Innovative Technology for Computer Professionals

Computer

02.15

COMPUTING IN HEALTHCARE

www.computer.org/computer

IEEE

IEEE computer society

39th Annual International Computers, Software & Applications Conference

www.compsac.org

CALL FOR PAPERS

July 1-5, 2015
Tunghai University
Taichung, Taiwan

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Authors are invited to submit original, unpublished research work and novel computer applications in full-paper format.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

MULTIMEDIA

Computer

FEBRUARY 2015
FEATURES

Qmags THE WORLD'S NEWSSTAND®

FEBRUARY 2015

CONTENTS

ABOUT THIS ISSUE
COMPUTING IN HEALTHCARE

Advances in computing technology are increasingly central to achieving truly personalized medicine.

FEATURES CONTINUED

COMPUTING PRACTICES

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

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 **MARK GALLAHER**
m.gallaher@computer.org

Computer Highlights Society Magazines

The IEEE Computer Society's lineup of peer-reviewed technical magazines cover cutting-edge topics in computing, including scientific applications, Internet computing, machine intelligence, pervasive computing, security and privacy, digital graphics, and computer history. Select articles from recent issues of other Computer Society magazines are highlighted below.

IEEE Software

Within the growing Internet of Things, sensing devices coupled with actuators through networked computer systems sense environmental cues to affect physical context in real time, with benefits ranging from greater physical well-being to a reduced energy footprint. Crucial to the IoT are **wireless sensor networks**, tiny battery-powered devices that interact by wireless communication. In "Debugging the Internet of Things: The Case of Wireless Sensor Networks," from *IEEE Software*'s January/February 2015 issue, a team from SensorHound and Purdue University propose software tools to enhance such networks' reliability.

IEEE Cloud Computing

Traditional security mechanisms are tailored for small-scale data, so they don't meet the needs of big data analytics and storage applications. The September 2014 special issue of *IEEE Cloud Computing* aims to stimulate discussion and research toward the innovation of security and privacy mechanisms for **big data applications in a cloud environment**.

IEEE Security & Privacy

Cryptography that keeps data secret, whether in transit or at rest, underlies secure communication, secure identities, and access control on the Internet; consequently, cryptographic algorithms and protocols are the subject of much ongoing research and refinement. *IEEE S&P*'s January/February 2015 issue focuses on current **trends in cryptography**, a field where increasing demand for secure communication is creating a vibrant, emerging landscape for privacy and authentication.

IEEE Internet Computing

Internet banking websites use security images as part of the login process, with the intent to foil phishing attacks. But do users notice when a security image is missing? "The Effectiveness of Security Images in Internet Banking," one of the "Best Conference Papers" featured in *IEEE Internet Computing*'s January/February 2015 issue, presents results from an extensive study on whether users notice and react to the absence of **Internet login security images**, finding that most participants enter their password even when a security image and caption aren't present.

IEEE Computer Graphics and Applications

Emotional regulation strategies determine how people feel, express, and regulate their emotions. Learning to adopt such strategies is especially important during adolescence, when deficiencies that can result in psychosocial and behavioral problems become evident. *IEEE CG&A*'s January/February 2015 Applications department, "A VR-Based Serious Game for Studying Emotional Regulation in Adolescents," demonstrates an **interactive virtual-reality psychotherapy game** to aid in the early detection of dysfunctional emotional regulation.

Computing in Science & Engineering

"If simulation is the third tier of science," write the guest editors of *CiSE*'s January/February 2015 issue, "then the communities that build the simulation software are the engine of innovation." This special issue presents the challenges and collective efforts of these **scientific software communities**, which produce "polished, sharable extensible software" but are often underappreciated by the average scientist.

0018-9162/15/$31.00 © 2015 IEEE

IEEE Intelligent Systems

According to the authors of "WaaS: Wisdom as a Service," from *IEEE Intelligent Systems*' November/December 2014 issue, recent advances in cloud computing, the IOT, human–computer interaction, big data, and other fields are contributing to a fusion of our social, cyber, and physical worlds—a **hyper-world** that uses data as a common bridge. The challenge they see is how "to realize the organic amalgamation and harmonious symbiosis among humans, computers, and things" that such a world might allow.

IT Professional

Can malware be exterminated? Pessimists believe that **complete malware detection** is an unsolvable and nonboundable problem; optimists argue for eventual solvability. In *IT Pro*'s November/December 2014 issue, the authors of "Can Malware Be Exterminated by Better Understanding Its Roots?" reveal pitfalls in malware research that, if addressed, could help move us in the right direction.

IEEE micro

The **miniaturization in electronic devices** has raised considerable reliability issues: smaller feature–sized transistors with both smaller capacitance and a lower supply voltage are increasingly vulnerable to cosmic rays, noises, and wear-outs. In "A Flexible, Self-Tuning, Fault-Tolerant Functional Unit Array Processor," from *IEEE Micro*'s November/December 2014 issue, a team from the Nara Institute of Science and Technology proposes a low-cost, self-tuning scheme to quickly locate defective processing elements or network connections.

IEEE MultiMedia

A widely studied topic in computer vision, **visual tracking** has many practical applications—automated surveillance, robot navigation, medical imaging, and traffic monitoring, to name a few—but remains challenging. In "Online Learning a High-Quality Dictionary and Classifier Jointly for Multitask Object Tracking," from *IEEE MultiMedia*'s October–December 2014 issue, the authors present a structured representation for visual object tracking that exploits label information strength and encourages images from the same class to have similar representations.

IEEE Pervasive Computing

Most wearable devices are powered by batteries that need frequent recharging, which can be difficult or even impossible for people who're traveling or living in remote areas. Two South China University of Technology researchers write in "Human Motion: Sustainable Power for Wearable Electronics," from *IEEE Pervasive*'s October–December 2014 issue, that wearable harvesters can be used to harness **kinetic energy from human motion**, providing sustainable power levels for wearable computing systems.

IEEE Annals of the History of Computing

In his introduction to the October–December 2014 special issue of *IEEE Annals*, "Algol Culture and Programming Styles," guest editor Gerard Alberts from the University of Amsterdam writes that for computer scientists and historians alike, the language born in 1958 as IAL (International Algebraic Language) and renamed **ALGOL (ALGOrithmic Language)** in 1959 is considered "a turning point in the development of programming languages and of software in general."

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

Denial-of-Service Attacks to UMTS

Elisa Bertino, Purdue University

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

Cellular communication networks are among today's most critical infrastructures, making possible important applications including location-based services, emergency management, and continuous healthcare monitoring. Consequently, cellular communication networks have been extensively analyzed to identify security threats and devise corresponding mitigation techniques. However, because achieving 100 percent security is impossible and new attacks are continuously being reported, it's crucial to research and identify new threats and vulnerabilities to improve network defenses.

In "A Denial of Service Attack to UMTS Networks Using SIM-Less Devices" (*IEEE Transactions on Dependable and Secure Computing*, vol. 11, no. 3, 2014, pp. 280–291), Alessio Merlo and his colleagues make an important breakthrough in cellular network security as they identify a novel *denial of service* (DoS) attack against universal mobile telecommunication system (UMTS) infrastructures. A DoS attack is disruptive and typically prevents legitimate users and applications from accessing networks.

The new attack operates at the user level and thus doesn't require hacking a network's intra-operator facilities. It also doesn't require that mobile devices be equipped with a valid subscriber identity module (SIM), making it fairly easy to carry out. This attack specifically targets the functionality of the home location register (HLR) database, which stores information about mobile subscribers and rules for call blocking and forwarding. Because the HLR is a central component of a cellular network, its inability to respond to legitimate users' requests makes the network's communication services unavailable to these users, disrupting network coverage. The article provides details about the attacking devices' design and an in-depth analysis showing that this new attack can reach cellular networks with an order of magnitude fewer resources than previous attacks.

> It's crucial to research and identify new threats and vulnerabilities to improve network defenses.

What distinguishes this new attack from previous attacks is that it doesn't require using a *botnet*. A botnet is a network of mobile devices owned by legitimate users that can be coordinated by a comman-and-control center to perform attacks unbeknownst to these users. The article shows that unlike botnet-based attacks, the new attack isn't impacted by user mobility, so the attack can be placed very precisely. The article doesn't propose a solution to protect against these attacks, so further research should look into identifying defenses, such as those based on anomaly detection.

This article is an important reference for researchers in academia and industry interested in securing cellular networks, as it demonstrates the importance of identifying all bottlenecks in a network infrastructure and making sure these bottlenecks can't be exploited by attackers. It's also a must-read for anyone interested in testing the robustness of HLR implementation solutions.

ELISA BERTINO is a professor in the Computer Science Department, Cyber Center, and the Center for Education and Research in Information Assurance and Security (CERIAS) at Purdue University. Contact her at bertino@cs.purdue.edu.

0018-9162/15/$31.00 © 2015 IEEE

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

Guido van Rossum: The Early Years of Python

Charles Severance, University of Michigan

Guido van Rossum discusses the initial development of Python, which has increasingly become the programming language of choice for many scientific fields due to its extensibility and ease of use.

Combining simplicity and ease of learning with powerful capabilities and strong performance, Python is one of the most popular programming languages in the world. Created by Guido van Rossum in 1989, Python is the leading application language in important emerging areas such as natural language processing and big data. It fills a very important gap between high-level applications like spreadsheets and statistical analysis packages and systems-oriented languages like C, C++, and Java. You can see my full interview with Guido at www.computer.org/computingconversations.

A NEW OPERATING SYSTEM

In the late 1980s, Guido was working at Centrum Wiskunde & Informatica (CWI), a mathematical and computer science research center in Amsterdam, developing a distributed operating system called Amoeba. Amoeba intended to make a network of computers appear as a single computer using a distributed kernel:

> *We wanted it to be self-hosting, and in order to do that, we realized we needed a large amount of user-level tools like an editor, a mail program, a login utility, or a backup tool. Because the file system model on Amoeba was very different from those on Unix systems, we couldn't use an existing suite of Unix utilities. We had a small team of people working on those tools, but it was very slow-going writing it all in C.*

Guido wondered if there might be a quicker way to get all of the OS utilities built:

> *I had this idea that given how much time we had available for Amoeba, I could actually build a whole new language, design and implement it from scratch, and then use it to implement our suite of tools and still be ahead of the game compared to a situation where we would have just clunked on writing the things we wanted to write in C.*

COMPUTING CONVERSATIONS

For an earlier project, Guido had worked on a language called ABC:

> *I wondered if the ABC language would be a much better language to write these utility tools for Amoeba. But ABC was very high level and abstract, and it wasn't well suited to talking to servers, file systems, and processes. In an alternate universe, ABC could have become the language of spreadsheets, as it was very good for talking about a user's data and doing all sorts of clever stuff using general-purpose data structures like lists and dictionaries. ABC also had nice code structuring devices, like a few simple statements that could be combined to create other constructs.*

Over a long holiday break in December 1989, Guido started developing an ABC-like language that could talk to the OS and would be suitable for quickly developing OS utilities for Amoeba. He named his nascent project Python, taking inspiration from the *Monty Python's Flying Circus* television program. After the holiday break, he continued to evolve the new language in his spare time:

> *For three months I did my day job, and at night and whenever I got a chance I kept working on Python. After three months I was to the point where I could tell people, "Look here, this is what I built." It had an interactive interpreter loop, so the first demos were all, "Let's assign an expression to a variable and print it back," or "Let's define a small function and call it," or "Let's put some things in an array and iterate over the array."*

Though Python wasn't yet ready to be used to develop Amoeba OS utilities, the language was very appealing for programmers who were tired of writing C programs or Unix shell programs for various tasks:

> *My two office mates were almost instantly taken with [Python] and started helping out. A few others within the institute were also excited about Python. We didn't use it on Amoeba right away because it wasn't mature enough to actually develop the system utilities that we wanted. But it worked well on our Unix system, and people outside my department at CWI started using it because it was fun and productive to use.*

Python is the leading application language in important emerging areas such as natural language processing and big data.

After a year of development and use within CWI, Guido and his colleagues decided that Python might have a broader use and so decided to release it as free software. This was before the term "open source" was coined, so they simply looked to the MIT X-11 license as an example of how to release free software. But they needed permission from CWI management to release the product. Guido's manager sent him to the legal affairs department:

> *She asked me questions like, "Did someone pay for this to be developed?," and my answer was "No, I started this all on my own time as part of some research projects and my manager's fine with it." So she said, "Sure, go ahead," and we did it. That was an incredibly lucky stroke.*

The first version of Python came out in February 1991 and was distributed using the alt.sources newsgroup on Usenet:

> *I immediately started getting useful, positive feedback from people who picked up Python from Usenet, and we quickly got into a routine of doing new Python releases. There were the obvious improvements to the language, and the library, and bug fixes. An important category of contributions were ports where people had different architectures and different compilers since the Unix world was much less homogenous at the time. There were a lot of small Unix vendors that had their own compilers or their own hardware, all sorts of incompatibilities.*

From the initial release of Python through the early 1990s, the size of the Python community grew and numerous organizations started depending on the language. With broadening adoption, there was a concern among users that "Guido might get run over by a bus." Some of the adopters were US government agencies that wanted to help grow and stabilize the Python community.

> *I got an invitation from NIST [National Institute of Standards and Technology] to come to the United States for a couple of months. We organized and hosted the first Python workshop at NIST in Gaithersburg, Maryland, in November 1994. Through the Python workshop, I met people*

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

from the Corporation for National Research Initiatives (CNRI), and they offered me a job working on Python. I went back to the Netherlands for a few months and then from 1995 to 2000 I worked in the US in northern Virginia at CNRI.

Now Guido could focus on building Python and evolving the user community with solid support from CNRI. During the late 1990s, Python moved through a series of 1.x releases:

When I started at CNRI, Python 1.3 was about to be released, and then while I was there we released several subsequent versions leading up to 1.5.2, which remained sort of the gold standard of Python for a long time afterward.

The growth of the Python community from its creation in December 1989 and the maturation of Python 1.x in the late 1990s laid the groundwork for the even broader expansion of Python 2.x and now 3.x. It's an excellent example of organizations like CWI, NIST, and CNRI making investments in an open source "commons," leading to significant positive value in computing.

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.

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

COMPUTING AND THE LAW

AUDIO

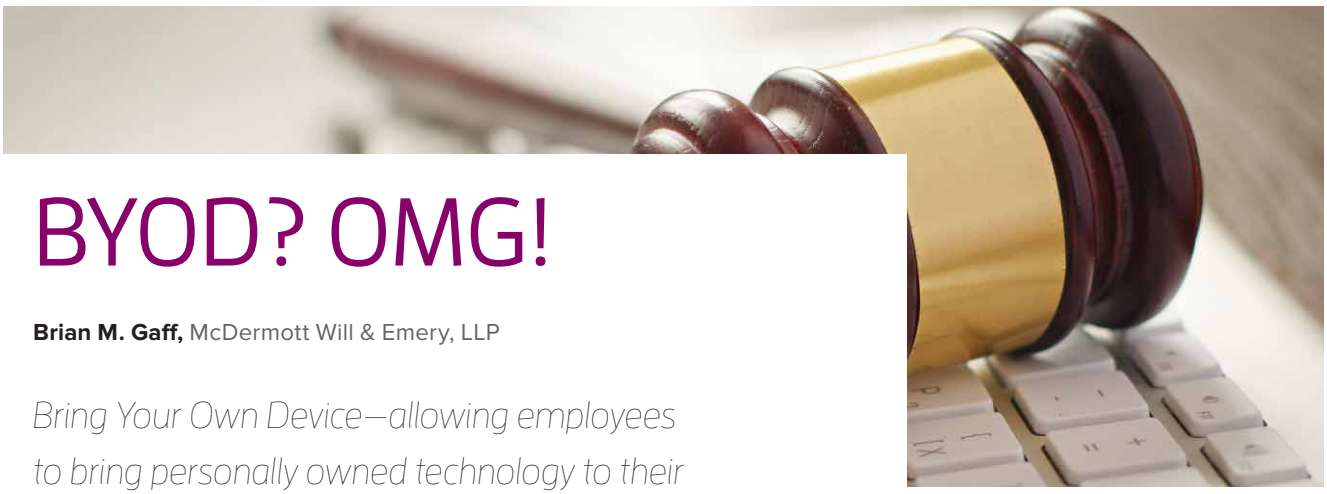

BYOD? OMG!

Brian M. Gaff, McDermott Will & Emery, LLP

Bring Your Own Device—allowing employees to bring personally owned technology to their workplaces might improve productivity, but it also creates risks for employers.

BYOD isn't just a new catchphrase, it's a growing trend in which companies permit—and even encourage—their employees to use their personally owned devices, such as smartphones and tablets, on the job. Although this might seem attractive, it can make a company's IT systems and data more vulnerable to malicious activity.

For an expanded discussion on this topic, listen to the podcast that accompanies this column at www.computer.org/computing-and-the-law.

WHAT IS BYOD?

It's common for employees to bring their own devices to their workplaces so they can maintain contact with family and friends. In those instances, employees use their devices exclusively for personal reasons. Employers usually don't take issue with that except, perhaps, if it becomes disruptive, hampers productivity, or creates a security risk (for example, if it's used where classified materials are present).

BYOD refers to the use of these personal devices for business purposes and reflects a blurring of the line between personal and business use on the same device. For example, an employee might use his smartphone to access the company's email system and read and respond to emails in connection with his job. He likewise might make telephone calls or send and receive text messages that are business related.

If a company allows the use of personally owned devices for business purposes, it needs to have a policy in place that, among other things, ensures the security of its systems, protects the confidentiality of its corporate materials, and respects the privacy of employees. The latter is particularly important, because employees undoubtedly have significant amounts of personal information on their devices.

REASONS FOR BYOD

Some argue that allowing employees to use their personally owned devices for work-related purposes will increase productivity. An underlying theory for this is that employees are more attuned their personal devices and will therefore use them more efficiently and even in some cases use them for work during their off-hours as well. Personally owned devices are typically more advanced compared to those that are employer issued, and most people usually prefer having newer-generation technology.

An employer that allows BYOD might create the perception that it's more accommodating to employees' needs and, therefore, a preferred place to work. This can help with recruitment and potentially improve employee morale. Also, there's the potential of reducing costs. Moving some or all of the device acquisition and usage costs from the employer to the employee could benefit the employer.

Another reason for allowing BYOD is that it's likely a fait accompli: many employers probably can't prevent employees from using their devices at work, so adapting to that reality might be the only reasonable option.

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

0018-9162/15/$31.00 © 2015 IEEE

EDITOR **BRIAN M. GAFF**
McDermott Will & Emery, LLP; bgaff@mwe.com

POLICIES TO SUPPORT BYOD

Employers need to have a BYOD policy in place, and employees need to understand it, to protect the interests of both.

At a minimum, a policy should define the devices that are and aren't acceptable for use in a BYOD context. Only certain brands, models, or configurations might be approved because of security or support requirements. If an employee isn't able to use one of the approved devices, then BYOD won't be an option.

The support issue is important. Unless the employer wants to make the investment to provide technical support for numerous and varied types of devices, setting limits on the types of devices is essential. The policy should set expectations as to the level of technical support available to employees using personally owned devices.

Even when an employee is using an approved device, the policy might need to restrict some of the device's capabilities when it's used for work-related purposes. For example, cameras could be prohibited in certain high-security environments. A policy might therefore require that the device's camera be automatically disabled when connected to the employer's network or when located in a certain area as determined by GPS.

To limit the likelihood of malware being introduced to the employer's systems, the policy might prohibit the installation of certain apps, or only allow the installation of apps from trusted sources. Further, the policy could limit how acceptable apps are used, which could include mandatory Web filtering.

The policy should define what employer systems and apps the employee is allowed to access on a personally owned device. This can range from limited access to just email to full access to R&D and production management systems. The employer needs to delineate the boundaries, which can be based on the employee's job responsibilities.

OWNERSHIP ISSUES

Having well-defined boundaries in place should help resolve the ownership issues. It's generally clear that the employee owns the device itself and the personal information stored on it. However, the policy should spell out that information stored on the device that's obtained from the employer's systems, as well as any of the apps installed on the device for work-related purposes, remain the property of the employer.

Ownership becomes important when an employee loses his device. It's important that the policy address this situation, as it can help limit the likelihood of a data breach. If the device can access the employer's systems, or if the device has the employer's information on it, the policy needs to allow the employer to remotely wipe the device to protect its systems and materials. As the term implies, "wipe" usually means deleting all information from the device, and that includes the employee's personal information, such as contacts, email, and photographs. That's potentially a significant loss for the employee, who will need to accept that consequence. Also, an employer might require a wipe in situations where the device isn't lost; for example, when the employee leaves the company and takes his personally owned device with him. The employee needs to accept that as well.

The employer should consider requiring passwords for and encryption of employees' devices. This doesn't eliminate the need to wipe a lost device. However, it provides an additional layer of security for a device that's lost but not yet wiped. Again, employees need to consent to this.

A BYOD policy might include terms that restrict how the device is used during personal time and off the employer's premises. For example, the policy might mandate that the device not be used such that it creates a dangerous distraction—for example, prohibiting its use while driving. The scope and enforceability of limitations like these should be discussed with a lawyer.

For companies that have been involved in litigation, the depth and complexity of the discovery phase of the litigation can be daunting. During this phase, documents, email, memoranda, and the like that are relevant to the dispute at issue are typically retrieved and closely examined. To the extent that employees have relevant information on their personally owned devices, it's possible that a litigation adversary will demand access to the devices. Whether the employee must consent to this should be part of the policy.

A BYOD policy is essentially a contract between an employer and employee. In return for granting access to the employer's systems and information, the employee agrees to abide by certain terms and conditions. For the policy to be enforceable, the terms and conditions should be reasonable and comply with the applicable laws and regulations, including those relating to employee privacy. That might include simultaneously complying with laws and regulations from multiple jurisdictions for employers that operate in different states and countries. Work closely with your lawyer to ensure that your BYOD policy is comprehensive and enforceable.

BRIAN M. GAFF is a senior member of IEEE and a partner at the McDermott Will & Emery, LLP law firm. Contact him at bgaff@mwe.com.

32 & 16 YEARS AGO

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

FEBRUARY 1983

www.computer.org/csdl/mags/co/1983/02/index.html

Introduction (pp. 8–9) "When designing a computer network, several sources of data insecurity need to be considered. Prominent among these are spurious message injection, message reception by unauthorized receivers, transmission disruption, and rerouting data to fake nodes. To maintain security against these hazards, a combination of encryption algorithms on the data and appropriate protocols for message exchanges is utilized. These techniques also facilitate the handling of other problems in computer communication networks, such as key distribution, authentication, privacy, digital signatures, network mail, and transaction verification."

Tutorial Survey (p. 15) "*Validation* and *authentication* refer to the methods of certifying the contents of a message and its originator, respectively. Both functions can usually be achieved through the use of a *digital signature*, which is appended to (or an integral part of) every message."

Key Protection (p. 27) "This article discusses the problem of protecting keys in a nationwide network using public-key cryptography for secrecy and digital signatures. Particular attention is given to detecting and recovering from key compromises, especially when a high level of security is required."

Data Security (p. 50) "Systematic research in secure protocols is still quite young; we expect, however, that the range of applications and methodologies will increase dramatically in the next few years. The separation of concerns (i.e., the strength of the cipher, the correctness of the protocol logic, and the adequacy of the implementation) evident in secure protocol design is, in fact, a microcosm of secure system design concerns."

Signing Mail (p. 55) "Computer-based message systems are likely to become the principal carriers of business correspondence. Unfortunately, with the efficiency of these systems come new possibilities for crime based on interference with digital messages. But the same technology that poses the threat can be used to resist and perhaps entirely frustrate potential crimes."

Voice Input (p. 91) "Mountain Computer has developed a plug-in speech digitizer for the IBM Personal Computer, enabling sound to be entered via a microphone, digitized, and stored on a diskette for later playback with a loudspeaker or the IBM computer speaker. According to the company, the Supertalker II reproduces the actual human voice, much like an audio tape recorder, providing better inflection than is ordinarily possible."

Burning ROM (p. 93) "A ROM simulator from P&E Microcomputer Systems is designed to allow faster development of microprocessor-based systems. New system development and debugging call for burning in a new PROM or EPROM every time a new code sequence, memory configuration, or timing profile is desired."

Selling Wool (p. 96) "Wool selling, a procedure ingrained with many time-honored practices in this South Pacific nation, took a step forward when the New Zealand Wool Board implemented two Sperry Univac System 80s at the start of wool buying season last year."

Standard LAN (p. 98) "Thirteen electronics companies have endorsed a single emerging standard for local-area networks that eventually will permit computers and office equipment—regardless of brand—to communicate with each other. The new IEEE P802.3 draft standard, CSMA/CD Carrier Sense Multiple Access with Collision Detection, represents convergence of the IEEE 802 working drafts, Ethernet specifications, and European Computer Manufacturers Association documents."

Workshop (p. 100) "Traditional views of microprogramming were contrasted with single-chip design and implementation in 'Psychology of Microprogramming,' chaired by [Will] Tracz. Commercial machine, bit-slice, microprogrammable machine, and single-chip 'cultures' were reviewed."

FEBRUARY 1999

www.computer.org/csdl/mags/co/1999/02/index.html

Chip Network (p. 9) "As the networking industry moves into the future, it is clear that the demand for more bandwidth

0018-9162/15/$31.00 © 2015 IEEE

with better performance and greater reliability is not going to abate. According to SwitchCore's [Kurt] Busch, internetworking chips represent a trend toward specialization in the networking industry that may be the only way vendors can develop products that will meet this ongoing demand."

OS Mobility (p. 13) "A number of operating systems with small footprints and reduced storage capacity have emerged to support the computing-related functions of handheld digital wireless devices, and OS developers are fighting for position in the marketplace."

Interlingual Internet (p. 18) "Computer and linguistics researchers throughout the world are hoping to develop a Universal Networking Language (UNL) that would let people who speak different languages communicate electronically."

International Y2K (p. 19) "The UN said it intends to meet again to discuss Y2K-related issues but has not specified a date. Regional working groups will be organized to discuss remediation of cross-border Y2K issues."

Frameworks (p. 24) "In this article, I describe our experience with developing an object-oriented framework for speech recognition applications that use IBM's ViaVoice speech recognition technology. I also describe the benefits of an object-oriented paradigm rich with design patterns that provide a natural way to model complex concepts and capture system relationships."

Innovation (p. 33) "Over the years, a particular blend of government, industry, and academia has been the foundation of computing innovation. If the US is to sustain its past growth in computing, researchers, business leaders, and policy makers need to understand the elements of this synergy."

Introduction (p. 45) "Public awareness of the Net as a critical infrastructure in the 1990s has spurred a new revolution in the technologies for information retrieval in digital libraries."

Federated Search (p. 51) "The Digital Libraries Initiative (DLI) project at the University of Illinois at Urbana–Champaign (UIUC) was one of six sponsored by the NSF, DARPA, and NASA from 1994 through 1998. The goal: develop widely usable Web technology to effectively search technical documents on the Internet. This article details their efforts."

The Informedia Project (p. 66) "Digital video presented a number of interesting challenges for library creation and deployment: the way it embeds information, its voluminous file size, and its temporal characteristics."

A Distributed Library (p. 74) "As the capabilities of distributed digital libraries increase, managing organizational and software complexity becomes a key issue. ... We have developed a software structure that successfully manages this complexity in our own digital library. ... Its novel, flexible macro language manages interfaces to many collections in many languages, and the structure allows development of experimental user interfaces."

The Stanford Infobus (p. 80) "Four years ago we set out to create a technical infrastructure to support the construction of digital libraries. In our view, a digital library comprises widely distributed resources that can be maintained autonomously by different organizations and will not require adherence to uniform interfaces. In defining an infrastructure, we wanted to consider all aspects, from user interface to low-level transport layers."

Online Learning (p. 115) "Our experiences have shown us that it is possible to design courses that effectively balance delivery of information with highly interactive, cost-effective learning in both corporate and college-level environments. This type of learning can apply to a range of other learning environments as well."

Embedded Systems (p. 116) "The low-cost, consumer-oriented, fast time-to-market mentality that dominates embedded-system design today forces design teams to use hardware-software codesign to cope with growing design complexities. New codesign methodologies and tools must support a key characteristic of next-generation embedded systems: the capability to communicate over networks and adapt to different operating environments."

Problem Solving (p. 122) "When we fixate on an elaborate requirements process, we try to persuade the powers that be to let us perform that process, and bemoan our lot when they don't support us. When we fixate on a quick process, we try to get it over with, doing the minimum needed to satisfy the process fanatics. There is a third way, however. We can reject process fixation and solve problems instead. We can reframe the requirements process as a goal-seeking dialogue whose purpose is to manage the risk of building the wrong product."

Open Source (p. 127) "Even if Linux proves the better product, Microsoft can always use tying—the creation of a mandatory dependency among applications—to maintain its monopoly. Isn't that really what COM+, Microsoft Transaction Server, Microsoft Message Queue Server, Active Directory, and Internet Explorer are designed to do?"

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

NEWS

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

Hackers Attack Internet Naming Authority

A spearfishing attack breached the internal systems of the Internet Corporation for Assigned Names and Numbers (ICANN), the organization announced recently.

The incident is important because the US-based international nonprofit group coordinates the Internet's critical domain-name and IP-address systems.

The attack began in late November 2014, when the unidentified hackers sent ICANN staff members email messages designed to look like they came from the organization's own domain.

The agency didn't specify how the incident unfolded, but the spearfishing messages could've linked to a fake ICANN page that asked the unsuspecting employees to log in with their email credentials.

Several employees provided their credentials, which the attackers then used to access internal emails and other ICANN systems.

For example, they gained administrative privileges to the centralized zone data system (CZDS), which helps users access websites by resolving letter-based domains to number-based IP addresses. The CZDS is also used by domain-name registries to manage the allocation of new generic top-level domains (TLDs).

Breaking into the system gave the hackers access to TLD zone files, which contain sensitive data including domain-name owners' names, usernames and encrypted passwords, street and email addresses, and phone and fax numbers.

ICANN says that it has deactivated all CZDS passwords and notified affected users, who can request new passwords at https://czds.icann.org.

The hackers also got into ICANN's members-only Government Advisory Committee wiki that contains various types of public information and user accounts on the ICANN blog and the ICANN WHOIS information portal.

The attackers might have accessed confidential information about the Internet's addressing system and plans for its future, depending on the jobs held by the ICANN employees whose accounts they compromised.

According to a statement posted on ICANN's website, "Earlier this year, ICANN began a program of security enhancements in order to strengthen information security for all ICANN systems. We believe these enhancements helped limit the unauthorized access obtained in the attack. Since discovering the attack, we have implemented additional security measures."

New Bluetooth Version Offers Online Connectivity, Could Be Used with the Internet of Things

A standards group has adopted a new version of Bluetooth that, for the first time, could be used to connect directly to the Internet.

This technology could prove popular as the Internet of Things (IoT)—in which many types of everyday devices connect online—becomes more widely used.

In addition to Internet connectivity, version 4.2 of the short-range, low-power wireless technology—which the Bluetooth Special Interest Group recently adopted—offers more speed and privacy than earlier versions.

Version 4.2 will let Bluetooth Smart sensors connect online via IPv6 over Low Power Wireless Personal Area Networks (6LoWPAN) technology. This approach enables even small objects—like those that could become part of the IoT—to access the Internet via a gateway.

Bluetooth 4.2 will use packets with 10 times the capacity of those in earlier versions of the technology. This will enable connection speeds that are 2.5 times faster and reduce transmission errors. In addition, the larger packets' increased efficiency will lower power consumption and increase devices' battery life.

The new standard will also enable users to encrypt transmissions, which is important for people who want to use their devices for secure activities such as locking and unlocking their home's front door.

To improve privacy, the specification will force *beacons*—Bluetooth-

0018-9162/15/$31.00 © 2015 IEEE

based technology that retailers use to transmit promotional messages to nearby shoppers' mobile devices—to obtain permission from users before tracking and contacting them.

Despite the improvements, industry observers say Bluetooth may not appeal to all users because it can be difficult to pair with devices.

Companies Look for Ways to Let Mobile Apps Work and Play Well with One Another

One frustrating issue for smartphone owners is their apps' inability to readily communicate with one another.

Users say this causes them to miss out on the benefits that would occur if, for example, an e-commerce program could talk to a consumer product-rating app so that people could learn more about potential purchases.

Now, though, large and small companies—including Facebook and Google—are trying to remedy this situation.

Allowing apps to communicate would not only benefit consumers but would also give search-engine operators, online-advertising sales companies, and other businesses access to considerable potentially useful information that is now isolated within apps.

However, doing so is a challenge, as mobile apps don't have links, Web addresses, or other mechanisms that make communication easy.

A key proposal addressing this issue is the use of *universal Web addresses* that work with programs—which could talk to each other via *deep linking*—as well as webpages.

Google is developing App Indexing technology to catalog and index application pages, whose data would then become available to the company's search engine. This could enable Google to better target its search-related advertising to individuals.

Twitter has developed Twitter Cards, which let users go from the company's app to other programs on their smartphones.

CALIFORNIA PRISONERS LEARN TO PROGRAM

Many companies search high and low to find a good programmer, and now, they may have a new place to look: prison.

Via a class called Code.7370, 18 inmates at California's San Quentin State Prison—near San Francisco—have begun learning software and Web development.

The program is sponsored by The Last Mile, a California nonprofit organization that provides business and technology training for incarcerated men and women. The purpose is to give them a way to make money while behind bars, to prepare them for jobs and other aspects of life after prison, and to reduce the chance they'll continue to commit crimes.

The six-month Code.7370 program is intensive, with classes held eight hours per day, four days per week.

The inmates were specially selected to participate in the program, based on their abilities and motivation to learn. Many have already taken courses in entrepreneurship and developed business plans for startups.

The coding class is held in a former prison print shop. In keeping with security procedures, inmates are strip-searched when they arrive and leave.

The prisoners face several obstacles that normal coding students don't confront. Some have been incarcerated for so long, they are almost completely unfamiliar with computer technology.

In addition, San Quentin regulations forbid Internet access, a must for most programming students. The Last Mile thus had to design Code.7370 with this in mind. The instructors from the Hack Reactor training company communicate with the students via the Google Hangouts instant messaging and video chat platform, the only outside connection permitted.

The Last Mile plans to work with the California Prison Industry Authority to help graduates get programming jobs—paying prevailing market wages—while still behind bars.

NEWS

Amputee Les Baugh, who lost his arms 40 years ago, uses his thoughts to control Johns Hopkins University's Modular Prosthetic Limbs and move objects from one shelf to another. Source: Johns Hopkins University Applied Physics Laboratory.

Famous Industries is taking a different approach by building websites—which are linkable and thus can communicate—that look and act like apps.

One concern is that some companies may not want their data available to other programs and may block access.

Also, if multiple companies develop their own application-communication technologies for different developers to adopt, many programs won't be able to talk to one another.

With this in mind, Facebook is working on an open standard for deep links for apps.

Users Can Control New Prosthetic Arms with Their Thoughts

For the first time, researchers have developed a pair of prosthetic limbs that amputees could control simultaneously, rather than just one at a time, with their thoughts.

The ability to create artificial limbs that can function much like real ones has been a longtime goal of prosthetic research

Researchers in the Revolutionizing Prosthetics (RP) program at Johns Hopkins University's Applied Physics Laboratory (APL) developed the Modular Prosthetic Limbs, which contain 100 sensors.

They tested the devices on Les Baugh of Colorado, who lost his arms 40 years ago in an electrical accident.

Before using the new arms, Baugh required reinnervation surgery, in which nerves in an amputee's shoulders or the remaining part of the arms are deactivated and then reactivated in ways that will let the muscles' thought-activated electromyographic signals control a prosthesis.

Explained Johns Hopkins trauma surgeon Albert Chi, "By reassigning existing nerves, we can make it possible for people who have had upper-arm amputations to control their prosthetic devices by merely thinking about the action they want to perform."

The researchers fitted Baugh with a socket that supports the new limbs, enables a lifelike range of movement, and connects with the reactivated nerves.

Then, they trained him to use his thoughts to control the devices. First, he worked with an APL system that uses pattern-recognition algorithms to identify the way his individual muscles communicate with other parts of his body. Researchers utilized this information to develop an approach that let him use his thoughts to send electrical signals through muscles in ways that control the prosthetic limb.

Baugh then practiced with a virtual-reality version of the limbs.

The researchers subsequently fitted him with the artificial arms, which he used to move several objects. For example, he could perform the complex motions required to transfer a cup from one shelf to another.

APL prosthetist Courtney Moran said, "This was significant because this is not possible with currently available prostheses. He was able to do this with only 10 days of training, which demonstrates the intuitive nature of the control."

"I think we are just getting started," noted RP principal investigator Michael McLoughlin. "There is just a tremendous amount of potential ahead of us, and we've just started down this road. I think the next 5 to 10 years are going to bring phenomenal advancement."

The Johns Hopkins researchers have sent Baugh home with a pair of the prosthetic arms he can use in his daily life.

Hackers Hit German Steel Factory and South Korean Nuclear Power Plant

A longtime fear of security experts is that hackers could attack power

plants, major factories, or other critical infrastructure systems and cause service disruptions, extensive damage, and even injuries or deaths.

Those fears were partially realized recently in incidents in Germany and South Korea.

Attackers broke into a German steel factory's production networks, manipulated a blast furnace's controls, and caused severe damage. This is one of the few cyberattacks that has ever caused physical destruction.

Hackers also breached the network of the company that runs South Korea's nuclear plants, while someone deliberately or inadvertently placed malware on one plant's control systems.

A report the German Federal Office for Information Security recently released revealed the attack on the unnamed steel factory, which occurred at an unspecified time in 2014.

The hackers used social engineering and spearphishing, in which they sent employees emails that looked like they came from within the company. This tactic is designed to, for example, convince workers to link to a fake company webpage and enter their office-network usernames and passwords.

Once inside the system, the attackers accessed production networks and began causing various components to fail. In one case, workers couldn't shut down a blast furnace properly, which damaged the entire plant.

In South Korea, hackers compromised the network of Korea Hydro & Nuclear Power, which operates the country's four nuclear power plants, in December 2014. The intruders—who security experts say could be from North Korea or a South Korean anti-nuclear power group—stole and posted online reactor designs, public-health monitoring data, and employee information.

Korean investigators identified a Chinese IP address in Shenyang, on the North Korean border, as the attack's source. Shenyang reportedly hosts North Korea's main Internet connection to the outside world. North Korea has denied involvement in the attack.

South Korea has asked the Chinese government to help with the investigation but hasn't received a response yet.

Korea Hydro & Nuclear Power said the breach didn't affect any of the nuclear plant's control systems, which reportedly aren't connected to externally facing networks.

During an audit related to the incident, the company identified a "low-risk" worm infection in devices connected to one plant's control systems, although none were found in the controls for the reactor itself. Security experts removed the worm, saying it was probably introduced by an employee's unauthorized attachment of a USB device to an in-house computer.

They speculate the worm is unrelated to past and ongoing cyberattacks on company systems.

Korea Hydro & Nuclear Power says it will enlarge its security staff.

South Korea is the world's fifth-largest user of nuclear power, which provides a third of the nation's energy.

The US Department of Energy is spending $325 million to build what could become the world's two fastest supercomputers.

US Plans World's Most Powerful Supercomputers

The US Department of Energy (DoE) is funding two supercomputers that could become the fastest in the world.

The DoE is using $325 million to develop the Summit supercomputer for the agency's Oak Ridge National Laboratory in Tennessee and the Sierra system at its Lawrence Livermore National Laboratory in California. The department plans to install them by 2017.

The DoE views the new systems as the next step on the path to exascale computing, in which systems could perform at least one exaflops (a quintillion floating-point calculations per second).

Summit and Sierra will each transmit more than 17 petabytes of data per second.

Summit would perform up to 300 petaflops, and Sierra would run 100 petaflops. Today's fastest computer is China's Tianhe-2, which performs 55 petaflops.

The new systems will use a supercomputing architecture that IBM has developed, in an effort to process today's huge data volumes better than current approaches. They will include IBM POWER CPUs, Nvidia Tesla GPUs, and Nvidia NVLink high-speed interconnects.

In traditional computing approaches, data moves regularly between processors and storage. However, this becomes time-consuming with large amounts of data.

IBM says its "datacentric" approach puts processing nodes wherever information resides within the system, eliminating the need for constant data movement. The company adds that this efficiency improves performance and reduces energy consumption.

According to IBM, Summit and Sierra will focus on Big Data processing.

DoE says the systems will be used for mission-critical applications and will transform research related to basic science, national defense, climate change, healthcare, manufacturing, the energy industry, and the environment.

NEWS

NASA built its RoboSimian robot to help after natural and man-made disasters by, for example, searching for and rescuing victims.

Sony will soon release the Single-Lens Display Module, which projects images or words onto a user's existing eyewear. Google Glass provides similar capabilities but comes with its own eyewear.

NASA Uses Apes as Model for Disaster-Response Robot

NASA's Jet Propulsion Laboratory (JPL) has developed an innovative ape-like robot designed to help with disaster-related activities such as victim searches and rescues.

RoboSimian has seven cameras that serve as its eyes, as well as wheels for coasting on flat surfaces. It has four limbs that can either grasp and manipulate objects or climb over debris and rugged terrain.

The machine also uses Lidar, a remote sensing technology that measures distances and then maps its surroundings by illuminating objects with a laser and analyzing the reflected light.

NASA is entering RoboSimian in the current $2 million DARPA Robotics Challenge, a competition in which teams from throughout the world design emergency-response robots.

Scientists and public officials consider such machines important because they could work in environments that are too hazardous for humans, such as nuclear-disaster zones or areas with extensive structural damage.

The robots could search for survivors, move rubble off of victims, or shut down utility systems' controls and thereby avert fires or other problems.

Many robots are designed to function like humans, walking upright and using mechanical arms for grasping and otherwise manipulating objects.

JPL, on the other hand, designed RoboSimian to move around like an ape, focusing on the ability to use all four limbs to provide stability and effectively traverse difficult terrain. However, this makes the robot slower than other DARPA Robotics Challenge competitors.

The NASA researchers are working with scientists from the University of California at Santa Barbara and Caltech to make it faster.

JPL designed RoboSimian's software to enable the robot to work largely on its own when communications with human controllers are interrupted, a frequent problem at disaster scenes.

In the past, JPL has built robots for work in space and has adapted some of this technology for use on Earth. With RoboSimian, they utilized some of the approaches they implemented in NASA's Mars Exploration Rovers and plan to employ some of the robot's new technology in its space-related projects for purposes such as assembly and exploration.

Sony System Converts Regular Glasses into Smart Spectacles

Sony has developed an attachable system that converts ordinary glasses into smart eyewear.

The Single-Lens Display Module includes a control board, processor,

This illustration by the Fraunhofer Institute shows how the organization's new passive coherent location (PCL) technology would enable officials to use cellular towers like radar to detect small boats, perhaps occupied by terrorists, trying to sneak into a harbor.

Bluetooth radio, and color organic-LED (OLED) display that projects images onto a user's existing glasses, goggles, or sunglasses. This turns the eyewear into heads-up displays of useful or entertaining information or images.

The module differs from Google Glass, which is standalone eyewear into which prescription lenses can be inserted.

The new display measures just 0.23 inches diagonally but provides images with a resolution of 640 × 400 pixels within a small window that doesn't obstruct wearers' vision.

Sony says its module and the enhanced information it provides could help users at work, while traveling, or during outdoor activities and sports.

The device is different from Sony's Smart EyeGlass, which is similar to Google Glass. The EyeGlass is self-contained, stand-alone eyewear that projects information from an Android phone onto the lens.

Industry observers say the new module could be particularly successful because many users might prefer to benefit from the capabilities provided while keeping the spectacles they already have or when getting new ones, something not possible with Google Glass.

Sony says another advantage is that the device is easily attachable and detachable and thus can be worn only when the user wants to do so.

The company says it plans to begin mass production of its device this year and will also release software developer's kits to partnering organizations.

New Approach Lets Cell Towers Function like Inexpensive Radar

Many small seaports can't afford expensive radar systems. However, this can make it difficult for officials to detect the secret entry of small ships operated by, for example, terrorists or smugglers.

Now, though, researchers at the Fraunhofer Institute for Communication, Information Processing, and Ergonomics have developed a way to use cellular towers to function like radar systems and detect such activity.

The approach works via passive coherent location (PCL) technology, which bounces cellular towers' radio signals off of objects and measures the return signals to determine the subjects' location.

The technology requires the

installation of relatively inexpensive PCL equipment rather than costly radar systems.

However, unlike radar, PCL doesn't produce strong signals. And the return signals are even weaker and thus tougher for the new system to distinguish from cellular and other transmissions in the area. The Fraunhofer researchers thus had to develop algorithms that better identify the return signals.

Their system is now sensitive enough to recognize small vessels from as far away as 4 kilometers and to even track them as they travel through the water.

Researchers Find Critical Flaw in Important Cloud Technology

Security researchers have found a major vulnerability in the increasingly popular open source Docker platform that is considered an important new cloud technology. Docker makes it easier to download applications over the Internet and run them on various types of machines via cloud platforms.

Florian Weimer, a member of open source software vendor Red Hat's Product Security team, and independent security researcher Taunis Tiigi recently found the flaw in all but the platform's newest iteration.

Through version 1.3.1, noted Docker Inc. in an online security advisory, "This vulnerability could be leveraged to perform remote code execution and privilege escalation" and ultimately steal files hosted in the cloud.

The company released version 1.3.2 to remedy the problem and advised users to download it as soon as possible, as there is no fix for the recently found bug. Docker subsequently released version 1.3.3 to deal with other issues.

Docker security is important for cloud users, as major technology companies such as Amazon and Google are supporting the platform.

Developers use Docker to put applications in software containers so that users can download them across the Internet or any private network and then run them on any cloud platform or Linux machine.

This would benefit cloud computing, which is typically used to make applications that are kept online available to all types of computing devices.

Proponents add that Docker also makes developers' lives easier by letting them focus on designing programs without worrying about the platform on which they will run.

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

GUEST EDITORS' INTRODUCTION

Technological Advances in Medicine: It's Personal

Alf Weaver, University of Virginia
Renée Bryce, University of North Texas

Current technological advances and those still being developed are poised to revolutionize medicine—creating tremendous opportunities for real-time, personalized patient monitoring and treatment, but also posing significant risks for medical data security.

GUEST EDITORS' INTRODUCTION

What achievements will we see in medicine over the next 5 to 10 years? Rapid technological advances, driven in part by our growing understanding of the human body and how it works, allow our caregivers—and us as patients—to interact with our bodies in ways previously unimaginable. We stand at the brink of truly personalized medicine that replaces a "one size fits all" approach with individualized attention to the specific characteristics that make a patient unique.

No longer is a breast cancer diagnosis the end of the inquiry; rather, it is merely the beginning. Sequencing a patient's genome enables physicians to determine what kind of breast cancer she has and then tailor effective, evidence-based treatment for the desired outcome. At the same time, low-cost health-monitoring devices and personal health records allow that patient to take a more active role in monitoring and managing her own overall well-being.

Accomplishing a grand vision for truly personalized medicine starts with patient data originating from many sources. This data must then be stored so that it is reliable, accessible, and sharable—yet, at the same time, secure. Required, then, are system architectures that guarantee these attributes. While such systems will continue to "live" in medical enterprise environments such as hospitals and clinics, personal wearable devices—whether generic along the lines of smartphones, or customized like Fitbit and its competitors—are making significant inroads into medical data reporting, collection, and storage.

Indeed, mobile devices have matured into intelligent data-gathering machines, able to collect and analyze physiological data and then report results to the wearer as well as to any remote observers or monitors—whether human, software, or some combination of both. Properly designed, such smart systems can be used within hospitals or link to remote locations, augmenting the equipment and personnel that make up current medical infrastructures. At the same time, we see an explosion in inexpensive health monitors and fitness trackers primarily for personal use. All of these trends create an incontrovertible need for new security requirements and privacy controls: given the Internet's infinite memory, once medical data has been promulgated publicly, it can never be effectively purged.

> **ACCOMPLISHING A GRAND VISION FOR TRULY PERSONALIZED MEDICINE STARTS WITH PATIENT DATA THAT IT IS RELIABLE, ACCESSIBLE, AND SECURE.**

IN THIS ISSUE

Focusing on the need for privacy and security, Jinquan Li's "Ensuring Privacy in a Personal Health Record System" draws an important, but often overlooked, distinction between electronic medical records (EMRs), typically generated and held by health professionals, and the personal health records (PHRs) that individuals can store electronically to manage and share their own health information. Based on how they originate, EMRs as well as PHRs "tethered" to a specific healthcare organization in the US are typically covered by the federal Health Insurance Portability and Accountability Act of 1996 (HIPAA), which governs privacy issues related to EMRs (among other matters). However, Web- and device-based PHRs have no such legal protection, so for them to enjoy widespread acceptance patients must be able to trust those companies and systems that collect, store, and disseminate PHR data. As Li suggests, changing the locus of control from providers to consumers will not be a simple matter. The risks and repercussions inherent in PHR data disclosure, whether intentional or accidental, can be serious; aggregated health information is highly valuable to third parties such as drug manufacturers, insurers, marketers, and employers—who may not always have patients' best interests in mind. System architects, take note!

Mobile devices and their increasingly significant role in health management is the topic of "Intelligent Disease Self-Management with Mobile Technology," by Marina Velikova, Peter J.F. Lucas, and Maarten van der Heijden. Just as today our minds can be connected 24/7 to the Internet, soon our bodies will be continuously reporting personal physiological status to software for recording, analysis, and prediction. Fully exploiting a smartphone's inherent instrumentation, such as using its microphone

to measure lung function or its camera to determine blood-oxygen saturation, will open new opportunities for disease self-management. The coming rush of customized, snap-in options will accelerate technological progress—but also exacerbate its perils. New hardware and new Web and mobile apps intended to enable effective, personalized medicine will inevitably create new problems as well.

One such potential problem involves quality of service (QoS). Consider this range of cases:

- a hospital patient on a post-op floor whose heart rate is being monitored;
- a patient in intensive care where multiple physiological quantities are monitored and analyzed;
- home-located patients or remote clinics that send and receive data consistently but on an unscheduled basis;
- doctors and other clinical personnel who oversee all these activities using mobile devices; and
- hospital monitoring stations that must carefully watch overall system performance.

Across such varied instances, can mobile devices adequately collect and display data in real time and maintain QoS? In "Medical-Grade Quality of Service for Real-Time Mobile Healthcare," Kyungtae Kang, Qixin Wang, Junbeom Hur, Kyung-Joon Park, and Lui Sha investigate the parameter space in which systems can supply "a level of [data] transmission speed, reliability, privacy, and security that provides real-time, confidential, and accurate service" for both in-hospital and remote applications.

Given the tremendous amount of medical data that can be generated, where should it all be kept? How should it be shared? Can it be adequately secured? If we solve those problems, what might actually be accomplished with the data? "Healthcare Data Integration and Informatics in the Cloud," by Arshdeep Bahga and Vijay K. Madisetti, proposes a new framework based on their prototype Cloud Health Information Systems Technology Architecture (CHISTAR) middleware to coordinate cloud-based data analytics for collecting, organizing, and securely exchanging healthcare data from a range of stakeholders in a range of formats. Software in the form of a Web and mobile app builder lets users implement multiple functions including epidemiological surveillance, adverse drug event prediction, and medical prognosis.

Even as mobile devices are poised to revolutionize personalized medicine, widespread and affordable whole genome sequencing (WGS) could one day go even further—presenting a novel set of benefits and challenges. In "Whole Genome Sequencing: Revolutionary Medicine or Privacy Nightmare?," Erman Ayday, Emiliano de Cristofaro, Jean-Pierre Hubaux, and Gene Tsudik explain that WGS will usher in a new era of "predictive, preventative, participatory, and personalized (P4) medicine." But like most technology, WGS presents a double-edged sword: on one hand, it can pinpoint and predict disease, allowing early-stage, life-saving treatments; on the other, a genome sequence's detailed biometric specificity offers opportunities to irrevocably compromise privacy.

Personalized medicine, already a reality, will inevitably grow in reach and impact as its diagnostic and predictive power expands. As technologists, our duty is to advance medical hardware and software in any way we can, while simultaneously enforcing data privacy—all with a spirit of passion and innovation.

ABOUT THE EDITORS

ALF WEAVER is a professor of computer science and founding director of the Applied Research Institute at the University of Virginia. His research interests include computer networks, network protocols, telemedicine, electronic commerce, medical data privacy and security, and crowdsourcing. Weaver received a PhD from the University of Illinois. He is an IEEE Fellow and served as an ACM National Lecturer. Contact him at acw@cms.mail.virginia.edu.

RENÉE BRYCE is an associate professor of computer science and engineering and director of the Software Testing Lab at the University of North Texas. Her research interests include software testing, specifically combinatorial testing in relation to Web and mobile applications. Bryce received a PhD from Arizona State University. Contact her at reneebryce@gmail.com.

AUDIO

Ensuring Privacy in a Personal Health Record System

Jingquan Li, Texas A&M University

Personal health records are integral to self-managed healthcare, but patient-controlled access comes with serious concerns that require a better balance of personalization, privacy protection, and security controls.

Personal health records (PHRs), which give patients the opportunity to store, manage, and share their personal health information, have wide-ranging implications for personal health maintenance and healthcare. Chronically or seriously ill patients can keep track of their disease and its symptoms and treatments and maintain an ongoing connection with their care providers (physicians, nurse practitioners, psychiatrists, laboratory personnel, and so on). For those with less serious, intermittent, or even no immediate health problems, PHRs can be invaluable in self-managing health, which could reduce healthcare costs. In treatment, PHRs improve patient–care provider communication and enable much more complete information than might be available in traditional treatment, particularly for emergency care, which might take place far from the patient's primary care provider.

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

Recognizing these benefits, more healthcare providers, insurers, and employers are offering PHRs, and, consequently, an increasing variety and amount of information flows in and out of a PHR database. As the sidebar "What Does a PHR Contain?" describes, a single PHR can contain information on the patient's medical conditions and histories, medications, mental health, genetic makeup, sexual behavior, lifestyle, beliefs, and habits. Some of this data must be shared for patients to receive proper medical care, but otherwise all data must remain private, since unauthorized disclosure could harm the patient.[1] Because the data in a PHR has high commercial value, PHRs make a tempting target for unscrupulous marketers, identity thieves, and corrupt organizations.

Thus, PHR systems offer new opportunities for personalized healthcare management, but they come with serious privacy and confidentiality risks. Patients are understandably worried about their data's secondary use, which erodes confidence in a PHR system and slows its acceptance. To reverse this trend, researchers must address both technical and legal challenges in preventing unauthorized PHR access and data use, which can

0018-9162/15/$31.00 © 2015 IEEE

TABLE 1. Comparison of personal health record (PHR) properties.

Attribute	Tethered PHRs	Untethered PHRs	
		Web-based	Device-based
Interoperability	Not interoperable	Interoperable	Interoperable
Accessibility	Portal or client server	Internet portal	PC-based device drive
Data sources	Electronic medical records (EMRs)	EMRs and consumer-added information	EMRs and consumer-added information
Completeness	Incomplete	Complete or partial	Complete or partial
Integrity	High	Depends on the accuracy and consistency of user-supplied data over its life cycle	Depends on the accuracy and consistency of user-supplied data over its life cycle
Major risks	Transfer to other PHR systems might not be feasible; consumer might not be able to enter data	Service provider and business partners might use PHR information for commercial or other secondary uses	Physical loss, theft, damage, and security risks
Compliance with legal mandates	Covered under the Health Insurance Portability and Accountability Act of 1996 (HIPAA)	Not covered under HIPAA	Not covered under HIPAA
Privacy control	Primary care site controls data	Consumer and service provider control data	Only consumer controls data
Security governance	Secure extranet portal	Acceptable with strong encryption and authentication	Acceptable with strong encryption and access control
Sample installations or trials	Mayo Clinic and Kaiser Permanente	Dossia Consortium's dossia.org and Microsoft's HealthVault	CapMed's HealthKey and MedicAlert's E-HealthKEY

lead to employer or insurer discrimination, medical identity theft, and commercial exploitation.[2]

PHR use is relatively new, so work to resolve trust and privacy issues has not investigated how to exploit a PHR system's architectural properties in providing patient-controlled data access and protection. Each PHR system type has specific advantages that vendors might merge into a hybrid architecture, which would give patients sole control over their PHR through privacy principles such as independent consent management, independent privacy and security audits, and regulatory compliance. Architectural design must consider privacy and security risks, some of which are unique to PHR access.

Although vendors have taken steps toward moving PHR control to the patient or consumer by offering Web-based PHR systems, many privacy issues remain open, such as how to enforce rules about the data's secondary use. Even so, PHRs have compelling advantages that motivate both patients and PHR system designers to persevere in creating a product that will satisfy the needs of both parties.

TYPES OF PHR SYSTEMS

PHR systems are either tethered or untethered to a healthcare provider, health plan, or other related entity. Table 1 gives some architectural properties for each PHR type. A *tethered* PHR is an extension of the system that manages a healthcare organization's electronic medical records (EMRs); the process of populating the PHR with EMR data is generally automatic and patients have only limited access to it.

In contrast, *untethered* PHR systems are Web-based, with PHRs stored on the Web or in the cloud, or device-based, with PHRs stored on a mobile device such as a smartphone. Both types of untethered PHRs give patients more control over the type and amount of PHR data and how to manage it, but security and privacy are also more open.

As Table 1 implies, interoperability, accessibility, completeness, integrity, and data sources define the PHR system's operational characteristics. Major risks, legal constraints, privacy, and security define the PHR system's limitations and possible acceptance barriers.

Because each PHR type has both advantages and limitations, a hybrid PHR system could be a promising new approach.

Tethered

A tethered PHR is a closed solution in that the hosting organization controls access and security, offering applications for the patient's access to the clinic's EMRs. Medical institutions that use tethered PHRs are typically prominent entities, such as the Mayo Clinic, or within health plans, such as Kaiser Permanente. The aims of tethered PHRs are to let patients access

COMPUTING IN HEALTHCARE

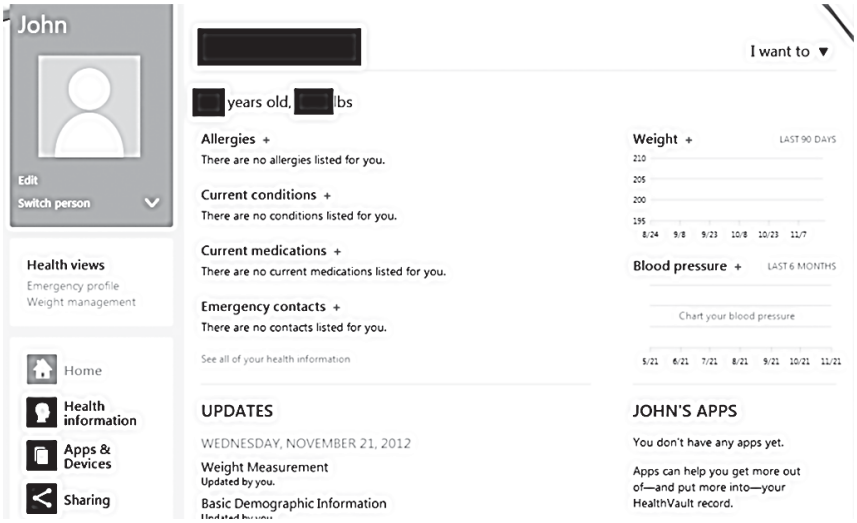

FIGURE 1. HealthVault homepage. HealthVault, a Web-based personal health record (PHR), gives consumers more control over their health choices by giving them the freedom to choose the information they want to include and with whom to share it.

medical records, share information more easily with their doctors, and better manage their and their families' health between scheduled visits. A strong advantage is controlled security and privacy, since providers must adhere to the Health Insurance Portability and Accountability Act of 1996 (HIPAA), the federal law that strictly limits access to an individual's health records.

Although the tethered PHR seems ideal on the surface, it is institution centered, which gives rise to major limitations. Because the provider controls information access, the amount of data in a PHR will vary considerably across institutions, and the data will be neither portable nor lifelong.[3] The PHR will have the hosting institution's EMR system format, which in most cases prevents the patient from sharing it with other medical providers and clinicians. When the patient's relationship with the institution ends, the PHR is no longer accessible.

The tethered PHR is also not a popular choice for most providers, since they receive no reimbursement for efforts to encourage patient care between visits. It is certainly a step toward giving patients access to their health information, but the mobile healthcare (m-health) trend is shifting the demand to more consumer-centric solutions.

Untethered: Web-based

At the other end of the PHR spectrum are Web-based PHRs, which are backed by technology vendors such as Dossia Consortium (www.dossia.org), Microsoft's HealthVault, WebMD, and PatientsLikeMe. The homepage of Microsoft's HealthVault in Figure 1 gives an idea of the control consumers have in managing their health information.

This PHR type is appealing because consumers have direct, portable access to personal health information, which they can then make available to a variety of care providers. Unlike the tethered PHR, the Web-based PHR lets consumers decide what data to include and with whom to share it. Care providers benefit from having a holistic window into the patient's entire medical history,[4] including information such as diet and exercise level, which could have a bearing on the patient's current condition. Some Web-based PHR systems, such as HealthVault and WebMD, are free downloads.

Portability is a primary advantage. For example, if a diabetic patient is rushed to an emergency room not affiliated with the patient's primary care provider, the emergency room physician still has access to information vital for treatment.

Countering these strong advantages is the potential for compromised privacy and security. HIPAA's privacy rule applies only to covered entities—healthcare providers, health plans, and healthcare clearinghouses[5]—not to technology vendors and their business partners. What happens when the information leaves the covered entities' files and ends up in a technology vendor's server?

Although Web-based PHR vendors vow to put users in control, they are also inserting themselves between patients and their most intimate data. The business model behind these services will depend heavily on advertisement revenue and partnerships with third-party suppliers of health-related products and services. This reliance raises important questions:

- Will the vendor's desire for more revenue result in sharing medical details either with marketing companies or with their numerous service partners?
- Will consumers have sole control over their data?
- If consumers want to close the account, can they dispose of the data as they see fit?

Not providing definitive answers to these questions will erode trust. Consumers might believe that the company offering PHRs is committed to privacy protection, but if they start receiving applications and services from third parties, they will realize that too many nonmedical companies are seeing their sensitive information.

Untethered: device-based

In contrast to the extreme openness of

WHAT DOES A PERSONAL HEALTH RECORD CONTAIN?

On 1 February 2009, President Barack Obama signed into law the American Recovery and Reinvestment Act of 2009 (ARRA), which established billions of dollars in incentives for clinicians and hospitals to use health information technology, including electronic medical records (EMRs). The law also clarifies that individuals have the right to receive copies of personal medical records from health practitioners in electronic formats and authorize their information to be stored in a service of the individual's choosing.

These records are the basis for the patient's personal health record (PHR), an electronic collection of personal health data from a variety of sources. A PHR system enables patients to electronically store, manage, and share their own health information apart from the electronic files or hard-copy records that healthcare providers maintain.[1] A seriously ill patient, for example, might visit a variety of care providers—each with separate records on that patient's treatment and health history, laboratory results, medications, and personal information. Remarkably, no healthcare institution is responsible for ensuring that a patient has complete records when that patent seeks care.

A PHR contains all or part of the information in Figure A, including medication records of pharmacies, test results from laboratories, claims data from insurers, and the patient's own notes about overall health, such as blood-sugar level or weight changes. The PHR system's primary goal is to empower the patient to access, understand, manage, and share health information with authorized parties in a private, secure, and confidential environment.[2]

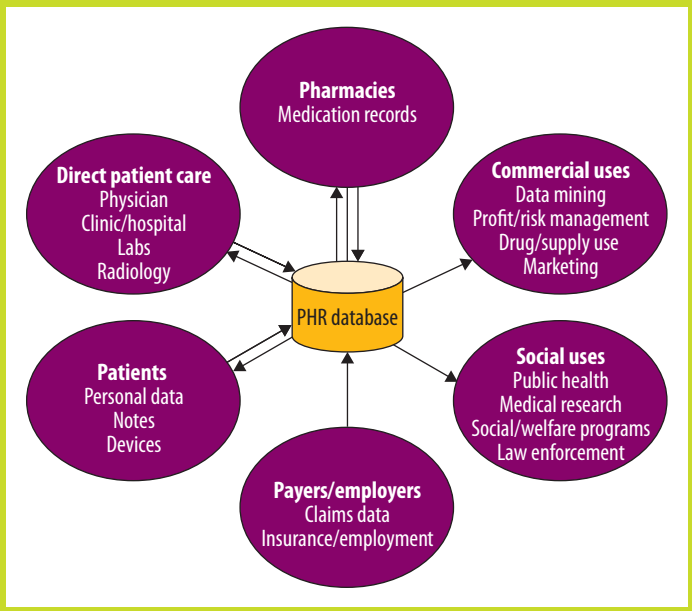

Figure A. Personal health information sources. A PHR can contain information from many sources, including data from the patient's smartphone or tablet.

Providing the patient with information control is the primary difference between PHRs and EMRs. Only healthcare professionals employed by medical institutions such as hospitals and doctors' offices can access EMRs, while patients have full access to their PHR and can decide what institutions will share that access.

References

1. J.M. Grossman, T. Zayas-Caban, and N. Kemper, "Information Gap: Can Health Insurer Personal Health Records Meet Patients' and Physicians' Needs," *Health Affairs*, vol. 28, no. 2, 2009, pp. 377–389.
2. "Connecting for Health: The Personal Health Working Group Final Report," 2003; www.providersedge.com/ehdocs/ehr_articles/The_Personal_Health_Working_Group_Final_Report.pdf.

a Web-based PHR, device-based PHRs are recorded and stored in a PC or mobile device, which typically has software with password protection, data encryption, data importing rules, and controlled data sharing. Examples include CapMed's Personal HealthKey and MedicAlert's E-HealthKEY. Some device-based PHR products allow PHR copying to a storage device such as a smart card or USB flash drive.

Mobile devices let consumers access a Web-based PHR system through their smartphones or tablets, for example. In an emergency or scheduled appointment, the care provider can open the device through any wireless network to gain immediate access to the patient's PHR and display the information on the institution's computer.

Despite software protection mechanisms, privacy and security assurance is still an issue. Tests have shown that the encryption and passwords in many device-based PHR products are weak.[6] Also, the device that holds the PHR is subject to damage or loss, as in a natural disaster, or theft, which can severely compromise the consumer. On the provider side, compromises can occur when a care provider connects the patient's device to a computer that is storing sensitive medical information on other patients.

Hybrid systems

To overcome the major risks and limitations of tethered and untethered PHR systems, a flexible PHR system must combine the best of both types. For example, the tethered or Web-based PHR might also allow consumers to download the PHR to a USB drive or a personal device. Hybrids that allow PHR redundancy by offering both local and remote storage will provide the greatest availability of health data under any circumstances and the most connection flexibility.

PRIVACY AND SECURITY RISKS

Maintaining data consistency and integrity in a hybrid PHR system will require mechanisms that protect the PHR across systems. A PHR product's success will depend heavily on its data access scheme: personal health information must be easily available to those who legitimately need it and strictly off limits to those who do not.

Threats include leaks due to patient confusion about access management, leaks from the healthcare provider, outside attacks, and data mining for commercial use.

Patient errors in access control

The duality of consumer empowerment and increased privacy risk can be problematic. Healthcare providers have long controlled medical records' storage and access, and moving that control locus to the patient is not straightforward. Overwhelmed by their new responsibility, patients might err in granting or revoking data access, resulting in information leaks or loss of data integrity. Individuals who lack the expertise to maintain their health data online or through a self-contained device might mistakenly delete parts of the PHR that are vital to treatment, for example. Poor PHR system design and improper use can cause patient errors in access control that endanger patient safety or decrease the quality of care.

Provider leaks

Healthcare providers can mistakenly or intentionally compromise patient information. A PHR system is complex, and care providers can choose options that lead to information disclosure. Insiders might leak patient information for spite or profit. Regardless of intent, the consequences of these leaks can have serious consequences: the literature has solidly documented the risks of releasing information about infectious conditions, mental health, chronic disease diagnoses, and genetic makeup.

PHRs contain highly sensitive health information that discloses the patient's identity.[5] Once care providers expose the PHR, the patient can do little to restore privacy. Chronically ill patients, in particular, are acutely concerned about privacy risks and generally unwilling to make their health information accessible. Yet these are the individuals who would benefit the most from PHR use.

Outside attacks

An Internet-accessible PHR is vulnerable to attacks from hackers and other unauthorized parties. Also, the physical device storing the PHR might be stolen, possibly exposing banking information and other sensitive data.

Commercial data mining

Some companies offer Web-based PHR systems on platforms that warehouse and mine personal health data for business and proprietary purposes. Aggregated personal health data has high value to pharmaceutical companies, marketers, insurers, and employers.

Although patients appear to have complete control over information access, as in the Web-based PHR system, they actually have little control over how the PHR system vendor and its numerous business partners will use that information. Some secondary uses are beneficial, such as the surveillance of adverse events from prescription drugs or health monitoring; others are marketing scams or result in employer and insurer discrimination.

Patients have little recourse to counter this less benign secondary use. The Health Information Technology for Economic and Clinical Health Act of 2009 (HITECH) significantly expands the financial risk of HIPAA violations and extends HIPAA provisions and penalties to business associates. However, to date, PHR vendors are bound by HITECH only through the Federal Trade Commission's rule that requires certain breach notification procedures; the rule does not define what constitutes the PHR vendor's appropriate use or disclosure of health information. Use or disclosure rules apply only if the PHR vendor enters into a business associate agreement with a covered entity. Even then, if the

vendor declares bankruptcy, any prior privacy agreements become void.

TOWARD A PATIENT-CONTROLLED PRIVACY POLICY

To meet the unique privacy challenges of PHR systems, vendors are working to create a formal policy that specifically addresses PHR information disclosure and use. One approach is based on the personally controlled heath records (PCHRs) platform developed at the Children's Hospital Informatics Program (CHIP) at Children's Hospital in Boston,[2,7] which lets patients assemble, maintain, and manage a secure copy of their medical data.

Indivo,[7] an Internet-based PCHR implementation, is in use worldwide and has served as a reference model for some Web-based PHR systems including dossia.org and HealthVault. Indivo uses security measures to minimize the risk that an unauthorized party could gain access to sensitive PHRs.

As a privacy protection strategy, the PCHR approach has limited value. Because both the patient and PHR vendor control the PCHRs, PCHR privacy depends on the vendor's business model and commercial interests. With revenue being critical to the success of any PHR product,[8] allowing market forces to determine a PHR system's nature and direction would be unwise. Giving patients sole access control is the only way to stop most PHR data exploitation.

Other privacy protection approaches, such as the W3C Platform for Privacy Preferences (P3P) and the Platform for Enterprise Privacy Practices (E-P3P),[9] attempt to block unauthorized access and data use. Again, however, without more patient control, authorized individuals and entities such as the PHR vendor and its business partners can use PHR data as they desire.

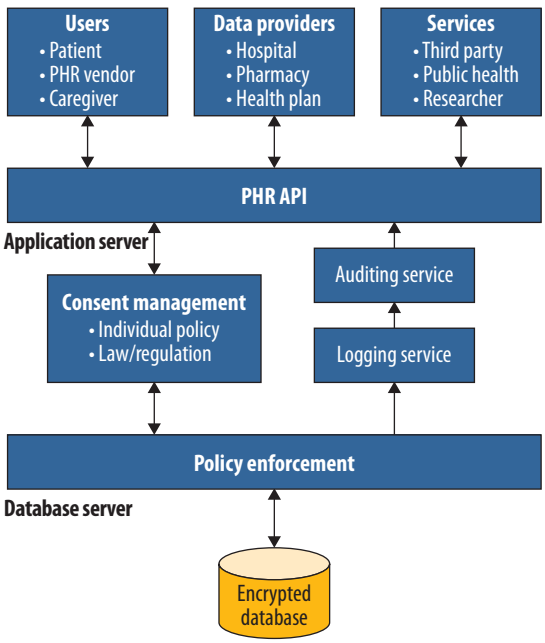

FIGURE 2. Consumer-controlled PHR system architecture. The presentation tier (API) helps patients specify a privacy policy, which the application server tier uses to control data access. The policy enforcement tier enforces the patient-defined privacy policy by storing every access request in a logging file.

These access and sharing loopholes beg the question, "What *is* PHR privacy?" Quite simply, it is the patient's ability to control PHR access and ensure the security and integrity of PHR information.[10] It is also the patient's trust in that ability. Designing a PHR system that puts patients in full control of their own data can go a long way toward mitigating their privacy concerns. Thus, patient-controlled privacy protection involves building independent consent management, compliance with regulatory requirements, and independent privacy and security audits into the PHR system.

Consent management

Patients want to control PHR access and data use in the same way that they control the funds in their bank account. Ensuring personal control over information access and use fosters trust between consumers and the technology provider, but letting technology companies manage such sensitive information makes some patients uncomfortable. On some level, they realize that they do not have full control over their data.

For many patients, the ability to authorize *any* use of their personal health information is the essence of privacy. Independent consent management—electronic data access consent apart from the PHR vendor—is a way to assure the patient of that ability. Unless the patient gives explicit consent, even the PHR vendor cannot share or use personal health information for nonmedical purposes. Electronic consent management must be a fundamental architectural component of any consumer-controlled PHR system. Embedding independent consent management into the system's design ensures that the PHR vendor goes through the patient before gaining knowledge about any patient-related data.[11] Patients can be more confident that they have complete control over who uses their data, when, and to what end.

Automated regulatory compliance

Not all patients have the knowledge and technical skills to maintain their

data's confidentiality. Even with informed consent, a default privacy and security policy that complies with regulatory requirements and fair information practices would provide considerable credibility and foster greater confidence in the system, which can lead to wider adoption.

Auditing access and use

Auditing every disclosure and use of PHR data empowers patients to identify any individuals or entities who have accessed or altered their PHR data. The PHR system must keep complete audit trails of every PHR data disclosure and use, including the person or entity accessing or using the data and the time and purpose of access and use. An independent verification or audit service ensures that the PHR system actually implements the privacy policy that the patient has created. By reviewing a list of PHR account accesses, the patient can detect a breach and take appropriate action.

Monitoring and auditing data use is a practical solution for determining the quality and integrity of a privacy-preserving PHR system. For example, Microsoft's commitment to independent third-party audits sets a new standard for privacy protection in PHR systems. Independent privacy and security logging and auditing mechanisms are critical to the patient's acceptance of and trust in the PHR system.

Implementation architecture

Most Web-based PHR systems use the three-tiered architecture in Figure 2, which consists of the presentation (API), application server, and database server tiers.

Presentation tier. The presentation tier handles data access requests, presents the information to data users and services, and helps patients specify the privacy policy that the application server and policy enforcement tiers will use. Any action the presentation tier tries to perform passes through the application server tier and is thus subject to the enforcements in the chosen privacy policy.

Application server tier. The application server tier—at the heart of patient-controlled privacy—is responsible for the business logic of serving the wide-ranging information needs of healthcare organizations in treating the patient and providing data stripped of the patient's identification to commercial and research enterprises. The application server tier extends the Web-based PHR system with three services: consent management, policy enforcement, and logging and auditing.

Consumers use the *consent management* service to control PHR access and data use. Even with explicit consent, the privacy policy must comply with regulatory requirements and be easy to audit.

As its name implies, the *policy enforcement* service implements the privacy policy, which the patient has created to address privacy risks from all stakeholders: consumers, healthcare professionals, the PHR vendor, the vendor's business partners, and even the patient (who might inadvertently compromise privacy).

Every access request creates an event that transfers from the policy enforcement service to the *logging* service, which stores every disclosure and access in log files. The *auditing* service shows a log file of access requests as proof that the PHR system is implementing the privacy policy.

Database server tier. The database server stores PHRs, which are encrypted to protect patients even during a data breach, hacker attack, or hardware theft. The encryption keys are kept on a separate physical server to prevent decryption of health data if the database server is ever compromised. Server storage and data network encryption work in concert to provide another layer of protection.

Few will debate the merits of PHRs in ushering in the age of m-health and patient-controlled disease management. However, privacy remains a threat to PHR systems' widespread adoption. Vendors offering PHRs want to keep their reputations by making good on their promises to put patients in control. Despite their intentions, privacy risks continue to sabotage the complete acceptance of PHR systems.

Although giving patients sole control over PHR access and data use is extremely difficult, it is the best way to ensure the appropriate use of PHR data. A patient-controlled privacy protection architectural scheme that is based on fundamental privacy principles and accounts for the unique characteristics of PHR systems is only part of the solution. Other components include detailed technology design, regulatory implementation mechanisms, and consumer education about the PHR benefits and data privacy management.

Another critical issue for PHR research is how to align the PHR vendor's commercial interests and incentives for protecting the privacy and confidentiality of PHR data. The PHR vendor requires some way to balance its need for revenue or other return on investment with the patient's need for

privacy. A balance that works to the advantage of both parties is key to PHR system acceptance.

REFERENCES

1. T. Rindfleisch, "Privacy, Information Technology, and Health Care," *Comm. ACM*, vol. 40, no. 8, 1997, pp. 93–100.
2. J. Li, "Privacy Policies for Health Social Networking Sites," *J. Am. Medical Informatics Assoc.*, vol. 20, no. 4, 2013, pp. 704–707.
3. J.M. Grossman, T. Zayas-Caban, and N. Kemper, "Information Gap: Can Health Insurer Personal Health Records Meet Patients' and Physicians' Needs," *Health Affairs*, vol. 28, no. 2, 2009, pp. 377–389.
4. N. Archer et al., "Personal Health Records: A Scoping Review," *J. Am. Medical Informatics Assoc.*, vol. 18, no. 4, 2011, pp. 515–522.
5. J. Li and M. Shaw, "Protection of Health Information in Data Mining," *Int'l J. Healthcare Technology and Management*, vol. 6, no. 2, 2004, pp. 210–222.
6. A. Wright and D.F. Sittig, "Encryption Characteristics of Two USB-Based Personal Health Record Devices," *J. Am. Medical Informatics Assoc.*, vol. 14, no. 4, 2007, pp. 397–399.
7. K.D. Mandl et al., "Indivo: A Personally Controlled Health Record for Health Information Exchange and Communication," *BMC Medical Informatics and Decision Making*, vol. 7, no. 25, 2007, pp. 1–10.
8. L. Martino and S. Ahuja, "Privacy Policies of Personal Health Records: An Evaluation of their Effectiveness in Protecting Patient Information," *Proc. 1st ACM Int'l Health Informatics Symp.* (IHI 10), 2010, pp. 191–200.
9. G. Karjoth, M. Schunter, and M. Waidner, "Platform for Enterprise Privacy Practices: Privacy-Enabled Management of Customer Data," *Privacy Enhancing Technologies*, LNCS 2482, 2003, pp. 69–84.
10. "Guidance from CDC and the US Department of Health and Human Services," Centers for Disease Control and Prevention (CDC) HIPAA Privacy Rule and Public Health, *MMWR Morbidity and Mortality Weekly Report*, May 2, 2003, vol. 52, no. S-1, pp. 1–12; www.cdc.gov/mmwr/preview/mmwrhtml/su5201a1.htm.
11. S. Haasa et al., "Aspects of Privacy for Electronic Health Records," *Int'l J. Medical Informatics*, vol. 80, no. 2, 2011, pp. e26–e31.

ABOUT THE AUTHOR

JINGQUAN LI is an associate professor of information at Texas A&M University—San Antonio. His research interests include information security and privacy and data mining. Li received a PhD in information systems from the University of Illinois at Urbana–Champaign. He is a member of the IEEE Computer Society. Contact him at jli@tamusa.tamus.edu.

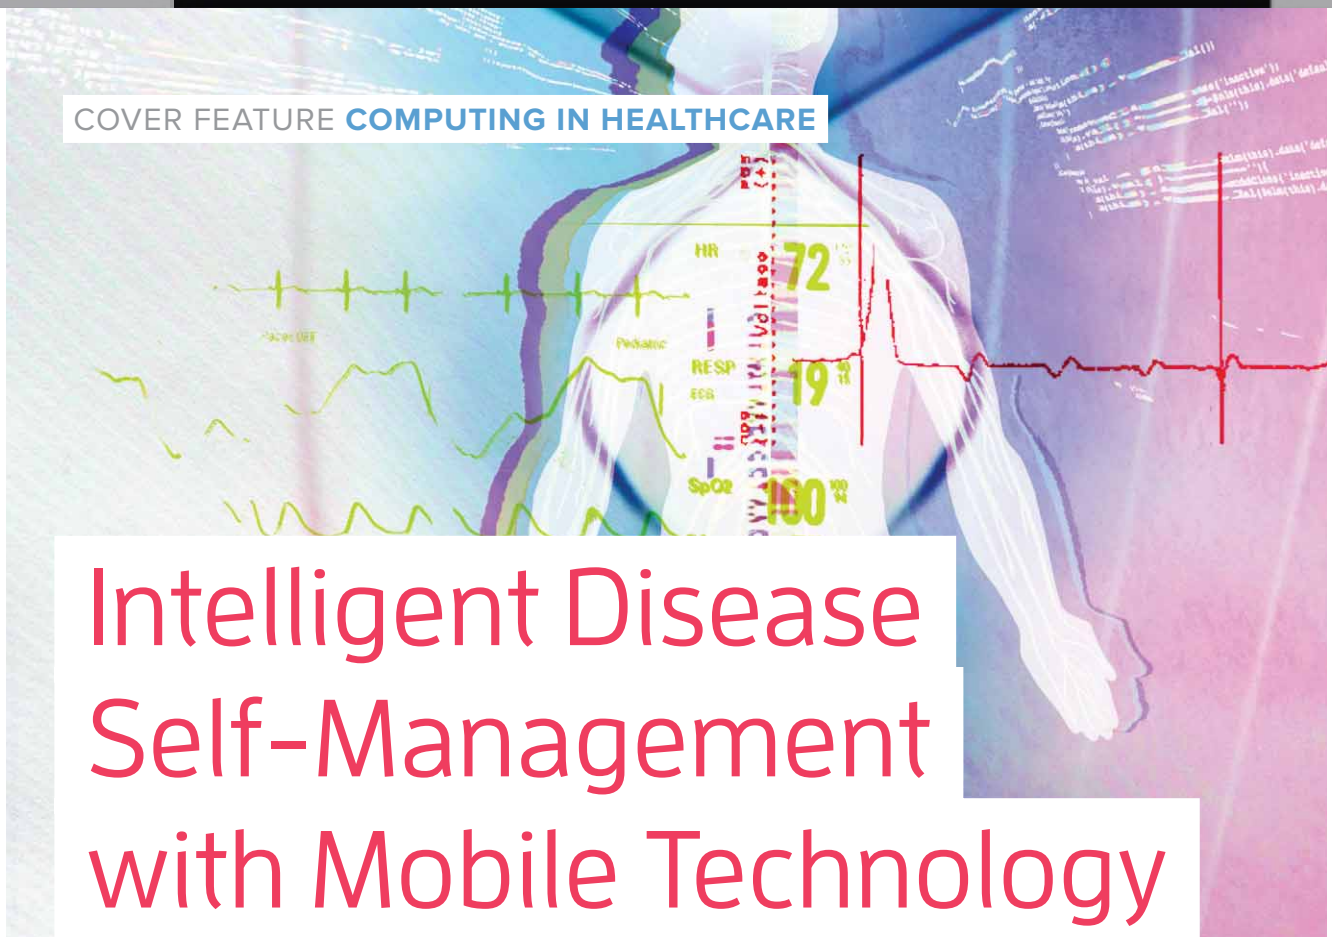

COVER FEATURE **COMPUTING IN HEALTHCARE**

Intelligent Disease Self-Management with Mobile Technology

Marina Velikova, Embedded Systems Innovation by TNO

Peter J.F. Lucas and Maarten van der Heijden, Radboud University, Nijmegen

Cost-effective mobile healthcare must consider not only technological performance but also the division of responsibilities between the patient and care provider, the context of the patient's condition, and ways to implement patient decision support and tailored interaction.

Mobile technology is a promising way to transfer aspects of clinical care support from the caregiver—physician, nurse, nurse practitioner, or physical therapist—to the patient, thus enabling disease self-management. It has almost limitless potential to inform and engage the patient in treatment decisions, to monitor the patient's condition, and to alert caregivers about any unexpected changes.[1]

However, continuous and active patient involvement in mobile healthcare (m-health) requires considerable insight into disease self-management as a process, which requires accommodating a shift in the patient's role. As Figure 1 shows, in traditional care, the patient is largely passive, but in an m-health system, the patient is actively engaged in decision making about treatment. Involvement to this degree requires sufficient knowledge about disease-related conditions and problem-solving skills that lead to behavioral changes and coping strategies. Active day-to-day self-management means regular monitoring and reporting of signs and symptoms, enhanced medication adherence, and appropriate responses to health changes. Inherent in all these requirements is a close patient–caregiver partnership.

Most research to date emphasizes the technological aspects of m-health and other forms of patient-centric healthcare[2,3] or focuses on specific m-health applications.[4]

0018-9162/15/$31.00 © 2015 IEEE

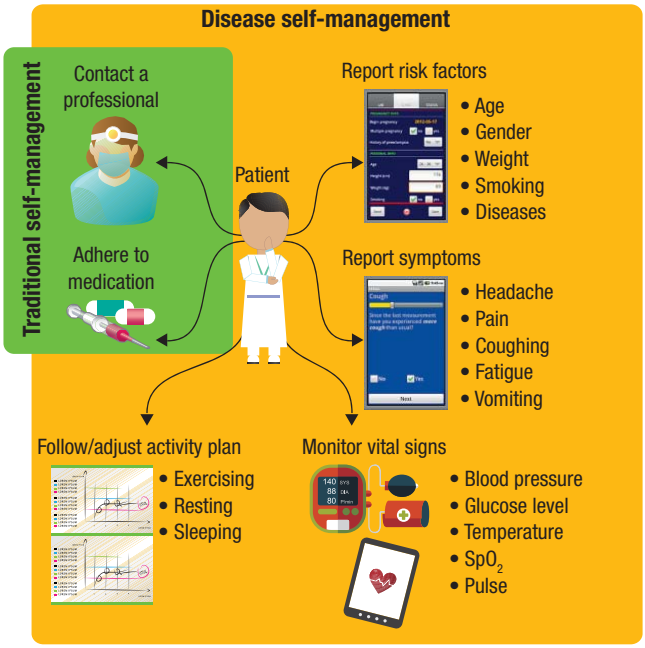

FIGURE 1. Traditional disease management versus disease self-management. In the traditional model (green), the patient passively follows a treatment plan, but in a mobile healthcare (m-health) system that supports disease self-management (yellow), the patient is actively involved in monitoring and decision making, and the caregiver receives continuous updates on the patient's condition. Communication is tailored to patient preferences, whether through voice, text, video, or email.

In contrast, we built two m-health systems—one for patients with chronic obstructive pulmonary disease (COPD)[5] and one for patients with pregnancy complications—primarily to explore how mobile technology can support the patient. Our focus was on embedding disease-specific Bayesian network models within a smartphone to enable on-the-fly patient data interpretation. In this way, the system could assess the patient's health status and advise the patient on what action to take—all without the care provider's direct involvement.

Experiments showed positive results for both m-health models. The COPD model, which we developed from clinical data in close cooperation with clinical experts, correctly detected 91 percent of COPD exacerbations from only patient-measured biosignals and reported symptoms. Tests of the pregnancy-related model, which we developed in cooperation with gynecologists,[6] on actual pregnancy data reliably predicted hypertensive complications for 60 percent of pregnant patients at least four weeks before they received an actual diagnosis.

As part of our work, we identified four foundational aspects of m-health for disease self-management—support for shared care, context awareness, embedded intelligence, and personalized interaction—and explored how best to integrate these in mobile technology.

A major challenge in disease self-management is determining the optimal complexity level for representing and interpreting context-specific clinical data to support patient-tailored interaction, decision making, and responses. We believe that an integrated approach like ours is the best way to evolve cost-effective m-health systems that can augment or even replace traditional disease management and thus lead to greater numbers of individuals who enjoy improved health at a lower cost. We also believe an integrated approach will stimulate the scientific, technological, and business development of mobile, patient-oriented decision-support systems.

SUPPORT FOR SHARED CARE

The success of intelligent disease self-management in m-health relies on how effectively the patient and caregiver share healthcare responsibilities. Systems with embedded intelligence must both automate patient data interpretation and feedback and support patient–provider interaction and decision making.

Data delivery speed is a critical shared-care service, since real-time data transmission is essential to monitoring the patient's condition. In a personalized m-health system, the patient could decide whether or not to transmit data and, if so, how much data and how often. If regularly transmitted data ceases unexpectedly, the caregiver might take steps to check if this was intentional and investigate the need for intervention.

Another service to support shared care is direct patient–caregiver communication through voice, video, text, or email, enabling a timely discussion of potential problems and possible treatment-plan adjustments. Patients would have the initiative to contact the caregiver, but the system would also advise the patient about the necessity of contact, thus lowering interaction degree and cost. Tailoring interaction in this manner will encourage both parties to partner in managing the disease.

An m-health system that supports shared care is likely to motivate patients to self-manage their disease and possibly achieve a better overall outcome than is possible with traditional treatment. Easier access to personal clinical information, for example, can increase the patient's desire to

explore evidence-based guidelines and relevant scientific studies to gain more knowledge about disease treatment.

CONTEXT AWARENESS

Because personalization is a strong benefit of m-health systems, mobile devices must be able to recognize, correctly interpret, and adapt to changes in patient, disease, and environmental contexts.

- *symptoms*, subjective experiences that the patient reports—for example, coughing, fatigue, and headache;
- *signs*, biomedical data based on observations from tests and measurement devices—for example, blood pressure via a sphygmomanometer, and glucose or hemoglobin levels from biochemical tests; and
- *biosignals*, repeatedly measured biological signals collected using noninvasive technologies—for example, signals from electrocardiography, which reflect physiological processes.

TO SUPPORT DISEASE SELF-MANAGEMENT, MOBILE DEVICES MUST EMBED DECISION AIDS WITH ENOUGH INFORMATION TO MAKE INTELLIGENT HEALTH DECISIONS.

Patient context

The patient's age, gender, and personal and family disease history, as well as lifestyle choices about diet, alcohol consumption, smoking, and activities, can affect the risk of contracting a disease or the progression of disease that the patient already has. A mobile device is a convenient way to gather such patient-specific information—for example, by having the patient fill out a questionnaire to gain information beyond what health records contain.

Personal and interpersonal factors such as prior health-related behavior, socioeconomic status, and social attitudes and support also influence the likelihood of health-promoting behavior and self-care. These factors determine the patient's physiological and psychological status, which in turn determine both current and future health. Health status is based on three data types:

Symptom reports are easy to collect through questionnaires on a mobile device, and measuring devices, such as body sensors, can collect signs and biosignals and wirelessly transfer them to a computing device. These measuring devices are compact, mobile, cheap, and easy to use, and measurements can be more representative than those obtained in a hospital or clinical environment. Remote measurements eliminate the white-coat effect—the patient's altered state from being in a clinic or hospital—as well as decrease costs, since the patient makes fewer visits to the hospital.

On the down side, ensuring that remote measurements are reliable might be more difficult, since the patient might not follow procedures or have the required skills. For example, an elderly patient with impaired renal function could find it difficult to quickly process the color analysis of a urine strip test, which could lead to the incorrect reporting of that measurement.

Disease context

Whether or not a mobile platform can manage a disease depends on the personal, societal, and economic burden that the disease imposes. Disease management can be short or long, depending on the disease type and severity.

Long-term disease management. Studies have shown that 75 to 85 percent of healthcare costs go to the management of chronic diseases, such as COPD, hypertension, and diabetes mellitus types 1 and 2.[7] Many chronic diseases are well understood and also preventable, since they are closely linked to the patient's lifestyle choices. For example, about 90 percent of COPD cases stem from the patient's smoking history. Treatment obviously begins with smoking cessation, which makes COPD an excellent candidate for judging the cost-effectiveness of m-health systems in dealing with long-term disease management linked to lifestyle choices.

Short-term disease management. Some diseases do not require indefinite management; for example, a few weeks of guided therapy are sufficient to treat many sports injuries. Pregnancy-related disorders, such as gestational hypertension and pre-eclampsia, require management only during pregnancy.

Disease severity. Disease severity is a concern primarily in long-term management. In COPD, for example,

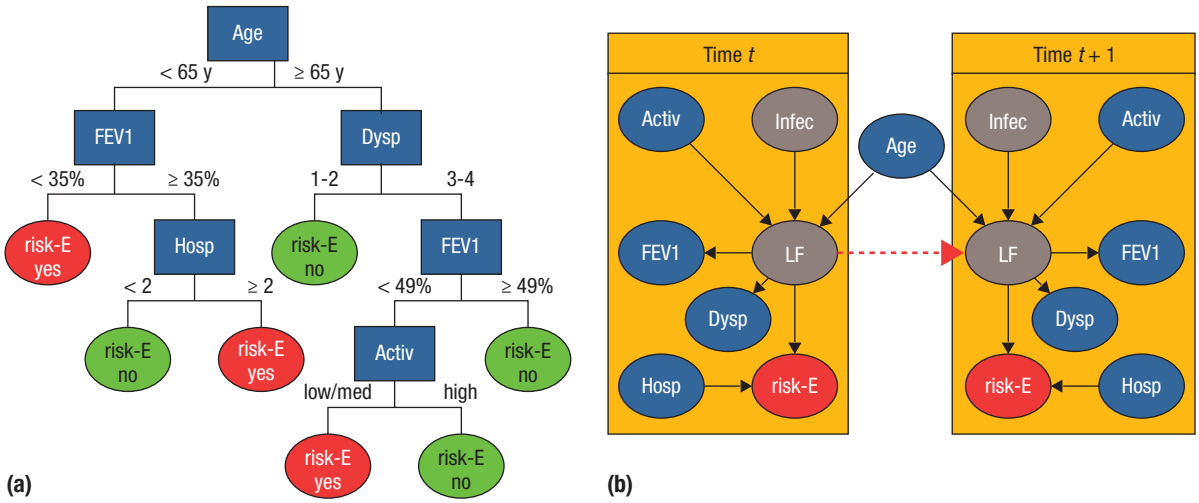

FIGURE 2. Two models for predicting the risk of chronic obstructive pulmonary disease (COPD) exacerbations: (a) decision tree and (b) temporal Bayesian network. The solid lines in (b) represent causal dependencies within a time slice; the dashed line represents dependencies among the organ's functioning states over time. Activ: activity; Dysp: presence of dyspnea; risk-E: risk for exacerbation; FEV1: forced expiratory volume in 1 second; Infec: presence of infection; Hosp: previous hospitalizations; LF: lung function.

stages run from mild to severe, and the stage determines the extent to which self-management is possible and desirable. M-health solutions support patients by recommending specific actions appropriate for the particular stage. Patients in the severe stage might often receive medication instructions, reminders, and prompts to contact the caregiver, while patients in the mild stage might receive alerts only occasionally to assess their health status by taking some measurement.

Environmental context

Current mobile devices have built-in sensors to collect data about their or the user's environment. Accelerometers and three-axis gyrometers recognize movement and orientation, while a microphone captures acoustic data. Recent research on these built-in sensors is identifying new applications, such as the use of a smartphone microphone to measure lung function[8] and a smartphone camera to measure the blood's oxygen saturation.[9]

Communication between mobile devices and environmental technologies such as smart home appliances and wearable sensors can further assist isolated patients in maintaining health and safety standards.

EMBEDDED INTELLIGENCE

Although collecting data on the patient and the patient's environment is now less of a technical challenge, the interpretation of that data in supporting disease self-management remains problematic. Interpretation must consider the clinical context of the patient's condition, which is challenging even for clinicians with extensive training. Medical decision-support systems—even those with AI-based algorithms—can assist in managing only one or two disorders. Reliably dealing with a wide disease spectrum is still not feasible because disease interaction is inherently highly complex. In this respect, not much progress has been made since the 1990s when researchers demonstrated that such broad-range systems were not accurate enough.[10]

Thus, to adequately support disease self-management, mobile devices must embed patient-operated decision aids so that patients have enough information about available options and predicted outcomes to make intelligent health decisions. The exact problem-solving capabilities of such decision aids will depend on

- the desired degree of decision-making complexity,
- the method of handling uncertainty in medical reasoning,
- the choice of patient health status modeling in a particular clinical context, and
- the method of learning personalized characteristics from patient-specific data.

Complexity tradeoffs

Simple decision aids are based mostly on if-then rules, often coded in a scripting language. The rules raise a warning when required data (clinical, location, and so on) is missing or when certain patient laboratory results exceed the normal range.

Personalizing disease management for chronically ill patients, who are not likely to conform to a population norm, will require a separate, highly complex rule set, which can rapidly become unwieldy. Decision trees offer a compact representation, such as the one in Figure 2a, that is feasible to implement on a mobile device.

Decision trees are a straightforward, intuitive way to enable personalized decisions, and they are easy to explain to nonprofessionals. However, they cannot identify the prognostic trends that are foundational to determining the patient's future treatment.

Handling uncertainty

The medical knowledge and data that informs decision making are inherently uncertain, implying that deterministic approaches such as decision trees are less suitable in handling medical reasoning than probabilistic methods, such as Bayesian networks and logistic regression.[11] Bayesian networks can tackle multiple cause–effect relationships, such as those among symptoms, signs, and diseases, and quantify uncertainty with probability distributions.[12,13] With their ability to handle both complexity and uncertainty, Bayesian networks are a promising choice for building clinical decision models in m-health systems.

Although the Bayesian model describes general relations among variables of interest, modelers can personalize predictions by entering patient-specific data and run what-if scenarios by entering virtual evidence. They can also use Bayesian updating to tune probability distributions to the patient's characteristics.

Bayesian networks are suitable both when no data is available—in which case, the modeler can use expert knowledge and information from literature—or when a great deal of data is available, in which case the modeler relies on machine learning. Finally, when available data consists of only several hundred patient cases, a typical goodness-of-fit statistical method, such as a logistic regression, is preferable to Bayesian networks in building prediction models.

Static versus dynamic health status capture

Intelligent m-health systems for disease self-management can describe and capture the patient's health status statically or dynamically. In static capture, the organ's functioning at a specific moment determines the outcome of laboratory tests and the presence or absence of symptoms. Static capture is generally sufficient for diagnosis.

Prognostic tasks, however, require a dynamic model that analyzes health states while accounting for temporal evolution and can thus make the predictions that disease self-management relies on to tailor treatment. A dynamic model must consider how patient characteristics such as preexisting diseases, age, and genetics affect organ functioning, as well as the use of therapeutic drugs to cure the disease or slow its progress.

Dynamic capture is important in disease self-management because it allows the patient and caregiver to take prompt remedial action. In Figure 2b, for example, the Bayesian network models interrelationships among a series of factors that influence COPD exacerbation during a particular time. Following clinical practice, reasoning compares the functioning of organ X at previous check-ups (1, ..., t-1) and at the present time t to determine if major functional changes have occurred. Given the history and current status, the system can predict the organ's functioning and the potential development of the disease for subsequent time slices (t + 1, ..., T).

> PROGNOSTIC TASKS REQUIRE A DYNAMIC MODEL THAT ANALYZES HEALTH STATES WHILE ACCOUNTING FOR TEMPORAL EVOLUTION.

Learning personalization

The ability of m-health systems to continuously collect patient and environmental data paves the way for extracting personalized data to tailor model structure and fine-tune statistical model parameters or update them to account for changes in patient behavior or condition. Recent examples include models based on decision trees or support-vector machine models that recognize patient activity and warn of potential danger or advise the patient to adopt healthier lifestyle choices.

By pooling data from many patients and applying discretization techniques, modelers can obtain meaningful discriminative features. For example, forced expiratory volume in 1 second (FEV1) measures airway obstruction from volume-time curves that a spirometer provides. FEV1's predictive value is expressed as a percentage, which modelers often discretize into ranges to define clinically meaningful values. In m-health systems, which typically have many sensors, certain sensor types can learn personalized features by age or gender, for example.

PERSONALIZED INTERACTION

Personalized interaction enables the patient to rely on a mobile device agent that knows the patient's preferences and can adjust its behavior accordingly. The more personalized

FIGURE 3. Android smartphone interface snapshots for m-health systems to monitor COPD exacerbation or pregnancy complications. (a) Both systems use an alert module to prompt patients to perform care tasks. (b) Questionnaire for the COPD system, which is relatively simple to accommodate COPD patients, who are typically older. Pregnant patients in contrast are usually younger and are comfortable with more elaborate displays, such as (c) clinical data, (d) measurement data, (e) current status and advice, (f) a prognostic chart, and (g) measurement analysis. Tailoring the interface to patient preferences is critical to the patient's acceptance of the mobile agent.

the patient–agent interaction becomes, the more motivated patients will be in integrating the agents into their daily lives. Interaction must account for the patient's preferences in mobile technology as well as the interface's adaptability.

Patient preferences

Different patient groups can have vastly different needs and preferences for mobile self-management, which underlines the importance of a tailored interaction approach. A user who often interacts with smartphone applications would probably not mind being interrupted by phone alerts, such as the one in Figure 3a, but another patient might dislike that intrusion. In Figure 3b, the interface for a questionnaire displays only one question at a time, which accommodates most patient preferences. The questionnaire module in the figure is part of the COPD-monitoring m-health system, but with different questions, it could easily work in other m-health systems.

Setting preferred days and times for reminders about taking measurements or medications or filling out questionnaires will give the patient more flexibility and encourage the patient to persevere in using the mobile agent instead of turning it off, which could result in missed measurements or medications.

Adaptive interface

The patient must be able to easily and intuitively interact with the mobile agent, so the complexity and form of displayed information should depend on the patient's education level, age, and interests. Interface adaptation ranges from a simple setting adjustment—for example, to raise the volume for elderly users—to tailoring the entire interface for a particular age group.

COPD patients tend to prefer simple, easily understood feedback and

uncluttered screens, so a visual depiction of their health status works better. Our experiments with the COPD-monitoring system confirmed these preferences.[5]

Younger patients, on the other hand, tend to be more familiar with modern technology and thus can tolerate a more complex interface and more functions. They also enjoy icons, such as a smiley face. Consequently, we gave patients using the pregnancy system the option of displaying a prognostic chart for disease development until the end of the pregnancy, which is based on the probabilities from the embedded Bayesian network model as well as detailed measurement analysis. Figures 3e through 3g show the respective interfaces. Overall, the interfaces in Figure 3 are based on generic display elements tailored to accommodate the preferences of a particular patient group. With different tailoring, the displays would be equally effective for m-health systems with other target patient groups.

The data-acquisition rate is also adjustable, depending on the risk to the patient. Low-risk patients can need fewer check-in times, so they can reduce the rate. If the patient's risk increases, the system will automatically increase the check-in rate so that it can detect any health deterioration in a timely manner and alert the patient and the caregiver to take the appropriate action.

Information and communications technologies are shaping the future of healthcare as a system in which patients will have more involvement in care-management choices. These technologies must provide not only continuous health monitoring but also personalized patient decision support and disease-management advice.

APPROPRIATE TRAINING ENHANCES THE WILLINGNESS OF CLINICIANS TO ADAPT THEIR PRACTICES TO MOBILE DECISION-SUPPORT TECHNOLOGIES.

It is now much easier to run large models on mobile devices,[14] but healthcare is complex and requires solutions that move beyond the technical. Decision support must be clinically valid as well, since patients using an m-health system are often outside a controlled clinical environment. Appropriate training programs can enhance the willingness of health professionals to adapt their clinical practices to mobile decision-support technologies. This nexus of technical, clinical, and psychological elements underlines the importance of multidisciplinary cooperation among clinicians, computer scientists, engineers, and patients early in system development. Only then can prototype systems transition to practical daily care products.

ACKNOWLEDGMENTS

We thank the anonymous reviewers for their constructive comments. This article is based on work supported in part by The Netherlands Organisation for Health Research and Development (ZonMW), the Technology Foundation STW, and the Foundation IT Projects (STITPRO).

REFERENCES

1. Global Observatory for eHealth (GOe), *M-Health: New Horizons for Health through Mobile Technologies*, tech. report, Worldwide Health Org., 2011.
2. S. Kumar et al., "Mobile Health: Revolutionizing Healthcare through Transdisciplinary Research," *Computer*, vol. 46, no. 1, 2013, pp. 28–35.
3. H. Viswanathan, B. Chen, and D. Pompili, "Research Challenges in Computation, Communication, and Context Awareness for Ubiquitous Healthcare," *IEEE Comm.*, vol. 50, no. 5, 2012, pp. 92–99.
4. A. Triantafyllidis et al., "A Pervasive Health System Integrating Patient Monitoring, Status Logging, and Social Sharing," *IEEE J. Biomedical and Health Informatics*, vol. 17, no. 1, 2013, pp. 30–37.
5. M. van der Heijden et al., "An Autonomous Mobile System for the Management of COPD," *J. Biomedical Informatics*, vol. 46, no. 3, 2013, pp. 458–469.
6. M. Velikova et al., "Exploiting Causal Functional Relationships in Bayesian Network Modeling for Personalized Healthcare," *Int'l J. Approximate Reasoning*, vol. 55, no. 1, 2014, pp. 59–73.
7. Continua Health Alliance, *The Next Generation of Healthcare: Personal Connected Healthcare*, tech. report, 2010; www.continuaalliance.org/sites/default/files/Continua%20Overview%20Presentation%20021513.pdf.

8. E. Larson et al., "SpiroSmart: Using a Microphone to Measure Lung Function on a Mobile Phone," *Proc. 14th ACM Int'l Conf. Ubiquitous Computing* (Ubicomp 12), 2012, pp. 280–289.
9. C. Scully et al., "Physiological Parameter Monitoring from Optical Recordings with a Mobile Phone. *IEEE Trans. Biomedical Eng.*, vol. 59, no. 2, 2012, pp. 303–306.
10. E.S. Berner et al., "Performance of Four Computer-Based Diagnostic Systems," *New England J. Medicine*, vol. 330, no. 25, 1994, pp. 1792–1796.
11. P. Szolovits, "Uncertainty and Decisions in Medical Informatics," *Methods Information in Medicine*, vol. 34, nos. 1–2, 1995, pp. 111–121.
12. J. Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, Morgan Kaufmann, 1988.
13. D. Koller and N. Friedman, *Probabilistic Graphical Models: Principles and Techniques*, MIT Press, 2009.
14. S. Evers and P.J.F. Lucas, "A Framework for Development, Teaching, and Deployment of Inference Algorithms," *Proc. 6th European Workshop Probabilistic Graphical Models* (PGM 12), 2012, pp. 99–106.

ABOUT THE AUTHORS

MARINA VELIKOVA is a research fellow in computer science with Embedded Systems Innovation by TNO, The Netherlands. Her research interests include knowledge representation, decision-support systems, and intelligent data analysis. Velikova received a PhD in operations research from Tilburg University. She is a member of IEEE and ACM, and serves on the editorial board of *BMC Medical Informatics and Decision Making*. Contact her at marina.velikova@tno.nl.

PETER J.F. LUCAS is a principal investigator with the Institute of Computing and Information Sciences at Radboud University, Nijmegen, and a professor at the Leiden Institute of Advanced Computer Science, Leiden University, The Netherlands. His research interests include probabilistic logics, probabilistic graphical models, decision-support systems, and m-health solutions. Lucas received an MD from Leiden University and a PhD in mathematics and computer science from Free University, Amsterdam. Contact him at peterl@cs.ru.nl.

MAARTEN VAN DER HEIJDEN is a postdoctoral researcher in computer science with the Institute of Computing and Information Sciences at Radboud University, Nijmegen. His research interests include probabilistic graphical models and m-health solutions. Van der Heijden received a PhD in artificial intelligence from Radboud University, Nijmegen. Contact him at m.vanderheijden@cs.ru.nl.

COVER FEATURE **COMPUTING IN HEALTHCARE**

Medical-Grade Quality of Service for Real-Time Mobile Healthcare

Kyungtae Kang, Hanyang University
Qixin Wang, Hong Kong Polytechnic University
Junbeom Hur, Chung-Ang University
Kyung-Joon Park, Daegu Gyeongbuk Institute of Science and Technology
Lui Sha, University of Illinois at Urbana–Champaign

A wireless electrocardiogram case study suggests that current CDMA2000 cellular technology has considerable potential in medical telemetry. Modifications to the network protocol stack ensure the highest data integrity and lowest service delay.

Wireless communications are rapidly becoming an essential component of modern healthcare, enabling patients to enjoy a greater level of mobility. The compelling benefits of patient freedom from wired medical equipment and the bureaucracy associated with medical treatment are major motivators for the recent explosion of mobile healthcare (m-health) systems and applications.[1–3] Because m-health heavily depends on collaborative interaction among mobile medical devices wirelessly connected to back-end clinical systems, the key challenge in m-health is how to achieve medical-grade quality of service (QoS)—a level of transmission speed, reliability, privacy, and security that provides real-time, confidential, and accurate service. A wireless system that can deliver medical-grade QoS must have broad coverage, a low error rate, and a low upper limit on service latency.

Several proposed systems have advocated both IEEE 802.11 wireless LANs (WLANs) and cellular technology as platforms for medical telemetry applications.[4–6] WLANs are cost-effective and are the platform of choice in many hospitals. Cellular networks are inherently more expensive, but they offer broader coverage, are a proven infrastructure to support high mobility, and use licensed frequency bands to avoid interference from other networks. Moreover, implementation on an existing cellular infrastructure can greatly reduce cost, thus mitigating the cellular platform's main disadvantage. These features make cellular technology a strong candidate for medical telemetry and other m-health applications.[4,5]

To explore the practical concerns associated with selecting a cellular platform, we evaluated the QoS requirements of some key medical applications and then developed a wireless system architecture based on CDMA2000 1xEV-DO[7] cellular technology, which uses code-division multiple access (CDMA) and time-division multiple access (TDMA) multiplexing to maximize throughput. Evolution-Data Optimized (EV-DO) is shorthand for a version of EV, a telecommunications standard for the wireless transmission of data through

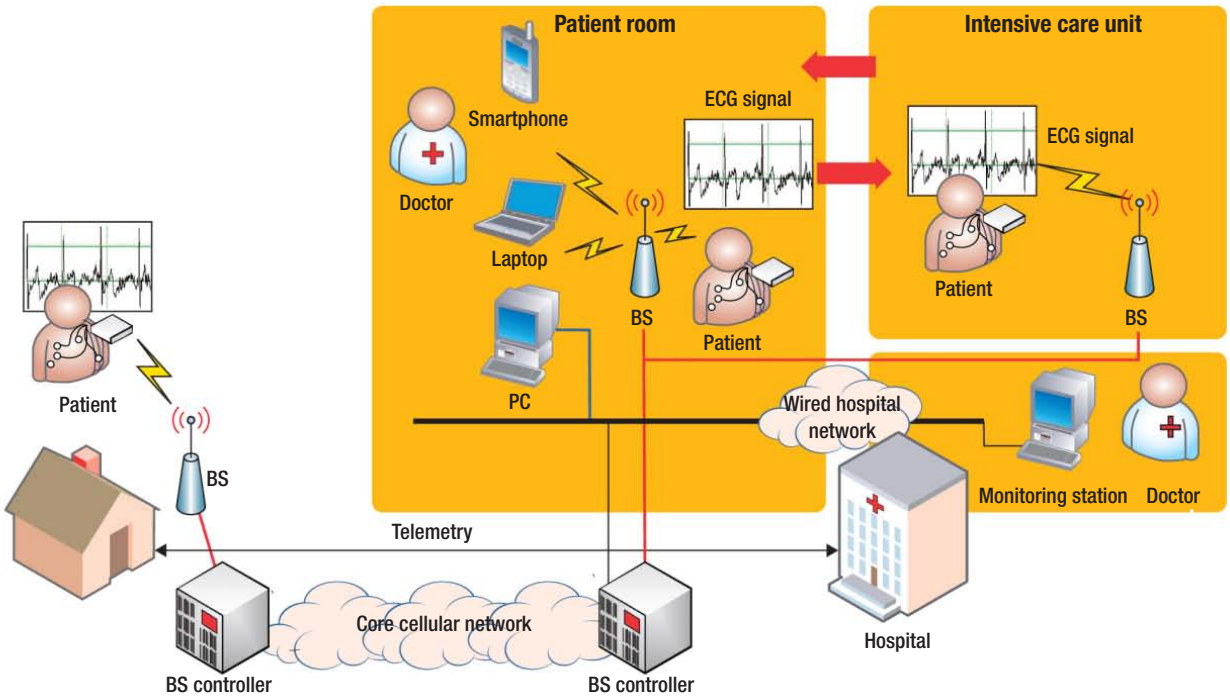

FIGURE 1. Continuous wireless electrocardiogram (ECG) service using cellular technology. Short-range wireless communication technologies such as Bluetooth receive signals from on-body sensors that measure the voltage difference between electrodes attached to the outer surface of the patient's skin. The signals then go to the patient's mobile device, which forwards data to the base station (BS) via an intermediate base station controller. The BS then sends the data to the appropriate hospital monitoring station in the patient's room or intensive care unit.

radio signals, typically for broadband Internet access.

To gauge the effectiveness of our architecture and to estimate key QoS metrics, we simulated continuous wireless electrocardiogram (ECG) monitoring across layers of our modified network protocol stack. In this focus, our work differs from recent studies of QoS in cellular-based m-health systems, which concentrate more on high-level system design than on network protocols.[2,6]

To test our proposed architecture, we chose an application that pushes the QoS envelope. As its name implies, continuous real-time ECG monitoring requires continuous real-time data transmission. The dropout rate can be at most a few seconds per hour; the packet error rate must be less than 0.1 percent, and service latency must be under 2 s. In addition, security and privacy must be at acceptable levels.[8,9]

Although continuous ECG monitoring requires a higher data rate than many m-health applications,[4] it is consistent with these other applications in its need for reliable transmission and low service latency,[8,9] since lost or corrupted data or late delivery obviously compromises any m-health application's effectiveness. Consequently, packet delivery ratio, security, service latency, and jitter are crucial QoS parameters not just for continuous wireless ECG monitoring but for m-health applications overall.

Regardless of the particular application, our study shows the need to examine the tradeoffs in implementing cellular networks in this domain. On the one hand are higher patient mobility and lower deployment cost; on the other are possible reductions in reliability and security along with increased transmission delay. Understanding the optimal tradeoff requires careful analysis.

REQUIREMENTS OF CONTINUOUS ECG MONITORING

In a continuous wireless ECG monitoring system, patients move around as they would normally, and clinicians have immediate access to their ECG data. In this sense, continuous monitoring differs from systems that collect data but store it for later downloads or systems that transmit data only if an event, such as arrhythmia or cardiac arrest, occurs.

Monitoring with a cellular network

Figure 1 shows a possible continuous wireless ECG scenario that uses a cellular system.

The system must have a way to sample, digitize, and group the analog ECG signals into packets. The chosen sampling frequency and digitization method determine traffic characteristics during transmission. The data rate depends on the number of sensors, bits, and sampling frequency. If L sensors sample the electrical signals from a patient's heart, and the system digitizes each sensor's output at r samples per second with a resolution of S bits, the resulting data rate is (LrS) bits

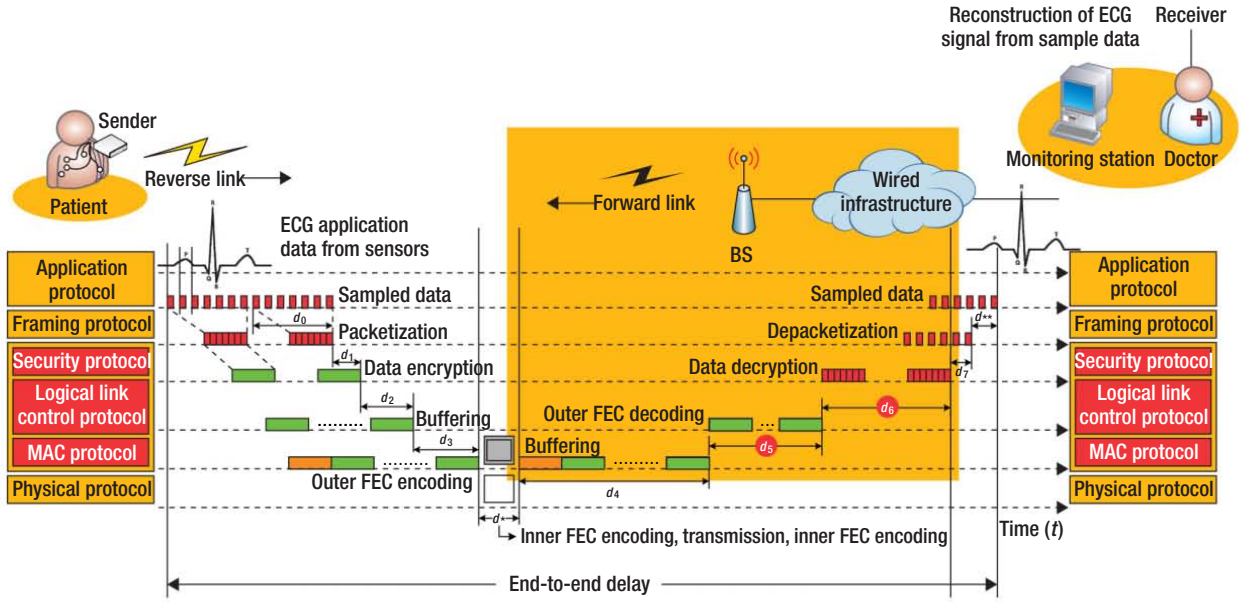

FIGURE 2. Modified protocol stack and delay sources. The stack is based on the use of the CDMA2000 1xEV-DO (code-division multiple access, Evolution-Data Optimized) cellular technology in continuous ECG monitoring. Delay sources (d_n) stem from functions across layers. Delays in black circles vary depending on the conditions of wireless channel. FEC: forward error correction; MAC: media access control.

per second. Thus, 12 sensors providing 500 samples per second at a resolution of 16 bits yields a data rate of 96 Kbps—a typical value for cellular ECG monitoring applications. This example illustrates the significance of a target data rate in tailoring cellular infrastructure to an ECG application.

Providing medical-grade QoS

The stringent QoS requirements of continuous wireless ECG monitoring require modifications to the network protocol stack, particularly in the data-link and physical layers.

Data-link layer. Link control provides reliable and secure communication through three main functions.

A *security function* authenticates devices and networks, ensures that data is not modified in transit, and encrypts transmitted data to ensure privacy.

Logical link control compensates for the bursty error processes inherent in a wireless channel. Error control can be based on retransmission, in which the system retransmits only errors and adapts easily to changing channel conditions. However, data delivery time can be unpredictable, which is unacceptable in m-health applications. The obvious alternative is forward error correction (FEC), which can both maintain throughput and control delay.

The third function in the data link layer is *media access control* (MAC), which determines when devices can access the communication medium. Either CDMA or TDMA techniques will satisfy an m-health application's timing requirements, offering a more predictable delay than the random-access or contention-based methods that IEEE 802.11 networks currently use.

Physical layer. Channel coding and digital signal modulation in the physical layer aim to optimize throughput and reliability. Channel coding introduces redundancy to cope with bits that are flipped because of noise or interference. FEC techniques such as turbo and convolutional coding typically provide this redundancy. Turbo coding is more effective when physical layer packets are several hundred bits or more.

The tradeoff between transmission bit rate and reliability drives the choice of modulation type and FEC code rate: a more robust modulation type that can tolerate higher levels of interference has a lower transmission rate. A higher FEC code rate increases redundancy, allowing the system to tolerate higher interference levels, which improves transmission reliability but erodes the effective bit rate. Generally, m-health applications fit better with modulation and coding schemes that favor reliable transmission rather than a high data rate, although adaptive modulation and coding is suitable if the system architect knows the precise tradeoffs required.

WIRELESS ECG WITH CDMA2000 TECHNOLOGY

With QoS requirements in hand, we looked at how to adapt the protocol stack for continuous ECG monitoring in the CDMA2000 1xEV-DO cellular system.

Protocol layers

Figure 2 shows our proposed protocol stack and the interactions among layers. This arrangement is based on the standard air specification for CDMA2000 multicast services.[10] The framing protocol, which packs higher layer data into frames, helps determine higher layer data boundaries.

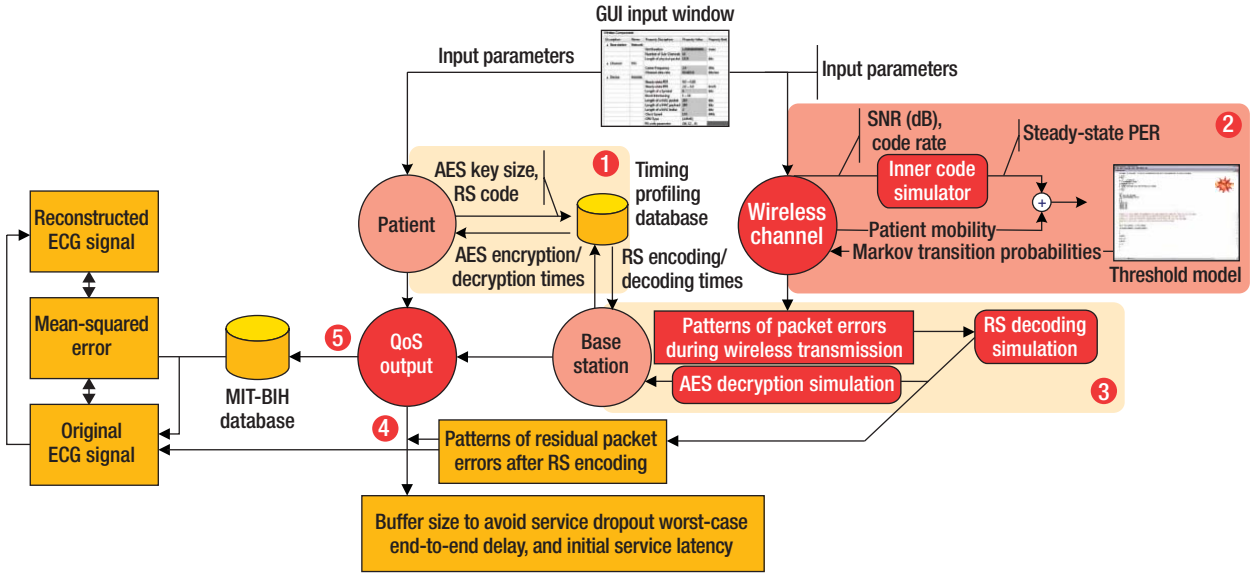

FIGURE 3. Simulation of a wireless continuous ECG monitoring system. (1) The system obtains the delays from Reed-Solomon (RS) coding and Advanced Encryption Standard (AES) encryption/decryption and (2) derives packet errors during wireless transmission using a threshold mode. The base station then (3) simulates RS decoding and AES decryption. Finally, the system (4) analyzes quality of service (QoS) metrics and uses the results to (5) estimate the mean-squared error. PER: packet error rate; SNR: signal-to-noise ratio.

However, the lower layers require significant modifications.

Security and logical-link layers. In the security layer, data encryption uses the Advanced Encryption Standard (AES),[10] which operates on 128-bit blocks using a 128-bit key. In the logical-link layer, the outer FEC code combines with the inner FEC (turbo) code in the physical layer to form an effective product code. Our stack uses Reed-Solomon (RS) outer codes because of their superior performance at low error rates—ideal for continuous ECG monitoring.

We specify an RS code by (N,K), which captures the idea that the encoder takes K data symbols and adds $(N - K)$ parity symbols to create an N-symbol codeword. A decoder can correct up to t symbol errors (in an unknown location), or up to $2t$ symbol erasures (errors in known locations), where $2t = N - K$. The location of erased symbols is often available in a digital communication system because the demodulator commonly flags incoming symbols that are likely to contain errors.

Because error bursts can defeat FEC codes designed for random errors, we added interleaving,[11] a simple process that spreads error bursts over time, allowing FEC to correct them and thus enabling RS codes to deal with error bursts.

To handle error bursts at the sender side, the system inserts security layer data into a buffer memory, or on *error control block* (ECB), a table of N rows and M columns, in a left-to-right, top-down fashion, and applies RS encoding along the columns. A reverse process of decoding and unpacking takes place at the receiver side. ECB width determines the length of error burst that interleaving can correct. Increasing the M value can correct longer error bursts but at the cost of a larger buffer and increased latency.[11]

MAC layer. The CDMA2000 MAC protocol specifies the procedures to transmit on both forward and reverse traffic channels. The reverse link uses CDMA techniques and Walsh codes to transmit data and employs a long pseudonoise sequence to identify mobile devices. The forward link uses both CDMA and TDMA techniques to transmit data.[7]

Physical layer. A CDMA2000 1xEV-DO Revision A system supports reverse-link data rates from 4.8 to 1,843.2 Kbps. The data rate is based on payload, modulation type, and how much redundancy turbo coding introduces.

Delay sources

As Figure 2 shows, an ECG application running on a CDMA2000 network has multiple sources of end-to-end delays. The accumulation delay in filling a frame (d_0) is 10.4 ms with a 1,000-bit frame and an ECG data rate of 96 Kbps. Encryption introduces an additional delay (d_1). Given this frame size, an RS code of (16,12), and two security layer packets in each ECB row ($M = 2$), writing to the ECB causes a 250 ms delay ($12 \times 2 \times 1{,}000/96$) ($d_2$).

Entering the security-layer packets into the ECB row by row and performing outer FEC encoding along the columns incurs an additional delay (d_3). Naturally, further coding and propagation delays (d^*) occur in the physical layer's hardware, including the inner FEC encoder and decoder, but these delays are small enough to be insignificant in our analysis.

The base station performs inner FEC decoding and buffers packets into the ECB, which fills at a fixed rate because CDMA2000 uses a TDMA technique to access the medium. Thus, the buffering delay (d_4) is constant. Once the ECB is full, the base station performs outer FEC decoding and decrypts security-layer packets, and the monitoring station depacketizes them to obtain the original ECG data. These actions incur three additional delays (d_5, d_6, and d_7). (We ignored the application processing delay (d^{**}), since it is unrelated to the wireless system.)

All these delays are fixed, except those from RS decoding (d_5) and AES decryption (d_6), which vary with the physical channel's condition. Implementing these processes in software greatly reduces this delay variation. In such an implementation, the system would require a buffer to absorb the delay variation, or jitter, which has a maximum amplitude of

$$2 \times \left(\overline{d_5 + d_6} - \underline{d_5 + d_6}\right),$$

where

$$\overline{d_5 + d_6}$$

is the worst-case time for RS decoding and the AES decryption of an ECB and

$$\underline{d_5 + d_6}$$

is the best-case time. The jitter buffer turns this maximum jitter into a fixed delay of the same magnitude (d_7).

QOS ANALYSIS

Optimizing a network for wireless ECG requires considering reliability and latency tradeoffs within a layer and between layers. To analyze these tradeoffs, we conducted a cross-layer simulation of a continuous ECG monitoring application in Java. Figure 3 shows the simulation process in which we varied channel and network parameters to obtain QoS metrics.

Our simulation uses software implementations of the RS erasure encoder and decoder (http://rscode.sourceforge.net) and Gladman's reference implementation of the AES algorithm (www.gladman.me.uk). We assume that the patient's mobile device has a low-power ARM9E processor core, but we consider a base station equipped with a more powerful ARM11 core as a possible alternative. We chose ARM processors because many 3G phones for the GSM and CDMA2000 networks use these processors.

We also chose the IAR embedded workbench for ARM (www.iar.com), which enabled us to profile execution times of AES encryption and decryption in the security layer and of RS encoding and decoding in the logical-link layer.

TABLE 1. System parameters for wireless ECG.

Parameter	Value
Network	
Inner turbo code rate (r)	1/4
Physical-layer packet length	1,024 bits
Sample outer RS code (N,K)	(16,12)
Symbol size in RS erasure codes (s)	8 bits
Block interleaving level (M)	1–8
Security layer packet length	1,000 bits
MAC-layer packet length	1,002 bits
AES cipher block unit size	128 bits
Cipher key size	128 bits
Physical channel	
Carrier frequency	1.8 GHz
Modulation	Phase-shift keying
Reference channel data rate	153.6 Kbps
Worst channel signal-to-noise ratio	–2.77 dB
Estimated wireless channel mobile velocity	2–4K mph
Wireless ECG	
Number of sensors	12
Sampling frequency per ECG sensor per second	500 Hz
Sample size	16 bits
Data rate	96 Kbps
Medical devices	
Mobile patient device	ARM9TDMI at 250 MHz
Base station (modem proc.)	ARM11 at 400 MHz

MAC: media access control; RS: Reed-Solomon; AES: Advanced Encryption Standard

Simulation parameters

Table 1 lists the parameters in our simulation, which we took from CDMA2000 1xEV-DO Revision A.[7,10] Because the mobile device has limited RF power, it requires coding rates and modulation techniques that are less affected by poor channel conditions. Therefore we chose a reverse-channel data rate of 153.6 Kbps and the (16,12) RS code, which offers a particularly good tradeoff between coding gain and processing delay.[10]

Process

The simulation first obtains the delay caused by AES encryption (d_1) and RS encoding (d_2) in the patient's

TABLE 2. Quality-of-service metrics from simulation: delays and latency (ms) and jitter buffer and ECB sizes (bits) .

M	$\overline{PER_r}$	d_0	d_1	d_2, d_4	d_3	$\overline{d_5 + d_6}$	$\underline{d_5 + d_6}$	d_7	Jitter buffer size	Worst-case service latency	ECB size
1	3.3E-3	10.4	7.04	125	10.9	44.91	27.77	34.29	3,291	347.14	16K
2	4.8E-4	10.4	7.04	250	21.8	89.82	55.42	68.81	6,606	687.48	32K
3	1.4E-4	10.4	7.04	375	32.7	131.99	83.10	97.78	9,386	1,019.51	48K
4	4.1E-5	10.4	7.04	500	43.6	174.15	110.81	126.68	12,161	1,351.47	64K
5	3.8E-5	10.4	7.04	625	54.5	214.93	138.55	152.78	14,666	1,679.25	80K
6	2.7E-5	10.4	7.04	750	65.4	257.10	166.25	181.70	17,443	2,011.24	96K
7	2.5E-5	10.4	7.04	875	76.3	297.88	193.93	207.92	19,960	2,339.14	112K
8	2.4E-5	10.4	7.04	1,000	87.2	337.30	221.64	231.32	22,207	2,662.86	128K

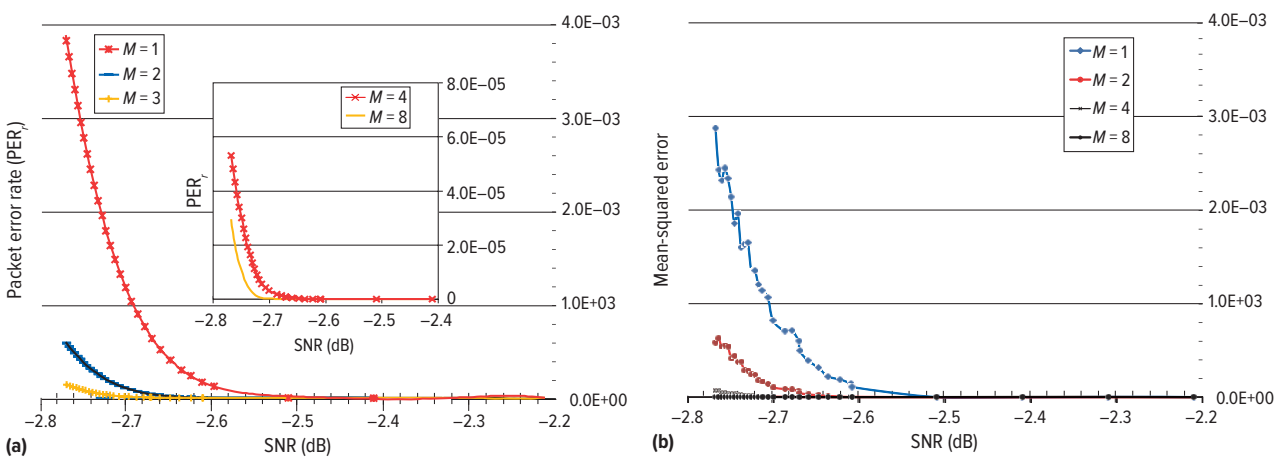

FIGURE 4. Average mean-squared error of reconstructed ECG signals. As M increases, error rate lessens only marginally, but latency grows, which suggests that beyond $M = 4$, performance gains are minimal.

(sending) device. It then uses the coded modulation library (www.iterativesolutions.com) to simulate inner turbo coding in the physical layer, which produces the channel's packet error rate (PER) with a given signal-to-noise ratio. This rate in combination with the patient's movement speed is the basis for simulating the occurrence of bursty packet errors in the wireless channel model.[12]

Having a model of the packet-error process in the physical layer enables the simulation of base station events. Simulating RS decoding in the logical-link layer and AES decryption in the security layer yields the pattern and number of residual errors in each ECB, together with the delays from RS decoding and AES decryption (d_4 and d_5).

Results go to an output control module that injects the residual errors into the MIT-BIH arrhythmia database,[13] which contains 48 half-hour excerpts from various ECG recordings. The control module then estimates the mean-squared error, which quantifies the difference between the ECG signals reconstructed at the remote monitoring station and the original signals from the patient; the required ECB and jitter buffer sizes; and the service latency.

SIMULATION RESULTS

Our simulation revealed some interesting patterns. For example, it showed that the delay from the AES decryption of a single security layer packet (d_1) will be 7.04 ms, while packetizing ECG data will create a delay (d_0) of 10.4 ms. Consequently, by the time the system packetizes a frame, the previous frame's encryption will already be complete, which means that AES encryption in the security layer will delay buffering by only 7.04

ms. For a particular ECB size, only the ECG data sampling rate will determine the buffering delays (d_2 and d_4).

Packet errors

Table 2, column 2, shows the upper bound on the residual packet error rate

$$\left(\overline{PER_r}\right)$$

in the data-link layer for different values of M, the block interleaving parameter. It also shows the worst- and best-case execution times for RS decoding and AES decryption and the delay corresponding to maximum resulting jitter. In the last columns are the buffer size to cope with this jitter and the corresponding service latency.

We also discovered an interesting pattern relating M values, the error rate, and latency. Although a higher value of M broadens the time scope of error bursts that FEC can handle, that flexibility comes with memory and buffering delays. However, our simulation showed that error-correction performance starts to saturate as M increases, until almost all the errors are corrected. Thus, as Figure 4 shows, increasing M further lowers the error rate only minimally, yet latency continues to increase, as Table 2 shows.

In our ECG application, setting M to 4 or 5 provided a reliability of 99.99 percent, an average mean-squared error of below 7.7E–7, and latency below 2 s. We assumed that the hospital had enough base stations to keep the signal-to-noise ratio above –2.77 dB.

Energy efficiency

Wireless technology for ECG monitoring must not only accurately deliver ECG data on time, but it must also manage energy in a power-constrained device. AES encryption and RS encoding in the data-link layer consume the most energy in our proposed wireless system.

We used an XEEMU simulator[14] to measure the average energy consumption of the reference software implementation of the AES cipher and RS encoder. XEEMU provides accurate energy data for the ADI 80200EVB XScale board, which uses the ARM9E instruction set.

We assumed that the core of our target processor for AES encryption and RS encoding (patient device) operates at 250 MHz; has separate instruction and data caches, both 32 Kbytes; and has a Micron 128-Mbyte SDRAM with a clock speed of 100 MHz (www.micron.com).

Table 3 shows the energy required to encrypt 128-bit data block and encode a codeword for common RS codes. Figure 5 shows energy consumption as the amount of energy in the data-link layer to handle a 1-byte data payload.

Table 3 shows that more energy is used to handle the data payload of one byte in the data-link layer as the number of parity symbols increases to improve the RS code's error-correcting capability. We expected this result, since

TABLE 3. Energy required to encrypt a 128-bit data block and encode an RS codeword.

Encryption (μJ)		CPU	MEM	Total
		35.412	**22.925**	**58.337**
Encryption (μJ)	(16,12)	15.130	9.371	24.501
	(16,13)	13.190	8.242	21.432
	(16,14)	11.090	7.143	18.233
	(32,24)	43.276	24.235	67.511
	(32,26)	36.768	20.911	57.679
	(32,28)	28.226	16.314	44.54

CPU: central processing unit, MEM: main memory

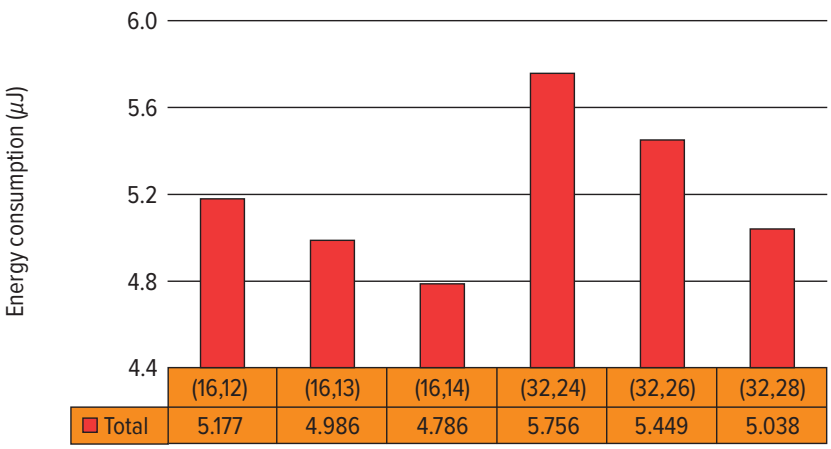

FIGURE 5. Energy consumption in the data link layer of the proposed cellular system for continuous ECG monitoring. Consumption reflects the energy required to handle a 1-byte data payload. Bars represent values with different RS code pairs.

ABOUT THE AUTHORS

KYUNGTAE KANG is an assistant professor in the Department of Computer Science and Engineering at Hanyang University, Korea. His research interests include operating, mobile, distributed, and cyber-physical systems. Kang received a PhD in electrical engineering and computer science from Seoul National University. He is a member of IEEE and ACM. Contact him at ktkang@hanyang.ac.kr.

QIXIN WANG is an assistant professor in the Department of Computing at Hong Kong Polytechnic University. His research interests include cyber-physical systems, real-time and embedded systems, real-time networking, wireless technology, and applications of these technologies in industrial control, medicine, and assisted living. Wang received a PhD in computer science from the University of Illinois at Urbana–Champaign. He is a member of IEEE and ACM. Contact him at csqwang@comp.polyu.edu.hk.

JUNBEOM HUR is an assistant professor in the School of Computer Science and Engineering at Chung-Ang University, Korea. His research interests include information security, mobile computing, and wireless network security. Hur received a PhD in computer science from the Korea Advanced Institute of Science and Technology. He is a member of IEEE. Contact him at jbhur@cau.ac.kr.

KYUNG-JOON PARK is an associate professor in the Department of Information and Communication Engineering at Daegu Gyeongbuk Institute of Science and Technology, Korea. His research interests include modeling and analysis of cyber-physical systems and design of medical-grade protocols for wireless healthcare systems. Park received a PhD in electrical engineering and computer science from Seoul National University. He is a member of IEEE. Contact him at kjp@dgist.ac.kr.

LUI SHA is the Donald B. Gillies Chair Professor of computer science at the University of Illinois at Urbana–Champaign. His research interests include cyber-physical systems, formalized reduced complexity architecture patterns, distributed real-time fault-tolerant computing systems, and dynamic real-time architecture. Sha received a PhD in electrical and computer engineering from Carnegie Mellon University. He is a fellow of IEEE and ACM. Contact him at lrs@illinois.edu.

the complexity of computing the syndromes and erasure evaluator polynomials increases with the amount of parity information. For example, the RS code (16,12) requires 36.4 percent more energy for RS encoding relative to the (16,14) RS code. The implication is that an ECG application needs an RS code that saves significant energy while still guaranteeing the required QoS level. This tradeoff depends on the channel status that mobile medical devices might be expected to experience and the required battery life.

The results of simulating our sample cellular-based wireless ECG monitoring system suggest that current CDMA2000 cellular technology has considerable potential in real-time medical telemetry and m-health applications in general. Consequently, we expect to see more advanced 4G technologies that will support the emerging growth of mobile, personalized medical care.

REFERENCES

1. R.S.H. Istepanian, E. Jovanov, and Y.T. Zhang, "Beyond Seamless Mobility and Global Wireless Healthcare Connectivity," *IEEE Trans. Information Technology in Biomedicine*, vol. 8, no. 4, 2004, pp. 405–414.
2. R.S.H. Istepanian, S.P. Costantinos, and S. Laxminarayan, "Ubiquitous M-Health Systems and the Convergence Towards 4G Mobile Technologies," *M-Health: Emerging Mobile Health Systems*, R.S.H. Istepanian, S. Laxminarayan, and C.S. Pattichis, eds., Springer, 2006, pp. 3–14.
3. Y.M. Fang, "Wireless Healthcare: Technologies for Bettering Our Life," *IEEE Wireless Comm.*, vol. 17, no. 1, 2010, pp. 2–3.
4. S.D. Baker and D.H. Hoglund, "Medical-Grade, Mission-Critical Wireless Networks," *IEEE Eng. in Medicine and Biology*, vol. 27, no. 2, 2008, pp. 86–95.
5. Q. Wang et al., "Building Robust Wireless LAN for Industrial Control with the DSSS-CDMA Cell Phone Network Paradigm," *IEEE Trans. Mobile Computing*, vol. 6, no. 6, 2007, pp. 706–719.
6. D. Vouyioukas, I. Maglogiannis, and D. Komnakos, "Emergency M-Health Services Through High-Speed 3G Systems: Simulation and Performance Evaluation," *Simulation*, vol. 83, no. 4, 2007, pp. 329–345.
7. N. Bhushan et al., "CDMA2000 1xEV-DO Revision A: A Physical Layer and MAC Layer Overview," *IEEE Comm.*, vol. 44, no. 2, 2006, pp. 37–49.
8. *IEEE Std. 11073: Health Informatics—PoC Medical Device Communication, Part*

00101: Guidelines for the Use of RF Wireless Technology, Dec. 2008.
9. L. Skorin-Kapov and M. Matijasevic, "Analysis of QoS Requirements for E-Health Services and Mapping to Evolved Packet System QoS Classes," *Int'l J. Telemedicine and Applications*, 2010; www.hindawi.com/journals/ijta/2010/628086.
10. P. Agashe, R. Rezaiifar, and P. Bender, "CDMA2000 High Rate Broadcast Packet Data Air Interface Design," *IEEE Comm.*, vol. 42, no. 2, 2004, pp. 83–89.
11. K. Kang, "Probabilistic Analysis of Data Interleaving for Reed-Solomon Coding in BCMCS," *IEEE Trans. Wireless Comm.*, vol. 7, no. 10, 2008, pp. 3878–3888.
12. M. Zorzi, R.R. Rao, and L.B. Milstein, "Error Statistics in Data Transmission over Fading Channels," *IEEE Trans. Comm.*, vol. 46, no. 11, 1998, pp. 1468–1477.
13. G.B. Moody and R.G. Mark, "The Impact of the MIT-BIH Arrhythmia Database," *IEEE Eng. in Medicine and Biology*, vol. 20, no. 3, 2001, pp. 45–50.
14. Z. Herczeg et al., "Energy Simulation of Embedded XScale Systems with XEEMU," *J. Embedded Computing*, vol. 3, no. 3, 2009, pp. 209–219.

ACKNOWLEDGMENTS

This research was supported in part by the Ministry of Science, ICT, and Future Planning, Korea, under the Information Technology Research Center support program (NIPA-2014-H0301-14-1044) supervised by the National ICT Industry Promotion Agency; and in part by the Basic Science Research Program through the National Research Foundation of Korea (NRF) funded by the MSIP (NRF-2013R1A1A1059188); and in part by the Hong Kong Research Grants Council (RGC) Early Career Scheme (ECS) PolyU 5328-12E, Hong Kong PolyU A-PJ80, A-PK46, and A-PL82.

COVER FEATURE **COMPUTING IN HEALTHCARE**

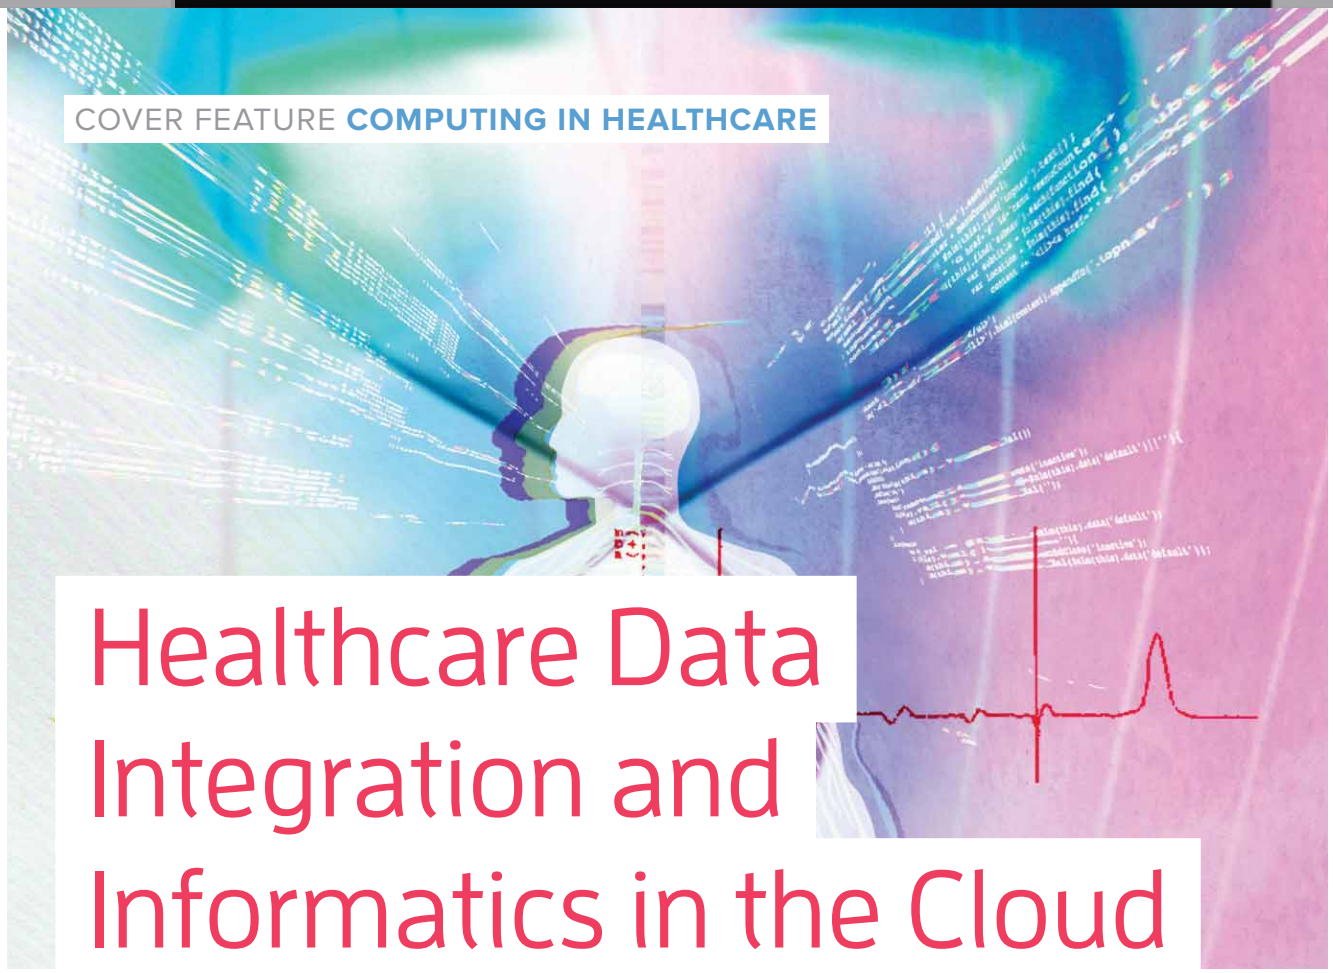

Healthcare Data Integration and Informatics in the Cloud

Arshdeep Bahga and Vijay K. Madisetti, Georgia Tech

An information integration and informatics framework for healthcare applications leverages the parallel computing capability of a cloud-based, large-scale distributed batch-processing infrastructure built with commodity hardware. The result is new flexibility for developers of advanced healthcare applications.

Because of the volume and variety of their data, healthcare applications providers and population health researchers face major challenges in integrating and effectively analyzing healthcare information. Traditional health information technology (IT) systems, such as electronic health record (EHR) and personal health record (PHR) systems, use different technical and semantic standards to represent and store data and are based on proprietary architectures. These client–server systems depend on local hardware, software, and data storage, and each system can have a different language and database technology. All these characteristics make it extremely difficult to accurately and easily integrate data from multiple, often conflicting, systems—yet such integration is key to developing advanced healthcare applications.

In contrast, cloud-based systems allow data storage on external servers[1] that developers can easily access, but interoperability remains a challenge, as the sidebar "Interoperability in Electronic Health Record Systems" describes. In earlier work,[2] we addressed this challenge by creating a cloud-based approach for the design of interoperable EHR systems and incorporating it in the Cloud Health Information Systems Technology Architecture (CHISTAR), a prototype system that enables semantic interoperability. CHISTAR's generic design methodology uses a reference model that defines a general-purpose data structure set and an archetype model that defines the clinical data attributes. CHISTAR enables secure access to healthcare data, supporting features such as authorization, identity management, and authentication services.

0018-9162/15/$31.00 © 2015 IEEE

We have extended that work to include a cloud-based information integration and informatics (III) framework that uses proven open source, cloud-based technologies to facilitate the collection and analysis of clinical data from diverse geographical locations. Its features include

- integration of data from distributed and heterogeneous sources into a common nomenclature,
- easier access to healthcare data stored in the cloud,
- efficient analysis of massive healthcare data collected in the cloud, and
- healthcare data storage and life-cycle management.

A significant advantage of our III framework is its use of technologies for clinical data integration and analysis that leverage the benefits and economies of cloud computing environments already in use in other domains. As the volume of clinical data continues to grow exponentially from distributed, heterogeneous sources, data analysis is becoming increasingly more problematic and a bottleneck to more sophisticated healthcare applications. Data integration approaches for collecting clinical data from distributed and heterogeneous health IT systems will contribute to more effective healthcare applications. Approaches for massive-scale clinical data analytics will facilitate the development of more efficient healthcare applications, improve prediction accuracy, and help in timely decision making.

Our III framework, together with the CHISTAR middleware, allows the collection, organization, and secure exchange of healthcare data from a range of stakeholders, including patients, hospitals, therapists, and insurers, in a range of formats (databases, structured and unstructured, and so on). Data exchange on this scale can help ensure more timely and accurate care delivery and thus potentially lower healthcare costs while raising quality. In addition, the framework will aid in the development of advanced healthcare applications, such as epidemiological surveillance and adverse drug events prediction.

Epidemiological surveillance, in particular, is representative of these advanced applications because it involves studying the distribution and determinants of health-related states or events in specified populations and applying any findings to diagnose diseases under national surveillance.[3] For that reason, we chose it as a use case to demonstrate how our framework can support developers in creating advanced healthcare applications. Our subsequent evaluation revealed that for applications such as epidemiological surveillance, our III framework provides several benefits relative to client–server EHR systems, including better scalability, faster development, and lower cost.

ARCHITECTURAL OVERVIEW

As Figure 1 shows, the III framework is part of a stack that includes Informatics App Builder and CHISTAR

INTEROPERABILITY IN ELECTRONIC HEALTH RECORD SYSTEMS

Because EHR data and system interoperability is the cornerstone of advanced healthcare applications, research is increasing into possible ways to provide it. Semantic interoperability is the focus of OpenEHR (www.openehr.org), for example. Mirth Connect (www.mirthcorp.com/products/mirth-connect), an open source integration engine, supports a variety of messaging standards and protocols for connecting to external systems and numerous databases for storing message data.

One study[1] described the potential of EHR data for epidemiological surveillance, but current approaches for tracking large-scale regional trends are not compatible with a healthcare system that operates in silos. To promote interoperability in this context, the study team proposed a tuberculosis detection algorithm for EHR data followed by implementation in a live surveillance and reporting system.

Another research effort[2] proposed methods of EHR-based surveillance to determine the risk factors of coronary heart disease and described a model of population-level EHR-based surveillance for those factors.

References

1. M.S. Calderwood et al., "Real-Time Surveillance for Tuberculosis Using Electronic Health Record Data from an Ambulatory Practice in Eastern Massachusetts," *Public Health Reports*, vol. 125, 2010, pp. 843–850.
2. J.J. VanWormer, "Methods of Using Electronic Health Records for Population-Level Surveillance of Coronary Heart Disease Risk in the Heart of New Ulm Project," *Diabetes Spectrum*, vol. 23, no. 3, 2010, pp. 161–165.

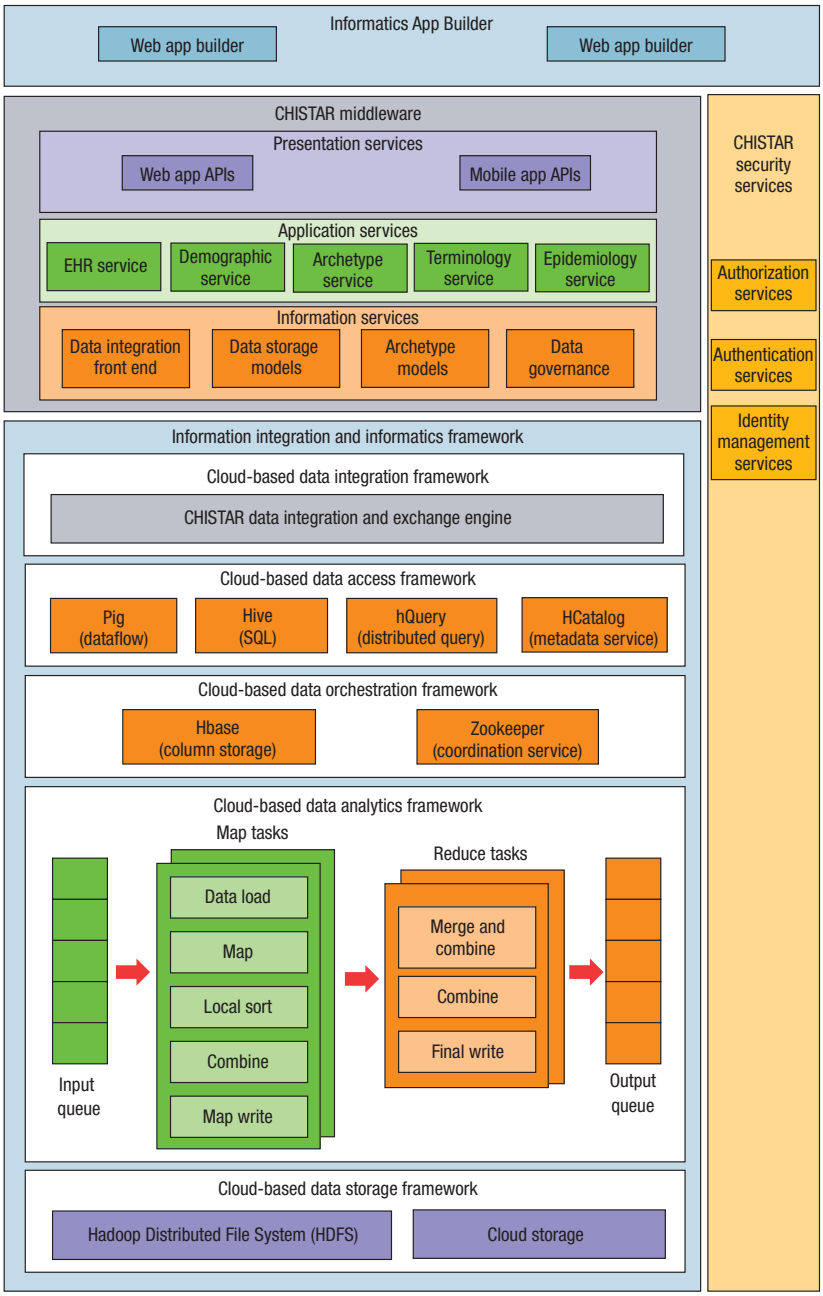

FIGURE 1. Technology stack for building healthcare applications that require massive clinical data from distributed, heterogeneous electronic health record (EHR) systems. Informatics App Builder provides application development tools, the Cloud Health Information Systems Technology Architecture (CHISTAR) middleware enables semantic interoperability, and the information integration and informatics (III) framework enables data integration, access, analytics, and storage.

middleware. To facilitate user interface development, the middleware has platform tools for mobile operating systems, including Android and iOS, and Windows for desktop applications. The framework also provides tools and APIs for data integration, access, analytics and storage.

Application building tools

As its name implies, Informatics App Builder is a tool suite for building mobile and Web-based healthcare applications. Because the III framework and CHISTAR middleware take care of underlying cloud infrastructure management, deployment configuration, and data management, developers can use Informatics App Builder without considering configuration and maintenance activities in the cloud.

Secure interoperability

CHISTAR middleware provides a variety of presentation, application, and information services to support the development of advanced informatics applications.[2] Its security services address the key requirements of the Health Insurance Portability and Accountability Act (HIPAA) and Health Information Technology for Economic and Clinical Health Act (HITECH), both of which require covered entities (those that create, maintain, transmit, use, and disclose an individual's protected health information) to protect the integrity, confidentiality, and availability of information they collect, maintain, use, or transmit.

CHISTAR's application services are platform-independent, and developers can customize tools to support specific platform features, enabling healthcare application development for a wide variety of platforms and devices.

Table 1 lists CHISTAR's security services.

Data integration

Data integration methods vary primarily in the level on which they focus. At the *application level*, integration involves integrating data from individual applications by reimplementing them in one domain-wide application. In integration at the *API level*, applications expose their APIs so that other applications can access

their data. Finally, at the *data level*, integration establishes a common domain model or global schema so that independently developed applications can exchange information. Data-level integration, which is what our III framework incorporates, has more development flexibility than the other approaches because it allows independent development and focuses on common exchange.

FRAMEWORK IMPLEMENTATION

Our III framework consists of multiple layers, starting with the data integration and exchange engine in CHISTAR, which makes up the *data integration* layer. The *data access* layer uses a series of open source technologies to enable access to cloud data.

The framework enables data interoperability by using a data storage model to represent data structure and a domain model to represent clinical knowledge.

The *data orchestration* layer consists of HBase and Zookeeper. HBase is a distributed nonrelational column-oriented database that runs on top of the Hadoop Distributed File System (HDFS; http://hadoop.apache.org/mapreduce), which is part of Hadoop, a cloud-based distributed batch-processing infrastructure. HBase provides a fault-tolerant way of storing large quantities of sparse data. Zookeeper is a distributed coordination service for maintaining configuration information, naming, and providing distributed synchronization and group services. The data analytics layer is built on top of Hadoop.

Finally, the *data storage* layer consists of HDFS and a cloud storage for raw files such as images. Implementation of the integration, access, and analytics layers was the most challenging.

TABLE 1. Security services in CHISTAR middleware.

Service	Basis
Authentication	Security Assertion Markup Language - Single Sign On (SAML-SSO)
Authorization	Open standard to authorization (OAuth)
Identity management	Federated identity
Data-at-rest security	Advanced Encryption Standard (AES) encryption
Data-in-motion security	Secure Sockets Layer (SSL)
Key management	Use of separate keys for key storage, rotation, and encryption
Data integrity assurance	Message authentication codes
Auditing	Logs of all reads and writes

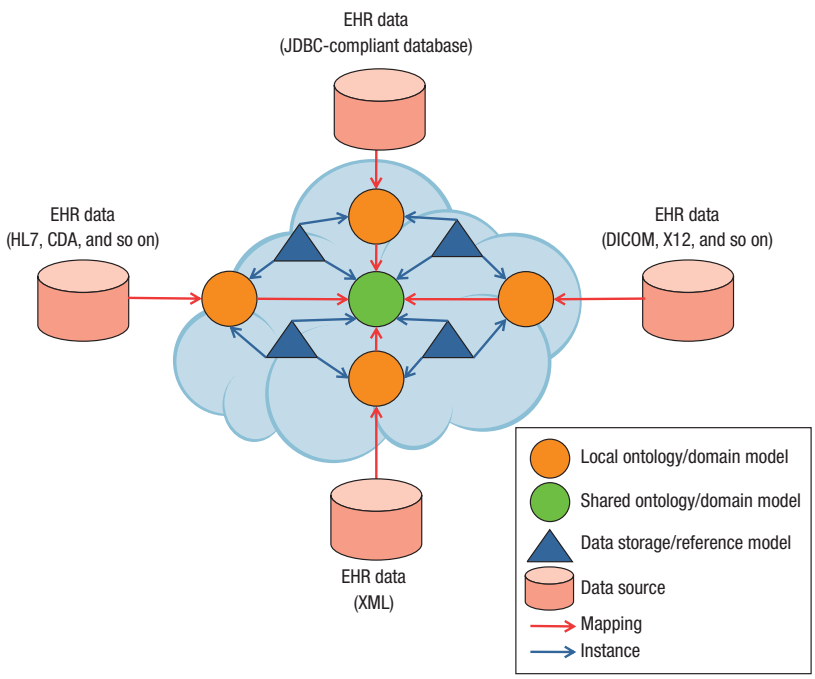

FIGURE 2. Proposed approach for data integration. The domain model is a conceptual representation of the application domain, the data storage model is the logical structure for data storage, and the mapping relates the source data to the domain model. Using two separate models (data storage and domain models) avoids the need to change software to accommodate changes in clinical knowledge. CDA: Clinical Document Architecture; DICOM: Digital Imaging and Communications in Medicine; JDBC: Java Database Connectivity; HL7: Health-Level 7.

Data integration

As Figure 2 shows, data integration allows integrating healthcare data that exists in various forms (structured or unstructured) on different data storage systems such as relational database management systems (MySQL and Oracle, for example), file servers (with text, image and video files), and

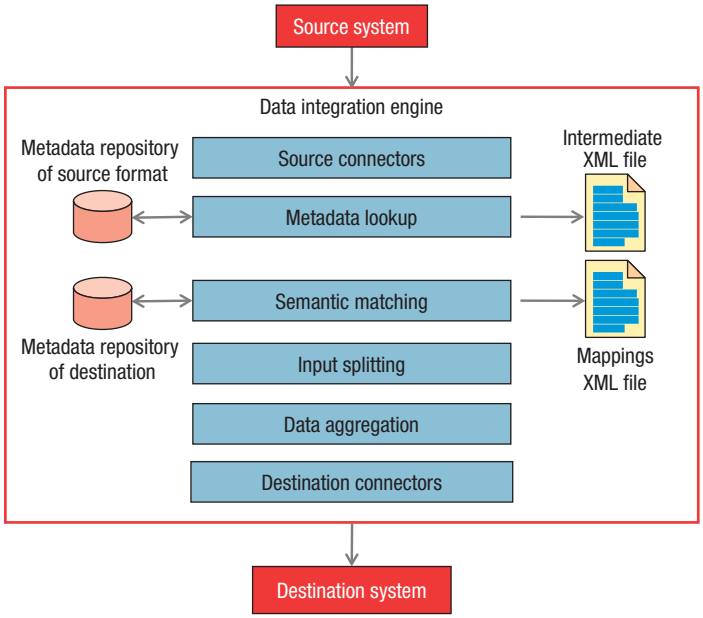

FIGURE 3. Data integration engine. Integration consists of connecting to an external system, parsing the input file to identify source data formats, matching semantics to source data formats stored in a metadata repository, splitting input to exploit parallelism, aggregating data and transforming it to the destination system's data format, and finally connecting to a destination system to write the results to storage.

EHR standards, such as Health Level-7 (HL7) messages.

Separating structure and knowledge. Using separate models to represent data structure and clinical knowledge has several benefits. The data storage model defines entities for data storage and represents the semantics of storing data, regardless of the domain. The domain model represents clinical knowledge and constraints on the generic data structures that the data storage model defines. By making the domain representation explicit, the domain model captures system-independent specifications for that domain. Consequently, there is no need to change the software when clinical knowledge changes, and the entire system becomes more robust.

Mapping. The CHISTAR data integration and exchange engine maps the data in source-specific format to local or global domain models. Three strategies are possible in mapping the source data format to the domain model:

- a global domain model for all data sources,
- a local domain model for each data source, or
- a hybrid approach that consists of local domain models for each data source and a shared global domain model.

In the first approach, one domain model expresses the shared semantics among all data sources, which makes the domain model highly susceptible to changes with the addition of data sources. Moreover, finding a minimal domain specification (or ontology) for all data sources might be extremely difficult.

In the second approach, each data source has a separate (local) domain model, which simplifies domain model construction. However, the lack of a common ontology complicates the comparison of domain models, making it harder to define mappings among them.

In the third approach, a hybrid of the first two approaches, each data source has its own domain model built on a single (global) domain model. The global domain model is less susceptible to changes with the addition of new data sources, and it is easier to find a minimal domain specification (ontology) because the domain knowledge that contradicts different data sources can be separate local domain models for each source.

Moving from source to destination. Figure 3 shows the data integration engine's architecture. Supported standards include HL7, Clinical Document Architecture (CDA), Continuity of Care Record (CCR), Digital Imaging and Communications in Medicine (DICOM), X12, Continuity of Care Document (CCD), XML, National Council for Prescription Drug Programs (NCPDP), Electronic Data Interchange (EDI), Delimited Text, and raw ASCII.

Once the source connector is in place, the data integration engine performs metadata lookup to discover the semantics of the source data elements. The lookup process consists of parsing the input file and looking up the source data's metadata repository to retrieve the semantics of all the source file's data elements. The data integration engine maintains metadata repositories for all the data types it supports.

The metadata lookup process is data driven and produces an intermediate XML file that has all the source file's data elements as well as the annotations from the repository lookup. The intermediate XML file eliminates the source data syntax and retains the data elements' hierarchy and properties along with the annotations.

To match semantics, which is the next integration step, the framework searches the metadata repository of the destination format and retrieves a list of candidate mappings for each data element in the intermediate file. To guide

the search, it uses the annotations of source data elements from the intermediate file. Semantic matching can be either automated or manual. In automated matching, the framework retains the most similar candidate mappings for all the source data elements.

The framework splits the input to parallelize data importing and uses jobs written in the MapReduce parallel programming model (described later) to aggregate data and transform it to the destination format. Finally, destination connectors write the created data files to HDFS storage.

Data access

The data access layer provides APIs for querying and retrieving healthcare data from the cloud. The data access layer is based on Pig, Hive, and hQuery—all open source technologies.

Pig, a platform for analyzing large datasets, consists of a high-level language for expressing data analysis programs and an infrastructure for evaluating them. Application developers can write procedural scripts in Pig Latin (Pig's language), and Pig's compiler produces sequences of MapReduce programs, which enable parallel data processing.

Hive provides a data warehousing infrastructure on top of Hadoop, which facilitates data queries and the analysis of large datasets stored in Hadoop-compatible file systems. Hive, which uses the SQL-like Hive Query Language (HQL), allows data querying in HDFS or HBase. hQuery (http://projecthquery.org) is an open source framework for the distributed querying of healthcare data.

HCatalog is a metadata service for Hadoop that is built on top of Hive metastorage and wraps additional layers around it. Because HCatalog provides a shared schema and data model for Pig, Hive, and MapReduce, developers need not worry about the data's location in the cloud. More important, HCatalog insulates data analytics applications from schema, location, or format changes, so if such changes occur, developers need not rewrite their data analytics applications.

Data analytics

Through its data analytics layer, the III framework supports a wide variety of data analysis algorithms within a cloud architecture. The analytics layer uses MapReduce to formulate healthcare data analysis jobs. The data analytics layer allows efficient data analysis of the massive healthcare data collected in the cloud.[4] The data analytics layer is based on Hadoop, a framework that provides an open source implementation of the MapReduce parallel processing model. MapReduce has two phases: *map* and *reduce*. In the *map* phase, MapReduce reads data from a distributed file system, such as HDFS; partitions the read data among computing nodes in a cluster; and sends the data to the nodes as key-value pairs. MapReduce processes each input record independently and stores key-value pairs as intermediate data on the local disk of the node running the map task.

When all the map tasks are complete, the reduce phase begins, which involves aggregating intermediate data with the same key. An optional combine task aggregates intermediate data with the same key for the map task's output *before* transferring the output to the reduce task.

SUPPORTING ADVANCED HEALTHCARE APPLICATIONS

As our III framework functions imply, data integration and analysis are powerful features that give developers the tools needed to exploit massive amounts of diverse data, empowering them to take healthcare applications to the next level. EHR systems include individual-level laboratory results, and diagnostic, treatment, and demographic data. Although EHRs were designed for clinical interactions between patient and provider, the EHR data is useful for population-level health surveillance efforts, disease detection, outbreak prediction, and public health mapping.

USING SEPARATE MODELS TO REPRESENT DATA STRUCTURE AND CLINICAL KNOWLEDGE MAKES THE SYSTEM MORE ROBUST, SINCE DOMAIN KNOWLEDGE CAN CHANGE WITHOUT AFFECTING SEMANTICS.

Because EHR system data is continuously updating, a framework that integrates data from multiple EHR systems will enable applications that can effectively and accurately predict outbreaks. Examples of more advanced healthcare applications that our III framework can support include

- epidemiological surveillance to predict outbreaks;

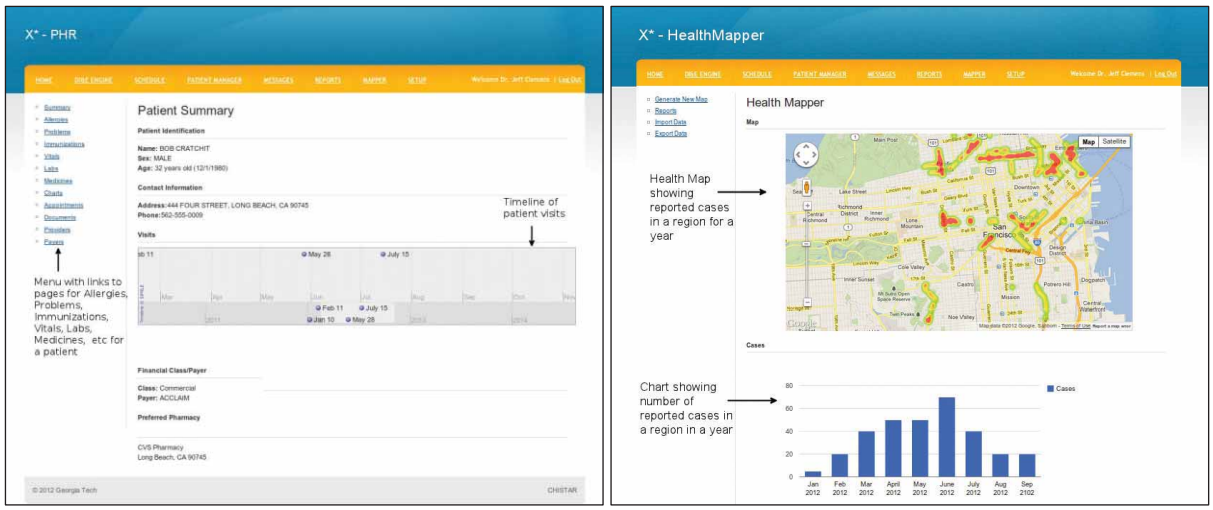

FIGURE 4. Results of applications developed using the proposed technology stack. (a) Screenshot of the personal health record (PHR) application, which demonstrates the ability to represent data from multiple heterogeneous health IT systems. (b) Screenshot of HealthMapper, which demonstrates the ability to analyze large amounts of data.

- patient similarity-based decision intelligence, which analyzes EHR data to extract a cluster of patient records most similar to a specific target patient;
- adverse drug events prediction, which predicts the patients most at risk for having an adverse response to a certain drug on the basis of other patients' adverse drug reactions;
- medical reconciliation, which detects omissions in a patient's medication list and identifies drugs that the patient might be taking that are not on the list; and
- medical prognosis, which predicts the likely outcome of an illness for a patient based on outcomes for similar patients.

USE CASE EVALUATION

To demonstrate the effectiveness of our technology stack (Information App Builder, CHISTAR middleware, and III framework), we created a PHR application that integrates data from 100 patients simultaneously. We also built the HealthMapper application, which analyzes massive amounts of EHR data to determine reported cases of a particular illness for a specific region during a specific period. To deploy our technology stack, we used the Amazon Elastic Compute Cloud (EC2) infrastructure (http://aws.amazon.com/ec2).

PHR application

Figure 4a shows a screenshot of the PHR application's patient summary page, a unified representation of data from multiple health IT systems.

HealthMapper application

Figure 4b shows a screenshot of HealthMapper, which groups cases according to address and zip code. To cluster reported cases, we ran MapReduce offline hourly or daily and stored the aggregated results in HBase. HealthMapper displays the aggregated results.

Range queries can be created from the application to display the reported cases over a specific period. HealthMapper uses the APIs provided by the framework's data access layer and the epidemiology service to query and analyze patient health records. The application demonstrates that with the proposed III framework it is possible to analyze massive healthcare data integrated from different health IT systems.

Response time

Figure 5a shows the average response time for the PHR application for three deployment configurations and a range of PHRs with 100 simultaneous users. Notably, response times improve with both vertical and horizontal scaling.

Figure 5b shows the offline clustering time for HealthMapper. The 5H(large) deployment can cluster 10 million records in only a few minutes, which means that clustering jobs can run hourly if necessary.

Integrating healthcare data from a variety of providers—not just traditional medical institutions and physicians, but also dentists, nurse practitioners, physical therapists, and psychologists—will improve both current healthcare delivery and research into disease prevention and public health issues. In the short term, integration will improve care coordination, decrease clinical mistakes from missing or incomplete data, reduce duplicate testing, and increase patient safety through more timely and accurate medication reconciliation. In the long term, integration will open channels through which to conduct public health research and population surveillance to pinpoint health issues.

Our proposed technology stack—Informatics App Builder, the CHISTAR middleware, and the III framework—combine to provide developers with

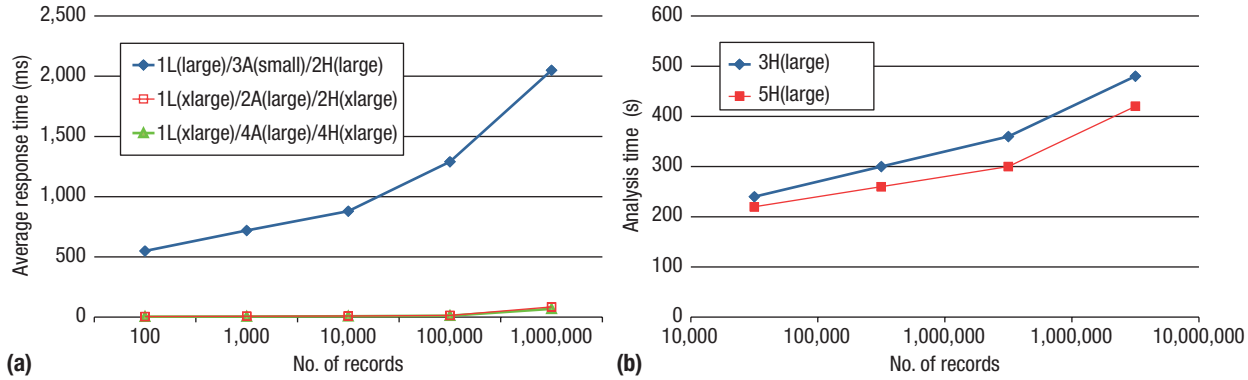

FIGURE 5. Application response times. (a) Average response time for the PHR application for up to a million patient records with 100 simultaneous users. (b) Offline clustering time for the HealthMapper application for up to a million patient records. #L: number of instances running load balancers and Web servers; #A: number of instances running application servers; #H: number of instances running the Hadoop/HBase cluster.

massive-scale data integration and analysis, which has significant benefits for application development.

Relative to client–server architectures, application development is faster and cost is lower, since data integration, storage, and analysis use open source, cloud-based technologies.

Applications built with our III framework also provide more opportunities to collaborate, since they have superior accessibility. Users can securely log into the system from anywhere with an Internet connection. Consistent data representation, data access, and interpretation of integrated data also enhance collaboration.

The security features of CHISTAR middleware, such as role-based access control, allow doctors, physicians, and specialists to collaborate more freely within a secure environment, thus improving the continuity of provided care.

Another important benefit is better scalability. Cloud-based applications built with the III framework have better scalability relative to client–server EHRs. The computing resources used by such applications can be scaled up on demand as more data is integrated and new users are added. Applications built with the III framework can leverage both horizontal (scaling out) and vertical scaling (scaling up). Because the III framework uses HBase storage, it can scale storage linearly and automatically by adding nodes.

Finally, cloud-based applications built with the III framework have reduced infrastructure and operation costs. Client–server applications require a team of IT experts to install, configure, test, run, secure, and update hardware and software. In cloud-based applications, the cloud provider takes care of all those functions.

ABOUT THE AUTHORS

ARSHDEEP BAHGA is a research scientist at Georgia Tech. His research interests include cloud computing, big data analytics, digital signal processing, and embedded software systems. Arshdeep received an MS in electrical and computer engineering from Georgia Tech. Contact him at arshdeep@gatech.edu.

VIJAY K. MADISETTI is a professor in the Department of Electrical and Computer Engineering at Georgia Tech. His research interests include digital signal processing, embedded computing systems, chip design, wireless and telecommunication systems, and systems engineering. Madisetti received a PhD in electrical engineering and computer sciences from the University of California, Berkeley. He is an IEEE Fellow, and executive director of Georgia Tech's India Initiative. Contact him at vkm@gatech.edu.

REFERENCES

1. A. Bahga and V. Madisetti, *Cloud Computing: A Hands-on Approach*, CreateSpace, 2013.
2. A. Bahga and V. Madisetti, "A Cloud-Based Approach for Interoperable Electronic Health Records (EHRs)," *IEEE J. Biomedical and Health Informatics*, vol. 17, no. 5, 2013, pp. 894–906.
3. C. Hajat, "An Introduction to Epidemiology," *Methods Molecular Biology*, Jan. 2011, pp. 27–39.
4. A. Bahga and V. Madisetti, *Internet of Things: A Hands-on Approach*, CreateSpace, 2014.

COVER FEATURE **COMPUTING IN HEALTHCARE**

Whole Genome Sequencing: Revolutionary Medicine or Privacy Nightmare?

Erman Ayday, Bilkent University
Emiliano De Cristofaro, University College London
Jean-Pierre Hubaux, EPFL, Lausanne
Gene Tsudik, University of California, Irvine

Whole genome sequencing will soon become affordable for many individuals, but thorny privacy and ethical issues could jeopardize its popularity and thwart the large-scale adoption of genomics in healthcare and slow potential medical advances.

In the past decade, whole genome sequencing (WGS) has evolved from a futuristic concept to a realistic technology that yields an individual's complete genome. Each genomic sequence contains a vast amount of information that enables significant progress in understanding, treating, and preventing disease. As such, WGS has the potential to revolutionize healthcare.

However, a genome also contains highly sensitive information that uniquely identifies an individual. When technology advances eventually make WGS affordable for the general population, individuals will need assurances about access to their genomic information. For example, who will store the digitized genome and where? How will access be controlled such that no one can inadvertently or deliberately leak genomic information to third parties? What will keep a healthcare provider's service partners from using genomic information in ways other than medical research or personalized medical treatment?

With DNA sequencing cost dropping below $1,000 per genome, these questions have become pressing. Both throughput gains and the cost reductions of new-generation sequencing platforms have defied Moore's

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

0018-9162/15/$31.00 © 2015 IEEE

ONGOING WORK TO PROTECT GENOMIC DATA

Over the past few years, research in genomic privacy has accelerated and now falls into four main categories:

- » string searching and comparison,
- » release of aggregate data,
- » alignment of raw genomic data, and
- » clinical use of genomic data, such as for personalized medicine.

Work in the first category is experimenting with the use of medical tools and private string comparison for privacy-preserving paternity tests, personalized medicine, and genetic compatibility tests.[1] More recently, researchers have extended that work to implement the GenoDroid toolkit.[2] which provides paternity and ancestry testing via a smartphone.

In the second category, researchers are focusing on privacy risks of releasing aggregate genomic data.[3] Others have explored the application of *differential privacy* to the publication of aggregate genomic trial statistics.[4,5] Their work aims to ensure that two genomic databases, which differ only by one individual's data, have indistinguishable statistical features. Hence, the published result from a genomic dataset does not reveal the existence of a particular individual in that dataset.

Research in the third category is looking at secure and efficient algorithms for read mapping (aligning millions of short sequences to a reference DNA sequence). One recent attempt on this direction works in a hybrid (public and private) cloud environment.[6] In this work, authors outsource the computationally intensive steps of the operation to a public (untrusted or commercial) cloud; they propose doing sensitive and lightweight computations on a private (trusted) cloud to protect the privacy of sensitive DNA information.

In the last category is work to preserve the patient's privacy in medical tests and personalized medicine. One approach uses homomorphic encryption and secure multiparty computation to protect patients' genomic data in this context.[7,8]

Some of these efforts have already materialized into practical genomic testing. However, it is hard to foresee the range and complexity of future genetic operations: some tests might be too computationally intricate to be performed on a personal device, or genetic tests might involve multiple genomes. Consequently, we expect the scope and nature of genomic data protection work to change as researchers make new discoveries and shift their focus to address a new set of needs. At the same time, the efforts already in progress are important stepping stones to solutions that address the multifaceted challenge of protecting genomic data.

References

1. P. Baldi et al., "Countering GATTACA: Efficient and Secure Testing of Fully-Sequenced Human Genomes," *Proc. 18th ACM Conf. Computer and Communications Security* (CCS 11), 2011, pp. 691–702.
2. E. De Cristofaro et al., "Genodroid: Are Privacy-Preserving Genomic Tests Ready for Prime Time?" *Proc. ACM Workshop Privacy in the Electronic Society* (WPES 12), 2012, pp. 97–108.
3. X. Zhou et al., "To Release or Not to Release: Evaluating Information Leaks in Aggregate Human-Genome Data," *Proc. 16th European Conf. Research in Computer Security* (ESORICS 11), 2011, pp. 607–627.
4. F. Yu et al., "Scalable Privacy-Preserving Data Sharing Methodology for Genome-Wide Association Studies," *J. Biomedical Informatics*, Feb. 2014, pp. 133–141.
5. A. Johnson and V. Shmatikov, "Privacy-Preserving Data Exploration in Genome-Wide Association Studies," *Proc. 19th ACM Int'l Conf. Knowledge Discovery and Data Mining*, 2013, pp. 1079–1087.
6. Y. Chen et al., "Large-Scale Privacy-Preserving Mapping of Human Genomic Sequences on Hybrid Clouds," *Proc. 19th Network and Distributed System Security Symp.* (NDSS 12), 2012; www.informatics.indiana.edu/xw7/papers/ndss2012.pdf.
7. E. Ayday et al., "Privacy-Preserving Computation of Disease Risk by Using Genomic, Clinical, and Environmental Data," *Proc. Usenix Security Workshop Health Information Technologies* (HealthTech 13), 2013; www.usenix.org/conference/healthtech13/workshop-program/presentation/ayday.
8. E. Ayday et al., "Protecting and Evaluating Genomic Privacy in Medical Tests and Personalized Medicine," *Proc. ACM Workshop Privacy in the Electronic Society* (WPES 13), 2013, pp. 95–106.

COMPUTING IN HEALTHCARE

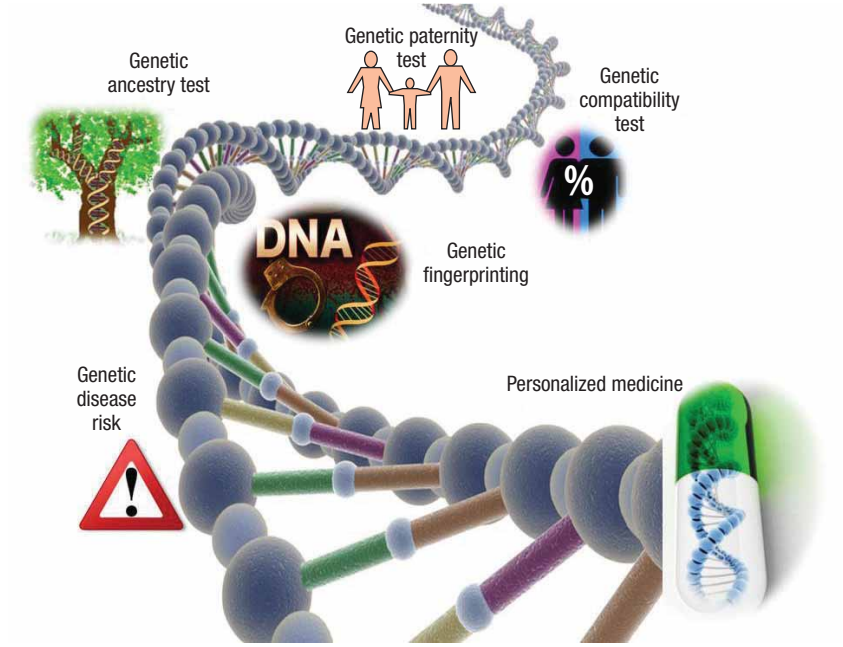

FIGURE 1. Genomics applications. Whole genome sequencing will enable personalized genomic medicine and facilitate testing for genetic disease risk and ancestry.

law. Thus, it is safe to assume that, in a few years, most individuals in developed countries will be able to obtain their digitized genomes for any number of purposes—from personalized medicine to paternity testing. Commercial entities, such as Knome and Illumina, already offer services that create reports from raw genomic data, which doctors use to guide treatment.

However, without a deeper understanding of the complex interplay between genomes and healthcare, WGS applications will be limited. Achieving progress in this research will require patients (or volunteers) who are willing to share their genetic data—an agreement that raises privacy protection, ethical use, and legal rights concerns. For example, in the Personal Genome Project (www.personalgenomes.org), participants agree to make their genomic data and other personal information publicly available on the Internet. Such pilot projects offer a glimpse into the future concerns of handling large-scale genomic data.

DNA sequencing greatly exacerbates data exposure and exploitation issues that social media and personal health records (PHRs) have already brought to the forefront. The genome represents an individual's biological identity and thus contains rich information about that person's ancestry. By combining the genomic data with data on the person's environment or lifestyle, a third party can infer the individual's phenotype, including predisposition to physical and mental health conditions (such as Alzheimer's disease, cancer, or schizophrenia).

If a genomic information leak occurs, revoking or replacing an individual's DNA sequence is impossible, which has serious implications for applications that depend on accurate genomic information. The use of DNA analysis in law enforcement and healthcare, for example, is already prompting ethical questions, such as how to guarantee the genomic information's integrity.

Until researchers address these open problems, the much anticipated benefits of personalized medicine could remain on hold.

GENOMICS 101

The human genome is encoded in double-stranded DNA molecules that consist of two complementary polymer chains. Each chain is a series of nucleotides, represented as the letters A, C, G, and T. Technicians collect DNA samples from a person's saliva, hair, skin, or blood, among other sources, and extract genetic material for sequencing. The resulting genome is a unique string of approximately 3.2 billion letter pairs (an arrangement of A, C, G, and T).

The reference genome, which scientists have assembled as a representation of the human genome, makes up 99.5 percent of a human's DNA sequence. The remaining 0.5 percent represents the individual's genetic variation. Although it might seem insignificant relative to the reference genome, this minuscule 0.5 percent corresponds to several million nucleotides.

The genetic variation can take several forms, the most common being single nucleotide polymorphism (SNP, pronounced "snip"). In simplest terms, a SNP is a position in the genome sequence with a nucleotide that varies between individuals. For example, in two sequenced DNA fragments from different individuals, AAGCCTA and AAGCTTA, the fifth nucleotide is C in one and T in the other.

Researchers have confirmed that humans have approximately 50 million unique SNPs,[1] a number that becomes more exact as more individuals consent to sequencing.

SNPs can help determine an individual's predisposition to certain disorders or diseases. For example, recent genome-wide association studies show that the presence of three genes with 10 particular SNPs can indicate susceptibility to Alzheimer's disease.[2,3]

Interdependent SNPs sometimes result in linkage disequilibrium (LD)[4]—the nonrandom association of alleles at two or more loci. The alleles descend

from single, ancestral chromosomes, so LD makes it possible to infer the nucleotide of a SNP from the contents of other SNPs. This relationship obviously complicates privacy protection.

PERSONALIZED MEDICINE AND BEYOND

WGS has the potential to bring about a new era of predictive, preventive, participatory, and personalized (P4) medicine[5] and enable applications such as those in Figure 1. P4 represents a significant healthcare paradigm shift[6] from the current trial-and-error treatment because it enables medication tailored to a patient's precise genetic makeup. P4 applications include assessments of disease and treatment risk, and paternity and ancestry testing, and the evaluation of genetic compatibility between potential partners to reduce the possibility of passing genetic diseases to their offspring.

Pharmacogenomics

Experiments have shown that certain genetic mutations alter drug metabolism and that genomic tests can help predict a patient's response to particular drugs. This experimentation and testing is part of *pharmacogenomics*—the study of how genetic variations affect an individual's response to medications. Examples of pharmacogenomics include testing for SNP mutations in the tpmt gene of children with leukemia and pretreatment testing for the correlation of the BRCA1/BRCA2 genes to familial breast and ovarian cancer syndromes.

Genomic tests to determine drug response are expected to become more widespread in the near future. Experts estimate that about a third of the 900 cancer drugs now in clinical trials could soon come to market with an enclosed recommendation for a DNA or another molecular test.[7]

Programs are underway to support pharmacogenomics. For example, Vanderbilt University's Pharmacogenomic Resource for Enhanced Decisions in Care and Treatment (Predict) program[8] evaluates patients' genetic characteristics to help physicians determine which drugs are most likely to work, thus avoiding the long trial-and-error period characteristic of traditional drug evaluation. In one case,[9] Predict program researchers used the genetic profile of a patient with coronary artery disease to help doctors select a specific cholesterol-lowering drug and successfully treat the patient in a fraction of the time with a conventional approach.

Testing for genetic disease risk

Low-cost WGS will give individuals direct access to their genomic information, which they could share with sites that test for genetic disease risks. One such site, 23andMe, already provides relatively low-cost genetic ancestry and disease risk tests for 960,000 specific SNPs, although it does not yet offer WGS. Since November 2013, the US authorities have suspended the health-related 23andMe tests, pending FDA investigation; however, such tests are still offered in the UK.

In parallel to direct-to-consumer services, national and regional efforts are attempting to introduce genomics into the clinical setting. Examples include the UK's 100,000 Genomes Project (www.genomicsengland.co.uk) and University Hospital Lausanne's biobank (www.chuv.ch/biobanque/bil_home/bil-patients-famille/bil-la_bil.htm).

Although researchers are enthusiastically exploring the relationship of genetics and personalized medicine, biomedical experts have expressed doubts about the extent to which gene mapping can predict the likelihood of developing a disease.[10] They argue that, although scientists have a list of genetic features that correlate to certain diseases,[2] they do not know whether (and to what extent) environmental factors also come into play.

Paternity and ancestry testing

The availability of a patient's fully sequenced genome will enable clinicians, doctors, and testing facilities to run complex, correlated genetic tests in a matter of seconds. Compared with the more expensive in vitro tests, these specialized computational algorithms enable faster and more accurate testing while preserving legal acceptance.

Commercial entities already offer ancestry and genealogical testing in which software compares an individual's genomic information with publicly available genomic data from a particular ethnic group to determine how the individual relates to the group. Online services also offer genetic compatibility tests that assess the risk of Mendelian inheritance[11]—the chance of transmitting genetic diseases to any offspring—in the couple being tested.

THREATS TO GENOMIC DATA PRIVACY

Many view genomic privacy with skepticism, since every individual constantly leaves behind biological material, such as hair, skin, or saliva—evidence that a third party can collect even days later and use to construct a DNA sequence. However, this threat is credible only for a targeted individual or a small group, not for a large

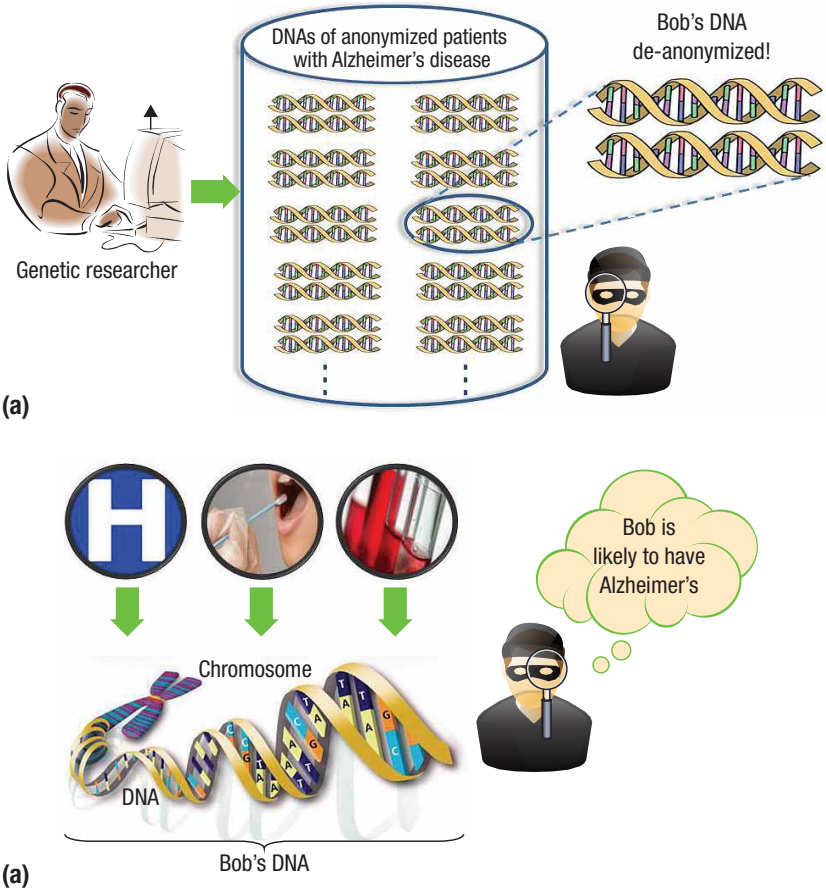

FIGURE 2. Two main threats to human genomic data privacy. (a) DNA donors in a public research database lose anonymity (de-anonymization), and (b) partial genomic data leakage allows outsiders to infer sensitive information. Figure used with permission from the US Department of Energy Genomic Science program (https://public.ornl.gov/site/gallery/detail.cfm?id=398&topic=&citation=&general=dna&restsection=all).

number of digitized genomes, such as in a research database.

Genomes in the latter setting face two main threats, as Figure 2 illustrates. Although existing laws protect data privacy in general, genomic data has certain characteristics that require more restrictive provisions to address unique privacy threats.[12]

Loss of donor anonymity

The primary traditional approaches to privacy protection are data de-identification or aggregation. Common de-identification strategies, which include deleting or masking identifiers, such as names and Social Security numbers, are ineffective for genomic data because the genome is the ultimate identifier.[13]

Aggregation—a strategy that combines data for a population—is also ineffective because enough published information is available to identify the individual from a case study and, in some instances, to recover parts of the genome sequence. For example, a 2009 study[14] shows that even the test statistics (such as *p*-values, *r*-squares) calculated from allele frequencies and published papers give away enough information to identify genetic trial participants. A 2013 study[15] demonstrated that third parties can use information from popular genealogy websites along with other available personal data to re-identify (counter de-identification of) DNA donors from a public research database.

Data leaks

Because the genomes of two closely related individuals are highly similar, the disclosure of a person's genome can possibly leak significant genomic information about that person's close relatives. This disclosure is a problem regardless of whether it was voluntary, accidental, or malicious.

The possibility of revealing others' identities makes genomic data privacy a unique issue, since, in most other sensitive scenarios, only the individual's data is at stake. Depending on the number of siblings and children, disclosure can affect a large group.[16] Failing to consider this possibility can have severe consequences, as the recent controversy about Henrietta Lacks' genome sequence attests. In researching Lacks' disease nearly five decades ago, scientists discovered cell properties in her cancerous tissue that made the cells highly suitable for biogenetic research. They harvested more cells without the family's knowledge and began using the HeLa cell (in honor of Lacks' first and last names) in studies. It eventually became so popular in genetics research that Lacks' surviving family members began receiving requests for tissue and blood samples. After several court cases to address privacy violations, in 2013, the National Institutes of Health (NIH) agreed to give the family some control over the HeLa cells' use.

Exacerbating the data leak problem is the genome's immutability and longevity. An individual can change passwords, account numbers, and even public key certificates. The same is not true of a genome. Moreover, future generations will inherit most of their ancestor's DNA, so genomic information disclosure can become an endless curse.

PRIVACY PROTECTION LAWS

Clearly, privacy concerns represent a formidable obstacle to assembling large human genomic databases and can delay (or derail) genome-wide association studies, which in turn could thwart advances in medicine and subsequent healthcare improvements. In law enforcement, which increasingly uses DNA-based identification, the need for genomic data security and reliability is also evident.

Existing laws protect genomic data privacy to some degree. In 1990, the National Human Genome Research Institute established the Ethical, Legal, and Social Implications Research Program to explore the repercussions of advances in genetic and genomic research on individuals, families, and communities. In 2008, the US government established the Genetic Information Nondiscrimination Act (GINA), which prohibits health insurance and employment discrimination on the basis of genetic information. Also, the Health Insurance Portability and Accountability Act (HIPAA) provides a general framework for protecting and sharing health information, and the State of California has begun to consider DNA privacy laws.[17] Meanwhile, in Europe, legislators are taking similar precautions.[18]

Discrimination through genetic data is not a new idea. As far back as 1997, *Gattaca*, a popular science fiction movie, touched on the notion of genism—the theory that genes determine distinctive human characteristics and abilities—and explored the idea that genetic discrimination could be as pernicious as overt racism.

THE CASE FOR STRICTER POLICY

Although current legislation provides guidelines for genomic data use, it does not contain enough technical information about safe and secure ways to store and process digitized genomes. One reason is that security and privacy issues for genomic data—both individual genomes and the genome collections in genomic databases—are not well understood.

Privacy practitioners and consumer organizations are strongly advocating the need for more restrictive legislation to close current policy gaps. A recent report from the US Presidential Commission for the Study of Bioethical Issues[19] analyzed WGS advances, highlighted growing privacy and security concerns, and made a few privacy and security recommendations.

We believe these recommendations reflect a general lack of understanding about the associated open technical problems. For example, one recommendation was to use de-identification, which is clearly unsuitable. The recommendations also fail to address several important points. For example, to guard against surreptitious DNA testing, any genomic data protection policy must recognize the need for informed consent. The policy should set forth procedures for authorities and companies to obtain written permission from an individual before collecting, analyzing, storing, or sharing that person's genetic information, such as hair or saliva samples—thus ensuring that no individual will be a victim of unauthorized sequencing.

A measure such as this will not be popular with those who view privacy-friendly measures as hindrances to genomic research. Scientists typically sequence DNA from large groups to determine genes associated with particular diseases. The informed consent restriction would mean that they cannot reuse large genomic datasets to study a different disease. Rather, they would have to destroy the data after each study or track down all previously enrolled study participants and secure a new authorization from each for the next study. Also, because related individuals have similar genomes, the participant's relatives might have to give consent as well.

GUIDELINES FOR GENOMIC DATA PROTECTION AND USE

The individual who requests and likely pays for genome sequencing should own the result, as is already the case for any other personal medical information. However, genomes are a new kind of personal health information, which raises numerous issues that technical approaches alone cannot address. Rather, technology must work with legal and professional guidelines that govern how to transmit, store, process, and eventually dispose of genomic information.

Storage and long-term protection

Storing and protecting the genome raises several important questions:

- Should the genome be stored on the individual's personal device? What special hardware security features are needed to prevent tampering?
- Should genome storage be outsourced to a cloud provider?
- Should the genome be encrypted? If so, what organization will generate and store the encryption keys?

Although encryption might seem the ideal answer to many of these questions, it has drawbacks. Encryption schemes that many consider strong at present might gradually weaken,

but the genome's sensitivity will not. Thus, a third party that cannot decrypt an encrypted genome might be able to do so years later. The Advanced Encryption Standard (AES) scheme supports key lengths up to 256 bits. Although several standardization bodies and intelligence agencies believe it will be secure for several decades,[20] computational breakthroughs or unforeseen weaknesses might allow early decryption.

One option is to periodically reencrypt the genome, assuming it cannot be copied. Another option is to use secret-sharing techniques to split the genome and partition it among several providers. However, efficient reassembly is problematic, as is the guarantee that providers do not collude in genome reconstruction. Moreover, the providers themselves must have sufficient longevity.

Finally, encryption will not prevent leaks of a long-deceased individual's genomic data, which can affect the privacy of that person's living progeny.

> ENCRYPTION SCHEMES THAT MANY CONSIDER STRONG AT PRESENT MIGHT GRADUALLY WEAKEN, BUT THE GENOME'S SENSITIVITY WILL NOT.

Accessibility

Given the genome's sensitivity, an individual should never disclose any genomic information, which would certainly prevent access to any genomic application except within the individual's secured personal device. Although it sounds ideal, such a restriction might be possible if operations were represented in some standardized form that some trusted agency has certified. For example, if testing for a genetic disease requires matching a well-known pattern in some approximate location in the genome, the US Food and Drug Administration (FDA) might certify that pattern and its parameters. Individuals would then be assured that the operation is a legitimate test for a specific genetic disease and that they will receive the results, which they then can opt to keep private.

Other questions about accessibility are more complicated:

- Should the sequencing facility keep an escrowed copy of the genome?
- Should the individual entrust a genome copy to his personal physician or health insurance provider?
- Is it possible to guarantee the digitized genome's integrity and authenticity? If so, how?
- If backups are made, how often and where should they be kept?
- Is it possible to securely erase a genome?
- Should individuals periodically request a new genome sequence to keep pace with more accurate technology?

Testing guidelines

To effectively replace their in vitro counterparts, computational genomic tests must be accurate, efficient, and usable for individuals who are not geneticists.

Accuracy. A computational genomic test should guarantee accuracy that is at least equivalent to the in vitro test. For example, a computational paternity test should provide the same confidence as the in vitro test, which is currently admissible in a court of law. Computational tests should also strive for accountability by furnishing guarantees of correctness for both execution and input information.

Efficiency. Computational genomic tests should incur minimal communication and computing costs. Patients might be used to waiting several days to obtain genetic test results. However, in a computational setting, long runtimes on personal devices might hinder the test's practicality.

Usability. Computational genomic tests are likely to involve the general population, which raises several usability questions:

- How much should the user know about genomic test aspects?
- What information about the test and results is appropriate, and at what granularity should it be presented?
- Do individual's privacy perceptions and concerns match the scientific community's expectations?

The last question is particularly complex. Some users might be willing

to forego their genomic privacy. For example, the expectation is that patients will reveal their genomes to their doctors so that they can benefit from tests that can possibly save them from a life-threatening disease, such as cancer. However, the same individual might not wish to reveal that information to an online service or pharmaceutical company.

These considerations are for the most part educated guesses, since few efforts have focused on users' concerns. Therefore, one research focus should be on exploratory user studies[21] to elicit insights into this issue and address the open problem of how to effectively communicate the potential privacy risks associated with genomic information and its disclosure.

Affordable, readily available WGS will stimulate thrilling opportunities, but it will also raise privacy concerns; addressing both sides of WGS will require long-term collaboration among geneticists, other healthcare providers, ethicists, lawmakers, and computer scientists. To this end, we helped organize the first multidisciplinary Dagstuhl seminar on genomic privacy, which took place in 2013[22] and will be held again in October 2015. We also helped launch an international workshop on genomic privacy, which took place in 2014 and will be held again in conjunction with the 2015 IEEE Symposium on Security and Privacy (www.genopri.org). Finally, we have set up www.genomeprivacy.org, a site that offers computer scientists tutorials and links to genome privacy research groups.

Long-term collaboration will require targeted funding support. In the US, genomic privacy has fallen into funding gap between agencies. The NIH funding, for example, solidly covers both bioinformatics and WGS ethical issues, but only sparsely supports research on genomic data privacy. The National Science Foundation's (NSF's) Smart and Connected Health program includes integrative projects that require collaboration among computer and health sciences, but the program may or may not engender long-range genomic privacy research.

Other US funding agencies have not, thus far, explicitly addressed genomic privacy. In Europe, numerous EU and nationally funded projects are focusing on e-health, and some consider data protection, but they largely overlook genomic data privacy. In addition, although most officials in charge of data protection typically have a strong legal background, they lack computer science expertise. Consequently and not surprisingly, they tend to rely on legislation more than on technology.

Our work is thus a call for research collaboration to specifically and vigorously address the privacy issues we have identified. Overcoming these obstacles will free WGS to reach its full potential to revolutionize medicine and allow individuals and society overall to reap the considerable benefit.

REFERENCES

1. Nat'l Center for Biotechnology Information, "dbSNP," Dec. 2014; www.ncbi.nlm.nih.gov/projects/SNP.
2. Eupedia, "Genetically Inherited Traits, Conditions, and Diseases," 2014; www.eupedia.com/genetics/medical_dna_test.shtml
3. S. Seshadri et al., "Genome-Wide Analysis of Genetic Loci Associated with Alzheimer Disease," *J. Am. Medical Assoc.*, vol. 303, no. 18, 2010, pp. 1832–1840.
4. D.S. Falconer and T.F. Mackay, *Introduction to Quantitative Genetics*, 4th ed., Addison Wesley, 1996.
5. L. Hood and D. Galas, "P4 Medicine: Personalized, Predictive, Preventive, Participatory: A Change of View That Changes Everything," 2009; www.cra.org/ccc/files/docs/init/P4_Medicine.pdf.
6. A. Weston and L. Hood, "Systems Biology, Proteomics, and the Future of Healthcare: Toward Predictive, Preventive, and Personalized Medicine," *J. Proteome Research*, vol. 3, no. 2, 2004, pp. 179-196.
7. A. Burke, "Foundation Medicine: Personalizing Cancer Drugs," 2012; www.technologyreview.com/featuredstory/426987/foundation-medicine-personalizing-cancer-drugs/.
8. My Drug Genome, "Using Genetics to Personalize Medication Treatment," 2014; www.mydruggenome.org/overview.php.
9. K. Whitney, "PREDICT Helps Pinpoint Right Statin for Patient," *Vanderbilt Univ. Medical Center Report*, 4 Oct. 2012; http://news.vanderbilt.edu/2012/10/predict-helps-pinpoint.
10. G. Naik, "Gene Maps Are No Cure-All," *Wall Street J.*, 3 Apr. 2012; www.wsj.com/articles/SB10001424052702304023504577319604245325644.
11. V. McKusick and S. Antonarakis, *Mendelian Inheritance in Man: A Catalog of Human Genes and Genetic Disorders*, John Hopkins Univ. Press, 1994.
12. Y. Erlich and A. Narayanan, "Routes for Breaching and Protecting Genetic Privacy," *Nature Reviews Genetics*, vol. 15, no. 6, 2014, pp. 409–421.
13. N. Homer et al., "Resolving Individuals Contributing Trace Amounts of DNA to Highly Complex Mixtures Using High-Density SNP Genotyping

ABOUT THE AUTHORS

ERMAN AYDAY is an assistant professor of computer science at Bilkent University, Ankara, Turkey. While conducting the research reported in this article, he was a postdoctoral researcher in the School of Computer and Communication Sciences at EPFL, Switzerland. His research interests include privacy, genomics, trust and reputation systems, and network security. Ayday received a PhD in electrical and computer engineering from Georgia Institute of Technology. He is a member of IEEE and ACM. Contact him at erman@cs.bilknet.edu.tr.

EMILIANO DE Cristofaro is a senior lecturer (associate professor) at University College London (UCL). While conducting the work reported in this article, he was a research scientist at Xerox's Palo Alto Research Center (PARC). His main research interests are privacy-enhancing technologies and applied cryptography. De Cristofaro received a PhD in networked systems from University of California, Irvine. Contact him at me@emilianodc.com.

JEAN-PIERRE HUBAUX is a professor in the School of Computer and Communication Sciences at EPFL, Switzerland. His research interests include privacy protection, notably in mobile networks and genomics. Hubaux received a DrEng in electrical engineering from Politecnico di Milano. He is a Fellow of IEEE and ACM. Contact him at jean-pierre.hubaux@epfl.ch.

GENE TSUDIK is a Chancellor's Professor of Computer Science at University of California, Irvine. His research interests include security, privacy, and applied cryptography. Tsudik received a PhD in computer science from University of Southern California. He is a Fellow of IEEE and ACM. Contact him at gts@ics.uci.edu.

Microarrays," *PLoS Genetics*, vol. 4, no. 8, 2008, pp. 1–9.
14. R. Wang et al., "Learning Your Identity and Disease from Research Papers: Information Leaks in Genome-Wide Association Study," *Proc. 15th ACM Conf. Computer and Communications Security* (CCS 09), 2009, pp. 534–544.
15. M. Gymrek et al., "Identifying Personal Genomes by Surname Inference," *Science*, vol. 339, no. 6117, 2013, pp. 321–324.
16. M. Humbert et al., "Addressing the Concerns of the Lacks Family: Quantification of Kin Genomic Privacy," *Proc. 20th ACM Conf. Computer and Communications Security* (CCS 13), 2013, pp. 1141–1152.
17. H. Shen, "California Considers DNA Privacy Law—Academic Researchers Fear Measures Would Prohibit Work with Genetic Databases," *Nature*, 18 May 2012; www.nature.com/news/california-considers-dna-privacy-law-1.10677.
18. Council of Europe, "Additional Protocol to the Convention on Human Rights and Biomedicine, Concerning Genetic Testing for Health Purposes," 2008; http://conventions.coe.int/Treaty/EN/Treaties/html/203.htm.
19. Presidential Commission for the Study of Bioethical Issues, "Privacy and Progress in Whole Genome Sequencing," 2012; www.bioethics.gov/cms/sites/default/files/PrivacyProgress508.pdf.
20. Nat'l Inst. Standards and Tech., "Cryptographic Key Length Recommendation," 2014; www.keylength.com/en/4.
21. E. De Cristofaro, "An Exploratory Ethnographic Study of Issues and Concerns with Whole Genome Sequencing," *Proc. 8th Network and Distributed System Security Symp. (NDSS) Workshop Usable Security* (USEC 2014), 2014; http://arxiv.org/abs/1306.4962.
22. K. Hamacher, J.-P. Hubaux, and G. Tsudik, "Dagstuhl Seminar on Genomic Privacy," Oct. 2013; www.dagstuhl.de/en/program/calendar/semhp/?semnr=13412.

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

COMPUTING PRACTICES

Creating Substance from a Cloud: Low-Cost Product Generation

Adam P. Spring, University of Plymouth

With a more controlled way to capture images, services to generate and process 3D data point clouds, and a 3D printer, amateur photographers can convert 2D photos into physical 3D objects.

In his 2011 best seller, *The Third Industrial Revolution*, economist Jeremy Rifkin explains how design and manufacturing have become decentralized processes that feed into the digital economy.[1] By operating through an infrastructure built on the Internet of Things (IoT), he argues, next-generation information and communications technology (ICT) is replacing factory-based production methods. Delocalized production techniques have developed around computer-aided design (CAD) and manufacturing (CAM), replacing Henry Ford's conveyor-belt assembly line with an anywhere, anytime production line connected through digital devices and ICT.

Affordable 3D printing is one aspect of this transformation, and 3D imaging technologies are an excellent way to understand the Third Industrial Revolution's impact on product development. For example, by implementing more controlled photographic techniques along with an easy-to-use 3D imaging service, anyone with rudimentary photographic skill can turn 2D photos into 3D solid-surface objects.

An application to replicate a cultural heritage monument from a popular 2D camera illustrates the mechanics of this production approach as well as the potential of CAD/CAM practices that align with concepts such as the IoT. The workflow derives from *photogrammetry*, which retrieves measurable data from photos, and uses a service based on *structure from motion* (SfM), a product of computer vision research. However, neither the equipment nor the workflow in the proposed approach is exclusive to experts in those domains; rather, they are accessible to any amateur photographer. The technology is readily available to anyone with access to a camera, the Internet (to work with open source or low-cost user-friendly programs), and a retail 3D printer.

As part of a cultural heritage project to examine connections between Great Britain and Germany during the Industrial Revolution, my colleague Caradoc Peters and I investigated several monuments in both countries. As part of this investigation, I applied the proposed production approach to convert 2D photos of one of the monuments to a 3D physical replication. I used the following items in my work:

- a Canon EOS 60D digital single-lens reflex (DSLR) camera and 50-mm Canon prime lens;

WORKFLOW BASED ON POINT CLOUDS

A point cloud documents a scene in 3D in the same way a photo documents a scene in 2D. Its *z*-coordinate adds depth to the scene collected in 2D that contains *x* and *y* coordinates. Services like Photo from Autodesk ReCap 360 exploit the idea of 3D documentation, adding the *z* coordinate in the 3D generation of conventional 2D photos.

The point cloud signified a move from linear 2D processes to nonlinear workflows centered on a 3D model,[1] and by the late 1990s, a commercial market had started to form around the point cloud as a usable framework in computer-aided design software, similar to AutoCAD.[2] Point clouds can be generated via structure from motion–based services and solutions using little to no calibration. Agisoft PhotoScan (www.agisoft.ru) and Autodesk 123D Catch (www.123dapp.com/catch), as well as Photo, use SfM and multiview stereo (MVS) image-matching solutions to turn 2D photos into 3D point clouds and surfaces.[3] SfM produces usable 3D results because it is looking primarily for matching points in a scene. Photos and images can come from any source, from a digital single-lens reflex (DSLR) camera to a search engine, and even crowdsourcing from multiple cameras and users worldwide.[3]

3D imaging technologies now let users collect as-built information at unprecedented resolution and scale. Because of this paradigm shift, it has become easier to document a building or cityscape with subcentimeter accuracy.[2] This ability to collect in situ information with 3D imaging technologies is an important part of a cultural heritage workflow that captures an artifact's or site's remains without risk of contamination. 3D imaging technologies give heritage practitioners the ability to examine sites or artifacts in the same amount of detail as a forensic scientist would a crime scene.

In terms of Jeremy Rifkin's *Third Industrial Revolution*, the point cloud is a visual medium and an agent for social and economic progress. It is comparable to the steam printing press or early forms of electronic communication in that it can transmit an idea or provide a better sense of context. The Building Information Model (BIM), for example, is an infrastructure and asset management strategy that uses point clouds to create intelligent objects linked to numerous forms of information.[4] Industry cultivates point cloud use through workflow models based on information and communications technology (ICT). The Rip-Mod-Fab or Capture-Compute-Create workflow models employed by Autodesk are good examples.

Figure A shows one such model, in which services and apps replace centralized fixed production modes. Users with no specialty skills collect scenes through sensors in inexpensive equipment and send them to cloud-processing services, export and refine them through low-cost or open source solutions, and disseminate results through an ordinary mobile device or 3D print out.

In this new workflow model, data informs design and manufacturing, and performance—the way in which the user and equipment interact with the object being scanned—dictates practices. Performance in this context is important because even

- Photo—an Autodesk ReCap 360 service that creates 3D point clouds and solid surface meshes from photos (https://recap360.autodesk.com);
- CloudCompare, a program that processes the resulting 3D point cloud and solid surface mesh (www.danielgm.net/cc);
- MeshLab, a program to process and edit unstructured 3D triangular meshes (http://meshlab.sourceforge.net);
- MeshMixer, an Autodesk 123D–related tool for 3D mashups and remixes from custom 3D designs or provided models (www.123dapp.com/meshmixer); and
- a Dimension 1200es 3D printer from Stratasys (www.stratasys.com/3d-printers/design-series/dimension-1200es).

This list suggests that point clouds—data points in 3D space—and associated solid-surface meshes could become the new template for design and manufacturing, as the sidebar "Workflow based on Point Clouds" describes.

DOCUMENTING THE SUBJECT

The subject of the sample application is a steam cylinder used in a Watt engine water pump that once serviced copper mines in the Saxony-Anhalt region of Germany and is situated in Löbejün. The Mansfeld Museum in nearby Hettstedt houses a replica of the engine that contained the cylinder. Figure 1 shows the monument and its proximity to the museum.

The cylinder was cast in 1788 in Glamorganshire, Wales, and later exported to Germany, where it formed part of that country's first steam-powered pumping engine and was in use until 1848. In 1934, the cylinder was mounted on a pedestal to commemorate Germany's first use of steam power.

Documenting the cylinder's history is part of a cultural heritage project that is examining the impact of British

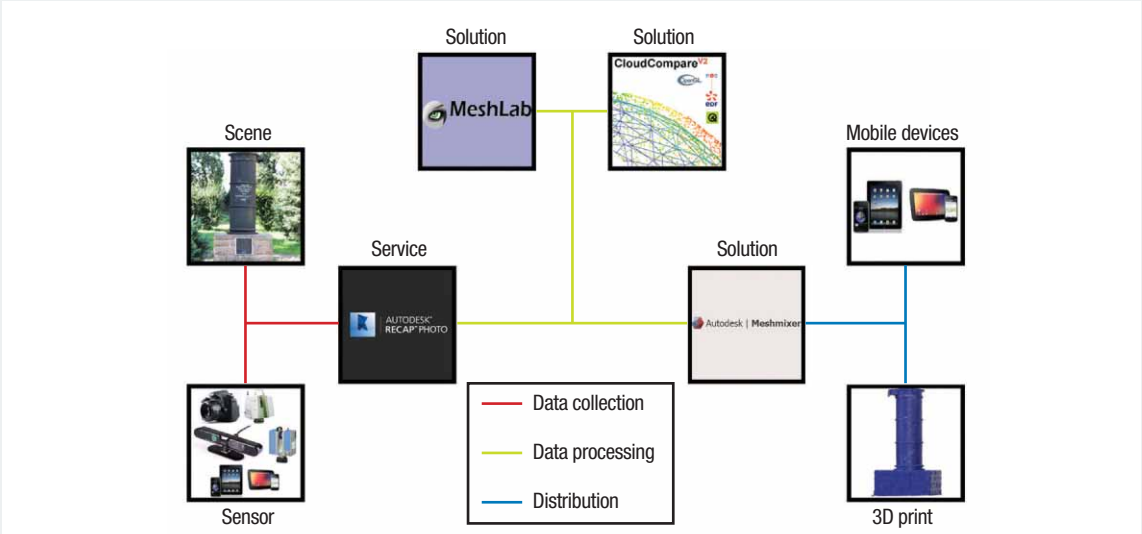

Figure A. Service- and solution-based 3D production line. Collected data enters a service to convert 2D images to a 3D point cloud and mesh, which solutions like CloudCompare and MeshMixer refine. Users can then export and process data in a form suitable for a 3D printer or mobile device.

individual users never produce exactly the same set of data every time.[5] The Löbejün monument workflow followed this model in that the preservation state informed the production and reproduction cycle. The only modifications were to delete unwanted data and fill any gaps in the generated 3D mesh; both steps were necessary to reproduce the monument as an authentic 3D printed facsimile.

References

1. A.P. Spring, "The Third Industrial Revolution," *GeoInformatics*, vol. 15, no. 7, 2012, pp. 32–34.
2. A.P. Spring, C. Peters, and T. Minns, "Using Mid-Range Laser Scanners to Digitize Cultural Heritage Sites," *IEEE Computer Graphics and Applications*, vol. 30, no. 3, 2010, pp. 15–19.
3. D. Crandall et al., "SfM with MRFs: Discrete-Continuous Optimization for Large-Scale Structure from Motion," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 35, no. 12, 2013, pp. 2841–2853.
4. G.A. Van Nederveen and F.P. Tolman, "Modeling Multiple Views on Buildings," *Automation in Construction*, vol. 1, no. 3, 1992, pp. 215–224.
5. C. Anderson, *Makers: The New Industrial Revolution*, Random House, 2012.

mining technologies and workforce emigration in Germany in the late 18th to early 19th centuries. The project's goal is to record site features and artifacts in a way that merges geographically disconnected information and creates a single channel for disseminating results. It is one of many projects that are exploiting the new relationship between IoT and 3D imaging technology. The sidebar "Technology, 3D Imaging, and the Internet of Things" describes this relationship in more detail.

DATA CAPTURE

The photogrammetric specifications for documenting the monument, shown in Figure 2, came from Adam Technology (www.adamtech.com.au), an Australian company that converted its analog stereo plotter solution into a software package called 3DM Analyst in the mid-1990s. The specifications are for a camera using the Advanced Photo System type-C (APS-C) image-sensor format—specifically, the Canon EOS 60D—with focal length and aperture considered alongside their equivalents for a full-frame 35-mm camera. APS-C is approximately equivalent in size to negatives of 25.1 × 16.7 mm (aspect ratio of 3:2), which is much smaller than the 36 × 24 mm for standard 35-mm film.

Sensor size and type are directly related to pixel size and resolution and thus help determine the optimum resolution for creating a 3D point cloud or solid surface mesh from the camera's output. Because most standard lenses are designed for full-frame camera sensors, when a DSLR camera sensor is cropped, the lens' focal length increases.

Figure 3 shows adjusted Canon EOS 60D settings to improve subsequent 3D image translation. Photogrammetrists refer to this process as *interior orientation*—linking known parameters to the sensor. For the Löbejün monument shoot, I put the camera's APS-C sensor in aperture priority mode (AV) with auto rotate switched off and used an

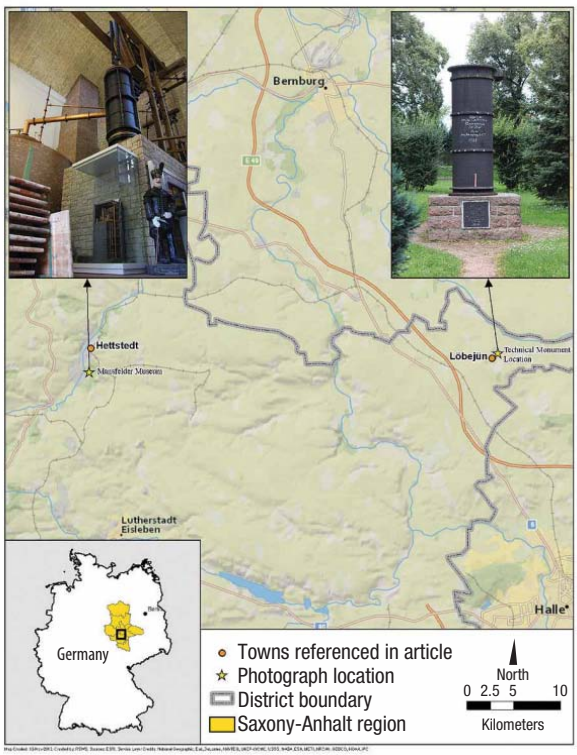

FIGURE 1. Steam cylinder monument in Löbejün, Germany, and replica of the steam engine that contained the cylinder at the Mansfeld Museum in Hettstedt.

ISO setting fixed to Canon's ideal sensor setting of 100 for sunny daylight. I set the aperture to 5.1—the equivalent of 7 or 8 on a full-frame camera—and, after automatically adjusting the focus, switched the focus setting to manual.

Although Photo automatically calibrates the images, regulating pixel information according to variables like surface reflectance and depth of field improves the quality of the meshes generated.[2]

The need for such regulation stems from the difference between photogrammetry and computer vision. Photogrammetry uses known distances and variables to derive metric value from images. 3DM Analyst, for example, creates 3D models by evaluating each stereo pair. In contrast, solutions based on computer vision concepts, such as SfM, generate 3D projections from photos using image resolution, the shapes within the image, and bundle adjustment based on multiview stereo scene reconstruction.[3] SfM-based generation does not require calibrated images,[2] and, in this respect, it is in line with the solid-surface modeling techniques in computer graphics.[4] Before it crossed into surveying-based applications, SfM was popular in disciplines such as robotics as a relative navigation tool that used scene reconstruction to guide vehicles or robots.

In photogrammetry, camera settings and shooting conditions are the camera model's interior and exterior orientations. Each camera model links known variables, such as focal length, to the creation of multiple photos, or photo batch. Photogrammetry adjusts the camera lens and sensor settings according to the scene's environmental conditions. For example, the base-to-distance ratio (the distance between each image in relation to the camera's distance from the object) can improve the quality of a 3D projection from a photo batch.

SfM-based solutions like Photo are similar to photogrammetry software like 3DM Analyst in that the 3D production cycle is linked to the camera location's six degrees of freedom: x, y, z, omega (ω), phi (φ), and kappa (κ).[3] Because it is adjustable, scale is arguably the seventh degree of freedom in a camera-based 3D workflow.

Observance of the common camera parameters in Table 1 and adherence to the shooting recommendations in Table 2 will improve image continuity in an SfM-based projection.

GENERATING THE 3D OBJECT

I photographed the monument moving left to right in a circular pattern, making sure that the depth of field was consistent across scenes and that I documented as much of the monument as possible. Having some of the surrounding area in each photo improves the likelihood of successful batch processing in Photo by matching surfaces for an otherwise geometrically regular object to the correct viewpoint and image.

I ran 102 images through Photo along with measurements from placing two standard rulers on the monument. The goal of this extra step was to add metric information to the scene through Photo's registration function in the Advanced Tools setting (Photo-to-3D interface), which also allows users to choose the photos for adding texture to the surface of the newly created 3D mesh.

Before running images through Photo, users can save the resulting 3D scene in the .rcm, .obj, .fbx, .rcs, or .ipm format. I chose to export the Löbejün monument scene in .rcm and .obj at maximum resolution. Photo took 2.8 hours to generate the scene and upon completion sent an email notification to my account.

PROCESSING AND EXPORTING THE SCENE

Figure 4 shows the 471-Mbyte .obj file that contained a 3D mesh consisting of

Camera Name:	Canon 60d				
Camera Details	**Width**	**Height**			
Number of pixels:	5184	x 3456		Image size:	17.9 megapixels
Image sensor dimensions:	22.3	x 14.9	mm	Field of View Crop/Lens multiplier:	1.6 x 1.6
Actual focal length of lens x adapter:	50	mm		Equivalent 35mm camera focal length:	81 mm
Actual aperture:	f/ 5.1			Equivalent 35mm camera aperture:	f/8.2
Focus distance:	3.00	m		Depth of field:	2.82m - 3.21m
Desired circle of confusion (diameter):	2.5	pixels		Hyperfocal distance:	45.63 m
Size of each pixel in CCD array:	4.30	x 4.31	um	3DM Analyst camera calibration settings:	0.00430 x 0.00431

FIGURE 2. Canon EOS 60D digital single-lens reflex (DLSR) camera specifications linked to photo rectification and photo resolution in accordance with the Adam Technology shooting framework in 3DM Analyst.

FIGURE 3. Adjusted Canon EOS 60D settings to improve 3D image translation. Using these settings maintains continuity between each image and can thus improve the resulting 3D projection.

TABLE 1. Common camera parameters and their relation to 3D projection.

Resolution	Depends on sensor and lens	Environmental factors can affect photo resolution, such as lighting, surface reflectance, and weather
Focal length	Measure from lens to sensor	For each shoot, maintain same focal length and focus setting to rectify the photo batch and ensure undistorted images
Sensor size	Depends on camera	Affects 3D projection resolution
Lens	Each lens has a perspective center that determines camera position	Better information is collected at the center of the lens—this should be considered when shooting stereo pairs, as should angle of incident of scene shot in relation to the object
Auto rotate	Automatic rotation	Disable to lock down photo's *x* and *y* coordinates
Pixel size	Optimum resolution for collecting data	The Canon EOS 60D has a photodiode pitch of 4.3 µm, which is a resolution linked directly to its internal sensor
Native ISO setting	Depends on sensor	Shooting at this setting in ideal lighting reduces photo's digital noise
Aperture size	Depends on application	Affects depth of field

TABLE 2. Photo-shoot recommendations to improve 3D projection.

Base-to-distance ratio	Manipulating the ratio adds scale to a photo shoot; for example, a scene 1 km away from the camera has a 1:1 scale if each camera position is 1 km apart
Scene composition	Placing targets in a scene provides scale and known geometry; make sure as much of the object or scene as possible is in each photo
Capture mode	Using a monopod, tripod, or automatic trigger can improve a photo shoot; for example, all these items can help blur reduction in low-light shooting
Scope	Where possible, take each photo to ensure a 60 percent overlap; take stereo pairs in landscape and portrait views with auto rotate switched off

FIGURE 4. 3D scene of the Löbejün monument in Photo. The scene represents a triangular mesh of 5.63 million triangles. Users can view a generated mesh at reduced resolution before exporting it as a .rcm, .obj, .fbx, .rcs, or .ipm file.

FIGURE 5. The Löbejün monument printed on a Dimension 1200es 3D printer from Stratasys. The final print was 19 cm high with a 10 × 10 cm base.

5.63 million triangles. I ran the .obj file in CloudCompare and MeshLab before exporting it to MeshMixer as an .stl file.

CloudCompare

CloudCompare began as part of a PhD project at Télécom ParisTech in 2004 to improve 3D imaging workflows in industrial plants in France and became open source software in 2009. It uses an octree file structure that splits large datasets into eight manageable interconnected sections. Users can then highlight a particular section and remove redundant or otherwise unwanted data as needed.

For the Photo scene in Figure 4, CloudCompare pared the file to 97.1 Mbytes. I simply chose the Segment function and deleted unwanted data outside the highlighted area. If datasets prove too large, users can process them through Photo at a lower resolution. An alternative is to export the datasets to CloudCompare through Photo at the maximum resolution and then select CloudCompare's meshing function to reduce the number of triangles or segment the file. Autodesk is also beta-testing Project Memento (http://labs.autodesk.com/utilities/memento), a meshing solution for larger files.

MeshLab

MeshLab is open source 3D modeling software that originated as a course assignment at the University of Pisa in 2005. Although it has difficulty handling more than 300 Mbytes of 3D data, its functions and file export options make it a useful processing tool. Among these options are the ability to export files in .stl, the standard 3D printing format, and to present 3D data on a smartphone or tablet running either Android or iOS (available from Google Play or iTunes, respectively).

MeshLab exported the 97.1-Mbyte CloudCompare .obj file as a 47.3-Mbyte .stl file, which was then ready for the software that would prepare it for 3D printing.

MeshMixer

MeshMixer, developed at the University of Toronto, is an Autodesk tool for working with unstructured polygonal meshes, providing functions such as mesh cleanup and surface sculpting and modeling, including surface extrusion.

I imported the MeshLab file of the Löbejün monument into MeshMixer primarily to fix holes in the mesh—essential to printing the artifact in 3D. MeshMixer automatically highlights any holes, which users can fill by selecting the Analysis function and clicking on Inspector. MeshMixer identifies holes as various colored balls, which

TECHNOLOGY, 3D IMAGING, AND THE INTERNET OF THINGS

Since its conception as the Parliament of Things in 1991, the Internet of Things (IoT) has slowly penetrated mainstream computer science to the point that researchers are beginning to see it as the foundation of the Third Industrial Revolution; indeed, Cisco Systems recently dubbed it the Internet of Everything. Despite its seeming rebranding, however, the underlying philosophy—humans and objects are independent agents in knowledge distribution and communication—has remained constant.[1]

The IoT acknowledges that people interpret, synthesize, and communicate their perceptions of the world through any device at hand. It is part of the Symbiotic Web, or Web 4.0, which describes how connectivity via the Internet makes it easier to crowdsource and network data. Workflows such as that used for the Löbejün monument can help people make better sense of the considerable data now at hand.

Because information generated in this manner is subject to human error, processes like 2D and 3D image production always have room to improve. Smartphone and tablet technologies, for example, have made it easier to apply the IoT to 3D imaging.[2] Cloud-based services like Photo leverage the IoT and Third Industrial Revolution concepts to make ICT easier to understand and use, supporting notions such as as-built and as-designed information. By eliminating the parts of a process that used to be accessible only to domain experts or large manufacturers, services like 123D Catch and Photo are reshaping computer-aided design and manufacturing on a daily basis.[3]

At the same time, data is becoming increasingly easier to collect, use, and distribute, which is creating gray areas in this industrial revolution. Cultural heritage documentation is one such area. Solutions such as those used to document and replicate the Löbejün monument impose structure on this freely obtained data—in this case, by assigning parameters to 2D images that make them easy to replicate in 3D. At the same time, the workflow emphasizes affordability and adaptability and represents the best fit for the given application. Other applications might use different parts of the workflow, depending on the services and solutions desired.

Attempts to impose structure are also evident in numerous product refinements. 123D Catch appeared in May 2011 as a desktop-based portal that sent photos to the cloud for processing and in 2012 extended its services to smart technology users through an iOS application. Photo's vendors released it in 2013 as ReCap Photo to serve as the professional equivalent of 123D Catch. By February 2014, they had refined ReCap Photo to support GoPro cameras, texture and mesh display, smart cropping, and boundary selection, and they had given it a new user interface.

Affordable depth cameras from companies like Prime Sense, makers of the active triangulation sensors inside Kinect for Xbox 360, have enabled 3D imaging, point clouds, and solid surface meshing for a broad range of users. Acquired by Apple in November 2013, PrimeSense also supports middleware and application development through OpenNI, software development kits, and websites.

References

1. B. Latour, *We Have Never Been Modern*, Harvard Univ. Press, 2012.
2. K. Ashton, "That 'Internet of Things' Thing," *RFID J.*, 22 June 2009; www.rfidjournal.com/article/view/4986.
3. J. Chandler and J. Frier, "Autodesk 123D Catch: How Accurate Is It?" *Geomatics World*, 2013; www.pvpubs.com/archives/index.php?article=1326&magazine=205&search=.

users can see and resolve by double clicking on the Auto Repair All option.

The Löbejün monument mesh had several holes, including a large one in the monument base, which MeshMixer represented as a red ball. Although Auto Repair All worked fine in this case, I could have also filled the base by clicking on the red ball, selecting Edit, and clicking on Erase and Fill. The program would then insert a predetermined shape from the selection on the screen's left side. Users can adjust the shape's form and size by reducing it in the Tools Properties window or smooth the shape by clicking on the Modify Selection and Smooth Boundary options.

Finally, users can extrude flat surfaces through the Edit function, which is useful in printing otherwise flat or paneled scenes, such as wall inscriptions or rock art.

The latest version of MeshMixer also links to Autodesk's 3D printer utility (http://apps.123dapp.com/3dprint/install.html), which is compatible with affordable systems like MakerBot, Stratasys's Objet500 Connex, and Objet Geometries' Alaris30. In May 2014, Autodesk also launched Spark (http://spark.autodesk.com), its open source platform for developing 3D printers. The first system to come out of the Spark

ABOUT THE AUTHOR

ADAM P. SPRING is a consultant and visiting lecturer in applied technologies and reality capture in the Department of Archaeology at the University of Plymouth, UK. His research interests include the user experience, 3D information capture, and multidimensional visualization and documentation. Spring received an MS in landscape archaeology from the University of Bristol, UK. He is a member of the ACM Special Interest Group on Computers, Information, and Society (SIGCIS). Contact him at adam@remotely-interested.com or through www.remotely-interested.com.

initiative was Ember. Other open source 3D printing projects include RepRap (http://reprap.org), which came out of the University of Bath, UK, in 2005.

3D PRINTING

I imported the .stl file into the CatalystEX operating software of the Dimension 1200es 3D printer and manually input surface dimensions. Figure 5 shows the output.

Although a 3D printer was readily available, I could have used other services to create the replica. For example, Autodesk's 123D Make (www.123dapp.com/make) can convert an .stl file into slices that a regular printer can print on cardboard. Online services like Shapeways (www.shapeways.com) also let users upload a file onto the company's server, specify the print medium from various materials, and request shipping. Users of any 3D print program can disseminate their results on community websites such as MakerBot's Thingiverse (www.thingiverse.com).

Tools like Autodesk ReCap 360, CloudCompare, MeshLab, and MeshMixer are enabling low-cost, easy-to-use 3D imaging and CAM solutions based on data that can come from any source or location and from users with a variety of skills. Leveraging technology–workflow combinations such as that used to document and replicate the Löbejün monument, any camera, smartphone, or tablet can reproduce an artifact's actual preservation state. Whether it is through crowdsourced or networked data from the IoT, as-built plans, nonlinear models, or 3D prints, strategies to structure increasing dataflow must become part of creating products in this collaborative era of document, design, and manufacture.

In this respect, the Third Industrial Revolution feeds into the IoT. Both are being presented to general audiences as user-friendly terms—even in the context of 3D imaging and manufacturing. Fundamentally, however, they are used to describe changes that fall in line with the Symbiotic Web or Web 4.0. Internet evolution and the way in which people interact with digital technologies are ultimately freeing up design and manufacturing processes to a wider population, shifting focus from centralized production modes to more dispersed collaborative efforts, which are often the source of true innovation.

REFERENCES

1. J. Rifkin, *The Third Industrial Revolution: How Lateral Power is Transforming Energy, the Economy, and the World*, Palgrave Macmillan, 2011.
2. R.I. Hartley and J.L. Mundy, "The Relationship between Photogrammmetry and Computer Vision," *Proc. Int'l Conf. Optics and Photonics: Integrating Photogrammetric Techniques with Scene Analysis and Machine Vision* (SPIE 93), vol. 1944, 1993; http://users.cecs.anu.edu.au/~hartley/Papers/SPIE-93/joint-paper/joint2.pdf.
3. N. Matthews and T. Noble, "Aerial and Close-Range Photogrammetric Technology: Providing Resource Documentation, Interpretation, and Preservation," Bureau of Land Management Tech. Note 428, 2008; www.blm.gov/nstc/library/pdf/TN428.pdf.
4. M. Levoy, "The Early History of Point-Based Graphics," *Point-Based Graphics*, M. Gross and H. Pfister, eds., Morgan Kaufman, 2007; http://graphics.stanford.edu/papers/points/levoy-pointbook-ch2.pdf.

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

STUDENT DESIGN SHOWCASE

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

Spotlighting Student Innovation

Greg Byrd, North Carolina State University

This new column provides a space for undergraduates in computer engineering and science to share their capstone project designs.

Welcome to the inaugural installment of a new bimonthly column for *Computer*. Student Design Showcase is intended as a venue for young computer engineers and scientists to show off their stuff. As column editor, my goal is to act as a conduit between these creative budding professionals and the broader IEEE Computer Society membership. To this end, I'll mostly get out of the way and let students tell their own stories.

FOCUS ON SENIOR DESIGN PROJECTS

A primary focus of the column will be the capstone project of the undergraduate computing curriculum, generally called "Senior Design." With this project, students take one or two semesters to pull together everything they've learned and apply their knowledge to a substantial design problem. Along the way, they practice other skills that will serve them well in the workplace: project management, group dynamics, and communication, to name a few.

Senior Design experiences vary as widely as the schools that offer them. In many cases, local industries provide projects and mentors. In others, faculty members sponsor such projects as an extension of, or introduction to, their own research. Sometimes, student projects are inspired by a desire to benefit society; the National Academy of Engineering's Grand Challenges contests, for example, encourage participants to show how "engineering will create a more sustainable, healthy, secure, and/or joyous world in the future" (www.engineeringchallenges.org). And in programs that emphasize entrepreneurship, students may be required to propose a product themselves, then research the market and design and build a prototype.

Whatever their structure, we want to honor the hard work, creative juices, and long caffeinated hours that go into these projects.

SOME GROUND RULES

I have only a few ground rules about submitting projects:

- The design and implementation parts of the projects must be performed by students only. The idea can be generated elsewhere, and professionals can serve as mentors and advisors, but the bulk of the work must be done by students.
- As noted, I'm primarily interested in undergraduate student work. I won't reject an interesting graduate project out of hand, but that's not my focus for this column.

SUBMISSION INFORMATION

To submit a project for consideration, visit the Student Design Showcase page at https://computingnow.computer.org/web/computingnow/computer/student-showcase. There you'll find a summary of the requirements, along with a link to a submission form. The form is easy to complete—just some basic information and a brief description of the project.

Once I've selected a project, I'll ask the author or authors to provide a detailed design document and be willing to collaborate actively over the following three or four weeks, with me and *Computer*'s editorial staff, to produce the final column.

- The project must include a significant computational component. It can be software only, or a combination of software and hardware, but this is a magazine for the Computer Society, and we want to see some computing!
- The project must be complete (preferably, within the past year) or nearly complete at the time it's submitted.
- The submission must be initiated by the student team, and at least one team member must be willing to be the primary contact.
- While Computer Society membership isn't a requirement, I certainly encourage it.

Given a set of submissions, how will I choose the projects to showcase in the column? Here are some criteria:

- Most importantly, the problem, the solution, or both must be interesting to *Computer* readers.
- Because I don't want projects all from the same institutions or in the same disciplines, I'll consider geographic and topical diversity in the selections.
- I'll favor projects that have some multimedia content: photos, videos, animations, software demos, and the like.

CALL FOR PROJECTS

So, students: I want to hear your stories, your inspirations, your challenges, and your innovative solutions. This is a technical magazine, so feel free to "let your geek flag fly"—give us some details, and show us old fogeys the skills and tools we need to learn to keep up with you.

And, members: spread the word to your students, colleagues, interns, and university contacts. This column will be driven by submissions, and I need your help to beat the bushes for the most creative and interesting design projects.

See the sidebar for submission information.

Student Design Showcase will begin in earnest with the April issue (sneak peak: elephant collars), and will appear every other month thereafter. I look forward to hearing about lots of innovative designs—and sharing the best of them with *Computer* readers.

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

OUT OF BAND

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

Cyber Chutzpah: The Sony Hack and the Celebration of Hyperbole

Hal Berghel, University of Nevada, Las Vegas

There's nothing about the recent Sony hack that withstands close scrutiny. The story began as bunk, took a spin around blather and hooey, and then seems to have come to rest on drivel.

I've restrained myself from commenting on the Sony hack until now. This entire story has been stuck on stupid, but after Sony CEO Kazuo Hirai's underwhelming talk at the 2015 Consumer Electronics Show in Las Vegas (my fair city), I can hold back no more. It's time to pull what little common sense is left of this story out of the Orwellian memory hole and try to get the narrative back on track.

Hirai said, "[Sony employees] were unfortunately the victims of one of the most vicious and malicious cyberattacks that we've known certainly in recent history .... And I have to say that freedom of speech, freedom of expression, freedom of association, those are [the] very important lifeblood—lifelines—of Sony and our entertainment business" (http://time.com/3655462/sony-chief-executive-hacking). This is hyperbole and drama befitting a Mickey Spillane novel—the Sony hack is *not* in the upper echelon of cyberattacks! It's not even in the second or third tiers. As a matter of fact, apart from the embarrassing executive emails that were leaked, it's not even very interesting.

Furthermore, if Sony really believed in freedom of expression, it wouldn't have fired its corporate communications executive Charles Sipkins over an alleged snub of cochairman Amy Pascal (www.rttnews.com/2430666/sony-executive-leaves-after-e-mail-reportedly-sought-his-firing.aspx). Sony's corporate stance on this offends the senses.

In the grand scheme of things, the Sony hack seems to be a rather pedestrian compromise of a security-challenged computer network. Examples of "vicious and malicious cyberattacks" are easy to find: consider the Trojan horse software hack by the US that led to the 1983 Trans-Siberian Pipeline explosion—reportedly the largest non-nuclear explosion in recorded history.[1] Now *that's* vicious and malicious.

Or, one might point to the Operation Olympic Games attack that used the Stuxnet worm to destroy uranium centrifuges at an Iranian fuel enrichment facility (www.nytimes.com/2012/06/01/world/middleeast/obama-ordered-wave-of-cyberattacks-against-iran.html). Once again, this qualifies as a vicious and malicious cyberattack. Since both of these examples involve cyberkinetic attacks on sovereign nations, they remain politically charged, so we'll pass over the geopolitical motives in silence.

For something to qualify as vicious and malicious, an action must have consequences that are savage, brutish,

violent, or fatal. Detestable and spiteful conduct usually won't qualify. Attacks against sovereign nations? Yes. Hacks of corporate computer networks? Not so much. The Sony hack is closer to MafiaBoy, the Google Gmail hack, the Solar Sunrise hack, and Albert Gonzalez's compromise of T.J. Maxx and Heartland Payment Systems than it is to the Siberian and Stuxnet examples. It's just another installment in the never-ending evolution of digital crime.

There's plenty of wiggle room in the continuum of state-involved criminal activity: state-sponsored, state-proxied, state-tolerated, state-aware, kleptocratic, narco-kleptocratic, and so on. But we need to be circumspect when we start assigning these tags to the countries involved. We didn't threaten and sanction Nigeria for its connection to the Nigerian 419 phishing scams, nor did we threaten and sanction Russia for the Gameover ZeuS botnet and Cryptolocker ransomware, even though both countries knew, or should have known, that these cybercriminal activities took place on their soil.

WHAT DO WE KNOW AND WHEN DID WE KNOW IT?

So why was Sony targeted? We've been led to believe that it was the North Korean supreme leader's reaction to the plot of the Sony motion picture *The Interview*, which involves the assassination of Kim Jong-un. By most accounts, the perpetrator is an anonymous hacking group called the Guardians of Peace, which is speculated to be a cyberattack group acting on behalf of the North Korean government. But if this *were* a North Korea–sponsored hack, wouldn't they have instructed their agents to conceal this connection? To borrow a phrase from Thomas Hobbes's *Leviathan*, history has shown that the lives of the perpetrators may become "nasty, brutish, and short." History has also shown that when nation-states are involved in cyber-conflicts, any clues left behind are most likely false flags. Over the past 65 years, the US Central Intelligence Agency has shown the entire global community the value of plausible deniability.

I'm not saying that Kim Jong-un is incapable of cyberwarfare. But how much would he gain by drawing attention to himself over an ego-motivated incident like this? This doesn't seem to be a sensible occasion for a "nana-nana-boo-boo" moment.

Let's look at the reported evidence. The US Federal Bureau of Intelligence (FBI) initially reported that North Korea was the likely source of mischief (www.politico.com/story/2014/12/fbi-briefed-on-alternate-sony-hack-theory-113866.html). But the time stamps of some of the recovered files showed that downloads might have been done at USB speeds, suggesting an inside job (www.4thmedia.org/2014/12/breaking-we-can-conclusively-confirm-north-korea-was-not-behind-sony-hack). The FBI then revised its account to suggest that the North Koreans may have subcontracted freelance hackers to do their bidding (http://in.reuters.com/article/2014/12/30/northkorea-cyberattack-idINL1N0UD1IB20141230). So, the source and rationale at this point seems to be a moving target. However, FBI Director James Comey still holds firm that North Korea must somehow be to blame. When asked upon what solid evidence this hypothesis is based, we get the time-worn shibboleth "trust me."

Consider two of Comey's statements in a recent *Wired* article (www.wired.com/2015/01/fbi-director-says-north-korean-hackers-sometimes-failed-use-proxies-sony-hack). He initially states, "I want to show you, the American people, as much as I can about the why, but show the bad guys as little as possible about the how. ... This will happen again and we have to preserve our methods and our sources." Then, in an effort to neutralize the critics, he says, "They don't have the facts that I have. They don't see what I see."

First, let's deal with the issue of how the FBI came to "know" what it claims. According to *Wired*, "Comey now says that the hackers in the attack failed on multiple occasions to use the proxy servers that bounce their Internet connection through an obfuscating computer somewhere else in the world, revealing IP addresses that tied them to North Koreans." Really? Are we to believe that hackers with the full financial and military backing of the North Korean government—the same government that has resources enough for a missile program (www.bbc.co.uk/news/world-asia-17399847)—doesn't have sufficient resources to hire competent hackers who know how to spoof IP addresses and use proxy servers? Does this sound reasonable to you? Script kiddies know this much!

If this is true, Kim Jong-un is getting ripped off by his cybermercenaries. Such claims should be viewed with considerable suspicion. Also, I have no idea what, if anything, Comey means by "preserving our methods and sources," if it doesn't involve subpoenas and warrants. The technical "methods" for analyzing network attacks are taught in SANS (www.sans.org) classes. Any claim that FBI network forensics specialists have a monopoly on network traffic analysis is preposterous.

As for the "facts," I seriously doubt that Comey did the network traffic analysis himself. The facts in his possession would probably be better characterized as reportage. Perhaps it might have been more accurate for Comey to say, "The summary that was presented to me [by ...] seems compelling." But I think the suggestion that Comey has possession of and is in a position to interpret ground-truth data is a bit of a stretch. I, for one, would feel much more comfortable relying on the opinions of those who have appropriate backgrounds in digital forensics. For all we know, Comey is making representations that have been filtered by layers of mid-level management with little or no understanding of the technological issues, or, worse yet,

through political filters to ensure that the leadership stays on message. Recall also that Iraq's supposed possession of weapons of mass destruction, uranium yellowcake from Niger, aluminum tubes for centrifuges, and the Prague connection with Al-Qaeda were all reported as certainties.

That said, unlike some of the other leaders of the military–industrial–intelligence community, Comey is a bureaucrat. He was the deputy attorney general that appointed Patrick Fitzgerald to investigate the outing of Valerie Plame as a covert CIA officer (a violation of federal law). Nothing much came from the investigation (note that Scooter Libby's sentence, resulting from his conviction for making false statements and obstructing justice, was commuted), but we can't fault Comey for that.

Comey also refused to re-certify the National Security Agency's (NSA's) domestic bulk metadata collection program in 2004, which sent shockwaves through the White House. Comey, along with Attorney General John Ashcroft, Assistant Attorney General Jack Goldsmith, FBI Director Robert Mueller III, and others, threatened to resign if George W. Bush didn't bring the NSA's program in line with the law.[2,3] Again, nothing much came of this due to subsequent decisions by the FISA (Foreign Intelligence Surveillance Act of 1978) Court and the passage of the 2007 Protect America Act. But in both cases, Comey et al. positioned themselves on the right side of history, at least in terms of these issues. So let's try to give Comey the benefit of the doubt (although he's making it difficult with his pronouncements).

Doubts about the North Korean connection aren't without substance (http://marcrogers.org/2014/12/21/why-i-still-dont-think-its-likely-that-north-korea-hacked-sony; http://blog.norsecorp.com/2014/12/29/ex-employee-five-others-fingered-in-sony-hack; and www.theatlantic.com/international/archive/2015/01/we-still-dont-know-who-hacked-sony-north-korea/384198). Bruce Schneier also has links to relevant data on his blog (www.schneier.com/blog/archives/2014/12/more_data_on_at.html).

Of course, if a connection between an adversary and a hostile act is never proved, bureaucracies might appeal to cognitive dissonance theory and confirmation bias. Taking this into account, the "absence of evidence is evidence of clever deceit." Logicians refer to this as a variety of "the argument from ignorance." However you wish to characterize the phenomena, it has been used masterfully for 50 years by neoconservatives—for example, Team B's claims of Soviet economic and military superiority while the country was imploding, and Donald Rumsfeld's dismissal of the failure to find weapons of mass destruction in the second Iraq war as irresponsible impatience by the media. Don't be surprised to see this kind of illogical belief perseverance resurface again in this context.

> Accusing attribution during an ongoing investigation is like painting falling leaves: the results are sloppy and unlikely to have enduring value.

EMAIL PROPRIETY 101

Some of you are old enough to remember the first principle of email propriety: don't include things in email that you're not willing to post on your office door. Apparently, some of Sony's ill-mannered executives never embraced this refrain. A choice selection of leaked email from Sony co-chair Amy Pascal and producer Scott Rudin were found to be injudicious and of questionable taste. (A summary timeline may be found at www.usmagazine.com/celebrity-news/news/sony-hack-key-events-from-leaked-emails-terror-threats-20141812.) Could it be that entertainment executives are occasionally petty, imprudent, and ill-tempered? Color me surprised! For over a century, entertainment executives have given substance and form to the expression "warm, caring, sensitive, and fair-minded"—especially when dealing with talent (actors, directors, artists, screenwriters, and the like). So who would have thought that an occasional racist thought might creep into their light-headed correspondence? Why, even a cursory review of the list of Academy Award winners will dispose of any thought of bias and discrimination in Hollywood. There's no more minority or gender bias in the entertainment industry than in, say, professional sports or politics, for goodness' sake. And no less, either! There's nothing remotely newsworthy in the leaked email that I can see. Gossipy? Yes. Newsworthy? No.

Now, if I haven't yet convinced you that this story is stuck on stupid, I've got a hole card. Politicians and bureaucrats pushed the story over the event horizon of dumb. First, President Obama made accusations that were apparently based only on the fungible intelligence I mentioned. These days, such accusations are predictable ingredients of an intelligence-state narrative. Obama castigated North Korea for the apparent "act of cybervandalism" (www.theguardian.com/us-news/2014/dec/21/obama-us-north-korea-state-terror-list-sony-hack) as he promised a "proportional response" (drones?)—even with the absence of concrete evidence. Then Sony decided to withhold the holiday release of *The Interview*. Obama criticized this action (www.theguardian.com/us-news/2014/dec/19/obama-sony-the-interview-mistake-north-korea). Not willing to concede the last point, Sony Entertainment CEO Michael Lynton responded that Sony sought advice from the White House without effect

UPDATE

The official "truth" continues to be a moving target (www.nytimes.com/2015/01/19/world/asia/nsa-tapped-into-north-korean-networks-before-sony-attack-officials-say.html, and www.theregister.co.uk/2015/01/19/nsa_saw_sony_hack). The latest FBI and NSA revelation is that they were monitoring North Korean network traffic all along, so they could easily trace the traffic back to the source in real time. Although this explanation seems the more likely than earlier ones, it still doesn't seem complete. For one thing, it leaves open the question of why the government didn't use a tech company intermediary to inform Sony that it should tweak its firewalls.

I look at these latest revelations as a last attempt to find some story that simultaneously satisfies the public curiosity, deflects media criticism, and doesn't make the agencies involved look incompetent. The intelligence agencies don't seem to understand that when a story arc begins with an absurdity or falsehood, the audience will never willingly suspend disbelief through to the final act. I still claim that critical pieces of the Sony hack narrative are missing, and that there's more to this than meets the proverbial eye. Stay tuned.

(www.theguardian.com/film/2014/dec/18/fbi-north-korea-sony-pictures-hack-the-interview). And so it goes. I'm confident that, were he still with us, Aldous Huxley would have said that this story does little more than feed mankind's almost infinite appetite for distraction from the more important affairs of our times.

KNOWN KNOWNS?

Someone hacked Sony. At this point, the finger-pointing and narrative is dominated by agendists who seek to create a usable history for themselves and their patrons. I'm not claiming that Kim Jong-un and North Korea aren't involved in the Sony hack. I'm claiming that it's irresponsible to make such accusations until verifiable proof is determined. Certainly the July 2009 distributed denial-of-service attacks against US and South Korean interests point to North Korean involvement, so we know that North Korea is capable of cybertransgressions. But in this case, the incomplete and unreliable evidence that's being offered amounts to little more than smoke and mirrors. The Sony hack story has all the substance and veracity of Nessie and Sasquatch sightings.

But let's be realistic. Searching for Nessie, Sasquatch, and the Guardians of Peace carries no penalty for the media. If the filmed search didn't find Nessie where expected, that's one more place we can rule out. We then get a few talking heads to follow up: "I never believed that Nessie would go there," "We're reviewing our evidence to see where we went wrong," and so on. Even if we can't conclusively prove that the Guardians of Peace are working for Kim Jong-un, we can find some senior government official to report that they probably are. That's almost the same thing as saying they might be, which is just one semantic smidge away from having no idea. But reporting that we have no clue won't sell much advertising. And, after all, we can always use some variant of the argument from ignorance to retroactively cover sloppy reporting.

In the meantime, Sony gets some much-needed free advertising for a film with an arguably tasteless plotline. This may be the real story: political satire works best when the audience isn't bludgeoned with crude character assassination, suggestions of cruelty, and comical disrespect. Making films that make sport of killing political leaders is just poor form and relies more on shock value than creativity. Moviegoers would be better served by a re-screening of Charlie Chaplin's 1940 classic *The Great Dictator* and using their imagination to port the concepts over to current affairs.

It's up to enlightened audiences to reject this background noise for what it is; mass media has every incentive to tilt toward coverage of the inane. And governments would be well advised to avoid attaching military and economic consequences to crimes against corporations, especially when such crimes have no national security implications. It's also a good idea to avoid prejudging the outcome of an ongoing investigation that involves world leaders. The tough talk and bogus claims from all directions, the threats and sanctions based on spotty evidence, and the accusations and counter-accusations serve us all poorly. Accusing attribution during an ongoing investigation is like painting falling leaves: the results are sloppy and unlikely to have enduring value. Thus far, reporting on the Sony hack has been banal in the extreme. C

REFERENCES

1. T.C. Reed, *At the Abyss: An Insider's History of the Cold War*, Presidio Press, 2004.
2. B. Gellman, *Angler: The Cheney Vice Presidency*, Penguin Press, 2009.
3. M. Isikoff and D. Corn, *Hubris: The Inside Story of Spin, Scandal, and the Selling of the Iraq War*, Crown, 2006.

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.

SOFTWARE TECHNOLOGIES

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

KnowLang: Knowledge Representation for Self-Adaptive Systems

Emil Vassev and Mike Hinchey, Lero—the Irish Software Research Centre

The KnowLang framework models uncertainty into the development of software-intensive systems to create enhanced possibilities for self-adaptation.

Software-intensive systems require considerable knowledge provided by software engineers and others to help explain the problem domain. Still, computers talk in a binary language that's simple, logical, and sound, with none of the ambiguity that characterizes human language. So, we can't simply give computers manuals and textbooks and expect them to know what they should do.

Instead, the knowledge given to computers must be expressed in well-founded computational structures that programs can translate into binary computer language. Such knowledge representation structures generally take the form of primitives: rules, frames, semantic networks, concept maps, ontologies, and logical expressions.[1]

KNOWLEDGE REPRESENTATION AND REASONING FOR SELF-ADAPTATION

Following a self-adaptation paradigm, software-intensive systems can respond to changing operational contexts, environments, or system characteristics and thus achieve greater versatility, flexibility, and resiliency, and also become more robust, energy-efficient, and customizable. Consequently, developing self-adaptive systems with an eye toward knowledge representation and reasoning (KR&R) has been an area of increasing interest over the years: examples include research in semantic mapping, aspects of planning and control, and, most notably, human–robot interaction (HRI). In general, KR&R methodologies strive to solve complex problems characteristic of nondeterministic operational environments and of systems that must reason at runtime to discover missing answers.

Decision making is a complex process, often based on more than just logical conclusions. In representing degrees of belief about knowledge that is necessarily uncertain or changing, probability and statistics can provide the basis for reasoning. For example, statistical inferences

Figure 1. KnowLang multitier knowledge specification model. KnowLang provides a formal language that integrates ontologies with rules and Bayesian networks based on logical and statistical reasoning to build a knowledge base (KB) via three main tiers: a *knowledge corpus* that explicitly represents domain concepts and relationships; *KB operators* that represent particular and general factual knowledge; and *inference primitives* that use additive probabilities to represent degrees of belief in uncertain knowledge.

might help us draw conclusions about a city's overall traffic patterns based on data obtained from relatively few streets.

Bayesian networks are often used to represent a *belief probability*, which summarizes a potentially infinite set of possible circumstances. Belief probability influences decision making based on a system's past experiences, associating future success with prior actions generated in the execution environment. Maintaining an execution history for such actions helps the system compute and re-compute the success probability of action execution. In this way, the system may learn (that is, infer new knowledge) and adapt so as not to execute actions that traditionally have had a low success rate.

KNOWLANG

To operate efficiently and reliably in open-ended environments, systems must have some initial knowledge as well as the ability to learn based on knowledge processing and awareness.[2] Moreover, a system's knowledge must be structured to provide an essential awareness of both its internal and external worlds. To meet these and other challenges, Lero—the Irish Software Centre (www.lero.ie), has developed the KnowLang framework within the ASCENS Project mandate (www.ascens-ist.eu).

KnowLang (http://knowlang.lero.ie) is a KR&R framework for efficient and comprehensive knowledge structuring that's intended to support both logical and statistical reasoning. At its very core, the framework is a formal specification language providing a comprehensive, yet multitier model where knowledge can be presented at different levels of abstraction and grouped by following both hierarchical and functional patterns.

Knowledge specified with KnowLang takes the form of a knowledge base (KB) incorporating an ontology using concepts organized through concept trees, object trees, relations, and predicates. Each concept is specified with particular properties and functionality and is hierarchically linked to other concepts. For reasoning purposes, every concept specified with KnowLang has an intrinsic "state" attribute that can be associated with a set of possible state expressions. Moreover, concepts and objects can be connected via relations. Relations are binary and can have *probability-distribution* attributes.

Probability distribution, which is used to support probabilistic reasoning, presents a belief probability about relations between different knowledge concepts—time, situation, action, event, and so forth—that are often in competition. By specifying KnowLang relations through their probability distributions, we're actually creating Bayesian networks that connect concepts and objects within the ontology.

MODELING KNOWLEDGE WITH KNOWLANG

Modeling knowledge with KnowLang occurs in three stages:

- *Initial knowledge gathering*, when domain experts determine the interest domain's basic notions, relations, and functions or operations.
- *Behavior definition*, during which domain-specific situations and behavior policies are identified as control data to help determine important self-adaptive scenarios.
- *Knowledge structuring* to encapsulate the identified domain entities, situations, and behaviors

into KnowLang structures—that is, concepts, objects, relations, facts, and rules.

This knowledge modeling process results in KnowLang's multitier specification model, illustrated in Figure 1.

A KB specified with KnowLang outlines a KR (knowledge relationship) context that is specific to the targeted system's domain. A special KnowLang Reasoner operates in this context to allow for knowledge querying and updating. The KnowLang Reasoner is conceived as a component hosted by a self-adaptive system; thus, it runs in the system's operational context as any other system component does. By operating on the KB, the reasoner can *infer* special self-adaptive behavior.

KnowLang provides a predefined set of "ask" and "tell" operators that allow communication with the KB. Tell operators feed the KR context important information—driven by errors, executed actions, new sensory data, and the like—thus helping the KnowLang Reasoner update the KR with recent changes in both the system state and the execution environment. The system uses ask operators to elicit recommended behaviors, with prior knowledge compared to current outside input to generate appropriate actions that comply with determined goals and beliefs. In addition, ask operators may provide the system with awareness-based conclusions about the current system and environment states and, ideally, with behavior models for self-adaptation.

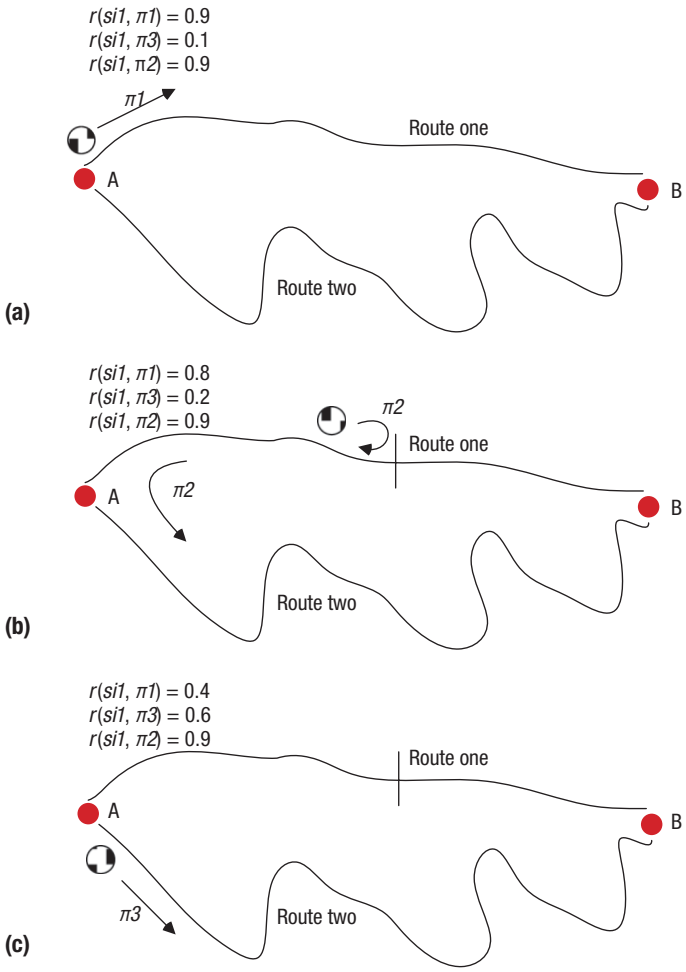

Figure 2. Through self-adaptive behavior based on applying KnowLang framework policies, a robot progammed to carry items from point A to point B via route one (a) will alter its course when route one is blocked (b) and instead follow route two (c). If route one is blocked in the future, the higher probabilistic belief rate regarding route two will lead the robot to change its behavior, choosing route two as its primary route.

KNOWLEDGE REPRESENTATION FOR SELF-ADAPTIVE BEHAVIOR

In summary, KnowLang employs special knowledge structures and a reasoning mechanism to model self-adaptive behavior. Such behavior can be expressed via KnowLang's structure policies: events, actions, and situations, as well as relations between policies and situations.[3] Policies are at the core of any KR for self-adaptive behavior. Ideally, KnowLang policies are specified in a way that will allow a system to pursue a specific goal within a specific situation via actions generated in the environment or in the system itself.

Specific conditions determine the specific actions to be executed. These conditions often differ from the past situations triggering a policy. Thus, self-adaptive behavior depends not only on the specific situations a policy is specified to handle, but also on additional conditions and probabilistic beliefs.

In order to initiate self-adaptive behavior, relations must be specified between policies and situations vis-à-vis a belief probability: a policy may be related to multiple situations and vice versa. A belief probability supports probabilistic reasoning, helping the KnowLang Reasoner choose the most probable situation–policy "pair"—that is, the most probable policy to be applied to a particular situation. Thus, we might specify several different relations connecting a specific situation with various policies that may be undertaken when

the system is in that situation; the probability distribution should help the Reasoner decide which policy to choose in each case.

At runtime, the KnowLang Reasoner maps situations to policies, and, for any actual situation, applies the policy with the highest possible belief probability. When a policy is applied, the Reasoner checks it against the particular conditions to be met and then performs actions that meet these particular conditions. Although initially specified, the belief probability is recomputed after any action is executed. The Reasoner maintains a history of these action executions, and re-computation is based on the consequences of the action execution, which allows for reinforcement learning within the system.

CASE STUDY

To illustrate self-adaptive behavior based on this approach, imagine a robot carrying items from point A to point B using two possible routes, route one and route two, as in Figure 2. Situation *si1:* "robot is at point A loaded with items" triggers policy $\pi 1$: "go to point B via route one" if the relation $r(si1,\pi1)$ has a higher probabilistic belief rate than other possible relations (for example, such a belief rate has been initially established for this relation because route one is shorter). Whenever the robot is in *si1*, it will continue applying the $\pi 1$ policy.

However, when the robot finds itself in situation *si2*: "route one is blocked," it will no longer apply that policy; *si2* will trigger policy $\pi 2$: "go back to *si1* and then apply policy $\pi 3$," with policy $\pi 3$ defined as $\pi 3$: "go to point B via route two."

The unsuccessful application of policy $\pi 1$ will decrease the probabilistic belief rate of relation $r(si1,\pi1)$, and the eventual successful application of policy $\pi 3$ will increase the probabilistic belief rate of relation $r(si1,\pi3)$. Thus, if route one continues to be blocked in the future, relation $r(si1,\pi3)$ will come to have a higher probabilistic belief rate than relation $r(si1,\pi1)$, and the robot will change its behavior by choosing route two as a primary route.

It is also possible for the situation to change in response to external stimuli—for example, the robot receives a "route two is blocked" message or "route one is obstacle-free" message.

Any long-running self-adaptive system must change behavior in response to stimuli from the execution environment, and all such environments are subject to uncertainty due to potential evolution in requirements, business conditions, available technology, and the like. Thus, it's important to capture and plan for uncertainty as part of the development process. Failure to do so may result in systems that are overly rigid for their purpose, an eventuality of particular concern for domains that typically use self-adaptive technology, such as unmanned space flight.

We hypothesize that KnowLang, by allowing developers to model uncertainty and create mechanisms for managing it as part of knowledge representation and reasoning, will lead to systems that are expressive of the real world, more fault tolerant because they can anticipate fluctuations in requirements and conditions, and highly flexible in managing dynamic change.

ACKNOWLEDGMENTS

This research was supported by the European project IP 257414 (ASCENS) and by Science Foundation Ireland grant 13/RC/2094 to Lero—the Irish Software Research Centre.

REFERENCES

1. E. Vassev and M. Hinchey, "Knowledge Representation and Reasoning for Intelligent Software Systems," *Computer*, vol. 44, no. 4, 2012, pp. 96–99.
2. E. Vassev and M. Hinchey, "Awareness in Software-Intensive Systems," *Computer*, vol. 45, no. 12, 2013, pp. 84–87.
3. E. Vassev, M. Hinchey, and B. Gaudin, "Knowledge Representation for Self-Adaptive Behavior," *Proc. 5th Int'l C* Conf. Computer Science and Software Eng.* (C3S2E 12), 2012, pp. 113–117.

EMIL VASSEV is a senior research fellow at Lero—the Irish Software Research Centre at the University of Limerick, Ireland. Contact him at emil.vassev@lero.ie.

MIKE HINCHEY is director of Lero—the Irish Software Research Centre and a professor of software engineering at the University of Limerick, Ireland. Contact him at mike.hinchey@lero.ie.

SECURITY

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

Attribute-Based Access Control

Vincent C. Hu, D. Richard Kuhn, and David F. Ferraiolo,
National Institute of Standards and Technology

Attribute-based access control (ABAC) is a flexible approach that can implement AC policies limited only by the computational language and the richness of the available attributes, making it ideal for many distributed or rapidly changing environments.

Traditionally, access control (AC) has been based on the identity of a user requesting execution of a capability to perform an operation (for example, read) on an object (for example, a file), either directly or through predefined attribute types such as roles or groups assigned to that user. Practitioners have noted that this AC approach is often cumbersome to manage given the need to associate capabilities directly to users or their roles or groups. In addition, the requester qualifiers of identity, groups, and roles are often insufficient in expressing real-world AC policies. An alternative is to grant or deny user requests based on arbitrary attributes of the user and selected attributes of the object, and environment conditions that could be globally recognized and more relevant to the policies at hand. This approach is often referred to as attribute-based access control (ABAC).

ABAC: A FLEXIBLE ACCESS CONTROL MODEL

ABAC is a logical AC model that controls access to objects by evaluating rules against the attributes of entities (subject and object), operations, and the environment relevant to a request. ABAC enables more precise AC by allowing for a higher number of discrete inputs into an AC decision and thereby providing a larger set of possible combinations of those variables to reflect a larger and more definitive set of possible rules to express policies, which are limited only by the computational language and the richness of the available attributes.

This flexibility enables creation of access rules without specifying individual relationships between each subject and each object. For example, a subject is assigned a set of subject attributes upon employment, such as Nancy Smith is a Nurse Practitioner in the Cardiology Department. An object is assigned its object attributes upon creation, such as a folder with Medical Records of Heart Patients. Objects may receive their attributes either directly from the creator or as a result of automated scanning tools. The administrator or owner of an object creates an AC rule using attributes of subjects and objects to govern the set of allowable capabilities—for example,

Computer | Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page

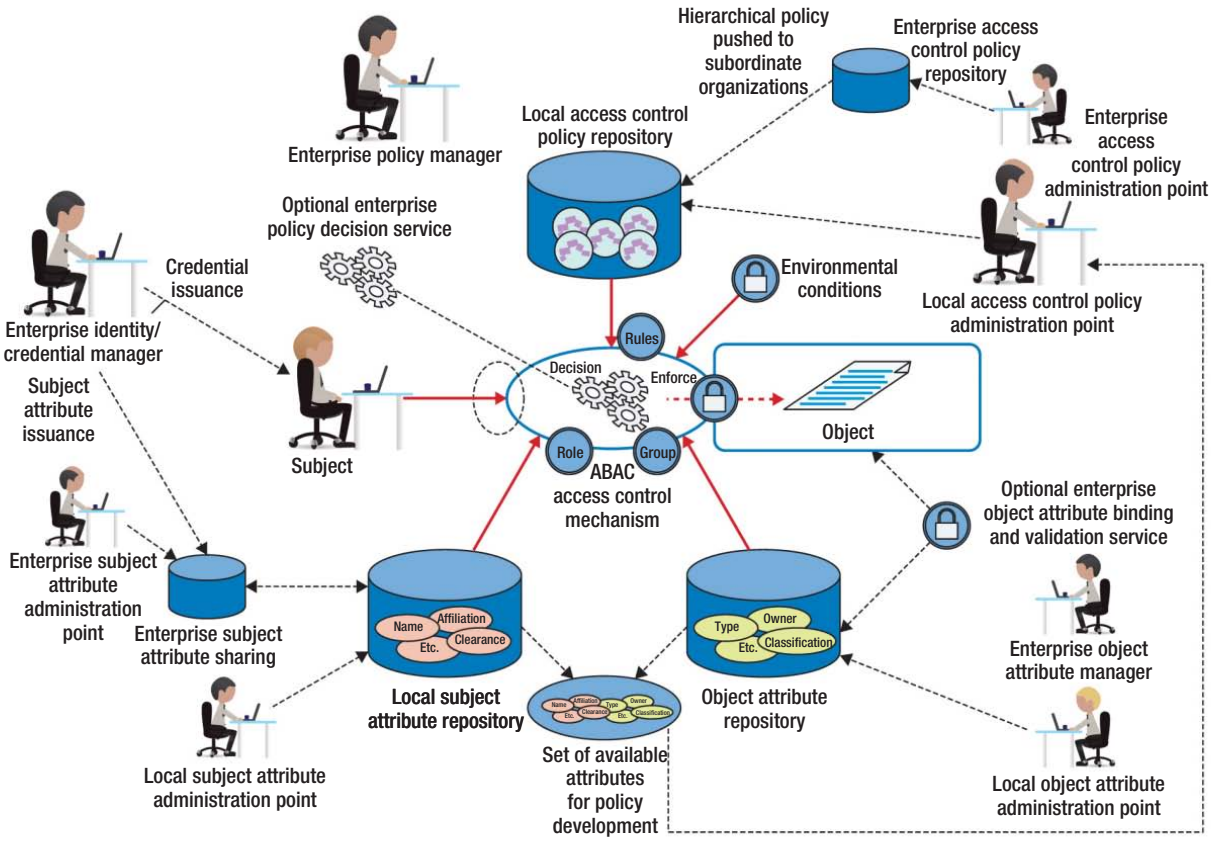

Figure 1. Attribute-based access control (ABAC) example. Adapted from V.C. Hu et al., *Guide to Attribute Based Access Control (ABAC) Definition and Considerations*, NIST Special Publication 800-162, Nat'l Institute of Standards and Technology, Jan. 2014.

all Nurse Practitioners in the Cardiology Department can View the Medical Records of Heart Patients.

Under ABAC, access decisions can change between requests simply by altering attribute values, without requiring changes to the subject/object relationships defining the underlying rule sets. This provides a more dynamic AC management capability and limits long-term maintenance requirements of object protections.

Further, ABAC enables object owners or administrators to apply AC policy without prior knowledge of the specific subject and for an unlimited number of subjects that might require access. As new subjects join the organization, rules and objects need not be modified, and as long as the subject is assigned the attributes necessary for access to the required objects—for example, all Nurse Practitioners in the Cardiology Department are assigned those attributes—no modifications to existing rules or object attributes are required. This accommodation of the external (unanticipated) user is one of the primary benefits of employing ABAC.[1,2]

As a result of this flexibility, ABAC has attracted interest across industry and government, and is the fastest-growing AC model today.[3] It has been integrated with other approaches, such as the International Committee for Information Technology Standards (INCITS) standard for role-based access control,[4] and has become the basis for an increasing range of products. But beyond the basic scheme of associating attributes with subjects, objects, and environments, there has been little consistency among ABAC implementations.

IMPLEMENTING ABAC IN THE ENTERPRISE ENVIRONMENT

Due to a lack of consensus on ABAC features, users can't accurately assess the benefits and challenges associated with the model. To help address this problem, the National Institute of Standards and Technology (NIST) released Special Publication (SP) 800-162, *Guide to Attribute Based Access Control (ABAC) Definition and Considerations*.[1] This document serves a twofold purpose. First, it provides federal agencies with a definition of ABAC and a description of its functional components. Second, it describes planning, design, implementation, and operational considerations for employing ABAC within an enterprise to improve information sharing while maintaining control of that information. The guide focuses on the

TABLE 1. Level of attribute assurance (LOAA) mappings example.

LOAA	Accuracy	Integrity	Availability
1	Attributes are properly verified for veracity through provision and management.	Secure attribute repository. Secure communication between attribute providers (APs) and relying parties (RPs).	Attribute refresh frequency meets the system performance requirement.
2	Includes level 1. Documented rule or standards for attribute value assignment and definition (syntax and semantic rule).	Includes level 1. Dedicated attribute repositories.	Includes level 1. Attribute caching during runtime meets the system performance requirement.
3	Includes level 2. Attributes cover all of the organization's protection policy requirements (semantically complete).	Includes level 2. Encrypted attribute values and communications between APs and RPs.	Includes level 2. Failover or backup attributes support.
4	Includes Level 3. Attributes under federated or unified governance.	Includes level 3. Formal rules or policy (or standards) for create, update, modify, and delete attributes.	Includes level 3. Log for attribute changes and access.

challenges of implementing ABAC rather than on balancing the cost and effectiveness of other capabilities versus ABAC.

When deployed across an enterprise to increase information sharing among diverse organizations, ABAC implementations can become complex, requiring an attribute management infrastructure, machine-enforceable policies, and an array of functions that support access decisions and policy enforcement. As Figure 1 shows, in addition to the basic policy, attribute, and AC mechanism requirements, the enterprise must support management functions for enterprise policy development and distribution, enterprise identity and subject attributes, subject attribute sharing, enterprise object attributes, authentication, and AC mechanism deployment and distribution.

Enabling these capabilities requires careful consideration of numerous factors that will influence the design, security, and interoperability of an enterprise ABAC solution. These factors can be summarized around a set of activities:

- establish the business case for ABAC implementation;
- understand the operational requirements and overall ABAC enterprise architecture;
- establish or refine business processes to support ABAC;
- develop and acquire an interoperable set of ABAC capabilities; and
- operate with efficient ABAC processing.

NIST SP 800-162 helps ABAC system planners, architects, managers, and implementers carry out these activities in four phases. The *initiation phase* includes building the business case for deploying ABAC capabilities; scalability, feasibility, and performance requirements; and developing operational requirements and architecture. The *acquisition/development phase* includes business process generation and deployment preparation, system development and solution acquisition considerations, and other enterprise ABAC capabilities. The *implementation/assessment phase* includes attribute caching, attribute source minimization, and ABAC interface specifications. Finally, the *operations/maintenance phase* includes availability of quality ABAC data.

ATTRIBUTE ASSURANCE

The metadata of ABAC attributes communicate aspects that are important for attribute standardization. By coupling a common set of mandatory and optional metadata with attribute assertions, ABAC systems can query attribute information to make their own risk-based decisions, especially when delivered via a broker connected to many systems.

In general, attribute metadata fall into three categories:

- *Accuracy* establishes the policy and technical underpinnings for semantically and syntactically correct use of these attributes

and environmental conditions, and ensures that the reported attributes are trustworthy, based on the trust established in the measurement and reporting processes.

- *Integrity* considers different standards and protocols used for secure sharing of attributes between systems in order to avoid compromising the integrity and confidentiality of the attributes or exposing vulnerabilities in attribute provider (AP) or relying party (RP) systems or entities.
- *Availability* ensures that the update and retrieval of attributes support the RP. In addition, attribute repositories' failover and backup capability must be considered. Note that some attributes might change regularly or over time.

An AP is any person or system that provides subject, object (or resource), or environmental condition attributes regardless of transmission method. The AP could be the original authoritative source or receiving information from an authoritative source for repacking and storing-and-forwarding to the ABAC system. Attribute values can be human generated (for example, an employee database) or derived from formulas (for example, a credit score). Regardless of the attribute source, the system should ensure that the attribute value received from an AP is accurately associated with the subject, object, or environmental condition to which it applies.[2] Table 1 illustrates example levels of attribute assurance (LOAA) based on the accuracy, integrity, and availability properties.

Attribute-based access control is a flexible approach that can implement AC policies limited only by the computational language and the richness of the available attributes. This flexibility enables the greatest breadth of subjects to access the greatest breadth of objects without specifying individual relationships between each subject and each object, making ABAC ideal for many distributed or rapidly changing environments.

ABAC has the potential to dramatically improve AC in modern applications such as e-commerce and the Internet of Things. In the meantime, a consensus definition of ABAC is needed, and work remains to be done in assuring attribute accuracy and reliability. For more information on ongoing efforts, see http://csrc.nist.gov/projects/abac/index.html.

REFERENCES

1. V.C. Hu et al., *Guide to Attribute Based Access Control (ABAC) Definition and Considerations*, NIST Special Publication 800-162, Nat'l Institute of Standards and Technology, Jan. 2014; http://nvlpubs.nist.gov/nistpubs/specialpublications/NIST.sp.800-162.pdf.
2. V.C. Hu, D.F. Ferraiolo, and D.R. Kuhn, *Assessment of Access Control Systems*, NIST Interagency Report 7316, Nat'l Institute of Standards and Technology, Mar. 2006; http://csrc.nist.gov/publications/nistir/7316/NISTIR-7316.pdf.
3. Avatier Corp., "Leveraging Today's Megatrends to Drive the Future of Identity Management," video presentation, Gartner Identity and Access Management (IAM) Summit, 2012; www.avatier.com/products/identity-management/resources/gartner-iam-2020-predictions.
4. D.R. Kuhn, E.J. Coyne, and T.R. Weil, "Adding Attributes to Role Based Access Control," *Computer*, vol. 43, no. 6, 2010, pp. 79–81.

VINCENT C. HU is a computer scientist in the Computer Security Division at the National Institute of Standards and Technology. Contact him at vhu@nist.gov.

D. RICHARD KUHN is a project leader and computer scientist in the Computer Security Division at the National Institute of Standards and Technology. Contact him at kuhn@nist.gov.

DAVID F. FERRAIOLO is a computer scientist and manages the Secure Systems and Applications Group in the Computer Security Division at the National Institute of Standards and Technology. Contact him at dferraiolo@nist.gov.

CS CONNECTION

THANKS TO VOLUNTEERS

The IEEE Computer Society thanks the following associate editors and editorial board members who retired at the end of 2014 for giving their valuable time and support to our publications.

Computer

Jean Bacon
Oliver Bimber
Albrecht Schmidt
Kelvin Sung
Kathleen Swigger

IEEE Annals of the History of Computing

Janet Abbate
Walter M. Carlson
Alan Clements
Thomas Haigh
Hunter Heyck
Eden Medina
Raul Rojas
Christopher H. Sterling
Wladyslaw M. Turski

IEEE Computer Architecture Letters

David August
Antonio Gonzales
Ramaswamy Govindarajan
Sudhanva Gurumurthi
Avi Mendelson
Mark Oskin
Norm Rubin

IEEE Computer Graphics and Applications

Anthony (Tony) DeRose
Rae A. Earnshaw
Dieter W. Fellner
Cindy Grimm
Diego Gutierrez
Anselmo Lastra
Gabriel Taubin
Luis Velho

IEEE Intelligent Systems

Silvia Coradeschi

IEEE Internet Computing

Siobhán Clarke
Samuel Madden
Pankaj Mehra
Torsten Suel
Doug Tygar

IEEE MultiMedia

Dan Ellis
Frank Nack
Doree Duncan Seligmann
Malcolm Slaney
Qibin Sun
Hari Sundaram

IEEE Pervasive Computing

Eyal de Lara

IEEE Security & Privacy

Elisa Bertino
Luanne Goldrich
Cynthia Irvine
Peter G. Neumann
Fred B. Schneider
Ray Vaughn
Tara Whalen
Alec Yasinsac

IEEE Transactions on Computational Biology and Bioinformatics

Tatsuya Akutsu
Rolf Backofen
Diego di Bernardo
Graziano Chesi
Arne Elofsson
Elena Marchiori
Sushmita Mitra
Vincent Moulton
Limsoon Wong

IEEE Transactions on Computers

D. R. Avresky
Yong-Bin Kim
Nagarajan Ranganathan
Kay Roemer
Alexander Shvartsman

IEEE Transactions on Dependable and Secure Computing

Peng Ning
Mukesh Singhal
Xinyuan Wang

IEEE Transactions on Haptics

Seungmoon Choi
Martha Flanders
Keyvan Hashtrudi-Zaad
Miguel Otaduy

IEEE Transactions on Knowledge and Data Engineering

Divyakant Agrawal
Christian Bohm
Brian Cooper
Juliana Freire
George Karypis
Gerome Miklau
Hiroyuki Kitagawa
Nikos Mamoulis
Joerg Sander
S. Sudarshan
Yufei Tao

IEEE Transactions on Learning Technologies

Paul De Bra
Vania Dimitrova
Jim Greer
Friedrich Hesse
Judy Kay
Chee-Kit Looi
Riichiro Mizoguchi
Thomas Ottmann
Demetrios Sampson
Timothy K. Shih
Simon Buckingham Shum
Marcus Specht

IEEE Transactions on Mobile Computing

Thomas Hou
Lili Qiu
Neal Patwari
Konstantinos Psounis
Ravi Sundaram
Andreas Terzis
Steve Weber
Edmund Yeh
Qian Zhang

CALL AND CALENDAR

IEEE Transactions on Parallel and Distributed Computing

Srinivas Aluru
Robert Baldoni
Olivier Beaumont
Jinjun Chen
Xiuzhen Cheng
Jiming Chenl
Mooi Choo Chuah
Dick Epema
Jose Flich
Hannes "Frank" Frey
Mahmut Taylan Kandemir
Keqiu Li
Xu Li
D. Manivannan
Jaime Lloret Mauri
Frank Mueller
Symeon Papavassiliou
Sanjay Ranka
Paolo Santi
WenZhan Song
Xueya Tang
My T. Thai
Dajin Wang
Jun Wang
Li Xiao
Dong Xuan
Jingyuan Zhang
Si Qing Zheng

IEEE Transactions on Pattern Analysis and Machine Intelligence

Shai Avidan
Fei-Fei Li
Greg Mori
Haesun Park
Ian Reid
Carsten Rother
Tobias Scheffer
Cristian Sminchisescu
Charles Sutton
Ben Taskar
Yee Whye The
Massimo Tistarelli
Eric P. Xing

IEEE Transactions on Scalable Computing

Hong Cai
Louise Moser

IEEE Transactions on Software Engineering

Antonia Bertolino
Marsha Chechik
Betty Cheng
Elisabetta Di Nitto
Harald Gall
Carlo Ghezzi
Dimitra Giannakopoulou
Jane Huang
Paola Inverardi
Michael Jackson
David Ross Jeffery
Marta Kwiatkowska
Bev Littlewood
Nenad Medvidovic
Martin Robillard
Gregg Rothermel
Wolfram Schulte
Dag Sjøberg
Margaret-Anne (Peggy) Storey
Paul Strooper
Walter Tichy
Frank Tip
Murray Woodside

IEEE Transactions on Visualization and Computer Graphics

Ronan Boulic
Wojciech Matusik
Dieter Schmalstieg

IT Professional

Thomas Costello
Edward J. Coyne
Thomas Jepsen
Phillip Laplante
Keith W. Miller
Jeffrey Voas

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

CALL AND CALENDAR

CALLS FOR ARTICLES FOR IEEE CS PUBLICATIONS

IEEE Computer Graphics and Applications plans a November/December 2015 special issue on cutting-edge visualization and computer graphics research in Asia.

In recent years, the world has witnessed rapid economic development and research progress in Asia. In the big data era, this has provided a rich base for the quick development of both research in and applications for visualization and computer graphics.

This special issue will cover recent advances in areas such as new methods, designs, and systems for scientific visualization; information visualization and visual analytics; new computer graphics algorithms and systems; and case studies describing success and failure in applying visualization to real-world problems in Asia.

Articles are due **1 March 2015**. Visit www.computer.org/web/computingnow/cgacfp6 to view the complete call for papers.

IEEE Transactions on Emerging Topics in Computing (TETC) plans a special issue on advances in mobile cloud computing for the first issue of 2016.

There has been a phenomenal burst of research in mobile cloud computing. Mobile applications demand greater resources and improved interactivity for a better user experience. Resources in cloud computing platforms such as Amazon Web Services, Google App Engine, and Microsoft Azure offer a way to remedy mobile devices' lack of local resources.

The objective of this special issue is to cover the most recent R&D on mobile cloud computing technologies and to give industry and academia an opportunity to showcase their recent progress in this area.

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

IEEE Transactions on Emerging Topics in Computing (*TETC*) plans a special issue on emerging trends in education for the first issue of 2016.

Technological advancements such as those used in cloud computing; mobile devices; and big, open, and linked data bring with them great opportunities for enriching and broadening the educational experience.

For instance, virtual learning environments are increasingly used in communications between students and teachers.

At the same time, mobile computing is expanding the reach of learning content and frameworks.

Although there are many future visions for education, great efforts will be needed to reach a profound integration between well-established and emerging technologies.

By building on a solid scientific and methodological foundation where theory and practice converge, this special issue aims to present both the current trends that characterize the learning and teaching domains of today, as well the expected evolution that will shape the education of tomorrow.

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

IEEE Security & Privacy plans a January/February 2016 special issue on the interconnected Web ecosystem.

The Web is being fused into the human experience in a number of interesting ways. For example, infant monitoring devices provide parents with real-time information on their smartphones about their baby's breathing, skin temperature, body position, and activity level. Unfortunately, this Web-enabled interconnectedness has opened a host of serious security, privacy, and dependability concerns. A single flaw or feature in Web technology could have significant, unforeseen negative consequences on myriad interconnected systems.

This special issue will address these challenges.

Abstracts are due **1 March 2015**. Articles are due **1 April 2015**. Visit www.computer.org/web/computingnow/spcfp-jan-feb-2016 to view the complete call for papers.

IEEE Software plans a November/December 2015 special issue titled "Refactoring: Accelerating Software Change."

Modern software is rarely written from scratch. It usually incorporates code from previous systems and is itself reincarnated in other programs. Software also constantly changes as bugs are fixed and features are added. These changes are usually performed by more than one programmer, and not necessarily by the code's original authors.

Refactoring—improving code's internal structure without altering its external behavior—supports this highly dynamic software life cycle.

This special issue will focus on the real-world application of research,

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.

ICIS 2015

The 14th IEEE/ACIS International Conference on Computer and Information Science (ICIS 2015) is sponsored by the IEEE Computer Society and International Association for Computer and Information Science (ACIS).

ICIS 2015 brings together scientists, engineers, computer users, and students to share their experiences, new ideas, and research results about all aspects—including theory, applications, and tools—of computer and information science. They will discuss the solutions they adopted to various problems and practical challenges they faced.

The conference will cover many topics, including communication systems and networks; mobile computing; parallel and distributed computing; software architectures, design patterns, and frameworks; data mining, data warehousing, and databases; speech and signal processing; image processing; pattern recognition; visual and multimedia computing; Web engineering and applications; intelligent agent technology; and agent-based systems.

ICIS 2015 will take place 30 June to 2 July 2015 in Las Vegas, Nevada. Visit www.acisinternational.org/icis2015 for complete conference information.

2015 EVENTS

MARCH 2015

9–13 ..... NetSys 2015
23–27 ..... PerCom 2015
23–27 ..... VR 2015

APRIL 2015

13–17 ..... NetSoft 2015
20–22 ..... BigMM 2015
29–30 ..... ENASE 2015

MAY 2015

4–8 ..... WICSA 2015
4–8 ..... FG 2015
16–24 ..... ICSE 2015

practical experiences, success stories, and lessons learned in this area.

Articles are due **1 April 2015**. Visit www.computer.org/software/cfp6 to view the complete call for papers.

IEEE Computer Graphics and Applications plans a January/February 2016 special issue on visual computing and the progress of developing countries.

Many developing countries' advancements are due in part to their technological development in fields like visual computing.

Visual computing can yield benefits in areas such as engineering, healthcare, industry, military, education, and government.

The guest editors welcome submissions from both commercial and academic sources and from both researchers and practitioners, particularly from developing countries.

Articles are due 1 May 2015. Visit www.computer.org/web/computingnow/cgacfp1 to view the complete call for papers.

IEEE Internet Computing plans a January/February 2016 special issue on Internet economics.

The breadth of economic activity on the Internet is exploding. The resulting economic systems lead to a plethora of new research questions, both theoretical and data-driven, touching on both analysis and design.

This special issue will address theoretical and applied research related to the modeling, analysis, and design of Internet-specific economic activities and incentive systems. The issue will be interdisciplinary in nature and will include any research related to economic aspects of the Internet.

Articles are due **1 May 2015**. Visit www.computer.org/web/computingnow/iccfp1 to view the complete call for papers.

MARCH 2015

9–13 March: NetSys 2015, Int'l Conf. Networked Systems, Cottbus, Germany; www.netsys2015.com

23–27 March: PerCom 2015, IEEE Int'l Conf. Pervasive Computing and Comm., St. Louis; www.percom.org

23–27 March: VR 2015, IEEE Virtual Reality, Arles, France; http://ieeevr.org/2015

APRIL 2015

13–17 April: NetSoft 2015, 1st IEEE Conf. Network Softwarization, London; http://sites.ieee.org/netsoft

20–22 April: BigMM 2015, 1st IEEE Int'l Conf. Multimedia Big Data, Beijing; www.bigmm2015.org

29–30 April: ENASE 2015, 10th Int'l Conf. Evaluation of Novel Approaches to Software Eng., Barcelona; www.enase.org

MAY 2015

4–8 May: WICSA 2015, 12th Working IEEE/IFIP Conf. Software Architecture, Montreal; http://wicsa2015.org

4–8 May: FG 2015, 11th IEEE Int'l Conf. Automatic Face and Gesture Recognition, Ljubljana, Slovenia; www.fg2015.org

16–24 May: ICSE 2015, 37th Int'l Conf. Software Eng., Florence, Italy; http://2015.icse-conferences.org

CAREER OPPORTUNITIES

SENIOR TERADATA WAREHOUSE CONSULTANT. Analyze and develop logical database designs, data models, and relational database definitions across multiple computing environments utilizing Teradata and SAP platforms and tools. Identify and implement new uses of information technologies to meet strategic objectives. Prepare business requirements reports, recommendations, feasibility studies and cost justification statements. Analyze and design the Teradata Warehouse Architectures to support the insurance business including the extraction, transformation and loading of data from Mainframe policy administration systems. Apply to: M. Gagne, MIP F110, Massachusetts Mutual Life Insurance Company, 1295 State St, Springfield, MA 01111.

SAP SECURITY ANALYST. Provide SAP security support to managers, perform security reviews and testing of the SAP systems, participate in audit and compliance activities. Perform security upgrades, manage transports, and troubleshoot access issues. Build and maintain SAP user roles across the SAP landscape including ECC, BI, PI, SM and Enterprise Portal. Apply to: M. Gagne, MIP F110, Massachusetts Mutual Life Insurance Company, 1295 State Street, Springfield, MA 01111.

ENGINEERING

QUANTUMSCAPE CORPORATION is accepting resumes for the position of Sr. Member of Technical Staff in San Jose, CA. Responsible for maintaining an atomic layer deposition (ALD) fabrication tool, running it to produce thin films of specified characteristics and properties, and measuring film properties. Mail resume to QuantumScape Corporation, Staffing Department, 1730 Technology Drive, San Jose, CA 95110. Must reference Ref. Code SMTS-CC.

SIEMENS PLM SOFTWARE INC. has the following openings: Software Engineer Adv/UGS159 in Milford, OH to design, develop, modify, & implement software programming for products. Requires up to 5% domestic travel to client sites. Software Engineer Adv/UGS165 in Troy, MI to develop data mgmt. solutions for Mechatronics in Teamcenter. Email resumes to PLMCareers@ugs.com & refer to Job title/code of interest. EOE.

SOFTWARE ENGINEER, San Mateo, CA sought by MDOTM, F/T, Dsgn & implmt user profiling, web systms for advertising platform in PHP/Java, Participate in reqmts review, functional spec, Technical Dsgn/Algorithms; participate in brainstorming sessions & contribute ideas to our tech, algorithms & products; help improve systm qlty through writing unit tests, automation & performing code reviews. Masters in Comp Sci or Info Systms & 2 yrs S/ware Engg exp. Knowl of: 1. Object Oriented Dsgn & coding skills in PHP/Java; 2. Web dvlpmt in Java/PHP; 3. Web svcs dsgn & dvlpmt; 4. SQL dbases. Resume to careers@mdotm.com.

NUVIEW TECHNOLOGIES INC. has the following opening in Orlando, FL:Programmer Analyst:- Analyze user requirements, procedures & problems to improve existing systems using Java based technologies. Test, maintain, monitor computer programs & systems. Knowledge of Databases & SQL. Req. 2 yrs exp. Send resume to Jobs@Nuviewtech.biz. This position will involve working in unanticipated locations.

CAREER OPPORTUNITIES

BMC SOFTWARE INC.. has an opening for Sr. Technical Support Analyst in Lexington, MA to install, configure, test & troubleshoot business software application releases in a customer support environ. Email resumes to Olivia_Delgado@bmc.com refer to Req# 15000025.

ENGINEERING. TRANSCRIPTIC, INC. is accepting resumes for the position of Member of Technical Staff in Menlo Park, CA. Utilize mechanical engineering knowledge to gather technical requirements, design, prototype and fabricate lab automation robotics for biological research. Mail resume to Transcriptic, Inc., Staffing Department, 3565 Haven Street, Menlo Park, Suite 3, CA 94025. Must reference Ref. MTS-JS.

COMPUTER SYSTEM ANALYST for a Medical and Surgical Clinic located in Anahuac, TX. Applicant must possess a Bachelor's Degree in Computer Information System. Job duties are to design and develop computer systems by configuring hardware & software & devise ways to apply existing systems to tasks. Must have knowledge of VB6.0, MySQL, Windows Server2000/ 2003, Revenue Mgt & Practice Mgt System & EMR system. Compensation based on experience. Respond by resume only to: Dr. Leonidas Andres, Job Code CSA001, Andres Medical Clinic, PO Box 1470, Anahuac, TX 77514.

.NET APPLICATION DEVELOPER (Chicago, IL) Write, review or re-write Microsoft Excel add-in prgm to support daily financial analyst & processing; Enhance, upgrade, existing financial reports publishing web application using Classic ASP, ASP.NET, JavaScript/JQuery, AJAX, XML and SQL. Reqs: MS in comp sci., s/ware tech. or closely rel., 24 mths exp. as comp prgmr/analyst. Resumes to Daniel Smereczynski, VP, First Analysis Securities Corp, One South Wacker Dr., Ste 3900, Chicago, IL 60606.

PURDUE UNIVERSITY. Tenure-Track/ Tenured Faculty Positions: The Department of Computer Science at Purdue University is entering a phase of significant growth, as part of a university-wide Purdue Moves initiative. Applications are being solicited for tenure-track and tenured positions at the Assistant, Associate and Full Professor levels. Outstanding candidates in all areas of computer science will be considered. Review of applications and candidate interviews will begin early in October 2014, and will continue until the positions are filled. The Department of Computer Science offers a stimulating and nurturing academic environment with active research programs in most areas of computer science. Information about the department and a description of open positions are available at http://www.cs.purdue.edu. Applicants should hold a PhD in Computer Science, or related discipline, be committed to excellence in teaching, and have demonstrated excellence in research. Successful candidates will be expected to conduct research in their fields of expertise, teach courses in computer science, and participate in other department and university activities. Salary and benefits are competitive, and Purdue is a dual career friendly employer. Applicants are strongly encouraged to apply online at https://hiring.science.purdue.edu. Alternatively, hardcopy applications can be sent to: Faculty Search Chair, Department of Computer Science, 305 N. University Street, Purdue University, West Lafayette, IN 47907. A

background check will be required for employment. Purdue University is an EEO/AA employer fully committed to achieving a diverse workforce. All individuals, including minorities, women, individuals with disabilities, and protected veterans are encouraged to apply.

SPRUCE TECHNOLOGY, INC., Clifton, NJ based IT firm, is seeking multiple candidates for following positions: **Sr. Programmer Analyst:** Develop & write computer programs to store, locate, & retrieve specific docs, data, & info; Analyze user needs & software req's to determine feasibility of design w/in time & cost constraints; Design, develop & implement the next gen platforms using tools & software w/back-end databases to provide an integrated management system. Will use a combination of T-SQL, PL/SQL, ProClarity, SSRS/SSAS/SSIS, Microstrategy, Informatica, OLAP, C++, ASP, Java Script, VB.Net, MDX. Masters in Engg (any) CS, Science (any), Comp Application w/ 1 year of related exp is req'd. Will accept Bachelor's degree w/5 yrs of related exp as equal to Master's degree. **Sr. Business Systems Analyst:** Design, test & conduct tech writing of software apps; Analyze, Plan & Develop Business Programs; Manage resources in accordance w/project schedule; Gather & synthesize business req's & translate the business req document; Design methodology & programs to ensure that the project deliverables meet industry best practices & standards; Review & Modify Software programs to fulfill desirable accuracy & reliability of programs; Coordinates & link computer systems within an organization to increase compatibility so information can be shared. Use Java, J2EE, JSP, Websphere, Weblogic, Oracle, SQL Server, PL/SQL. Bachelor degree in B. Admin (any), Science (any) or Comp. Sci. w/5 yrs of related occupation exp is req'd. Will accept combination of degrees that is equal to Bachelors. **Sr. Software Engineer:** Gather & analyze tech req's; Design & develop software programs; Develop interaction models & user flow diagrams; Resp for software devel life cycle incl analysis, design, devel, implementation, & support; Write database queries & involve in testing the application. Will use a combination of J2EE, JavaScript, Ajax, Servlets, Beans, Hibernate, Springs, ApacheAxis, JBoss-Seam, JBossCache, JBuilder, Eclipse, Flex Builder. Bachelor'sin Engg (any) CS, Science (any), Computer Application w/5 yrs of exp in the related field is req'd. Any combination of education equal to Bachelors is acceptable. **Database Engineer:** Resp for analysis, modeling, design, devel & implementation of relational database and data warehousing systems; Conceptualize & communicate enterprise data architecture frameworks for global enterprises; Resp for programming & developing database packages, functions, & triggers. Will use a combination of ETL tools DataStage, SSIS, Informatica, AbInition, TeraData, Business Objects, MicroStrategy, SQL*Loader, ODBC, Toad, Visual Basic, Java, Korn Shell and Unix Shell Scripting. Bachelor in Engg (any), CS, Science (any), Computer Application w/2 yrs of exp in the related field is req'd. Any combination of education equal to Bachelors is acceptable. **Sr. SAP Analyst:** Analyze, develop & maintain of software apps by using SAP HCM; Devel functional business process specs, document technical solutions & maintain business process procedures; Setup SAP-Success Factors Compensation activities; Configure SAP R/3 system for complete testing & internal order functionality; Involve in all phases of project life cycle using ASAP methodology; Involve in data migration by using SAP LSMW tools; Work on enhancements/interfaces and custom development with EDI interface. Will use SAP R/3, HPQC, SAP Service Market Place, Success Factors Cloud Support Portal. Bachelor in Science (any), B. Admin (any), CIS with 5 years of related occupation experience is required. Will accept combination of degrees that is equal to Bachelors. Apply with 2 copies of resume to Spruce Technology Inc., 1149 Bloomfield Avenue, Suite G, Clifton, NJ 07012.

UCF CENTER FOR RESEARCH IN COMPUTER VISION. Multiple Assistant Professor Positions. CRCV is looking for multiple tenure-track faculty members in the Computer Vision area. Of particular interest are candidates with a strong track record of publications. CRCV will offer competitive salaries and start-up packages, along with a generous benefits package offered to employees at UCF. Faculty hired at CRCV will be tenured in the Electrical Engineering & Computer Science department and will be required to teach a maximum of two courses per academic year and are expected to bring in substantial external research funding. In addition, Center faculty are expected to have a vigorous program of graduate student mentoring and are encouraged to involve undergraduates in their research. Applicants must have a Ph.D. in an area appropriate to Computer Vision by the start of the appointment and a strong commitment to academic activities, including teaching, scholarly publications and sponsored research. Preferred applicants should have an exceptional record of scholarly research. In addition, successful candidates must be strongly effective teachers. To submit an application, please go to: http://www.jobswithucf.com/postings/34681 Applicants must submit all required documents at the time of application which includes the following: Research Statement; Teaching Statement; Curriculum Vitae; and a list of at least three references with address, phone numbers and >>email address. Applicants for this position will also be considered for position numbers 38406 and 37361. UCF is an Equal Opportunity/Affirmative Action employer. Women and minorities are particularly encouraged to apply.

ENGINEERING. Zscaler, Inc. is accepting resumes for the position of Senior QA Engineer in San Jose, CA. Lead a project team of engineers to design, develop and test company software products, including protocols, such as TCP/IP, HTTP, SSL and Encryption technologies, to provide cloud security. Mail resume to Zscaler, Inc., Staffing Department, 110 Baytech Drive, Suite 100, San Jose, CA 95134. Must reference Ref. SQE-GB.

CAREER OPPORTUNITIES

SUNFIX TECH reqs (1) Data Warehouse BI Specialist to provide data management, data warehousing & business intelligence solutions. ETL/DataWarehousing Exp req. 2) Sr. Programmer Analyst to analyze user req's, design & devel custom software apps. Java/J2EE exp req. Positions require Masters (Engng/Comp. Sci) +1 year exp; BS +5 years exp can be substituted for the MS degree req. Any combination of foreign edu equiv to a US MS or BS will be considered. Position req's extensive travel to client sites. Send res to Sunfix Technologies, Inc. at 15 Corporate PL South, Suite 140, Piscataway, NJ 08854.

EXPEDIA, INC. currently has openings for the following opportunities in our San Francisco, CA office (various/levels/types:) Software Engineers: (728.SWE-SF) Design, implement, and debug software for computers including algorithms and data structures. Send your resume to: Expedia Recruiting, 333 108th Avenue NE, Bellevue, WA 98004. Must reference position and Job & Job ID# listed above.

MEMBER OF TECHNICAL STAFF/SERVER DEVELOPER. (San Francisco) S/ware co. seeks individual to join our server dvlpmt team & create scalable, high qlty, & secure server s/ware. Duties will incl: analyzing the product reqmts in order to come up w/best possible solution, while ensuring product reliability; bldg infrastructure to manage metadata for millions of user's social contacts & handling them in a way which maintains user's privacy; bldg highly scalable, realtime applics using the appropriate prgmg lang (e.g. Java, Go, Shell script, etc), & dsgn patterns, as reqd by the task; dsgng complex domain models which can be easily persisted to our long term data stores, incl various data store backends (e.g. MySQL, Apache Lucene, or other NoSQL backend); dvlp tools to automate & simplify both dvlpmt & production tasks; dsgn & dvlp end-to-end svcs & frameworks that can be used by our IOS & Android clients, or on the web; write comprehensive test cases & run them on Continuous Integration tools (e.g. Bamboo, Hudson, etc.) to ensure code qlty & reliability. Reqmts: Master's deg in Comp Sci or rltd disciplines, + 2 yrs exp or Bachelor's deg + 6 yrs. Alternate combo of education/exp will be considered. Send resume to Humin Inc., 655 Montgomery St, Ste 1950, San Francisco, CA 94111. Fax 212-619-5210.

EMPLOYMENT OPPORTUNITIES (UNIONDALE, NY): SR. PROGRAMMER ANALYSTS: Develop & write computer programs to store locate & retrieve specific documents, data & info; Implement web applications using ASP.NET w/VB.NET on an oracle back-end & writing PL/SQL. Develop & direct software system testing & validation procedures, programming, & documentation; Compile & write documentation of program development & subsequent revisions, inserting comments in the coded instructions. Will use ASP.NET, SQL Server, VB.NET, Visual Studio. Bachelor Deg in Engineering (any), Comp. Sci., Science (any), MIS w/ 5 yrs of exp in any related field is req'd. **SR. SYSTEMS ADMINISTRATORS:** Review & Test all EDI transaction processing output & other applications systems; Supervise EDI data flow & managing operation; Setup & configure EDI Systems & coordinate w/Utilities for tests & live production systems; Perform data backups & disaster recovery operations; Administer, configure & maintain system applications & network environment; Troubleshoot & resolve hardware, software, or other network & system problems. Master's Deg in Comp. Sci., Engineering (any), Science (any), MIS is req'd. Resumes to EC INFOSYSTEMS, INC. 50 Charles Lindbergh Blvd. Suite 411, Uniondale, NY 11553.

ENGINEERING. PIVOTCLOUD INC. is accepting resumes for the following positions in Sunnyvale, CA: Principal Software Engineer (Ref. PSE-SS): Architect, dsgn & dvlp compliance, content monitoring, supervision, surveillance, legal discovery, & litigation support applications on mult client-side devices, w/ extensive use of C/C++/Objective C/Java/C#, Windows Phone,.Net, SQL Server, SharePoint Server, Exchange Server, & other rltd Microsoft technologies. Software Engineer (Ref. SE-LK): Architect & dsgn high scale distributed cloud services. Dsgn & implem't security S/W for data storage & data mgmt in the cloud. Position based at co. headquarters in Sunnyvale, CA, or can be home based in continental U.S. Mail resume to PivotCloud, Inc., Staffing Dept, 530 Lakeside Dr, Ste 180, Sunnyvale, CA 94085. Must reference Ref. Code.

LEAD ARCHITECT. (Atlanta, GA & various unanticipated locations throughout U.S.): Plan engineering activities of global bus. integration sols. projs. REQ Bach. deg., (or for. deg. equiv.), in Mgt. Info. Syst. or Comp. Sci. & 4 yrs. exp. in prov. intgr., dev. & consulting for B2B sfwr apps. Stated exp. must incl. 2 yrs. exp. in customer-site implem. of sfwr projs. & app. intgr. Ext. dom. / int'l travel is req. (50%). (Job Code "CG".) **PROJECT MGRS.** (mult. openings) (Atlanta, GA): Dir. R&D & plan engineering activities of global bus. integration sols. projs. REQ Master's deg. in Industrial Engin. Aca. crswrk in Probabilistic Models; & Deterministic Optimization is req. Ext. dom. / int'l travel is req. (50%). (Job Code "PM".) Send resume w/Job Code "PM" or "CG" to Dan Marischuk, Seeburger, Inc., 1230 Peachtree St, NE, #1020, Atlanta, GA 30309.

NISUM TECHNOLOGIES. has multiple openings for the following positions at its office in Brea, CA. *Technical Lead: Design, develop and test software systems. *Sr. QA Engineer: plan and conduct analysis, inspection, design, test and/or integration to assure quality of projects. *QA Engineer: Assist in analysis, design, test and/or integration of projects. *Sr. Application Developer: design, develop, maintain & test software applications/systems. Analyze customer requirements & custom design systems as needed. *Database Administrator: Install, maintain, administer & troubleshoot databases. *Programmer Analyst: Analyze, develop & write codes to implement system applications. Job requires min. of M.S./foreign equiv. + exp., M.S./foreign equiv., or B.S./foreign equiv. + exp.. Education/Experience requirements vary depending on position level/type. Travel/relocation required. Send resume and salary history & position applied for to: Nisum Technologies, 500 S. Kraemer Blvd., Brea, CA 92821. Attn: H.R. Manager.

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

CLOUDERA, INC. is recruiting for our Palo Alto, CA office: Software Engineer: maintain a thorough understanding of the Hadoop eco-system that Cloudera ships with CDH. Create test cases for testing CDH components with Hue & run those tests to validate different CDH releases in a timely manner SE-DV to: Attn.: HR, Cloudera, 1001 Page Mill Rd., Bldg. 2, Palo Alto, CA 94304.

RIVERBED TECHNOLOGY, INC. has opening for Network and Application Performance Engineer (Job # 2015-1816) in Bethesda, MD (telecommuting from anywhere in U.S. is acceptable). Duties: Plan, design, deploy, operate and optimize service provider/large enterprise infrastructures. For more info and to apply, go to http://www.riverbed.com/us/careers/ and reference the Job#.

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

Information Systems Manager [REQ#9FRU9X]. Mng team of Architects & DBAs for Apple Information Sys.
Software Engineer Applications [Req#9D32DG]. Research & dev large-scale Cloud-based productivity app suite across mult platforms.
Software Development Engineer [Req#9JPMUV]. Conduct SW testing for all iOS apps.
Systems Design Engineer [Req#9DE444]. Resp for multi-radio co-existence perform eval, data analysis, & design optimization on a wide variety of projects for various wireless techs.
Software Development Engineer [Req#9CYPUL]. Develop & maintain test frameworks for HTTP Live Stream & progress dwload techs for iOS & OSX platforms.
Software Development Engineer [Req#9AQL67]. Design & dev sw & add new customer facing features.
Software Quality Assurance Engineer [Req#9AYSSN]. Develop, run & maintain SW tests for communication systems.
Software Engineer Applications [Req#9H4NPQ]. Design & develop end-to-end advanced analytic solutions on core Apple data sets.
Software Engineer Applications [Req#9H3TFZ]. Design & dev the next generation of Apple's Employee Syst platform & suite of products. Travel req'd: 25%.
Software Quality Assurance Engineer [Req#9K6SJ3]. Create & execute test plans for vid encoder, pre-processor, post-processor & vid algorithm module used in various Apple apps.
Information Systems Engineer [Req#9GJ364]. Design, dev, implement & maintain EAI, internal Cloud, DB & J2EE app syst.
SAP Performance Analyst [Req#9GFVWS]. Perform SAP performance testing, troubleshooting, and tuning.
Systems Design Engineer [Req#9GPR5M]. Perform RF System Design Validation and Debug for wireless telecommunication systems.
Computer Vision Research Engineer [Req#9DCSRS]. Design & dev algorithms & SW for Computer Vision sys.
Software Development Engineer [Req#9GJRL6] Lead team of build engs to sup ongoing SW builds of major rel of iOS & OS X operate sys.
Software Development Engineer [Req#9DNPSZ] Des, dev & imple SW for routing service.
Systems Design Engineer [Req#9DGTRY] Dev, opt & debug calibration & perform verif stations. Travel req. 20%.
Systems Engineer [Req#9E52CY] Build & trouble comp server sys.
Software Development Engineer [Req#9CXSES] Des & dev natural lang process tech for local Apple prod into int'l markets & help scale art intel apps including Siri to new lang.
Software Engineer Applications [Req#9P92BS] Define & evan OS X/1OS diag data analytic sys.
Operations Project Engineer, New Products [Req#98F49S]. Lead OEM oper. team. Plan & execute dev builds & new prod ramps. Travel req'd 30%.
ASIC Design Engineer (IC Packaging Engineer) [9E5VV6]. Perform IC pckg mech simulation & characterization.
Software Development Engineer [9GXRKD]Serve as a member of the core op sys network team. Provide SW dev in comp network tech.
Hardware Development Engineer [Req #9D2VXQ]. Respon for design & develop of new prods. Design display tech including new pixel & circuit. Travel req 20%
Software Development Engineer [Req#9F4SZH] Design complex distribute sys w/ an emphasis on high avail & perform.
Systems Architect [Req #9AYN8F]. Respon for providing tech guide to drive exec of multi projects. Dsgn high perform, scalable sw sols to supp high demand of daily trans.
Software Engineer Applications [Req#9J2TXF]. Perfrm data mining & rsrch of data w/ focus on developing automated processes & tools to surface actionable data.
Software Engineer [Req#9F4SQW]. Dsgn & dvlp app interfaces, database interfaces, & SW layer abstractions. Lead dvlpmnt of server-side SW components, data persistence, & caching components.
ASIC Design Engineer [Req#9CYUG2]. Prfrm regular STA runs & anlyz, triage & deliver timing results. Validate netlist, parasitics, constraints or other input collateral for quality.
Systems Design Engineer [Req#9E8Q4M]. Dsgn & evaluate Radio Frequency System for iPhone and iPad.
Software Development Engineer [Req#9CCN8P]. Design & develop Network Protocol Software for desktop & mobile platforms.
Software Architect [Req#9GWPE4]. Dsgn, dvlp, & deploy SW apps for Data-warehousing / Business Intelligence Projects in the Apple IST Marketing Dept.
Reliability Engineer [Req#9E23BV]. Guide design teams in creating reliable designs for novel HW technologies.
Hardware Development Engineer [Req#9AE2X8]. Dvlp new designs, panel processes & optics. Lead engineering investigation on new concepts. Travel Required 25%.
Hardware Development Engineer [Req#9DF2GM]. Respon for the design & develop of baseband power mgmt. sols for future wireless products.
Software Development Engineer [REQ#99WSUF]. Perform initial design & dev of apps. Share expertise in app & framework devlpmt.
Product Design Engineer [Req# 9EZW4P] Dsgn & dvlp materials & processing for consumer electronics products.
Software Development Engineer [Req#9FSPTK]. Manage end-to-end lifecycle(s) of complex Wireless functionality in Apple products. Write test plans, test cases, develop automation & ad-hoc testing.
Software Engineer Applications [Req #9GASLJ]. Respon for test sw that forms foundation for iCloud products & srvcs.
Software Development Engineer [Req#9K2VU3] Dev, des & implmt, architect for SW components. Write code used in maps search svs.

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

It's work that matters. It's what we do at Symantec. Symantec is the world leader in providing solutions to help individuals and enterprises assure the security, availability, and integrity of their information. In essence, we protect the free flow of information in a connected world. As the fourth largest independent software company in the world, Symantec has operations in more than 40 countries with 475 out of Fortune's global 500 companies using our solutions.

People look to us to safeguard the integrity of their information, ensuring it is secure and available. Achieving this ambitious goal is only possible through the combined efforts of the innovators and visionaries that Symantec continuously attracts. Symantec draws the very best people with a variety of backgrounds, experiences and perspectives and provides them with a work environment where uniqueness is valued and empowered. The creative people we attract help define the spirit of innovation at Symantec. Symantec is proud to be an equal opportunity employer.

We currently have openings for the following positions (multiple positions/various levels/types):

Columbia, Maryland

Software Engineers (SWEMD115) Responsible for analyzing, designing, debugging and/or modifying software; or evaluating, developing, modifying, and coding software programs to support programming needs.

Culver City, California

Computer Systems Analyst (CSACC115) Analyze engineering, business and/or other business intelligence issues for application to Symantec solutions; and/or provide operational support in the development and implementation process of computer software applications, systems or services.

Database Administrator(DBCC115) Support both development and production environments. Responsible for primary application/database and working with DBA team. The applications supported by these databases are from a wide variety of vendor provided to in-house developed apps.

IT Infrastructure Specialist (ITSCC115) Manage large complex infrastructure by designing, planning, and implementing complex infrastructure systems. Establish and recommend policies and standards on system use and services and automate monitoring or periodic preventative maintenance processes; or Analyze user requirements, procedures, and problems to automate or improve existing systems and review computer system capabilities, workflow, and scheduling limitations.

Knowledge Engineer (KECC115) Build, maintain and use knowledge-based systems. Collect and analyze data for projects and departmental needs, create reports, scorecards and dashboards to analyze performance and results of projects and on-going business. Work with and support projects that require the collection and analysis of data.

Security Infrastructure Administrator (SIACC115) Perform system and database administration for the ongoing maintenance of security network architecture and systems. Maintain multiple production, development and QA SQL Server environments.

Software Engineers: (SWECC115) Responsible for analyzing, designing, debugging and/or modifying software; or evaluating, developing, modifying, and coding software programs to support programming needs.

Software QA Engineers: (SQACC115) Responsible for developing, applying and maintaining quality standards for company products. Develop and execute software test plans. Analyze and write test standards and procedures

Herndon, Virginia

Computer Systems Analyst (Mgr, Ops) (1648.1513) Prioritizes & tracks MSS customer analysis issues. Requests and analyzes new functionality.

Engineering Managers (EMVA115) Direct and supervise team of engineers (QA and/or development teams); Develop standards for products and/or oversee development and execution of software and/or analysis of test results.

Houston, Texas

Software QA Engineers (1648.1589) Responsible for developing, applying and maintaining quality standards for company products. Develop and execute software test plans. Analyze and write test standards and procedures.

Lindon, Utah

Senior Technical Education Staff Member (1648.609) Deliver skills/technical training for Symantec products and conduct trainer events. Develop curriculum design for big picture view of learning and development. Must be available to work on projects at various, unanticipated sites throughout the United States. May Telecommute.

Submit resume to JOBADS@symantec.com . Must reference position & code listed above. EOE.
For additional information about Symantec and other positions visit our website at http://www.symantec.com.

Symantec™

It's work that matters. It's what we do at Symantec. Symantec is the world leader in providing solutions to help individuals and enterprises assure the security, availability, and integrity of their information. In essence, we protect the free flow of information in a connected world. As the fourth largest independent software company in the world, Symantec has operations in more than 40 countries with 475 out of Fortune's global 500 companies using our solutions.

People look to us to safeguard the integrity of their information, ensuring it is secure and available. Achieving this ambitious goal is only possible through the combined efforts of the innovators and visionaries that Symantec continuously attracts. Symantec draws the very best people with a variety of backgrounds, experiences and perspectives and provides them with a work environment where uniqueness is valued and empowered. The creative people we attract help define the spirit of innovation at Symantec. Symantec is proud to be an equal opportunity employer.

We currently have openings for the following positions (multiple positions/various levels/types):

Mountain View, California

Agile Product Owner (1648.1995) Work with a cross-functional team of UI developers, architects, designers and QA testers to ensure that the website is intuitive and designed with a long lifespan in mind in a mobile-friendly way.

Computer Systems Analyst (CSAHQ115) Analyze engineering, business and/or other business intelligence issues for application to Symantec solutions; and provide operational support in the development and implementation process of computer software applications, systems or services.

Data Scientists (DSHQ115) Design, develop, and program methods, processes, and systems to consolidate and analyze unstructured, diverse "big data" sources to generate actionable insights and solutions for client services and product enhancement.

Data Science Director(DADHQ115) Develop and code software programs, algorithms, and automated processes to cleanse, integrate, and evaluate large datasets from multiple disparate sources.

Engineering Manager(EMHQ115) Direct and supervise team of engineering (QA and/or development teams). Develop standards for products and/or oversee development and execution of software and/or analysis of test results. Plan, design, develop and implement processes.

Network Systems Engineer (NSEHQ115) Design, architect, and maintain network systems. Perform tasks related to network engineering, planning, and configuration.

Oracle Database Administrator (1648.2419) Responsible for handling Oracle EBS 11i and R12 environments along with related environments which depends on ERP systems. Support both development and production environments.

Product Managers (PDMHQ115) Develop company market requirements for technical products or product lines, including product strategy definition, requirements analysis, and/or pricing.

Product Marketing Manager (PMMHQ115) Develop product marketing strategy to drive product demand. Plan the launch of new products and releases and manage the cross-functional implementation of the plan.

Production Specialist (1648.1283) Involvement in all phases of eCommerce data entry, collection, generation and validation. Assist in data processing and maintenance of the catalog merchandizing & pricing rules for the online stores.

Search Engine Marketing (SEM) and Display Managers (SEMHQ) Responsible for search marketing strategies and plans for a portfolio of regions within the global team. Assist with developing pay-per-click (PPC) account strategies and roadmaps.

Senior Principle Project Manager Specialist (1648.1119) Manage the scrum activities of two eBusiness Teams. Serve and support the Product Owner and Development Team in their quest to do everything possible to delight customers.

Software Engineers (SWEHQ115) Responsible for analyzing, designing, debugging and/or modifying software; or evaluating, developing, modifying, and coding software programs to support programming needs.

Software QA Engineers (SQAHQ115) Responsible for developing, applying and maintaining quality standards for company products. Develop and execute software test plans. Analyze and write test standards and procedures.

Web Developers (WEBHQ115) Design and develop web applications and websites; create and specify architectural and technical parameters. Designing and implementing of the PC Tools website (and associated websites.

San Francisco, California

Program Manager (1648.2444) Work closely with engineering members, managers and leads, product managers, ensure rapid execution and on time, high quality delivery of complex Data Loss Prevention (DLP) projects.

Program Manager (Product Lifecycle Engineer) (1648.303) Participate in all software product development life cycle activities. Move software products through the product development cycle from design and development to implementation and testing.

Software Engineers (SWESF115) Responsible for analyzing, designing, debugging and/or modifying software; or evaluating, developing, modifying, and coding software programs to support programming needs.

Submit resume to JOBADS@symantec.com . Must reference position & code listed above. EOE.
For additional information about Symantec and other positions visit our website at http://www.symantec.com.

Symantec™

EDITOR DAVID ALAN GRIER
George Washington University, grier@gwu.edu

VIDEO

The Tyranny of Geography

David Alan Grier, George Washington University

The software industry is more regional than one would think; most jobs in the field are centered in nine US cities.

When Matt's job search reached its third year, his family became convinced that his studying computer science in college was a mistake. "If computer science is such a good thing to study," his grandmother sniffed, "why can't it get him a job?" I wasn't an intimate friend of Matt's, but as his neighbor, I'd seen him grow from toddler to teenager to tech geek. As a gesture to the family, I offered to meet with the young man to see what I could learn about his job prospects.

The idea that someone might have made a bad career choice by studying computer science is almost heretical in the modern world. But even though statistics point to strong demand and high salaries for software developers, I regularly see recent graduates move out of the field after a single job, or even fail to secure an initial position in the industry. As students, they talked freely about their career aspirations, but found something thwarting those ideals during their job search.

When Matt and I met, he blamed his jobless status on the economy. "Jobs are hard to find," he remarked. Indeed, the economy was still recovering from the recent recession and wasn't creating many new jobs.

As Matt described the kind of position he desired, I realized he was actually fighting the tyranny of geography. He outlined the qualities he wanted in a job, preferring one in a region that was close to family and friends, recreation, and his favorite sports teams. But as far as I could tell, that region had few jobs such as he described.

In spite of its global reach, the software industry is more regional than its advocates profess. If you look at software development positions in the US, you'll find that 40 percent are found in just nine cities. Even within those nine, the jobs are unevenly dispersed. Roughly 10 percent of all system software jobs are located around the San Francisco Bay Area. Approximately the same percentage of application programmers work near Puget Sound in Washington State. Finally, about 12.5 percent of all positions in computer security—the most rapidly growing category of software jobs—are a short drive from the Potomac River Basin in Washington, DC.

Perhaps it isn't surprising that computing jobs are concentrated in specific regions. Since the start of the industrial economy, we have seen regions become centers for specific kinds of products: Chicago for meat, Detroit for automobiles, North Carolina for furniture. However, these regions became industry centers because they provided some key input that wasn't readily available in other parts of the country: hogs in Chicago, steel in Detroit, wood in North Carolina. Software requires no such input. It depends only on ideas, skills, and access to computers, all of which are now widely available. Indeed, the early pioneers of the Internet argued that computer networks would "make every local resource available" to a worldwide audience.

Yet, for the moment, networks haven't broken the tyranny of geography. Regional centers are of "striking importance" to national wealth, argues economist Michael Porter. Programming skills don't seem to be as widely distributed as we might have thought, or perhaps people with such skills are attracted to certain kinds of environments.

When Matt finally found a software development job, he was able to avoid the nine US cities. But his job also fell into a cluster, though smaller than Silicon Valley or Seattle or Northern Virginia. His job was in a small city in the middle of the country. It was close to a university and next to a highway that leads to the lakes up north. It was also the only place in his state where development jobs might be found.

DAVID ALAN GRIER is a professor at the Center for International Science & Technology Policy at George Washington University. Contact him at grier@gwu.edu. A podcast version of these essays can be found at erranthashtag.dagrier.net.

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**.

0018-9162/15/$31.00 © 2015 IEEE