they used WS. Unlike the after-school center studies, these sessions were open-ended.

To normalize the data for analysis, we chose a random subset of 30 school-age participants and used a play observation scale[7] to code their behavior according to type of activity performed every 10 seconds. Through the play observation scale, we could measure whether or not the child was playing with WS, interacting with the computer, or exploring the device, as well as identify whether play was gross motor, rhythmic, or fantasy. We used 10 video coders, and validated the coding scale with a reliability test.

Results. Within our participant subset, WS encouraged free play among 66 percent (the remaining 34 percent did not interact with WS or the computer). Of those, 85 percent engaged in creative kinetic play, as shown in Figure 4a, and 65 percent engaged in fantasy play. In Figure 4b, for example, the children are pretending to be an octopus. Kinetic play included dancing, jumping, tumbling, or just imitation.

Most participants wore WS on the elbow, most likely because it is used in many histrionic movements. Exploring the different sounds challenged the children's imagination and allowed them to personally connect with the PA.

CORE DESIGN PRINCIPLES

In designing our PAs, we used a research-through-design process, in which researchers describe the problem and how their perception of a preferred state changes over time.[8] Our preferred state was that wearables promote free play, and from our observations we connected four key wearable qualities to free-play experiences. We saw that exploiting the *individuality* of interaction promotes children's creativity and spontaneous social interaction,

FIGURE 5. Natural interaction based on histrionic arm movements. Two boys use WS while pretending to (a) play a guitar and (b) drive a car. Natural interaction is integral to fantasy and kinetic free play.

building on *natural interaction* to support common play movements. *Ubiquity* enables this spontaneous interaction and constant use creates *intimacy*, a strong connection between user and device. To various degrees, all three of our PAs supported these core principles.

Individuality

Wearables encourage individuals to explore the device at their own pace and in their own way. In free play, this supports the natural transition from solitary to associative play[9] or collaboration, which is foundational to free play.

Evidence in our PAs. Although interaction with all three PAs was individual, children built on that individuality to create social games, defining their own interaction rules to create a wide range of social free-play activities. The children wearing Statue created new rules for blind man's bluff, for example, and evolved strategies based on Statue's feedback.

We also experimented with interconnected WS play and found that it actually reduced play value. Both FeetUp and WS use networking but not to interconnect users. Interaction remains individual; the children dictate the degree of social interaction. In a side experiment, we compared an individual WS version with one in which at least two participants had to interact with WS simultaneously or it would not emit sounds. Participants were in 12 groups of 2 to 4 per group. In contrast to the wide range of game types with the individual version (sensory, fantasy, construction, and challenge), the interconnected version showed a predominance of challenge games. Having one dominant game type reduces the play value and hence the potential to encourage free play because play value strongly relates to the ability to enable a variety of games.[10]

Lessons for designers. Taking advantage of the wearable object's individuality supports the natural transition from individual to parallel and group activities.[9] We recommend moving from digital networks that define basic interaction rules among players to enabling children to create all their interaction rules on the basis of individual device feedback. The games they create define the social interaction rules.

Individual interaction is concerned with the play between user and device as the starting point of group play. Traditionally, digital augmentation has focused on shared objects to encourage free play, but our studies show that interacting individually also promotes group play.

Natural interaction

Natural interaction involves communicating with the device through natural movements,[11] which in games means through common play movements such as pretending to be something else or suddenly standing still. It is not necessarily individual interaction, since users remain physically and socially active in common play activities. Because it encourages wearable integration in everyday games, natural interaction is an important wearables design consideration.

The degree and quality of support for natural movements depends on the wearable's sensors and actuators. Sensors on the skin can detect body state, motion, and even some feelings, but they can be uncomfortable or intrusive; other sensors and actuators detect vibrations far less intrusively.

Evidence in our PAs. Input to all three PAs came from body movements

through an accelerometer in Statue, a pressure sensor in FeetUp, and a bend sensor in WS. Children readily integrated Statue, however, because pretending to be a statue is part of many popular games. WS interaction supported histrionic arm movements, which is a common element of fantasy play. Figure 5 shows two examples.

By understanding how natural interaction supports games, we can see the extent to which a device encourages free-play games. Statue tended to support primarily rule-based and open-ended games, FeetUp supported mainly kinetic games, and WS encouraged both kinetic and fantasy games.

Lessons for designers. Supporting natural interaction by augmenting common play behaviors with the use of body sensors and unobtrusive actuators encourages children to move freely and to keep eye contact with their peers.

Wearables are causing a design shift from software- and hardware-based user interfaces to ways of augmenting natural play behavior, so sensors in these devices should support both movement and sound, for example. We used a bend sensor rather than a pressure or motion sensor because bend sensors, made of antistatic fabric, allow for more flexible interaction. Children can thus readily integrate natural free-play movements in the augmented experience. In our studies, bend sensors supported rhythmic and fantasy play, as well spontaneous movement not tied to any game.

> ADDING SENSING AND REACTIVE TECHNOLOGIES TO GARMENTS CAN EXTEND THEIR FUNCTIONALITY AND POTENTIAL FOR CREATIVE USE.

All three PAs provided sound feedback, which is central to natural interaction, and Statue and FeetUp also provided visual feedback as lights. Despite being widely regarded as the predominant channel, visual feedback was less relevant in our studies than sound. The children could not look at their blinking socks while jumping or track all the blinking Statues while running. Participants' comments included "I couldn't see the blinking but we all heard the sound" and "You can hide the lights but you still have the sound." The relevance of sound feedback encouraged us to focus on movement-to-sound interaction in developing WS.

Ubiquity

Wearables should be able to go anywhere and be used any time. Adding sensing and reactive technologies to everyday garments or accessories extends their functionality and provides the potential for creative use.

Although ubiquity is not an explicit interactive quality, it promotes interaction by enabling spontaneous use at unexpected moments. When children are without structured entertainment, they often complain about being bored. The ubiquity of wearable PAs enables boredom to become creativity. In this sense, wearables can help reverse the free-play decline.

Evidence in our PAs. All three concepts of our PAs could be applied to either accessories (Statue) or clothing (FeetUp and WS). Although we designed WS as a band, it is fabric based, so conceptually it could be worn as a shirt. Regardless of the form, when the PAs are worn, they are ubiquitous and go from ordinary clothing to toys when the child is bored.

Lessons for designers. The ubiquity of wearables enables the augmentation of everyday objects through the use of secondary playful options.

Intimacy

Users become personally connected to their devices because feedback is based on their unique characteristics. Sensors, for example, measure only that user's heart rate and calories burned, usually with some common goal like health or fitness. The device can be even be a coach, as in the popular fitness bands that measure miles walked and send encouraging messages and rewards. This direct personalized interaction builds intimacy.

According to Jean Piaget, the development of intelligence rests on exploiting individual creativity through reflective abstraction—creativity that stems from an individual's own interactions with objects.[12] Body sensors give children a new way to interact with their body, quantifying its functions for games and competition, and thus enabling it to be a play object. Feedback connects children with their own actions and that connection can

elicit positive emotional responses, such as pride and pleasure, which can trigger creativity.

Another aspect of intimacy is continued use over a long period, but our studies did not explore that idea.

Evidence in our PAs. In our studies, the user's connection with a quantified self was evident in both WS and FeetUp. Users chose sounds that pleased them and then performed actions to elicit those sounds. FeetUp's beeping with each jump gave them a sense of pride in the number of jumps they could make or pleasure when they could move their feet in a certain rhythm. WS provides a palette of sounds to explore the intimate connection with body movements. This leads to a large diversity of games; for instance, one boy used WS to pretend to be a bear with the sound of thunder.

Statue also provided a certain level of intimacy because, in certain games, children connected its sounds with being discovered. They enjoyed finding out that moving slowly up and down could keep them hidden.

Lessons for designers. By exploiting the intimate connection between wearables and children's bodies by involving bodily information that the children cannot quantify accurately in their interactive experiences,[13] they can be motivated to engage in "reflexive abstraction," which is key to encouraging creativity.[12]

Although sports, health, and fashion are the most common wearable applications, wearables that encourage free play can enhance child development and possibly become part of a larger proactive move to stop the decline of free play. Children's lives are increasingly tightly scheduled, and they are increasingly turning to a digital screen for entertainment. Expanding free-play opportunities through wearables can provide an enticing alternative.

Our design principles are just a stepping stone to more in-depth studies on wearable qualities. When, for example, might one quality, such as intimacy, dominate the play experience? We also believe that additional research efforts should heavily involve school-age children; we could not have identified the ways that wearables relate to free play without putting the children first in our PA design and development. Wearable toys and free play must be studied outside the laboratory.

Finally, our results might extend to other groups of wearers, not just children. Teenagers and adults, for example, could also benefit from wearables that support their interest in play and social interaction through technology.

Indeed, a deeper understanding of how wearable mechanics, interaction quality, and everyday experiences intertwine can help improve the everyday lives of any individual, regardless of age, gender, or physical and social abilities.

ABOUT THE AUTHORS

ANDREA ROSALES is a postdoctoral researcher in the Internet Interdisciplinary Institute IN3 at Open University of Calonia, Spain. Her research interests include user studies of wearable and playful digital interaction with children and older people. Rosales received a PhD in information communication and audiovisual technologies from Pompeu Fabra University, Spain. Contact her at arosalescl@uoc.edu.

SERGIO SAYAGO is a visiting tenure-track lecturer in the Department of Computer Science and Industrial Engineering at the University of Lleida, Spain. His research interests include human–computer interaction, the elderly population, and ethnography. Sayago received a PhD in computer science and digital communication from Pompeu Fabra University. He is a member of the Spanish Human–Computer Interaction Association (AIPO). Contact him at sergio.sayago@diei.udl.cat.

JOSEP BLAT is a professor in the Department of Information and Communication Technologies at Pompeu Fabra University and head of its Interactive Technologies group. His current research interests are 3D graphics, human–computer interaction, and technology-enhanced learning. Blat received a PhD in mathematics from Heriot-Watt University, Scotland. Contact him at josep.blat@upf.edu.

REFERENCES

1. L.E. Berk, *Child Development*, Pearson, 2006.
2. D.G. Singer et al., "Children's Pastimes and Play in Sixteen Nations: Is Free-Play Declining?," *Amer. J. Play*, vol. 1, no. 3, 2009, pp. 283–312.
3. V. Braun and V. Clarke, "Using Thematic Analysis in Psychology," *Qualitative Research in Psychology*, vol. 3, 2006, pp. 77–101.
4. A. Rosales, E. Arroyo, and J. Blat, "Evocative Experiences in the Design of Objects to Encourage Free-Play," *Proc. Ambient Gaming Workshop, Int'l Joint Conf. Ambient Intelligence* (AmI 11), 2011, vol. 277; doi: 10.1007/978-3-642-31479-7_39.

5. A. Rosales, E. Arroyo, and J. Blat, "FeetUp: A Playful Accessory to Practice Social Skills through Free-Play Experiences," *Proc. IFIP Conf. Human-Computer Interaction* (Interact 11), 2011, pp. 37–44.
6. A. Rosales et al., "On the Evocative Power and Play Value of a Wearable Movement-to-Sound Interaction Accessory in the Free-Play of School Children," *Ambient Intelligence Smart Environments*, vol. 6, no. 3, 2014, pp. 313–330.
7. K.H. Rubin, "The Play Observation Scale (POS)," 2001; www.rubin-lab.umd.edu/CodingSchemes/POS Coding Scheme 2001.pdf.
8. J. Zimmerman, E. Stolterman, and J. Forlizzi, "An Analysis and Critique of Research through Design: Towards a Formalization of a Research Approach," *Proc. Designing Interactive Systems* (DIS 10), pp. 310–319.
9. M. Parten, "Social Play among Preschool Children," *J. Abnormal Social Psychology*, vol. 27, 1933, pp. 430–440.
10. B.M. Kudrowitz and D.R. Wallace, "The Play Pyramid: A Play Classification and Ideation Tool for Toy Design," *Arts Technology*, vol. 3, no. 1, 2010, pp. 36–56.
11. N. Bianchi-Berthouze et al., "Does Body Movement Engage You More in Digital Game Play and Why?," *Proc. 2nd Int'l Conf. Affective Computing and Intelligent Interaction* (ACII 07), LNCS 4738, 2007, pp. 102–113.
12. J. Piaget, "Creativity," *The Learning Theory of Piaget and Inhelder*, J.M. Gallagher and D.K. Reid, eds., Brooks/Cole Pub., 1981, pp. 221–229.
13. L. Pagliarini and H.H. Lund, "Wearing the Playware," *Artificial Life Robotics*, vol. 15, no. 4, 2011, pp. 381–384.

See **www.computer.org/computer-multimedia** for multimedia content related to this article.

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

STC 2015

The 27th Annual IEEE Software Technology Conference (STC 2015)

STC
IEEE SOFTWARE TECHNOLOGY CONFERENCE

Call for Presentations

12 October - 15 October
Hilton Long Beach, Long Beach, California, USA

STC is a practitioner's conference focused on Meeting Real World Challenges through Software and Systems Technology. As technologists and as citizens, we are faced with myriad challenges from defending national security, to ensuring the robustness of our critical infrastructure, to sustaining and enhancing large portfolios of legacy systems - all within ever tighter resources constraints.

We invite you to connect with your colleagues and share your ideas, success stories, and lessons learned in applying software and systems technology to meet our common challenges, by submitting an abstract for consideration as one of our conference session or tutorial speakers.

Visit http://ieee-stc.org for abstract submission guidelines and general conference information.

SUBMISSIONS WILL BE ACCEPTED FROM January 25, 2015 - June 15, 2015.

Focus areas this year include:

- Critical Infrastructure Challenges
- Agile/Lean Development
- Cybersecurity
- Supply Chain Risk Management
- Cloud computing
- Mobile applications
- Big Data and Data Analytics
- Affordability
- Open source
- Systems engineering challenges for software-intensive systems
- Human capital and workforce knowledge transition

Important Dates (**Updated**)
June 15, 2015 - Abstract submission closes
July 31, 2015 - Acceptance notices sent
August 14, 2015 - Final abstract and Bio submission
September 11, 2015 - Final presentation due

IEEE

IEEE computer society

RESEARCH FEATURE

Selecting the Best Reliability Model to Predict Residual Defects in Open Source Software

Najeeb Ullah and Maurizio Morisio, Polytechnic University of Turin
Antonio Vetrò, Technical University of Munich

A proposed method evaluates eight popular software reliability growth models and selects the one that can best predict the software's remaining faults, providing practical support for project managers who are considering an open source component.

Software reliability—the probability of failure-free software operations for an extended period in a specific environment[1]—is a critical quality characteristic that affects both safety and cost. For open source software (OSS), for example, reliability is foundational to widespread adoption. Consequently, decades of research have focused on developing and refining various classes of software reliability models, one of which is the software reliability growth model (SRGM), which testers use to determine the cumulative number of expected defects in software before release or, less commonly, to predict its residual defects—the remaining faults or failures after release.

As the sidebar "Software Reliability Growth Model Basics" describes, each SRGM is a mathematical expression that specifies the general form of the software failure process as a function of factors—such as fault introduction and removal—and the operational environment.[1] The model assumes that the failure rate (failures per unit of time) generally decreases as testers identify and remove defects.

An SRGM estimates the failure rate's curve shape by statistically estimating certain parameters, which are specific to the chosen SRGM. With this shape, testers can estimate the extra time required to meet a specified reliability objective and identify the software's expected reliability after release.[1] However, because each application has its own failure-rate curve, no universally applicable SRGM is possible. Moreover, there is no consensus on how to select the best model, and no studies have yet proposed an empirical SRGM selection methodology that is suitable for OSS, which has different needs than software developed in house.

To address that gap, we developed a method to select the best SRGM among several candidates specifically for predicting residual defects in OSS—a measure of concern to project managers who must decide whether or not to include an OSS component. We tested our method empirically by applying 8 popular SRGMs to 21 releases of 7 OSS projects. We also quadrupled the number of

0018-9162/15/$31.00 © 2015 IEEE

SOFTWARE RELIABILITY GROWTH MODEL BASICS

Software reliability growth models (SRGMs) have either a concave or an S-shaped curve. S-shaped models assume that the cumulative failures occur in an S-shaped pattern. Testers are initially unfamiliar with the product, and the fault removal rate slowly increases as they become more familiar. As the testers' skills further improve, that rate increases more quickly and then levels off as residual defects become more difficult to find. The concave models do not include an initial learning curve. Rather, they assume that the failure-rate increase reaches a peak and then becomes stable.

All SRGMs use a nonhomogeneous Poisson process (NHPP) to model the failure process, which is characterized by its mean value function (MVF). If $\{N(t), t > 0\}$ denotes a counting process that represents the cumulative defects detected in t, then an SRGM based on an NHPP is defined as[1]

$$P\{N(t) = n\} = \frac{m(t)^n}{n!}e^{-m(t)}, \quad n = 1, 2\ldots,$$

where MVF is represented as $m(t)$ and is nondecreasing in t under the bounded condition $m(\infty) = a$, with a being the expected total number of defects that testers will eventually detect.

The value of a tells testers whether the software is ready for release or how much additional testing is required if it is not ready. Varying the MVF makes it possible to define different NHPP models.

More basic details on SRGM use are available at the Polytechnic University of Turin website (http://softeng.polito.it/najeeb/IEEE/QuickRefresher.pdf).

Reference

1. M.R. Lyu, *Handbook of Software Reliability Engineering*, McGraw-Hill, 1996.

datasets others have used for the evaluation, allowing us to observe our proposed method's output across a variety of projects and across releases of the same project, each with different defect data amounts.

Our method is different from other approaches because it emphasizes both OSS and the prediction of residual defects, not the cumulative defect expected number. Thus, it is suitable for stable released OSS projects that do not require formal testing. Unlike other methods, our approach is also suitable for practitioners with no background in statistics.

An empirical validation shows that our method is highly effective. Of the 21 releases, it chose the model with the highest estimation precision in 17 releases and with the second highest precision in the other 4 cases.

MODEL COMPARISON METHODS

With no universal SRGM, project managers need some way to choose the best model from myriad available candidates. Unfortunately, despite decades of research, the only consensus about selection is that it should be done on a case-by-case basis. There is no universally accepted selection criterion or metric, and the selection criteria that have been reported were evaluated on very few projects. Some work looks at why selection is difficult, rather than focusing on a selection method. One research group[2] observed that hidden design flaws are the main causes of software failures, which makes model selection problematic.

Models also differ in their intended application. Some SRGM applications look at the total number of cumulative defects at some point in time, which is evident when reliability starts to stabilize. Other SRGM applications are more interested in predicting the total number of defects that will eventually occur and, by extension, the residual defects. Although most research focuses on the former perspective, we believe that the latter perspective characterizes software reliability more concretely.

Metrics versus selection method

In the literature we reviewed, researchers typically applied comparison metrics to some number of SRGMs and noted patterns. Only Catherine Stringfellow and A. Amschler Andrews of Midwestern State University[3] proposed a selection method. Like our method, theirs attempts to predict residual defects, but Stringfellow and Andrews validated their method only on closed source software (CSS) for which testers have completed 60 percent of the planned tests.

Although Stringfellow and Andrews provided no guidelines for adapting their method to OSS, it can help testers decide whether or not to stop testing and release CSS.

Predictive quality

In cumulative defect prediction, many researchers have shown that some model types have higher predictive quality; for example, geometric models—those based on hypergeometric distribution—have better predictive quality than other models.[4] One research group in a different study found that different models work well only on certain datasets, so comparing models' predictive quality for a given application is the best selection approach.[5]

Another group of researchers, which analyzed the predictive quality of 10 models using 5 metrics, observed that the best predictive model depends on the metric used, since the different metrics in their study produced different model choices for the same dataset.[6] Two other approaches rank different models in terms of best fit but do not select the best predictor.[7,8]

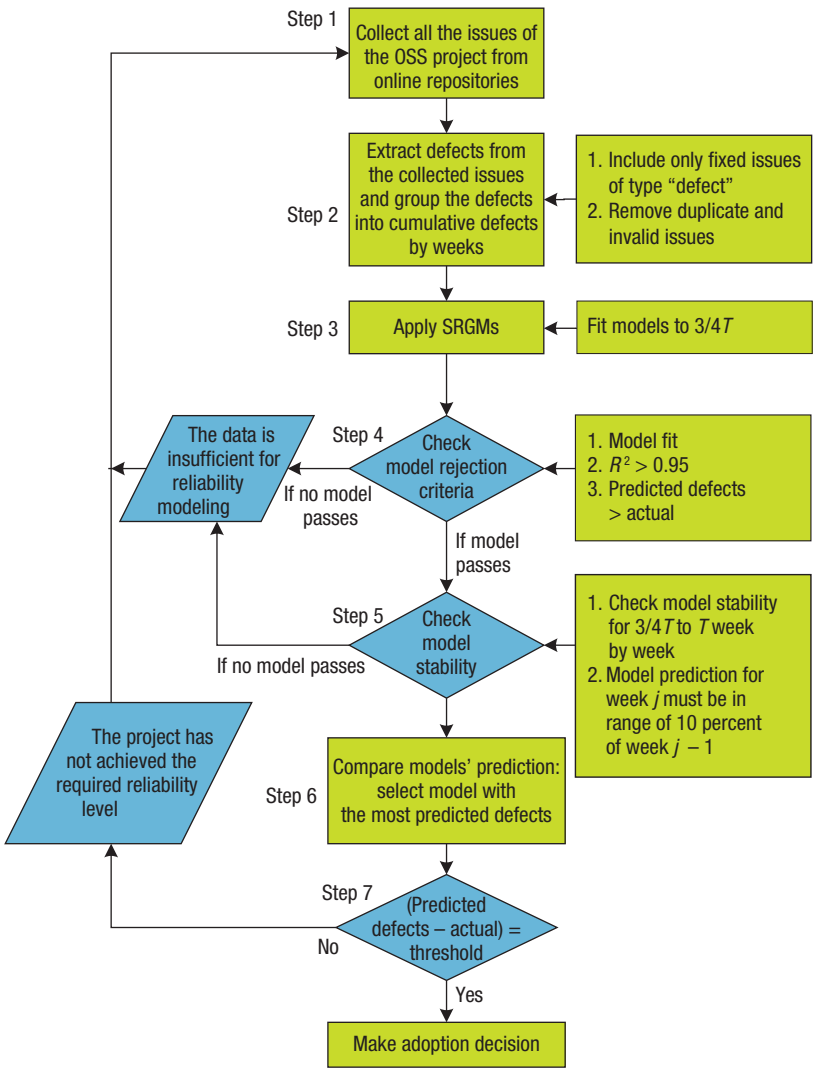

FIGURE 1. Steps in the proposed method of selecting the software reliability growth model (SRGM) with the best fit to data and the best predictive quality in estimating expected residual defects in open source software (OSS).

Of the work we reviewed, all researchers evaluated the predictive quality of SRGMs on the basis of fitting the models to one portion of the defect dataset and predicting only the second portion. Except for the study by Stringfellow and Andrews, all these efforts evaluated a model's predictive quality only in terms of the software's overall behavior, not its residual defects. Evaluations based on overall behavior show only which model outperforms the others, which is not useful insight for practitioners. These studies also used only five or fewer datasets to validate their methods.

METHOD OVERVIEW

We designed our method to select the SRGM that has the best fit to an OSS component's defect dataset and is the best predictor of the component's total number of expected residual defects. Our main purpose is to support project managers in deciding whether or not to adopt an OSS component.

Our method derives from Stringfellow and Andrews' work but must deal with two new problems that stem from the nature of OSS:

- Because many SRGMs have assumptions that might not apply to OSS, a particular model might not fit the data or have a low *goodness of fit* (GoF).
- The defect data for an OSS component is usually limited. The smaller the defect dataset, the longer it might take the models to stabilize.

To address these problems, our method uses several SRGMs and selects the models that best fit the data.

METHOD IMPLEMENTATION

Figure 1 shows a flowgraph of the steps in our method.

Step 1: Collect defect data

After selecting the OSS release of interest, the first step is to collect issues and defect data from online repositories of the OSS project, which are SourceForge, Apache, and Bugzilla.

Step 2: Extract defects

The issues collected in step 1 can be bugs, feature requests, improvements, or tasks, so they must be filtered to remove all issues except bugs or defects—no enhancements, feature requests, tasks, or patches—and only defects reported as closed or resolved. Filtering should exclude open, reopened, or duplicate defects. Part of step 2 also involves grouping the defect data of the entire release interval [0,*T*] into cumulative defects by weeks. Figure 2 shows this interval graphically.

Step 3: Apply models to data

Table 1 lists the eight SRGMs that our method uses. However, these models are a particular context for applying the method. The remaining steps, as well as the release interval, are meant

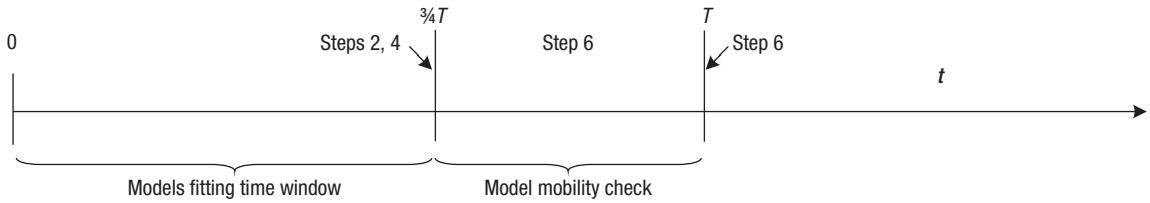

FIGURE 2. Timeline of method steps across two-thirds of an OSS release interval [0,*T*]. The other third of total release time *t* is for validation. Steps 3 through 4 represent the model fitting time window, which takes approximately three-fourths of the method application time, while step 5, the window for determining model stability, takes the other fourth. Steps 6 and 7, select model and compute residual defects, occur at the interval's end.

TABLE 1. Software reliability growth models used in our method.

Model name	Type	Mean value function $m(t)$
Musa-Okumoto*	Concave	$m(t) = a\ln(1+bt),\ a>0,\ b>0$
Inflection S-shaped†	S-shaped	$m(t) = a\dfrac{1-\exp[-bt]}{1+\psi(r)\exp[-bt]},\ \psi(r) = \dfrac{1-r}{r},\ a>0,\ b>0,\ r>0$
Goel-Okumoto†	Concave	$m(t) = a(1-\exp[-bt]),\ a>0,\ b>0$
Delayed S-shaped†	S-shaped	$m(t) = a(1-(1+bt)\exp[-bt]),\ a>0,\ b>0$
Generalized Goel†	Concave	$m(t) = a(1-\exp[-bt^c]),\ a>0,\ b>0,\ c>0$
Gompertz†	S-shaped	$m(t) = ak^{b^t},\ a>0,\ 0<b<1,\ 0<k<1$
Logistic†	S-shaped	$m(t) = \dfrac{a}{1+k\exp[-bt]},\ a>0,\ b>0,\ k>0$
Yamada exponential‡	Concave	$m(t) = a(1-\exp(-r(1-\exp(-bt)),\ a>0,\ b>0,\ r>0$

*J.D. Musa and K. Okumoto, "A Logarithmic Poisson Execution Time Model for Software Reliability Measurement," *Proc. 7th Int'l Conf. Software Eng.* (ICSE 84), 1984, pp. 230–238.
† M. Xie, *Software Reliability Modeling*, World Scientific Publishing, 1991.
‡ H. Pham, "Software Reliability and Cost Models: Perspectives, Comparison and Practice," *European J. Operational Research*, vol. 149, 2003, pp. 475–489.

to be a generic description. In this step, testers apply the models to defect data from step 2, placing them in the model fitting window—marked with the 3/4*T* point in Figure 2 (designating three-fourths of the time allotted for method application).

Because of the nature of defect data, testers can use a general technique, nonlinear regression (NLR), to fit a model to the data. NLR estimates parameters by minimizing the sum of the squares of the distances between the data points and the regression curve. It is an iterative process that starts with an initial estimated value for each parameter. The iterative algorithm then gradually adjusts these parameters until they converge on the best fit. Consequently, adjustments make virtually no difference in the sum of squares.

If the model cannot describe the data, its parameters cannot converge to the best fit, so there is no fit. If a fit is possible, our method evaluates the model's GoF on the basis of the R^2 value, which determines how well a curve fits the data. The value is defined as

$$R^2 = 1 - \frac{\sum_{i=1}^{k}\left(m_i - m(t_i)\right)^2}{\sum_{i=1}^{k}\left(m_i - \sum_{j=1}^{k}\frac{m_j}{n}\right)^2},$$

where k represents dataset size, $m(t_i)$

represents predicted cumulative failures, and m_i represents actual cumulative failures at time t_i.[9] R^2 takes a value between 0 and 1, inclusive. The closer the R^2 value is to one, the better the fit. We chose the R^2 value for its simplicity and because an evaluation of several statistical tests for GoF showed that this measure was at least as powerful as the other tests analyzed.[10]

The larger the R^2 value, the better the curve fits the data, and the easier it is to see any data variation. Fitting a model enables us to estimate the value for all model parameters, notably the expected number of total defects (the a parameter).

For model fitting, we used Prism, a commercial curve-fitting program that employs NLR techniques for curve fitting and supplies model equations and parameter constraints. The program then fits the model to the data and returns an estimate of the best-fitted values for all the parameters of the models along with the GoF value (R^2).

Step 4: Test against fit and prediction thresholds

The fitted models must pass a test based on a threshold GoF value. In this step, the method compares the fitted models' GoF values with the specified R^2 value threshold. The threshold is a subjective decision, and other applications of our method might have a different value. Our GoF value threshold of 0.95 is based on Stringfellow and Andrews' work.

This step also involves checking the fitted models' predictions against the actual number of defects found. Only the models whose prediction is greater than the actual number of defects pass the rejection test, as prediction is meaningless if the model predicts a number that is lower than the actual number of defects.

If no model passes this step, the collected defect data is insufficient for reliability modeling and all the models fail the test, or testers must supply additional data.

Step 5: Test against prediction stability threshold

In this step, the method evaluates the remaining models' prediction stability. (In our study, all models passed step 4, but other comparisons might have different results.) A prediction is stable if the prediction for week j is within ± 10 percent of the prediction for week j – 1. Again, threshold setting is a subjective decision. We used a threshold of 10 percent because of its successful use in other work.[11] If no model has a stable prediction within the threshold, the collected defect data is insufficient for reliability modeling, and testers must supply additional data.

Our method checks model stability within the window of 3/4T to T—the model stability check window in Figure 2. The method checks stability in the same manner for all models: it adds one week of defects to the cumulative defects of 3/4T, so 3/4T + 1 week. It then fits all the models that have passed the rejection step to the cumulative defects in 3/4T + 1 week. Next, it adds another week, so 3/4T + 2 weeks, and so on until it reaches the total number of weeks in the release T.

Step 6: Select the best model

The best SRGM is the one that has passes all the threshold tests and has the highest number of predicted defects. The choice is conservative, but models that overestimate the actual number of defects will be more suitable for a project manager's goal, as cost estimates will be worst case, lowering the risk of adoption (no cost surprises). Widely differing prediction values among models might require conducting an additional analysis with other quality assessment methods,[12] or choosing a subset of models based on the GoF indicator.

Step 7: Compute residual defects

Once the method has selected the best SRGM, it uses that model to compute the OSS's residual defects. With this number, the project manager can decide whether to adopt the OSS, wait until more defects are identified and fixed, or evaluate a different OSS.

METHOD APPLICATION

We applied our method to a cross-section of OSS projects, selecting seven projects with varying natures and large, well-organized communities: Apache, Gnome, C++ Standard Library, JUDDI, HTTP Server, XML Beans, and Enterprise Social Messaging Environment (ESME).

Gathering project data

For Apache and Gnome, we took defect data on three releases each from a published report that had already grouped them into cumulative defects by weeks from release date.[13] We identified the other five open source projects from Apache.org (https://issues.apache.org), which characterized the projects as stable in production with reported issues fixed and closed (66, 95, 68, 64, and 82 percent of reported issues, respectively).

We used Atlassian's JIRA issue tracker (www.atlassian.com/software/jira) to collect defect data for the five projects. JIRA tracks bugs and tasks, links issues to related source code,

TABLE 2. Model fitting for stability check for Gnome release 2.0.

Weeks after release	Actual defects	Delayed S-shaped		Logistic		Gompertz		Generalized Goel	
		Pred*	R^2	Pred*	R^2	Pred*	R^2	Pred*	R^2
12	58	68	0.974	59	0.9937	73	0.9889	105	0.9819
13	58	71(S)	0.9781	62(S)	0.9937	74	0.9909	100	0.9852
14	66	78	0.9764	69	0.9879	84(D)	0.9891	202(D)	0.9866
15	72	86	0.9747	78	0.9844				
16	74	90	0.9772	83	0.986				

*Pred: total defects predicted

plans agile development, monitors activity, and reports project status.

We downloaded all the issues for 3 releases of C++ Standard Library, 3 releases of JUDDI, 2 releases of HTTP Server, 4 releases of XML Beans, and 3 releases of ESME—a total of 15 releases.

For each issue, JIRA records the project name and useful information, such as

- key, summary, and issue type—the unique identity of each issue, a comprehensive description, and whether the issue is a bug, task, improvement, or new feature request;
- status and resolution—the current status can be resolved, closed, open, or reopened, and the resolution can be fixed, duplicate, or invalid;
- created and updated or fixed dates and times; and
- affected versions—the project releases that had the issue.

Following step 2, we manually filtered JIRA data to include only closed or resolved issues and then filtered those results again to include only defects or bugs.[14] We then grouped the refined data into cumulative defects by weeks. The full defect dataset of each release is available at the Polytechnic University of Turin website (http://softeng.polito.it/najeeb/IEEE/datasets.pdf).

Results

Method application took about two-thirds of the time window in Figure 2, with the remaining third devoted to validation. For example, the time interval for Gnome release 2.0 is 24 weeks, so method application was across 16 weeks (2/3 × 24). Of the 16 weeks, we used the last 4 weeks for model stability checking. We chose two-thirds as a window for estimating model parameters because previous studies implied that model parameters do not become stable until about 60 percent of testing is complete.[11]

Table 2 shows partial results for the Gnome 2.0 release for weeks 12 to 16. The Musa, Inflection, Goel, and Yamada SRGMs destabilized in week 13. The table shows results for the remaining four models.

Columns from left to right show the number of actual cumulative defects found in that week and, for each SRGM shown, the number of total defects predicted (Pred) and the GoF value (R^2) value. The results for all project releases are available at the Polytechnic University of Turin website (http://softeng.polito.it/najeeb/IEEE/Remaining_results.pdf).

The method fit each SRGM each week, marking a model with R, S, or D to denote causes for elimination:

- R: failed the fitting and prediction check (did not occur in this release),
- S: model stabilized that week, and
- D: status changed to unstable (destabilized).

All the models—including the four not shown—performed well in fitting and passed the rejection test (step 4), but their predictive quality differed considerably. Five of the eight models destabilized by week 14, and all five significantly overestimated the defect number ("Actual defects" column in Table 2).

Table 2 also shows that the Delayed S-shaped and Logistic models stabilized first at week 13 and remained stable up to week 16 (throughout the entire stability check window). The Delayed S-shaped and Logistic models predicted the number of defects at week 16 as 90 and 83, respectively. The Delayed S-shaped model predicted more residual defects than the Logistic model did, so the method selected it as the best (underlined values in the table).

VALIDATION

For validation, the remaining third of the release interval, we measured each model's prediction capability using *prediction relative error* (PRE):

$$PRE = \frac{\text{Predicted - Actual number of defects}}{\text{Predicted}},$$

where Predicted is the total number of defects a model predicted at two-thirds of the time interval, and Actual is the

TABLE 3. Validation results in choosing the best predictor model.

Project	Release	Model selected by prediction relative error	Model selected by our method
Gnome	V2.0	Delayed S-shaped	Delayed S-shaped
	V2.2	Goel-Okumoto, Yamada exponential	Goel-Okumoto
	V2.4	Inflection S-shaped, Gompertz	Inflection S-shaped
Apache	2.0.35	Delayed S-shaped, Logistic	Gompertz
	2.0.36	Delayed S-shaped, Logistic, Gompertz, Generalized Goel	Generalized Goel
	2.0.39	Inflection S-shaped, Goel-Okumoto, Generalized Goel	Goel-Okumoto
C++ Standard Library	4.1.3	Musa-Okumoto	Inflection S-shaped
	4.2.3	Musa-Okumoto	Gompertz
	5.0.0	Inflection S-shaped, Yamada exponential, Generalized Goel	Yamada exponential
JUDDI	3.0	Musa-Okumoto, Goel-Okumoto, Delayed S-shaped, Yamada exponential, Gompertz	Delayed S-shaped
	3.0.1	Musa-Okumoto, Delayed S-shaped	Delayed S-shaped
	3.0.4	Goel-Okumoto	Goel-Okumoto
HTTP Server	3.2.7	Goel-Okumoto	Goel-Okumoto
	3.2.10	Delayed S-shaped, Logistic	Logistic
XML Beans	2.0	Gompertz	Delayed S-shaped, Gompertz
	2.2	Logistic	Logistic
	2.3	Musa-Okumoto, Goel-Okumoto, Logistic, Yamada exponential, Generalized Goel	Logistic
	2.4	Delayed S-shaped, Gompertz	Gompertz
Enterprise Social Messaging Environment	1.1	Musa-Okumoto, Goel-Okumoto	Goel-Okumoto
	1.2	Logistic, Gompertz	Gompertz
	1.3	Delayed S-shaped, Logistic, Gompertz	Inflection, Goel-Okumoto, Generalized Goel

number of defects at the end of the time interval.

For each release and each project, we computed the PRE for each model, ranked the models accordingly, and considered the model with the lowest PRE as the best predictor for that release.

Table 3 compares the best predictors chosen by PRE and by our selection method for each project release. For 17 of the 21 releases, the PRE and our method chose the same model. In the remaining four releases, the best model had a negative PRE, so our method rejected the model. Even so, the method still chose the second-best model (the one with the lowest positive PRE) in all four of the releases.

OBSERVATIONS

Our results are promising and prompt several observations that could guide future work.

Validity of results

With 21 datasets, our validation is more extensive than any similar study to date. However, because we used datasets that others produced, we could not control the quality of the issues collected and reported. Issues, time to fix, or description could have been missing. We tried to mitigate this lack of control by selecting established OSS projects with large communities and using datasets that similar research efforts have employed.

One validity threat is the lack of generalization. Although we used the largest number of datasets with the largest variety of release intervals, which vary

from days to a year, we still cannot generalize our results to all project releases because a model could still fail to fit the data or fail the GoF or stability threshold tests.

Model ranking

Of the 21 datasets, no model ranked best in more than a few cases, and each of the 8 models was ranked best at least once. This result is consistent with related work and provides further evidence that any methodology must select the best SRGM on a project-by-project basis.

That said, models with certain characteristics might be better predictors. For example, in 14 of the 21 releases, the best model was S-shaped, not concave. Our previous studies[14] had similar results. One possibility is that the OSS project community members (users and reviewers) tend not to react immediately to a new release—behavior that is consistent with the learning phase in S-shaped models.

Model fit

In line with a common assumption in the literature, we considered each release to be its own independent project,[11,12,14,15] and our results show that different models fit different releases of the same projects. One explanation is that model selection is based only on defect history, not on any other project characteristic, and the factors that determine defect history are not well understood. Defects can stem from characteristics like code complexity or from the domain's inherent challenges, or even from the degree of coding skill. An organization's processes and management choices can also influence project quality.

Amount of defect data

Our method rejected models primarily because of prediction instability, not GoF value, which could signal the need for more defect data. The method overcomes this instability by 21 releases, but evaluating this much data might not be feasible with other applications. Additional research might focus on determining the minimum amount of defect data needed to select an SRGM.

Parameter choices

The thresholds and windows in our method seemed to perform well: GoF minimal fitting threshold of 0.95, stability threshold of 10 percent, three-fourths interval for method application, and one-fourth interval for fitting and stability check. A threshold sensitivity analysis was outside the scope of our study, but it could be a topic for future work.

Our work aims to support practitioners in characterizing OSS reliability in terms of residual defects, with the larger goal of helping project managers decide on OSS use. We believe that our method key contributions are its systematic approach and the extent of validation. In future work, we plan to refine our method and conduct another study to explore additional selection methods and possibly develop supporting tools to automate aspects of our method.

ABOUT THE AUTHORS

NAJEEB ULLAH is a PhD candidate in computer engineering at Polytechnic University of Turin, Italy. His research interests include software reliability, software process improvement, and empirical software engineering. Ullah received an MS in computer engineering from Polytechnic University of Turin. He is a member of IEEE. Contact him at najeebullahkhan1984@gmail.com.

MAURIZIO MORISIO is a professor and Software Engineering Group leader in the Department of Control and Computer Engineering at Polytechnic University of Turin. His research interests include empirical software engineering, software processes, and green IT. Morisio received a PhD in software engineering from Polytechnic University of Turin. He is a member of IEEE and associate editor in chief of *IEEE Software*. Contact him at maurizio.morisio@polito.it.

ANTONIO VETRÒ is a postdoctoral research fellow in the Software and System Engineering Department at the Technical University of Munich, Germany. His research interests include empirical methodologies, analyses of process and product data for software quality, and data quality assessment. He is a member of ACM. Contact him at vetro@in.tum.de or phisaz@gmail.com.

REFERENCES

1. *IEEE Std. 1633-2008, IEEE Recommended Practice on Software Reliability*, IEEE, 2008; http://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=4554206.
2. S. Brocklehurst et al., "Recalibrating Software Reliability Models," *IEEE Trans. Software Eng.*, vol. 16, no. 4, 1990, pp. 458–470.
3. C. Stringfellow and A.A. Andrews, "An Empirical Method for Selecting Software Reliability Growth Models," *Empirical Software Eng.*, vol. 7, no. 4, 2002, pp. 319–343.
4. J.D. Musa, A. Iannino, and K. Okumoto, *Software Reliability: Measurement, Prediction, Application*, McGraw-Hill College, 1987.
5. A.L. Goel, "Software Reliability Models: Assumptions, Limitations, and

Applicability," *IEEE Trans. Software Eng.*, vol. 11, no. 12, 1985, pp. 1411–1423.
6. A.A. Abdel-Ghaly, P.Y. Chan, and B. Littlewood, "Evaluation of Competing Software Reliability Predictions," *IEEE Trans. Software Eng.*, vol. 12, no. 9, 1986, pp. 950–967.
7. N. Miglani and P. Rana, "Ranking of Software Reliability Growth Models Using Greedy Approach," *Global J. Business Management and Information Tech.*, vol. 1, no. 11, 2011, pp. 119–124.
8. M. Anjum, A. Haque, and N. Ahmad, "Analysis and Ranking of Software Reliability Models Based on Weighted Criteria Value," *Int'l J. Information Technology and Computer Science*, vol. 5, no. 2, 2013, pp. 1–14.
9. K.C. Chiu, Y.S. Huang, and T.Z. Lee, "A Study of Software Reliability Growth from the Perspective of Learning Effects," *Reliability Eng. & System Safety*, vol. 93, no. 10, 2008, pp. 1410–1421.
10. O. Gaudoin, B. Yang, and M. Xie, "A Simple Goodness-of-Fit Test for the Power-Law Process based on the Duane Plot," *IEEE Trans. Reliability*, vol. 52, no. 1, 2003, pp. 69–74.
11. P. Zeephongsekul, G. Xia, and S. Kumar, "Software Reliability Growth Model: Primary Failures Generate Secondary Faults under Imperfect Debugging," *IEEE Trans. Reliability*, vol. 43, no. 3, 1994, pp. 408–413.
12. C. Stringfellow, "An Integrated Method for Improving Testing Effectiveness and Efficiency," PhD dissertation, Colorado State Univ., 2000.
13. X. Li et al., "Reliability Analysis and Optimal Version Updating for Open Source Software," *Information and Software Technology*, vol. 53, no. 9, 2011, pp. 929–936.
14. N. ullah and M. Morisio, "An Empirical Study of Reliability Growth of Open versus Closed Source Software through Software Reliability Growth Models," *Proc. 19th Asia-Pacific Software Eng. Conf.* (APSEC 12), 2012, pp. 356–361.
15. K. Sharma et al., "Selection of Optimal Software Reliability Growth Models Using a Distance Based Approach," *IEEE Trans. Reliability*, vol. 59, no. 2, 2010, pp. 266–276.

STUDENT DESIGN SHOWCASE

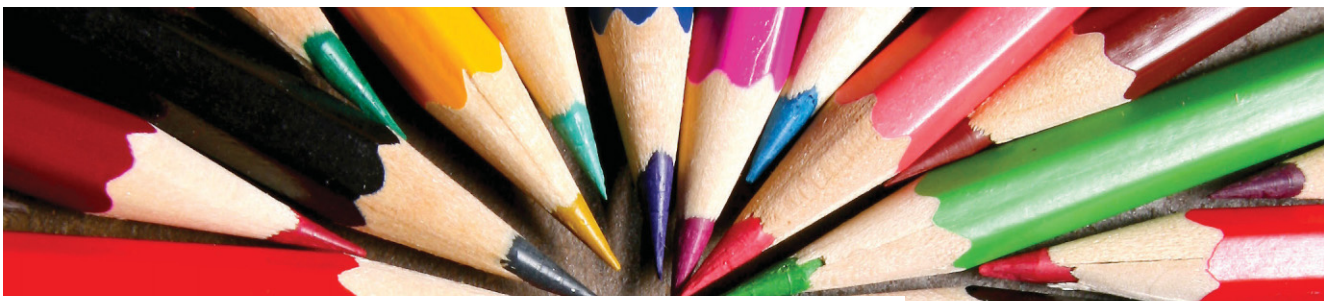

Tracking Cows Wirelessly

Greg Byrd, North Carolina State University

A student team from NC State designed and built a prototype wireless network to monitor the milking and weighing of cows.

To successfully operate any farm, effective livestock management is crucial. Efficient, affordable, and scalable livestock management solutions play an increasingly important role in modern farming, as the number of dairy farms in the US decreases, but the number of dairy cows on each increases. Dairy cows require careful monitoring for milking, weighing, and other activities, so the ability to reliably track these animals in large numbers is particularly important.

Dairy cows are typically identified by visible ear tags. Although tags with embedded RFID devices have been available—allowing them to be scanned electronically—because of cost, most tags use low-frequency (LF) RFID, so the scanner must be within a few inches of the tag. Consequently, farmworkers need to be "up close and personal" with each and every cow for reliable scanning.

Although RFID tagging of cattle has been widely adopted in Europe, US dairy farms are more reluctant to do so because of costs and the lack of national standards. To address the cost and convenience factors, a team of students in the Department of Electrical and Computer Engineering (ECE) at North Carolina State University (Figure 2) designed and built a prototype wireless network that combines long-range ultra-high-frequency (UHF) RFID tags with low-cost wireless and computing components. The long-range RFID allows unmanned scans of multiple tags, and the wireless network provides scalable data collection without costly infrastructure.

NETWORK OVERVIEW

Figure 1 shows an overview of the prototype network using the ZigBee wireless protocol to communicate. The RFID reader is connected to the ZigBee wireless networking node, so when a new RFID tag is detected, the ZigBee node sends a data packet to the controller node. As cows enter and exit the milking station, an RFID reader identifies the cows by their long-range RFID tags, and when the RFID tag is no longer in range, the ZigBee node sends a data packet indicating that the cow left the milking station. (Multiple cows may enter the milking stalls at the same time, so the prototype system can monitor the entry and exit times of up to eight cows simultaneously.)

As the cows leave the milking station, they pass through a weighing station. A floor scale and an RFID

0018-9162/15/$31.00 © 2015 IEEE

EDITOR **GREG BYRD**
North Carolina State University; gbyrd@computer.org

reader are attached to a microprocessor. The cow's weight and identity are passed to a ZigBee node, which sends this information to the controller node, collects the milking and weighing information and the amount of time each cow spends in the milking station, displays it on a user interface, and records it in a spreadsheet.

TAGS, READERS, AND SENSORS

The passive UHF RFID tag is the enabling technology for this project. In a passive RFID system, the reader antenna sends a radio signal at a particular frequency. The RFID tag (also called an RFID transponder) contains an antenna tuned to the same frequency and a microchip. The received signal powers the chip, which modulates a signal that is transmitted back to the reader, and the reader then translates the modulated signal into digital data.

Developed by RFID Sensor Systems, the project's sponsor, the prototype tags have been demonstrated to work from distances of up to 150 feet (45 meters). RFID tags measure approximately 6 inches by 0.75 inches and are mounted to the back of a standard plastic ear tag.

The RFID reader is SkyeTek's Nova module, which includes an ARM Cortex CPU and a UHF transceiver. SkyeTek Protocol software runs on the microprocessor, providing an API for standard RFID-read operations. A serial communications link connects the SkyeTek module to the ZigBee node, which is described in the next section.

For the milking station, the prototype system includes a separate short-range antenna for each stall; multiple antennae can be multiplexed into a single reader. Certain techniques can use a single antenna to distinguish signals from multiple tags, which would be an interesting benefit of the long-range tags provided by RFID

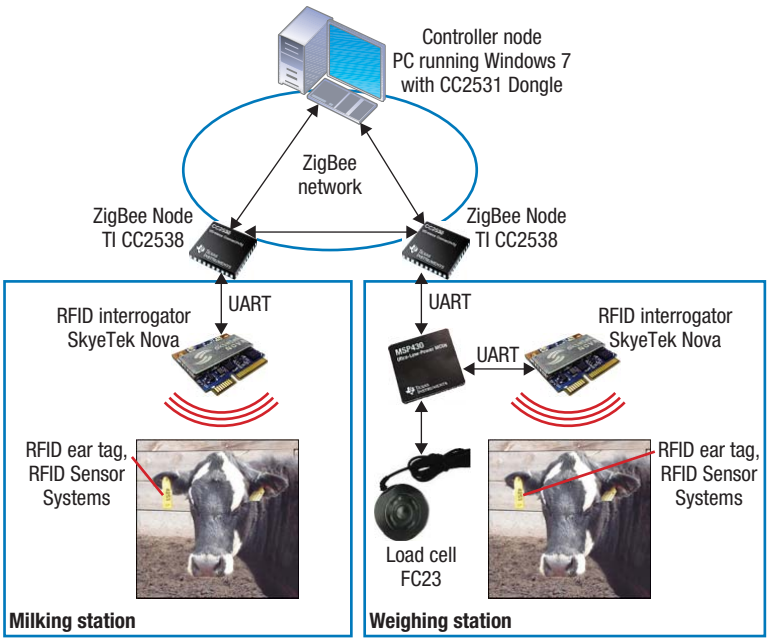

Figure 1. Overview of a prototype wireless network for tracking dairy cows. RFID tags indicate when a cow enters and leaves the milking or weighing station. Timestamps and weight are delivered to the controller node via the wireless mesh network. UART, universal asynchronous receiver/transmitter.

Figure 2. The Department of Electrical and Computer Engineering (ECE) at North Carolina State University design team. Left to right: Bryan Campbell, André Ramos, Youn Chu, and Anthony Laws.

PROJECT DETAILS

- Name: Howling Cow Wireless Network
- School: North Carolina State University
- Department: Electrical and Computer Engineering (ECE)
- Student Participants: Bryan Campbell, Youn Chu, Anthony Laws, André Ramos
- Faculty Mentor: Dr. Rachana Gupta
- Sponsors: Dr. William Carr, RFID Sensor Systems
- Special Thanks: Jim Carlson (ECE), Dr. Jake Adams (ECE), Dr. Dan Poole (Animal Science)

2015 IEEE/IBM WATSON STUDENT SHOWCASE

IEEE and IBM are in search of creative, forward-thinking students to participate in an exciting team-based showcase. Do you want an opportunity to work with IBM's Watson to develop an innovative app? Do you want to develop your cognitive computing skills while earning a chance to win cash prizes? Winning entries will also be featured in this column. The deadline for submitting completed entries has been extended to 30 September 2015, but register your intent to participate now at http://goo.gl/52WUlh so you don't miss this opportunity.

Sensor Systems but was beyond the scope of this project.

For the weighing station, both an RFID reader and a weight (load) sensor are needed. The load sensor is incorporated into a floor-mounted plate and connected to a low-power Texas Instruments MSP430 microprocessor. When the RFID reader detects a cow entering the weighing area, it sends a signal to the microprocessor via the serial port. The microprocessor uses an integrated analog-to-digital converter to read the weight from the load sensor. Both the identification and the weight are sent over a serial connection to a ZigBee node, which sends the appropriate data packet to the controller node.

Both the MSP430 microprocessor and SkyeModule Nova utilize low-power processing cores that have special sleep modes when the CPU is inactive. In addition, the MSP430 FR5738 microprocessor includes 16 Kbytes of integrated ferroelectric nonvolatile RAM (FRAM), which offers lower energy consumption than flash memory for storing instructions and constant data. Although milking station components are likely to have access to electricity, the weighing station and other future sensor stations may require solar power or other energy harvesting and storage solutions.

WIRELESS RADIO

ZigBee is a collection of protocols designed for low-power wireless networks, such as smart homes and sensor networks. The physical layer is based on the IEEE 802.15.4 standard for low-rate wireless networking. Upper-layer protocols add routing and other features to create a mesh network so that a device can relay information through multiple ZigBee nodes to a remote destination node. Other protocol features include security, device discovery, and messaging.

The ZigBee node chosen for this prototype is the Texas Instruments CC2538 system on chip. The chip includes an ARM Cortex-M3 processor, a 2.4-GHz radio transceiver, embedded RAM and flash memories, and a hardware cryptographic accelerator. Like the MSP430 and SkyeModule Nova, the CC2538 features low-power operating modes.

ZigBee is a good match for this application because only a small amount of data is sent per event. The latency also isn't critical. What's needed is a reliable, low-power, inexpensive, and secure network that's scalable to larger farms with multiple milking and sensor stations. An extensive ZigBee mesh network can be deployed, with transmission distances of up to a mile long for a single link.

CONTROLLER NODE SOFTWARE

The controller node is a standard laptop or desktop system equipped with a ZigBee network node. The choice for this prototype is the Texas Instruments CC2531 USB Dongle. It includes an 8051 microcontroller, which runs packet-sniffing software and delivers packets to the host through the USB interface.

As described above, the controller receives data packets from the milking and weighing station nodes. A Java application displays the information for each event (such as when a cow enters or leaves a station) on a user interface window. Events are also recorded in a spreadsheet-compatible file for postprocessing.

DEPLOYMENT ENVIRONMENT

The target deployment environment for the wireless network is the dairy

research and teaching farm at NC State. The farm spans 389 acres, supports a herd of 300 cows, and includes a 20-stall milking station, visitor's center, classroom, and museum. The farm's milk is used in Howling Cow dairy products—including ice cream, milk, and heavy cream—and is processed at the Feldmeier Dairy Processing Lab, which is operated by the Department of Food, Bioprocessing, and Nutrition Science. Howling Cow products are sold at various locations on campus.

The project plan also included an on-the-farm demonstration. Unfortunately, logistical challenges associated with working with live animals have delayed the demonstration. Nonetheless, the components were successfully tested together in the lab.

Our goal was to create a low-cost proof-of-concept system. The retail cost of the system components is around $1,200. With additional development and higher volumes, the cost can be further reduced. For example, the long-range RFID tags would allow multiple cows to be scanned at the milking station with a single reader and antenna. Also, the computational requirements for the load sensor, RFID reader, and ZigBee communication could be consolidated onto a single processor.

When deployed at the farm, researchers and students will use the Howling Cow network to demonstrate how RFID, wireless, and other information technologies can improve the efficiency and productivity of the farm as well as the health of the animals. The farm of the future will employ many emerging technologies: unmanned drones for crop and livestock monitoring (where the long-range RFID tags will be critical), self-driving tractors, 3D printing of replacement parts, and precision agriculture driven by data analytics. Pilot projects like this one will lead the way to show how high-tech can be affordable for farmers in North Carolina and around the world.

SUBMIT A PROJECT

As much as I enjoy writing about our students at North Carolina State University, I'm really interested in hearing about interesting student-led design projects in computer science and engineering everywhere. If you would like to see your project featured in this column, fill out the submission form at: www.computer.org/student-showcase.

GREG BYRD is associate head of the Department of Electrical and Computer Engineering at North Carolina State University. Contact him at gbyrd@computer.org.

A Farewell to Air Gaps, Part I

Hal Berghel, University of Nevada, Las Vegas

One of the most fundamentally misguided ways to protect a networked infrastructure is to introduce an air gap. The US has been mastering the art of crossing them for more than 30 years.

Eager to send détente to an early grave, US President Ronald Reagan wasted no time taking advantage of the intelligence provided by the French domestic intelligence agency's KGB spy, Colonel Vladimir Vetrov, codenamed "Farewell." In 1981, French President François Mitterrand offered Reagan the "Farewell Dossier," 4,000 KGB documents that by some accounts triggered a spectacularly kinetic CIA response.

Gus Weiss documented the entire affair in his capacity as a special assistant to the secretary of defense as well as director of International Economics for the National Security Council (NSC). According to Weiss,[1] the Soviets under Leonid Brezhnev viewed détente as some much needed economic breathing room that enabled them to improve the USSR's economy. To expedite this improvement, the Soviet Council of Ministers and the Communist Party of the Soviet Union's Central Committee established the KGB's Directorate T to find and retrieve Western targets of opportunity (read: stolen intellectual property). A new operating arm, Line X, was charged with the oversight and management of these acquisitions.

Going back at least 30 years before the Farewell revelation, the US and its NATO allies were keen to keep the keys to the kingdom out of Soviet hands. Through export and other economic policies, they opposed Line X and the forces behind it. Thus the sharing of nuclear secrets, certain sophisticated manufacturing techniques, semiconductor technology, and sensitive trade secrets was verboten, and the corresponding products, weapons, advanced computers, machinery, and the like were embargoed.

But that didn't stop the KGB. Line X populated visiting Soviet delegations with KGB agents to learn about agriculture, manufacturing, defense, and whatever else they could. What they didn't learn from direct inspection, they tried to buy. What they couldn't buy directly, they tried to purchase through third parties. Failing that, theft was always a viable option. This was Cold War technology transfer Soviet style: a continuation of the atom bomb spying effort, but with economic objectives that continued to some degree until Reagan found out about all of it from Mitterrand.

THE CIA's MARK TO MARKET

At the NSC, Weiss read the Farewell Dossier and determined that the Soviets were acquiring technology trade

0018-9162/15/$31.00 © 2015 IEEE

EDITOR **HAL BERGHEL**
University of Nevada, Las Vegas; hlb@computer.org

secrets at breakneck speed. "Our science was supporting their national defense," he reported.[1] From the dossier, William Casey's CIA developed a Line X shopping list in 1982 and, with the FBI's help, assisted the Soviets on their shopping spree by seeing that they got enhanced versions of the things they sought—these items "would appear genuine, but would later fail."[1] Defective computer chips, flawed parts, and misleading or bogus technical information were supplied to vendors known to sell to the Soviet Union. We know this from Weiss's report,[1] but things got much more interesting when Thomas Reed, also a member of Reagan's NSC, came into the picture. Weiss and Reed were with the NSC when the Reagan–Casey–CIA–FBI intrigue began in early 1982.

Reed picked up Weiss's storyline in *At the Abyss: An Insider's History of the Cold War* (Presidio Press, 2005). He tells how the Soviets were in need of some software for their natural gas pipeline that stretched from Siberia to Eastern Europe. They dispatched a KGB operative to a Canadian software supplier. The US intelligence folks were given the heads-up through Farewell, prompting the FBI to work with the Canadians to "enhance" the software, which the KGB then obtained.

According to Reed, the software included a Trojan horse allowing the West to regulate the pipeline's pump speeds and settings, valve openings, and internal pressures—pushing them well beyond safe operating limits. Needless to say, the flawed software produced the desired result of disrupting the pipeline's operations—in fact, it's alleged to have caused the largest non-nuclear explosion in history (equivalent to detonating three kilotons of TNT). In response to the NSC's concern about the resulting explosion, Weiss is reported to have told them not to worry about it, but he gave no explanation. Reed claims that Weiss told him the story and provided him with his notes shortly before his mysterious death on 25 November 2003. Although some dispute whether the hack resulted in the actual explosion, it wasn't for lack of effort on the part of the Reagan administration. So far as I can determine, no authoritative source disputes the rest of Reed's account.

Some of Weiss's notes appeared on the CIA website under his name on 14 April 2007,[1] more than two years after Reed's book had removed any cover of secrecy. The Trans-Siberian Pipeline part of the story isn't included in Weiss' published CIA notes; add this to Weiss's strong opposition to the Iraq war and you have the stuff of which great conspiracies are made. I'd be remiss if I didn't mention here that Reed remains a partisan loyalist to Reagan and his politics, so it could be that his account is biased and somewhat sanitized—perhaps there's much more to the story than even Reed reported.

MALWARE ARCHAEOLOGY AND STUXNET

For the next chapter on international cyber sabotage, we turn to Stuxnet. So that we're all on the same page, you need to know that the officially unconfirmed US/Israeli cyberattack against the Iranian uranium enrichment facility at Natanz—reported in 2010—was never called Stuxnet by those who allegedly deployed it. "Operation Olympic Games" was the codename used by the planners when they presented the idea to the George W. Bush administration originally. The Stuxnet moniker came much later, from investigators external to the project who juxtaposed fragments of contained filenames.

The Stuxnet archaeology produced sufficient digital artifacts from which several conclusions can be drawn. First, it actually shares some of the architecture and codebase with the remote-access Trojan and information stealer, W32.Duqu,[2] and the espionage hack, Flame.[3] In fact, early versions of Stuxnet (circa v0.5) are thought to derive from the Flame platform, whereas later versions (circa v1.0) also derive from the Tilded platform—so-called because contributors tended to start filenames with tildes. Other Flame- and Tilded-based malware are certain to remain in circulation in some form

In *Countdown to Zero Day: Stuxnet and the Launch of the World's First Digital Weapon*,[4] journalist Kim Zetter—who's covered cybercrime for *Wired* and other publications—convincingly documents (see "The Legend of Stuxnet" sidebar) that although Flame and Duqu are derived from the Flame and Tilded codebases, respectively, Stuxnet borrows from both, and in different proportions over time.[4] Kaspersky Lab's Costin Raiu believes that the Stuxnet cyber aggressors borrowed heavily from the Flame platform code and then shifted to Tilded later in development due to the simplicity and tightness of the latter code.[5]

In *Confront and Conceal: Obama's Secret Wars and Surprising Use of American Power* (Broadway Books, 2012), journalist David Sanger explains Operation Olympic Games in the context of the political climate during both the latter George W. Bush and early Barack Obama administrations. He speculates that Flame was a US artifact used when Stuxnet was in the experimental stage, and that the Duqu/Tilded code was primarily an Israeli product. Only after Bush authorized Operation Olympic Games did the two teams begin sharing code. There's considerable evidence for this because Flame and Duqu differ greatly in sophistication—pointing to the likelihood that different teams with different skill levels contributed to it. The Stuxnet payload—the part directed at specific industrial controllers—seems

THE LEGEND OF STUXNET

The discovery of the Stuxnet worm by VirusBlokAda (http://anti-virus.by/en/index.shtml), a small Belarus security company, is itself an interesting story that's now part of cyber lore. Stuxnet's creators went to considerable trouble to thwart discovery by major security software vendors, but VBA was apparently too small to be on their radar. An Iranian reseller for VBA reported events of interest from several customers that couldn't seem to rid themselves of malware infection. This alerted VBA, which in turn uncovered a kernel-level rootkit operating on its customers' computers. Further analysis discovered an injector based on Windows .LNK files carried by USB flash drives. VBA had unknowingly uncovered the first few layers of Stuxnet because its security monitoring was more effective in this case than more widespread security software.

.LNK files are simply file shortcuts or links to local files for use with Microsoft Explorer. If the local file is an executable, activating the link causes it to load and run. Actually, the Stuxnet v1.0 .LNK exploit started out as a Microsoft feature. Not content to allow users access to a distracting array of file extensions, Microsoft shipped Windows with the extensions to known file types (.doc, .ppt, and so on) suppressed. To see the common file extension, the user has to disable this option using Folder Options in the Control Panel. What's more, even if you disable this feature, Windows will still suppress system-reserved file extensions like—you guessed it—.LNK. These file extensions aren't file-system links and thus aren't suppressed in Explorer, but rather are handled by the Windows Registry. By default, the Registry entry for .LNK is NeverShowExt. So, this "feature" was already weaponized by Microsoft: to wit, if .LNK file extensions are suppressed, your file.doc.lnk appears instead as an innocuous file.doc. That this .doc file will link to something other than a Word file is inherently concealed.

Stuxnet's .LNK injector took advantage of a design flaw in IconHandler within the Windows Shell that incorrectly parses .LNK files. IconHandler allowed the execution of the executable linked to the icon instead of just displaying the icon on USB devices (see "Shortcut Icon Loading Vulnerability" at https://technet.microsoft.com/library/security/ms10-046). VBA analysts spotted this attack vector and hence were credited with the Stuxnet discovery, although they didn't reverse engineer the code enough to reveal the payload. That was done by Symantec[1] and others. Although the injector was designed for Windows, the payload was optimized for the industrial controllers.

The .LNK injector wasn't a zero-day exploit, but it was novel. Autorun was the default injector for removable storage and media before the mid-2000s. Since responsible security consultants had recommended disabling Autorun by 2000, it was becoming ineffective by the time that v0.5 was released, so Stuxnet's authors instead went with newer and more sophisticated exploits. Unlike v0.5, v1.0 used a multi-exploit.

In all, Symantec and Kaspersky Lab identified five different infection hacks: the Autorun exploit in v0.5, the .LNK hack in v1.0 and later, and three elevation-of-privileges exploits through vulnerabilities in the Windows keyboard file handler, print spooler, and task manager. In addition, Kim Zetter[2] documents eight different propagation tactics once the malware was installed on one computer:

1. the .LNK hack described above continuing to work on new USB targets,
2. infecting Siemens Simatic Step 7 project files,
3. exploiting a security-through-obscurity defect in the way Siemens handled user authentication,
4. injecting malware into shared Siemens databases,
5. using a peer-to-peer exploit on LAN file sharing that worked in much the same way as software updaters do,
6. installing a covert file-sharing server on each infected machine for redundancy,
7. spreading via network file shares, and
8. exploiting a Conficker-style Trojan-horse vulnerability that had been reported and for which a patch had been created but not necessarily installed.

Of these, (1) and (2) were the most heavily used. The novelty of the exploits, together with the complexity of the code and the multitude of attack vectors, distinguishes Stuxnet from other malware and made it the current gold standard of cyber warfare.

For a politician, the allure of Stuxnet-style cyber-kinetic attacks is that they don't immediately put American lives at risk. As such, Operation Olympic Games can be thought of as a tactical sibling to the current US drone war, the ideological ancestry of which dates back to Eisenhower-era covert CIA operations in the Middle East (Operation Ajax) and Central America (Operation PBSUCCESS).

References

1. N. Falliere, L.O. Murchu, and E. Chien, *W32.Stuxnet Dossier*, report v. 1.4, Symantec, Feb. 2011; https://www.symantec.com/content/en/us/enterprise/media/security_response/whitepapers/w32_stuxnet_dossier.pdf.
2. K. Zetter, *Countdown to Zero Day: Stuxnet and the Launch of the World's First Digital Weapon*, Crown, 2014.

to have remained relatively constant over time, suggesting that it was developed first and held in reserve while awaiting an opportunity to launch.

Absent leaks or whistleblowers, we might never know the full extent of the cyber mischief Stuxnet and its various sources have contributed to, although the Flame arsenal of attacks is now known to include a hack of Microsoft Windows Updater—a feature not present in the Stuxnet-era exploits. Much of the current analysis of Flame and its derivatives must end with some speculation, of course, since state-sponsored developers don't use networked version control repositories like GitHub.

CENTRIFUGE SUBTERFUGE

Details of the Stuxnet worm have been well documented by scores of security analysts.[6,7] It's well known that it targets a very narrow range of industrial SCADA (supervisory control and data acquisition) systems running Siemens Simatic Step 7 (S7) software. By corrupting programmable logic controllers (PLCs) that use two specific frequency converters, the worm causes the PLCs to direct the systems to operate outside safe limits to the point of self-destruction (sound familiar?).

The SCADA systems themselves are controlled by Windows-compatible computers, so the attack vector is indirect via Windows PLCs that are likely isolated from the Internet by means of an air gap. More specifically, Stuxnet targets the S7 PLC software running on Windows computers, which in turn regulates S7 controllers with S7-315 or S7-417 frequency converters used by the IR-1 family of uranium centrifuges. That last sentence is a mouthful, but it highlights the fact that Stuxnet's payload was precisely targeted. The particular S7 controllers used by these centrifuges were specifically introduced as a cost-saving measure to replace the hardwired and telemetry-controlled systems of the 1950s and 1960s. Although they effectively reduced the cost, it's difficult to discount the negative externalities from the total cost of use—as they were designed without robust security in mind.

That's *what* Stuxnet did. Of course, that doesn't explain *how* the feat was accomplished. Because the earlier Stuxnet version (v0.5 from 2005–2007[8]) wasn't Internet enabled, it only propagated through shared S7 project files on Windows computers. Unlike its more effective descendant, v0.5 manipulated the input and output valves on the centrifuges. Although excessive pressure would damage the centrifuges, v0.5 wasn't as effective as v1.0, which could actually manipulate centrifuge rotors until they spun out of control. v1.0 also used existing Windows vulnerabilities to more effectively propagate the worm through a LAN, removable USB storage devices, network file shares, Windows remote procedure calls, printer spools, or the Internet itself.

The initial v0.5 injection was accomplished using a Flame platform Autorun exploit through infected USB drives, which were carried to the Natanz facility and inserted into network computers by four Iranian subcontractors—a relatively simple injection strategy compared to v1.0. Interestingly, the payload didn't change in subsequent versions of the worm, suggesting that the weaponized S7/SCADA attack code was likely mature by the time President Bush approved proceeding to the attack's final stages. Symantec also assesses the payload code to be far superior to the injector and call-home code, which many have used to speculate on the nationality of the teams that developed each. As an interesting aside, v0.5 used four command-and-control servers to update the code, all of which claimed to be from a nonexistent advertising agency, Media Suffix, whose tag line, "Deliver What the Mind Can Dream," is an apt mantra for what I'll call postmodern, neoconservative, malware epistemology.

The detailed technical reports on Operation Olympic Games/Stuxnet are consistent in their recognition of its aggressiveness and capabilities—when measured in terms of breadth and depth, the destructive potential certainly qualifies as revolutionary. Not only was it a hydra-headed near-zero-day exploit—and the first known rootkit to target PLCs—the operation also deployed a sophisticated injector into the S7 processors. The worm was also a self-replicator par excellence—working well in LANs, peer-to-peer communications, removable storage devices, network shares, and I/O streams. It also used a sophisticated covert command-and-control interface, an advanced form of anti-malware sonar for most of the popular security products available, strong encryption to hide its binaries, and an embedded playback mechanism to spoof normal operation.

> We might never know the full extent of the cyber mischief Stuxnet and its various sources contributed to.

Although not as complicated as Flame, Stuxnet represents a serious contribution to the art of cyber warfare and is unique as a cyber-kinetic attack tool. Attacking a sovereign nation's infrastructure puts it in a league of its own. This hasn't escaped the attention of political leaders and state-supported technologists of every stripe. Not only did Stuxnet set a new standard for hacking into industrial control systems, it upped the ante in the global cyber-arms race.

Although the inevitable retaliatory attacks by anti-Western interests

OUT OF BAND

will be decried as naked cyber aggression in the mass media (as the bogus claims surrounding the Sony hack make clear[9]), the cyber-aware globalists will always regard the US and Israel as mentors. Future cyber-kinetic attacks will predictably involve failures of power grids, water supplies, and transportation. This cyber genie is out of the bottle.

If nothing else, the Trans-Siberian Pipeline hack and Stuxnet attack on Natanz demonstrate that air gaps have been ineffective for well over 30 years. The air gap joins security-through-obscurity as classic examples of faith-based security—a strategy for protection that's based on faith alone (see my column "Faith-Based Security"[10]).

Next month we'll investigate what, if any, lessons we've learned.

REFERENCES

1. G.W. Weiss, *The Farewell Dossier: Duping the Soviets*, report, CIA, posted 14 April 2007; www.cia.gov/library/center-for-the-study-of-intelligence/csi-publications/csi-studies/studies/96unclass/farewell.htm.
2. *W32.Duqu: The Precursor to the Next Stuxnet*, white paper v1.4, Symantec, 23 Nov. 2011; www.symantec.com/content/en/us/enterprise/media/security_response/whitepapers/w32_duqu_the_precursor_to_the_next_stuxnet.pdf.
3. *W32.Flamer*, report, Symantec, updated 5 Jun. 2012; www.symantec.com/security_response/writeup.jsp?docid=2012-052811-0308-99.
4. K. Zetter, *Countdown to Zero Day: Stuxnet and the Launch of the World's First Digital Weapon*, Crown, 2014.
5. Resource 207: Kaspersky Lab Research Proves that Stuxnet and Flame Developers Are Connected, report, Kaspersky Lab, 11 Jun. 2012; www.kaspersky.com/about/news/virus/2012/Resource_207_Kaspersky_Lab_Research_Proves_that_Stuxnet_and_Flame_Developers_are_Connected.
6. N. Falliere, L.O. Murchu, and E. Chien, *W32.Stuxnet Dossier*, report v. 1.4, Symantec, Feb. 2011; www.symantec.com/content/en/us/enterprise/media/security_response/whitepapers/w32_stuxnet_dossier.pdf.
7. K. Zetter, "An Unprecedented Look at Stuxnet, the World's First Digital Weapon," *Wired*, 3 Nov. 2014; www.wired.com/2014/11/countdown-to-zero-day-stuxnet.
8. G. McDonald et al., *Stuxnet 0.5: The Missing Link*, report v. 1.0, Symantec, 26 Feb. 2013; www.symantec.com/content/en/us/enterprise/media/security_response/whitepapers/stuxnet_0_5_the_missing_link.pdf).
9. H. Berghel, "Cyber Chutzpah: The Sony Hack and the Celebration of Hyperbole," *Computer*, vol. 48, no. 2, 2015, pp. 77–80.
10. H. Berghel, "Faith-Based Security," *Comm. ACM*, vol. 51, no. 4, 2008, pp. 13–17.

HAL BERGHEL is an ACM and IEEE Fellow and a professor of computer science at the University of Nevada, Las Vegas. Contact him at hlb@computer.org.

SECURITY

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

Vetting Mobile App Vendors

Jeffrey Voas, IEEE Fellow

Vetting a mobile app vendor's processes via a questionnaire is a poor substitute for vetting the app itself, but situations often arise when it's the only or most practical option.

The certification of software as a product has been an elusive goal in the software community as well as a long-time wish of mine.[1,2] Talk has gone on for decades, but with few results. The reasons are numerous, but at the core of the problem is the fact that software is nonphysical, as is the input data from the environment that software feeds on during execution.

All software isn't equal: there are numerous types, and software delivery channels are equally numerous. This column focuses on mobile apps, one genre of software that can be delivered in many ways—today, most commonly via app stores.

MOBILE APP SECURITY RISKS

Organizations deploy various apps to facilitate business processes and decision making. Apps provide unprecedented connectivity among employees, vendors, and customers and enable real-time information sharing, unrestricted mobility, and improved functionality. Despite these benefits, acquired third-party apps could have serious security risks. Whether purchased or free, mobile apps often contain code vulnerabilities. Moreover, ascribing blame to mobile apps for a security violation is difficult because they interact with other entities (including other apps) to accomplish their purpose.

To mitigate these risks, a prudent organization normally takes measures to ensure that acquired apps conform to its predefined security requirements. These measures would likely involve in-house testing of candidate apps using third-party tools such as static and dynamic code analyzers that search for the presence of known vulnerabilities.[3] (Such tools do little for unknown ones.) However, test tool licenses can be prohibitively expensive. Although an organization might be able to leverage test tool results for an app shared by others, this would be unusual and in some cases—depending on the license agreement's fine print—possibly illegal. Further, many organizations don't know how to configure tools to test apps against security requirements, assuming the right tools are available and the requirements exist.

MOBILE APP VENDOR QUESTIONNAIRE

As an alternative to third-party testing tools, I suggest employing a questionnaire to elicit security-related information directly from app vendors to predict the potential

28 QUESTIONS FOR MOBILE APP VENDORS

1. Does your app contain any known security issues that need to be mitigated or should be revealed? Are there current, publicly known vulnerabilities in the app? For example: Does it access phone, network (Wi-Fi and/or Bluetooth), and/or GPS resources; camera or other recording hardware or software; operating system resources or libraries? What data, if any, does the app transfer from the mobile device to a remote host, including during PC synchronization? If you believe that your app doesn't contain any security issues, explain why and ignore the remaining questions. If it does, please respond in detail to each of the questions that are pertinent to your app. Note that insufficient information may result in a delay or void adoption of the app.
2. Is the app source available? Is there an active user community providing peer review and evolving the app?
3. Does the app license restrict a licensee from discovering flaws or disclosing details about app defects or weaknesses to other organizations?—that is, is there a "gag rule" or limits on sharing information about newly discovered flaws?
4. Does the app have a reputation? What about relative to security? Are there reviews that recommend for and against the app?
5. Are there security requirements "structured" as part of a general release and others "as needed" or "custom" for a particular release?
6. What process was used to prioritize security-related enhancement requests from users?
7. Were security requirements developed independently or integrated into mainstream requirements engineering activities?
8. Were misuse/abuse cases derived from the app requirements? Were relevant attack patterns used to identify and document potential threats?
9. What threat assumptions were made, if any, when designing protections for the app and information assets processed?
10. What threat-modeling process, if any, was used when designing app protections?
11. What secure development standards and/or guidelines were provided to developers?
12. Are security risk management activities incorporated as part of the software development methodology?
13. Does the app validate inputs from potentially untrusted sources before being invoked?
14. Has the app been designed to execute within a constrained execution environment?
15. Does documentation explain how to install, configure, and use the app? Does it identify options that shouldn't normally be used because they create insecurities?
16. Does the app have any security-critical dependencies or need additional controls from other software, firmware, or hardware? If yes, please describe.
17. Does the app include content produced by suppliers other than the primary developer?
18. What policies and procedures were used to verify the functionality and security of nondevelopmental components?
19. What types of functional tests were performed on the app during in-house development—for example, spot checking, component-level testing, and integrated testing?
20. How often are security tests performed on the app? Are they performed by an internal test team, an independent third party, or both?
21. How much code did in-house testing cover?
22. Were test cases included for potential misuse/abuse scenarios?
23. Were static or dynamic software security analysis tools used to identify weaknesses in the app that can lead to exploitable vulnerabilities? If so, when in the software development life cycle are these scans performed?
24. Has the app been certified and accredited? If so, what release/version/configuration? When? By whom? What criteria or scheme was used to evaluate and accredit the app?
25. Will patches be distributed?
26. How are reports of defects, vulnerabilities, and security incidents involving the app collected, tracked, and prioritized?
27. What policies and processes are used to verify that app components don't contain unintended, dead, or malicious code?
28. How frequently will major versions of the app be released?

security risks of an app's code. In addition to vetting app vendors, such questionnaires make all stakeholders within the organization—acquirers, integrators, and users alike—more "security savvy."

My proposed questionnaire (see the sidebar) modifies select questions from the US Department of Homeland Security (DHS)'s *Software Supply Chain Risk Management & Due-Diligence* (https://buildsecurityin.us-cert.gov/sites/default/files/due-diligence-MWV12_01AM090909.pdf), which is designed to aid in the procurement of more secure all-purpose software of four types:

- commercial off-the-shelf (COTS) proprietary,
- government off-the-shelf (GOTS),
- COTS open source, and
- custom (bespoke).

My modifications make the DHS questions specific to third-party mobile apps. I excluded questions that solicited proprietary information or are too laborious to complete for most mobile app vendors—that is, requiring pages rather than a few sentences in response.

Readers shouldn't take my questionnaire as definitive: it doesn't apply to all apps, and many other DHS questions might be more relevant for a particular scenario. Nevertheless, it serves as a useful sanity checklist of a candidate vendor's security maturity.

Vetting a mobile app vendor's processes via a questionnaire is a poor substitute for vetting the app itself, but situations often arise when it's the only or most practical option. Still, many vendors won't welcome such questions or be truthful in responding to them. Reassure candidate vendors that your goal isn't to expose proprietary details about their app, and point out that in responding to your questions they'll learn something about their own security maturity. It's a red flag if they find the questions, which touch on the most basic principles of software security and testing, difficult to understand.

I acknowledge that most organizations buy and use laptops, smartphones, and other devices with multiple preloaded apps that can't always be vetted. Moreover, neither app acquirers nor vendors would consider using my proposed questionnaire for a 99-cent app. However, if an organization plans to use an unvetted app on numerous devices and deems it critical that the app not conflict with the organization's IT security policy, my approach is both pragmatic and cost-effective.

ACKNOWLEDGMENT

I thank Joe Jarzombek at the US Department of Homeland Security for facilitating a discussion on this topic years ago and encouraging me to write a column about it.

REFERENCES

1. J. Voas, "Certifying Off-the-Shelf Software Components," *Computer*, vol. 31, no. 6, 1998, pp. 53–59.
2. J. Voas, "Toward a Usage-Based Software Certification Process," *Computer*, vol. 33, no. 8, 2000, pp. 32–37.
3. S. Quirolgico et al., *Vetting the Security of Mobile Applications*, NIST Special Publication 800-163, Nat'l Inst. of Standards and Technology, US Dept. of Commerce, January 2015; http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-163.pdf.

JEFFREY VOAS, Security column editor, is an IEEE Fellow. Contact him at j.voas@ieee.org.

SOFTWARE TECHNOLOGIES

Designing Resource-Aware Cloud Applications

Reiner Hähnle, Technical University of Darmstadt
Einar Broch Johnsen, University of Oslo

Realizing the full potential of virtualized computation—the cloud—requires rethinking software development. Deployment decisions, and their validation, can and should be moved up the development chain into the design phase.

As data storage and processing move to the cloud, our interactions with computers are undergoing dramatic transformation. The cloud, as a network of virtual machines, has no fixed location and is only accessed remotely. Although many continue to use a desktop PC to access cloud-based data and applications, people increasingly rely on their smartphones, tablets, and other mobile devices.

Despite serious data-privacy concerns, several qualities make cloud computing compelling from a business standpoint.[1] One such quality is elasticity: businesses can pay for computing resources only when they're needed, avoiding major upfront investments for resource provisioning; the client application can add processing power, memory, and additional virtual machines on the fly as needed. The cloud also offers a scalable virtualized platform for data processing that can be shared among multiple devices—if a service uses cloud-based processing, its capacity can be adjusted automatically as new users arrive. Thus, another key cloud feature is agility: new services can be quickly and flexibly brought to market at minimal cost, without initial investments in hardware.

With these key advantages, cloud computing is rapidly gaining popularity throughout the world. Currently, there are more than 3.9 million cloud computing–related jobs in the US and more than 18 million worldwide.[2] In the EU, cloud computing is projected to create 2.5 million jobs and grow the economy annually by €160 billion by 2020.[3] However, inconsistent reliability and insufficient control

0018-9162/15/$31.00 © 2015 IEEE

EDITOR **MIKE HINCHEY**
Lero—the Irish Software Research Centre;
mike.hinchey@lero.ie

of resources present significant barriers to the industrial adoption of cloud computing. To overcome these barriers and gain control of the cloud's virtualized resources, client services must become resource aware.

RETHINKING SOFTWARE DEVELOPMENT

The virtualization of computation has transformed how we interact with, share, and manage data, but taking full advantage of cloud computing requires us to rethink how we design and develop software.

The elasticity of software executed in the cloud gives designers control over the execution environment's resource parameters, such as the number and kind of processors, memory, storage capacity, and bandwidth. Parameters can even be changed dynamically at runtime. Thus, cloud-service clients can not only deploy and run software, but they can also fully control the tradeoffs between incurred costs to run the software and the quality of service (QoS) delivered.

Traditionally, software design is based on specific assumptions about deployment, including the size of data structures and the amount of RAM and number of processors required to execute it. However, because of the cloud's transformative effect on computation, software now must be designed for scalability from the beginning or else it will later require extensive design changes.

Later changes to software can be fatal because of the relative increase in cost to fix defects during successive development phases.[4] As Figure 1 shows, the IBM Systems Sciences Institute estimated that, compared to fixing a defect during design, it costs 6.5 times more to fix during implementation, 15 times more during testing, and 100 times more during maintenance. These costs don't take into account additional economic losses associated with delayed time to market, lost customers, and bad public relations.

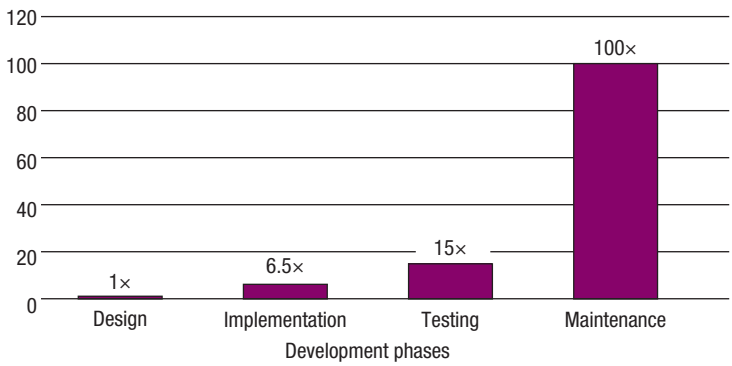

Figure 1. Relative costs to fix software defects at each phase of the development process in a static infrastructure. (Source: IBM Systems Sciences Institute)

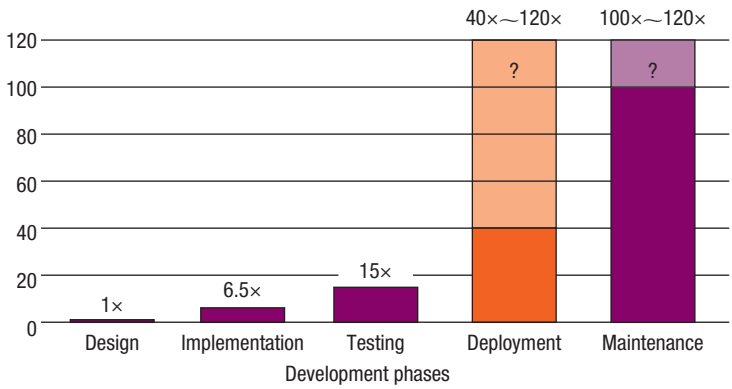

Figure 2. Relative costs to fix software defects at each phase of the development process in an elastic virtualized infrastructure. (Source: E. Albert et al., "Engineering Virtualized Services," *Proc. 2nd Nordic Symp. Cloud Computing Internet Technologies* [NordiCloud 13], 2013, pp. 59–63)

The ratios in Figure 1 are for static deployment, but it's reasonable to expect even more significant costs when fixing defects in elastic virtualized environments due to the additional complexity of resource management. Figure 2 presents a conservative estimate of the ratios in an elastic environment: it can cost between 40 and 100 times more to fix defects during deployment than during design. It's thus essential to detect and fix deployment errors—such as failure to meet service-level agreements (SLAs)—as early in the development process as possible, preferably during the design phase.

To maximize cloud computing's potential, software development demands a new methodology that models deployment during the design phase and facilitates detection of deployment errors early and efficiently through the use of software tools such as simulators, test generators, and static analyzers.

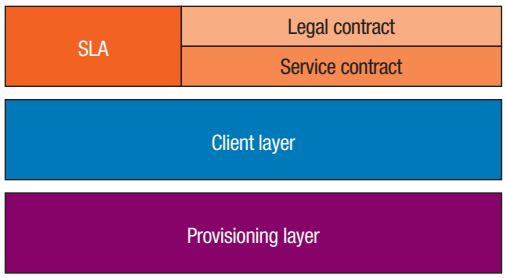

Figure 3. Conceptual layers of a deployed cloud service. Functionality is represented in the client layer. The provisioning layer makes resources available to the client layer and determines available memory, processing power, and bandwidth. The service level agreement (SLA) specifies what resources the provisioning layer should make available to the client service and includes both a legal and a service contract.

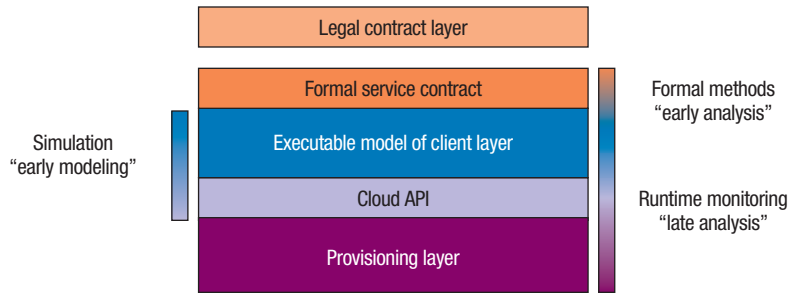

Figure 4. Conceptual layers of a resource-aware deployed cloud service. The service contract is formalized and statically checked against an executable client model at design time. The client model can access an abstract interface to the provisioning layer in the form of a cloud API, thus enabling simulation. The cloud API also permits instrumentation of deployed applications with monitors that ensure compliance with the service contract at runtime.

CONTROLLING DEPLOYMENT IN THE DESIGN PHASE

Our analysis presents a software engineering challenge: how to introduce deployment decision validation into the modeling phase of the software development chain without convoluting the design with deployment details.

With a cloud-based service, developers typically design its functionality first, then determine which resources are needed and how they'll ultimately be provisioned through an SLA, as Figure 3 shows. Functionality is represented in the client layer. The provisioning layer makes resources available to the client layer and determines available memory, processing power, and bandwidth. The SLA specifies what resources the provisioning layer should make available to the client service. A typical SLA has two components: a legal contract stating the mutual obligations of the client service and software developer, and consequences in the event of a breach by either party; and a service contract that details the technical parameters and cost figures for the offered services.

We propose connecting these currently distinct conceptual layers to better control deployment in the design phase.

First, the client should have access to the provisioning layer through a special cloud API to observe and modify resource parameters, as Figure 4 shows. The cloud API isn't identical to the APIs commonly offered by cloud providers, but is instead an abstract interface to the provisioning layer. Using an executable language, such as Abstract Behavioral Specification (ABS),[5] it enables design-phase modeling of client behavior to simulate different client-side provisioning schemes and observe their impact on cost and performance. Creating such an executable model is feasible during the design phase.

Second, to connect an SLA to the client layer, a formal semantics is given to service contract aspects of the SLA. This would allow, at design time, analysis of client behavior such as resource consumption and worst-case performance vis-à-vis the SLA; test-case generation; and functional verification.[6] This process is highly automated, for example, in ABS.[7] Design-time analysis makes the assumptions about the cloud API explicit. The cloud API also permits instrumentation of deployed applications with monitors that ensure compliance with the service contract at runtime.

BENEFITS OF RESOURCE AWARENESS

Making deployment decisions during the design phase can help shift control from the provisioning layer to the client layer and enables the client service to become resource aware. This provides numerous attractive opportunities.

Fine-grained provisioning

As in other metered-industry sectors such as telephony and electricity, business models for cloud-based resource provisioning are increasingly fine grained. Selecting the best software model is highly complex, which thus far has benefited resource providers. Design-time analysis and comparison of deployment decisions make it possible to optimize provisioning schemes for end users, such as spot pricing.

Tighter provisioning

Better profiles of the client layer's resource needs will help cloud providers avoid over-provisioning to comply with their SLAs. Better usage of resources means that more clients can be served with the same amount of hardware in the datacenter, without violating SLAs and incurring penalties.

Application-specific resource control

Through design-time analysis of scalability, the client layer can better leverage the cloud's elasticity. Knowing beforehand the load thresholds for scaling up deployment, cloud-services providers can avoid breaking SLAs and disappointing end users' expectations.

Application-controlled elasticity

Autonomous, resource-aware services that run their own resource-management strategy are possible. Such services could monitor loads on virtual machines as well as end-user traffic, and thus could make decisions about the tradeoffs between the QoS delivered and the incurred cost. The service could interact with the provisioning layer through an API to dynamically scale up or down, or it could even request or bid for virtual machine instances with given profiles in a future virtual resource marketplace.

The efficiency and performance of cloud-based services could be improved by moving deployment decisions up the development chain. In addition, resource-aware services could give the client better control of resource usage and thereby comply with SLAs more inexpensively.

Formal methods for early analysis, executable models of client behavior for early monitoring, and runtime monitors for late analysis are key elements to realizing this vision. Concrete examples of these elements are being implemented as part of the EU FP7 project "Envisage: Engineering Virtualized Services" and are described, along with other details about the project, at www.envisage-project.eu.

ACKNOWLEDGMENT

This work was supported in part by EU project FP7-610582 Envisage: Engineering Virtualized Services.

REFERENCES

1. R. Buyya et al., "Cloud Computing and Emerging IT Platforms: Vision, Hype, and Reality for Delivering Computing as the 5th Utility," *Future Generation Computer Systems*, vol. 25, no. 6, 2009, pp. 599–616.
2. L. Columbus, "Where the Cloud Computing Jobs Will Be in 2015," *Forbes*, 12 Dec. 2014; www.forbes.com/sites/louiscolumbus/2014/12/12/where-cloud-computing-jobs-will-be-in-2015.
3. European Commission, "European Cloud Computing Strategy," Digital Agenda for Europe: A Europe 2020 Initiative, updated 27 Feb. 2015; http://ec.europa.eu/digital-agenda/en/european-cloud-computing-strategy.
4. B.W. Boehm and P.N. Papaccio, "Understanding and Controlling Software Costs," *IEEE Trans. Software Eng.*, vol. 14, no. 10, 1988, pp. 1462–1477.
5. E.B. Johnsen et al., "ABS: A Core Language for Abstract Behavioral Specification," *Proc. 9th Int'l Symp. Formal Methods for Components and Objects*, LNCS 6957, Springer, 2011, pp. 142–164.
6. E. Albert et al., "Engineering Virtualized Services," *Proc. 2nd Nordic Symp. Cloud Computing Internet Technologies* (NordiCloud 13), 2013, pp. 59–63.
7. E. Albert et al., "Formal Modeling of Resource Management for Cloud Architectures: An Industrial Case Study Using Real-Time ABS," *J. Service-Oriented Computing and Applications*, vol. 8, no. 4, 2014, pp. 323–339.

REINER HÄHNLE is a professor and dean of the Computer Science Department at the Technical University of Darmstadt. Contact him at haehnle@cs.tu-darmstadt.de.

EINAR BROCH JOHNSEN is a professor in the Department of Informatics at the University of Oslo. Contact him at einarj@ifi.uio.no.

CS CONNECTION

EIC APPLICANTS SOUGHT FOR CS JOURNALS

The IEEE Computer Society seeks editor in chief applicants for two transactions journals: *IEEE Transactions on Cloud Computing* (*TCC*) and *IEEE Transactions on Sustainable Computing* (*TSUSC*). The term for *TCC* is two years, and the term for *TSUSC* is three years—each begins on 1 January 2016 and is renewable for one additional term.

Prospective candidates are asked to provide a complete curriculum vitae, a brief plan for the publication's future, and a letter of support from their institution or employer. The plan should include (1) the candidate's perspective on the challenges and opportunities for the enhancement and improvement of the publication's quality, (2) specific tasks to be undertaken if appointed EIC, (3) objective milestones associated with each task, and (4) a proposed schedule.

The EIC will be responsible for the publication's day-to-day volunteer leadership, including coordinating and overseeing the peer review process; recommending candidates for the editorial board; chairing the editorial board; developing editorial plans for the publication; serving as a nonvoting, ex-officio member of the publication's steering committee; serving as a nonvoting, ex-officio member of the IEEE Computer Society Publications Board; and working in general with volunteers and staff to ensure and maintain the timely publication of an exceptionally high-quality transactions.

Applications are due **30 June 2015**. The search committee prefers electronic submissions in Microsoft Word or PDF. Please direct all questions and submit completed applications to Kimberly Sperka at ksperka@computer.org. For more details, please visit www.computer.org/tcc for *TCC* or www.computer.org/tsusc for *TSUSC*.

DAVID PADUA RECEIVES HARRY H. GOODE AWARD

David Padua, the Donald Biggar Willett Professor of Engineering and a professor of computer science at the University of Illinois at Urbana–Champaign, has been named the 2015 IEEE Computer Society Harry H. Goode Award recipient.

Cited for his "basic and lasting contributions to parallel languages, compilers, and tools," Padua's research focuses on program analysis, transformation, and optimization strategies. His research and projects have focused on developing methodologies to facilitate the programmer's task of creating reliable, easy-to-maintain programs that achieve excellent performance. Padua's ongoing projects include the study of program optimization strategies and the design of compiler techniques for new parallel programming constructs.

The Harry H. Goode Award recognizes exceptional achievement in the information-processing field—either as a single contribution of theory, design, or technique of outstanding significance or as the accumulation of important contributions on theory or practice over an extended period. The award includes a bronze medal and a US$2,000 honorarium. To view the list of recipients, visit www.computer.org/web/awards/harrygoode.

UNDERGRADUATE TEACHING AWARD GOES TO HENRY CHAN

Henry Chan, associate professor in the Department of Computing at Hong Kong Polytechnic University (HKPU), has received the IEEE Computer Society's 2015 Computer Science and Engineering Undergraduate Teaching Award. Recognized for his "outstanding contributions to computing education through teaching, mentoring students, and service to the education community," Chan will receive a certificate and US$2,000 honorarium. The award honors outstanding contributions to undergraduate education through teaching and service.

A supporter of the Computing for Application, Research, Entrepreneurship, and Service (CARES) educational philosophy at HKPU, Chan has received three President's Awards and five Faculty Awards, including a President's Award in 2010–2011 for his achievement in teaching. Under his supervision, Chan's students received more than 30 external awards, and his R&D team won a silver medal at the 38th International Exhibition of

0018-9162/15/$31.00 © 2015 IEEE

Inventions in Geneva for their Flexible RFID Encoder and Decoder (FRED) software tool. He has coauthored a textbook on teaching e-commerce using an integrated computing and business approach, as well as numerous research papers.

Chan chaired the IEEE Hong Kong Section in 2012 and the IEEE Hong Kong Section Computer Society Chapter in 2008–2009. He served as founding general chair of the IEEE International Conference on Teaching, Assessment, and Learning for Engineering (TALE) in 2012, and founding cochair of the Computer App Programming (CAP) competition for secondary school students. He is vice chair of the IEEE CS Cloud Computing Special Technical Community, editor for Cloud-Link on the IEEE Cloud Computing website, and associate editor of the *Journal of Electronic Commerce Research*.

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

2015 IEEE COMPUTER SOCIETY PETITION CANDIDATES

At their 5 June meeting, the IEEE Computer Society Board of Governors will approve the candidate slate for the 2015 election.

As set forth in Computer Society Bylaws Article II—Nominations and Elections, candidates for the Board of Governors and the positions of president-elect and first and second vice president may be nominated by written petition of the membership.

Petition candidates shall meet the qualifications and follow the procedures for the office sought as specified by IEEE and IEEE CS bylaws.

Petition candidates shall personally notify, in writing or via email, the Society secretary and assistant secretary of his or her intention to run as a petition candidate, stating the position to be sought. Notification may be sent in the year during which the petition candidate seeks inclusion on the ballot, but no sooner than the day after the date on which the Society's Board of Governors approves the slate of candidates.

Individual Computer Society members eligible to vote in the election may nominate candidates by written petition, provided such nominations are made at least 28 days before the first published day of the election. In 2015, the deadline for receipt of written petitions is Monday, 6 July. The number of signatures required on a member petition is determined in accordance with Article II, Section 1.

For further information about the petition process, contact Anne Marie Kelly at amkelly@computer.org.

CALL AND CALENDAR

CALLS FOR ARTICLES FOR *COMPUTER*

Computer plans a March 2016 special issue on communications and privacy under surveillance.

Today, our daily communications are subject to surveillance by government agencies and private hackers. The purpose could be legal crime investigation or terrorism prevention, or it could be illegal privacy infringement or theft of confidential information. There are often gray areas between legal and illegal surveillance.

Despite the controversy, communications surveillance technologies continue to evolve. However, these sensitive technologies have seldom been reported on in the scientific literature, as developers have largely subscribed to the idea of "security by obscurity." Now, though, many realize that surveillance technologies could be improved much faster via "security by clarity," with the technologies examined by both good and bad guys. This could also reduce the gray area between legal and illegal activities, thereby protecting privacy.

This special issue will foster dissemination of the latest communications and privacy surveillance methodologies. The guest editors' aim is to publish high-quality articles presenting the state of the art in lawful, mass, wireless, and backdoor surveillance.

Articles are due **1 August 2015**. Visit www.computer.org/computer/cfp3 to view the complete call for papers.

Computer plans an April 2016 special issue on human augmentation.

Human augmentation encompasses technologies that enhance productivity and improve or restore capabilities of the human body or mind. Such technologies are designed to empower and improve health, quality of life, and performance.

Examples include devices that contribute to more advanced sensory capabilities such as glasses for viewing augmented visual content, or limb-like devices that enhance motion or muscle capabilities with minimal invasiveness.

Other types of human augmentation might work with specific IT resources such as big data assets. These include watches or other wearable electronics that link wearers to outside sources of information that might be visual, audio, or text-based.

This special issue will focus primarily on approaches that work with computing technologies rather than just pharmaceutical, psychological, or physiologic approaches. Articles could address examples of human-augmentation technologies and systems, an overview of a specific application area, or technical challenges.

Articles are due **1 October 2015**. Visit www.computer.org/computer/cfp4 to view the complete call for papers.

Computer plans an August 2016 special issue on supply chain security for cyberinfrastructure.

The design, fabrication, assembly, distribution, system integration, and disposal of today's electronic components, systems, and software involve multiple untrusted parties. Recent reports demonstrate that this long and globally distributed supply

SUBMISSION INSTRUCTIONS

The Call and Calendar section lists conferences, symposia, and workshops that the IEEE Computer Society sponsors or cooperates in presenting.

Visit www.computer.org/conferences for instructions on how to submit conference or call listings as well as a more complete listing of upcoming computer-related conferences.

0018-9162/15/$31.00 © 2015 IEEE

chain is vulnerable to counterfeiting (including cloning, overproduction, and recycling) and malicious design modification (such as Trojan attacks). The issues associated with counterfeit components include security and reliability risks to critical systems as well as profit and reputation loss for intellectual property owners, all of which discourage innovation in system development. Recent bugs such as Heartbleed have shown that flaws in open source and third-party code can have a tremendous impact, including the leakage of sensitive and personal data.

Although awareness in the hardware supply chain has increased in recent years, the scope of the problem continues to grow and evolve. Data from the Government and Industry Data Exchange Program (GIDEP) and Information Handling Services indicates a sixfold and fourfold increase, respectively, in reported counterfeit components over the last four years. Existing solutions fail to provide adequate protection against supply chain security issues, and many are too intrusive and expensive to be practical for industry use.

This special issue is intended to raise awareness of supply chain issues, highlight new attacks, point out existing solutions, and encourage fresh protection approaches. It will focus on supply chain security as well as comprehensive, cost-effective, and easy-to-use solutions.

We solicit articles on topics related to security in all parts of the hardware and software supply chain. Although articles that focus on specific supply chain security gaps are acceptable, the guest editors strongly encourage those that address problems with all steps of the supply chain and/or hardware–software integration.

Articles are due **1 February 2016**. Visit www.computer.org/computer/cfp8 to view the complete call for papers.

SEEKING PAPERS ON COMPUTATIONAL SOCIAL SYSTEMS

IEEE Transactions on Computational Social Systems welcomes submissions on topics such as modeling, simulation, analysis, and the understanding of social systems from the quantitative and/or computational perspective. Learn more at www.ieeesmc.org/publications/transactions-on-computational-social-systems/call-for-papers-and-special-issues.

STC 2015

The 27th Annual IEEE Software Technology Conference (STC 2015) is co-sponsored by the IEEE Computer Society, IEEE, and the Computer Society's Technical Council on Software Engineering.

Software is a primary enabler as well as a major cost driver for current and emerging systems. Effectively addressing software acquisition, development, and sustainability challenges, including safety and cybersecurity, is a key to system and mission success.

STC 2015 is focused on bringing attendees insight into current software technology issues, proven solutions, and lessons learned applicable to government, defense, and industry stakeholders. The goal is to provide immediately usable insights and practical solutions from software technology professionals.

STC 2015 will take place 12–15 October 2015 in Long Beach, California. Visit http://conference.usu.edu/STC for complete conference information.

CALLS FOR ARTICLES FOR IEEE CS PUBLICATIONS

IEEE Pervasive Computing plans an April–June 2016 special issue on domestic pervasive computing.

Pervasive computing crosses many domains, but perhaps none are as complex or important as the home. Households are full of challenges and opportunities for technologies and applications. These include setup and ease of use, multiple occupants with different needs and expertise, critical privacy and security concerns, and a range of social scenarios and deployment variations.

This special issue aims to explore technologies and applications related to all aspects of pervasive computing pertaining to domestic situations, entailing fields such as ubiquitous computing, ambient intelligence, the Internet of Things, and intelligent environments.

Articles are due **1 July 2015**. Visit www.computer.org/pervasive/cfp2 to view the complete call for papers.

IT Professional plans a March/April 2016 special issue on emerging IT trends in healthcare and well-being.

IT advances are driving and facilitating major transformations in

CALL AND CALENDAR

EVENTS IN 2015

JULY 2015

7–10 ..... ICAC 2015
13–15 ..... ICCCNT 2015
13–16 ..... CBI 2015

AUGUST 2015

6–7 ..... NAS 2015
22–25 ..... HotChips 2015
24–26 ..... FiCloud 2015

SEPTEMBER 2015

7–9 ..... EISIC 2015
16–18 ..... VS-Games 2015
29 Sept.–1 Oct. ..... ICSME 2015

OCTOBER 2015

12–15 ..... STC 2015
19–21 ..... SOCA 2015
19–21 ..... WIMOB 2015
25–30 ..... VIS 2015

healthcare and personal well-being, which are also being affected by factors such as the push for cost reduction and increased personalization, the increasing volume and variety of medical-related data, and the demand for data sharing and social presence.

IT can revolutionize and improve healthcare access and delivery and can help enrich the quality of life by supporting chronic-disease management, elder care, healthy lifestyle campaigns, disease prevention, remote care, telehealth, and rural care.

Articles are due **1 August 2015**. Visit www.computer.org/itpro/cfp2 to view the complete call for papers.

IEEE Internet Computing plans a May/June 2016 special issue on cloud storage.

Cloud computing has established itself as a key component of the modern computing infrastructure. The initial focus was on virtualizing computing resources, but it has since expanded to networking-resource virtualization via software-defined networking.

Virtualized and software-defined cloud storage, in which storage hardware is decoupled from the management software, is the last piece necessary to realize fully virtualized datacenters. However, cloud storage faces technical challenges, including those related to performance, privacy, and security.

This special issue will address these and other issues with articles from both industry and academia.

Brief article descriptions are due **1 August 2015**. Articles are due **1 September 2015**. Visit www.computer.org/web/computingnow/iccfp3 to view the complete call for papers.

IEEE Computer Graphics and Applications plans a May/June 2016 special issue on high-performance visualization and analysis.

The ability to collect and generate data has grown dramatically. Concurrently, computing technology has rapidly evolved from single-processor machines to large-scale, petaflops high-performance systems. The confluence of these factors introduces new research challenges and opportunities for high-performance scientific visualization and analysis.

Articles are due **1 September 2015**. Visit www.computer.org/web/computingnow/cgacfp3 to view the complete call for papers.

IEEE Transactions on Emerging Topics in Computing (*TETC*) plans a special issue on approximate and stochastic computing circuits, systems, and algorithms, slated for the third issue of 2016.

The last decade has seen renewed interest in nontraditional computing paradigms. Several of these paradigms leverage error resiliency by not requiring exactness in computing.

This special issue focuses on research into the novel design and analysis of two of these approaches: approximate and stochastic computing. Approximate computing is driven by energy efficiency, whereas stochastic computing achieves fault tolerance and area savings through randomness.

Articles are due **1 September 2015**. Visit www.computer.org/cms/Computer.org/transactions/cfps/cfp_tetcsi_asccsa.pdf to view the complete call for papers.

IEEE Transactions on Emerging Topics in Computing (*TETC*) plans a special issue on low-power image recognition, slated for late 2016.

Many of today's images and videos are captured using smartphones, and these devices' cameras can be used for multiple imaging applications, from high-fidelity location estimation to posture analysis. However, image processing is computationally intensive and consumes considerable energy, both of which present problems for resource-constrained mobile devices.

This special issue focuses on the intersection of image recognition and energy conservation. Articles should describe energy-efficient systems that perform object detection and recognition in images.

Because this issue aims to establish the state of the art in this area, all articles must include results achieved on a common core of datasets and training images, based on the same metrics.

Articles are due **1 September 2015**. Visit www.computer.org/cms/Computer.org/transactions/cfps/TETCLuBerg.pdf to view the complete call for papers.

IEEE Computer Graphics and Applications plans a May/June 2016 special issue on quality assessment and perception in computer graphics.

The computer graphics community has successfully exploited knowledge about the human visual system and its limitations for several years. This field is increasingly active across diverse research areas. Nonetheless, it has received relatively little attention thus far.

This special issue will focus on visual quality and perception for computer graphics. The guest editors are soliciting papers describing innovative, perception-related computer graphics techniques.

Articles are due **1 November 2015**. Visit www.computer.org/web/computingnow/cgacfp4 to view the complete call for papers.

IEEE Transactions on Emerging Topics in Computing (*TETC*) plans a special issue on defect and fault tolerance in VLSI and nanotechnology systems, slated for the last issue of 2016.

The scaling of CMOS devices and increased interest in emerging technologies make topics related to defect and fault tolerance in VLSI and nanotechnology systems increasingly important.

This special issue will explore design, manufacturing, test, reliability, and availability issues that are affected by manufacturing defects and system-operation faults.

Articles are due **1 December 2015**. Visit www.computer.org/cms/Computer.org/transactions/cfps/cfp_tetcsi_dft_vlsi_ns.pdf to view the complete call for papers.

IEEE Transactions on Emerging Topics in Computing (*TETC*) plans a special issue on "new paradigms in ad hoc, sensor, and mesh networks, from theory to practice," slated for the fourth issue of 2016.

Ad hoc, sensor, and mesh networks have attracted significant attention by academia and industry in the past decade. In recent years,

CALL AND CALENDAR

however, new paradigms in these areas have appeared as a result of the increase in the number and processing power of smartphones and other portable devices. Furthermore, new applications and emerging technologies—such as smart homes, body area networks, and the Internet of Things—have created new research challenges for ad hoc networks.

The focus of this special issue is on novel applications, protocols, and architectures; nontraditional measurement, modeling, analysis, and evaluation approaches; prototype systems; and experiments in ad hoc, sensor, and mesh networks.

Articles are due **1 December 2015**. Visit www.computer.org/cms/Computer.org/transactions/cfps/cfp_tetcsi_npahsmntp.pdf to view the complete call for papers.

IEEE Internet Computing plans a September/October 2016 special issue on cyber-physical security and privacy.

Cyber-physical systems (CPS) integrate computing and communication capabilities with the monitoring and control of physical systems such as power grids, water treatment plants, and factories. The widespread growth of wireless embedded sensors and actuators is creating new CPS applications and increasing the information infrastructure's role in control systems.

Many CPS applications are safety-critical. Their failure could cause irreparable harm to the physical system under control and the people who depend on it. These applications thus require strong security, as well as privacy protection in many cases.

To address these issues, the guest editors invite original research papers, including submissions from multiple interdisciplinary backgrounds.

Brief article descriptions are due **14 December 2015**. Articles are due **14 January 2016**. Visit www.computer.org/web/computingnow/iccfp5 to view the complete call for papers.

JULY 2015

7–10 July: ICAC 2015, 12th IEEE Int'l Conf. Autonomic Computing, Grenoble, France; http://icac2015.imag.fr

13–15 July: ICCCNT 2015, 6th Int'l Conf. Computing, Communications, and Networking Technologies, Irving, Texas; www.6icccnt.com

13–16 July: CBI 2015, 17th IEEE Conf. Business Informatics, Lisbon, Portugal; http://cbi2015.inesc.pt

AUGUST 2015

6–7 August: NAS 2015, 10th IEEE Int'l Conf. Networking, Architecture, and Storage, Boston; www.nas-conference.org

22–25 August: HotChips 2015, HotChips Symp., Cupertino, California; www.hotchips.org

24–26 August: FiCloud 2015, 3rd Int'l Conf. Future Internet of Things and Cloud, Rome; www.ficloud.org/2015

SEPTEMBER 2015

7–9 September: EISIC 2015, European Intelligence and Security Informatics Conf., Manchester, UK; www.eisic.eu

16–18 September: VS-Games 2015, 7th Int'l Conf. Virtual Worlds and Games for Serious Applications, Skövde, Sweden; www.his.se/en/Research/our-research/Conferences/VS-games-2015

29 September–1 October: ICSME 2015, 31st Int'l Conf. Software Maintenance and Evolution, Bremen, Germany; www.icsme.uni-bremen.de

OCTOBER 2015

12–15 October: STC 2015, 27th IEEE Software Technology Conf., Long Beach, California; http://conference.usu.edu/STC

19–21 October: SOCA 2015, 8th IEEE Intl Conf. Service Oriented Computing & Applications, Rome; www.dis.uniroma1.it/~soca2015/index.html

19–21 October: WIMOB 2015, 11th IEEE Intl Conf. Wireless & Mobile Computing, Networking, and Communications, Abu Dhabi, UAE; http://conferences.computer.org/wimob2015/#

25–30 October: VIS 2015, IEEE VisWeek 2015, Chicago; http://ieeevis.org

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

IEEE computer society

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.
COMPUTER SOCIETY WEBSITE: www.computer.org
OMBUDSMAN: To check membership status or report a change of address, call the IEEE Member Services toll-free number, +1 800 678 4333 (US) or +1 732 981 0060 (international). Direct all other Computer Society-related questions—magazine delivery or unresolved complaints—to help@computer.org.
CHAPTERS: Regular and student chapters worldwide provide the opportunity to interact with colleagues, hear technical experts, and serve the local professional community.
AVAILABLE INFORMATION: To obtain more information on any of the following, contact Customer Service at +1 714 821 8380 or +1 800 272 6657:

- Membership applications
- Publications catalog
- Draft standards and order forms
- Technical committee list
- Technical committee application
- Chapter start-up procedures
- Student scholarship information
- Volunteer leaders/staff directory
- IEEE senior member grade application (requires 10 years practice and significant performance in five of those 10)

PUBLICATIONS AND ACTIVITIES
***Computer*:** The flagship publication of the IEEE Computer Society, *Computer*, publishes peer-reviewed technical content that covers all aspects of computer science, computer engineering, technology, and applications.
Periodicals: The society publishes 13 magazines, 18 transactions, and one letters. Refer to membership application or request information as noted above.
Conference Proceedings & Books: Conference Publishing Services publishes more than 175 titles every year.
Standards Working Groups: More than 150 groups produce IEEE standards used throughout the world.
Technical Committees: TCs provide professional interaction in more than 45 technical areas and directly influence computer engineering conferences and publications.
Conferences/Education: The society holds about 200 conferences each year and sponsors many educational activities, including computing science accreditation.
Certifications: The society offers two software developer credentials. For more information, visit www.computer.org/certification.

NEXT BOARD MEETING
1–5 June 2015, Atlanta, GA, USA

IEEE

revised 27 Jan. 2015

CAREER OPPORTUNITIES

NAJANAJA LTD. seeks a PLM Technical Analyst for Irving, TX to gather/analyze requirements, create technical design for Windchill PLM projects, create proof of concepts using web stack, write technical docs, lead onshore/offshore team. Tools used - J2EE, Windchill PLM (Product Lifecycle Management), Configuration/Change/CAD Management. Master's degree, Knowledge of SDLC & Windchill PLM reqd. Apply to jobs@najanaja.com or NajaNaja Ltd, 625 Heritage Ln, Flower Mound, TX 75022.

TECHNICAL CONSULTANTS sought by GSPANN Technologies, Inc. in Milpitas, CA. Bach's deg in Comp Sci, Engr, or rltd w/ 5 yrs exp. Provide customers w/ service, support, problem solving, & escalation. Understand/assess web dvlpmt, production reqmts, goals, & problems/issues to propose/provide solutions. Provide knowl transfer, mentoring, & training. Dsgn/implmt workflows as well as solutions using Content Mgmt Systm. Build/deliver presentations/solution demonstrations. Also resp for build, delivery, & mgmt of proof of concept & pilot projects. Mail resumes to GSPANN, 362 Fairview Way, Milpitas, CA 95035, Attn: HR.

INTERSHOP DEVELOPER, Brightstar US, Inc. (Libertyville, IL) to Assist in the dvlpmt of Brightstar's next generation e-Commerce processes & svcs for customers. Align to dsgn reqmts of the Brightstar systms (applicable) to clients are managed & successful for applics locally, regionally or globally dvlpd. Identify technical issues & provide resolution steps. Dsgn & dvlp systms using Intershop platform. Reqmts: Bachelor's deg, or foreign equiv, in Comp Sci or rltd field. Min of 5 yrs exp in job or of overall s/ware dvlpmt exp w/a min of 3 yrs working w/Intershop Enfinity 6.x or Intershop 7. Solid understanding of SQL, Oracle SWL preferred. Exp w/project implmtn using technologies such as Java, JSP, Servlet, JDBC, XML/XSLT, & HTML. Multiple positions available. Resumes to: Sapna Rao, HR Mgr, 850 Technology Way, Libertyville, IL 60048.

PROGRAMMER ANALYST: Design, develop, integrate, test and implement application software utilizing knowledge in Agile environment utilizing technologies like .Net Framework 4.0/4.5, C#.Net, ASP.Net Web Applications, Web Services, Ado.Net, WinForms, LINQ, Visual Studio .Net, Sql Server Management Studio, JavaScript,VB Script, T-Sql,SQL server, Oracle, HTML, XML,CSS, TFS, IIS,MS-Unit. Must be willing to travel & reloc. Requires MS in comp sci, eng or rel. Mail resumes to Code Ace Solutions Inc. 50 Cragwood Road, Ste 217, South Plainfield, NJ 07080.

TECHNICAL LEAD/TEAM 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

0018-9162/15/$31.00 © 2015 IEEE

App, Comp Sci, or related w/5 yrs of progressive exp or a Master deg or the foreign equiv in Comp App, Comp Sci or related w/1 yr exp managing & leading a team of 3 developers in application analysis, design, development, implementation, testing and performance tuning of software applications through full product development life cycle and release process. Mentor the junior team members by providing technology knowledge transition. Provide comprehensive consultation to business unit & IT management and staff at the highest technical level on all phases of application. Provide subject matter expertise and implement the code using following tools/technologies: Java/ J2EE, JSP, Servlets, HTML, XML, JavaScript, JSF, Spring, Hibernate, EJB, UML, Rational Rose, RSA, JDBC, Oracle, DB2, Websphere MQ, JMS, JAAS, Web services, JAXB, JUnit , Portlets, WPS, WAS , Tomcat, RAD, Eclipse, SubVersion, VSS, and SCRUM Tools. Send resume: Indotronix Int.l Corp., Recruiting (AS), 331 Main St, Poughkeepsie, NY 12601.

Samsung Research America, Inc.
has the following opportunities
(various levels) available
in **San Francisco, CA**:

UI Engineer, Staff 2
(Ref# SF15E01)

Research Engineer, Staff 1
(Ref# SF15E02)

Industrial Designer
(Ref# SF15E03)

Research Engineer, Sr.
(Ref# SF15E04)

Specific requirements apply. All of these positions will involve developing technologies for company's computer, digital television, mobile telephone, printer, or other electronic products.

Mail your resume referencing job title and Ref# to farhat.k@samsung.com.

JUNIPER NETWORKS

Juniper Networks is recruiting for our Sunnyvale, CA office:

Resident Engineer #26985: Provide problem resolution related to Juniper routers and their integration to customers' networks. Carry out testing of new features and functionalities required by customers in a laboratory environment. May work at other undetermined worksites in the U.S. and relocation may be required.

Software Engineer #29393: Perform software development tasks related to software releases. Perform software configuration management administration and data mining for key performance indices.

Resident Engineer #33335: Provide post-sales support of Juniper's networking products and carry out testing of new designs, features, functionality, and software releases as required by the customer in a laboratory environment.

Resident Engineer Staff #2509: Support design, deployment, and operational readiness of Juniper network routing, switching and security products within the customer infrastructure. Troubleshoot equipment and network problems. May work at other undetermined worksites throughout the US and relocation may be required.

Software Engineer #17561: Design, develop, assist in deployment and maintenance of networking, kernel, TCP/IP solutions on Unix environment.

Software Engineer Staff #12921: Perform service gateway software design and maintenance, including designing, implementing, testing and debugging security features in flow based packet forwarding engine of Juniper's NGSRX platforms.

QA Engineer #15048: Perform manual and automated functional tests, write test plans on Juniper networks products, report defects in features and implement protocols such as TCP/IP and HTTP.

Software Engineer #32599: Define User Interface concepts and guide development that culminates in an excellent user experience. Collaborate with visual designers and corporate marketing to integrate visual and behavioral elements. Lead design process, drive decisions, create schedules and track user experience outcomes.

Test Engineer Staff #16163: Prepare and execute test plans for automated testing of product features. Architect and develop the infrastructure of company's testing program for contrail SDN functionalities.

Software Engineer Staff #24491: Develop detailed software functional and design specifications for new software features on company products. Work with Product Management team to get requirements for new software features.

Technical Support Engineer #17362: Deliver hands-on technical assistance and troubleshooting skills in resolving critical customer hardware and software issues. Identify and lead service support initiatives to improve the supportability of processes, products, and systems.

Juniper Networks is recruiting for our Herndon, VA office:

Technical Support Engineer Staff #15227: Provide high level expertise on Juniper products and deliver in-depth diagnostics and root-cause analysis for network impacting issues on Juniper routing products (internet backbone routers).

Technical Support Engineer Staff #15388: Provide high level expertise on Juniper products and deliver in-depth diagnostics and root-cause analysis for network impacting issues on Juniper routing products (internet backbone routers).

Juniper Networks is recruiting for our San Francisco, CA office:

Software Engineer #32224: Develop cloud-based security services. Work alongside technical leads on critical aspects of security software testing.

Mail single-sided resume with job code # to
Juniper Networks
Attn: MS 1.4.251
1133 Innovation Way
Sunnyvale, CA 94089

CLASSIFIED LINE AD SUBMISSION DETAILS:

Rates are $425.00 per column inch ($640 minimum). Eight lines per column inch and average five typeset words per line. Send copy at least one month prior to publication date to:

DEBBIE SIMS
Classified Advertising
Computer
10662 Los Vaqueros Circle,
Los Alamitos, CA 90720
(714) 816-2138; f (714) 821-4010
Email: dsims@computer.org

CAREER OPPORTUNITIES

Microsoft **Microsoft Corporation currently has the following openings (job opportunities available at all levels, including Principal, Senior and Lead levels):**

REDMOND, WA

Applied Scientist: Utilize knowledge in applied statistics and mathematics to handle large amounts of data using various tools. http://bit.ly/MSJobs_Data_Applied_Science

Artists, Art Leads and Animators: Responsible for designing and creating art assets that meet or exceed industry standards for quality while supporting Microsoft Game Studio (MGS) business goals. http://bit.ly/MSJobs_Art

Business Managers and Business Development Managers/Business Development and Strategy Analyst Manager: Develop business opportunities for sales of software and services. http://bit.ly/MSJobs_Business_Development

Consultants: Deliver design, planning, and implementation services that provide IT solutions to customers and partners. Roving Employee—requires travel up to 100% with work to be performed at various unknown worksites throughout the U.S. Telecommuting Permitted. http://bit.ly/MSJobs_Technical_Delivery

Consultants: Deliver design, planning, and implementation services that provide IT solutions to customers and partners. Requires domestic and international travel up to 25%. http://bit.ly/MSJobs_Technical_Delivery

Content Developer/Engineer: Responsible for the design, development, deployment, vision, and business strategy for content creation, acquisition, production, editorial, and publishing activities at Microsoft. http://bit.ly/MSJobs_Content_Publishing

Data Scientist: Manipulate large volumes of data, create new and improved techniques and/or solutions for data collection, management and usage. http://bit.ly/MSJobs_Data_Applied_Science

Design Verification/Validation Engineers: Responsible for ensuring the quality of Microsoft hardware products. http://bit.ly/MSJobs_Hardware_Design_Verification_Eng

Evangelists/Technical Evangelists: Secure future growth of the Microsoft platform by engaging a community of customers, partners, and academics to embrace and adopt Microsoft technology. http://bit.ly/MSJobs_Tech_Evangelist

Evangelists/Technical Evangelists: Secure future growth of the Microsoft platform by engaging a community of customers, partners, and academics to embrace and adopt Microsoft technology. Requires travel throughout the U.S. up to 25%. http://bit.ly/MSJobs_Tech_Evangelist

Game/Systems Designer: Create design documents for multiple major features on large projects and the entire design on smaller projects, ensuring consistency of design. http://bit.ly/MSJobs_Game_Design

Hardware Dev. or Design Engineers, Hardware Engineers, Electrical Engineers, and Design Engineers (all levels, including Leads and Managers): Design, implement and test computer hardware. http://bit.ly/MSJobs_Hardware_Dev_Eng

International Project Engineers/Managers: Ensure the successful internationalization or localization of software components for foreign markets. http://bit.ly/MSJobs_Intl_Proj_Eng

Machine Learning Scientist: Design and deliver general and/or domain-specific machine learning algorithms and systems. http://bit.ly/MSJobs_Data_Applied_Science

Premier Field Engineers: Provide technical support to enterprise customers, partners, internal staff or others on mission critical issues experienced with Microsoft technologies. Roving Employee—requires travel up to 100% with work to be performed at various unknown worksites throughout the U.S. Telecommuting permitted. http://bit.ly/MSJobs_Support_Delivery

Researchers/Scientists: Conduct research and lead research collaborations that yield new insights, theories, analyses, data, algorithms, and prototypes. http://bit.ly/MSJobs_Research

Service Engineers/Managers, Service Operations Engineers, and Systems/Operations Engineers: Plan, architect, deploy and/or support complex client/server or database software systems. (http://bit.ly/MSJobs_Service_Engineering) (http://bit.ly/MSJobs_IT_Serv_Eng) (http://bit.ly/MSJobs_IT_Serv_Ops)

Solution Managers: Identify and analyze internal client and partner business needs, and translate needs into business requirements and value-added solutions and solution roadmaps. http://bit.ly/MSJobs_IT_Solution_Mgmt

Solutions Sales Professional/Specialist: Enhance the Microsoft customer relationship from a capability development perspective by articulating the value of our services and solutions and identifying competition gaps in targeted accounts. http://bit.ly/MSJobs_Solution_Sales

Support Engineers / Escalation Engineers: Provide technical support on issues experienced with Microsoft technologies. http://bit.ly/MSJobs_Support_Eng

Designers: Develop user interface and user interaction designs, prototypes and/or concepts for business productivity, entertainment or other software or hardware applications. http://bit.ly/MSJobs_Design

Design Researchers: Develop user interface and user interaction designs, prototypes and/or concepts for business productivity, entertainment or other software or hardware applications. http://bit.ly/MSJobs_Design_Research

Business Process Manager: Responsible for the design, implementation, and release of programs or projects. http://bit.ly/MSJobs_Fin_Plan_Analy_Contr

Network Planner: Contribute to the network strategy and planning team for online services infrastructure. Travel up to 50% with work to be performed at various unknown worksites throughout the U.S. Telecommuting permitted. (http://bit.ly/MSJobs_1493) (http://bit.ly/MSJobs_1339)

Site Reliability Engineer: Plan, architect, deploy and/or support complex client/server or database software systems. (http://bit.ly/MSJobs_1392) (http://bit.ly/MSJobs_1421)

Compliance Service Engineer - Dynamics or Other: Define and deliver on the IT auditing/compliance needs required to operate Microsoft Online systems. http://bit.ly/MSJobs_1882

Franchise Business Manager-Operating Systems Engr Grp or Other: Responsible for optimizing product line(s) within a video game franchise. http://bit.ly/MSJobs_1873

Associate Architect MCS - COO or Other: Identify and analyze internal client and partner business needs and translate needs into business requirements. Requires domestic and international travel up to 50%. http://bit.ly/MSJobs_1972

Industrial Designer II - Devices Group or Other: Design for the Xbox core product, accessories, incubation, and vision hardware programs. http://bit.ly/MSJobs_1240

Electrical Engineer - Devices Group or Other: Design, develop and test antennas and RF circuitry in consumer electronics devices. http://bit.ly/MSJobs_1239

Security Analyst II - Operating Systems Engineer Group or Other: Design service specifications that meet business requirements. http://bit.ly/MSJobs_1233

Electrical Engineer - Devices Group or Other: Responsible for working with team members to evaluate audio and Ethernet options for performance, functionality, stability, cost, and risk. http://bit.ly/MSJobs_1418

Account Manager - Biz Development and Evangelism GRP or Other: Responsible for bringing games to all Microsoft platforms by utilizing strategic partnerships. http://bit.ly/MSJobs_1382

Audio Implementer - Operating Systems Engr Grp or Other: Work with team members to architect the audio content pipeline. http://bit.ly/MSJobs_1381

Data and Applied Scientist - Applications and Services Engineer Group or Other: Responsible for designing, prototyping, implementing and testing machine learning models and algorithms. http://bit.ly/MSJobs_1219

Advanced Analytics Manager – Devices Group or Other: Requires mastery of Production Planning Dashboard (PPD) to conduct data analysis, forecast, and generate reports. http://bit.ly/MSJobs_1457

Signal Integrity Test Engineer - Devices Group or Other: Perform Signal Integrity testing for high-speed signal interface standards such as DDR3, USB2/3, DP, HDMI, SATA. http://bit.ly/MSJobs_1456

Electrical Engineer - Devices Group or Other: Design, implement and test computer hardware products that add strategic value to the company. This position requires international and domestic travel up to 50%. http://bit.ly/MSJobs_1454

Systems Designer – Operating Systems Engr Grp or Other: Work with the Lead Designer to

understand the vision of the product and design gameplay systems. http://bit.ly/MSJobs_1415

Electrical Engineer II - Devices Group or Other: Work with team members to evaluate memory options for performance, functionality, stability, cost, and risk. http://bit.ly/MSJobs_1423

Sourcing Engineer – Devices Group or Other: Required to specialize in manufacturing processes and total cost management of assigned categories. http://bit.ly/MSJobs_1383

Premier Field Engineer – COO or Other: Work closely with Microsoft's Client OEM partners to advise on engineering enablement and device health. http://bit.ly/MSJobs_1388

Electrical Engineer: Design, implement and test computer hardware products. This position requires domestic and international travel up to 25%. (http://bit.ly/MSJobs_1452) (http://bit.ly/MSJobs_1465) (http://bit.ly/MSJobs_1510) (http://bit.ly/MSJobs_1354)

Producer – Operating Systems Engineer Group or Other: Serve as the production lead for game development projects that include external partners. This position requires domestic, regional and international travel up to 25%.http://bit.ly/MSJobs_1491

SW/FW Engineer: Responsible for developing firmware for embedded devices. This position requires domestic, regional and international travel up to 25%. (http://bit.ly/MSJobs_1534) (http://bit.ly/MSJobs_1542)

Risk Integration Manager- Operating Systems Engineer Group or Other: Responsible to bringing new partners / new scenarios to Risk Management platform. http://bit.ly/MSJobs_1547

Senior Applied Researcher - Applications and Services Engr Grp or Other: Solve ambiguous problems, design and implement machine learning algorithms and analyze big data. http://bit.ly/MSJobs_1444

Sourcing Manager, Optics - Devices Group or Other: Support the Microsoft global procurement process by managing the development, implementation, and alignment of global sourcing strategies. Requires international travel up to 50%. http://bit.ly/MSJobs_1754

Principal Group Service Engineering Manager - GFS - Bandwidth or Other: Plan, architect, deploy and/or support complex client/server or database software systems. Requires international travel up to 25%. http://bit.ly/MSJobs_1386

Tech Evangelist – Developer Experience and Evangelism or Other: Secure future growth of the Microsoft platform by engaging a community of customers, partners, and academics. Requires international and domestic travel up to 25%. http://bit.ly/MSJobs_1344

Senior Art PM - Studios or Other: Coordinate program development of high-end, three dimensional art asset creation and computer software applications. Requires international and domestic travel up to 25%. http://bit.ly/MSJobs_1338

UX Researcher (Design Researcher) - OSG Core PM or Other: Develop user interface and user interaction designs, prototypes and/or concepts. Requires international and domestic travel up to 25%. http://bit.ly/MSJobs_1281

Service Engineer II - GFS - Networking and Shared Service or Other: Plan, architect, deploy and/or support complex client/server or database software systems. Telecommuting permitted. http://bit.ly/MSJobs_1913

Security Project Manager - GFS - Security or Other: Support and drive all physical security systems project implementation, testing, commissioning, and acceptance. Requires domestic and international travel up to 50%. http://bit.ly/MSJobs_1899

Category Manager Windows - Retail S&M Opex or Other: Responsible for reinvigorating the PC market, succeeding in tablets and improving the PC purchase experience. Requires Domestic and International travel up to 25%. http://bit.ly/MSJobs_1892

Senior Architect, Imaging - MDG Phones or Other: Design, implement and test computer hardware products that add strategic value to the company. Requires international and domestic travel up to 25%. Telecommuting permitted. http://bit.ly/MSJobs_1943

Hardware Engineer (Reliability Engineer): Design, implement and test computer hardware products that add strategic value. Requires international and domestic travel up to 25%. http://bit.ly/MSJobs_1982

Senior International Site Manager: Ensure the successful internationalization or localization of software components for foreign markets. http://bit.ly/MSJobs_1871

Security Analyst: Gather system and host data across Microsoft's network. http://bit.ly/MSJobs_1529

Business Analytics & Insights Manager: Apply principles and techniques of finance, marketing, and/or business to articulate business problems. http://bit.ly/MSJobs_1506

Senior Producer: Coordinate the development, execution, and release of gaming products and services. http://bit.ly/MSJobs_1714

Senior IT Service Engineer: Design service specifications that meet business requirements by analyzing solution priorities. http://bit.ly/MSJobs_1499

Senior Business Process Manager – CFS – G&A or Other: Analyze, design and deliver business information and financial reporting solutions for internal business groups http://bit.ly/MSJobs_1621

Senior Security Analyst: Design, develop, and conduct global security investigations and tests on attacks targeting Microsoft. http://bit.ly/MSJobs_1710

Principal IT Enterprise Architect: Responsible for working with team members to provide business-to-technology mapping for moderate risk, moderate complexity software development engagements. Requires travel throughout the United States up to 10%. http://bit.ly/MSJobs_1681

Business Analytics Specialist - Shared CSS Delivery or Other: Responsible for the design, implementation, and release of programs or projects. http://bit.ly/MSJobs_1288

Principal IT Enterprise Architect-Corporate Functions IT or Other: Create information strategies and models and master data management to enable businesses to consolidate. http://bit.ly/MSJobs_1293

Sales Excellence Manager – SMB Coverage or Other: Responsible for developing business opportunities and driving business processes for sales of software or services. http://bit.ly/MSJobs_1519

Solution Specialist DCMOD-EPG Core Account Coverage or Other: Enhance the Microsoft customer relationship from a capability development perspective. Requires travel up to 25% with work to be performed at various unanticipated worksites throughout the U.S. http://bit.ly/MSJobs_1525

C+E Business Intelligence Developer - Corp-Fin-G&A or Other: Responsible for developing or testing computer software applications, systems or services. http://bit.ly/MSJobs_1296

Solution Specialist Devices-EPG Core Account Coverage or Other: Enhance the Microsoft customer relationship from a capability development perspective. Requires travel up to 50% with work to be performed at various unanticipated worksites throughout the U.S. http://bit.ly/MSJobs_1299

Senior Business Planner-Product Marketing or Other: Responsible for developing product licensing strategy for products and services. http://bit.ly/MSJobs_1235

Sales Excellence Manager – EPG Core Account Coverage or Other: Support the Productivity business with data gathering and business insights. http://bit.ly/MSJobs_1635

Senior Learning and Development Specialist – WW M&O or Other: Design, develop and deliver tailored training programs for the field sales and services organization. Requires domestic and international travel up to 25%. Telecommuting permitted. http://bit.ly/MSJobs_1462

Senior Solutions Sales Specialist, Customer Advocacy and Technology Management - EPG Core Account Coverage or Other: Enhance the Microsoft customer relationship from a capability development perspective. Requires domestic and international travel up to 25%. Telecommuting permitted. http://bit.ly/MSJobs_1171

Software Architect-Transform the Datacenter Pillar or Other: Engage with strategic enterprise customers to solve their business needs. Requires travel throughout the U.S. up to 25%. http://bit.ly/MSJobs_1727

Senior Business Development Manager - Business Development or Other: Develop business opportunities for sales of software and services. Requires travel throughout the U.S. up to 25% of the time. http://bit.ly/MSJobs_1772

Multiple job openings are available for each of these categories. To view detailed job descriptions and minimum requirements, and to apply, visit the website address listed. EOE.

CAREER OPPORTUNITIES

Microsoft® **Microsoft Corporation currently has the following openings (job opportunities available at all levels, including Principal, Senior and Lead levels):**

MOUNTAIN VIEW, CA

Applied Scientist: Utilize knowledge in applied statistics and mathematics to handle large amounts of data using various tools. http://bit.ly/MSJobs_Data_Applied_Science

Data Scientist: Manipulate large volumes of data, create new and improved techniques and/or solutions for data collection, management and usage. http://bit.ly/MSJobs_Data_Applied_Science

Design Verification/Validation Engineers: Responsible for ensuring the quality of Microsoft hardware products. http://bit.ly/MSJobs_Hardware_Design_Verification_Eng

Machine Learning Scientist: Design and deliver general and/or domain-specific machine learning algorithms and systems. http://bit.ly/MSJobs_Data_Applied_Science

Researchers/Scientists: Conduct research and lead research collaborations that yield new insights, theories, analyses, data, algorithms, and prototypes. http://bit.ly/MSJobs_Research

Service Engineers/Managers, Service Operations Engineers, and Systems/Operations Engineers: Plan, architect, deploy and/or support complex client/server or database software systems. (http://bit.ly/MSJobs_Service_Engineering) (http://bit.ly/MSJobs_IT_Serv_Eng) (http://bit.ly/MSJobs_IT_Serv_Ops)

Principal Electrical Engineer - Devices Group or Other: Responsible for developing product and system level design. http://bit.ly/MSJobs_1230

Systems Engineer – Devices Group or Other: Work with team members to implement hardware and embedded software across a variety of products and technologies. http://bit.ly/MSJobs_1406

Technology Solutions Professional-EPG Core Solution Specialist - CnE or Other: Drive product win rates by proving the value of products to customers and partners. http://bit.ly/MSJobs_1672

SAN FRANCISCO, CA

Sr. Technical Evangelist-DX Evangelist Spec-Corp SMSG (S&T) or Other: Secure future growth of the Microsoft platform by engaging a community of customers, partners, and academics. Telecommuting from home office permitted. Travel required up to 35% to undetermined client sites throughout the U.S. http://bit.ly/MSJobs_1177

SUNNYVALE, CA

Applied Scientist: Utilize knowledge in applied statistics and mathematics to handle large amounts of data using various tools. http://bit.ly/MSJobs_Data_Applied_Science

Data Scientist: Manipulate large volumes of data, create new and improved techniques and/or solutions for data collection, management and usage. http://bit.ly/MSJobs_Data_Applied_Science

Machine Learning Scientist: Design and deliver general and/or domain-specific machine learning algorithms and systems. http://bit.ly/MSJobs_Data_Applied_Science

Researchers/Scientists: Conduct research and lead research collaborations that yield new insights, theories, analyses, data, algorithms, and prototypes. http://bit.ly/MSJobs_Research

Service Engineers/Managers, Service Operations Engineers, and Systems/Operations Engineers: Plan, architect, deploy and/or support complex client/server or database software systems. (http://bit.ly/MSJobs_Service_Engineering) (http://bit.ly/MSJobs_IT_Serv_Eng) (http://bit.ly/MSJobs_IT_Serv_Ops)

Design Researchers: Develop user interface and user interaction designs, prototypes and/or concepts for business productivity, entertainment or other software or hardware applications. http://bit.ly/MSJobs_Design_Research

SAN DIEGO, CA

Hardware Dev. or Design Engineers, Hardware Engineers, Electrical Engineers, and Design Engineers (all levels, including Leads and Managers): Design, implement and test computer hardware. http://bit.ly/MSJobs_Hardware_Dev_Eng

Specialist, Cellular Modem Software - NDS Engineering Group or Other: Responsible for configuration, debugging and implementation of cellular modem software. http://bit.ly/MSJobs_1988

Design Validation Engineer - MDG Phones or Other: Perform wireless device testing in lab, debug device and test system related issues. http://bit.ly/MSJobs_2008

Hardware Engineer: Design, implement and test hardware products that add strategic value. Domestic and International travel required up to 25%. (http://bit.ly/MSJobs_1999) (http://bit.ly/MSJobs_1971)

GM, Hardware- Surface or Other: Lead the overall Strategy and Planning Team for the Surface organization within Microsoft. Requires domestic and international travel up to 25%. http://bit.ly/MSJobs_2068

FORT LAUDERDALE, FL

Business Managers and Business Development Managers/Business Development and Strategy Analyst Manager: Develop business opportunities for sales of software and services. http://bit.ly/MSJobs_Business_Development

Sales Excellence Manager – EPG Core Account Coverage or Other: Measure and track an efficient Segment Sales Excellence function spanning scorecard, Rhythm of Business, and business management. http://bit.ly/MSJobs_1619

Chief Security Advisor - AMS-EPG PS or Other: Formulate and promote external company strategy regarding computer security issues. Requires domestic and international travel up to 40%. Telecommuting is permitted. http://bit.ly/MSJobs_1496

Partner Technical Consultant-Global Business Support or Other: Help Microsoft partners accelerate their sales cycles and design, develop and deploy solutions. http://bit.ly/MSJobs_1313

Account Executive-Field Global Sales or Other: Responsible for selling online advertising solutions. http://bit.ly/MSJobs_1394

Account Executive, Advertising & Online New Markets LATAM - Field Global Sales or Other: Develop business, grow revenue and increase overall sales partner satisfaction and loyalty. Requires local travel up to 50%. Telecommuting permitted. http://bit.ly/MSJobs_1429

Technical Advisor-Consumer or Other: Provide technical advice and support on issues experienced with Microsoft technologies. Telecommuting permitted. http://bit.ly/MSJobs_1704

DOWNERS GROVE, IL

Technical Account Manager: Assure productive use of Microsoft technologies, focusing on delivery quality through planning and governance. Requires travel up to 40% to various client sites in the Midwest region. Telecommuting permitted. http://bit.ly/MSJobs_1502

CAMBRIDGE, MA

Data Scientist: Manipulate large volumes of data, create new and improved techniques and/or solutions for data collection, management and usage. http://bit.ly/MSJobs_Data_Applied_Science

International Project Engineers/Managers: Ensure the successful internationalization or localization of software components for foreign markets. http://bit.ly/MSJobs_Intl_Proj_Eng

Service Engineers/Managers, Service Operations Engineers, and Systems/Operations Engineers: Plan, architect, deploy and/or support complex client/server or database software systems. (http://bit.ly/MSJobs_Service_Engineering) (http://bit.ly/MSJobs_IT_Serv_Eng) (http://bit.ly/MSJobs_IT_Serv_Ops)

CHARLOTTE, NC

Support Engineers / Escalation Engineers: Provide technical support on issues experienced with Microsoft technologies. http://bit.ly/MSJobs_Support_Eng

WASHINGTON DC

Business Managers and Business Development Managers/Business Development and Strategy Analyst Manager: Develop business opportunities for sales of software and services. http://bit.ly/MSJobs_Business_Development

Senior Business Strategy Manager: Analyze market and technology trends and data to develop creative solutions to key business strategy issues. Requires travel throughout the United States up to 10%. http://bit.ly/MSJobs_1518

Multiple job openings are available for each of these categories. To view detailed job descriptions and minimum requirements, and to apply, visit the website address listed. EOE.

Microsoft

Microsoft Corporation currently has the following openings (job opportunities available at all levels, including Principal, Senior and Lead levels):

NEW YORK, NY

Support Engineers / Escalation Engineers: Provide technical support on issues experienced with Microsoft technologies. http://bit.ly/MSJobs_Support_Eng

Solution Specialist DC-EPG Core Account Coverage or Other: Enhance the Microsoft customer relationship from a capability development perspective. Requires travel up to 20% with work to be performed at various unanticipated worksites throughout the U.S. http://bit.ly/MSJobs_1283

Premier Field Engineer-Global Business Support or Other: Provide technical support on mission critical issues experienced with Microsoft technologies. Requires travel up to 50% with work to be performed at various unknown worksites throughout the U.S. Telecommuting permitted. http://bit.ly/MSJobs_1349

Solutions Sales Specialist – EPG Core Account Coverage or Other: Enhance the Microsoft customer relationship from a capability development perspective. Telecommuting permitted. http://bit.ly/MSJobs_1653

FARGO, ND

Technical Account Managers: Assure productive use of Microsoft technologies, focusing on delivery quality through planning and governance. http://bit.ly/MSJobs_Delivery_Relationship_Mgmt

IRVING, TX

Support Engineers / Escalation Engineers: Provide technical support on issues experienced with Microsoft technologies. http://bit.ly/MSJobs_Support_Eng

HOUSTON, TX

Technical Account Manager: Assure productive use of Microsoft technologies, focusing on delivery quality through planning and governance. Requires travel throughout the U.S. up to 75%; telecommuting permitted. http://bit.ly/MSJobs_1868

LOS ANGELES, CA

Technical Account Manager: Assure productive use of Microsoft technologies, focusing on delivery quality through planning and governance. Requires travel throughout the U.S. up to 25%. http://bit.ly/MSJobs_1540

EDINA, MN

Technical Account Manager – Premier COGS or Other: Assure productive use of Microsoft technologies, focusing on delivery quality through planning and governance. Requires travel up to 75% to various client sites in the Minneapolis region. Telecommuting permitted. http://bit.ly/MSJobs_1630

BENTONVILLE, AR

Senior Premier Field Engineer - Global Business Support or Other: Provide technical support to enterprise customers, partners, internal staff or others on mission critical issues. Requires travel up to 50% with work to be performed at various unknown worksites throughout the U.S. http://bit.ly/MSJobs_1301

ISELIN, NJ

Premier Field Engineer-Global Business Support or Other: Provide technical support to enterprise customers, partners, internal staff or others on mission critical issues. Requires travel up to 75% with work to be performed at various unknown worksites throughout the U.S. http://bit.ly/MSJobs_1251

Solution Sales Professional, PDW-EPG Core Account Coverage or Other: Enhance the Microsoft customer relationship from a capability development perspective. Requires travel to various unanticipated locations up to 50% of the time. Telecommuting permitted. http://bit.ly/MSJobs_1752

AUSTIN, TX

Account Technology Strategist-EPG Core Account Coverage or Other: Provide pre-sales technical and architectural support for sales of software, solutions, and related products. http://bit.ly/MSJobs_1483

Multiple job openings are available for each of these categories. To view detailed job descriptions and minimum requirements, and to apply, visit the website address listed. EOE.

SOFTWARE DEVELOPER - Assist in development of software for cloud based scheduling systems and desktop management suits for multiple industries. Consult with clients to determine business requirements. Design and execute tests to ensure quality control. Master's degree in Computer Science, Systems Eng. or related field, or foreign equivalent. Must be proficient in C#.Net; ASP.Net; SQL Server; Java script; Visual Studio; Win Forms; Visual Basic; testing frameworks such as Selenium, Sikuli, and HP Quality Center; Hudson Continuous Integration Server; and GIT. Resumes to: Elite Software Inc., 4001 W. Newberry Rd, Gainesville, FL 32607.

SIEMENS PLM SOFTWARE INC. has an opening in Troy, MI for Sr. Application Engineer to support technical sales, determine customer specs & demo related LMS Imagine.Lab software to customer. Requires 25% domestic/international travel. Email resumes to PLMCareers@ugs.com & refer to Req#143906. EOE

SPLUNK INC. has the following job opportunities in San Francisco, CA: Senior Software Engineer in Test (REQ#9ASSS8). Test, design & develop automated testing framework for distributed sys upgrades. Senior UI Software Engineer (REQ#9CXTSZ). Design & implement new front end for Co.'s Mobile Intelligence product. Refer to Req# & mail resume to Splunk Inc., ATTN: J. Aldax, 250 Brannan Street, San Francisco CA 94107.

PROGRAMMER ANALYST: Design, develop, test & implement Web and Client/Server Technologies using knowledge in C#.Net, ASP.Net, ASP.Net MVC, WCF Services (SOAP & REST), HTML5, CSS3, XSD, XML, SQL Server 2008 R2/2012,SSRS, SSIS, JavaScript, JQuery, ADO.Net, ADO.Net Entity framework, Windows XP/VISTA/7/8 Visual Studio 2010,2012. Must be willing to travel & relocate. Requires MS in Computer science, Engineering or related. Mail resumes to Strategic Resources International, 777 Washington Rd, Ste 2, Parlin, NJ 08859.

SYSTEMS ENGINEER (Mult. Openings) Allen, TX, Nexius Insight, Inc. W/a BS in Comp Info Sys. or rltd + 30 mos exp as ITO Service Delivery Consult. or rltd. Dsgns & dvlps solutions to complex applics prgms, syst. admin. issues, or network concerns. Dsgns, modifies, dvlps, writes, & implmts s/ware prgmg apps. to adhere to dsgns supporting internal bus. reqs. for SalesForce platform. Standardizes qlty assurance procedures for s/ware apps. Identifies s/ware reqs. & works w/suppliers in meeting customer objectives by implmtg changes. Supports &/or installs s/ware apps. in SalesForce. Coords a team of s/ware dvlprs in SalesForce s/ware dvlpmt. Mail resumes to Nexius Insight, Inc., 1301 Central Expressway S., Ste 200, Allen, TX 75013. Attn: HR

CLOUDERA, INC. is recruiting for our Palo Alto, CA office: Software Engineer: Architect, design, develop and test large distributed systems based on the Hadoop ecosystem capable of processing multiple petabytes of data. Mail resume w/job code #36215 to: Cloudera, Attn.: HR, 1001 Page Mill Rd., Bldg. 2, Palo Alto, CA 94304.

CAREER OPPORTUNITIES

NEW YORK UNIVERSITY ABU DHABI POSTDOCTORAL RESEARCHER established the Center for Interdisciplinary Studies in Security and Privacy (CRISSP-AD) to meet the growing challenges that are faced in securing networked information technology systems that have become pervasive and address broader political, economic and policy issues that help understand and mitigate cyber risk. CRISSP-AD is seeking outstanding postdoctoral candidates to join our team to further support its research activities. The areas of expertise aimed at, but not limited to, are: cyber threat intelligence and network analytics; hardware and critical-infrastructure security; and user-centric security. Candidates must hold (or be close to completing) a Ph.D. in Computer Science, Electrical and Computer Engineering or quantitative social sciences. Ph.D. holders with a strong publication record and hands-on skills are encouraged to apply. Applicants with a M.S. in the same fields may apply for a Research Assistant or Research Engineer position. To be considered submit a cover letter, curriculum vitae, and a statement of research interests online at: http://nyuad.nyu.edu/en/about/careers/faculty-positions.html

SENIOR SIEBEL CRM CONSULTANT -Cambridge MA. Design, develop, test and deploy Siebel CRM technology solutions. Analyze customer requirements to design technical architecture for business process solutions. Translate business requirements into Siebel CRM functional and configuration requirements. Work on Siebel Tools, Siebel Workflows, Siebel Reports, OATS, Selenium, EAI, Siebel Integration and Siebel upgrade. Required a Master's Degree in Computer Science, Engineering, Math, CIS or MIS and 1 year of work experience. A Bachelor's degree in a related field and 5 years of work experience would be acceptable in lieu of Master's degree in related field. Any suitable combination of education, training or experience would be acceptable. Resume to Dipali Trivedi, Director, CLOUDFOUNTAIN, INC.,125 Cambridge Park Drive, Suite 333, Cambridge, MA 02140.

PROGRAMMER ANALYST: Analyze, develop, implement, migrate and test software applications utilizing knowledge of C,C++, PL/SQL, Java, .NET, Oracle Applications(11i/R12) , reports & form builder (6i/10g), SDLC, OA Framework, SQL * Loader,Unix. AIM. Knowledge in Microsoft Windows, Visual Basic .NET,Linux, SOA, System administration and Web Methods reqd. Must be willing to travel & reloc. Reqs MS in comp sci, eng, bus or rel. Mail resumes to Nitya Software Solutions Inc. 9690 South 300, Ste 319, Salt Lake City, Utah 84070.

BUCKNELL UNIVERSITY, GIS/WEB APPLICATION SPECIALIST: DUTIES: LEWISBURG, PA. Develop, integrate, implement GIS/digital scholarship applications that will advance the University's academic mission. Research & recommend infrastructure solutions

INFOSYS LIMITED is in need of individuals to work full-time in Plano,Texas and various unanticipated locations throughout the U.S. Must be willing to work anywhere in the U.S.as all job opportunities may involve relocation to various and unanticipated client site locations; any relocation to be paid by employer pursuant to internal policy. We have multiple openings for each job opportunity, and are an Equal Opportunity Employer M/F/D/V. Please apply online at: http://www.infosys.com/careers/apply-now/apply.asp. Select 'Americas' under 'Job Opportunities' and follow the link for 'Experienced Professionals.' Once a user account has been created, please follow the link for 'Search Openings' and enter reference ID(s) for the position(s) of interest in the 'Auto Req ID' box.

ASSOCIATE ENGAGEMENT MANAGER(S) needed to contribute to competitor analysis and prospect identification; provide ground intelligence to pursuit teams, as well as account context and client introductions required for opening diverse service offerings in account(s). Travel required. **(REF ID 8880BR).**

CONSULTANT(S) (DOMAIN) - US needed to help conduct IT requirements gathering, define problems, provide solution alternatives, create detailed computer system design documentation, implement deployment plan and help conduct knowledge transfer with the objective of providing high quality IT consulting solutions. **(REF ID 8878BR).**

CONSULTANT(S) (PRODUCTS AND PACKAGES) - US needed to help conduct IT requirements gathering, define problems, provide solution alternatives, create detailed computer system design documentation, implement deployment plan, and help conduct knowledge transfer with the objective of providing high-quality IT consulting solutions. **(REF ID 8879BR).**

LEAD CONSULTANT(S) (DOMAIN) – US needed to anchor different phases of IT engagement including business process consulting, problem definition, discovery, solution generation, design, development, deployment and validation. **(REF ID 8851BR).**

LEAD CONSULTANT(S) (PRODUCTS AND PACKAGES) - US needed to anchor different phases of the IT engagement including business process consulting, problem definition, discovery, solution generation, design, development, deployment and validation. **(REF ID 8852BR).**

PRINCIPAL(S) - Business Consulting needed to lead small proposals and multiple streams on complex proposals. Develop best in class proposals that present Infosys Point of View, approach and IT solution. Help identify clients and opportunities for the practice, present preliminary ideas and proposals to clients, lead engagements from launch to closure. Travel Required. **(REF ID 8854BR).**

PRINCIPAL CONSULTANT(S) (DOMAIN) US needed to lead the engagement effort for IT assignments, from business process consulting and problem definition to solution design, development and deployment. Lead proposal development. Travel required. **(REF ID 8855BR).**

SENIOR TECHNOLOGY ARCHITECT(S) – US needed to provide IT architectural solutions for one or more projects. Provide input to create technology and architectural frameworks. Understand and analyze client business & IT problems, technology landscape, IT standards, and enterprise roadmaps. **(REF ID 8876BR).**

TECHNOLOGY ARCHITECT(S) - US needed to provide inputs on IT solution architecture based on evaluation/understanding of solution alternatives, frameworks and products. Will interact with clients to elicit architectural and non-functional requirements like performance, scalability, reliability, availability, maintainability. **(REF ID 8877BR).**

for GIS use on campus and upgrade of GIS software. Provide technical support for GIS use across campus. REQUIREMENTS, Must have Bachelor's Degree or foreign equivalent in Geomatics Engineering or related field and 6 months post-baccalaureate experience in GIS System development & support in an academic or professional setting. Experience must include 6 months experience using ArcGIS desktop and Image processing software to conduct applied spatial analysis in academic research projects; building GIS systems or modules using Python, C#, SQL, and C/C++/ Delphi; using ArcGIS Server, Google Maps API, Java Script, HTML, and CSS to build GIS web applications; and knowledge of software algorithms on both vector and raster data. Please send cover letter and resume including job history to Alison Epting Razet, Bucknell University, One Dent Drive, Lewisburg, PA 17837; ae015@bucknell.edu. Equal Opportunity Employer.

MPHASIS CORP. has multi openings at various levels for the follow'g positions at its office in NY, NY & unanticipated client sites thr/o the US 1. Info. Sys. Anyst* - Ana. & provide sys req & spec. 2. SW Dvlper* - Design, dvlp & modify SW sys. 3. Sys. Architect Dvlper* - Dvlp IT architecture 4. Graphic UI Desgr* - Design UI & perform UAT 5. N/W Infra Eng* - Maintain & TRBL n/w, design, dvlp, install n/w infra appl. 6. Business Operation Anyst* - Ana bus process thru app of s/w sol. 7. IT Mgr* - Plan & manage the delivery of IT proj. 8. Enterprise Svc Engagem't Mgr* - E2E sale of IT svc/prod. 9. Eng Engagem't Mgr* - Manage & direct business integration of proj activities.10. Mkt Dvlpt Mgr* - Promote IT svc/prod. & impl bus plans.Must have a Bachelor/equiv and prior rel. exp, Master/equiv, or Master/equiv and prior rel. exp. Edu/exp req vary depending on position level/type. *Lead positions in this occupation must have Master/ equiv+2yr or Bach/equiv+5yr progressive exp. Travel/relo req. Send resume & applied position to: recruitmentus@mphasis.com or 460 Park Ave. S., Ste# 1101, New York, NY 10016 Attn: Recruit.

SAMSUNG SEMICONDUCTOR INC. has a Director, Software Engineering (job code: 5CJ2512) job opportunity in Menlo Park, CA: Define, consult and contribute to the software architecture of SSIC mHealth Simband platform. Mail resume to 2440 Sand Hill Rd., Ste. 302, Menlo Park, CA 94025, Attn: S. Tan. Must reference job code to be considered. EOE.

Microsoft Corporation currently has the following openings (job opportunities available at all levels, including Principal, Senior and Lead levels):

REDMOND, WA

Software Engineers and Software Development Engineers in Test (all levels, including Leads and Managers): Responsible for developing or testing computer software applications, systems or services. (http://bit.ly/MSJobs_SDE) (http://bit.ly/MSJobs_IT_SDE)

Software Engineers and Software Development Engineers in Test (all levels, including Leads and Managers): Responsible for developing or testing computer software applications, systems or services. Requires domestic and international travel up to 25%. (http://bit.ly/MSJobs_SDE) (http://bit.ly/MSJobs_IT_SDE)

Research Software Development Engineers (all levels): Responsible for conducting applied research into new products and services through software engineering techniques. http://bit.ly/MSJobs_Research_Software_Engineer

MOUNTAIN VIEW, CA

Software Engineers and Software Development Engineers in Test (all levels, including Leads and Managers): Responsible for developing or testing computer software applications, systems or services. (http://bit.ly/MSJobs_SDE) (http://bit.ly/MSJobs_IT_SDE)

PALO ALTO, CA

Software Engineers and Software Development Engineers in Test (all levels, including Leads and Managers): Responsible for developing or testing computer software applications, systems or services. (http://bit.ly/MSJobs_SDE) (http://bit.ly/MSJobs_IT_SDE)

SAN FRANCISCO, CA

Software Engineers and Software Development Engineers in Test (all levels, including Leads and Managers): Responsible for developing or testing computer software applications, systems or services. (http://bit.ly/MSJobs_SDE) (http://bit.ly/MSJobs_IT_SDE)

SUNNYVALE, CA

Software Engineers and Software Development Engineers in Test (all levels, including Leads and Managers): Responsible for developing or testing computer software applications, systems or services. (http://bit.ly/MSJobs_SDE) (http://bit.ly/MSJobs_IT_SDE)

SAN DIEGO, CA

Software Engineers and Software Development Engineers in Test (all levels, including Leads and Managers): Responsible for developing or testing computer software applications, systems or services. (http://bit.ly/MSJobs_SDE) (http://bit.ly/MSJobs_IT_SDE)

CHICAGO & DOWNERS GROVE, IL

Software Engineers and Software Development Engineers in Test (all levels, including Leads and Managers): Responsible for developing or testing computer software applications, systems or services. (http://bit.ly/MSJobs_SDE) (http://bit.ly/MSJobs_IT_SDE)

CAMBRIDGE, MA

Software Engineers and Software Development Engineers in Test (all levels, including Leads and Managers): Responsible for developing or testing computer software applications, systems or services. (http://bit.ly/MSJobs_SDE) (http://bit.ly/MSJobs_IT_SDE)

DURHAM, NC

Software Engineers and Software Development Engineers in Test (all levels, including Leads and Managers): Responsible for developing or testing computer software applications, systems or services. (http://bit.ly/MSJobs_SDE) (http://bit.ly/MSJobs_IT_SDE)

ALISO VIEJO, CA

Software Engineers and Software Development Engineers in Test (all levels, including Leads and Managers): Responsible for developing or testing computer software applications, systems or services. (http://bit.ly/MSJobs_SDE) (http://bit.ly/MSJobs_IT_SDE)

RESTON, VA

Software Engineers and Software Development Engineers in Test (all levels, including Leads and Managers): Responsible for developing or testing computer software applications, systems or services. (http://bit.ly/MSJobs_SDE) (http://bit.ly/MSJobs_IT_SDE)

NEW YORK, NY

Software Engineers and Software Development Engineers in Test (all levels, including Leads and Managers): Responsible for developing or testing computer software applications, systems or services. (http://bit.ly/MSJobs_SDE) (http://bit.ly/MSJobs_IT_SDE)

HUMACAO, PUERTO RICO

Software Engineer, Principal IT-Ops Services or Other: Responsible for developing or testing computer software applications, systems or services. http://bit.ly/MSJobs_1930

FARGO, ND

Software Engineers and Software Development Engineers in Test (all levels, including Leads and Managers): Responsible for developing or testing computer software applications, systems or services. (http://bit.ly/MSJobs_SDE) (http://bit.ly/MSJobs_IT_SDE)

Multiple job openings are available for each of these categories. To view detailed job descriptions and minimum requirements, and to apply, visit the website address listed. EOE.

CAREER OPPORTUNITIES

SAMSUNG SEMICONDUCTOR INC. has a Security Architect job opportunity in Milpitas, CA: Security Architect (job code: 5MD1320) Work closely with business partners and other stakeholders such as legal, HR, IT, facilities etc. to provide information security oversight where required and assist in building corporate policies and procedures. Mail resume to Samsung Semiconductor, c/o Staffing – PTCL, 601 McCarthy Blvd., Milpitas, CA 95035. Must reference job code to be considered.

PROGRAMMER ANALYST - dsgn, dvlp, maintain, test & implmt applic s/w databases ,ETL Mappings, BI reports utilizing knowledge of Sql, Plsql, Unix Shell Scripting, Oracle 8i/9i/10g/11g, SQL Server, Informatica 8.6/9, ERWIN data modeler, TOAD, SQL Plus, Oracle OBIEE 10g/11g; Must be willing to travel & reloc. Reqs MS Comp Sci, Engg or rel. Mail resumes to Strategic Resources International, Inc, 777 Washington Rd, Ste 2, Parlin , NJ 08859.

SR. IP NETWORK ENGINEER. (Mult. Openings) Nexius Insight, Inc. Allen, TX. W/BS in Comp Sci/Electronics Engg or rltd + 5 yrs exp. Network Engr or rltd. Undrstnd existing data center Switching & routing architect. at high & low level. Prov. architect. recommend. on different dimensions incl reliability, security, & scalability. Dvlp detailed plans for configuring, testing, & upgrading ntwrk elements. Use dynamic routing protocols (BGP & OSPF) & policy-based routing. Implmt. Pv6 networks in data center environ. incl assoc. routing architect. (BGPv6). Utilize expertise w/Cisco Nexus family of switches (2000, 5000 & 7000 series). Use exp. w/layer 2 protocols such as VTP, link aggreg. (LACP & Cisco VPC) & UDLD. Utilize exp w/first hop redund. protocols (e.g. HSRP, VRRP). Use exp. w/ MPLS at telco scale. Utilize exp w/Nexus modules & assoc. 10, 40 & 100 gbps interfaces. Utilize exp w/data center firewalls, load balancers & security apps., part. Juniper &/or Cisco. Use 3GPP mobility stds & interfaces & high level EPC architect. Mail resumes to: Nexius Insight, Inc., 1301 Central Expressway S, Ste 200, Allen, TX 75013. Attn: HR

COMPUTER PROFESSIONALS. Central NJ IT Consulting company requires candidates for following positions at their primary New Brunswick, NJ location 1) Biztalk Developer: design & implement EAI & SOA business solutions using Biztalk App Integration Bus, Web Services, Real Time Messaging, Data Virtualization Layer, etc. Candidates must have min 1yr exp in SOA, Biztalk ESB, SAP/PeopleSoft Adapters, HIPPA/ EDI, .NET,BAM, BRE, etc. 2) Java Developer: To Design, develop & implement business software apps using Java & J2EE, Candidates must have 1 yr. of mandatory exp in Java Flex, Blaze Server, Web Services, UI Development, integration w/legacy systems, Hadoop, etc. All positions require: MS in CS/Engineering/Info systems/Business or related. BS degree + 5yrs exp can be substituted for the MS reqirmt 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 reqirmts & State & Local reqirmts. Mail your resume to: eVantage Solutions Inc., Inc, Attn: HR, 317 George St., Ste # 205, New Brunswick, NJ 08901.ttn: HR, 317 George St., Ste # 205, New Brunswick, NJ 08901.

EVENTBRITE, INC. is looking Sr. Software Engineers – Front End in San Francisco, CA to build software solutions & application features. Resume to HR, Job #EB05, Eventbrite, Inc., 155 5th St. Fl 7, San Francisco, CA 94103.

Microsoft®

Microsoft Corporation currently has the following openings (job opportunities available at all levels, including Principal, Senior and Lead levels):

REDMOND, WA

Program Managers: Coordinate program development of computer software applications, systems or services. (http://bit.ly/MSJobs_ProgMgr) (http://bit.ly/MSJobs_HW_ProgMgr) (http://bit.ly/MSJobs_ProdQlty_Supp) (http://bit.ly/MSJobs_IT_ProgMgr)

Program Managers: Coordinate program development of computer software applications, systems or services. Requires domestic and international travel up to 25%. (http://bit.ly/MSJobs_ProgMgr) (http://bit.ly/MSJobs_HW_ProgMgr) (http://bit.ly/MSJobs_ProdQlty_Supp) (http://bit.ly/MSJobs_IT_ProgMgr)

Business/Operations Program Managers: Responsible for the design, implementation, and release of programs or projects. (http://bit.ly/MSJobs-Buss_Oper_Prog_Mgmt) (http://bit.ly/MSJobs_Ops_PM)

Program Manager II - Operating Systems Engr Grp or Other: Define and execute the security plan for a Microsoft consumer product combining hardware and software technologies. http://bit.ly/MSJobs_1411

Program Manager: Coordinate program development of computer software applications, systems or services. This position requires international and domestic travel up to 50%. (http://bit.ly/MSJobs_1476) (http://bit.ly/MSJobs_1364)

NPI/SCC Engineering Program Manager - Devices Group or Other: Coordinate program development of hardware products / systems. This position requires domestic, regional and international travel up to 25%. http://bit.ly/MSJobs_1485

Director, Program Manager – Product Marketing or Other: Manages program development of computer software applications. http://bit.ly/MSJobs_1880

Business Program Manager (Learning & Development Specialist MPN) - SMSP Partner Coverage or Other: Responsible for the design, implementation, and release of programs or projects. Requires Domestic and International travel up to 25%. http://bit.ly/MSJobs_1489

Senior Hardware Program Manager - Large Screen Devices (LSD) or Other: Responsible for leading a cross-functional team from initial conception to high volume production. Requires Domestic and International travel up to 25%. http://bit.ly/MSJobs_1922

Senior Director Business Operations & Program Management- Mobile Devices Sales or Other: Manage the operations and business-critical function of the sales of smartphones and mobile phones. Requires Domestic and International travel up to 25%. http://bit.ly/MSJobs_1967

Business Program Manager: Responsible for the design, implementation, and release of programs or projects. Telecommuting Permitted. Position allows employee to reside anywhere in the U.S. and telecommute to perform work exclusively from home. (http://bit.ly/MSJobs_1526) (http://bit.ly/MSJobs_1622)

Principal Program Manager: Coordinate program development of computer software applications, systems or services. Requires international travel up to 50%. Telecommuting permitted when not required at factory sites. http://bit.ly/MSJobs_1869

MOUNTAIN VIEW, CA

Program Managers: Coordinate program development of computer software applications, systems or services. (http://bit.ly/MSJobs_ProgMgr) (http://bit.ly/MSJobs_HW_ProgMgr) (http://bit.ly/MSJobs_ProdQlty_Supp) (http://bit.ly/MSJobs_IT_ProgMgr)

PALO ALTO, CA

Senior Program Manager - Skype Engineering or Other: Coordinate program development of computer software applications, systems or services.Requires domestic and international travel up to 25%. http://bit.ly/MSJobs_2264

SAN FRANCISCO, CA

Program Managers: Coordinate program development of computer software applications, systems or services. (http://bit.ly/MSJobs_ProgMgr) (http://bit.ly/MSJobs_HW_ProgMgr) (http://bit.ly/MSJobs_ProdQlty_Supp) (http://bit.ly/MSJobs_IT_ProgMgr)

SUNNYVALE, CA

Program Managers: Coordinate program development of computer software applications, systems or services. (http://bit.ly/MSJobs_ProgMgr)(http://bit.ly/MSJobs_HW_ProgMgr) (http://bit.ly/MSJobs_ProdQlty_Supp) (http://bit.ly/MSJobs_IT_ProgMgr)

SAN DIEGO, CA

Program Managers: Coordinate program development of computer software applications, systems or services. (http://bit.ly/MSJobs_ProgMgr) (http://bit.ly/MSJobs_HW_ProgMgr) (http://bit.ly/MSJobs_ProdQlty_Supp) (http://bit.ly/MSJobs_IT_ProgMgr)

Business Program Manager- Mobile Device Sales or Other: Responsible for the design, implementation, and release of programs or projects. Requires domestic and international travel up to 25%. http://bit.ly/MSJobs_1587

Operations Program Manager (Tech Solutions Professional, MMDS)- Mobile Device Sales or Other: Responsible for the design, implementation, and release of programs or projects. Requires domestic and international travel up to 25%. http://bit.ly/MSJobs_1589

CAMBRIDGE, MA

Program Managers: Coordinate program development of computer software applications, systems or services. (http://bit.ly/MSJobs_ProgMgr) (http://bit.ly/MSJobs_HW_ProgMgr) (http://bit.ly/MSJobs_ProdQlty_Supp) (http://bit.ly/MSJobs_IT_ProgMgr)

DURHAM, NC

Program Managers: Coordinate program development of computer software applications, systems or services. (http://bit.ly/MSJobs_ProgMgr)(http://bit.ly/MSJobs_HW_ProgMgr) (http://bit.ly/MSJobs_ProdQlty_Supp) (http://bit.ly/MSJobs_IT_ProgMgr)

RESTON, VA

Program Managers: Coordinate program development of computer software applications, systems or services. (http://bit.ly/MSJobs_ProgMgr)(http://bit.ly/MSJobs_HW_ProgMgr) (http://bit.ly/MSJobs_ProdQlty_Supp) (http://bit.ly/MSJobs_IT_ProgMgr)

IRVING, TX

Business/Operations Program Managers: Responsible for the design, implementation, and release of programs or projects. (http://bit.ly/MSJobs-Buss_Oper_Prog_Mgmt) (http://bit.ly/MSJobs_Ops_PM)

RENO, NV

Business/Operations Program Managers: Responsible for the design, implementation, and release of programs or projects. (http://bit.ly/MSJobs-Buss_Oper_Prog_Mgmt) (http://bit.ly/MSJobs_Ops_PM)

Multiple job openings are available for each of these categories. To view detailed job descriptions and minimum requirements, and to apply, visit the website address listed. EOE.

CAREER OPPORTUNITIES

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

Systems Design Engineer (Req# 9JQNT9) Eval latest iPad, iPhone, &iPod HW systems in field. Travel Req 30%

ASIC Design Engineer (Req# 9D82KR) Work w/ SoC design, system group & OSATs to define & execute the state-of-art assembly technology roadmap including PoP, SiP, CSP, wirebond & flip chip BGA.

Engineering Project/Program Manager (Req#9ML2Y6) Plan & execute Apple Power Eng Programs through building strategic vision, implementing processes, & providing precise guidance to team while leveraging eng & operation partners to meet & exceed programs' critical milestones. Travel Req 25%.

Senior Software Engineer (Req# 9QZTQV) Dsgn & dvlp software for web apps.

Software Engineer Applications (Req#9L5TEK) Des & imple user friendly, highly function & secure PTS sys that meet global req of reg, zero down, reliable & scal.

Software Engineer Systems (Req# 9HDNBF). Respon for writing sw to develop & maintain mission critical portal app including writing code for new feats & sup existing srvcs.

Software Development Engineer (Req#9FTNYV). Respon for design & develop of algorithms & sw for comp vision sys.

Software Engineer Applications (Req#9KXSB3). Respon for design & develop of web/mobile based sols for Apple Retail Bus.

Software Engineer, Applications (REQ#9M2RXL). Des & dev SW sys based on machine learning & algorithms.

IST Technical Project Lead (REQ# 9LZVCC). Anlyz bus users' reqs to provide conceptual & detailed des to meet bus users' needs by combining SAP stand. config & custom dev.

Information Systems Manager (Req# 9BGQ73) Admin, install, config, trbleshoot & write spprt doc for IBM/AIX & act as lead for team of AIX sys engs.

Hardware Development Engineer (Req# 9F4T6X) Dsgn, intgrt & validate analog & mixd signl hrdwre subsys for prtble dvcs.

Software Engineer Applications (Req#9NBSC4). Respon for develop, design & debug of sw for iOS retail sys.

Information Systems Engineer (Req#9H3UX6). Respon for design, implement & deploy for the Apple communities used to sup prod questions & issues.

Engineering Program Specialist (Req#9CYVHL). Assist hw eng team to coordinate key areas of hw eng progs w/ exposure to cross-func teams & overseas vendors.

Technical Content Analyst (REQ# 9HBVRY). Resp for diagnosing, troubleshooting, analyzing CMS/ data processing problems to implnt & imprv comp systems

User Interface Designer & Prototyper (REQ#9BVW2D). Des & prototyp user interf for new consmr prod

Software Development Engineer (REQ#9J2U2T). Create test plans for new & ex'ting 1st-party apps & OS functn'lity.

Software Development Engineer (REQ#9LDPT6). Des & impl large-scale, high vol, high avail, queue-based backend pipeline SW using Java.

Software Development Engineer (REQ#9HKTB3). Res for validat. of WiFi SW stack on iOS & OS X platforms.

ASIC Design Engineer (REQ# 9AAUUK). Dev tests & test environments for GPU designs.

Software Engineer Applications (REQ#9D2W42). Des & dev large scale machine learning platform w/ mission critical perform req.

Technical Project Coordinator (Req# 9Q3S2W) Coord & drive build deliver & cross function'l dependencies.

Software Engineer, Applications (REQ#9RXPXE). Des & dev web/ mobile based solt'ns for Apple Retl Busns.

Software Engineer Applications (Req#9UXPWL). Des & dev SW for Apple internet servers.

Software QA Test Engineer (Req# 9QEUU6) Test iOS device integ & compatib w/Apple CarPlay & non-CarPlay auto head units.

Hardware Development Engineer (Req#9FD26J). Support the des, constrctn, & validation of board/ sys level electronics for consumer electronic devices.

Software Development Engineer (Req#9NYVDS). Des & dev sys SW apps & frameworks for Apple Watch prod w/ well-designed solutions.

Software Engineer, Applications (Req#9H6THF). Spprt lrg-scale retail POS sys utlzng exp in comp apps sftwre dsgn, dvlpmnt, implmntion & maintnce.

Software Quality Assurance Engineer (Req#9LG3YS). Responsible for testing code for iTunes products and services.

ASIC Design Engineer (Req# 9FLQ3Y). Respon for design verification of complex SOCs. Create direct & random test for hw design.

Hardware Development Engineer (Req#9K4Q7D). Lead proj from concept phase prod in colab de-

veloping audio features. Travel Req'd 25%

Software QA Engineer (Req# 9U4VHM). Def & create test assurance, plans & automation.

Software Development Engineer (Req#9D2W4F). Test cellular telephony functionality of iOS devices. Travel req'd 30%.

Layout/Mask Designer (Req#9FC4KS). Work on the latest technology nodes to create world-class custom digital macros, libraries, etc.

Software Development Engineer (Req#9FC4R8). Des & dev SW to assist with the evaluation of Maps services.

Hardware Development Engineer (Req#9D9W3W). Respon for design & develop of hw for OLED displays. Travel req 20%.

Systems Design Engineer (Req# 9F52Y4). Respon for eval of iPad & iPhone hw/sw wireless systems.

Hardware Development Engineer (Req#9UD48M). Dsgn & intgrte solid-state fngrprint snsrs for mbl devices. Travel req'd 20%.

Software Engineer Systems (Req# 9NLSL8). Architect, dev & mntain full-stack SW solutions that include a front-end, serv. & storage.

Software Engineer Applications (Req#9JTTQC). Des & dev Datastore abstraction services, framew & web based app that can rapidly scale & serve variety of usage patrns & data form's @ PB scale .

Software Engineer Applications (Req#9K2NHY). Des & dev Datastore abstract. services, frmewk & web based app that can rapidly scale & serve variety of usage patt's & data form's @ PB scale.

Mechanical Quality Engineer (REQ#9E5UWC). Dev & implem't inspection tools in high vol manuf. enviromnt. Travel req: 35%

Software Engineer Applications (Req#9RDT85). Des, dev, test & deploy high volume, highly scalable, & highly fault tolerant email infrstrctre.

Software Engineer, Applications (REQ#9REUAV). Monitor, maintain, & automate all data ctr oper for Apple ISO for all hosts & dvcs by perf puppet config mgt & bldg custom modules

Software Development Engineer (Req#9CHME3) Dev & implmnt SW to improve modem performance in iOS devices.

Engineering Project Manager (Req# 9DPPBR) Mng new prdct ramp to launch Apl's prdcts on tme w superior prdct qual & max cust avail. Travel req'd 30%.

Software Quality Assurance Engineer (REQ#9LUS65). Create, dev & lead the exec of all tst strategies to ensure highst qual embed prdct SW solu's. Travel req'd 15%

Systems Design Engineer (Req# 9FZN77). Work w/ cross-funct teams designing, testing & debugging products for compliance w/ worldwide EMC regs.

Software Development Engineer (Req#9F4VHN) Identify areas of weakness in existing prod. Prop sol's & improv to existing prods.

Hardware Development Engineer (REQ#9LD3MH). Des, dev & validate baseband HW circuits & syst for wireless comm devices.

Application Engineer (REQ# 9E2PSD). Des & impl'nt Enterprise product data & release sys'm.

ASIC Design Engineer (REQ# 98LV5L). Des App. Specific Integrated Circuits (processor design) .

Software Engineer Applications (Req#9PRUG9). Des & build user interface for Apple's support and service bus.

Software Development Engineer (Req#9BV3H9). Des & implmnt audio drivers and audio SW tools for OS X and iOS.

Software Development Engineer (REQ#9FZ25R). Des, dev & debug drivers involved w/ radio comm.

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

Systems Programmer (Req# 9CZPUX) Dfne, crte & implmnt cstm dvlpm integrations, user intrfces & lrg scl intrnl tool sets.

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

ASIC Design Engineer (REQ# 9GX2FB). Imp complx, high perform & low power unts of a CPU involving logic des, HDL synth & place-and-route.

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

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

CAREER OPPORTUNITIES

The Department of Computer Science and Engineering at the University of Notre Dame

seeks candidates for a teaching faculty position to teach courses primarily in the CS&E undergraduate curricula. This is a full-time, continuing position in the Special Professional Faculty track. Competitive candidates will have the training and experience necessary to teach effectively in a range of courses in accredited degree programs in Computer Science and Computer Engineering. Candidates with backgrounds in all areas of Computer Science and Computer Engineering will be considered. Relevant industry experience is also valued.

The University of Notre Dame is a private, Catholic university with a doctoral research extensive Carnegie classification, and consistently ranks in USNWR as a top-twenty national university. The South Bend area has a vibrant and diverse economy with affordable housing and excellent school systems, and is within easy driving distance of Chicago and Lake Michigan.

Qualified candidates should have at least a Masters degree, and preferably a doctoral degree, in Computer Science, Computer Engineering, or a related area.

Applications should include a cover letter, curriculum vitae, statement of teaching experience and philosophy, and names of at least three professional references, at least two of whom must be able to comment on the applicant's teaching experience. Review of applications will begin on June 1 and continue until the position is filled.

Applications should be submitted at http://apply.interfolio.com/29569.

ADVERTISER INFORMATION • JUNE 2015

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

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

Advertising Sales Representatives (display)
Central, Northwest, Far East:
Eric Kincaid
Email: e.kincaid@computer.org
Phone: +1 214 673 3742
Fax: +1 888 886 8599

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

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

Southeast:
Heather Buonadies
Email: h.buonadies@computer.org
Phone: +1 973 304-4123
Fax: +1 973 585 7071

Advertising Sales Representative (Classified Line & Jobs Board)
Heather Buonadies
Email: h.buonadies@computer.org
Phone: +1 201 887 1703

CAREER OPPORTUNITIES

39th Annual International Computers, Software & Applications Conference

CALL FOR PAPERS

July 1-5, 2015
Tunghai University
Taichung, Taiwan

www.compsac.org

Mobile and Cloud Systems - Challenges and Applications

COMPSAC is the IEEE Signature Conference on Computers, Software, and Applications. It is one of the major international forums for academia, industry, and government to discuss research results, advancements and future trends in computer and software technologies and applications. The technical program includes keynote addresses, research papers, industrial case studies, panel discussions, fast abstracts, doctoral symposium, poster sessions, and a number of workshops on emerging important topics. With the rapidly growing trend in making computations and data both mobile and cloud-based, such systems are being designed and deployed worldwide. However, there still exists several challenges when they are applied to different domains or across domains. COMPSAC 2015 will provide a platform for in-depth discussion of such challenges in emerging application domains such as smart and connected health, wearable computing, internet-of-things, cyber-physical systems, and smart planet.

Technical Symposia | **Workshops Program** | **Special Sessions**

COMPSAC 2015 will be organized as a tightly integrated union of several symposia, each of which will be focusing on a particular technical segment. **Please visit www.compsac.org for full information on symposia organization.**

* **Symposium on Embedded & Cyber-Physical Environments**
* **Symposium on Software Engineering Technologies & Applications**
* **Symposium on Technologies and Applications of the Internet**
* **Symposium on Security, Privacy and Trust Computing**
* **Symposium on Mobile, Wearable and Ubiquitous Computing**
* **Symposium on Web Technologies & Data Analytics**
* **Symposium on Human-Machine and Aware Computing**
* **Symposium on Novel Applications and Technology Advances in Computing**
* **Symposium on Computer Education and Learning Technologies**
* **Symposium on IT in Practice**

Authors are invited to submit original, unpublished research work and novel computer applications in full-paper format. Simultaneous submission to other publication venues is not permitted. The review and selection process for submissions is designed to identify papers that break new ground and provide substantial support for their results and conclusions as significant contributions to the field. Submissions will be selected that represent a major advancement in the subject of the symposia to which they are submitted. Authors of submissions with a limited contribution or scope may be asked to revise their submissions into a more succinct camera-ready format; e.g., a short paper, workshop paper, fast abstract, or poster.

COMPSAC 2015 will also feature a workshops program for topics closely related to the conference theme, *Mobile and Cloud Systems - Challenges and Applications.* Special sessions such as Fast Abstract and Industry Papers will be applicable especially for researchers and engineers who would like to present a new, early and work-in-progress ideas, method, and analysis. The Doctoral Symposium will provide a forum for doctoral students to interact with other students, faculty mentors, industry and government. Students will have the opportunity to present and discuss their research goals, methodology, and preliminary results within a constructive and international atmosphere.

Important Dates for Authors:

January 17, 2015: Paper submissions due
March 15, 2015: Paper notifications
April 28, 2015: Camera ready and registration due

Contact COMPSAC

For full information and CFP please visit **www.compsac.org**
Contact COMPSAC organizers at **cs_compsac@iastate.edu**

COMPSAC Sponsors:

IEEE computer society | IEEE | IEEE CLOUD COMPUTING

COMPSAC Technical Co-Sponsors:

Information Processing Society of Japan | CCF | Korean Institute of Information Scientists and Engineers

COMPSAC 2015 Local Host:

东海大學 TUNGHAI UNIVERSITY

EDITOR DAVID ALAN GRIER
George Washington University; grier@gwu.edu

VIDEO

Disruptive Computing

David Alan Grier, George Washington University

To expand e-commerce, we shouldn't try to fit computing technology into markets, but instead fit markets into technology.

I wondered if I had surrendered needlessly. I had changed tables in a coffee shop to avoid overhearing a group from an e-commerce firm celebrate their triumphs in some small consumer marketplace. I was waiting for an appointment with that firm, and felt I could be more productive if I didn't have to listen to the group's bellowing. I picked up my computer as the group's leader bragged, "We were hired to disrupt the market with technology, and disrupt we did."

Claiming that you're skilled in business or technology because you're disruptive is akin to claiming that you're a good athlete because you're ruthless. A certain amount of aggression might be valuable in sports, but other skills are needed to put points on the board. I was left wondering how and when computing had aligned itself with global capitalism and market disruption.

I quickly searched the IEEE CS Digital Library for early connections between computing technology and markets. The initial results were promising, including a 1953 report on a meeting about computers and accountants, articles from the early 1960s on the SABRE airline reservation system, discussions about market simulations from the 1970s, and an intriguing 1987 piece by MIT researchers who argued that computers would expand the use of markets in societies and companies.

After my appointment, I found that my initial impressions were wrong. The literature actually treated markets and economics as another application for the traditional aspects of computing: data structure, algorithms, and digital communication. The articles kept a clear distinction between the technology and the application, merely suggesting that computer scientists could collaborate more closely with accountants, bankers, and business leaders.

Only the 1987 article suggested that computing might be intimately connected to markets. We "should not expect the electronically interconnected world of tomorrow to be simply a faster and more efficient version of the world we know today," it read. "Instead, we should expect fundamental changes in how firms and markets organize the flow of goods and services in our economy."

If anything, the literature suggested that computing shifted in 1997, when the world began thinking seriously about e-commerce. After that point, we started to see the rise of the technical literature that would support e-commerce, including the now-common topics of search, trust, privacy, recommender systems, auction strategies, and algorithmic game theory. All of these fields existed in some form before 1997, but none had the focus they acquired that year.

In its e-commerce expansion strategy, the US government described a new way of doing business using ideas it drew from computing technology. "We need to learn what can be done to expand the use of this technology," explained the principal author of a 1997 e-commerce report, former government official Ira Magaziner, "because the economic prosperity we are enjoying will increasingly depend upon it."

The government strategy concluded that the "genius and explosive success of the Internet can be attributed in part to its decentralized nature and to its tradition of bottom-up governance." The report argued that this technology would encourage entrepreneurship, expand markets and global trade, lower costs, support long-distance collaboration, and give greater choice to consumers. To expand e-commerce, we shouldn't try to fit computing technology into markets, but instead fit markets into technology. "Business models must evolve rapidly to keep pace with the breakneck speed of change in the technology," the report continued. "Unnecessary regulation of commercial activities will distort development of the electronic marketplace."

Even if we didn't choose to align ourselves with e-commerce, we can't stop the market from choosing us. Computing has proudly disrupted existing knowledge, transformed work, and restructured organizations. If business needed a partner, it couldn't have chosen better.

See **www.computer.org/computer-multimedia** for multimedia content related to this article.

DAVID ALAN GRIER is a professor of international science and technology policy at George Washington University. Contact him at grier@gwu.edu.

0018-9162/15/$31.00 © 2015 IEEE