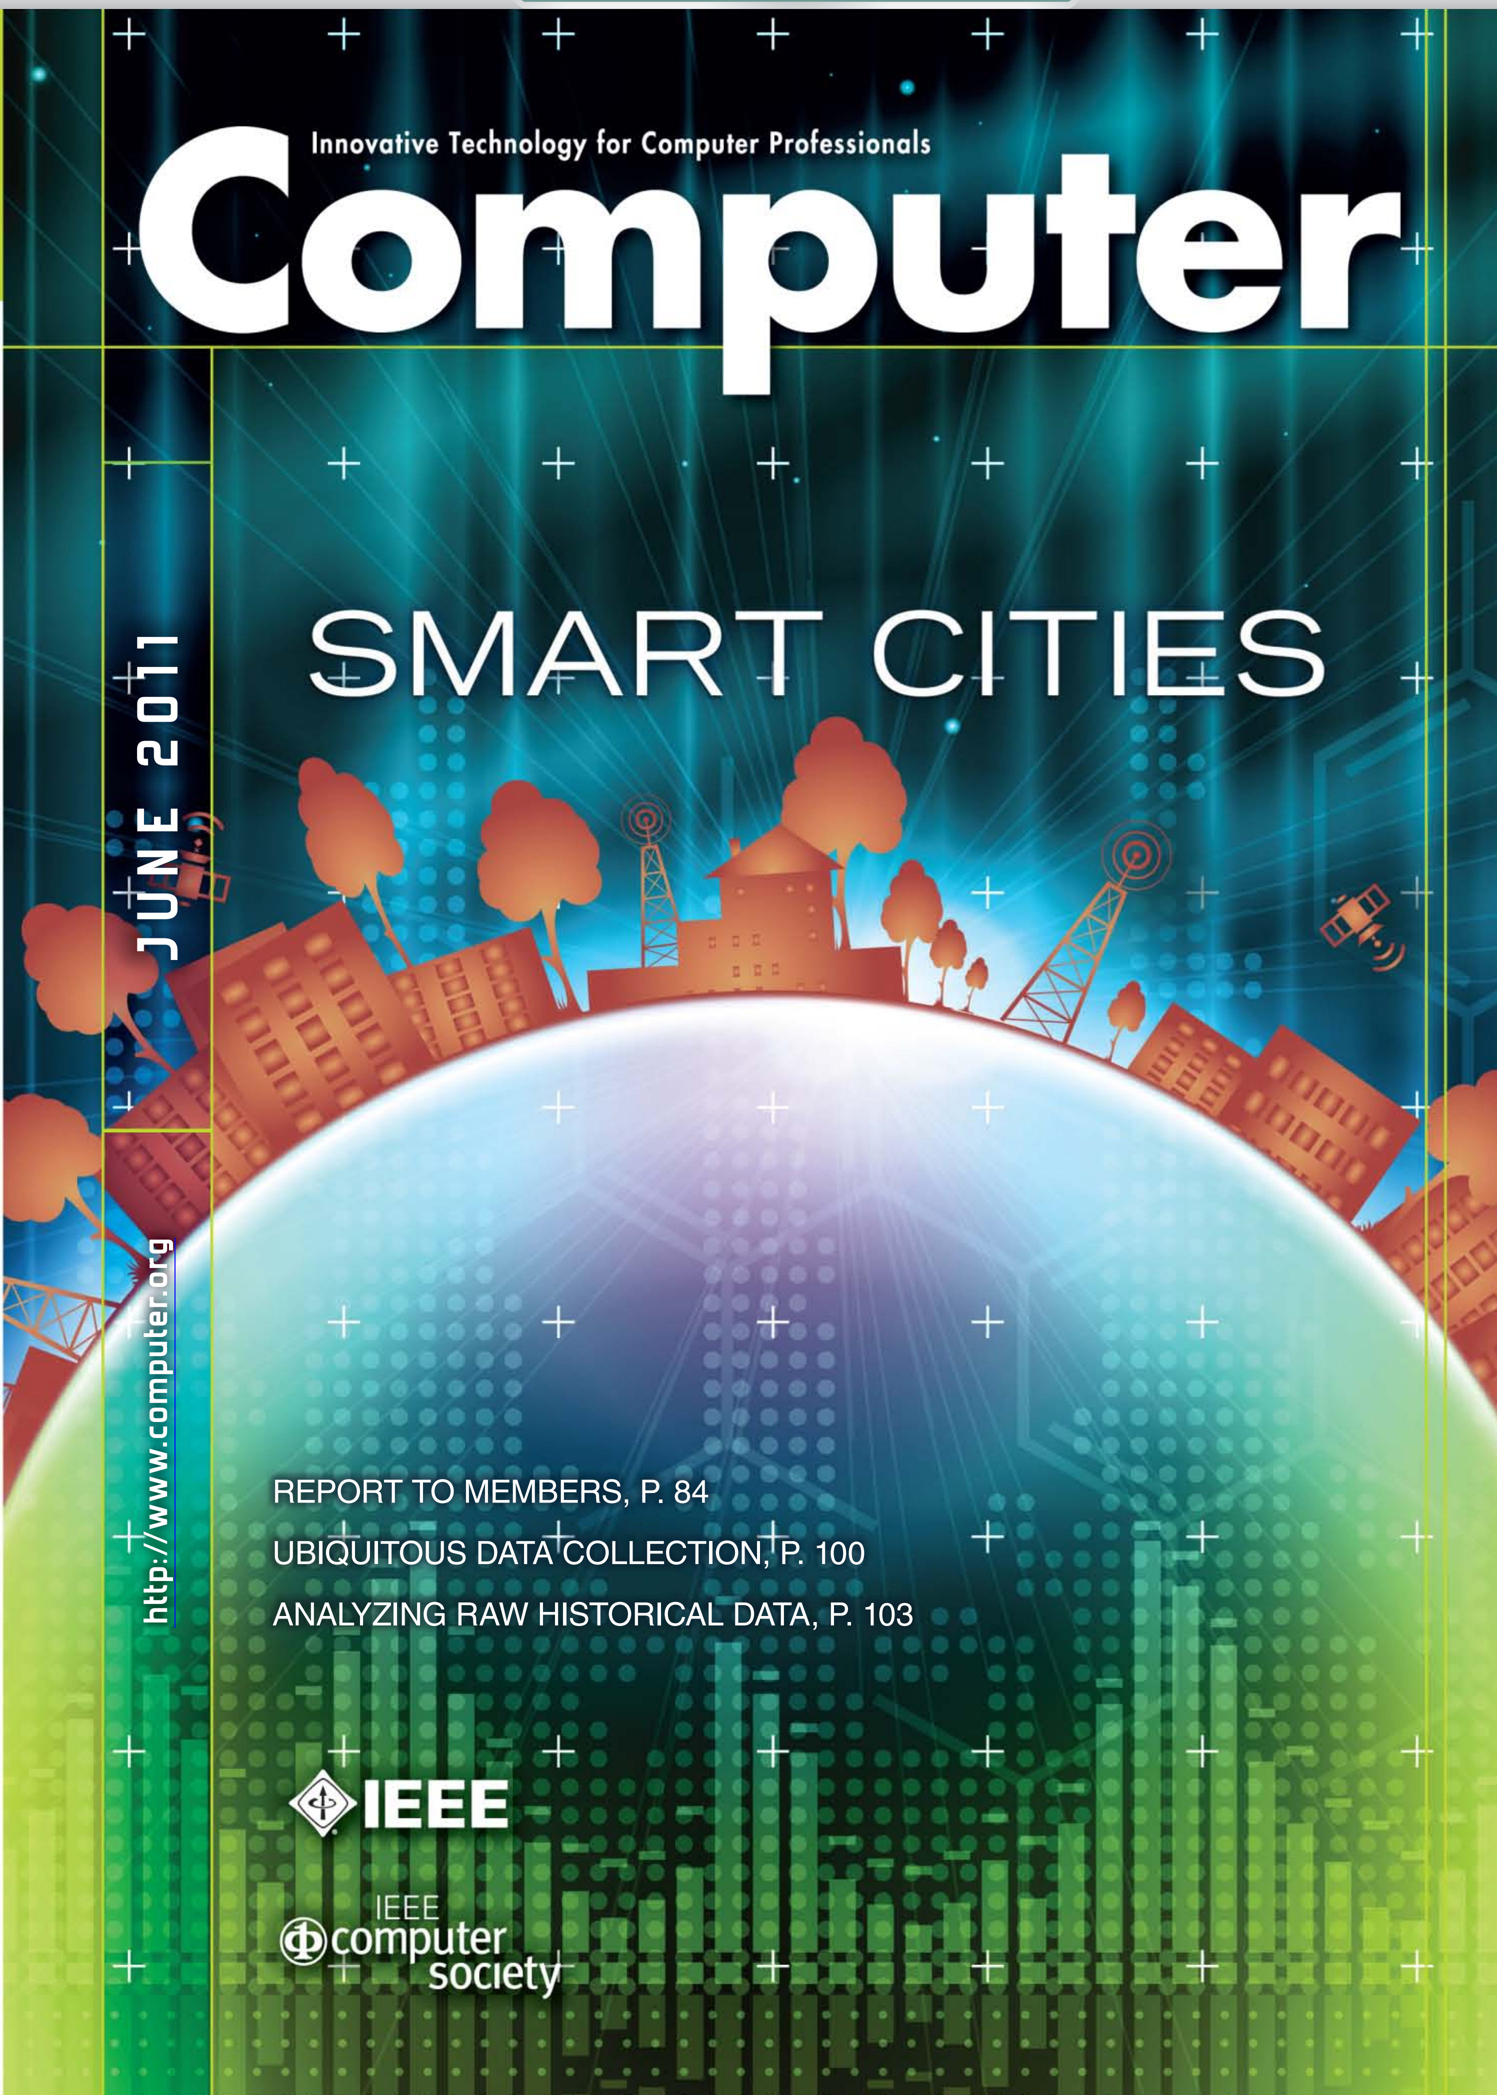

Innovative Technology for Computer Professionals

Computer

JUNE 2011

http://www.computer.org

SMART CITIES

REPORT TO MEMBERS, P. 84

UBIQUITOUS DATA COLLECTION, P. 100

ANALYZING RAW HISTORICAL DATA, P. 103

IEEE

IEEE computer society

Computer JUNE 2011 Crowdsourcing//Smart Cities//Ring Generators Volume 44 Number 6

Contents | Zoom in | Zoom out | For navigation instructions please click here | Search Issue | Next Page

You're a smart, experienced technology professional. You take pride in your work. You pay attention to detail.

And still you're at risk of a professional liability lawsuit that could result in enormous costs, as well as compromise your reputation and livelihood. Protect your assets and reputation against the cost and serious consequences of a liability claim with the Professional Liability Insurance Plan for IEEE members in the United States.* And now, IEEE Professional Liability insured members in the United States who operate a computer-related professional services business out of their home are eligible for General Liability insurance so they can be protected against bodily injury, personal injury and advertising injury claims as well.

Are you wondering if this insurance is right for you? Read more about what fellow IEEE members had to say about this plan by visiting www.ieeeinsurance.com/liability today.

IEEE

Snapshot Comparison of Professional Liability vs. General Liability

Type of insurance	Your coverage benefits	Types of claims
PROFESSIONAL Liability: Coverage for the professional services you provide	Reimbursement of legal fees & compensatory damages Access to legal assistance	Negligence Misrepresentation Violation of good faith Inaccurate advice Error or omission Failure to render professional services
GENERAL Liability: Coverage for your business' property, including your home	SAME AS ABOVE	Bodily injury Personal injury Advertising injury

"The PRICE WAS BETTER than the other policies we looked at. I've had questions about the policy and rates and each time I made an inquiry, I received a prompt response."

Stephen Richter, Richter & Associates

To learn more, call 1.800.375.0775 or visit www.ieeeinsurance.com/liability

*The IEEE Professional Liability and General Liability Insurance Plans are underwritten by Certain Underwriters at Lloyd's of London and are only available to active members who reside in the United States of America.

This program is administered by Marsh U.S. Consumer, a service of Seabury & Smith, Inc., d/b/a in CA Seabury & Smith Insurance Program Management. CA Ins. Lic. #0633005. AR Ins. Lic. #245544. IEEE prohibits discrimination, harassment and bullying. For more information, visit www.ieee.org/web/aboutus/whatis/policies/p9-26.html.

IEEE members in the U.S. can scan to request a FREE Risk Assessment and Review Toolkit.

Computer
Innovative Technology for Computer Professionals

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 90720-1314; voice +1 714 821 8380; fax +1 714 821 4010; **IEEE Computer Society Headquarters**, 2001 L Street NW, Suite 700, Washington, DC 20036. IEEE Computer Society membership includes $19 for a subscription to *Computer* magazine. Nonmember subscription rate available upon request. Single-copy prices: members $20.00; nonmembers $99.00.

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.

Innovative Technology for Computer Professionals
Computer
www.computer.org/computer

CONTENTS

ABOUT THIS ISSUE

This special issue offers four articles on advances in the development of smart cities, covering the IT perspective, applications and services, interaction and user experience enrichment, and a fast-progressing smart city initiative in South Korea.

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

IEEE Computer Society: http://computer.org
Computer: http://computer.org/computer
computer@computer.org
IEEE Computer Society Publications Office: +1 714 821 8380

Flagship Publication of the IEEE Computer Society

June 2011, Volume 44, Number 6

IEEE

IEEE computer society

Printed with inks containing soy and/or vegetable oils

Sustainable Forestry Initiative — Certified Fiber Sourcing www.sfiprogram.org

BPA Worldwide

Computer Highlights Society Magazines

The IEEE Computer Society offers a lineup of 13 peer-reviewed technical magazines that cover cutting-edge topics in computing including scientific applications, design and test, security, Internet computing, machine intelligence, digital graphics, and computer history. Select articles from recent issues of Computer Society magazines are highlighted below.

IEEE Software

The basic definition of refactoring—the process of making small behavior-preserving transformations—has remained the same over time, but its intent has varied considerably from the original purpose of improving code readability, extensibility, and maintainability. For example, an HTML refactoring such as "Turn on autocomplete" doesn't improve the code, but it does make a Web form easier to use. This shifts the refactoring's intent toward improving software product usability. In *IEEE Software*'s May/June issue, the authors of "Refactoring for Usability in Web Applications" assert that it's important to link refactorings not only to the code problems they can eliminate but also to the specific quality attributes they aim to improve.

IEEE Intelligent Systems

Users often have no direct knowledge of their intelligent devices' algorithms, data requirements, limitations, and representations. Poor user interface design, users' lack of understanding, inadequate mental models, or incorrect expectations can cause problems. Training can help. The Experiential User Guide described in "Improving Users' Mental Models of Intelligent Software Tools" in the March/April 2011 issue of *IS* addresses the genuine cognitive challenges that both novice and experienced users have with today's complex, intelligent software tools.

IEEE Computer Graphics and Applications

In "Practical Game Design and Development Pedagogy," in the May/June issue of *CG&A*, author Paul J. Diefenbach of Drexel University notes several recurring observations about student tendencies that have influenced the structure and focus of Drexel University's Game Development Workshops. Originally, the game-design part of the workshops built upon previously taught basic game theory, gameplay critique, and nonlinear-story skills. Now, it also addresses vaguely envisioned student game concepts and pitches. The game-development part of the workshops, which has employed agile development, is now structured to guide students with little experience in long-term project planning.

Computing in Science & Engineering

"Trends in High-Performance Computing," in the May/June issue of *CiSE*, looks at recent advances in supercomputer technologies. Supercomputer architectures have evolved from early custom-designed systems to the current clusters of commodity multisocket, multicore systems. Supercomputers are usually designed to achieve the highest possible performance in terms of the number of 64-bit floating-point operations per second. Twice a year, the supercomputing community ranks the systems and produces the Top-500 list, which shows the world's 500 highest performing machines. The technologies used in the top-ranked machines give a good indication of architecture trends.

IEEE Security & Privacy

Building Dependability, Reliability, and Trust

In "Shouldn't All Security Be Usable?," authors Mary Frances Theofanos of the US National Institute of Standards and Technology and Shari Lawrence Pfleeger of Dartmouth College introduce *S&P*'s March/April special issue on usability and security and provide a description of how usable security has been transformed from a desirable system property to a rich area of serious research. They explain how usability testing differs from other kinds of software testing, describe some of the key papers on which current research builds, and summarize the four articles in the special issue.

 Published by the IEEE Computer Society

IEEE Pervasive Computing

The retail experience is undergoing significant changes due to a confluence of pervasive computing technologies, such as affordable smartphones with a plethora of retail applications, social media, sensing and analytics, and wireless technologies. In "Pervasive Retail," in the April-June issue of *PvC*, guest editors Chandra Narayanaswami, Antonio Krüger, and Natalia Marmasse introduce three articles that examine some of the opportunities and challenges in the pervasive retail space.

IEEE Internet Computing

The May/June *IC* is a special issue on security and privacy in social networks. "Social networks' security and privacy requirements still aren't well understood or fully defined," write guest editors Gail-Joon Ahn of Arizona State University, Mohamed Shehab of the University of North Carolina at Charlotte, and Anna Squicciarini of Pennsylvania State University in "Security and Privacy in Social Networks." "Nevertheless, it's clear that they'll be quite different from classic security and privacy requirements because social networks involve user-centric concerns and allow multiple users to specify security policies on shared data." The issue includes four articles encompassing research advances and state-of-the-art technologies for addressing these requirements.

IEEE Micro

Micro's annual March/April "Hot Chips" issue features articles based on papers from the August 2010 Hot Chips 22 conference, a leading forum for presenting new processor architectures and system-enabling aspects of silicon such as software, I/O, and packaging. In their introduction, "Hot Chips 22" guest editors Jose Renau of the University of California, Santa Cruz, and Will Eatherton of Juniper Networks highlight eight articles—three on new processor architectures from IBM and AMD, four on high-performance computing chips and systems, and one on a 45-nm multicore research prototype that uses dark silicon to execute general-purpose smartphone applications that use 11 times less energy than today's most energy-efficient designs.

IEEE MultiMedia

Mobile and cloud computing are relatively new computing platforms. In "Virtualized Screen: A Third Element for Cloud–Mobile Convergence," in *MultiMedia's* April-June issue, authors Yan Lu, Shipeng Li, and Huifeng Shen of Microsoft Research Asia illustrate the advantages of an approach with two applications, Cloud Browser and Cloud Phone. The article describes an approach to rendering computer screens in the cloud and delivering them as images to the client for interactive display.

IT Professional

Guest editor Bruce Potter of Ponte Technologies summarizes the evolution and challenges of securing information systems over the past few decades in "Coming to Grips with Security," his introduction to *IT Pro*'s May/June special issue on IT security. "The one constant," Potter writes, "is that attackers will change tactics based on our evolving defenses." The issue features three articles addressing new ways of thinking about IT security, including virtual machine introspection, a community knowledge base, and a forensic Web services framework.

IEEE Design & Test of Computers

"Functional verification is widely acknowledged as a significant bottleneck in today's SoC design methodology," write guest editors Prabhat Mishra of the University of Florida, Zeljko Zilic of McGill University, and Sandeep Shukla of Virginia Tech in "Multicore SoC Validation with Transaction-Level Models," their introduction to *D&T*'s May/June issue. They present four articles that highlight challenges and recent trends in multicore SoC validation using transaction-level models. The articles cover theoretical as well as practical aspects of high-level validation to capture as many bugs as possible early in the design process and thereby drastically reduce the overall validation effort.

IEEE Annals of the History of Computing

The April-June issue of *Annals* presents five feature articles that editor in chief Jeffrey Yost of the Charles Babbage Institute and editorial board member William Aspray of Indiana University invited in conjunction with a 2010 workshop. Their purpose was to extend *Annals*' content to include authors who hadn't previously written for the magazine but who were conducting cutting-edge research on largely unstudied, but significant, topics in computing history. The diverse articles include histories of evidence-based medicine, slot-machine digitization, hearing-aid contributions to electronics miniaturization, and PC "tinkering." Attention to computer users and user-driven innovation emerges as a common theme among the articles, offering a new way of looking at the history of computing.

To access these articles, visit
http://computingnow.computer.org/elsewhere

Dumb Grids and Smart Markets

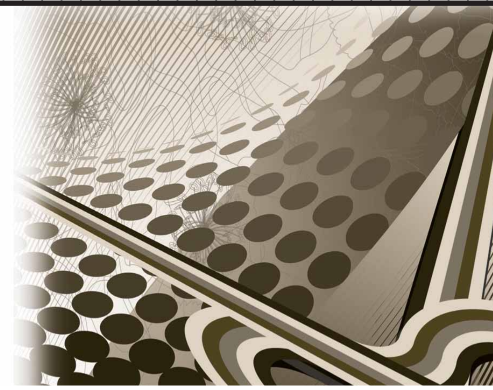

David Alan Grier, *George Washington University*

The task of building a smart grid requires us to solve two kinds of problems: technical and social.

Nothing demands a more detailed accounting of our actions than a roaring prairie fire, a glowing wall of energy that stands ready to take all that we might have to offer. Such fires occasionally occur in the region that holds my family's farm, a picturesque 12 acres in the dry foothills of the American West.

The farm has an emergency electrical generator to provide water to fight the flames. Even a modest fire could cut the connection to the power grid and make it impossible to pump anything from the aquifer that lies some 30 feet below the surface.

The generator can deliver 15 kW in a matter of seconds, but it can't begin to satisfy the full demand from the farm. The water pump takes almost 3 kW by itself. That load is compounded by the power required by other essential equipment: the air conditioner, barnyard lamps, and hay loader. If you want a hot shower and a cold drink after defeating the fire, you'll need to add a couple more kW for the refrigerator and the water heater. Without an accurate account of the electrical load, the generator will grind to a halt long before you can bring its power to bear on the fire.

In that moment before you throw the switch, when the circuits are still part of the local power grid, nobody has to make a detailed load calculation for the farm. The power demand for every family in the hills is aggregated into a single quantity. The local utility cares nothing about your individual needs until it comes time to settle the bill at the end of the month. You might need a little more electricity or your neighbor might require a little less. From the utility's point of view, these fluctuations are lost in the aggregate demand. It manages the local infrastructure for the aggregate, not for the individual. This may not be the most efficient way to manage electrical power, but it has worked fairly well for almost a century.

AGGREGATE DEMAND

More than any other issue, the ideas of aggregation and markets lie behind the current discussions of electrical infrastructure, the discussions that focus on the concept we now call the smart grid. We tend to view this as a technological problem. "Using digital sensors, computing modeling and real-time data, a smart grid would revolutionize our antiquated" electrical power system, explained a recent article in *The New York Times*. Yet, the smart grid really involves a radical change in how we view the market for electrical power, a change that will require us to be more responsible for our own actions and more dependent upon our neighbors.

For most of its history, the power market has involved a dialogue between two parties: centralized generators and aggregate populations. In its earliest days, this dialogue was best described by utility magnate Samuel Insull. After starting his career as an assistant in Thomas Edison's company, Insull launched his own firm to create the large electrical utilities. In forming those utilities, he recognized that aggregate demand was more uniform than the demand from any individual household and thus "the fundamental basis of the profit-making of an energy-selling company."

Insull promoted his ideas with a set of charts and diagrams that became well-known in the early electrical industry. Aggregating demand from small consumers produced a more uniform and predictable demand for electricity, he argued. "The characteristics of an individual are therefore of interest only in that they form a contribution to the characteristics of the group," he explained.

Consequently, "we can sell these small customers at a profit as a whole whereas any engineer who knew the facts could demonstrate to me that each one by himself is a loss to us."

The financial markets lavishly rewarded Insull's ideas. His utilities were the largest and best financed

Published by the IEEE Computer Society

of his age. He used these resources to build the control structure that we currently associate with electrical markets. This structure divided the electrical supply into three parts: generation, transmission, and distribution. He moved generation into large, centralized plants because they were more efficient. By contrast, his competitors were generally building regional or neighborhood plants.

To aggregate demand, Insull created large distribution grids that were controlled by regional substations. These stations kept statistics on electrical usage, projected demand, and monitored the actual consumption of power. The stations were connected to generators and to one another by dedicated phone lines. The stations could use these phones to call for additional power or report that the demand for electricity wasn't as great as anticipated.

The last element, the transmission lines, linked generators to grids. Initially, the utilities owned these lines as well as the grids and generators. However, by slow and steady steps, they became part of a market for bulk electricity. By the 1930s, this market included the generators in government hydroelectric projects. By the late 1970s, it had expanded to embrace a substantial collection of government, private, and public entities. However, this market still operated according to Samuel Insull's principle of aggregate demand.

A NEW PRINCIPLE FOR ELECTRIC MARKETS

As currently conceived, the smart grid will bring us the tools to start thinking about individual electrical consumption rather than aggregate consumption. Just as social networking has given people the ability to exchange ideas as individuals rather than as members of institutions, the smart grid allows entities to participate in the electrical markets as individuals. These markets will provide new information for controlling the grids, calling for power from transmission lines, and ultimately projecting the requirements for electrical generation.

The proponents of smart grids argue that this approach to power management will let them accommodate new technologies, such as electric automobiles. The "idle capacity of today's electric power grid," reports a study by the Department of Energy, "could supply 70 percent of the energy needs of today's cars and light trucks without adding to generation or transmission capacity—if the vehicles charged during off-peak times."

Many of the social and managerial problems come from the segmented structure of the nation's electrical infrastructure.

This aspect of the smart grid discussion parallels the ideas of Insull, who used statistical methodology to incorporate the demand from electrical railroads into his utilities. "If you consider it merely as a fraction of the supply of energy required by a community for all kinds of purposes," he observed, "the demand from railroads is found to be simply an incident."

TWO KINDS OF PROBLEMS

The task of building a smart grid requires us to solve two kinds of problems: technical and social. Of these two classes, the technical problems may be easiest to address. Much of the fundamental technology is based on well-tested ideas, such as those found in the Internet's hardware and software. Some have called for radically new technology to support the smart grid, including new means for transmitting and storing power. However, most of the research builds on existing technology.

By contrast, addressing the social and managerial problems may be much more difficult. Many of these problems come from the segmented structure of the nation's electrical infrastructure. Although this infrastructure involves several large government entities, such as the generation plants of the Bureau of Land Reclamation or the Tennessee Valley Authority, local agencies control most of the system. As they should, these local agencies consider only their own interests when they make policy. As a consequence, they tend to reject plans that require them to make sacrifices on behalf of the greater good. This is perhaps best illustrated by the problems of building new transmission facilities. "It is becoming increasingly difficult to site new conventional overhead transmission lines," explained a recent US government report, "particularly in urban and suburban areas experiencing the greatest load growth."

Yet, the construction of new facilities is only one aspect of the social problem. The plans for the smart grid propose an infrastructure that can operate in a unified manner. Such plans will require large investments and force the public debate to concentrate on the problem of controlling financial risk rather than on the value of radical innovation. No one, as Samuel Insull knew well, wants an expensive infrastructure, no matter how large a market it creates, if that infrastructure has a high risk of creating a financial loss.

THE RIGHT FOCUS

The task of focusing investment on the right aspects of the smart grid will be tricky, as the new electrical infrastructure might create markets that can't deliver all the benefits their supporters claim. Many reports note that the smart grid will be able to handle new suppliers as well as consumers; thus, any small investor could purchase a generator and add power to the network. However, that idea will work only if the market can set a price that will reward small investments. The current electrical infrastructure can't reward a family that wants to put solar cells on its roof or chooses to place a spinning turbine next to its barn. It's far from obvious that the smart grid will do a better job of providing incentives to such projects unless it substantially increases the price of electricity.

Any market, can, of course, be shaped by initial investment, technical and operational standards, subsidies, and early demonstrations. Currently, much of the technical work on the smart grid involves these issues. "The Smart Grid will ultimately require hundreds of standards, specifications, and requirements," explains a report by the National Institute of Standards and Technology. In addition to standards, NIST is creating a plan—a road map—to shape the industry and encourage investment. The Department of Energy is supporting research to develop the smart grid's basic technical elements. Agencies in Europe and Japan are addressing similar issues as they work to create their own version of the smart grid.

Among those hundreds of specifications and requirements, one item is repeatedly identified as central to the task of creating a smart grid: securing network information. If markets run on information and if bad information can destroy markets, then the markets that the smart grid forms will need to protect their data streams. "Cyber security must address not only deliberate attacks, such as from disgruntled employees, industrial espionage, and terrorists," notes the NIST report, "but also inadvertent compromises of the information infrastructure due to user errors, equipment failures, and natural disasters."

With almost 25 years of experience in dealing with the problems of data security, we've learned that cyber attacks can embarrass organizations, damage credit ratings, disrupt government operations (as such attacks did in Estonia), and destroy industrial machinery (as seems to have been the case with the recent Stuxnet malware). Records of electrical usage may not be as sensitive as our medical history or our credit information, but they still build an intimate record of our lives. The utilities know when we're sleeping, and they know when we're awake. They know when we watch television or open the refrigerator to cheat on our diet. Conceivably, they'll know about the unexpected trip that requires us to charge the electric automobile.

Of course, we have a long experience with trading information for better industrial products and more efficient services. Overall, the exchange seems to have been beneficial, producing more value than it has claimed. Still, we don't really know all the ramifications of living in a more active and individualized electricity market. It might have little impact upon our day-to-day lives or put us in a position of revealing to our neighbors more details than we'd like them to know.

Even though he managed his utilities with aggregate statistics, Samuel Insull would certainly have wanted to know more about his customers. "At his fingertips," noted one biographer, "he always had an impressive array of statistical data." He argued that the value a community received from a stable source of electrical power was well worth the cost of the information. As Insull explained in 1912, "There is no greater problem in the industrial world today than the proper method of producing energy and distributing it in a given area."

David Alan Grier is an associate professor of international science and technology policy at George Washington University and a senior member of IEEE. Contact him at grier@computer.org.

32 & 16 YEARS AGO

JUNE 1979

MEMBERSHIP (p. 4) "The Membership Committee began its work on December 11, 1951 with the formation of the Professional Group on Electronic Computers of the Institute of Radio Engineers. This group, with two chapters and an initial membership of 1000, was the beginning of the Computer Society. … Today we are approaching 40,000 members and 86 chapters. We hope for 45,000 members and 100 chapters by 1980. …"

CIRCUIT SWITCHING (p. 9) "Currently there is a considerable interest in the development of digital communication techniques. This interest is due to the development of distributed processing concepts, to changing regulatory policies of the US Government …, and to the development of large networks. One widely used communication technique that has tremendous growth potential is circuit switching. There are two primary reasons for computer-oriented people to be interested in circuit switching. First, the computer plays a major role in the control and implementation of a circuit switch. Second, circuit switches have been used as the interconnection structure for systems containing (or hypothesized to contain) a multiplicity of processing elements."

NETWORKS (p. 10) "In circuit switching, the term 'network' is used in two different senses. One is that of the network of switches (in the broader sense) tied together with relatively long-distance transmission facilities; a public telephone network is a network of switches, where there is a circuit switch or 'office' at each node. The other sense is that of a network of switching devices employing metallic contacts or semiconductor logic gates. …"

PRIVATE BRANCH EXCHANGES (p. 24) "The designers and manufacturers of modern PBX equipment have three primary objectives for their systems. First, the life-cycle cost of the system should be as low as possible. …

"Second, the design of the PBX should allow for future enhancements in features and switching capability. … Third, because the PBX is central to the entire customer communications system, it must be extremely reliable and provide for redundant operation in the event of critical component failures."

CROSSPOINT PATTERNS (p. 32) "Circuit switching networks are systems which provide a set of interconnecting circuits from a set of inputs to a set of outputs by opening and closing switches, or crosspoints. As a discipline, circuit switching networks deal fundamentally with the design and analysis of crosspoint patterns. At first thought, the idea of designing a system for simply interconnecting terminals might seem too basic to constitute a research area. However, the vast majority of the area of circuit switching networks lies far beneath the surface."

COMMUNICATION DESIGN (p. 57) "To solve the problem of providing fast, efficient communications at a reasonable cost, many different networks have been proposed in the literature, a number of which we discuss here. However, no single network is generally considered 'best,' since the cost-effectiveness of a particular design varies with such factors as the computational tasks for which it will be used, the desired speed of interprocessor data transfers, the actual hardware implementation of the network, the number of processors in the system, and the cost constraints on the construction."

THE INTEL 8089 (pp. 67-68) "Microprocessors of today (such as the Intel 8086) have attained respectable performance levels by innovative architectural and technological advances. However, such advances in microprocessor performance may be seriously overshadowed by the constraints of traditional non-intelligent I/O subsystems. The Intel 8089 I/O processor is designed to solve such problems by providing the necessary intelligence and capability to microcomputer I/O subsystems. …"

ARITHMETIC (p. 81) "The FPS-100 is a mini-sized programmable array processor manufactured by Floating Point Systems, Inc. When attached to a small mini-computer host, it provides the OEM with floating-point computational performance previously available only on expensive super-computing systems, the manufacturer states. It performs up to 8 million floating-point operations per second and features 38-bit precision and a synchronous architecture."

SOLAR ENERGY (p. 88) "The Solar Energy Information Data Bank, a system being developed at SERI [The US Department of Energy's Solar Energy Research Institute], will contain information on US government, industry, scientific, and public policies and activities as they relate to renewable resources and technologies. While the data bank's primary role will be to serve domestic users, information of use to developing countries will be contained in many of its data files. The International Programs Branch of SERI's Academic and International Programs Division will assemble and store basic energy profiles for a number of developing countries as part of the international thrust in SERI's data base systems."

32 & 16 YEARS AGO

JUNE 1995

EDITORIAL (p. 7) "The size of the VCR tape rental market far exceeds the box office theater market, and just about everything else. This 'content' will quickly find its way onto long-playing CD-ROMs for your home computer. It's one step away from being really cool with your friends; everyone who is anyone in 2002 will own a multifunctional PC. That, my friends, will influence the design of microprocessors more than all the computer scientists in the firmament."

SOFTWARE PIRACY (p. 11) "Piracy rampaged the software publishing and distribution industries last year. Losses due to the illegal copying of software hit a worldwide all-time high of $15.2 billion, reported the US Business Software Alliance in its 1994 annual survey of 77 countries, 'The Impact of Software Piracy on the International Marketplace,' published in April."

INTERRUPT PROCESSING (p. 36) "As the complexity of modern processors grows, effective strategies for handling interrupts become increasingly important. To meet the demands of more applications, peripherals, and functions, such strategies must feature different types and levels of interrupts that have been selectively combined. While many strategies are possible, it's a challenge for the designer to determine which interrupt-processing techniques and methods are best for optimal system performance."

VIRTUAL ENVIRONMENTS (p. 57) "Different and more interesting work, learning, and play—these are the key attractions behind virtual-environment technologies, as researchers explore new ways for people to interact with and use computers. The virtual-environment field spans a broad range of technologies, many of which are still in early phases of development. Extensive research is ongoing worldwide to develop these technologies and to discover new methods to integrate them into useful applications."

GOVERNMENT INVESTMENT (p. 79) "The success of the ARPA VLSI program lends support to the [National Research Council] committee's view that long-term government investment is necessary 'to sustain the innovation and growth needed for enhancing the information infrastructure.' The committee notes that few companies are willing, or able, to invest for a payoff 10 to 15 years down the road, especially if a complex technology requires several engineering iterations and thereby risks disclosure of key insights. Further, pointing to IBM's failure to exploit John Cocke's early work on RISC technology until after government-sponsored university research programs embraced it, the committee suggests that 'it is often easier for a start-up to form, raise venture capital, and succeed than for an established firm to abandon a currently successful direction in favor of a new approach just when the old approach is most financially successful.'"

TECHNOLOGY EVALUATION (p. 86) "Unfortunately, the software industry lacks standard measurements and benchmarks for evaluating the effectiveness of programming tools and languages, design approaches, or almost any other kind of technology. Purchasing and acquisition decisions are often made on the basis of unsubstantiated vendor claims. Moreover, once a new tool or methodology is acquired, deployment is often slow. Tools are acquired without considering training needs, or if training is considered, it's not readily available due to schedule pressures."

MICE (p. 92) "When photographs of computer mice first surfaced in publications, everyone began to ask, 'What are those things? Are they for real?' Now, after many years of faithful performance, mice are standard on most machines. In fact, they have multiplied and flourished to the point of taking on many new forms and characteristics. …"

PERSONAL ORGANIZERS (p. 97) "Texas Instruments introduced the PS-6800 and PS-6700 personal organizers and an optional PC connectivity kit. Both organizers store reminders, addresses, and notes, and they include a calculator, clock with alarm, and a 100-year monthly calendar. The organizer has a high-contrast liquid crystal display that shows six lines of information with 24 characters per line. The organizer's scan feature provides quick access to a telephone directory and a daily planning calendar. Users can set up to four daily reminder alarms, and task reminders can be set to carry over to the next day. …"

COMPUTER-BASED SYSTEMS (p. 100) "The integration of information technologies dramatically increases the interactions among physical components and processes, generates complex dynamics, and creates component interdependencies unknown in earlier systems. Although development processes for the individual components exist, they remain largely fragmented. Thus, we need a system-level discipline and interdisciplinary engineering specialty—engineering of computer-based systems (ECBS)—that addresses the qualities and challenges of CBSs."

PDFs of the articles and departments from Computer's June 1979 and 1995 issues are available through the IEEE Computer Society's website: www.computer.org/computer.

Editor: Neville Holmes; neville.holmes@utas.edu.au

Mobile Security: Finally a Serious Problem?

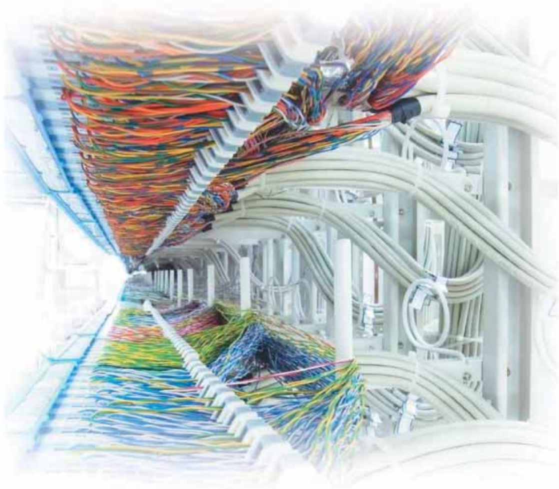

Neal Leavitt

The growing popularity of wireless technology may have finally attracted enough hackers to make the potential for serious security threats a reality.

In the world of computers and communications, the more widely a technology is used, the more likely it is to become the target of hackers.

Such is the case with mobile technology, particularly smartphones, which have exploded in popularity in recent years. According to market analysis firm ABI Research, 370 million smartphones were in use globally last year.

Many users download mobile applications with little regard to whether they're secure, providing a ready way for hackers to attack the devices.

In addition, said Gustavo de Los Reyes, executive director for AT&T Security R&D, "These phones are being used frequently for sensitive transactions like banking, mobile payments, and transmitting confidential business data, making them attractive targets if not protected."

"The payoffs—financial and personal information—could be huge," noted Purdue University computer science professor Richard P. Mislan.

Smartphones generally connect to the Internet, as well to PCs for software updates or media synchronization, providing convenient attack vectors.

Device makers and wireless-service providers have long focused on communications and other services, with security remaining an afterthought.

Referring to the two most popular smartphone platforms, Ed Moyle, senior analyst with market research firm Security Curve, said, "Security is now playing catch-up with the rapid adoption of Android and iPhone, both of which are hard for enterprises to manage."

Thus, after years of warnings about mobile security, there finally appears to be a reason to worry.

In fact, the number and types of mobile threats—including viruses, spyware, malicious downloadable applications, phishing, and spam—have spiked in recent months.

For instance, McAfee Labs' threat report for 2010's fourth quarter reported a 46 percent increase in malware targeting mobile phones over the same time period the previous year.

"We're seeing more than 55,000 new pieces of [mobile] malware on a daily basis," said Dave Marcus, McAfee Labs' director of security research and communications.

THREATS ON THE MOVE

Mobile devices increasingly face various types of threats, as Figure 1 shows.

Botnets

Attackers form a botnet by infecting multiple machines with malware that victims generally acquire via e-mail attachments or from compromised applications or websites. The malware gives hackers remote control of the "zombie" devices, which can then be instructed to perform harmful acts in concert.

"These command channels could also provide a way to update the malicious code so that it will communicate or act differently," said Juniper Networks research engineer Troy Vennon.

The easiest way for an attacker to benefit from a mobile zombie network is to send SMS or multimedia message service (MMS) communications to a premium phone account that charges victims fees per message, explained Vennon.

The scammers act as the premium-account owner's affiliates, receiving some of the money that their attacks generate, noted Bradley Antsis, vice

0018-9162/11/$26.00 © 2011 IEEE Published by the IEEE Computer Society

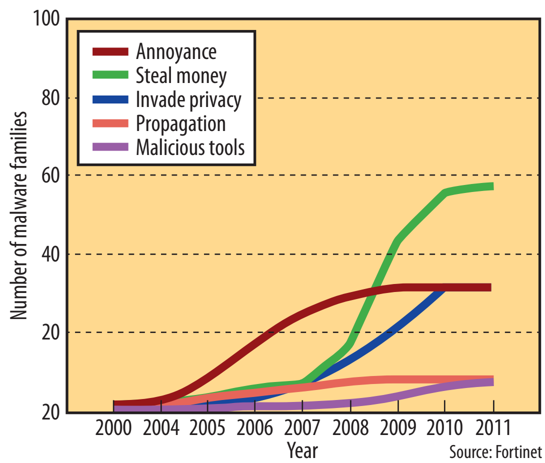

Figure 1. The number of threats to mobile devices, particularly those designed to steal money, has increased steadily during the past few years.

president of technical strategy for security vendor M86.

The Yxe malware family that hit China last year caused this problem.

Also in 2010, malware originating in Holland exploited a vulnerability in jailbroken smartphones—those that owners have modified to gain OS root access and remove manufacturers' usage limitations—to create a botnet. The network sent SMS messages to premium numbers.

Last year, another mobile botnet targeted European customers of a Dutch online bank. The malware used in the attack included command logic that gave the hacker remote control of victims' smartphones.

With PCs, hackers often use zombies within botnets to launch denial-of-service attacks. Thus far, though, there have been no major mobile DoS incidents.

Malicious applications

In some cases, hackers have uploaded malicious programs or games to third-party smartphone-application marketplaces—such as those for Apple's iPhone and Google's Android devices—or have otherwise made them available on the Internet.

"These malicious apps are usually free and get on a phone because users voluntarily install them," said Pierluigi Stella, chief technology officer for Internet security vendor Network Box USA.

Once on a handset, the programs steal personal information such as account passwords and logins and send it back to the hacker. They also open backdoor communication channels, install additional applications, and cause other problems.

Most mobile application marketplaces don't require that code in applications be cryptographically signed by the developer before it can be uploaded, noted Kurt Stammberger, vice president of market development for security vendor Mocana.

"Until this becomes common," he said, "malicious apps will proliferate quickly on mobile platforms."

Social networking

As smartphone use has grown, so has mobile social networking.

Malicious links on social networks can effectively spread malware. Participants tend to trust such networks and are thus willing to click on links that are on "friends'" social networking sites, even though—unknown to the victim—a hacker may have placed them there, said M86's Antsis.

Clicking on a link could download a malicious application on a victim's computer, said Network Box USA's Stella. This could let a hacker place Trojans, spyware, and backdoors on the machine and even conduct identity or information theft, he added.

Some schemes use a sensational headline or promise information on a current hot topic to grab readers' attention and encourage them to click on a malicious link.

Spyware

Hackers can use spyware available online to hijack a phone, allowing them to hear calls, see text messages and e-mails, and even track a user's location through GPS updates.

Most commercial mobile spyware applications send an update of captured communications or location data to a website where the spy logs in to view the data, noted Juniper Networks' Vennon. In some cases, SMS communications inform the spy that the system has obtained new data.

The software can even create a hidden access point inside a mobile phone that lets a hacker turn on the device without it ringing, in essence converting it into a microphone, said Purdue University's Mislan. The spy could then hear nearby conversations.

While some malware writers sell or give away mobile spyware, there are also online vendors—such as ClubMZ, FlexiSPY, and Retina-X Studios—that sell the software commercially.

These companies say their products are only for legal uses and can be helpful in finding a stolen mobile device or in monitoring the activities of children, as well as employees using company phones.

Mobile phone spyware is illegal in the US but is sold by websites hosted elsewhere, noted Simon Heron, principal with Network Box's UK office.

Bluetooth

Bluetooth enables direct communication, including the sharing of content, between mobile devices.

Wireless devices can broadcast their presence and allow unsolicited connections and even the transmission of executables if users don't configure their Bluetooth operations appropriately.

On rare occasions, mobile malware—such as the Cabir worm—has used Bluetooth to propagate.

Wi-Fi

Hackers can intercept communications between smartphones and Wi-Fi hotspots.

The fundamental vulnerability is hotspot architecture with no encryption to protect transmitted data.

"If a user connects to [such] a hotspot for the first time, the end-to-end connection between the user's device and the hotspot provider is not secured, so the [hacker] can intercept and control the user's traffic," said Carnegie Mellon University computer science professor Patrick Tague. In this scenario, the hacker gets between the user and the hotspot provider and hijacks the session via a man-in-the-middle attack.

A hacker can also set up a peer-to-peer network that mimics a Wi-Fi hotspot offering a high-quality connection, which entices users to connect. The hacker then intercepts victims' transmissions without their knowledge.

Phishing

Phishing poses the same risk on smartphones as it does on desktop platforms.

In fact, many users trust their mobile device more than their computers and thus are more vulnerable to phishing.

Additionally, said Juniper Networks' Vennon, the lack of maturity in phishing filters and reputation-based services in mobile browsers, combined with the immediacy and portability of telephone communications, makes the platform attractive for phishers.

Mobile phishing is particularly tempting because wireless communications enable phishing not only via e-mail, as is the case with PCs, but also via SMS and MMS, noted AT&T's de Los Reyes.

Social media phishing is becoming a major issue as social networking sites contain an increasing amount of personal information that phishers can use to make their attacks more effective, said Paul Henry, security and forensics analyst for market research firm Lumension Security.

The number and types of mobile threats—including viruses, spyware, malicious downloadable applications, phishing, and spam—have increased in recent years.

TRADITIONAL SECURITY APPROACHES

Mobile communications can use the same types of security—including antivirus and firewall products—as fixed communications. Vendors include Fortinet, F-Secure, Juniper Networks, Kaspersky Lab, Lookout Inc., Mocana, NetQin, Trend Micro, and Trusteer.

Most of these products work much like their PC counterparts. For example, mobile antivirus products scan files and compare them against a database of known mobile malware code signatures. Noted Mocana's Stammberger, this approach is compute-intensive and "eats batteries for lunch."

Mobile security software is also more likely to use the cloud to offload some of the processing typically associated with PC-based products, said Chris Perret, CEO of security vendor Nukona.

There are only a few mobile encryption software products, including SecurStar's Phonecrypt and Credant Technologies' Mobile Guardian for Handhelds. They're scarce primarily because they're challenging and expensive to develop.

"Restricted [OS] kernel access means you can't put the cryptographic processes sufficiently low down in the stack, close to the silicon. Processor limitations, memory constraints, and battery-life issues make some of these apps as slow as molasses," explained Stammberger.

OTHER MEASURES

Security vendor MobileIron recently launched a storefront so that businesses can deliver mobile applications directly to employees without posting them publicly. This lets businesses enforce security policies about which users and devices can access specific corporate applications.

This summer, trials will start for the AT&T Smart Mobile Computing platform, which will include features such as mobile security, mobile device management, a virtual private gateway, encryption, policy controls, a virtual desktop, and cloud computing capabilities. Customers will also be able to apply their own security policies to this platform.

The AT&T Security Research Center recently opened in New York City. Employees have expertise in a broad range of areas, including security, cellular systems, networking, and data mining.

Under a distribution and marketing agreement signed last year, Verizon Wireless will promote Lookout Inc.'s mobile security products to customers.

Handset and chip makers are also addressing mobile threats. For example, Mocana's Stammberger said his company is working with Freescale Semiconductor, IBM, Intel, LG, Motorola, and Nokia to "better leverage and improve their on-chip crypto-acceleration hardware."

Mocana also sells Acceleration Harness, a technology for connecting OS-based security software with on-chip hardware acceleration.

As the mobile ecosystem evolves and hackers probe for vulnerabilities, devices will face a growing number of the types of attacks traditionally launched against desktop systems, said Trusteer chief technology officer Amit Klein.

"We need to implement mobile security solutions now to protect against these new threats," he added.

"The greater visibility of these attacks will place an increasing importance on mobile device makers having enterprise-grade security features and configuration options in place. It will become necessary for security to be considered in all phases of application development to ensure that resiliency against attacks is built into mobile devices from the start," said Adrian Stone, director of security response for BlackBerry vendor Research in Motion.

"Our dependence on an always-on, connected, mobile device environment is going to be profound in critical contexts that we can't imagine today," said Stammberger. "We have to be able to trust these devices, but we can't now. There's still a lot of work that needs to be done to get to the point where that trust is warranted."

Neal Leavitt is president of Leavitt Communications (www.leavcom.com), a Fallbrook, California-based international marketing communications company with affiliate offices in Brazil, China, France, India, and the UK. He writes frequently on technology topics and can be reached at neal@leavcom.com.

Editor: Lee Garber, *Computer*; l.garber@computer.org

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

IEEE computer society

PURPOSE: The IEEE Computer Society is the world's largest association of computing professionals and is the leading provider of technical information in the field.
MEMBERSHIP: Members receive the monthly magazine *Computer*, discounts, and opportunities to serve (all activities are led by volunteer members). Membership is open to all IEEE members, affiliate society members, and others interested in the computer field.
COMPUTER SOCIETY WEBSITE: www.computer.org
OMBUDSMAN: To check membership status or report a change of address, call the IEEE Member Services toll-free number, +1 800 678 4333 (US) or +1 732 981 0060 (international). Direct all other Computer Society-related questions—magazine delivery or unresolved complaints—to help@computer.org.
CHAPTERS: Regular and student chapters worldwide provide the opportunity to interact with colleagues, hear technical experts, and serve the local professional community.
AVAILABLE INFORMATION: To obtain more information on any of the following, contact Customer Service at +1 714 821 8380 or +1 800 272 6657:

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

PUBLICATIONS AND ACTIVITIES

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

IEEE

revised 5 May 2011

COMPUTER SOCIETY OFFICES

Washington, D.C.: 2001 L St., Ste. 700, Washington, D.C. 20036-4928
Phone: +1 202 371 0101 • **Fax:** +1 202 728 9614
Email: hq.ofc@computer.org
Los Alamitos: 10662 Los Vaqueros Circle, Los Alamitos, CA 90720-1314
Phone: +1 714 821 8380
Email: help@computer.org

MEMBERSHIP & PUBLICATION ORDERS

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

Now, manage both your UPS and your energy proactively

Only APC Smart-UPS saves money and energy without sacrificing availability

Today's more sophisticated server and networking technologies require higher availability. That means you need more sophisticated power protection to keep your business up and running at all times. But that's not all. In today's economy, your UPS must safeguard both your uptime *and* your bottom line. Only APC by Schneider Electric™ helps you meet both of these pressing needs. Specifically, the APC Smart-UPS™ family now boasts models with advanced management capabilities, including the ability to manage your energy in server rooms, retail stores, branch offices, network closets, and other distributed environments.

Intelligent UPS management software

PowerChute™ Business Edition, which comes standard with Smart-UPS 5 kVA and below, enables energy usage and energy cost reporting so you can save energy and money by tracking energy usage and costs over time; CO_2 emissions monitoring to reduce environmental impact through increased understanding; and risk assessment reporting so you can identify and proactively manage threats to availability (e.g., aging batteries).

Best-in-class UPS

Our intelligent, interactive, energy-saving APC Smart-UPS represents the combination of more than 25 years of Legendary Reliability™ with the latest in UPS technology including an easy-to-read, interactive, alphanumeric LCD display to keep you informed of important status, configuration, and diagnostic information, a unique battery life expectancy predictor, and energy-saving design features, like a patent-pending "green" mode. Now, more than ever, every cost matters and performance is critical. That's why you should insist on the more intelligent, more intuitive APC Smart-UPS.

Why Smart-UPS is a smarter solution

Intuitive alphanumeric display

Get detailed UPS and power quality information at a glance – including status, about, and diagnostic log menus in up to five languages.

Configurable interface

Set up and control key UPS parameters and functions using the intuitive navigation keys. On rack/tower convertible models, the display rotates 90 degrees for easy viewing.

Energy savings

A patent-pending "green" mode achieves online efficiencies greater than 97 percent, reducing heat loss and utility costs.

Download White Paper #24, "Effect of UPS on System Availability," and register to WIN APC Smart-UPS 1500VA rack/tower LCD 120V, a $779 value!

Visit www.apc.com/promo Key Code b614v Call 888-289-APCC x6272

APC™ by Schneider Electric

Fast New Version of I/O Interface Introduced

Just when it looked like USB had won the I/O interface war, a trade organization has announced plans for a new, faster version of a competing technology that provides ways to connect peripherals to systems.

The Serial ATA International Organization recently said it will complete work on its SATA Universal Storage Module approach, geared toward consumer-storage applications, later this year. SATA-IO says a single USM unit could simultaneously and quickly provide content to multiple devices, such as a TV, laptop, and PC.

The organization says the technology will offer data rates up to 6 gigabits per second, with support for earlier SATA versions that offer speeds of 1.5 or 3 Gbps.

USB has become the dominant I/O interface and is used in many settings. The latest version, USB 3.0—also known as SuperSpeed USB—began appearing in consumer products in late 2009. It provides data rates up to 5 gigabits per second, has been widely adopted, and is backward-compatible with earlier USB versions.

Man-in-the-Browser Attacks Afflict Corporate Networks

Cybercriminals are beginning to use a new approach to launch attacks against increasingly well-defended corporate networks. With these networks protected by firewalls, antivirus software, and other forms of security, hackers are increasingly turning to what they now see as the most vulnerable point: the browser.

They are thus burrowing into the networks via man-in-the-browser attacks, which start with a website containing malicious software. When visitors arrive, the malware is installed on the browser, giving hackers control of it.

Attackers could steal login credentials, account numbers, financial information, or other personal material. They could also modify pages, content, or transaction data presented to the user. During a financial transaction, the hackers could use the malware to surreptitiously send transfer or payment requests. The customer wouldn't realize there is a problem until the next account statement arrives.

China Develops Technology for Ultrafast Supercomputers

China recently provided information on a microprocessor that its Institute of Computing Technology has designed and that is slated to power a multipetaflops supercomputer in the not-too-distant future.

China's Dawning Information Industry will use about 3,000 Godson-3B multicore multiprocessors in a high-performance 300-teraflops system it plans to release this summer. The company plans a future version that will deliver at least 1 petaflops.

The eight-core, 1.05-GHz Godson-3B chip, built with 65-nm feature sizes, performs 128 gigaflops and is already being manufactured.

Planned for release in two years, the 2-GHz Godson-3C will have 16 cores and 28-nm feature sizes, and will offer 512 Gflops.

The Chinese have become important players in the supercomputer industry. On the most recent list of the top 500 high-performance computers—published in November 2010 by several researchers in the US and Germany—China has the fastest and third-fastest systems.

The most powerful supercomputer is the Tianhe-1A, developed by the Chinese National University of Defense Technology and located at the National Supercomputing Center in Tianjin. The third-fastest is the Nebulae, developed by Dawning and located at the National Supercomputing Center in Shenzhen.

Amazon's Massive Cloud Hosting Site Crashes

Amazon Web Services—a huge, cloud-based Web-hosting system—crashed recently, causing problems for the many large online operations it serves.

Industry observers say the temporary crash raises questions about the dependability of AWS and perhaps even the cloud itself. Other large, cloud-based Web services have also experienced temporary failures.

Many organizations—including FourSquare, Hootsuite, the New York Times, Quora, and Reddit—pay to run at least some of their websites on Amazon's service, which works with the company's Elastic Compute Cloud.

EC2 is a distributed system of servers—located in five global regions, with additional subsystems in each region—designed to provide flexible, scalable, and redundant Web services.

Amazon has touted its system's reliability. However, the recent crash pushed some hosted websites offline for many hours. Amazon says the crash started in a subsystem within its US East Region that became unable to service read and write operations.

According to the company, the following occurred:

- Amazon was changing system configurations to upgrade the capacity of a primary network in the subsystem. Instead of shifting its traffic to a redundant router, the company mistakenly moved the communications to a

Published by the IEEE Computer Society

RESEARCHERS SPIN SILK INTO MATERIAL FOR ELECTRONICS

Taiwanese academic researchers have discovered a way to use silk to improve performance in electronic devices. Silk is plentiful and significantly less costly than amorphous silicon and other materials typically used for manufacturing electronic components, such as those used in displays.

The researchers developed a way to turn liquid silk into a membranous material called *silk fibroin* that can be used to make flexible organic thin-film transistors (OTFTs).

Silk fibroin offers optical transparency, electrical insulation, and flexibility, noted professor Jenn-Chang Hwang of National Tsing Hua University. The soft, low-cost silk fibroin could be used in flexible e-book readers, LED displays, and RFID tags, said Hwang.

Some manufacturers are experimenting with using OTFTs to make better flexible, active-matrix, flat-panel displays, said Jennifer Colegrove, vice president of emerging display technologies for DisplaySearch, a market research firm.

Hwang said his team's process marries silk protein, used as a gate dielectric for insulation, to pentacene, a semiconducting material.

The researchers extract the protein by boiling silkworm cocoons in a liquid and then processing the result, yielding silk fibroin. They employ spin coating to apply the fibroin to a polyethylene terephthalate substrate and use gold as the gate electrode. When dry, they deposit pentacene onto the silk fibroin.

The resulting OTFT has very high electron mobility, which could improve a transistor's performance at least 20 times over that of typical amorphous-silicon transistors.

Most of today's OTFTs have lower electron mobility, which inhibits performance, noted Colegrove. Hwang said his experiments suggest that his team's OTFT's pentacene layer is better arranged molecularly to enable higher electron mobility when on top of silk fibroin.

Colegrove stated that the National Tsing Hua research looks promising and that substituting silk protein for conventional silicon dioxide enables the pentacene to perform better. However she expressed uncertainty about whether silk fibroin is less expensive than silicon dioxide when the production process is taken into account and said she would like to see more published work demonstrating the material's reliability and stability.

Hwang noted that the researchers' goal is to have silk-related electronics products on the market soon and that they are already speaking with some firms—which he declined to identify—about using the material.

Silk has long interested academic researchers for its strength and flexibility, as well as its potential uses in new settings, said University of Alabama at Birmingham assistant professor of chemistry Eugenia Kharlampieva.

Taiwanese scientists have developed a method for using silk to make organic thin-film transistors that could improve performance in electronic devices. Silk is plentiful and significantly less costly than other materials typically used for manufacturing electronic components.

lower-capacity network, which could not handle the volume.

- This left the primary and secondary networks unable to function, causing a loop of service requests that could not be satisfied, thereby overloading the system.
- Initially, additional attempts to resolve the issues caused them to spread, thereby affecting the redundancy designed to avoid system unavailability.

After fixing the problems, Amazon said it will audit its configuration process, change its procedures as necessary, and increase automation to prevent a recurrence of the issues that caused the crash.

Could a Proposed Wireless Network Harm GPS Services?

A proposed fast wireless network could overwhelm GPS signals in the US, interfering with important services such as police and air-traffic communications, as well as navigation devices, according to some industry observers.

The issue involves permission the US Federal Communications Commission (FCC) gave LightSquared to construct and implement a nationwide broadband network using frequencies close to those utilized by GPS.

GPS vendors say that strong signals from LightSquared's system could jam their transmissions. LightSquared and the FCC contend the proposed network would not cause these problems.

US officials say they won't let LightSquared switch on its network as scheduled later this year unless they feel confident that GPS systems will still function properly.

Earlier this year, the FCC approved LightSquared's network proposal, saying it would increase competition among wireless providers and thereby make US mobile services faster and less expensive.

The company plans to operate

a hybrid network that would work with both Long Term Evolution (LTE) wireless and satellite technologies. Its services—which would be provided to the public by wireless carriers, not directly by LightSquared—would compete with those offered by companies such as AT&T and Verizon Wireless.

Various sources say interference from LightSquared's service could threaten, for example, automobile-navigation systems, public-safety communications networks, and a proposed US Federal Aviation Administration plan to use GPS to improve the nation's air-traffic-control operations. These sources include car-navigation equipment manufacturers, aviation-related organizations, police officials, and the US military.

In the past, GPS receivers have been able to filter out low-power signals in nearby frequencies, most of which have come from satellites. However, vendors say, a large ground-based system would cause interference issues.

GPS units could add filters for LightSquared signals, but some sources say that this approach might be very expensive.

LightSquared and the FCC say more testing will occur to determine whether the new system would cause interference problems. The company is slated to participate in the testing.

Study Shows Significant Downtime for Critical APIs

A study has shown that some of the public APIs that major online operations use to access various important third-party services experience high downtime levels, which could cause problems for these operations.

Website- and application-performance monitoring company WatchMouse checked the availability of 50 major cloud-based APIs every five minutes from 16 February to 17 March 2011. The testing included a simple API call and a check for a valid result. The lack of a proper or prompt response contributed to downtime ratings.

According to WatchMouse's website, "In accordance with industry standards, availability of [at least] 99.9 percent is regarded as good, while anything below 99 percent is regarded as poor. (99 percent uptime equals over 80 hours downtime per year, or about one business day per month.)"

Ten of the interfaces tested—five from Google and the Basecamp, Delicious update, eBay shopping, Quora, and SimpleGeo APIs—had 100 percent uptime.

The lowest score was 94.32 percent uptime, by the MySpace Open Search API. Others scoring below 99 percent were the Digg (98.66 percent), Gowalla (98.52 percent), GeoNames (97.47 percent), Eventful (97.20 percent), and Posterous (97.17 percent) APIs.

WatchMouse's own API showed 99.98 percent uptime.

New Organization to Develop Networking Standards

A group of technology companies has formed an organization to develop and manage standards based on a new approach called software-defined networking.

The Open Networking Foundation (ONF) says it will standardize SDN technologies—pioneered at Stanford University and the University of California, Berkeley—designed to make large and small networks programmable.

Proponents say SDN could yield flexible, secure networks that have fewer traffic issues and that could be less costly to construct and run.

The group says the ONF's activities could provide network owners with a standardized way to have flexible control over their operations.

ONF members include Broadcom, Cisco Systems, Dell, Deutsche Telekom, Ericsson, Facebook, Google, Hewlett-Packard, IBM, Juniper Networks, Microsoft, Verizon, VMware, and Yahoo.

Typically, most of the Internet's intelligence resides in the computers functioning as end points, not in routers or elsewhere within the network.

In cloud computing, though, information and applications are on network-based computers. For this system to function properly, SND proponents say, smarter networks would be necessary to, for example, orchestrate the behavior of the thousands of routers that would be involved.

This approach would enable, for instance, the configuration of networks to handle heavy traffic or the prioritization of certain types of data to achieve quality of service for time-sensitive traffic.

Proponents say SDN would permit these types of activities to happen by making previously proprietary hardware and software systems that control packet flow more open.

Several companies—including Cisco, Hewlett-Packard, and Juniper Networks—have produced prototype SDN systems.

Some industry observers say SDN appears to have wide industry support but still must prove itself in actual use and in the marketplace. C

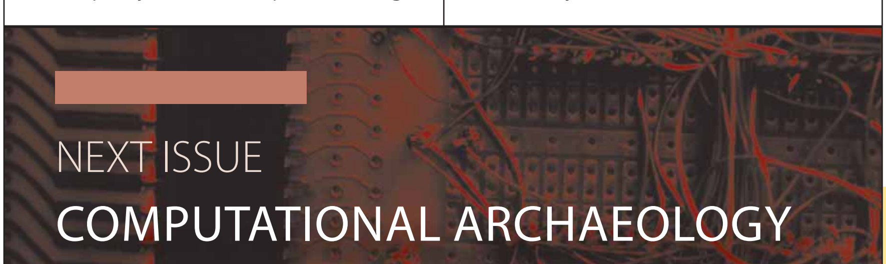

Editor: Lee Garber, *Computer*; l.garber@computer.org

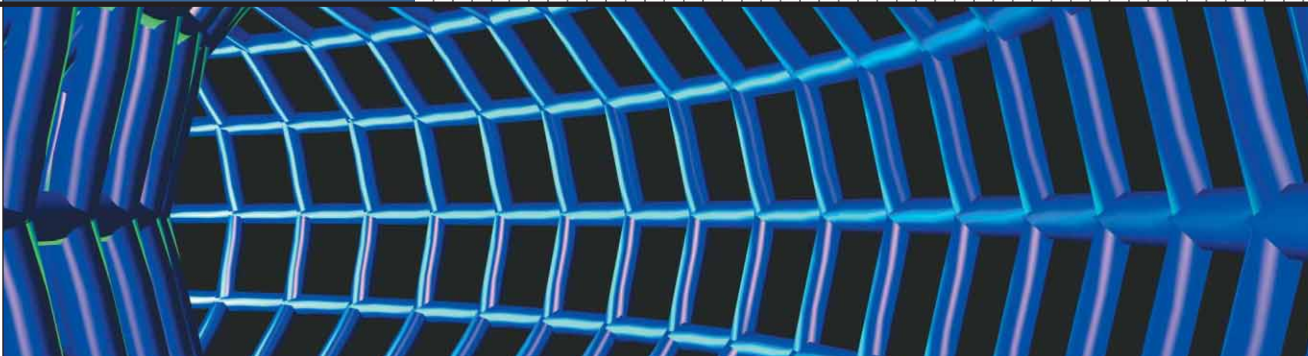

Understanding and Using Rendezvous to Enhance Mobile Crowdsourcing Applications

R.J. Honicky, *Nokia Research Center, Berkeley*

Researchers are exploring the rendezvous concept—bringing sensors close to one another in space or time—as a way to make sense of disparate data collected by individuals. Applications such as participatory atmospheric sensing illustrate the potential of rendezvous to help create powerful mobile applications.

As mobile phones increasingly become general-purpose devices, phone manufacturers have integrated various sensors into them. For example, even low-end phones now include cameras, and mid-range phones often include GPS capability for use with popular mapping applications. Accelerometers and magnetometers provide a level of immersion and integration between the virtual and physical worlds that enable next-generation augmented-reality applications, such as Wikitude, Star Walk, Dish Pointer, Work Snug, and Car Finder. The "Categories of Mobile Sensing Applications" sidebar defines a proposed classification for mobile sensing applications.

At the same time, crowdsourced information—such as Wikipedia and online reviews—now dominates the Internet. Distributed processing projects such as SETI@home and Folding@home illustrate the power of distributed processing and illuminate people's willingness to contribute their own idle resources toward social and scientific advancement. Applications such as cell tower location databases, traffic flow monitoring, and pothole detection have the added characteristic of automatically integrating sensor data from mobile devices.

Bringing these two trends together with *rendezvous*—when two or more sensors come close to one another in space or time—offers the potential for powerful new applications that exploit these characteristics.

MODELING RENDEZVOUS

People rendezvous quite frequently, but only in the past few years has the ubiquity of mobile phones, particularly those with GPS technology, enabled applying the concept to the large-scale study of human mobility. Although GPS receivers have long been available, coupling them with mobile phones has drastically simplified the data collection and storage process, putting studies with hundreds of users over years within reach of researchers with modest budgets. These studies continue to shed light on the nature

Published by the IEEE Computer Society

of human mobility, and most importantly, rendezvous. For example, mobile-phone calling records have recently revealed Lévy walk mobility patterns in human mobility,[1] making it possible to better model and understand mobile systems.

A study by Kyunghan Lee et al.[2] described an algorithm based on the principle of least action. According to this principle, people plan their trips so that they visit nearby destinations at around the same time. These investigators describe how this algorithm not only corresponds well to real human mobility patterns, but also has important consequences for the frequency and distribution of rendezvous between users. This work provides an important tool for evaluating rendezvous-based algorithms.

Much delay-tolerant networking (DTN) literature builds on assumptions about rendezvous, which provides opportunities for mobile devices to exchange data when they are within communication range. The accuracy of mobility models[2] and social relationship models[3] has a significant impact on routing-algorithm design and performance. The DTN literature thus offers significant analysis and insights into human mobility and rendezvous.

The most mature and well-studied application of rendezvous in the DTN literature is opportunistic, peer-to-peer routing (sneaker nets), in which data passes between mobile devices with the goal of eventually reaching a destination.[4] Each rendezvous represents an opportunity for mobile devices to exchange data. These data exchanges might take place either among unrelated devices or within a federated set. Rendezvous among a user's laptop, phone, and set-top box, for example, might provide opportunities for data transfer over low-power, high-bandwidth connections.

DTN-based routing, however, is only one of several potential applications. In most cases, rendezvous is also essential to data fusion. In the physical world, proximity in space and time often corresponds to a high correlation and systematic relationship among sensor readings. Understanding the relationships among sensor readings from nearby sources can enable us to combine the readings, increasing the quality and quantity of data available in the system.

For example, consider atmospheric sensing for pollution monitoring. If two sensors are close to each other, they should read almost the same values. If they do not, at least one is miscalibrated. The closer the sensors are, the more closely they should agree. Thus, rendezvous between sensors can help determine the relative miscalibration among sensors. Knowing the true calibration of some sensors in the system, or the statistical distribution of miscalibrations, makes it possible to determine the sensors' absolute calibration. Well-calibrated sensors can significantly increase the system's accuracy. Both the CaliBree system from the MetroSense project at Dartmouth,[5] and research efforts at Berkeley, Intel, and Nokia demonstrate this capability.

CATEGORIES OF MOBILE-SENSING APPLICATIONS

A useful way to look at mobile-sensing applications classifies them in terms of four levels of integration of the data from different users.

- ***Personal sensing*** **applications typically serve only the needs of one user and do not require interaction with other devices. An application that records notes, lectures, or concerts for an individual's later use is one simple example.**
- ***Shared sensing*** **takes such applications one step further by permitting easy data sharing. For example, a shared-sensing application could record a concert and then automatically post, stream, or send it to someone else via messaging.**
- ***Contextualized sensing*** **applications automatically augment sensor data using other information. In the concert example, that data might be the recording device's location and orientation.**
- ***Sensor fusion*** **applications use contextual data to combine measurements from different, possibly heterogeneous, sensors together. Such applications could, for example, improve an audio recording's quality by automatically combining it with other recordings, based on location data.**

Today, many mobile-sensing applications fall into the personal, shared, or contextualized categories. When people converge in space or time, however, opportunities arise to truly make the data's value greater than the sum of its parts. Developers and researchers can use this convergence—also referred to as rendezvous—to enhance mobile applications and experiments, increasing the value of collected information.

In sensing systems, precision roughly corresponds to the number of bits of (meaningful) information a sensor provides; accuracy corresponds with the correctness of those bits. Calibrated sensors let systems take advantage of rendezvous to increase precision by averaging readings from sensors located close to one another.[6,7]

When a system wants to sense a quantity for which a phone does not have an appropriate sensor, it can "borrow" sensors from nearby phones. Shane Eisenman et al. discussed secure mechanisms for letting users borrow sensors from appropriate nearby devices.[8]

These examples describe powerful ways in which researchers have already exploited rendezvous to intelligently combine information from different sensors. These are first steps in understanding how to take advantage of rendezvous to build more situated, immersed, and representative computer systems.

Readers interested in a more in-depth mathematical understanding of mobility should examine a study by Marta Gonzalez et al. in which they used mobile-phone records to analyze human mobility,[1] the SLAW algorithm described by Kyunghan Lee et al.,[2] and a social-based forwarding algorithm investigation conducted by Pan Hui et al.[3] in which they used several different rendezvous datasets.

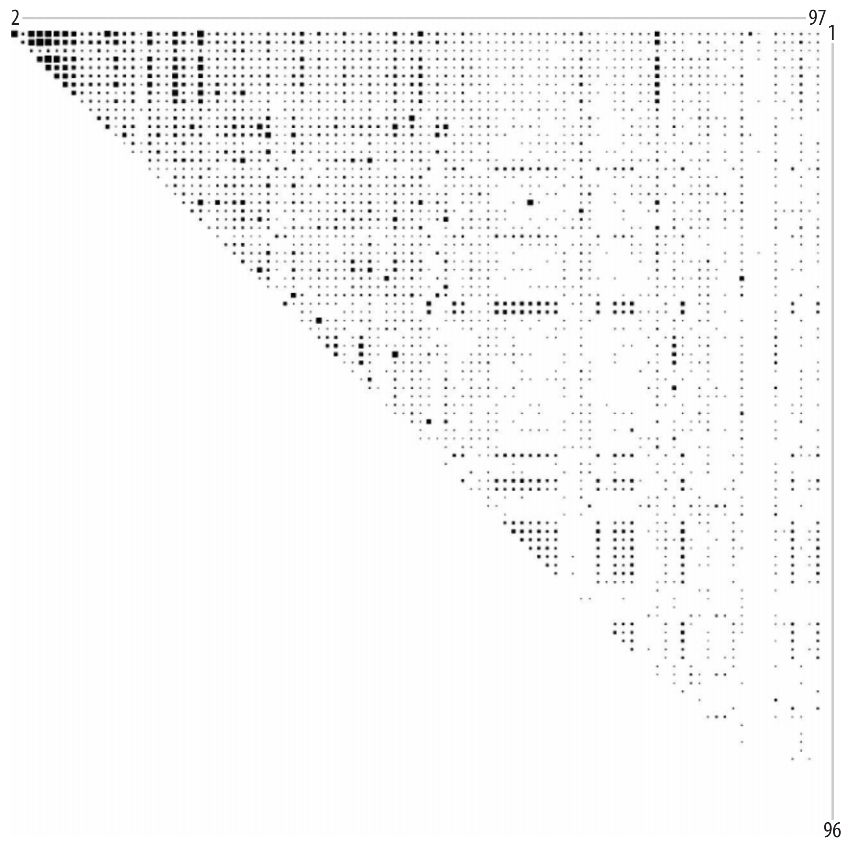

Figure 1. Each square represents the number of times a user rendezvoused with another user in the study, ordered by the total number of rendezvous for a given user. Surprisingly, most users interacted at least slightly with most other users, although some users in this cohort clearly interact more frequently with others. Only 466 user pairs (about 11 percent) did not interact with each other at all; several such pairs contain users who did not participate in the study for very long.

QUANTIFYING RENDEZVOUS: REALITY MINING

The widely available dataset from the Reality Mining project offers a concrete picture of how often and where people rendezvous.[7] This project used software on mobile phones to track user behaviors over the course of approximately nine months from 2004 to 2005. Among the 100 study participants, 75 were students or faculty at the MIT Media Lab, and 25 were incoming students at the MIT Sloan School of Management. The software collected statistics on incoming and outgoing communication, as well as the Bluetooth and GSM beacons detected by the phone.[9] The results exclude data from three of the 100 users because they did not participate significantly in the study. The study's large scope had a significant impact on our quantitative understanding of social interaction.

Because of a Bluetooth beacon's limited range, the Reality Mining data indicates when participants were close to one another. Since these users worked nearby, they provided a mechanism for studying the effects of increasing user density in a small area.

Note that the participant selection process for the Reality Mining study is unclear and almost certainly biases the data. However, the objective was to gather data in an area where a reasonable percentage of people use a system that leverages rendezvous. Therefore, a dataset that selects 100 users randomly from, say, all of Boston, would be far too sparse to simulate a high penetration rate for a rendezvousing application.

The MIT Media Lab consists of 300 to 400 people, and the Sloan School of Management, 1,000 people. Thus the Reality Mining data's gross density of users in these two groups is 20 to 25 percent and 2 to 3 percent, respectively. Obviously, biases in the selection process might cause the results to skew toward a much higher penetration, but that is not problematic for these results because the goal was to understand macroscopic rendezvous characteristics and how they might apply in a situation in which a mobile application has high penetration.

In Figure 1, each square represents the number of times a user rendezvoused with another user in the study, ordered by the total number of times a given user rendezvoused. Although some users in this cohort clearly interacted more frequently with others, most users interacted at least slightly with most other users. The most connected user rendezvoused 6,256 times during the study, and the most connected pair of users rendezvoused 585 times. The median number of rendezvous was 862, and the median number of rendezvous between two users was one.

This data suggests that most users rendezvous quite frequently. It also raises the question of how many people congregate at one time, and for how long people rendezvous. For each user in the study, Figure 2 shows the number of other nearby users versus time. The maximum number of nearby users was 15. Interestingly, many users frequently congregated with other users.

Figure 3 shows the time that a user is in proximity to a given number of other participants, ranked by fraction of time in proximity to at least one other participant. The 5th, 50th, and 95th percentile users were at 0.0039, 0.052, and 0.16. That is, the median user spent about 5.2 percent of the recorded study time in proximity to at least one other user in the study. This figure shows that, in fact, many participants spent a significant amount of time close to other participants, and even the least-connected participants rendezvoused occasionally.

This data does not, however, immediately reveal where it's possible to increase precision. Instead, it's necessary to consider GPS locations.

PARTICIPATORY ATMOSPHERIC SENSING

Consider an application that can benefit from detecting and exploiting rendezvous between users: *participatory atmospheric sensing*. Rendezvous provides opportunities for automatic sensor calibration, and for increasing system precision. Furthermore, personal mobile sensors will sample the most in areas of high user density, so for atmospheric compounds that affect people, the sampling density will, in some sense, approximate an ideal sampling distribution.

Participatory sensing has other important benefits. For example, since the sensors move with people, they track individual exposure. In our study conducted in Accra, Ghana, West Africa,[10] several users changed their daily routes and habits as a result of carrying pollution sensor hardware. Some users expressed concern and outrage over the pollution levels in their communities and sought ways to actively reduce pollution.

Whether motivated by the desire for social change, scientific inquiry, or personal interests, researchers have paid

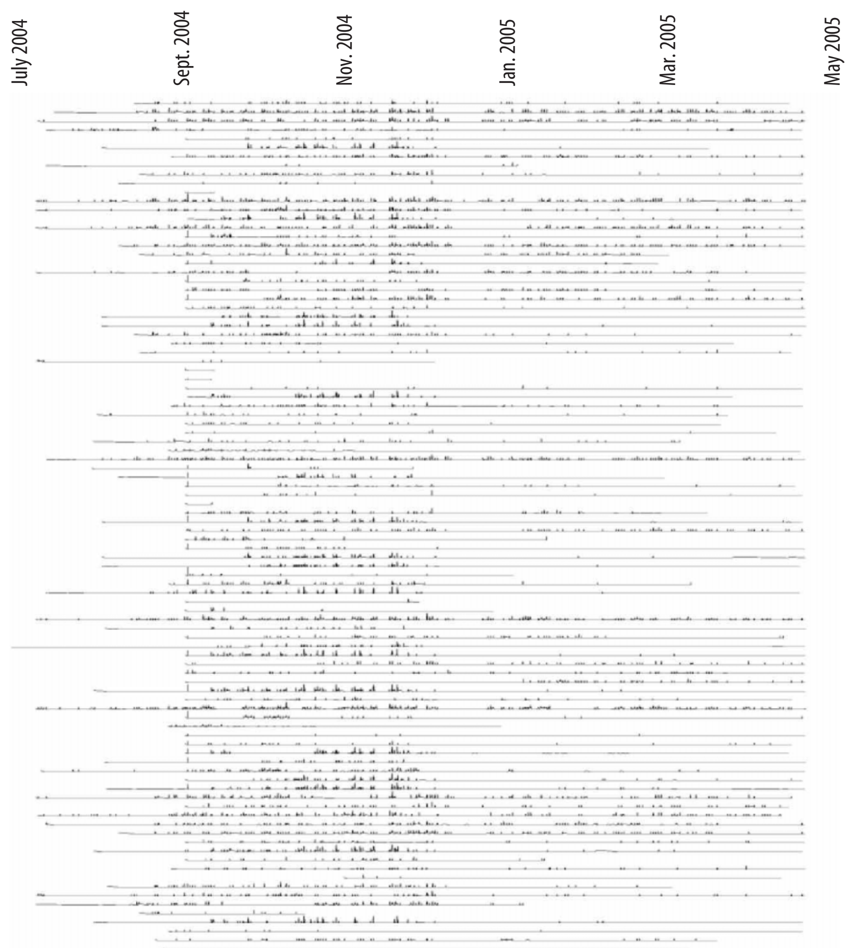

Figure 2. Each line shows the number of users in proximity to a given user, for each user in the study. The height of each line represents between 0 and 15 users in simultaneous rendezvous. Many users frequently congregated with several other users.

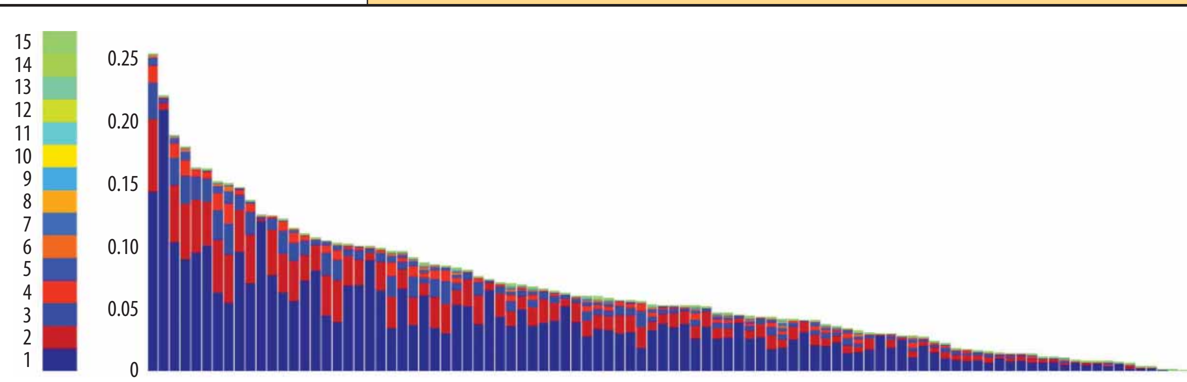

Figure 3. Fraction of time that a given user is in proximity to other participants, ranked by fraction of time in proximity to at least one other participant.

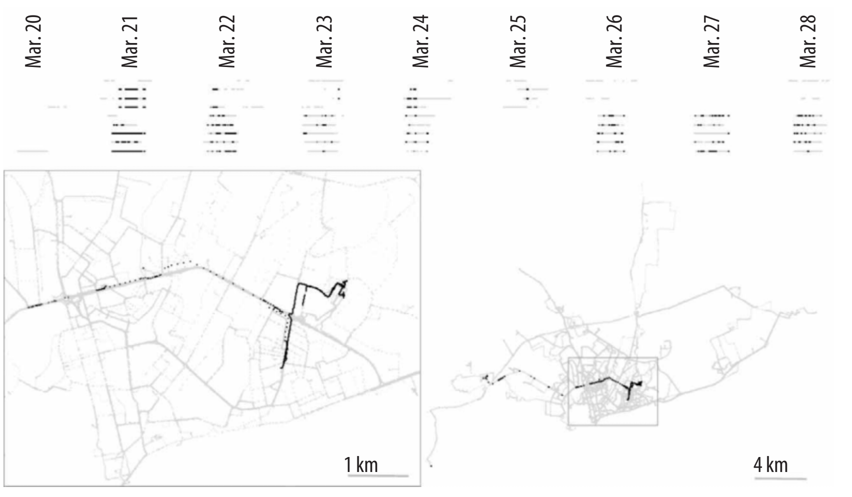

Figure 4. Rendezvous points (black) superimposed on all locations (gray) in the Ghana study. Sample times for each sensor appear at the top, again with rendezvous points in black and all sample times in gray. This figure shows that rendezvous occurred in hotspot locations, rather than being distributed throughout the map. They are not isolated to a few lucky coincidences, but distributed throughout the study in time and location.

increasing attention to the potential for individuals and scientists to harness information gathered by members of society.[6] For example, Jason Corburn examined the larger idea of *citizen science*, in which community members take part in the observation or analysis of their environment or in other scientific studies.[11] Corburn also discussed *street science*, in which participants use data to effect social change.

Our Common Sense project (www.communitysensing.org) has performed three major sensor deployments since 2007. Besides Ghana, the project has also deployed sensors on street sweepers in San Francisco[12] and badge devices carried by community activists in West Oakland.[13] These badges transmit sensor and location data into a database, but they do not contain voice equipment. Because the Ghana and West Oakland studies included user location information, they provided a resource for examining real-world user rendezvous.

The study in Ghana consisted of 10 users, one of whom left the study early due to equipment failure. Six users were taxi drivers, and four were university students. The taxi drivers ensured good geographic coverage, although their movement is not necessarily representative of typical people. As with the Reality Mining participants, these users came from a geographically dense subset of the population, and thus very roughly simulated an application having high user penetration. Shanghai Jiao Tong University recently completed a more extensive data collection campaign focused on taxi mobility (http://wirelesslab.sjtu.edu.cn).

Figure 4 shows the location histories of nine participants in the Ghana study, along with a timeline of samples. Again, the selection process for participants was heavily biased toward a small geographic area. The gray dots represent sample locations, and the black dots represent rendezvous between users. The users rendezvoused throughout the study's geographic area, although they tended to rendezvous in specific hot spots.

Figure 5 shows a similar plot for the West Oakland study. Again, users rendezvoused throughout the study period, supporting the intuition that rendezvous is a frequent occurrence rather than a lucky coincidence or abnormal event. The stars represent GPS error, which typically occurred when the devices were indoors, so they actually reduce the probability of rendezvous, because if users are in the same location, they have a lower chance of rendezvous in the event of a GPS error. The location data from this study is extremely biased because some participants spent significant time together, making the data valuable for calibrating sensors, rather than for illustrating typical rendezvous behavior.

Nonetheless, at a macroscopic level, these results are consistent with the findings of Lee et al., who analyzed approximately 150 GPS traces from 66 users in five separate locations and found that people moved among clusters of locations.[2]

Real-world location and sensor data provided a basis for automatically calibrating sensors in the Oakland study. Figure 6 shows the raw analog-to-digital converter data, smoothed ADC data, and the data after automatic calibration.[6,7] Since there is no absolute reference sensor in the data, the y-axis scale is unlabeled. We only know the relative calibration unless one of the sensors has absolute calibration information. If we had either the true calibration of one of the sensors, or knew the distribution of sensor miscalibrations, we could infer the absolute calibration of each sensor.

Although study participants moved independently, they spent the most time in the same general area. Thus, the smoothed sensor signals, which roughly correspond to the ambient pollution concentration, should match. Figure 6 shows that this is indeed the case after calibration.

As the immersion of sensors and communication into daily life continues, applications will emerge that can combine data that individuals acquire with data from other sources. These important applications will help make sense of oceans of noisy data.

Mapping applications that contextualize user data have taken a first step in this direction, but leave data fusion for users to perform. Other potential applications in this area might include an audio capture and enhancement application, indoor localization through time-of-flight calculation of sound events recorded by different devices, and forming ad hoc phase array antennas for higher-gain GPS signal reception. Privacy and energy are important issues that warrant separate treatment.[14,15]

Of course, economics will dictate what sensors make their way into phones and other mobile devices. Although altruism might suffice to spur the development of such applications, the objective is to find a compelling economic model for collaborative sensing applications in general, and for mobile atmospheric sensing in particular. Regardless, sensors will continue to penetrate further into everyone's lives and communities, and rendezvous will play an important role in aggregating and enhancing data from individuals to produce information that benefits society.

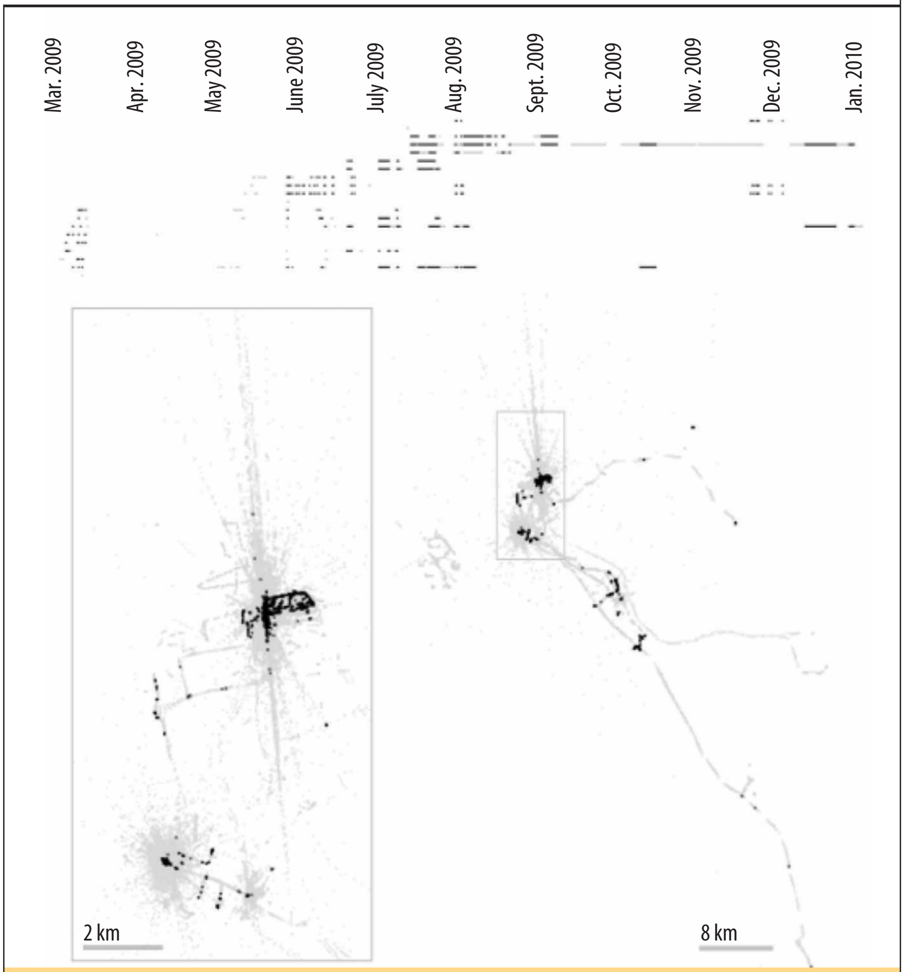

Figure 5. Rendezvous points (black) superimposed on all locations (gray) in the Oakland study. Sample times for each sensor appear at top; rendezvous points are in black, and all sample times in gray. As in the Ghana data, rendezvous happen in hot-spot locations, rather than distributed throughout the map. They are not isolated to a few lucky coincidences, but distributed throughout the study in time and space. The stars represent GPS error.

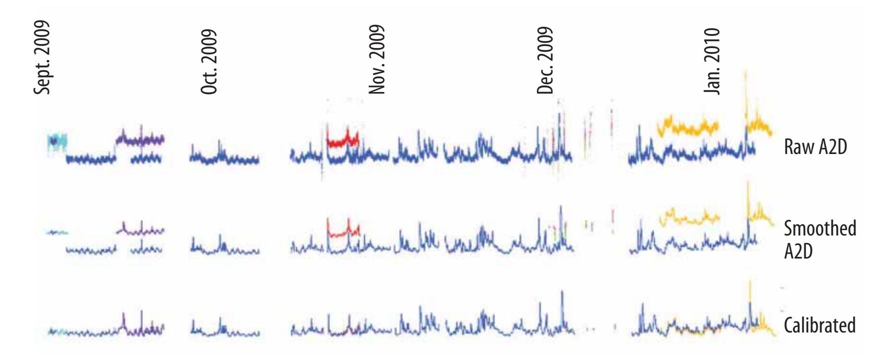

Figure 6. The top row shows raw data from the Intel badges' ADC versus time. Colors represent different sensors. The bottom row shows the same data after calibrating each sensor using the research algorithm, and smoothing the data to distinguish the different sensors. As expected, this data shows roughly the same readings, since the ambient pollution concentration is relatively constant over the study area. The middle plot shows uncalibrated data after smoothing. The calibrated data is missing an absolute reference, so no concentration value is indicated.

Acknowledgments

This research was conducted at the University of California, Berkeley, Intel, and Nokia, and was supported in part by the National Science Foundation.

References

1. M.C. Gonzalez, C.A. Hidalgo, and A.-L. Barabasi, "Understanding Individual Human Mobility Patterns," *Nature*, June 2008, pp. 779-782.
2. K. Lee et al., "SLAW: A New Mobility Model for Human Walks," *Proc. 28th Conf. Computer Comm.* (Infocom 09), IEEE Press, 2009, pp. 855-863.
3. P. Hui, J. Crowcroft, and E. Yoneki, "Bubble Rap: Social-

Based Forwarding in Delay-Tolerant Networks," *Proc. 9th ACM Int'l Symp. Mobile Ad Hoc Networking and Computing* (MobiHoc 08), ACM Press, pp. 241-250.

4. Z. Zhang, "Routing in Intermittently Connected Mobile Ad Hoc Networks and Delay-Tolerant Networks: Overview and Challenges," *Comm. Surveys & Tutorials*, vol. 8, no. 1, 2006, pp. 24-37.
5. E. Miluzzo et al., "CaliBree: A Self-Calibration System for Mobile Sensor Networks," *Proc. 4th Int'l Conf. Distributed Computing in Sensor Systems* (DCOSS 08), ACM Press, 2008, pp. 314-331.
6. E. Canessa and M. Zennaro, eds., *mScience: Sensing, Computing and Dissemination*, ICTP—The Abdus Salam Int'l Centre for Theoretical Physics, 2010; www.m-science.net/book.
7. R.J. Honicky, "Towards a Societal Scale, Mobile Sensing System," doctoral dissertation, Dept. Computer Science, UC Berkeley, 2010.
8. S. Eisenman, N. Lane, and A. Campbell, "Techniques for Improving Opportunistic Sensor Networking Performance," *Distributed Computing in Sensor Systems*, S. Nikoletseas et al., eds., LNCS 5067, Springer, 2008, pp. 157-175.
9. N. Eagle and A. Pentland, "Reality Mining: Sensing Complex Social Systems," *Personal and Ubiquitous Computing*, May 2006, pp. 255-268.
10. E. Paulos, R.J. Honicky, and E. Goodman, "Sensing Atmosphere," *Proc. 5th ACM Conf. Embedded Networked Sensor Systems* (SenSys 07), ACM Press, 2007; http://repository.cmu.edu/hcii/203.
11. J. Corburn, *Street Science: Community Knowledge and Environmental Health Justice*, MIT Press, 2005.
12. P.M. Aoki et al., "A Vehicle for Research: Using Street Sweepers to Explore the Landscape of Environmental Community Action," *Proc. Conf. Computer-Human Interaction* (CHI 09), ACM Press, 2009, pp. 375-384.
13. W. Willett et al., "Common Sense Community: Scaffolding Mobile Sensing and Analysis for Novice Users," *Proc. 8th Int'l Conf. Pervasive Computing* (Pervasive 10), LNCS 6030, Springer, 2010, pp. 301-318.
14. A. Kapadia et al., "Anonysense: Opportunistic and Privacy-Preserving Context Collection," *Pervasive Computing*, J. Indulska et al. eds., LNCS 5013, Springer, 2008, pp. 280-297.
15. Y. Wang et al., "A Framework of Energy Efficient Mobile Sensing for Automatic User State Recognition," *Proc. 7th Int'l Conf. Mobile Systems, Applications, and Services* (MobiSys 09), ACM Press, 2009, pp. 179-192.

R.J. Honicky, *a researcher at Nokia Research Center, Berkeley, studies mobile architectures and interfaces for improved energy consumption and performance, as well as mobile sensing. He also focuses on appropriate technologies for people in economically developing countries. Honicky received a PhD in computer science from the University of California, Berkeley. Contact him at honicky@gmail.com.*

GUEST EDITOR'S INTRODUCTION

IT Footprinting—Groundwork for Future Smart Cities

Sumi Helal, *University of Florida*

The goals for developing smart cities are clear and convincing, and the technology is promising and exciting, but achieving these goals requires a massive IT footprinting process.

Recent advances in pervasive and ubiquitous computing provide a glimpse into the future of our planet and reveal exciting visions of smart cities, homes, workplaces, hotels, schools, and much more.

Driven by a technological evolution offering "low-power many things and wireless almost everything"—for example, IEEE 802.15.4 radio, wireless sensor networks, and sensor platforms—we could, in only a decade, envision and prototype impressive smart-space systems that improve quality of life, enhance awareness of resources and the environment, and enrich user experiences.

For most of these systems, the goals are clear and convincing. But prototyping is one thing—commercial proliferation and creating a new industry are another.

THE PATH TO SMART CITIES

In particular, many governments around the world are pursuing the development of smart cities. Driven by increasing urbanization and serious economic and environmental challenges, smart cities are emerging as a way to offer technology solutions to bridge the widening gap between supply and demand while reducing urbanization's impact on the environment. This effort also offers an opportunity to recover from the effects of the global recession by creating new green industries and businesses.

Yoshiaki Kushiki, corporate advisor to Panasonic Corporation, points to an important requirement for such new industries to succeed: consumers must also change by choosing "spiritual richness" over "material affluence." This is a valid observation because most of the products and services in smart cities will offer intangible benefits such as energy savings, sustainability, and reduced CO2 emissions in contrast to a "materials only" approach—for example, acquiring the newest and fanciest iPhone on the market.

In principle, the path to smart cities is obvious—embed sensors, actuators, or computers into objects and spaces that make up the smart city's important elements. This "smartening" approach is not new, having driven the embedded systems industry for more than two decades. What's new here are the massive scale and the new ecology that the smartening process requires.

MASSIVE SCALE

Smartening an entire city is a massive IT footprinting process. To appreciate the scale, think about these questions. How many people live in a large city? How many homes, apartment buildings, and office buildings are there? How many electricity meters? How many wall power plugs? How many cars and parking spots? This is a sample set, but the list could grow very large, and the total number of objects and spaces that will have to be smartened could easily number in the hundreds of millions.

The smart grid alone—an essential part of most envisioned smart cities—will require installing a smart meter in every structure drawing electric power, thereby injecting tens of millions of sensing devices into the cyber-physical infrastructure. The smart grid will also require other controllers that integrate the electrical grid with the IT infrastructure, starting from the supply side, passing through the distribution networks, and ending at the demand side.

 Published by the IEEE Computer Society

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

Developing smart cities may start with the smart grid, but it doesn't end there. CO_2 zero-emission homes and net-zero-energy buildings may necessitate additional IT footprints. Plug-in hybrid and electric vehicles will also require IT footprinting in the form of in-dash PCs. In addition to energy, other concerns such as healthcare, elder care, education, security, public safety, smart transportation, entertainment, and emergency response will require massive IT footprinting.

NEW ECOLOGY

But do we have all the ingredient technologies for such massive IT footprinting? Unfortunately, the answer is no. Achieving this objective requires rethinking embedded computing's role and goals. Existing embedded operating systems (OSs), for example, will need to change. They're currently focused only on smartening an object or a device—they're smart but introverted. They don't know about the cloud or the edge devices in between.

To put it another way, current embedded OSs are for smart cities what DOS was for the PC in the early 1980s. We need a new breed of embedded OSs capable of natively and autonomically integrating, connecting, and programming enormous numbers of sensors to the cloud. Without this fundamental change, the IT footprint will simply look like a massive array of stovepipes that needs an army of engineers to integrate it, no matter how smart each individual pipe is.

Advances in hardware are also required. Common components and subsystems will need to enter the marketplace at a faster pace. Memory cards are an example of an industry that has successfully followed this model. Developers will need to tightly integrate more "brain" and communication elements into common modular components in useful packages and form factors. Perhaps what's needed is an ecosystem parallel to the auto parts industry, offering the same standardized functional part in various forms, sizes, and shapes.

Once IT footprinting is enabled, adding point-of-service delivery and interaction platforms such as set-top boxes, smartphones, and public displays will complete the necessary infrastructure for smart-city rollouts. Then, a new world of powerful and serious applications and services will be unleashed, making the app store as commonplace as the local grocery store or Facebook.

This special issue offers four articles on smart cities covering the IT perspective, applications and services, interaction and user experience enrichment, and a fast-progressing smart-city initiative in South Korea.

The topic will obviously be revisited in the future to assess progress and refine our understanding of the changing landscape of goals and requirements. We look forward to your active participation in this area and to receiving your future contributions on this topic. C

Sumi Helal *is a professor in the Department of Computer and Information Science and Engineering at the University of Florida. His research interests include pervasive computing, mobile computing, and networking. Helal received a PhD in computer sciences from Purdue University. Contact him at sumi.helal@gmail.com.*

COVER FEATURE

Smarter Cities and Their Innovation Challenges

Milind Naphade, Guruduth Banavar, Colin Harrison, Jurij Paraszczak, and Robert Morris, *IBM*

The transformation to smarter cities will require innovation in planning, management, and operations. Several ongoing projects around the world illustrate the opportunities and challenges of this transformation.

Cities are experiencing unprecedented socioeconomic crises. Urban growth and migration are putting significant stress on city infrastructure as demand outpaces supply for water, energy, transportation, healthcare, education, and safety. To reduce costs, improve efficiencies, and deliver the quality of life citizens expect while balancing budgets, cities are increasingly looking to information and communications technology (ICT) and new working practices.

The transformation to smarter cities will require innovation in planning, management, and operations. Several ongoing projects in Brazil, the US, Denmark, South Korea, and other countries illustrate the opportunities and challenges of this transformation.

NEED FOR SMARTER CITIES

There is an urgent need for cities worldwide to become smarter in how they manage their infrastructure and resources to cater to the existing and future needs of their citizenry. Concurrent trends in urbanization, economic growth, technological progress, and environmental sustainability are the drivers for this newfound urgency.

Urbanization

More than 50 percent of the world's population now lives in cities. By 2050, the UN forecasts this number to increase to 70 percent due to growth in the current urban population and migration from rural areas.[1] Some of this growth will be in 27 megacities with greater than 10 million people, but more than half of this growth will occur in cities that currently have fewer than 500,000 people. Urban infrastructure already experiencing stress will be hard-pressed to provide even basic services, while emerging cities will face greenfield development challenges.

Economic growth

The top 100 urban conglomerations currently account for 25 percent of the worldwide gross domestic product. By bringing people together, cities stimulate creativity and entrepreneurship, which further spurs economic activity. While the developed world has underinvested in its cities, the developing world by some estimates will need $40 trillion by 2030 for its new urban infrastructure, which presents tremendous innovation opportunities.

Technological progress

ICT advances have revolutionized all aspects of life. Two billion people use the Internet, and more than five billion are mobile subscribers.[2] There are 30 billion RFID tags embedded in our world and a billion transistors per human, each costing one ten-millionth of a cent.[3] This convergence of pervasive sensing and networking, wireless connectivity, and cheaper, faster, smaller computers has made it easier to intelligently control systems and empower people.

Environmental sustainability

There is evidence that human activity has caused unprecedented environmental change, and population growth will soon stress the world's natural resources to the breaking point. Global warming, air pollution, land degradation, declining per-capita availability of fresh water, food shortages, and reduced biodiversity are some of the starkest challenges. Top priorities for cities include sustain-

 Published by the IEEE Computer Society

ing water, energy, and food supplies, managing waste (95 percent of cities still dump raw sewage into their waters), and reducing greenhouse gas emissions.

Contrary to popular opinion, urban life is often greener than suburban and rural life: inhabitants consume less energy and space for living and use less energy for transportation. In developed societies, urban inhabitants are less dependent on fossil fuels and make more journeys on foot than those living in less densely populated areas.

SMARTER CITY TRANSFORMATION

It can take a decade for a city to become truly smart. Sometimes the impetus for transformation is recovery from a natural disaster, an impending large-scale event, or a sizable government investment. At other times visionary city leaders galvanize the citizenry and business community to channel their energy and resources into such a project.

Assessment

Once it has decided on transformation, a city must evaluate its needs and innovation opportunities, set clear objectives, prioritize development efforts, and establish metrics that let city planners, ICT consultants, and residents assess progress.

IBM and other organizations have created several tools to facilitate this process[4] including the Smarter City Assessment Tool, the Actionable Business Architecture, and the Municipal Reference Model. Researchers can use IBM's Component Business Model to partition a city into independent operational units, and they can apply its Smarter City Maturity Model to various domains ranging from resident-oriented services, such as social services and public safety, to structural functions, such as road maintenance and traffic management.

Assessment should be flexible enough to let cities choose the domains most important to them and provide some means of projecting costs and measuring progress—for example, how well various agencies are integrating operations and sharing data. Applying maturity models to numerous cities has also revealed common patterns that can help guide investment strategies and predict the outcome of adopting a particular approach. For example, investing heavily in road expansion without a smarter transit policy can reduce the use of public transportation.

A 'system of systems'

As Figure 1 shows, at the highest level, a smarter city integrates and optimizes a set of interdependent public and private systems to achieve a new level of effectiveness and efficiency. These systems are increasingly both producers of information and consumers of one another's information, although interactions can also be indirect. Hence, a smarter city can be viewed as a "system of systems."

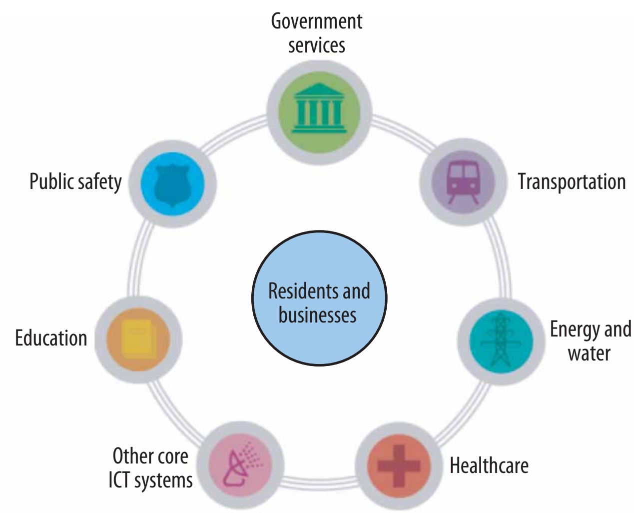

Figure 1. A smarter city constitutes a "system of systems"—a set of interdependent public and private systems that the city can integrate and optimize to achieve a new level of effectiveness and efficiency.

Smarter city transformation relies on the use of powerful analytical techniques to extract insights from real-world events to improve urban business processes. These processes can be broadly divided into planning, management, and operations.

Planning. A smarter city provides urban planners with tools to exploit various sources of information about human behavior to aid in the allocation of resources—land, water, transportation, and so on—as the city evolves. Holistic modeling of the city's ecosystem provides quantitative support for strategy development, performance evaluation, identification of emerging best practices, and integration of initiatives. Analyzing data from other comparable cities can help planners calibrate urban dynamics models and compare their relative progress.

Management. A smarter city can coordinate infrastructure management activities—the creation and maintenance of roads, equipment, and other assets—by providing cross-agency visibility of planned interventions. For example, the electrical utility's replacement of a cable under a street intersection might offer traffic managers an opportunity to save money by replacing a signal at the same location. By providing a time dimension, smarter city data can reveal historical views of each domain and enable managers to project its evolution.

Operations. A smarter city integrates multiple data sources to represent the interdependence of urban domains in real time. For example, electrical utilities can combine sophisticated models of near-term demand based on historical usage patterns (day of week, holidays, local weather, major events, and so on) with real-time traf-

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

COVER FEATURE

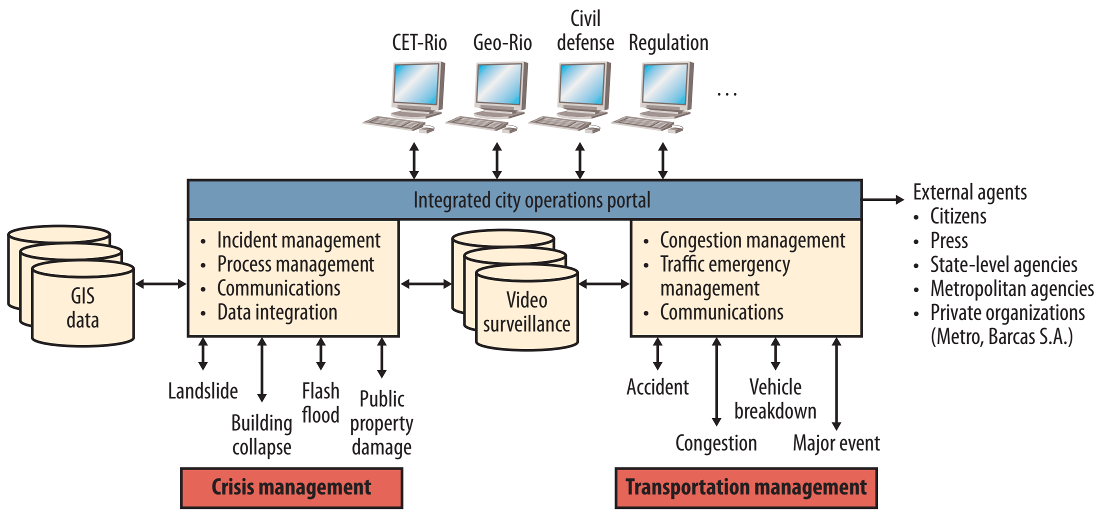

Figure 2. Rio de Janeiro's Operations Center integrates dynamic, real-time data and multiple systems to improve crisis and transportation management.

fic information that could impact future demand. Thus, awareness of a major delay in outbound commuter traffic in the early evening could let the utility project a delay in demand because those commuters will arrive home late. The utility could likewise use real-time weather data to predict the location of cables damaged in a rainstorm.

EXAMPLE SMARTER CITIES

Examples of cities of various size, geography, and economy illustrate different aspects of smarter city transformation.

Rio de Janeiro, Brazil

Every summer, Rio de Janeiro faces the consequences of intense rainfall, including landslides and flooding. In April 2010, the region endured one of the worst series of torrential rainstorms in decades, in which mudslides killed more than 200 people, left tens of thousands homeless, and caused more than $13 billion in damage (http://en.wikipedia.org/wiki/April_2010_Rio_de_Janeiro_floods_and_mudslides). The resulting chaos, loss of life, and destruction of property motivated state and city authorities to implement advanced ICT capabilities in Rio to better manage disasters and emergencies as well as planned events of national importance.

Using a substantial monetary investment resulting from the city's selection to host the 2014 World Cup and the 2016 Summer Olympics, and under the visionary leadership of Rio's mayor, Eduardo Paes, the city has embarked on an ambitious program to connect multiple systems to improve crisis and transportation management.

The Rio Operations Center[5] opened in December 2010. As Figure 2 shows, it includes an incident management system that will help the city to prepare for and respond to flood-related incidents and a process management system through which multiple agencies can make coordinated and intelligent decisions based on dynamic data from weather sensors, video surveillance, and field personnel, overlaid on a comprehensive geographic information system (GIS). As this system evolves, researchers can integrate data from transportation systems, buildings, and possibly energy, water, and other subsystems into the Rio Operations Center to create a true closed-loop system.

Dubuque, Iowa

The city of Dubuque is an example of a smarter city in which an economic crisis motivated a transformation. Following the demise of the wood-milling industry in the 1980s, this city of 60,000 has evolved into a vibrant hub of sustainable development by using ICT to optimize resources and operations for its citizens and city management. This has led to one of the fastest urban economic turnarounds in the US, with a diversified workforce and local industries working to make water, energy, and buildings sustainable.

In 2006, Dubuque created a sustainability model with three major themes: economic prosperity, sociocultural vibrancy, and environmental integrity. The model was based on 11 core principles including smart use of energy, water, and other resources; green buildings; reasonable mobility; and greater community knowledge.

Three years later, Dubuque partnered with IBM to become a living lab for smarter city sustainability management that monitors water and electricity consumption in homes using smart meters and provides guidelines and incentives to residents to optimize their individual consumption.[6] Figure 3 shows the citizen dashboard for integrated sustainability management.

Dubuque has been a pioneer in putting citizens at the center of smarter city transformation, a vital aspect that is often ignored. The city also integrates planning, management, and operations both to minimize its carbon footprint and to encourage economic growth.

Bornholm, Denmark

On Bornholm, Denmark, an island of 40,000 inhabitants, international researchers are developing a system that integrates electric vehicles (EVs) into the local power grid, which relies heavily on renewable wind power. As Figure 4 shows, the EDISON system—Electric Vehicles in a Distributed and Integrated Market Using Sustainable Energy and Open Networks (www.edison-net.dk)—includes a network of public and personal charging stations and integrated technologies to manage the charging of EVs as well as load balancing, billing, and so on.

In addition to reducing carbon dioxide emissions, EVs can act as supplemental storage devices that send power back to the grid when needed. Offshore wind power provides 20 percent of Denmark's electricity, and as part of a long-term strategy to increase that figure to 50 percent, the government hopes to deploy some 200,000 EVs nationwide by 2020.[7]

Songdo IBD, South Korea

Seoul, the capital of South Korea, already has an advanced ICT infrastructure that provides inexpensive, superfast broadband access to its more than 24 million citizens from virtually anywhere within the city limits. In addition to accelerating the country's economic growth, this universal connectivity has transformed governance, resulting in the world's most advanced and efficient e-government.[8] Among the many benefits of this transformation is greater transparency—the online disclosure of all bids for government contracts has significantly reduced corruption.

Work is now under way on the Songdo International Business District (www.songdo.com), a ubiquitous ecological (u-eco) city 40 miles west of Seoul on reclaimed land in the Incheon Free Economic Zone. When completed in 2015, Songdo IBD will include 80,000 apartments, 50 million square feet of office and retail space (including the 68-story Northeast Asia Trade Center, the tallest building in the country), a hospital, arts and convention centers, an

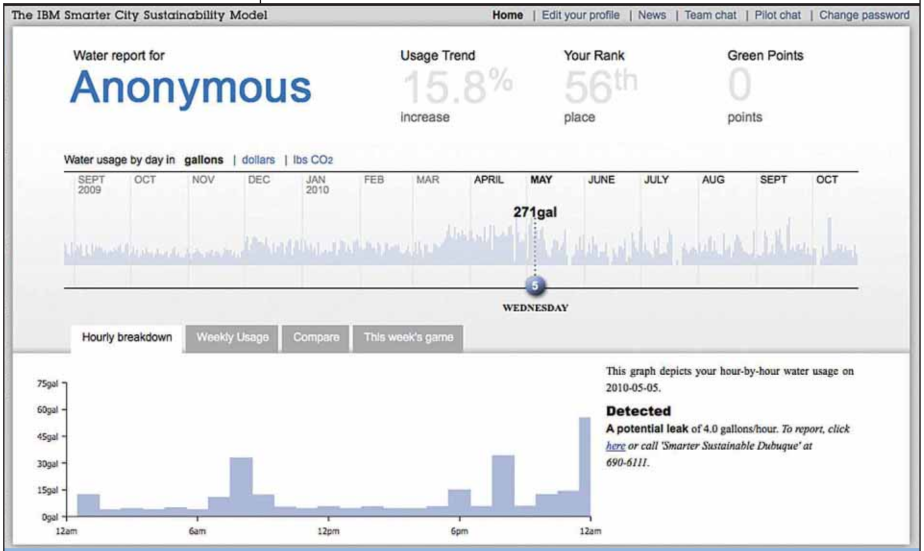

Figure 3. Citizen dashboard for integrated sustainability management in Dubuque, Iowa.

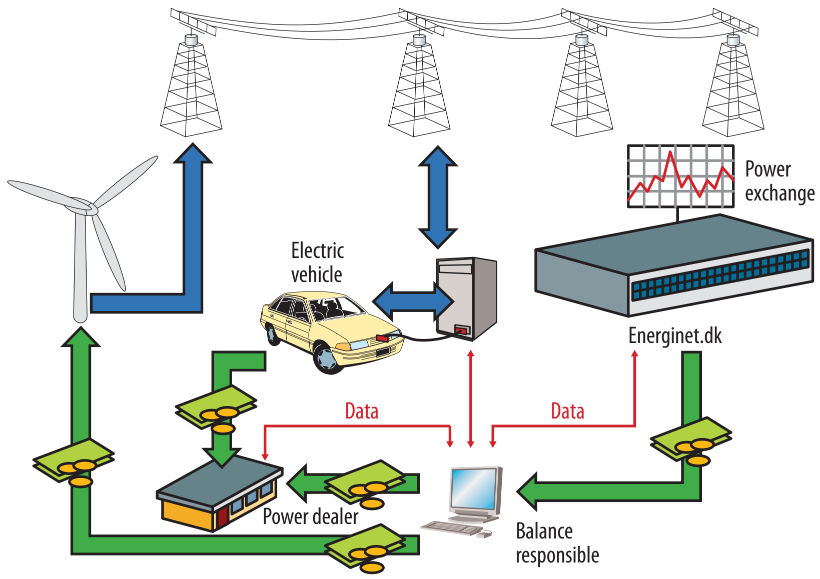

Figure 4. On the Danish island of Bornholm, the EDISON project—Electric Vehicles in a Distributed and Integrated Market Using Sustainable Energy and Open Networks—is developing a system that integrates electric vehicles into the local power grid, which relies heavily on renewable wind power.

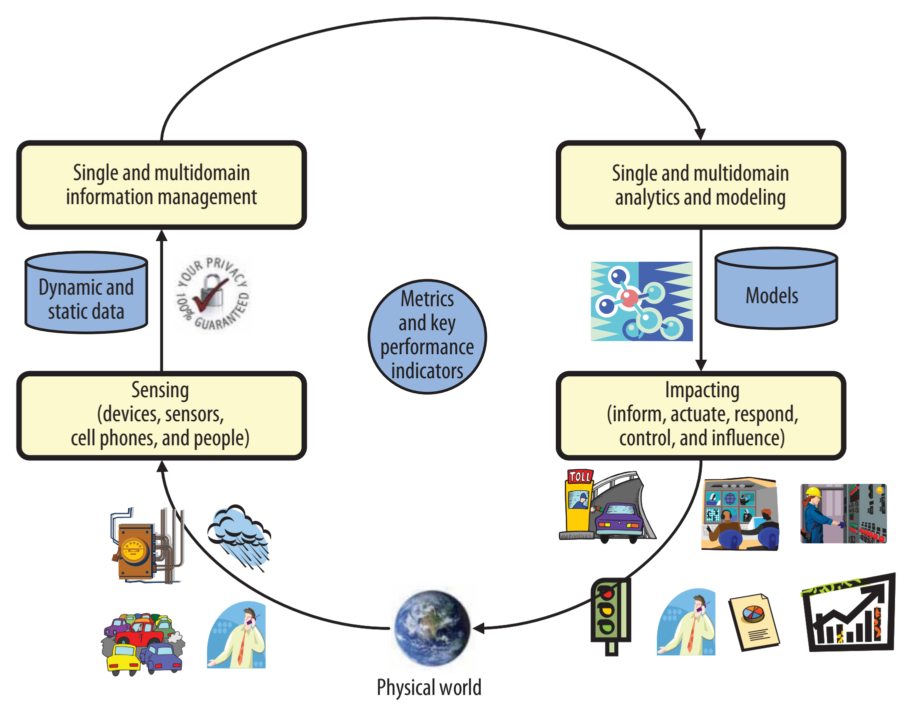

Figure 5. Smarter cities technology innovation framework. To optimize key metrics and performance indicators, all systems must be integrated in a closed loop.

academic complex, and 600 acres of open space, making it the largest private real-estate development in history.

A wide-area network of computers will link Songdo IBD's structures and offer citizens and businesses various digital services. For example, individual apartments feature panels in each room that control lighting, temperature, and access to media; 20,000 residential units will feature telepresence technology. A green, state-of-the-art datacenter will help manage all aspects of urban life, from traffic control to water and energy use to recycling.

Other smarter cities

Many other cities and regions around the world are using technology innovation to improve their planning, management, and operations.

Malta is creating the world's first countrywide smart grid, using ICT to optimize its water and energy systems.[9] Stockholm, Sweden, has implemented a system that automatically charges drivers a fee based on how much they drive (using control points outfitted with lasers and cameras) to reduce congestion and greenhouse gas emissions.[10] India's Gujarat International Finance Tec-City (http://giftgujarat.in/index.aspx) is a greenfield development using advanced technology to, among other things, eliminate discharge waste and achieve 90 percent use of public transit. Portland, Oregon, is using a system dynamics model to discover and coordinate dependencies among key city systems. And Singapore is aspiring to become an innovative technology leader by using ICT to achieve water sustainability and to manage traffic and energy.

INNOVATION CHALLENGES

Figure 5 illustrates the ideal of a smarter city as a closed loop of interconnected city systems. These systems can be characterized by function: sensing, information management, analytics and modeling, and influencing outcomes. To optimize key metrics and performance indicators, all of these systems must be tightly integrated.

Sensing a city and its inhabitants

A city is full of sensors—smart water and electric meters, mobile phones, GPS devices, traffic sensors, parking meters, pipe sensors, weather sensors, building sensors, and so on. Even people can be sensors—using crowdsensing to gather intelligence on city operations is an emerging research area. The main innovation challenges in sensing a city and its inhabitants are trading off cost with quality and dual usage, and ensuring privacy and security.

Cost versus quality. Cheap, ubiquitous sensors can be used in large numbers, but their noisy, low-quality signals impose a nontrivial burden on analytics systems and might also require frequent calibration and diagnostic evaluation. In contrast, high-cost sensors with embedded intelligence can make analysis simpler and more accurate, and may be self-calibrating and -diagnosing, but cannot be installed in the quantities needed to cover large areas.

Cost versus dual usage. A sensor used for purposes for which it was not originally intended will not yield high-quality data, but replacing it with the proper sensor or augmenting it with another sensor can be prohibitively expensive.

Dubuque, for example, uses smart electric meters in homes to aggregate resource consumption at 15- or 60-minute intervals. Originally designed for billing and dynamic pricing, these meters are not as accurate as having a sensor on every electrical device. Disaggregating energy use for each house requires sophisticated analytics but is currently cheaper than instrumenting homes with a large number of localized sensors. This problem could be alleviated if future appliances have built-in functionality to report their energy consumption to home-area networks.

Another example is the use of mobile phones for location sensing. Mobile phones can be used to estimate users' locations through the data generated at cell towers. This is an inexpensive way to measure positional data and enables a plethora of smarter city applications such as traffic management and emergency response—a particularly attractive option for a developing economy—but its high range of error would require sophisticated analytics.

Privacy. The sensors that produce the best data and enable the most effective modeling are also the most intrusive, and thus more likely to make inhabitants of the environment being sensed uncomfortable. This is another reason why putting energy-monitoring devices on every home appliance is currently impractical.

Security. Significant research is needed to ensure that both sensors and actuators are secure when it comes to the acquisition, storage, and transmission of information. Tampering with or snooping on sensed data could result in nightmarish scenarios ranging from thieves knowing when residents are not at home to terrorists turning off a city's power and water. Cybersecurity must be taken to a new level before sensors can be deployed on a massive scale.

Managing information across all city systems

The main research challenges in managing smarter city data are the need for common information models and the ability to safely share information across multiple agencies within a city and among multiple cities in a metropolitan region.

Information models. To ensure end-to-end visibility while managing smarter city infrastructure and services, it is necessary to integrate data from disparate sources, each with its own sampling frequency, latency characteristics, and semantics. For example, information related to roads is scattered across many agencies including transportation, urban planning, public works, emergency services, public safety, and environmental management.

Creating and applying a unified information model makes it possible to obtain a more complete picture of urban activity. The ability to understand how combinations of factors contribute to, say, a rapid increase in the demand for water or an unusually high accident rate on a stretch of road in turn facilitates better operational decisions.

A principal technology base for managing spatial information and entities in cities is a GIS with associated databases and mapping tools. In the future, this model will be extended to include not only static data such as topography, land use, and the built environment but also dynamic information such as service delivery, resource consumption, and the movement of people (which can be used to infer behavior), vehicles, and freight.

Privacy, security, and access control. Securing data from sensed urban environments is a major research challenge. City datacenters constitute the largest potential single point of failure and thus require the most stringent security. Authorities must also design and implement privacy policies to prevent unauthorized access to data. In addition, access control mechanisms must be in place to ensure that visualization, analytics, and modeling applications do not misuse data. For example, mobile phone data used to sense traffic congestion with the explicit approval of the devices' owners should not be used to issue speeding citations.

Securing and controlling access to a plethora of data streams and applications will require a fairly agile, real-time implementation engine. New models are also needed to encourage the utmost transparency within required constraints.

Creating and applying a unified information model makes it possible to obtain a more complete picture of urban activity.

Standards and interoperability. Different cities and even different agencies within the same city adopt different models to manage information. Some industries, such as the electrical utility industry, have a well-developed set of standards, while others, such as the water utility industry, seem unaware of the concept. Initiatives like the Municipal Reference Model,[11] developed by the Municipal Information Systems Association of Canada, aim to incorporate the perspectives of the many stakeholders involved in urban infrastructure and services. More such models should be developed and refined through an open industry process.

Observing and understanding city activity

The availability of massive amounts of sensed information opens up fascinating opportunities to understand city activity through modeling and analytics. This will require real-time and batch analysis of heterogeneous data with cross-domain dependencies in the presence of significant uncertainty and variability.

System models. Models of city systems must be statistical as well as physical. For example, Dubuque and Malta model resource consumption at both the individual and community level, Stockholm models traffic congestion, and Denmark's EDISON project models wind power generation and EV-based energy storage. Traffic management, emergency response, and other such services in Rio rely on weather prediction models.

The challenge in building such models is the lack of ground-truth data to accurately calibrate and train them. Models need to be bootstrapped with existing data and adaptable to changes in the state of the city. Another set of challenges revolves around the difficulty of predicting human behavior—for example, how residents will respond to green incentives.

Analytics. By studying how its systems interact with one another and with users over a long period of time, a city can operate its infrastructure and services more effectively—the equivalent of a company applying enterprise resource planning. Recurring patterns, anomalies, and evolving events in models can reveal important insights, and "what if" scenarios can help in planning.

For example, Stockholm's traffic congestion models enable better road-use charging policies, Rio's fine-grained rainfall prediction models facilitate emergency planning and traffic management, Singapore is using sophisticated traffic models to optimize transportation operations, and Dubuque's energy consumption models help Iowa's utility board and energy independence office perform statewide energy policy planning.

Influencing outcomes

The ultimate goal of introducing smarter city technology is to optimize control of city systems and to offer citizens both a wider range of choices and real-time feedback to influence their behavior and thus obtain better outcomes.

Optimal system control. Conceptually, a smarter city can be viewed as an interdependent collection of closed-loop systems. At the physical engineering level, a closed-loop system can be likened to a process control system automatically (except in extreme cases) regulated through appropriate feedback and control. At the decision support level, a closed-loop system is automated to a limited degree through the application of rule-based models.

Optimizing resources across all systems as well as within each subsystem is a challenging problem that must address cross-domain and cross-system dependencies that may also be cyclic. In the EDISON project, for example, grid electricity capacity in part depends on how much power EVs store, but the energy they feed back to the grid depends on their usage, which in turn depends on traffic congestion. Managing energy generation and storage and managing vehicle use are each difficult, but optimizing the two systems together is a highly complex cross-domain resource optimization problem. Ensuring fail-safe operations of interdependent systems is also critical.

Human-city interaction. People are at the center of a city's transformation into a smarter city. They are important sources of data, both about themselves and about the physical world, and can become willing participants if they can easily provide the data—for example, through mobile phones—and if they see the value of the information resulting from their contribution. While information models and analytical algorithms can organize data, provide insights, and make simple decisions, it will be up to humans to make the complex decisions.

Modeling, understanding, and influencing human behavior, providing location-aware information and feedback, and cultivating trust in smarter city technologies are key challenges that will draw on work in psychology, user experience design, urban systems, location-aware services, game design and game theory, and social computing.

Research will focus on what information to present to users, when and how to present it, and what reactions to expect. It will also explore how to motivate positive behavioral change. For example, Dubuque matches 50 percent of the costs of fixing household water leaks if users report the leaks and follow through on repairs, and its insightful feedback on water consumption has increased conservation by 10 percent. Likewise, Seoul is using its citywide Internet connectivity to incentivize use of public transit by providing online information to drivers about insurance discounts, reduced-cost parking, and a tax break for leaving their cars at home one business day a week.

Cities must get smarter to address an array of emerging urbanization challenges, and as the projects highlighted in this article show, several distinct paths are available. The number of cities worldwide pursuing smarter transformation is growing rapidly. However, these efforts face many political, socioeconomic, and technical hurdles.

Changing the status quo is always difficult for city administrators, and smarter city initiatives often require extensive coordination, sponsorship, and support across multiple functional silos. The need to visibly demonstrate a continuous return on investment also presents a challenge. The technical obstacles will center on achieving system interoperability, ensuring security and privacy, accommodating a proliferation of sensors and devices, and adopting a new closed-loop human-computer interaction paradigm. C

References

1. United Nations, *World Urbanization Prospects: The 2009 Revision—Highlights*, 2010; http://esa.un.org/unpd/wup/doc_highlights.htm.
2. Int'l Telecommunication Union, *The World in 2010: ICT Facts and Figures*, 2010; www.itu.int/ITU-D/ict/material/FactsFigures2010.pdf.
3. S. Carter, "Smart SOA: SOA Imperatives for Today … and for the Future," keynote presentation, 2009 SOA Symp.,

2009; http://soasymposium.com/pdf/2_Sandy%20Carter_SOA%20Symposium_For%20Publication.pdf.

4. M. Kehoe et al., *Smarter City Series: A Foundation for Understanding IBM Smarter Cities*, IBM Corp., 2011; www.redbooks.ibm.com/redpapers/pdfs/redp4733.pdf.
5. "Film Operations Center Rio with English Subtitles," uploaded by maurosegura 5 Jan. 2011; www.youtube.com/watch?v=qWiCJvXHZjY.
6. "IBM and Dubuque, Iowa Partner on Smarter City Initiative," 17 Sept. 2009; www-03.ibm.com/press/us/en/pressrelease/28420.wss.
7. A. Christensen, "In the Shadow of Copenhagen: Combining Wind Power and Electric Vehicles in Denmark to Address Climate Change," blog, 1 Dec. 2009; http://asmarterplanet.com/blog/2009/12/wind-power-and-electric-vehicles-in-denmark.html.
8. S. Kim and B. Powell, "Seoul: World's Most Wired Megacity Gets More So," *Time*, 24 Aug. 2009; www.time.com/time/magazine/article/0,9171,1916302-1,00.html.
9. H. Goldstein, "Malta's Smart Grid Solution," *IEEE Spectrum*, June 2010; http://spectrum.ieee.org/energy/environment/maltas-smart-grid-solution.
10. IBM Corp., "How It Works: Driving Change in Stockholm," 3 Apr. 2007; www.ibm.com/podcasts/howitworks/040207/images/HIW_04022007.pdf.
11. R. Wiesman, "The Municipal Reference Model: Smarter Government by Design," 2010; www.misa.on.ca/en/municipalinterface/municipalinterface.asp.

Milind Naphade *is the program director in the Smarter City Services Group at the IBM T.J. Watson Research Center in Hawthorne, New York, and leader of the Smarter Sustainable Dubuque living lab. He previously served as IBM Research's liaison on IBM's Global Innovation Outlook team, which examined areas of innovation focus for IBM. Naphade received a PhD in electrical engineering from the University of Illinois at Urbana-Champaign. He is a senior member of IEEE and a member of the IEEE Circuits and Systems Society's Multimedia Systems and Applications Technical Committee. Contact him at naphade@us.ibm.com.*

Guruduth Banavar *is vice president and chief technology officer of IBM's Global Public Sector, with a focus on smarter cities initiatives. He received a PhD in computer science from the University of Utah. Banavar is a member of the IBM Academy of Technology. Contact him at banavar@us.ibm.com.*

Colin Harrison *is a Distinguished Engineer with IBM's Corporate Strategy team, working to develop technical strategies for smarter cities. He received a PhD in electrical engineering from Imperial College London. Harrison is a senior member of IEEE and a member of the IBM Academy of Technology. Contact him at harrisco@us.ibm.com.*

Jurij Paraszczak *is director of IBM Research Industry Solutions and leader of IBM's smarter cities program as well as relationship manager to the IBM Venture Capital Group. He received a PhD in chemistry from the University of Sheffield, UK. Contact him at jurij@us.ibm.com.*

Robert Morris *is vice president of Services Research and Global Labs at IBM. He received a PhD in computer science from the University of California, Los Angeles. Morris is a fellow of IEEE and a member of the IBM Academy of Technology. Contact him at rjtm@us.ibm.com.*

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

COVER FEATURE

From Space to Stage: How Interactive Screens Will Change Urban Life

Kai Kuikkaniemi, Giulio Jacucci, Marko Turpeinen, and Eve Hoggan, *Helsinki Institute for Information Technology*

Jörg Müller, *Deutsche Telekom Laboratories*

Framed digital displays will soon give way to walls and facades that creatively motivate individual and group interaction. A stage serves as an apt metaphor to explore the ways in which these ubiquitous screens can transform passive viewing into an involved performance.

In many ways, cities have become utilitarian networks that people traverse on their way to somewhere more interesting. Individuals pass one another with rarely a second look, and billboards or other public displays are usually predictable, rectangular frames that garner only cursory attention.

City venues were not always so uninteresting. Before the digital age, the city was a highly interactive stage on which society played out a variety of creative experiences. European traders invested in elaborate facades to ensure that their houses stood out in an already impressive town square. Great town halls created displays of public life. Each person had a social role, and public venues were the place for all to act their parts, typically by being seen in clothing and accessories that reflected how they wanted others to view them.

With the digital age, much of this stage has become virtual, with the Internet serving as a powerful platform for creative and interactive play. Facebook, YouTube, Twitter, and Flickr have replaced the bustle of a city square. Like their ancestors, modern actors and actresses invest huge amounts of time to enhance their images and foster their community roles. The stage has moved from the town square to the living room and office, and public places have become shadows of their former color—spaces with very little opportunity for self-expression.

Screen technology is promising to return the interactive experience to urban spaces. Economies of scale are enabling any surface to become a display, offering a window to the Internet. No longer are displays limited to the conventional frame shape but can accommodate arbitrary shapes and sizes, walls, and even entire squares and streets.

As a result, surfaces and environments need no longer be passive. With the aid of large-scale automation and ubiquity, pervasive computing can transform urban areas into the most impressive stages the world has yet seen. Through this metamorphosis, city spaces become more lively, and inhabitants become more aware of their community.

THE STAGE IN EVERYDAY LIFE

Creating and shaping these impressive arenas requires first understanding what constitutes a stage. A play usu-

Published by the IEEE Computer Society

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

ally takes place inside some kind of magic circle, where performers enact roles using props and interactions to elicit a fictional space. The performance relies on a socially interpretative act, "translating real bodies, words, and movements into the objects of another, hypothetical world."[1]

In a sense, everyday social life is also a play in which people can take one of many roles.[2] In this interpretation, social space can be front stage (what the audience sees) and backstage (what the actors see), and an individual can be both audience and actor. For a salesperson, for example, the salesroom is the front stage, where he performs the specific role of showing customers the products. However, when he enters the storage room—the backstage—he plays a different role. In this play, the salesperson can switch roles and change behavior instantly, and even take on new roles, becoming a husband, father, or another persona once he leaves the job.

This role changing implies that the main power of stages and props stems from their ability to frame situations. If the salesperson changes into shorts, grabs a racket, and moves to a tennis court, his role changes from salesperson to tennis player. Similarly, if a man wearing a police uniform is shouting at another man in rags, the impression is that the other man might have committed some crime. If both men are wearing business attire, the impression might be quite different. Location can also influence impressions, as can spotlight and shadow. Thus, control over the stage implies the ability to control a situation's framing.

Framing is equally important in imaginative play. People readily use local props and media to experience fictional situations, such as those in pervasive games.[3] Performance behavior and playfulness are emerging in various strands of everyday life such as entertainment, advertising, and general social interaction. All these efforts aim to tell a story, evoke emotions, and sometimes even move the audience to some action.

TOWARD UBIQUITOUS SCREENS

Digital signage—the use of digital displays to replace analog signs—is a promising way to augment the everyday stage and make it more flexible and engaging. According to several industry reports, digital signage is already a billion-dollar business, and the markets are expected to grow more than 20 percent annually through 2013. Most of this growth is for signage in public or semipublic spaces.[4]

At present, most digital screens show videos or graphics according to some predetermined schedule, and users cannot manipulate them. However, such passive screens will soon give way to interactive displays such as touch screens, fueling the growth of the interactive public display market at a pace that is expected to exceed that of the market for public noninteractive screens.[4]

Most digital signage installations operate remotely: the screens connect through the Internet to a digital content management system, which operates scheduling and updates videos and graphics. Theoretically, the same system could deliver updates in real time, allowing installations to adapt to their environments.

Currently, LCD and LED technology drive the digital screen market. Alternative technologies, such as electronic ink, are on the horizon. Such novel technologies are expected to provide new design economies as displays become paper-thin, transparent, and flexible. In a decade or so, some researchers expect more advanced technologies such as displays based on carbon nanotubes to replace conventional technologies, pushing the market even more aggressively toward increasingly less expensive and larger installations.

If these projections and expectations even approach reality, a significant part of the wall space in popular public and semipublic spaces will be covered with digital screens. This proliferation will undoubtedly redefine how the space will look as well as how people will use it.

Digital signage—the use of digital displays to replace analog signs—is a promising way to augment the everyday stage and make it more flexible and engaging.

INTERACTING WITH DISPLAYS

So far, human-computer interaction research has involved primarily scenarios in which a single user interacts with the system through a personal display. Design paradigms and usability factors differ considerably when multiple users access the display simultaneously[5,6] or when people interact with a screen only in passing.[7] Relative to the single-user scenario, these design and usability shifts require the study of many more design aspects and considerations, such as those in Table 1.

Basically, public display interaction paradigms require either *touch interaction*, which can also be a tangible interface, or *distance interaction*, through, for example, a mobile device or camera installed with the display. Public displays most commonly provide information or advertising, but other use cases are certainly possible. For example, people can use the display interface to post photos or just as a way to have fun. Euclide's virtual puppets[8] and Magical Mirrors are evidence of how a display can offer entertainment value.

COVER FEATURE

Table 1. Current and future design challenges for interactive public displays.

Design issue	Considerations
Multiple users	Can social interaction occur through the display or around the display? Does the interaction design and visualization allow simultaneous use?
Implicit interaction	Does the display support implicit interaction based on behavioral cues from video and acoustic features? These are generally difficult because of noise and lighting conditions.
Adaptive screen	Does the display adapt to general conditions (such as weather, time, and events), social dynamics around the display, and person or role identification?
Interaction sessions and life cycle	Does the display support different interaction phases: passing by, viewing and reacting, subtle interaction, direct interaction, repeat interaction, follow-up, and memorabilia?
Screen form and shape	Does the screen's form suit its purpose? Many shapes and form factors are possible, from cylindrical to wall and facade displays, including multiple display orchestration.
Environmental factors	Have developers taken architecture and urban design, lighting, sound, safety, and connectivity into account?
Privacy	Does the system protect the data created in interaction, to avoid exposing identity by managing private and public interfaces?

Interaction perspectives

Researchers view human-display interaction from both the user and system perspectives. From the user perspective, interaction occurs in phases: passing by, viewing and reacting, subtle interaction, direct interaction, multiple interactions, and follow-up actions.[9] From the system perspective, the display has different behavior and thus requires varied interaction techniques, depending on the user-interaction phase.

In the first three phases—passing by, viewing and reacting, and subtle interaction—camera-based techniques can facilitate interaction. With camera-based detection, designers can exploit many natural interfacing techniques, such as pointing, gazing, facial expression, and posture. Most important, such detection enables the screen to recognize that a person is in front of it. Speech and sound recognition can also provide presence information, although the interfacing techniques are more limited.

Camera-based detection might include identification techniques such as face recognition. Mobile phone mechanisms, such as Bluetooth identification, radio frequency identification, and optical markers can also identify

EUCLIDE'S VIRTUAL PUPPETS

Visitors to the Cittá della Scienza science museum in Naples can interact with virtual puppets that entertain while explaining various exhibits, adapting dialog to a particular visitor, such as a child. The Euclide system uses a data glove, standard and MIDI keyboards, and a mouse to operate puppets at five museum stations. The puppeteer at a control system sees what the puppet sees through a camera and microphone at each station and can rapidly switch between stations and choose to interact with passersby.

Pupils interact with the virtual puppet display while a puppeteer controls the puppet using multimodal interfaces at the hidden station.

MAGICAL MIRRORS

In 2006, SAP and the Berlin University of Arts installed Magical Mirrors as a temporary art project in the street-level windows of a Berlin office building. The mirrors were four large public displays that showed a mirror image of the environment in front of them and used optical effects, such as aura, flexibility, and luminescence, to react to audience movement. Users could observe the gesture-based interaction of others or be an audience to their own gesturing image. The degree of interaction depended on age and role. For example, children were highly interactive, playing freely, while policemen tended to avoid interaction.

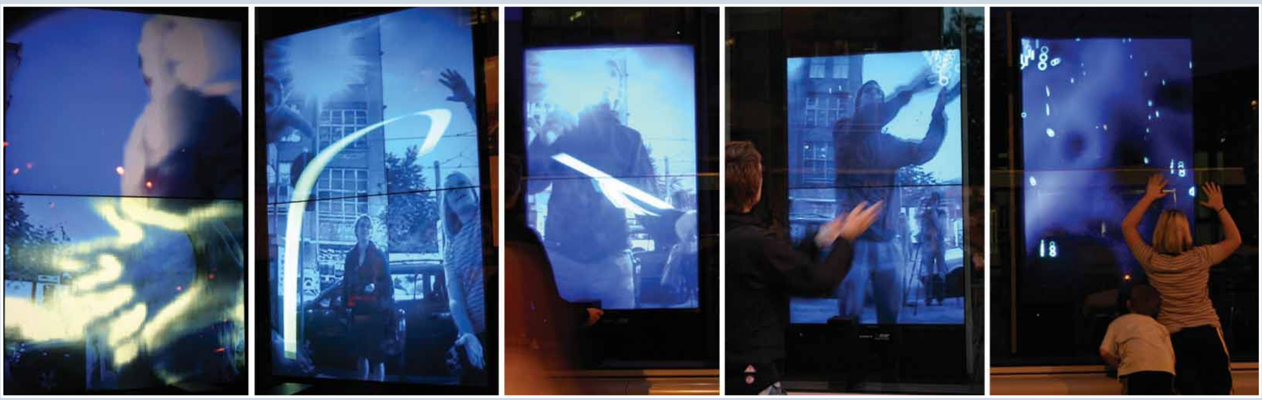

From left to right: hand motion creates an aura; a flexible band (next two photos); and luminescent numbers (last two photos).

an individual. Identification personalizes the interface, making it a powerful medium for optimizing usability, tailoring a marketing message, and enabling more complex social interaction and communication through a display. For example, a screen might recognize a first-time user and provide feedback that a frequent user wouldn't need. This usability optimization ensures that the screen adapts to each user's background. The first-timer isn't lost and frustrated, and the experienced user doesn't have to wade through unnecessary components and endure redundancy.

There is a delicate balance between personalization and privacy threats, however. Adaptive and sentient displays that can constantly store interaction data pose privacy issues. Anonymizing methods attempt to address these issues. Another approach is to use interaction data only locally and then delete it. Regardless of the method, the displays must put people in control of their screen interaction and any revelation of personal data and provide a way to transparently convey this control.

Environmental variables

An alternative to user identification is to base adaptivity on environmental variables, such as weather, time, traffic, people flow, news, or Web activity. This begs an interesting research question: How do you know when a flow of people is an environmental variable versus a crowd of interacting individuals? Similar to implicit interaction and adaptivity, cameras, sounds, or mobile devices can aid in implementing crowd or large-group interaction. For example, during breaks in Finnish hockey games, simple audio- or camera-based crowd interaction games on large stadium displays engage the spectators and enhance their shared experience (www.uplause.com/what_is_it). In crowd interaction, the public screen cannot provide individual feedback for single users, which significantly limits the possible screen designs. On the other hand, multiple users create a significant amount of social content, so even a simple design can quickly become lively and complex when many people are interacting.[10]

Shape and size

Many people assume that public displays must be framed, which is not surprising since current manufacturing technology limits production to flat, rectangular displays with a physical frame. Like a painting, these displays hang on the wall, except that the viewer has a window to the virtual world. And, similar to a painting, a framed display clearly differentiates the virtual world inside the frame from the real world outside and defines an optimal viewing position some distance from the frame's center.

New display technologies will change all this. Any surface, regardless of its shape, will be a candidate for serving as a frameless display. Freed from its restrictions, screen content will become part of the real world instead of being always virtual.

Between these two extremes are semiframed displays of various shapes. Cylindrical displays, for example, have a frame at the top and bottom, but not at the left and right, which gives them extreme flexibility in the horizontal view and removes the need for a single dominant viewing po-

DIGITAL ADVERTISING COLUMN

The digital advertising column is a cylindrical display that can detect user movements and react accordingly: for example, by drawing flowers in relation to the user's movements. The round shape invites users to stay in motion more than a flat screen. Furthermore, a flat screen has a single preferred standing position in the center front. Possession of this position identifies a single main user or performer. In contrast, all the positions around a cylindrical display are equal, making it easier for others to join the performance.

A passerby discovers the column and explores it from several positions. Unlike a flat display, all vantage points are equally important. Because it isn't necessary to occupy the center-front position to get the display's full value, more people can join in the viewing.

sition. The digital advertising column is an example of a semiframed display.

WHAT DOES A MEDIA STAGE OFFER?

Clearly, public-display interaction has multiple design dimensions. Interaction modalities and social dynamics, adaptation and personalization, screen layout, and multiple-screen orchestration all influence design. General space variables, such as lighting and sound, safety, and indoor-outdoor considerations are also critical design factors. In a sense, the design of a public digital display is akin to staging a theater play or an art installation, requiring careful attention to visual details. Also, gaining the community's trust and engagement is a delicate matter that often involves public relations work. And deploying a physical public-display installation requires special attention to practicalities from avoiding vandalism to moderating user creativity.

The tradeoff for all this extra work is the amazing potential of a media stage: it can create a social place, increase an event's visibility, foster a collective awareness that can change behavior, and even serve as an outlet for spontaneous creative performances.

Social place

Ubiquitous displays offer a way to transform an urban space into a sociable place. A place is a space with meaning: spaces are merely constructed areas, while places include the practices and cultural understanding of the people in the space.[11] Positioning interactive screens that encourage participation in an urban landscape can make places into spaces that serve as a stage for social interaction.

Configurability is an environment's capacity to become adaptable to different uses and provide varied experiences.[12] Developers can use ubiquitous media technologies to facilitate configurability, creating places that open people's eyes to the unique features of their living space.

Event visibility

Cities are filled with special events, but many go unnoticed. In addition to promoting transformative events such as the Olympic Games, public displays can make even small events more visible, giving the impression of a more lively and social urban life. As people discover that old familiar places have something new to offer and explore, the city can regain its capacity to surprise.

Public displays offer an alternative to receiving event news through a broadcast e-mail or webpage. Instead, users learn about the event through a location familiar to them. Traditional posters serve a similar purpose, but with an interactive digital screen, advertisers can change the message more rapidly and target it more accurately to the intended audience. Ultimately, a digital screen can extend the actual event—for example, with live streaming—or event organizers can design specific remote interfaces for those outside the event to participate in the event and interact as a spectator crowd. An event might even be distributed, using public displays to connect remote locations.

Indeed, the extended event stage is already in practice. All large concerts feature a screen that increases the artist's visibility, and some of the larger events stream a live broad-

CITYWALL, A MEDIA EXPLORATION

CityWall, an experiment we conducted in one of Helsinki's central pedestrian areas, turns a store window into an interactive display. The area is between bus and railway stations and major shopping centers, thus connecting two key locations. It consists of several small shops and cafes and often hosts temporary attractions that attract random groups.

Repurposing an existing architectural element—the store window to an interactive display—is very different from introducing an architectural element into the space as a new construct. CityWall effectively transforms an architectural object into an interaction of space and events. In the window, it is, of course, highly visible, yet it also changes how people use that space. For example, passersby initially grouped around CityWall's shop window to seek shelter from the rain, but then started using the space in front of the display to engage in playful and social media explorations.

Clockwise from left: passersby play with images on the CityWall timeline; the display from a street view; and two users share comments on pictures of Helsinki.

cast to another location. During the 2010 World Cup finals in Netherlands and Spain, massive crowds gathered to participate in an event that took place on the other side of the world. They could have watched the same stream in their homes, but chose instead to join the crowd experience.

At present, such event visibility enhancements involve only massive events and temporary infrastructures, but as permanent public-display infrastructure proliferates, similar arrangements will become feasible for lesser events. The technical setup is not complicated, making it suitable for short events as well. Connecting remote locations to a place of central importance is one basic use case: the action hotspot is the main stage, but public displays can partially reproduce that stage's physical and social presence for people at another location.

Collective awareness and group behavior

People constantly adapt their behavior according to what they see others do and what feedback they receive for their actions. A major goal in this social process is to maintain role consistency—to ensure that actions don't contradict the desired roles, but strengthen them. In the context of this goal, a media stage can support behavioral change in three significant ways. The first is by letting users be the audience, observing how others (their role models) behave so they can change their behavior accordingly.

A media stage can also support behavioral change by providing immediate and continuous feedback about the users' actions. Nuage Vert, or Green Cloud (http://hehe.org2.free.fr/?language=en), is a city-scale display that illuminates the vapor emissions of Helsinki's Salmisaari power plant to show local residents their current electricity consumption level. A laser ray traces the cloud during the night, turning it into a citywide neon sign.

Such persistent feedback on collective behavior can be an efficient way to help communities achieve some common goals, such as cutting energy consumption or riding bicycles instead of driving cars. Feedback is essential to learning how to change behavior, and it is relevant to a range of applications, from safety and security to antisocial behavior prevention. Using public displays as a feedback channel in a city space can transform the way its inhabitants approach responsible living and use the city's resources.

Finally, the media stage can encourage participatory civic discourse and critique,[13] particularly in light of the increasing reliance on social media platforms to express opinions. For example, by providing a forum for posting pictures about Helsinki, our CityWall experiment effectively uses a media stage as content for public debate. The stage can also serve its conventional purpose, which is to allow creative expression. Spontaneous playful interaction can nurture self-esteem and embolden individuals to be more proactive in other walks of life.

Almost any surface might soon become a digital display, which will transform the appearance of cities, bring the Internet out of homes and offices, and return the center of social life to the public places where it originated. Turning billboards into media stages can benefit both organizations, which have a new influence channel, and individuals, who can present themselves in desired roles.

A challenging open issue is how cities, individuals, and organizations will share influence on the media stage. Many cities have forced organizations to adopt a shared brand appearance, where the organization must adapt its logo to the city's style. It is likely that cities will grant the right to deploy display technology only if they can influence how deployment affects public appearance. Traditionally, cities prescribe only the number of billboards and their location. Digital displays allow much more fine-grained negotiation on what appears on the screen, when, and how. Public and private organizations are likely to negotiate influence, such as screen time and mixed content, but individuals must also have a say.

Regardless of the perplexities surrounding any sharing arrangement, public-display installations will fuel the next wave of social change. Broadcasting media made people more globally aware. Social media has made them more socially aware. Public display-based ubiquitous media will make them more community aware.

References

1. C. Counsell and L. Wolf, *Performance Analysis: An Introductory Coursebook*, Routledge, 2001.
2. E. Goffmann, *The Presentation of Self in Everyday Life*, Anchor Publishing, 1958.
3. M. Montola, J. Stenroos, and A. Waern, *Theory and Design of Pervasive Games—Experiences on the Boundary Between Life and Play*, Morgan Kaufmann, 2009.
4. S. Khatri, "Emerging Display 2009 Report—Touchscreen Interfaces in Signage & Professional Applications," *iSuppli*, 2009; www.isuppli.com/Display-Materials-and-Systems/Pages/touch-screen-interfaces-continue-to-drive-growth-in-signage-and-professional-applications.aspx.
5. P. Peltonen et al., "It's Mine, Don't Touch!: Interactions at a Large Multitouch Display in a City Center," *Proc. 26th Int'l Conf. Human Factors in Computing Systems* (CHI 08), ACM Press, 2008, pp. 1285-1294.
6. G. Jacucci et al., "Worlds of Information: Designing for Engagement at a Public Multitouch Display," *Proc. 28th Int'l Conf. Human Factors in Computing Systems* (CHI 10), ACM Press, 2010, pp. 2267-2276.
7. J. Müller et al., "Requirements and Design Space for Interactive Public Displays," *Proc. Int'l Conf. Multimedia* (MM 10), ACM Press, 2010, pp. 1285-1294.
8. C. Coutrix et al., "Engaging Spectators with Multimodal Digital Puppetry," *Proc. 6th Nordic Conf. Human-Computer Interaction: Extending Boundaries* (NordiCHI 10), ACM Press, 2010, pp. 138-147.
9. D. Michelis and J. Müller, "The Audience Funnel: Observations of Gesture-Based Interaction with Multiple Large Displays in a City Center," *Int'l J. Human-Computer Interaction*, vol. 27, no. 6, 2011, pp. 562-579.
10. D. Vogel and R. Balakrishnan, "Interactive Public Ambient Displays: Transitioning from Implicit to Explicit, Public to Personal, Interaction with Multiple Users," *Proc. 17th Ann. Symp. User Interface Software and Technology* (UIST 04), ACM Press, 2004, pp. 137-146.
11. P. Dourish, "Re-space-ing Place: 'Place' and 'Space' Ten Years On," *Proc. 20th Conf. Computer-Supported Cooperative Work* (CSCW 06), ACM Press, 2006, pp. 299-308.
12. T. Binder et al., "Supporting Configurability in a Mixed Media Environment for Design Students," *Personal and Ubiquitous Computing*, vol. 8, no. 5, 2004, pp. 310-325.
13. M. Ananny, C. Strohecker, and K. Biddick, "Shifting Scales on Common Ground: Developing Personal Expressions and Public Opinions," *Int'l J. Continuing Eng. Education and Life-Long Learning*, vol. 14, no. 6, 2004, pp. 484-505.

Kai Kuikkaniemi *is a research project manager at the Helsinki Institute for Information Technology (HIIT) and a doctoral student in the School of Arts at Aalto University. His research interests include the design of playful public displays, social interaction around public displays, sensor-enhanced pervasive mobile gaming, and the use of future Web technologies in the media industry. Kuikkaniemi received an MS in industrial engineering and management from Helsinki University of Technology. He is a member of the ACM. Contact him at kai.kuikkaniemi@hiit.fi.*

Giulio Jacucci *is a professor in the Computer Science Department at the University of Helsinki, director of HIIT's Network Society Research Program, coordinator of the BeAware project (www.energyawareness.eu) at Aalto University, and cofounder and chairman of the board of MultiTouch Ltd. (www.multitouch.fi). His research interests include human-computer interaction in ubiquitous computing with a focus on multimodality and surface and adaptive computing. Jacucci received a PhD in information processing science from the University of Oulu. He is a member of the ACM. Contact him at giulio.jacucci@helsinki.fi.*

Marko Turpeinen *leads HIIT's Digital Content Communities research group and is the director of the Helsinki node in the European Institute for Innovation and Technology ICT Laboratories. His research interests include customized media, computer-mediated communities, the future of gaming, and the role of media technology in promoting sustainable development. Turpeinen received a PhD in com-*

puter science from the Helsinki University of Technology. He is a member of IEEE and the ACM. Contact him at marko.turpeinen@hiit.fi.

***Eve Hoggan** is a postdoctoral researcher at the University of Helsinki and HIIT. Her research interests include human-computer interaction, multimodal interaction, and novel interaction techniques. Hoggan received a PhD in cross-modal audio and tactile interaction from the University of Glasgow, UK. She is a member of IEEE and the ACM. Contact her at eve.hoggan@helsinki.fi.*

***Jörg Müller** is a senior researcher at Deutsche Telekom Laboratories, Berlin. His research interests include interaction and context adaptivity with digital signage and advertising in pervasive computing, particularly the use of sensors to enable interaction and audience measurement. Müller received a PhD in context-adaptive digital signage in transitional spaces from the University of Münster, Germany. He is a member of the ACM. Contact him at hans-joerg.mueller@telekom.de.*

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

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

COVER FEATURE

Experiences inside the Ubiquitous Oulu Smart City

Felipe Gil-Castineira, Enrique Costa-Montenegro, Francisco J. Gonzalez-Castano, and Cristina Lopez-Bravo, *University of Vigo, Spain*

Timo Ojala, *University of Oulu, Finland*

Raja Bose, *Nokia Research Center, Palo Alto*

The UrBan Interactions (UBI) research program, coordinated by the University of Oulu, has created a middleware layer on top of the panOULU wireless network and opened it up to ubiquitous-computing researchers, offering opportunities to enhance and facilitate communication between citizens and the government.

Ubiquitous-computing research traditionally has concentrated on smart homes and smart offices, where the physical dimensions of the smart space are constrained within a building or a small geographical area. However, with the deployment of urban wireless infrastructures in cities all over the world, the vision for developing smart cities is finally taking shape.

Moreover, unlike other ubiquitous-computing research topics that are primarily motivated and driven by academic or industrial groups, the research on smart cities has brought together a diverse group of participants, including governments, urban planners, sociologists, and traditional ubiquitous-computing researchers. Several governments have already undertaken ambitious programs to build smart cities by augmenting existing city infrastructure with embedded sensing as well as communication and interaction technologies.

One such endeavor, the First International Open Ubiquitous City Challenge, offered international ubiquitous-computing researchers the opportunity to implement and evaluate innovative applications and services for the Open Ubiquitous Oulu testbed, a real smart-city environment located in Oulu, Finland.

SMART CITIES AROUND THE WORLD

The vision for smart cities originated as a natural evolution of research in smart homes and other smaller-scale smart spaces. The Universal City project[1] envisioned smart cities as a means of extending ubiquitous-computing services beyond the traditional smart home environment, forming the basis of interaction within an entire community on a much larger geographical scale. In that environment, smaller smart spaces such as smart homes interoperate with each other and interact with other external infrastructure components such as embedded outdoor sensors and mobile terminals while providing a unified and personalized interface for their residents.

Unlike other projects in which smart cities are being built from the ground up, the Amsterdam Smart City project's goal (www.amsterdamsmartcity.com) has been to integrate smart technologies into an existing historical city to make it more eco-friendly and energy-efficient. A key achievement of this project has been the installation of green energy infrastructure such as solar panels, household wind turbines, and electric-car charging stations in an attempt to reduce the city's carbon footprint. These in turn become a part of a smart grid system that includes the virtual power plant concept, in which households can sell excess energy generated from domestic solar panels, wind turbines, and biomass plants to the city to generate income.

 Published by the IEEE Computer Society 0018-9162/11/$26.00

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

This not only provides a medium for citizens to have a say in citywide initiatives to encourage environmentally friendly behavior but also gives them an incentive to actively participate in such efforts.

The SmartCity initiative (www.smartcity.ae) is attempting to create a global network in which each smart city is a self-sustained business township with advanced ICT infrastructure. Currently, the SmartCity network consists of three smart cities, one each in Dubai, Malta, and India. The goal is to build each smart city as a cluster of knowledge and talent pools while sharing intelligence, talent, and opportunities with other smart cities in the network.

South Korea has been at the forefront of the development of the *ubiquitous city* concept. A u-city is a smart city in which knowledge and services are available to residents through the use of ubiquitous computing, with sensing and communication resources embedded in urban elements such as residences, building infrastructure, and open spaces.[2] These cities use the information and communications technology infrastructure to provide interlinked services organized within specific domain areas. For example, u-life services provide related functionalities for residents such as home automation and monitoring. U-business provides work-related services such as videoconferencing, information management, and document sharing. U-government provides services that empower citizens by making them aware of traffic and health hazards and enables them to participate in the governance process anytime, anywhere.

A u-city is designed from the ground up to be completely user-centric, providing more personalized yet nonintrusive services for its residents. The first Korean u-city, Hwaseong Dongtan, has been partially completed and is currently operational. Another u-city, New Songdo City, is being constructed on a 1,500-acre man-made island off the coast of Incheon, South Korea. This city is centered on u-life, a proposed lifestyle that utilizes smart cards with RFID technology to provide personalized services and user interfaces for residents anywhere within the city.

Figure 1. A UBI hotspot in the Ubiquitous Oulu Smart City.

OPEN UBIQUITOUS OULU

Open Ubiquitous Oulu enables ubiquitous-computing research in authentic urban settings with real users on a broad scale and in a sufficient time span. Such studies are important because real-world ubiquitous-computing systems are culturally situated and can't be reliably assessed with lab studies detached from the real-world context. By deploying a system for a sufficiently long time, researchers can establish technical and cultural readiness and identify the critical mass of users needed to determine whether the system can be deemed either successful or unsuccessful.[3]

The UrBan Interactions (UBI) program is developing UBI-hotspots, interactive public displays embedded with computing resources such as two cameras, an NFC/RFID reader, panOULU wireless access points, and high-speed Internet access to provide rich interaction between physical, virtual, and social spaces.[4,5] In addition to the UBI-hotspots, the testbed comprises a wide variety of pervasive computing resources deployed across the city, including a panOULU LAN/Bluetooth/wireless sensor network and middleware resources.

The UBI-hotspots alternate between two states. In the passive broadcast state, the entire screen is allocated for the UBI-channel, a digital signage service. In the interactive state, the screen is split between the UBI-channel and the UBI-portal, which can embed any Web service found on the Internet. As Figure 1 shows, the UBI-portal provides access to a wide range of interactive services such as directories, games, a street gallery of new media art exhibitions, sending of UBI-postcards, or uploading of personal photos and videos. Currently, all interactive events such as face detection and launching of specific services by users are logged for reporting and research purposes.

The panOULU wireless LAN is a citywide Wi-Fi network comprising approximately 1,270 IEEE 802.11 access points.[6] The access points provide open and free wireless Internet access to all public users without any limitations. However, comprehensive real-time network traces currently are archived for reporting and research purposes. For example, in February 2011, 26,013 devices used the

Figure 2. A panOULU wireless LAN access point (WLAN AP), wireless sensor network edge router (WSN ER), and Bluetooth access point (BT AP).

network, for a total of 658,742 sessions, culminating in usage of 19.3 million online minutes.

The panOULU Bluetooth network includes 30 access points scattered across the city center. Eighteen of these access points are installed on traffic lights, and they use the panOULU wireless LAN for Internet access. Twelve additional access points are placed inside UBI-hotspots. All access points sniff nearby Bluetooth radios, and the real-time traces are suitably anonymized and used for modeling pedestrian and vehicular flows and networks. The 12 access points inside the UBI-hotspots are also used for pushing multimedia content to mobile devices via Bluetooth.[7]

The panOULU WSN is an IP-based wireless sensor network comprising 13 edge routers (ERs) located across the city. The ERs are equipped with an IEEE 802.15.4 radio on the 2.4-GHz and 868-MHz bands, and the 6LoWPAN protocol stack, which offers low-power wireless connectivity. Twelve ERs are installed inside the WLAN mesh access points, as shown in Figure 2, and one ER is also placed inside a UBI-hotspot. An ER has a line-of-sight in the 500-m range with 1 mW transmission power. The MediaTeam Oulu research group is currently developing the panOULU WSN infrastructure to automatically meter energy consumption in homes[8] and for environmental monitoring using low-power sensors.

These heterogeneous computing resources constitute a large distributed system that's organized with a middleware layer. It provides various resources for supporting technology experiments, developing ubiquitous-computing applications, and managing and monitoring the applications and the testbed.

This kind of a large-scale testbed deployment presents many challenges, including establishing financial and technological viability and sustainability. After the initial capital investment, mostly from public sources, sufficient funding must be obtained to cover operational expenses and renewal. For example, a portion of the capacity of the UBI-hotspots is sold commercially to generate revenue, which in turn conflicts with research use. A practical challenge is the operational execution of maintenance, which is expensive and time-consuming and is an area in which research organizations typically aren't efficient.

Another important challenge is the measurement of success, whether it's assessing the socioeconomic impact of the testbed or the merit of any application or service deployed within it. These assessments are difficult due to the shortage of comprehensive data over a sufficiently long period and the lack of a universally accepted methodology for evaluating real-world deployments. Finally, any deployment in a city center is subject to daily scrutiny by the general public and media, which can become ill-tempered and impatient at times.

To strengthen the testbed's long-term prospects, the UBI program is integrating it into OULLabs (Oulu Urban Living Labs; www.oullabs.fi), a regional living lab initiative. This project brings together a range of testbed and human resources to facilitate various activities such as user-driven open innovation in developing and testing new technologies and applications in authentic urban settings with real users.

THE UBI CHALLENGE

Because they realized that few researchers have access to such a versatile u-city testbed for development, deployment, testing, and learning purposes, the MediaTeam Oulu research group wanted to make this testbed openly available to as many researchers as possible. The First International Open Ubiquitous City Challenge (www.ubioulu.fi/en/UBI-challenge), or UBI Challenge, is being organized for this purpose in collaboration with several leading international ubiquitous-computing experts. This competition challenges the global R&D community to design, implement, deploy, and evaluate novel applications and services in real-world settings in the city of Oulu.

The motivation for the challenge is to stimulate global research collaboration on urban informatics in a concrete manner and provide the international research community with an opportunity to transfer ideas from labs into real-world urban environments where they can make an impact. The goal is also to support the development of metrics for evaluating urban computing infrastructure and applications in real-world settings.

Participation was encouraged by advertising that up to five proposals would be selected as finalists for deployment in Oulu, with each group receiving grants of up to 10,000 euros and having the opportunity to present their research at the International Conference on Mobile and Ubiquitous Multimedia (MUM 2011). Out of the numerous proposals received, the international jury invited four to compete in the finals. All four finalist teams will arrive in Oulu in mid-2011 to finalize the implementation and deployment of their services. Thereafter, the jury's Oulu-based members will meet and assess each of the services in situ according to various performance and usability metrics.

MOTIVATING CITY RESIDENTS TO EXERCISE

A sedentary lifestyle increases the likelihood of developing obesity, diabetes, and cardiovascular disease. Estimates indicate that the healthcare cost of these chronic diseases is rapidly approaching US$1 trillion and that a strong association exists between increased physical inactivity and the emergence of chronic diseases in 21st-century industrialized societies.[9] Private companies have found that they can reduce healthcare costs by offering employees cash to quit smoking or by serving healthy food in the cafeteria.

Some governments are also following a similar model. For example, the government of Nova Scotia, Canada, offers economic incentives to parents who register children in sports or recreation activities that offer health benefits. Recently, the US Department of Health and Human Services announced that it's devoting economic resources to encourage citizens to participate in health improvement programs (www.letsmove.gov).

Smart cities can become the cornerstone for promoting healthier lifestyles by using their urban wireless infrastructures to implement mechanisms that directly encourage residents to play sports and exercise or serve as tools to measure actual involvement in exercise programs.

During recent years, several tools have been developed for tracking performance while exercising, especially while running, walking, hiking, or biking. Tools like Nike+, RunKeeper, Endomondo, Strands, and Nokia Sports Tracker help users keep records of information pertaining to their sports activities, such as duration, distance, pace, speed, elevation, calories burned, and the course traveled on a map. This information is uploaded to a Web-based cloud service where participants can share statistics and make detailed comparisons. Most of these applications seek to keep different kinds of users motivated by offering personal or collective challenges—for example, men versus women, aiming to complete 10 miles a week or run a total of 500 miles, or competing against the user's own previous workouts or against other people's workouts along a route. However, these applications usually require carrying a GPS-enabled smartphone, which discourages many potential users due to the cost of the device. Moreover, a user who wants to update his or her position and statistics on the Web-based service in real time must maintain a persistent data connection to the cloud.

One way to keep participants motivated to exercise is to create social networks in which they can share their performance with their friends and also compete against them. Developers can enhance existing tools with related features supported by smart city infrastructures. For example, instead of using a GPS device, it's possible to use the urban wireless infrastructure to keep track of a user's performance. Moreover, smart cities are designed to be adapted to their specific environment and the target user base. Exploiting this characteristic can stimulate people to engage in sports activities because their friends and acquaintances who use the system are located in the same area. Participants also can be encouraged to compete against others in a similar physical condition or to follow a regimen that's suitable for their level of fitness.

The goal is also to support the development of metrics for evaluating urban computing infrastructure and applications in real-world settings.

RunWithUs

We're designing RunWithUs, a service that will be deployed in the Ubiquitous Oulu Smart City, with the goal of motivating citizens to practice jogging, letting them

- select a route according to length, difficulty, groups or friends currently following it, number of runners in a similar physical condition who usually follow this route, number of runners in the area, and weather conditions or pollen levels;
- compare differences in performance between runners following the same route, competing with other users' performances, or competing against one's own previous workouts (phantom runner); and
- promote sports practice and social life by establishing rankings: for example, biggest group of the month, most regular group of the month, and so on.

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

Figure 3 shows the RunWithUs system architecture, which includes Wi-Fi-enabled phones or tags that are connected to the wireless infrastructure and tracked by the UBI middleware, weather or other environmental sensors, and user interfaces such as a UBI-display or a Web browser on a PC or a smartphone.

Because RunWithUs was selected as one of the four finalists in this year's UBI Challenge, we'll have the opportunity to implement the system in a real-world smart city, thereby identifying difficulties and user concerns related to the location methodology's precision as well as issues related to privacy.

Location methodology

Implementing these features requires estimating each user's location. For this purpose, we're using the panOulu wireless infrastructure to locate personal Wi-Fi devices such as mobile phones and MP3 players. We're also taking advantage of existing Ubiquitous Oulu and social networking APIs to offer several user interfaces (top layer in Figure 4) to configure the system, including

- city displays that show a map with information about groups of runners, statistics, or promotional videos and animations, and users can join a group at any moment, even when they're already running;
- a website providing the same information as in the urban displays, including a special version for mobile devices; and
- a Facebook interface to share results, comments, and so forth with friends or other runners.

In a typical use case, a new user who wants to access RunWithUs must register using a UBI-hotspot. The user provides a login, a password, and a MAC address, usually from the mobile phone's Wi-Fi, an MP3 player, or a Wi-Fi badge. This MAC address tracks the user throughout the city. Once registered, the user can log into the system through any of the available interfaces.

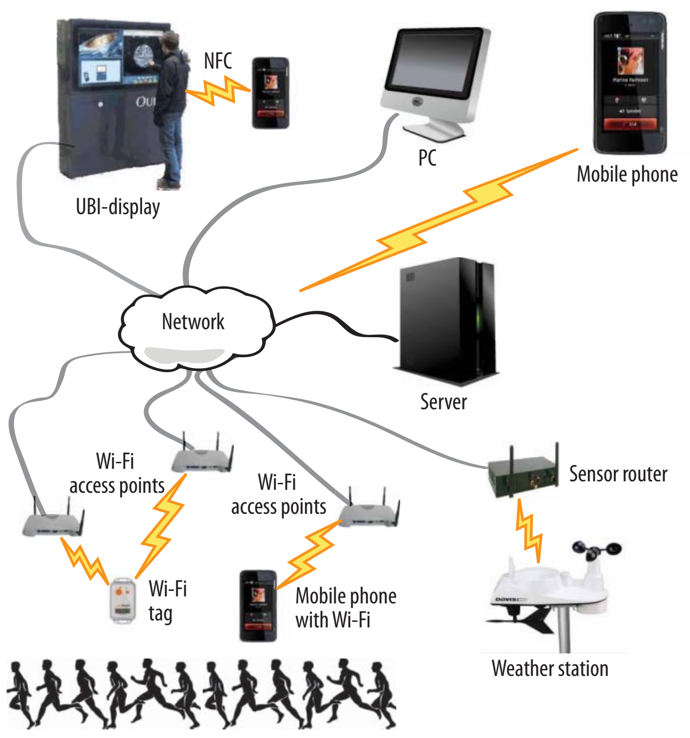

Figure 3. RunWithUs system architecture.

Privacy

Naturally, users might be concerned that such a system could track them while they aren't running or when they don't desire to. Initially, we implemented a procedure to activate and deactivate the tracking service, for example, pressing a "Start Running" button in the RunWithUs user portal. This action allows the infrastructure to check the MAC address location and start tracking until the user presses "Stop Running." These actions are needed to differentiate when users are running versus when they're engaged in other nonexercise-related activities. When users are engaged in other activities, although the infrastructure could still detect their MAC address, the system won't update their location information. But since it's highly likely that users will forget to stop the tracking service, we implemented a timer to stop it by requiring users to select the amount of time they're going to run when they start the service.

Currently, the system uses panOULU Wi-Fi access points to track runners. Following an increasing trend, a recent report indicates that approximately 15 percent of the total number of mobile phones in Finland has Wi-Fi support.[10] A user who doesn't own a Wi-Fi-enabled phone or doesn't want to carry it while running could use a smaller Wi-Fi tag such as AeroScout, Ekahau, or RedPine. By using Wi-Fi access points for tracking, a runner's location can only be roughly estimated based on an access point's coverage area, but the hypothesis is that this will be adequate for this use case. Nevertheless, the deployment in Oulu will reveal whether or not this supposition is correct. For example, because we expect poor Wi-Fi coverage in parks, we'll need to adapt RunWithUs to overcome this limitation. Bluetooth access points could provide a finer-grained location, but the low number of currently installed units will likely reduce their effectiveness.

Figure 5 depicts the various elements that comprise the RunWithUs service. The different user interfaces provide unregistered users information about the system such as routes, statistics, or public events. Unregistered users can also get some minimal and anonymous information about who's running right now or what the most popular routes are at a given time. There are two main reasons for making this information available. The first is to make the service attractive to potential future registered users; the second is to help tourists find a place to exercise comfortably and safely by finding appropriate routes suitable for their physical condition in safe areas.

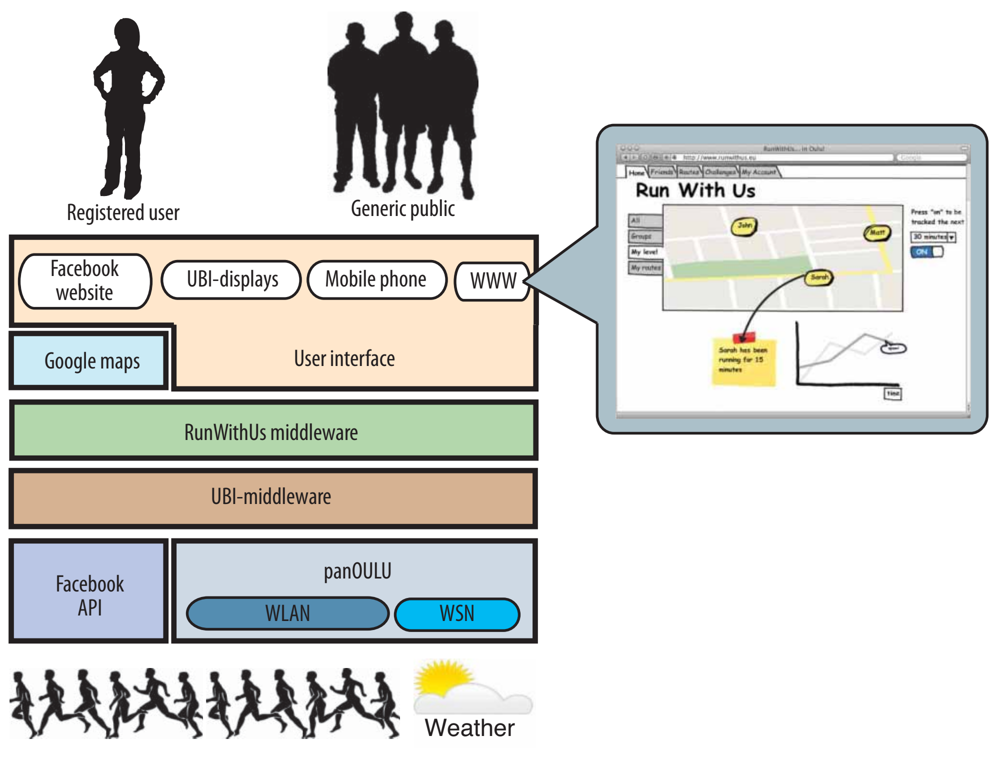

Figure 4. RunWithUs software stack.

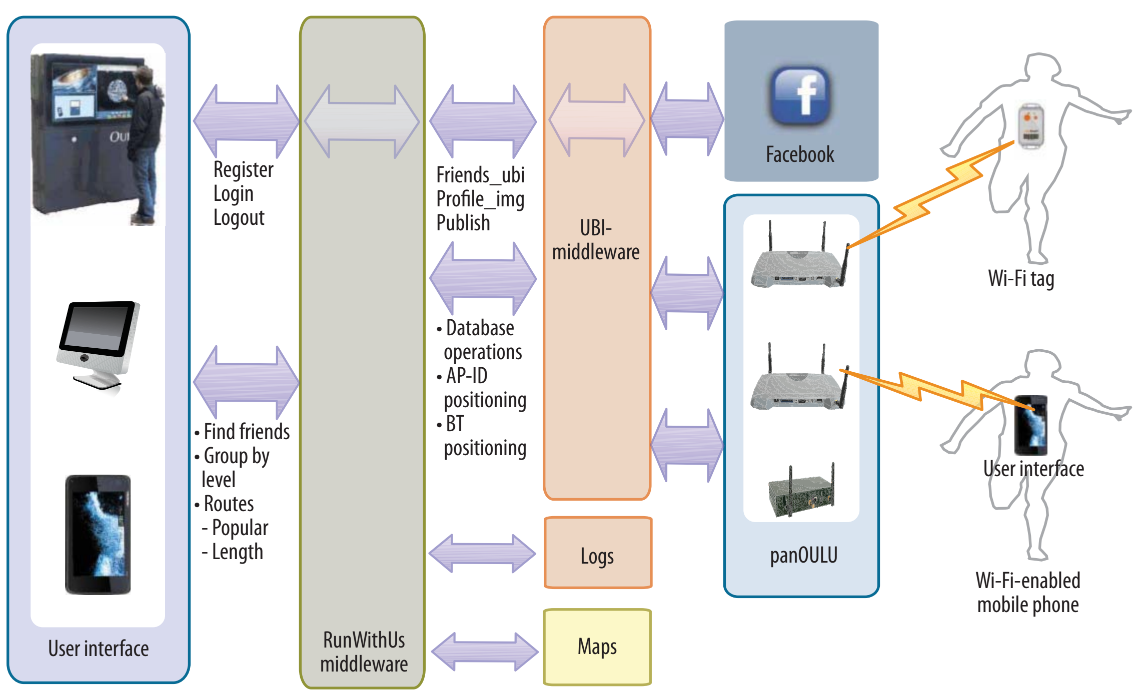

Figure 5. RunWithUs in operation.

In addition to the information provided to unregistered users, registered users can get information their friends want to share, such as their favorite routes or their performance. Registered users can also view who among their friends is running now, including their names, physical condition level, and corresponding location. Users also can determine who intends to go jogging at a particular time in case they want to join their friends.

The RunWithUs service's middleware layer tracks runners and uses that information to create routes and recommendations. This layer is also responsible for creating dynamic webpages for the different user interfaces. Routes and runner positions are overlaid on maps using the Google Maps API. For registered users, the system can filter results from their Facebook friends or publish statistics as text or images. Personal statistics such as total distance traveled, fastest kilometer, and so on can also be published according to user preferences using the Facebook API.

RunWithUs is still in its early stages, but we've already confronted several challenges. Implementing a multiscreen application designed to operate in devices with different sizes, resolution, and input interfaces isn't a trivial endeavor. Furthermore, applications developed for smart cities should be adapted to the environment in which they're going to be installed, making it necessary to perform onsite tests, especially when using APIs and interfaces interacting with citizens or mobile devices. But for RunWithUs, the main challenges are related to using Wi-Fi to track users. This technology can only offer precision at several hundred meters, and the fastest person in the world needs only roughly 10 seconds to cover 100 meters, making it necessary to find the optimal periodicity of location requests to offer maximum precision without overloading the servers. The low precision also makes it necessary to use postprocessing and additional information to determine whether users are following a predefined route, running together, or creating teams, or to store precise information about the track.

As the number of smart cities continues to increase, the number of related research initiatives is also increasing rapidly. However, it's still difficult for researchers to get access to a full-featured real-world testbed. Thanks to initiatives such as the UrBan Interactions research program, it's now possible to create and test smart-city applications and services designed to provide benefits to the community at large, discovering limitations that were not foreseen during the design phase and iteratively refining them until they become a ubiquitous and helpful part of people's daily lives.

References

1. T. Yamazaki, "Beyond the Smart Home," *Proc. Int'l Conf. Hybrid Information Technology* (ICHIT 06), vol. 2, IEEE Press, 2006, pp. 350-355.
2. S. Lee et al., "Towards Ubiquitous City: Concept, Planning, and Experiences in the Republic of Korea," *Knowledge-Based Urban Development: Planning and Applications in the Information Era*, T. Yigitcanlar et al., eds., IGI Global, 2008, pp. 148-170.
3. S. Greenberg and B. Buxton, "Usability Evaluation Considered Harmful (Some of the Time)," *Proc. Conf. Human Factors in Computing Systems* (CHI 08), ACM Press, 2008, pp. 111-120.
4. T. Ojala et al., "Open Urban Computing Testbed," *Proc. 6th Int'l ICST Conf. Testbeds and Research Infrastructures for the Development of Networks and Communities* (TridentCom 10), ICST, 2010, pp. 1277-1288.
5. T. Ojala et al., "UBI-Hotspot 1.0: Large-Scale Long-Term Deployment of Interactive Public Displays in a City Center, *Proc. 5th Int'l Conf. Internet and Web Applications and Services* (ICIW 10), IEEE Press, 2010, pp. 285-294.
6. T. Ojala et al., "Supporting Session and AP Mobility in a Large Multi-Provider Multi-Vendor Municipal WiFi Network," *Proc. 3rd Ann. Int'l ICST Conf. Access Networks* (AccessNets 08), ICST, 2008, pp. 89-101.
7. H. Kukka et al., "BlueInfo: Open Architecture for Deploying Web Services in WPAN Hotspots," *Proc. Int'l Conf. Web Services* (ICWS 09), IEEE Press, 2009, pp. 984-991.
8. T. Ojala et al., "UBI-AMI: Real-Time Metering of Energy Consumption at Homes Using Multi-Hop IP-Based Wireless Sensor Networks," *Proc. 6th Int'l Conf. Grid and Pervasive Computing* (GPC 11), to be published in 2011.
9. F.W. Booth et al., "Waging War on Modern Chronic Diseases: Primary Prevention through Exercise Biology," *J. Applied Physiology*, vol. 88, 2000, pp. 774-787.
10. A. Riikonen, "Mobile Handset Population in Finland 2005-2009," 2010; www.netlab.tkk.fi/tutkimus/momi/publications/Riikonen_2010_Mobile_Handset_Population_2005-2009.pdf.

Felipe Gil-Castineira *is an assistant professor in the Department of Telematic Engineering at the University of Vigo, Spain. His research interests include wireless communication technologies, embedded systems, and ubiquitous computing. Gil-Castineira received a PhD in telecommunication engineering from the University of Vigo. Contact him at xil@gti.uvigo.es.*

Enrique Costa-Montenegro *is an associate professor in the Department of Telematic Engineering at the University of Vigo. His research interests include wireless networks, car-to-car communication technologies, multiagent systems, and recommendation technologies. Costa-Montenegro received a PhD in telecommunication engineering from the University of Vigo. He is a member of IEEE. Contact him at kike@gti.uvigo.es.*

Francisco J. Gonzalez-Castano*, a full professor in the Department of Telematic Engineering at the University of Vigo, where he leads the Information Technologies Group, is also the R&D director at Gradiant, Spain. His research inter-*

ests include telecommunications and computer science, and he holds three Spanish patents, a European patent, and a US patent. Gonzalez-Castano received a PhD in telecommunication engineering from the University of Vigo. Contact him at javier@gti.uvigo.es.

***Cristina Lopez-Bravo** is an associate professor in the Department of Telematic Engineering at the University of Vigo. Her research interests include high-performance networking, with a focus on the design and evaluation of packet scheduling. Lopez-Bravo received a PhD in telecommunications engineering from the Polytechnic University of Cartagena, Spain. Contact her at clbravo@gti.uvigo.es.*

***Timo Ojala** is a professor of computer engineering and director of the MediaTeam Oulu research group in the Computer Science and Engineering Laboratory at the University of Oulu, Finland. His research interests include human-computer interaction, distributed systems, and ubiquitous computing. Ojala received a Dr.Tech in electrical and information engineering from the University of Oulu. He is a member of IEEE and the ACM. Contact him at timo.ojala@ee.oulu.fi.*

***Raja Bose** is a senior researcher at the Nokia Research Center, Palo Alto. His research interests focus on the theme of mobile interoperability, including creating new devices, systems, and services that provide innovative user experiences by leveraging the interoperability of mobile devices with other smart environments. Bose received a PhD in computer engineering from the University of Florida. Contact him at raja.bose@nokia.com.*

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

COVER FEATURE

Building an Integrated Service Management Platform for Ubiquitous Ecological Cities

Jungwoo Lee, *Yonsei University*
Songhoon Baik, *KT Corp.*
Choonhwa Lee, *Hanyang University*

As one of the frontrunners in the race to build smarter cities, South Korea is pushing the envelope by promoting the development of a standard architecture for a service management platform that integrates ubiquitous computing and green technologies.

Dramatic technological progress in recent years has brought unprecedented changes to every corner of our society and transformed daily life. In particular, wireless and mobile communications, radio frequency identification (RFID), and wireless sensor networks have paved the way for ubiquitous networking, and Internet-enabled devices are increasingly being used for accessing and processing information as well as communications.

Smart mobile devices along with high-speed, far-reaching access networks and sensors embedded in the environment provide the technical foundation for a ubiquitous city—or *u-city*—where objects and people are intimately connected. As advances in information and communications technology (ICT) open up opportunities for more effective and efficient urban management, innovations will be needed to provide new infrastructure services to cope with the changes.

FROM U-CITIES TO U-ECO CITIES

Current urban development trends emphasize the use of ICT to build smarter cities and, ultimately, a smarter planet.[1] Ubiquitous computing technology plays an increasingly important role in these efforts, enabling intelligent transportation systems (ITSs), geographic information systems (GISs), smart homes and workplaces, and environmental monitoring.[2-6]

A u-city is a smart city in which physical infrastructure instrumented with various sensors, such as power grids and oil pipelines, and mobile objects, such as humans and vehicles, are connected through ICT. In this dynamic and evolving ecosystem, everyone—from citizens to facilities managers to emergency responders to traffic control operators—can access a wide variety of advanced technologies and services, like those shown in Figure 1, using any device anywhere, anytime.

Countries around the globe have launched u-city projects, with South Korea at the forefront. In fact, the South

 Published by the IEEE Computer Society

Korean government is pushing the envelope by advocating ubiquitous ecological cities. A *u-eco city* combines core u-city technologies such as integrated city management/operations and citizen services with green technologies to increase convenience, safety, and quality of life while reducing carbon emissions—in short, a place where people, technology, and the environment coexist in harmony. The "South Korea's U-Eco City Initiative" sidebar describes the evolution of this national urban development effort.

The nation's first u-eco city, Hwaseong Dongtan (www.udongtan.or.kr/english/cyber/cyb_01_1.aspx), was completed in 2008. As Figure 2 shows, six additional u-eco cities are under construction, and 18 are at the design stage.

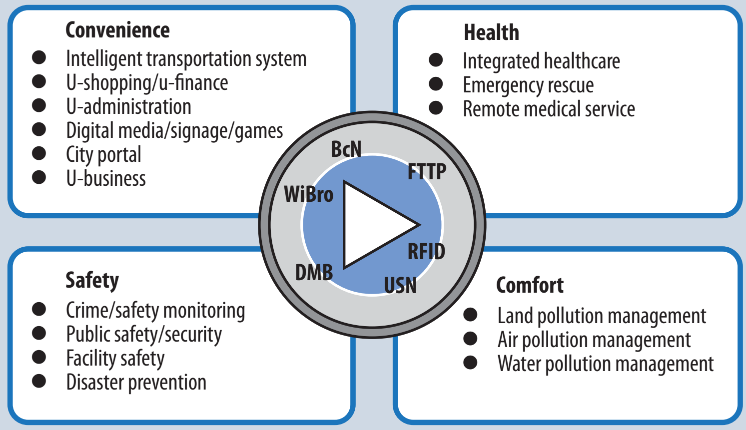

Figure 1. Technologies such as wireless broadband (WiBro), broadband convergence networks (BcNs), fiber to the premises (FTTP), radio frequency identification (RFID), ubiquitous sensor networks (USNs), and digital multimedia broadcasting (DMB) enable a wide range of ubiquitous city (u-city) services.

As u-eco cities began to emerge around the country, policymakers expressed concern about incompatibility among various operation and management platforms developed by South Korean ICT companies, which could lead to duplicated investments at the national level. Consequently, U-Eco City R&D Center researchers are developing a standard architecture for an integrated service management platform designed to enhance the efficiency of u-eco city management and communications.

INTEGRATED OPERATIONS CENTER

A comprehensive survey of South Korean u-eco cities under development[7] reveals a total of 228 potential services, which can be grouped into 11 categories:

- administration,
- transportation,
- welfare (health and medical services),
- environmental management,
- crime and disaster prevention and response,
- facility management,
- education,
- culture and tourism,
- logistics,
- labor and employment, and
- other services.

These services vary widely in scope and functionality, and the lack of coordination among u-eco city development efforts has hindered progress.[8] No reference model is available, let alone a national standard.

The first step in addressing this problem will be the creation of a u-eco city integrated operations center (IOC-UC) that will gather, process, and store information on all services and make appropriate recommendations to service operators.

Among the 228 services identified in the survey, each IOC-UC will initially focus on 31 commonly found across u-eco cities in South Korea. Figure 3 shows these 31 services organized into five urban management domains, with supportive operational models for services, control, infrastructure, and data management. By developing individual services within this framework, a u-eco city can provide a coherent view of its services and administrative organization to residents.

The IOC-UC also provides support for day-to-day monitoring, managing, and provisioning of a u-eco city's infrastructure as well as its services.

INTEGRATED SERVICE MANAGEMENT PLATFORM

The various u-eco city service platforms developed in South Korea have been designed and implemented as isolated stovepipe systems in which sensors and devices are connected to servers dedicated to a particular application domain, and networks are separated from one another. There are no defined application programming interfaces—adding a new service requires ad hoc, hard-wired customization by a specific vendor. In addition, it is not possible to synergistically use data collected from different services, and there is no support for system integration of neighboring cities or communications and control handovers. All of these characteristics significantly limit the benefits of u-eco city services.

As Figure 4 shows, a u-eco city integrated service management platform (ISMP-UC) can overcome the limitations of closed stovepipe systems and enable

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

COVER FEATURE

SOUTH KOREA'S U-ECO CITY INITIATIVE

Policymakers and technological leaders in South Korea conceived the vision of a u-eco city around 2003, when the deployment of ICT infrastructure such as code division multiple access (CDMA) and fiber to the premises (FTTP) was sufficient to accommodate every individual and business in the country.

In 2004, after soliciting public input, the Ministry of Information and Communication and the Ministry of Land, Transportation, and Maritime Affairs began coordinating efforts on a national-scale u-eco city development program. Two years later, the U-Eco City R&D Center (www.ueco.or.kr) was established under the auspices of the Korea Land Corporation (later merged with the Korea Housing Corporation) to fund and oversee the program, which was formally launched in August 2007.

As Figure A shows, South Korea's u-eco city development program consists of one overall group responsible for strategic planning, marketing, research, and testbed construction and two core groups charged with developing infrastructure and services. The projected total budget for all u-eco city projects from 2007 through 2013 is US$130 million.

South Korea's private sector has also recognized the potential opportunities of u-eco cities and has created a forum to promote and discuss the initiative (www.ucta.or.kr/en/ucity/background.php). Participating members include major ICT players such as KT, Samsung SDS, and LG CNS.

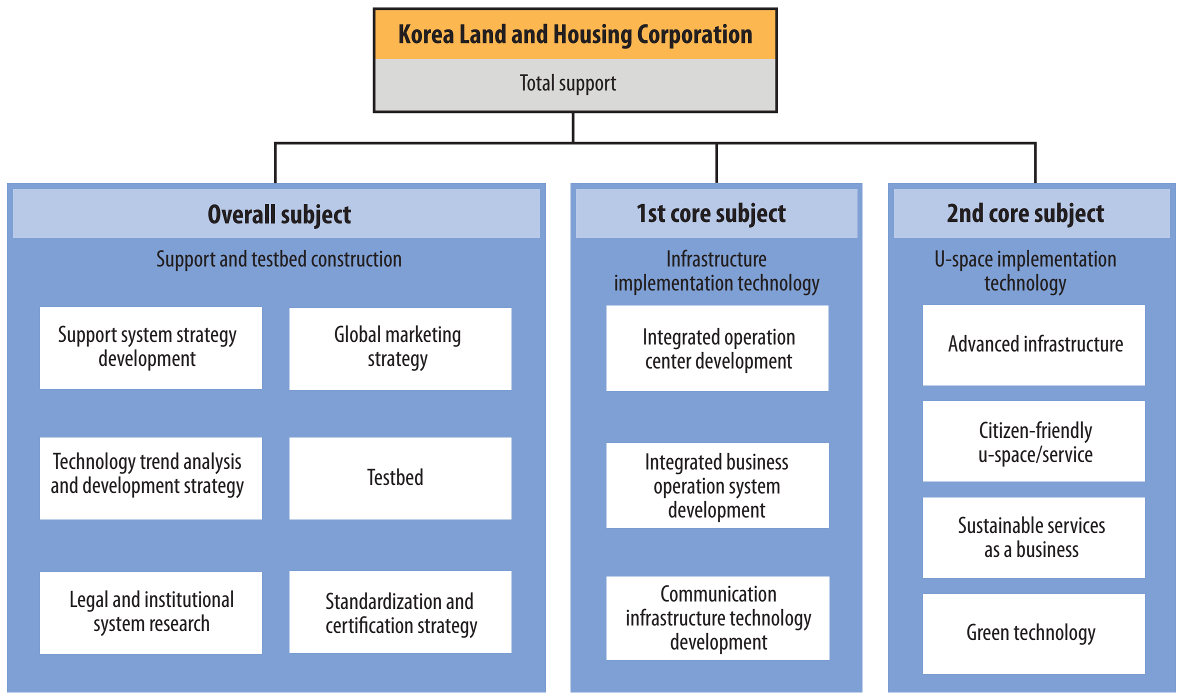

Figure A. South Korea's u-eco city development program consists of one overall group responsible for strategic planning, marketing, research, and testbed construction and two core groups charged with developing infrastructure and services.

synergistic service collaboration. The ISMP-UC makes operational decisions based on input from sensors and other networks and sends control commands to system components or external entities. For example, if sensors on the road detect an automobile accident, the system can notify a nearby police station and hospital. The system could also automatically identify involved vehicles and drivers from data collected by streetlight-mounted cameras. The ISMP-UC can also create new services by combining information from different sources.

To handle the diversity and dynamism inherent in a u-eco city, the service architecture should be as flexible and extensible as possible.[9] The ISMP-UC has three basic layers. The bottom layer consists of various types of sensors, actuators, and other devices distributed about the city. On the top layer is a range of u-eco city services. Between these layers lies the middleware that collects and processes data and contextual information. The middleware's service-oriented architecture enables services to be developed independently and invoked through standardized Web services interfaces.

As Figure 5 shows, the ISMP-UC middleware includes a gateway service, a ubiquitous information service, a mobility manager, an operations management service, and an integrated database.

Gateway service

The gateway service is a collection of interface adapters for connecting middleware components and entities that are either internal or external to the system. By keeping open its interface to information systems and various user devices and sensors, the GS can cope with the ever-changing needs of an evolving urban space and its residents by accommodating new devices and services without the need for any change to existing services or other parts of the architecture.

The GS has three components. The *internal link service* relays information and control signals between the ubiquitous information service/mobility manager modules and the IOC-UC and its 31 common services. The *external link service* transfers information and control commands between these mod-

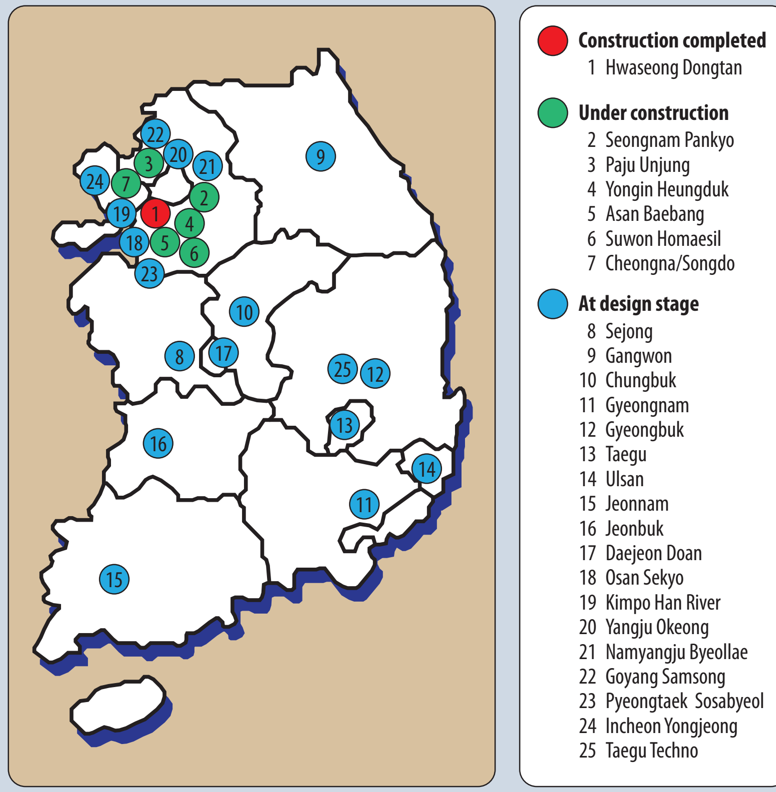

Figure 2. In South Korea, one u-eco city has been completed, six are under construction, and 18 are at the design stage.

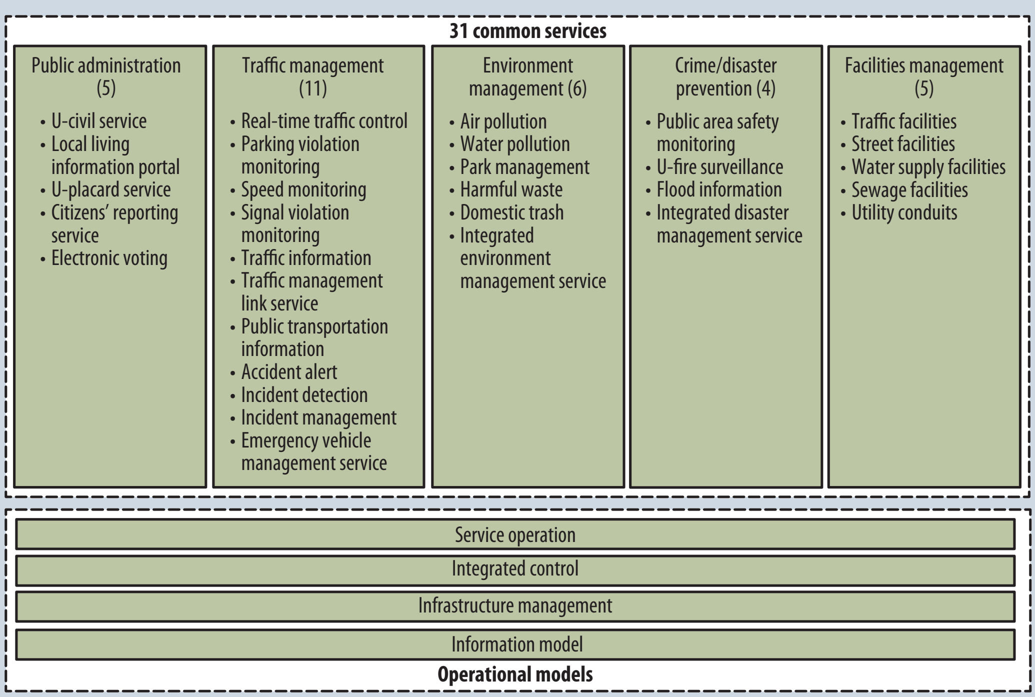

Figure 3. Each u-eco city's integrated operations center (IOC-UC) will initially focus on 31 common services organized into five urban management domains, with supportive operational models for services, control, infrastructure, and data management.

COVER FEATURE

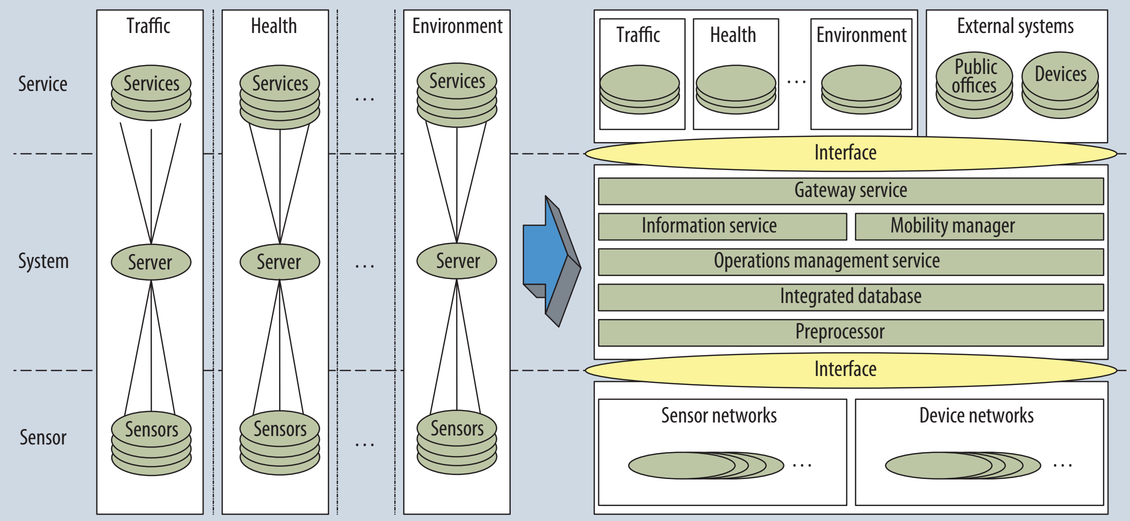

Figure 4. A u-eco city integrated service management platform (ISMP-UC) can overcome the limitations of closed stovepipe systems and enable synergistic u-eco city service collaboration.

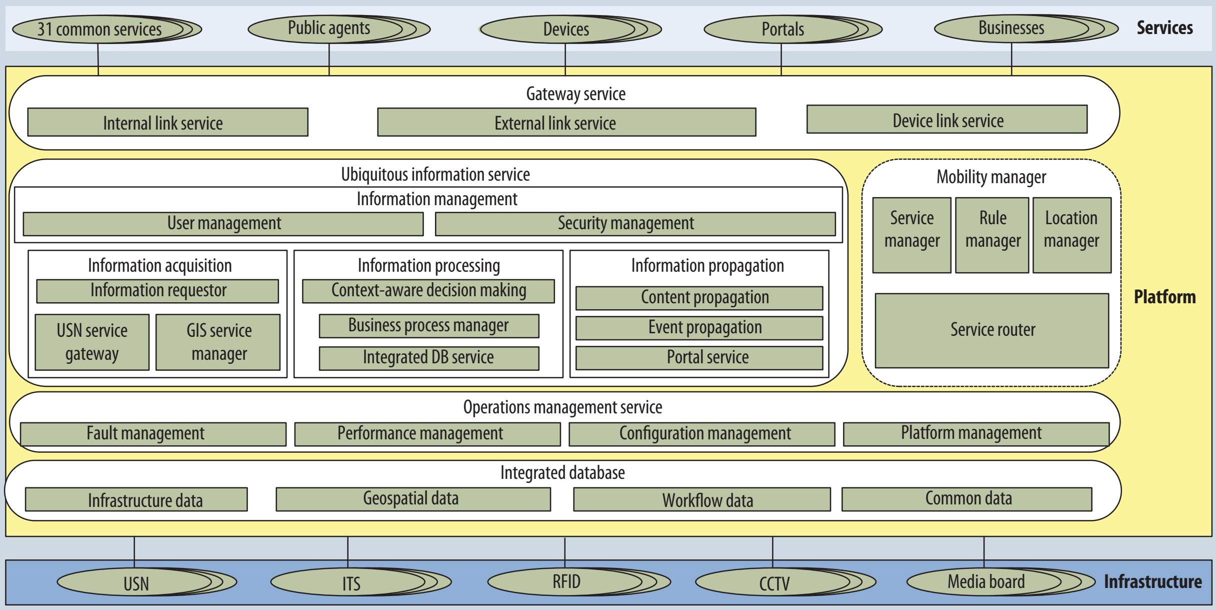

Figure 5. The ISMP-UC middleware includes a gateway service, a ubiquitous information service, a mobility manager, an operations management service, and an integrated database.

ules and systems run by outside agencies, institutions, and companies. The *device link service* interacts with mobile devices, service networks, and sensor nodes with computation capability, providing communication paths to the information service and mobility manager at a lower level.

Ubiquitous information service

At the heart of the ISMP-UC architecture is the ubiquitous information service, which includes four basic building blocks for u-eco city services.

Information acquisition. Upon receiving a request for sensor or geographical data from external entities or internal components, the information acquisition service passes it through the link services to either the *USN gateway* or the *GIS service manager*. The integrated database temporarily saves the responses from these networks for further processing.

Information processing. This service is responsible for processing all data in the system. With the help of the *context-aware decision-making service*, a business process manager orchestrates the interactions of component services and procedural logic within and across services. In addition, an *integrated database service* manages various data, whether inside or outside the platform, collected from different sources and created by context-aware analyses and other business processes.

Information propagation. This service has three components: the *content propagation* and *event propagation* modules deliver processed data and event data, respectively, to higher-level objects, while the *portal service* is responsible for the operator-friendly display of processed results.

Information management. This service contains two modules. The *user management service* provides information about users and groups requested by other components and services, and controls access privileges to middleware components and application services. The *security management service* prevents data from being transferred to unauthorized entities and supports encryption functions.

Mobility manager

The proliferation of smartphones and other Internet-enabled devices as well as digital signage makes it possible to deliver rich content to users anytime, anywhere. The mobility manager (MM) provides for the delivery of information such as traffic, weather, and news to mobile devices. The MM can also convey data via digital billboards, kiosks, and variable message signs.

The MM effectively separates the service and device layers so that u-eco city applications can be developed independently of particular hardware devices and platforms. Applications in the service layer can use Web APIs to display information on a target device. The MM can identify a group of devices in a particular geographical area and broadcast data only to that group. This allows for efficient information dissemination to, for example, an accident site or disaster area.

The MM combines user profiles, device capabilities, and adaptation processes such as transcoding and resizing to deliver context-aware information. It supports both always-connected devices and intermittently connected devices or mobile terminals for emergency situations.

Operations management service

The operation management service (OMS) monitors faults in system components and devices, performance of system components, service composition and device networking, and proper functioning of the platform itself. The OMS has four primary modules. The *fault management* module analyzes failures or errors in services and sensors connected to the system and recommends solutions when available. The *performance management* module tracks system performance and, in real time, periodically analyzes and manages various system capabilities. The *configuration management* module configures system components and application services and manages network resources. The *platform management* module provides administrative support for the IOC-UC including faults, performance, backups, and security.

Integrated database

The integrated database serves as the repository for data from all system component modules and application services. It contains infrastructure and workflow data in addition to common business data about users and organizations. The database also stores geographic and spatial information about roads and urban facilities such as sewage treatment plants.

IMPLEMENTATION AND ROLLOUT

A prototype implementation of the ISMP architecture is expected to be completed by the first half of 2011. Performance fine-tuning and further customization will begin soon after that. In addition, the U-Eco City R&D Center is developing and testing sample u-eco city services.

A prototype implementation of the ISMP architecture is expected to be completed by the first half of 2011.

Researchers are creating a full-scale, realistic testbed in the new Cheongna u-eco city in the Incheon Free Economic Zone (IFEZ) on the west coast of South Korea. Established in 2004, the IFEZ is an eco-friendly, mixed-use environment situated midway between major East Asian cities like Beijing, Shanghai, Tokyo, and Hong Kong. Aside from its locational advantages for business, the IFEZ will incorporate ubiquitous computing technologies for all aspects of urban life.

Figure 6 shows an artist's rendering of Cheongna, where the 77-story World Trade Center with neighboring business and commercial facilities, along with a high-tech industrial park, are already under construction. The master plan for Cheongna includes an IOC-UC based on the ISMP-UC architecture. Thirteen u-eco city services are being considered for first-phase deployment.

Sejong, a u-eco city that will be built south of Seoul beginning in 2013, will serve as another testbed for the ISMP-UC as well as interoperability with Cheongna's IOC-UC. Other candidate cities for ISMP-UC deployment include Busan, the second largest city on the southern coast. Unlike Cheongna and Sejong, u-eco city services in Busan will be part of a downtown rejuvenation project.

COVER FEATURE

Figure 6. Artist's rendering of Cheongna, Incheon Free Economic Zone, which will contain a full-scale, realistic testbed for u-eco city services.

CHALLENGES

The U-Eco City R&D Center must address several key challenges to fulfill the vision of South Korea's ambitious u-eco city initiative.

Service interoperability

Interoperability is a major concern when developing ubiquitous services that evolve independently yet rely on one another to accomplish a larger purpose. For example, two closed-circuit TV (CCTV) broadcasting system manufacturers might use different time stamps. Service developers and device manufacturers must agree on message-passing protocols, naming conventions, service invocation requests, and results. There must also be agreement on higher-level service representations such as tickets for traffic violations, as local authorities could have various management formats and structures.

Beyond syntactic data interoperability is the more challenging issue of service-level interoperability. As semantic interactions among different government agencies and manufacturers are not yet standardized, adapting u-eco city service modules will require significant effort. The solution calls not only for technical expertise but also for social and institutional consensus building.

Service developer concerns

Before the ISMP-UC, which makes service interfaces public, third-party city service developers had developed their own proprietary solutions for sewage management, water supply monitoring, and so on in isolation. For example, a CCTV network monitoring traffic in a business district was separate from a CCTV network monitoring crime in a residential district. Services on the two networks could not communicate with each other.

These legacy systems must be modified and incorporated into the ISMP-UC platform. The ISMP-UC provides common functions and a Web services API that third-party developers can use to compose their application services. These Web services will be able to communicate with each other so that, for example, a crime-prevention CCTV network can interact with a traffic-monitoring CCTV network, providing appropriate authorities with better information to track criminals across the city.

However, third-party service developers have raised concerns that these changes would reduce their role and thus impact their profits. They have also resisted providing technical details about their proprietary systems that are essential for ISMP-UC implementation. The U-Eco City R&D Center has spent considerable time persuading developers that the ISMP-UC will benefit their business in the long run by making it easier to develop new services using existing service modules. The proliferation of new services would in turn increase economies of scale for third-party developers.

Institutional resistance

Conceived as a central monitoring and control point, the ISMP-UC must interact with other systems already in place. For example, an automobile accident monitoring and report service would need to access a vehicle registration database, resident registration information, insurance company databases, and possibly hospital patient management systems. However, sharing databases across different government agencies and companies requires complex approval processes. As u-eco city services spread over various domains, involved institutions must reach agreement on information sharing.

The U-Eco City R&D Center is developing a secure brokering system for stakeholders, but this is taking much longer than initially expected. Different institutions have their own requirements established in accordance with government regulations and business logic. Because resolution of the institutional boundary problem could take a long time, the Cheongna testbed will use interface emulators rather than actual databases and information collected from real devices.

Supply-push versus demand-pull

U-eco city development efforts in South Korea have largely been driven from the top down rather than by consumer needs. As ISMP-UC development progressed, however, U-Eco City R&D Center researchers began to realize that user acceptance and other socioeconomic issues are a bigger hurdle to the rollout of u-eco city services than technological obstacles.

It is not yet clear who will be responsible for the cost of u-eco city service deployment and maintenance. Residents will be reluctant to foot the bill. Several studies are exploring the business opportunities of ubiquitous services, but the market is not yet mature enough to attract numerous entrepreneurs. A new business model and new incentives are needed to sustain the initiative.

U-eco city efforts have been conceived primarily by technology experts who foresee the needs of future ICT-based city management. During this period of interpretive flexibility, however, residents and other social groups must be involved in exploring alternative designs.[10] Efforts are needed to build public awareness and understanding of the benefits of u-eco city services, nurture a market for the services, and eventually establish them as essential components of urban social infrastructure.

Over the past several years, u-cities have begun to emerge around the globe. As one of the front-runners in the race to build smarter cities, South Korea is taking the extra step of integrating cutting-edge ubiquitous computing and green technologies into the development of new cities. The ISMP-UC middleware architecture for u-eco city operation centers is the first step toward realizing this vision. Once successfully completed and tested for performance and interoperability, it will be a solid reference model for u-eco city services and could serve as a standard reference model for worldwide u-city developments.

Acknowledgments

We are grateful to Youngjoon Kim, SK C&C, Ltd., for his input on service mobility support, and Sumi Helal, University of Florida, for his insightful comments on an early draft of this article. We also thank Mac Son, Hyejung Lee, and Jaesoo Jang for their research assistance and diagram preparation. This research was supported by a grant (07 High Tech A01) from the High Tech Urban Development Program, funded by the South Korean government's Ministry of Land, Transportation, and Maritime Affairs.

References

1. J.M. Eger, "Smart Growth, Smart Cities, and the Crisis at the Pump: A Worldwide Phenomenon," *I-WAYS*, Jan. 2009, pp. 47-53.
2. W.J. Mitchell, *E-topia: "Urban Life, Jim—But Not as We Know It,"* MIT Press, 1999.
3. A. Macias-Diaz, "The u-City Index: Integrated Planplementation of Future Ubiquitous Cities," *Proc. 15th Int'l Conf. Urban Planning and Regional Development in the Information Society* (REAL CORP 10), Competence Center of Urban and Regional Planning, 2010; www.corp.at/archive/CORP2010_178.pdf.
4. N. Komninos, *Intelligent Cities and Globalisation of Innovation Networks*, Routledge, 2008.
5. P. van den Besselaar and S. Koizumi, eds., *Digital Cities III. Information Technologies for Social Capital: Cross-cultural Perspectives*, LNCS 3081, Springer, 2005.
6. G.S. Yovanof and G.N. Hazapis, "An Architectural Framework and Enabling Wireless Technologies for Digital Cities & Intelligent Urban Environments," *Wireless Personal Comm.*, May 2009, pp. 445-463.
7. S.-H. Lee et al., "Ubiquitous Urban Infrastructure: Infrastructure Planning and Development in Korea," *Innovation: Management, Policy & Practice*, Dec. 2008, pp. 282-292.
8. D.-H. Shin, "Ubiquitous City: Urban Technologies, Urban Infrastructure and Urban Informatics," *J. Information Science*, Oct. 2009, pp. 515-526.
9. K.-W. Nam and J.-S. Park, "Software Platform Architecture for Ubiquitous City Management," *Proc. 5th Int'l Conf. Digital Society* (ICDS 11), Int'l Academy, Research, and Industry Assoc., 2011, pp. 178-181.
10. T.J. Pinch and W.E. Bijker, "The Social Construction of Facts and Artefacts: or How the Sociology of Science and the Sociology of Technology Might Benefit Each Other," *Social Studies of Science*, Aug. 1984, pp. 399-441.

***Jungwoo Lee**, a professor in the Graduate School of Information and director of the Information Technology Strategy and Policy Research Institute (ITSPRI), Yonsei University, Seoul, South Korea, currently leads the U-Eco City R&D Center's ISMP-UC development project. Lee received a PhD in computer information systems from Georgia State University. He is a member of the IEEE Computer Society and the ACM. Contact him at jlee@yonsei.ac.kr.*

***Songhoon Baik** is a research engineer with KT Corp., South Korea. He led a project team that designed and developed KT's u-eco city service platform, Ubi-Cahn, and is currently participating in the U-Eco City R&D Center's ISMP-UC development project. Baik's research interests include geospatial information systems and construction-ICT convergence technologies. He received a PhD in civil engineering from Chungnam National University, South Korea. Contact him at baiksh01@paran.com.*

***Choonhwa Lee** is an assistant professor in the Division of Computer Science and Engineering at Hanyang University, Seoul, South Korea. His research interests include mobile networking and computing, peer-to-peer protocols, and middleware and services computing technology. He received a PhD in computer science from the University of Florida. Contact him at lee@hanyang.ac.kr.*

RESEARCH FEATURE

Ring Generator: An Ultimate Linear Feedback Shift Register

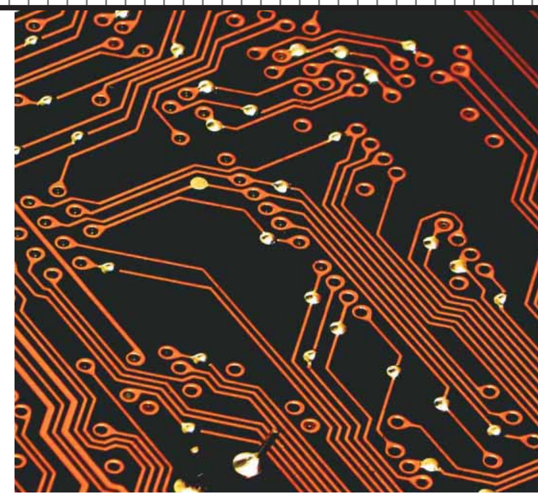

Nilanjan Mukherjee and Janusz Rajski, *Mentor Graphics*
Grzegorz Mrugalski and Artur Pogiel, *Mentor Graphics Polska*
Jerzy Tyszer, *Poznan University of Technology*

Because they are universal devices that provide unprecedented speed of operation and a layout-friendly structure, ring generators offer a superior option to using traditional solutions for handling pseudorandom and deterministic binary sequences.

In 2001, the electronic design automation industry witnessed the introduction of embedded deterministic test (EDT), a compression technology that addresses the problem of managing escalating test data volume and test application time.[1] Compression allows applying more test patterns to integrated circuits in less time and with fewer tester pins. EDT adds simple logic to a circuit under test to transform compressed test inputs into desired internal test bits, both deterministic and random. This on-chip decompression circuitry uses a *ring generator* architecture.

Although originally designed for nanometer test, ring generators are universal devices with a potentially countless number of applications in such diverse areas as computer engineering, mobile telephony, cryptography, and discrete event counting. The distinctive engineering advantage of ring generators lies in their regular structure, extremely short propagation paths, and small internal fan-outs.[2]

In communications and digital broadcasting, ring generators randomize the transmitted bitstreams, which prevents short repeating sequences from forming spectral lines that can complicate symbol tracking at the receiver or interfere with other transmissions.

The GPS can use ring generators to rapidly produce a sequence indicating high-precision relative time offsets. Cellular telephony and Bluetooth systems can use ring generators as shrinking or alternating step generators in stream ciphers. As extremely fast sources of pseudorandom sequences, ring generators can be deployed in a direct-sequence spread-spectrum radio or in various programmable sound generators.

Finally, HDTV, digital audio broadcasting systems, gigabit Ethernet scramblers, and satellite communication systems might also adapt ring generators due to their high performance and generic design flexibility.

LINEAR FEEDBACK SHIFT REGISTERS

Linear feedback shift registers have been the most popular devices to date for generating and handling pseudorandom sequences. Figure 1a shows an n-bit LFSR consisting of n memory elements with connected XOR gates.

The modulo 2 sum of the selected stages goes back to the leftmost stage of the register, forming a Fibonacci LFSR. This LFSR is represented by the characteristic polynomial $x^n + h_{n-1}x^{n-1} + \ldots + h_1x + h_0$, where h_k refers to the register's kth flip-flop. If $h_k = 1$, a feedback structure encompasses this flip-flop. A characteristic polynomial that causes an n-bit LFSR to go through all possible $2^n - 1$ nonzero states before entering the seed state is a *primitive* polynomial. The corresponding LFSR, or *maximum-length*

Published by the IEEE Computer Society

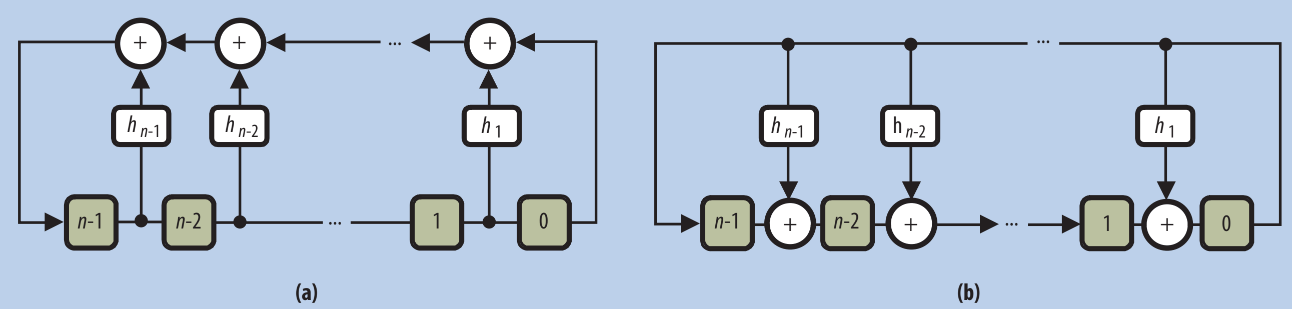

Figure 1. Two types of linear feedback shift registers: (a) Fibonacci LFSR and (b) Galois LFSR. LFSRs are commonly used to generate and handle pseudorandom sequences.

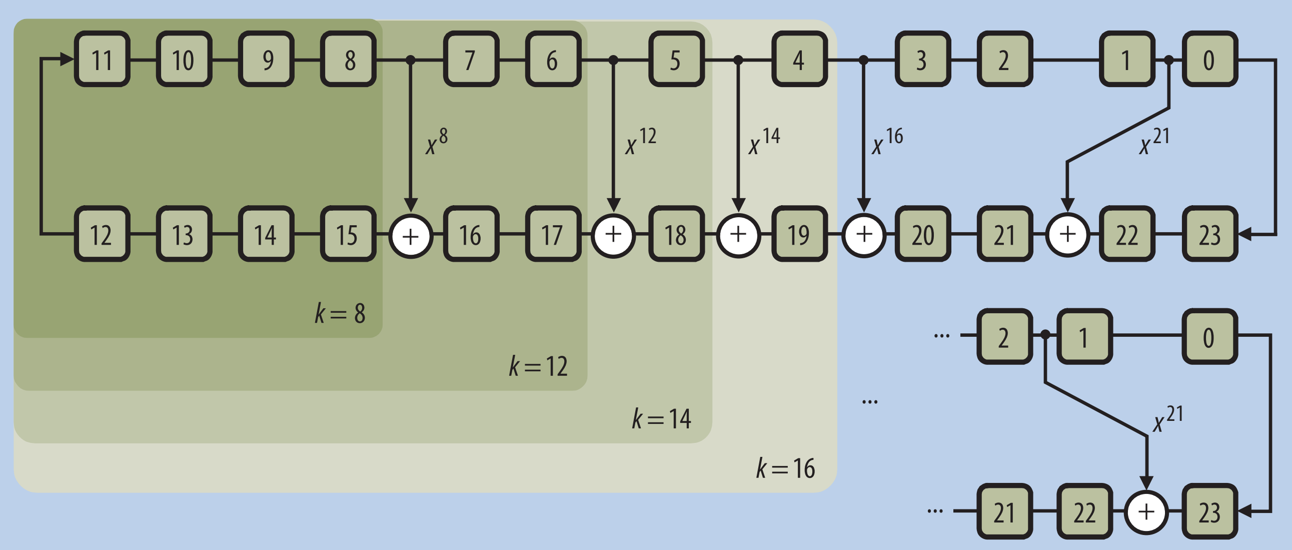

Figure 2. A ring generator implementing the polynomial $x^{24} + x^{21} + x^{16} + x^{14} + x^{12} + x^8 + 1$. Feedback lines form a regular rack frame. As long as the lines do not cross each other, we can rearrange feedback connections to accommodate a variety of additional requirements.

LFSR, produces the output sequence—a maximum-length sequence or *m*-sequence.[3]

Figure 1b shows a Galois LFSR, an alternative LFSR implementation with interspersed XOR gates. In this LFSR, the output of the rightmost stage feeds the stages indicated by the characteristic polynomial.

The depth of the linear logic in feedback paths limits the operating speed of Fibonacci LFSRs. Similarly, Galois LFSRs can cause severe frequency degradation due to long feedback lines, and they can take up a considerable area as they involve a large fan-out on the rightmost stage output. These limitations are pronounced for polynomials with numerous terms. Finally, the irregularity of their interconnections precludes LFSRs from becoming modular designs.

Several designs attempt to improve the performance of conventional LFSRs. In particular, hybrid LFSRs reduce the number of XOR gates by using both forms of the feedback logic in the same register;[4] windmill machines elevate a state transition rate but need additional registers;[5] and decimation allows summing up several *m*-sequences produced by independent devices with a multiphase clock generator.[6]

RING GENERATOR ARCHITECTURE

A ring generator is a linear finite state machine created by forming a circle (ring) consisting of memory elements, and then gradually adding feedback connections that correspond to successive terms of a characteristic polynomial. Given tap x^k, we create a feedback loop by encompassing k adjacent flip-flops, always beginning with the leftmost ones. Figure 2 demonstrates this technique for the polynomial $x^{24} + x^{21} + x^{16} + x^{14} + x^{12} + x^8 + 1$.

In this approach, two feedback lines cannot cross each other. Instead, they form a regular rack frame provided that a suitable characteristic polynomial is deployed.[7] Because it is acceptable to choose different subsets of k adjacent flip-flops as long as the resultant feedback line does not cross any other feedback line, the ring generator offers a certain degree of flexibility in forming its structure. Figure 2

RESEARCH FEATURE

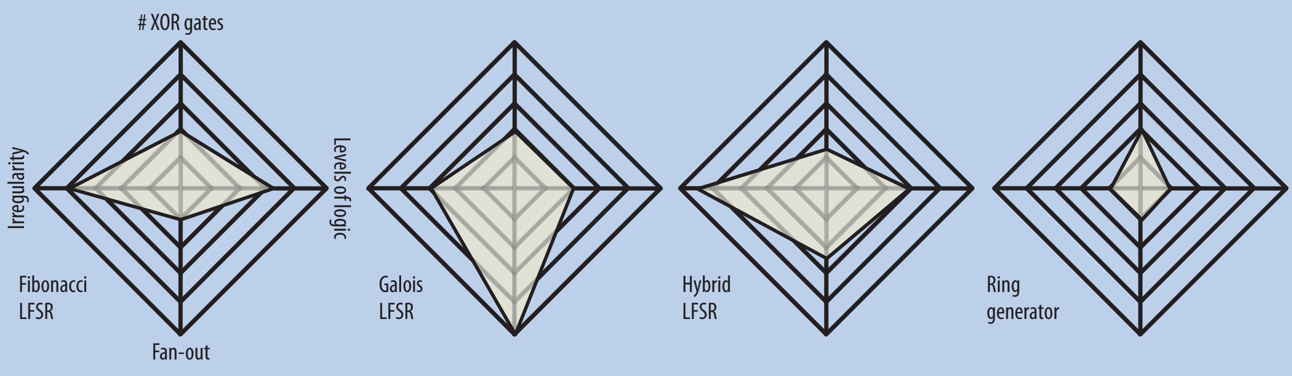

Figure 3. A comparison of three types of LFSRs and a ring generator:. Points on each axis indicate performance associated with four variables: number of XOR gates, fan-out counts, degree of irregularity, and number of logic levels. Once these points are connected, the resulting area indicates the design's ranking (the best performance is represented by the diagram's center).

shows two implementations of a feedback tap x^{21} to illustrate this phenomenon. The resultant ring generators will feature different state trajectories but still produce the same *m*-sequence as the corresponding LFSR, although differently phase-shifted in each case. Thus, given a primitive polynomial, the ring generator remains a maximum-length finite state machine.[8]

The ring generator features significantly reduced levels of XOR logic—only one 2-input XOR gate occurs between designated flip-flops if feedback logic is to be inserted. This design maximizes the operating speeds because a single 2-input XOR gate delay and short feedback lines cause no frequency degradation.

In addition to simple feedback connections, the reduced internal fan-outs are a ring generator's most noticeable feature. Because a switching device supplies a limited amount of current, the number of devices a single output of that switch can drive is also limited. Often, additional buffers boost and sharpen signals that might otherwise degrade or be distorted. Because these buffers deteriorate performance, the presence of large fan-outs is disadvantageous. Fortunately, any stem in the ring generator feeds (at most) two devices.

A ring generator's simplified circuit layout and routing make it highly modular. In many designs, connecting any device to other nodes requires routing many wires across a chip. To alleviate this problem, developers can use different segments of a ring generator to drive different regions of the circuit. They can easily place the ring generator because its successive components are cleanly separated and have a simple interface. In particular, they can use any connection (a wire) in a generator as a cut between segments because connections link only two memory elements and have no extra fan-outs.

Figure 3 compares three types of LFSRs with a ring generator. The radar diagrams' four axes represent the total number of XOR gates, the fan-out counts, the degree of irregularity, and the number of logic levels (representative of a device speed). The best performance for a specific variable is represented by points located very close to the center of the diagram. Linking the points on each axis creates a composite profile of each circuit. Its area directly corresponds to a given design's final ranking. As the figure shows, ring generators clearly outperform the LFSR solutions.

PHASE SHIFTER DESIGN

To reduce structural dependencies between sequences leaving the generator, a ring generator that needs to drive many nodes is often connected to a phase shifter. In principle, logic synthesis of a phase shifter entails finding a set of linear combinations of the generator outputs such that each resulting sequence, obtained by summing up these outputs, will be shifted with respect to every other sequence by at least a predefined number of bits.

Designers can use *dual ring generators* to obtain large and fast phase shifters. As Figure 4 shows, in a ring generator's dual form, the direction of all feedback connections is reversed. Hence, a dual ring generator features XOR gates placed on the outputs of those flip-flops that the original circuit used to drive feedback taps, while the feedback lines originate at the former locations of the respective XOR gates.

Assuming that a dual generator's initial state is a single logic 1 in the *i*th stage, its state after *q* clock cycles is of particular interest. Indeed, the locations of logic 1s in this vector identify the outputs of the original ring generator to be added modulo 2 to produce a sequence spaced *q* shifts up a reference.[8] The reference is a sequence originating from a stage pointed out by the logic 1 in the dual circuit's initial state. In other words, logic simulation of a dual ring generator creates a channel with a phase delay equal to *q* clock cycles with respect to the generator's *i*th stage. The following example further demonstrates this approach.

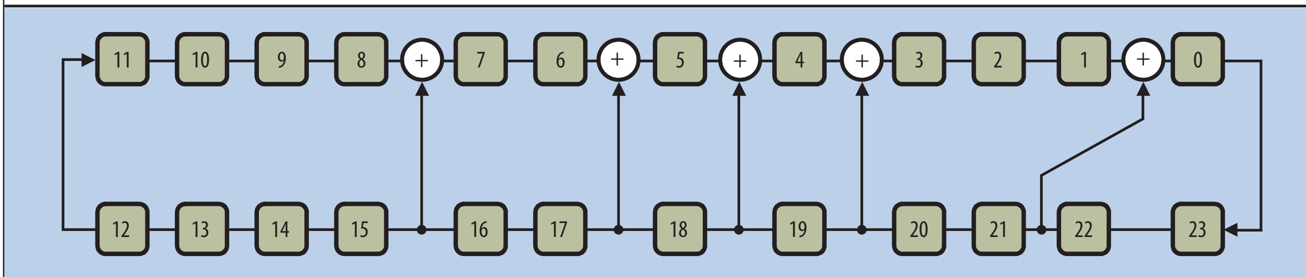

Figure 4. Dual ring generator structure. In this type of ring generator, the direction of all feedback connections is reversed.

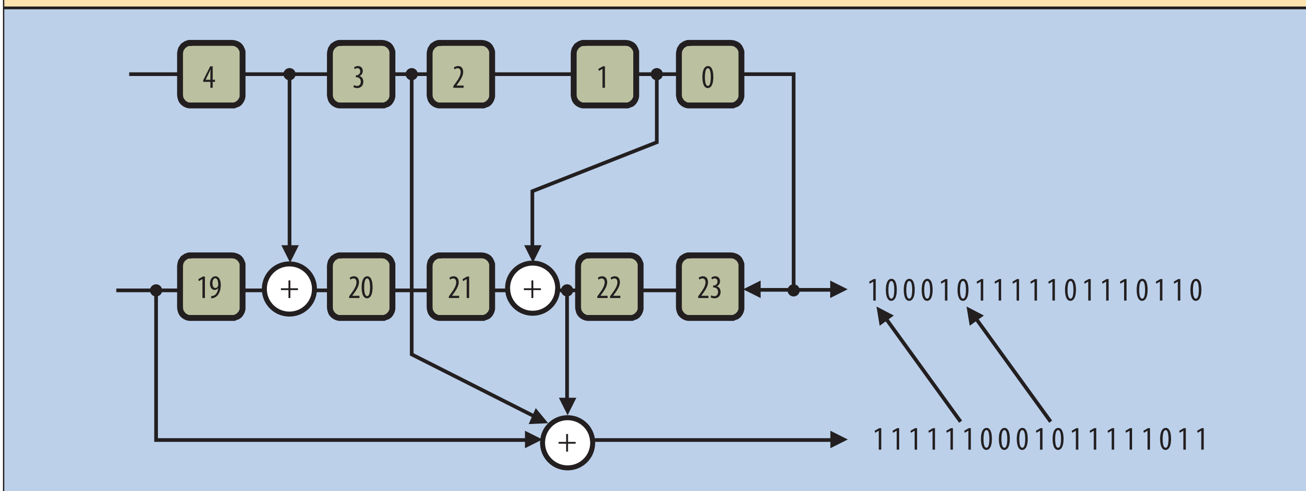

Figure 5. Phase shifter. Additional XOR gates produce replicas of an *m*-sequence shifted by a predetermined number of cycles. The XOR gate's inputs are selected through logic simulation of a dual ring generator.

Example 1. Assume we initialize the dual ring generator in Figure 4 with the following pattern:

0 0 0 0 0 0 0 0 0 0 0 1

0 0 0 0 0 0 0 0 0 0 0 0

This vector's content indicates that a sequence produced by flip-flop 0 (where logic 1 is located) of the original circuit will serve as a reference sequence for bitstreams a phase shifter produces. After five clock cycles, the dual ring generator reaches the following state:

0 0 0 0 0 0 0 0 1 0 0 0

0 0 0 0 0 0 0 1 0 0 1 0

This pattern indicates that adding modulo 2 the content of flip-flops 3, 19, and 22 of the original ring generator in Figure 2 generates a sequence that is a shifted (by 5 bits) replica of a sequence produced by the same generator's flip-flop 0. Figure 5 shows the corresponding phase shifter circuitry, a sequence it produces, and a reference sequence.

DENSE RING GENERATORS

As the number of nodes that a ring generator has to feed increases, a deployed phase shifter's complexity also increases. The large number of interconnects and their placement can result in on-chip routing congestion. Exploiting certain generic ring generator properties can alleviate these problems.

Given a ring generator, we might expect that outputs of many of its flip-flops could serve as parallel primary outputs with significant mutual phase shifts, assuming that an XOR gate is placed at every flip-flop output to be observed. Clearly, many such outputs can only be guaranteed by *dense* polynomials—that is, polynomials with many nonzero coefficients. However, to avoid increased fan-outs, the best polynomials for this purpose are those that feature nonzero coefficients for virtually every other term. The actual number of usable outputs then equals approximately $n/2$, where n is the generator's size.

Consider the 32-bit ring generator in Figure 6. It implements the dense primitive characteristic polynomial $x^{32} + x^{29} + x^{26} + x^{24} + x^{22} + x^{20} + x^{18} + x^{16} + x^{14} + x^{12} + x^{10} + x^{8} + x^{4} + x^{2} + 1$. The polynomial has 15 nonzero coefficients. Consequently, assuming that outputs of lower-row flip-flops serve as parallel outputs, this ring generator can feed up to 14 nodes without resorting to a phase shifter. We can verify that all the outputs shown in Figure 6 produce bitstreams mutually separated by at least 10,000 clock cycles. Dense ring generators similar to the one in

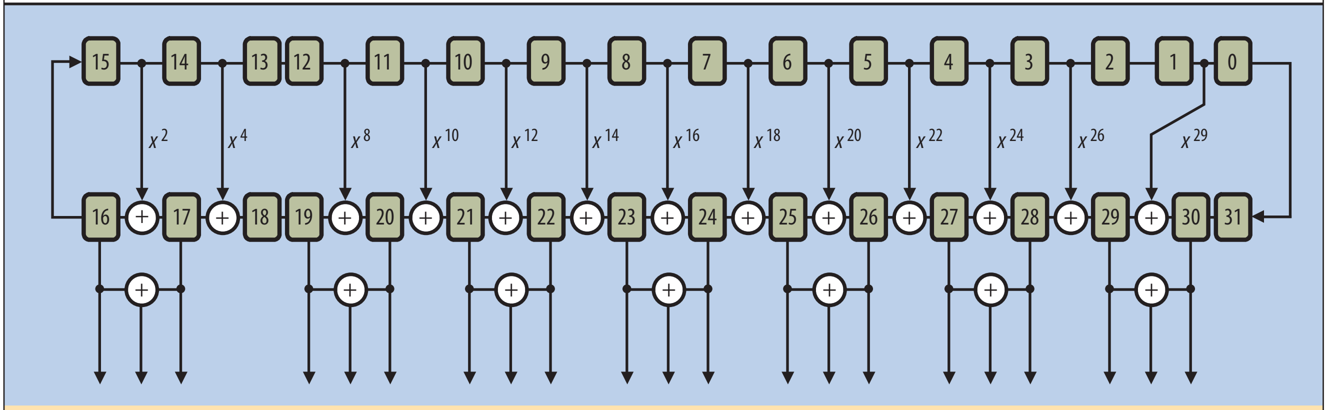

Figure 6. Dense ring generator. Outputs of many lower-row flip-flops serve as sources of phase-shifted *m*-sequences. Additional XOR gates further increase driving abilities of the ring generator with no negative impact on its overall performance.

Figure 6 maintain all advantages of ring generators. That is, in addition to a regular structure, each stem in a circuit has a maximum fan-out of 2, and the maximum delay is equal to the propagation time through a single 2-input XOR gate. Therefore, despite numerous feedback coefficients, this type of circuitry experiences no frequency degradation.

Developers can further increase the number of parallel outputs of a dense ring generator without affecting internal fan-out counts and propagation delays. As Figure 6 shows, they do this by adding extra 2-input XOR gates driven only by signals arriving from adjacent flip-flops located in the lower row of a circuit.

The successful adoption of high-performance dense ring generators requires a proper selection of feedback polynomials.[9] These polynomials allow a highly regular and modular implementation with the maximum parallel outputs. Furthermore, a mutual separation of 10,000 bits is guaranteed for every pair of bitstreams produced on the outputs of generators implementing those polynomials.

EVENT COUNTERS

State-of-the-art semiconductor digital circuits routinely require some form of discrete event counting performed at speed. Observability of their internals is crucial for embedded systems debugging, testing, and monitoring. Some useful events to observe are bus idle and data cycles, the number of grants or retries, device acquisition time or ownership time, and anomalies in gigabit Ethernet interfaces.

Counters are typically part of an on-chip facility that handles occurrences and durations, often comparable to switching times of single logic components. A conventional binary counter's content is easy to interpret and use. Its successive stages simply form a binary number corresponding to the number of clock cycles applied once the counter is reset and activated. Unfortunately, a critical path spanning all stages of a sequential counter severely affects the counting speed.

Although binary ripple-carry counters scale to virtually unlimited size yet increment in a constant time, this design style is quite complex, requires significant hardware real estate, and introduces unacceptable delays.[10] Other solutions, including prescaled counters or counters with a next-state generator, are relatively cumbersome to set up and use in automated synthesis as they have less-regular structures.[11]

Because ring generators feature much shorter critical paths, they can act as extremely fast event counters. Recovering the actual number of events recorded by a given *n*-bit ring generator is a nontrivial task, however. To count events, we initialize a ring generator with the value of 10 … 0, shifting it anytime an event occurs. Successive states visited by the ring generator will then correspond to a cumulated number of monitored events. Information produced by such a counter cannot be used directly, however, and the approach requires additional postprocessing.

A discrete-logarithm-based method computes the number of shifts (events) concealed in Galois LFSRs.[12] We can use a similar approach for a ring generator as long as its states are first properly converted into the corresponding states of the Galois LFSR.

Although a Galois LFSR and the corresponding ring generator feature different state trajectories, both circuits produce identical (though shifted) *m*-sequences on their outputs. In particular, they will always produce outputs with completely aligned *m*-sequences if proper initial states are chosen. This allows a mutual conversion between states of both devices. Such a conversion uses symbolic simulation to create linear expressions in the initial variables and for *n* successive clock cycles. In fact, these expressions form a set of *n* equations associated with two flip-flops, producing aligned *m*-sequences. In each equation, the expression representing an LFSR is equated with the corresponding expression obtained for a ring generator. Using Gaussian elimination with the

LFSR variables as pivots derives the actual mapping. As a result, given a ring generator's contents, we can easily compute the corresponding state of the LFSR, as the following example shows.

Example 2. Consider a 10-bit Galois LFSR and the corresponding ring generator, as shown in Figure 7. Let $x_9 \ldots x_0$ and $z_9 \ldots z_0$ be these devices' initial states, respectively. Linear equations over GF(2) in these 20 variables representing values of the two unshaded flip-flops in the figure (they produce fully aligned m-sequences) in 10 successive clock cycles are as follows:

$$x_0 = z_0$$
$$x_1 = z_1$$
$$x_2 = z_2$$
$$x_0 + x_3 = z_3$$
$$x_0 + x_1 + x_4 = z_4$$
$$x_1 + x_2 + x_5 = z_5$$
$$x_0 + x_2 + x_3 + x_6 = z_3 + z_6$$
$$x_1 + x_3 + x_4 + x_7 = z_3 + z_4 + z_7$$
$$x_2 + x_4 + x_5 + x_8 = z_4 + z_5 + z_8$$
$$x_9 = z_9$$

This set of equations reduces to:

$$x_0 = z_0, x_1 = z_1, x_2 = z_2, x_3 = z_3 + z_0,$$
$$x_4 = z_4 + z_1 + z_0, x_5 = z_5 + z_2 + z_1,$$
$$x_6 = z_6 + z_2, x_7 = z_7, x_8 = z_8 + z_0, x_9 = z_9$$

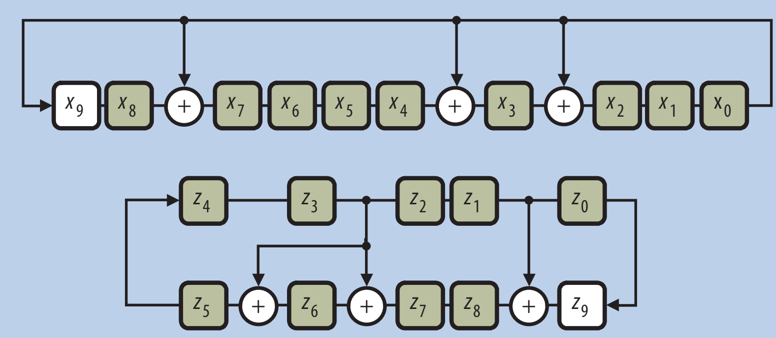

Figure 7. Galois LFSR and the corresponding ring generator. The *m*-sequences observed on the outputs of the unshaded flip-flops are completely aligned.

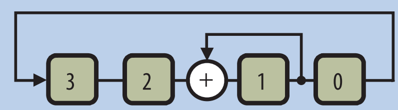

Initial	1000	0100	0010	0001
after $2^0 = 1$	0100	0010	0011	1000
$2^1 = 2$	0010	0011	1011	0100
$2^2 = 4$	1011	1111	1101	0011
$2^3 = 8$	0101	1010	0111	1110

Figure 8. Lookup table for ring simulation. Once constructed using the principle of superposition, the table allows logic simulation of very large linear finite-state machines.

Suppose the ring generator has reached state $z_9 \ldots z_0$ = 0110001011. Using this mapping, we can now compute values of successive bits of the corresponding Galois LFSR state. For instance, $x_4 = z_4 + z_1 + z_0 = 0 + 1 + 1 = 0$. Thus, $x_9 \ldots x_0$ = 0010100011.

After determining a state of the LFSR that corresponds to a recorded state of the ring-generator-based counter, we can easily recover the number of shifts concealed in the LFSR content. A discrete-logarithm-based counting[12,13] might serve this purpose.

FAST SIMULATION OF RING GENERATORS

The ability to quickly arrive at a state that a ring generator reaches after applying a given number of clock cycles is instrumental in many applications—simulation, sampling, testing, and so on. It would be possible to determine such a state by simply running the generator for up to 2^n clock cycles. However, this naive approach clearly fails for large values of n. A more subtle technique handles hundreds of bits-long finite state machines and provides an answer in at most $O(n^2)$ time.[13]

Consider a 4-bit ring generator implementing the primitive polynomial $x^4 + x + 1$ (the procedure is applicable to any linear finite-state machine). To determine a state that this circuit reaches after a given number of clock cycles, we can use a lookup table consisting of four rows and four columns, as Figure 8 shows. The entry located in the ith row and jth column represents a state that the ring generator enters after 2^i steps, assuming that the initial state features a single 1 on the jth position in its binary representation.

Single simulation steps suffice to determine the contents of the table's first row. Recall that this row contains states reachable after a single clock cycle. As a result, the first row entry in column j stores the state that immediately follows a state comprised of all zeros but a single 1 on position j. For instance, the ring generator in Figure 8 initialized with state 0010 moves to state 0011 in one step. This is represented by the entry in row one, column three of the table. We can easily determine entries in the remaining rows using the principle of superposition. Thus, other than in the first step, we do not need to use any form of logic simulation.

RESEARCH FEATURE

Suppose we want to determine the state that the ring generator in Figure 8 reaches after two clock cycles, assuming that its initial state was 0010. This problem reduces to finding a state that the same circuit reaches after one clock cycle, provided that its current state is 0011 (see the first row of the lookup table in Figure 8). Superposition allows further decomposition of the problem into two simpler tasks—finding immediate successors of states 0010 and 0001—as the superposition of these states yields state 0011.

As the first row of the lookup table shows, the ring generator moves in one step from states 0010 and 0001 to states 0011 and 1000, respectively. Their bit-wise sum results in state 1011. This is exactly the state that should be placed in row two, column three, of the lookup table. We obtain entries in the remaining rows similarly using data computed in the previous steps. Thus, we determine any state that a linear circuit reaches after 2^k cycles by adding respective states reachable after 2^{k-1} clock cycles.

Using tables like the one shown in Figure 8, we can find a state reachable after an arbitrary number of cycles (*C*) in at most *n* basic steps, each comprising as many as *n* lookups. The computational complexity of this process is thus $O(n^2)$. The method starts by expressing *C* as a sum of powers of 2, and then follows the rules presented earlier. The following example illustrates this technique.

Example 3. Let the ring generator in Figure 8 start in state 1010. Suppose we seek a state that this circuit reaches after the next 11 clock cycles. Because $11 = 2^0 + 2^1 + 2^3$, we proceed in three steps. We can find an immediate (after $2^0 = 1$ step) successor of state 1010 by determining immediate successors of states 1000 and 0010, which are, according to the lookup table in Figure 8, states 0100 and 0011, respectively. Their sum yields state 0111. Next, we decompose this state into three components, look up states that can be reached from them in $2^1 = 2$ steps, and find their sum, that is, 1100. Eventually, we apply the same approach to state 1100, this time by retrieving states reachable in $2^3 = 8$ clock cycles from states 1000 and 0100. These states are 0101 and 1010, respectively. Hence, we get the final state 1111.

Addressing challenges of nanoscale designs necessitates novel solutions. The semiconductor industry has already successfully implemented ring generators in hundreds of millions of integrated circuits. Although the ring generators were born in a digital test realm, they are well positioned to widely replace conventional linear feedback shift registers or linear cellular automata in other areas. Although they offer unprecedented performance features, they do not compromise the quality of advanced devices. Furthermore, designers can easily adopt rings' instrumentation, such as their fast simulation technique, to tackle problems related to verification of high-degree characteristic polynomials or computationally intensive logic simulation. With these capabilities, the automated design of microelectronics structures can further increase the capacity integrated into modern chips. ■

References

1. J. Rajski et al., "Embedded Deterministic Test for Low-Cost Manufacturing," *IEEE Design and Test of Computers*, vol. 20, no. 5, 2003, pp. 58-66.
2. J. Rajski et al., *Method for Synthesizing Linear Finite State Machines*, US patent 6,353,842, Patent and Trademark Office, 2002.
3. R. Lidl and H. Niederreiter, *Introduction to Finite Fields and Their Applications*, Cambridge Univ. Press, 1994.
4. L.-T. Wang and E.J. McCluskey, "Hybrid Designs Generating Maximum-Length Sequences," *IEEE Trans. Computer-Aided Design*, vol. 7, no. 1, 1988, pp. 91-99.
5. W.W. Warlick and J.E. Hershey, "High-Speed *m*-Sequence Generators," *IEEE Trans. Computers*, vol. 29, no. 5, 1980, pp. 398-400.
6. A.C. Arvillias and D.G. Maritsas, "Toggle-Registers Generating in Parallel *k* *k*th Decimations of *m*-sequences $X^p + X^k + 1$: Design Tables," *IEEE Trans. Computers*, vol. 28, no. 2, 1979, pp. 89-101.
7. J. Rajski and J. Tyszer, "Primitive Polynomials over GF(2) of Degree up to 660 with Uniformly Distributed Coefficients," *J. Electronic Testing: Theory and Applications*, vol. 19, 2003, pp. 645-657.
8. G. Mrugalski, J. Rajski, and J. Tyszer, "Ring Generators—New Devices for Embedded Deterministic Test," *IEEE Trans. Computer-Aided Design*, vol. 23, no. 9, 2004, pp. 1306-1320.
9. G. Mrugalski et al., "High-Performance Dense Ring Generators," *IEEE Trans. Computers*, vol. 55, no. 1, 2006, pp. 83-87.
10. J.E. Vuillemin, "Constant Time Arbitrary Length Synchronous Binary Counters," *Proc. IEEE Symp. Computer Arithmetic*, 1991, pp. 180-183.
11. M.R. Stan, A.F. Tenca, and M.D. Ercegovac, "Long and Fast Up/Down Counters," *IEEE Trans. Computers*, vol. 47, no. 7, 1998, pp. 722-735.
12. D.W. Clark and L.-J. Weng, "Maximal and Near-Maximal Shift Register Sequences: Efficient Event Counters and Easy Discrete Logarithms," *IEEE Trans. Computers*, vol. 43, no. 5, 1994, pp. 560-568.
13. N. Mukherjee et al., "High-Speed On-Chip Event Counters for Embedded Systems," *Proc. IEEE Int'l Conf. VLSI Design*, IEEE Press, 2009, pp. 275-280.

***Nilanjan Mukherjee** leads a technical group in the Design-to-Silicon Division at Mentor Graphics. His research focuses on developing next-generation test methodologies, test data compression, test synthesis, memory testing, and fault diagnosis. Mukherjee received a PhD in electrical engineering from McGill University, Montreal, Canada. He is a member of IEEE. Contact him at nilanjan_mukherjee@mentor.com.*

***Janusz Rajski** is a chief scientist and director of engineering at Mentor Graphics. His research interests include design automation and testing of VLSI systems, design for testability, built-in self-test, test data compression, and logic synthesis. Rajski received a PhD in electrical engineering from the Poznan University of Technology, Poland. He is a senior member of IEEE. Contact him at janusz_rajski@mentor.com.*

Grzegorz Mrugalski *is a software development manager at Mentor Graphics Polska, Poland. His research interests include design for testability, embedded test, and software quality assurance. Mrugalski received a PhD in electrical engineering from the Poznan University of Technology, Poland. He is a member of IEEE. Contact him at grzegorz_mrugalski@mentor.com.*

Artur Pogiel *is a software development engineer at Mentor Graphics Polska, Poland. His main research interests include design for testability, built-in self-test, embedded test, and fault diagnosis. Pogiel received a PhD in electrical engineering from the Poznan University of Technology, Poland. He is a member of IEEE. Contact him at artur_pogiel@mentor.com.*

Jerzy Tyszer *is a professor on the Faculty of Electronics and Telecommunications at the Poznan University of Technology, Poland. His research focuses on various aspects of VLSI test, design for testability, built-in self-test, embedded test, fault diagnosis, and computer simulation of discrete-event systems. Tyszer received a PhD in electrical engineering from the Poznan University of Technology, Poland. He is a senior member of IEEE. Contact him at tyszer@et.put.poznan.pl.*

CAREER OPPORTUNITIES

SYSTEMS ANALYST, Parsippany, NJ: Please contact employer for details. Reply to: VBS Data Resources, Inc. 28 Hale Court, Parsippany, NJ 07054.

SR. TECHNICAL CONSULTANT sought by GSPANN Technologies, Inc. in Milpitas, CA. Resp for providing professional comp consulting services in the form of systems analysis, dsgn & dvlpmt, systems integration &/or testing consulting. Min. Req. MS Comp Sci. Attn: HR 362 Fairview Way, Milpitas, CA 95035.

NORTHERN ILLINOIS UNIVERSITY, Postdoctoral Research Associate in Proton Computed Tomography. The Computer Science Department at Northern Illinois University has a DoD-funded opening for a postdoctoral research scientist to assist and develop software code for image reconstruction in proton computed tomography (pCT). Project goals include reconstructing 10 7 image voxels from 10 9 proton traces in under 10 minutes using a high performance CPU/GPU computational cluster. The researcher's responsibilities will include taking an active development role in the development and implementation of image reconstruction algorithms, building and maintaining virtual models of the pCT detector system and test objects using GEANT4, implementation testing with simulated and real proton trace data, and performing reconstructed image error analyses produced by varying system parameters. Desirable areas of skill, knowledge, and experience include as many of the following areas as possible: considerable programming experience in C/C++ and scripting in a Unix environment, Monte Carlo simulation techniques, GPU programming, MPI programming, tomographic image reconstruction, solution of large linear systems (10 9 equations with 10 7 unknowns), optimization of large non-linear systems, and GEANT4 programming. Candidates should also possess good communication and interpersonal skills. Applicants should possess either 1) a Ph.D. in Physics or related field with strong programming skills, 2) a Ph.D. in Computer Science with a strong physics background, or 3) a Ph.D. in Applied Mathematics with strong programming skills and physics background. The position is for two years pending continuation of funding, with a third year requested from DoD. Review of applications will begin on 1 Aug. 2011 and continue until the position is filled. Vitae and two letters of reference should be sent to Department of Computer Science Attn: Postdoc Search Northern Illinois University DeKalb, IL 60115. Pre-employment criminal background investigation required. NIU is an AA/EEO Employer.

WORKERS' COMPENSATION INSURANCE RATING BUREAU seeks Sr. App SW Eng in San Francisco, CA. Responsible for data conversion from legacy sys to new sys & required interfaces. Mail resume to: WCIRB. Attn: D. Smith. 525 Market St, Ste. 800, San Francisco, CA 94105. Ref job code: SASE.

SOFTWARE ENGINEER: Operational-level programming, sys anal., design for routine app sys or parts of moderately complex or complex sys. Determine source data (input), proc. reqs (output), output formats, timing estimates. Verify sys or netw. meets perf. criteria. Work w/ offshore team to provide alternate pro-

JUNIPER NETWORKS

— Juniper Networks, Inc. is recruiting for our Sunnyvale, CA office: —

Software Engineer #12857: Work with Product Management to hash out feature requirements. Work with cross-functional software teams to finalize solution design and implement the solution.

Test Engineer #20037: Test different storage protocols and functionalities of company storage products. Test case execution and validation.

Advanced Services Engineer #20154: Deliver risk mitigation and optimization services to Advanced Services Customers. Inform customer of major risks in current and target software releases to help customer make educated decision regarding software.

Engineering Services Engineer #19624: Analog electrical engineering design of power conversion, data acquisition, and power control. Schematic capture and component selection.

IT Network Engineer #13530: Design and implement IT infrastructure, including servers, storage, operating systems and networks to support company ERP and CRM applications.

Test Engineer #6083: Test complex routing and switching products, write test automation scripts, use third party network test equipment, write libraries and programming interfaces to interact with various test devices.

Technical Marketing Specialist Staff #20051: Provide technical expertise to customers and Juniper's field sales organization and aid in closing strategic opportunities.

Quality Assurance Engineer #18792: Work closely with development team to execute, debug and develop test suites in Perl.

Software Engineer #7376: Design, dev. and maint. features related to IPSEC VPN on Junos, and be respon. for writing technical specs of the features being implemented.

Hardware Engineer #8425: Involve in the design of hardware, documentation related to design specification, test document, bringup, testing and working with cross functional team members.

Services Program Manager #19828: Serve as the expert on the functionality and data managed within the Clarify/Case Manager solutions. Provide daily data administration support for the case management platforms, including business notifications, product classifications and custom attributes.

Technical Support Engineer #12465: Analyze customer requirements and design functional specifications. Develop interactions and solutions to optimize the ISP network to meet customer's specific needs.

Quality Assurance Engineer #7328: Lead system test team and drive deliverables, benchmark new software and hardware performance to provide insight to developers and architects on improvements and potential issues.

— Juniper Networks, Inc. is recruiting for our Herndon, VA office: —

Technical Support Engineer #4118: Provide high level technical expertise on specific product(s) to company's customers. Work closely with the customer's support and account team.

Mail resume with
job code # to
Attn: MS A1.2.1.435
Juniper Networks
1194 N. Mathilda Avenue
Sunnyvale, CA 94089

Published by the IEEE Computer Society

0018-9162/11/$25.00 © 2011 IEEE

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

gram design approaches. Softw. Install. & testing, DB admin, user support, hardware eval. MS CS or related + 3 yrs relevant exp. or BS + 5 yrs. Catalyst Rx, 800 King Farm Blvd Rockville MD 20850. To apply email hr@catalystrx.com w/ job# SE-1 in subj. line. EOE.

SR. CONSULTANT, Austin, TX, Ascendant Technology. QA Test Analysis. BA (or foreign equiv) in Comp Sci or related +3 yrs exp. in QA test environ. Resumes only to C. Jones, HR Mgr., Ref# 091253, 16817 167th Ave. NE, Woodinville, WA 98072.

ENBRIDGE EMPLOYEE SERVICES, INC. in Minot, ND seeks Senior Engineer – Engineering Services (Job Code 506335) to perform engineering duties providing engineering support to management, administration, & field personnel on operational matters & design, construction, project management of capital & expense projects, relating to crude liquid pipeline & facility gathering systems. Submit resumes by mail referencing job code to HR, 1409 Hammond Avenue, Superior, WI 54880 - EOE

invent

Hewlett-Packard State & Local Enterprise Services, Inc.
is accepting resumes for the following positions:

IT Developer/Engineer

Sacramento, CA. (Ref. #RSLSACITDE11)

Research, design, develop, configure, integrate, test, and maintain existing and new business applications and/or information systems solutions including databases through the integration of technical and business requirements. Extensive travel required to various unanticipated locations throughout the US.

Services Information Developer

Olympia, WA. (Ref. #SLOLYSID11)

Conceptualize, design, develop, unit-test, configure, or implement portions of new or enhanced (upgrades or conversions) business and technical software solutions through application of appropriate standard software development life cycle methodologies and processes.

Technology Consultant

Palo Alto, CA. (Ref. #RSLPALTC11)

Provide technology consulting to customers and internal project teams. Provide technical support and/or leadership in creation and delivery of technology solutions designed to meet customers' business needs and, consequently, for understanding customers' businesses. Extensive travel required to various unanticipated locations throughout the U.S.

Mail resume to Hewlett-Packard State & Local Enterprise Services, Inc., 5400 Legacy Drive, MS H1-6F-61, Plano, TX 75024. Resume must include Ref. #, full name, email address & mailing address. No phone calls please. Must be legally authorized to work in the U.S. without sponsorship. EOE.

SR. SOFTWARE ENGINEER: Analyze, dsgn, dvlp, test & implmt real time messaging applic systems using knowl of & exp w/Core Java, J2EE,Spring, Tibco/JMS, Oracle 11, TCP/IP, FIX parsers, SOAP & Rest web services, Purify, Quantify, Linux & Windows, RUP & IConix. Should have exp in Test driven dvlpmt, Extreme prgmg &/or Scrum. Reqs MS Comp Sci, Eng or rel. Mail resumes to Informatic Technologies Inc., 900 Oak Tree Ave., Ste. C2, South Plainfield, NJ 07080.

SIEBEL ADMINISTRATOR, Memphis, TN. Travel to multiple client locations nationwide to analyze, design, develop, administer, support sophisticated SIEBEL Applications in multiplatform environment. Work with PERL scripting, SQL, Java, (Func&Tech), Shells (SH, CSH) on Sun Solaris, Unix platforms, eScript, JAVA scripts. Test, troubleshoot, administer, perform QA in production environment. Reply to: Comspark International Inc., 3265 W. Sarazen Circle # 201, Players Crossing Office Park, Memphis, TN 38125.

NETWORK ADMINISTRATOR: Dsgn, dvlp & test WAN n/works utilizing TCP IP n/work w/OSI model using Cisco, Juniper, routing/switching protocols BGP, OSPF, MPLS, ATM, Frame relay, IP Sec encryption. Installation & admin of sniffer, packet analysis tools on n/work. CCNA , fiber channel n/work, SAN is reqd. Knowl of VLAN, DNS, DHCP, HTTP, SSL, FTP, STP, RSTP preferred. Reqs MS Comp Sci, Eng or rel. Mail resumes Saphire Solutions Inc., 523 Green St., 2nd Fl., Iselin NJ 08830.

COMPUTER OUTSOURCING DATA SERVICES/BI CORPORATE ARCHITECT (Unisys; Rockwall, TX): Leadership & creativity in dvlpmnt & implementation of data services & business intelligence strategy, service & solutions covering wide range of outsourcing sw systems. Design, definition & realization of data related projects. Apply extensive bkgrnd in ticketing systems in ITSM-based ticketing. Create reqs for db & data stores. Use advanced BI methodologies such as OLTP, OLAP, Cube & Dimension model-

ing. Apply knowledge of SDLC dvlpmnt methodologies. Utilize numerous related techs. Req. BS CS; 10 yrs. exp in multiple areas of tech or app specialization w/strong bkgrnd in managed services & outsourcing domain. Resumes to IEEE Computer Society, 10662 Los Vaqueros Circle, Box #COM53, Los Alamitos, CA 90720.

SENIOR SAP DEVELOPMENT LEAD: provide SAP support and programming in all areas of operations such as Sales/Distribution, S&OP, Inventory, Purchasing and Production issues and new enhancements. BS in Computer Science/Engineering and 5 years of SAP experience. Job location: Nashville, TN. Mail resume to Renée C. Drinnon, Akzo Nobel Coatings Inc, 20 Culvert St., Nashville, TN 37210. Refer to job code MB001 when applying.

SAP BW ANALYSTS sought by 21st Century Professionals, Corp, an emerging IT consulting co. providing cost effective & business integrated solutions to Customers, in Westminster, CO. Provide analysis & support for SAP BW tasks incl dvlp BW reqmts & dsgn deliverables per ASAP Methodology ensuring that SAP BW dsgn is compatible w/business reqmts. Min req MS Engg Tech or rltd deg. 3116 W 112th Ct, Ste. D, Westminster, CO 80031.

FUJITSU NETWORK COMMUNICATIONS INC. has the following job opportunities available in Sunnyvale, CA: •Software Development Engineer (Req # FNC1437): Design new user interface client components for the NETSMART 1500 product using Java and Java 2 enterprise technologies. •Software Development Engineer (Req # FNC1386): Develop Object Oriented software in Java/C++ using CORBA middleware technology and design Applications and Reusable Software libraries. Submit resume to Fujitsu Network Communications, Staffing Department, 2801 Telecom Pkwy, Richardson, TX 75081. Req # must be noted or referenced when submitting resume.

CLASSIFIED LINE AD SUBMISSION DETAILS: Rates are $400.00 per column inch ($500 minimum). Eight lines per column inch and average five typeset words per line. Free online listing on careers.computer.org with print ad. Send copy at least one month prior to publication date to: Marian Anderson, Classified Advertising, Computer Magazine, 10662 Los Vaqueros Circle, Los Alamitos, CA 90720-1314; (714) 821-8380; fax (714) 821-4010. Email: manderson@computer.org.

invent

Hewlett-Packard Company has an opportunity for a

IT Developer/Engineer

Austin, TX
Reference: AUSITDE51

Research, design, develop, configure, integrate, test, and maintain existing and new business applications and/or information systems solutions including databases through integration of technical and business requirements.

Mail resume to Hewlett-Packard Company, 5400 Legacy Drive, H1-6F-61, Plano, TX 75024. Resume must include Ref. #, full name, email address & mailing address. No phone calls please. Must be legally authorized to work in the U.S. without sponsorship. EOE.

invent

Hewlett-Packard Company has an opportunity for the following position in **Palo Alto, CA** and at various unanticipated sites throughout the U.S.

Technology Solutions Consultant

Reference: PALPCA2

Reqs: Knowledge of UNIX & Networking. Exp. must include at least: 2 yrs installing & using HP's Server Automation SW; using Oracle DB tools; & 5 yrs performing customer support to internal users or external customers.

List full name, address & email address on resume. Ref Job# PALPCA2 & send resumes to Hewlett-Packard Company, H1-6E-28, 5400 Legacy Drive, Plano, TX 75024. No phone calls pls. Must be legally authorized to work in the U.S. w/o sponsorship. EOE.

invent

HP Enterprise Services, LLC is accepting resumes for the following positions:

Information Systems Architect

Thousand Oaks, CA • Ref: ESTHOUISA11

Architect effective information systems solutions that address the customer's business problems, needs and opportunities, in a manner consistent with the company's strategic and business goals.

Technology Consultant

Oakland, CA • Ref: ESOAKSDI1

Provide technology consulting to customers and internal project teams. Provide technical support and/or leadership in creation and delivery of technology solutions designed to meet customers' business needs and, consequently, for understanding customers' businesses.

Mail resume to HP Enterprise Services, LLC, 5400 Legacy Drive, MS H1-6F-61, Plano, TX 75024. Resume must include Ref. #, full name, email address & mailing address. No phone calls please. Must be legally authorized to work in the U.S. without sponsorship. EOE.

Apple is looking for qualified individuals for following 40/hr/wk positions. To apply, mail your resume to 1 Infinite Loop 84-REL, Attn: AEC Staffing-LJ, Cupertino, CA 95014 with Ref # and copy of ad. **Job site & interview, Cupertino, CA**. Principals only. EOE.

Electrical Engineer [Ref#7443364]

Architect, design and validate critical aspects of electrical behavior of displays for Apple's handheld form factor consumer line (iPod and iPhone). Requires Bachelor's degree, or foreign equivalent, in Electrical Engineering, or related and 2 years professional experience in job offered or in a related occupation, including analog design fundamentals; PCB design and layout; circuit simulation; power conversion design; D/A converter design and validation; signal conditioning design fundamentals; and fundamentals of Integrated IC design and troubleshooting.

Location Software Engineer [Ref#7575381]

Shape the next generation of location-aware iPhone applications. Requires Bachelor's degree, or foreign equivalent, in Computer Science, or related field and 5 years professional experience in job offered or in a related occupation. Professional experience must be post-baccalaureate and progressive in nature. Must have professional experience with: embedded platform driver development in Unix; MEMS sensors or GPS hardware integration in embedded mobile devices; system/software architecture and design; C++; object oriented design; communication protocols and IPC; unit test design and debugging.

Radio System Integrator [Ref#7519837]

Evaluate future cellular chipsets and technologies and drive their System Architecture, proof-of-concept, planning, and execution into the iPhone/iPad platform. Req.'s Bachelor's degree, or foreign equivalent, in Electrical Engineering, Computer Engineering, or related. Five (5) years professional experience in job offered or in a related occupation. Professional experience must be post-baccalaureate and progressive in nature. Also experience with: computer architecture, advanced logic/digital design; hardware design of high-volume Smartphones; wireless communications protocols such as GSM, CDMA2000, WCDMA, GPS, and WLAN; experience with integration and testing of RF components and circuits; design and validation of power supplies such as linear regulators and switched-mode regulators, DC/DC converters; design power management of radio on iPhone/iPad.

Software Engineering Manager [Ref# 7533008]

Responsible for the Quality Assurance for Apple's iOS products. Requires Master's degree, or foreign equivalent, in Electrical Engineering, or related and 4 years professional experience in job offered or in a related occupation, including Bluetooth Stack and Bluetooth Profiles; technical leadership and project management; Bluetooth and WiFi Physical layer testing and coexistence algorithms; Bluetooth and WiFi Logging and Debugging Tools including FTS Bluetooth and WiFi Sniffer; testing and quality assurance methodologies; C; C++; Python; Perl; audio quality tools; and engineering tools including oscilloscopes, Spectrum Analyzers, and Logic Analyzers.

Wireless System Test Engineer [Ref# 7599005]

Perform interoperability testing of Apple wireless products (iPad, iPhone, Mac CPUs, portable/desktops, and Base stations or Access points) with 3rd party WLAN/BT devices. Requires Master's degree, or foreign equivalent, in Network Communications, or related and 2 years professional experience in job offered or in a related occupation, including IEEE 802.11n, 802.11e and WMM; 802.11i including 802.1x IAS/ACS/FreeRadius/SBR authentication technologies; associated protocol sniffers like tcpdump, ethereal, WildPackets Omnipeek and Wireshark; Controller based WLAN roaming, mobility and high density wireless client stress testing using mad-wifi based client simulators like vsta and wpa-sim; WLAN automation using C; PERL; Tcl; Bash.

Software Development Engineer [Ref# 7610816]

Design, develop and implement data warehousing and client-server applications spanning business, financial and other related areas. Requires Bachelor's degree, or foreign equivalent, in Computer Science and Engineering, or related and 5 years professional experience in job offered or in a related occupation. Professional experience must be post-baccalaureate and progressive in nature. Also experience with data warehousing; Business Intelligence, Data Warehouse and/or Data Management space; MicroStrategy 9; MicroStrategy data warehouse; performing database tuning and optimization; MicroStrategy application that crosses multiple business lines; ETL tools and UNIX batch processes; SQL/Excel/VB toolsets or languages; ROLAP/MOLAP/OLAP Tools; Relational Reporting Tools.

Systems Engineer [Ref#7543098]

Design, architect, develop, and maintain high performance systems. Requires Bachelor's degree, or foreign equivalent, in Computer Science and Engineering, or related degree and 5 years professional experience in job offered or in a related occupation. Professional experience must be post-baccalaureate and progressive in nature. Must have professional experience with architecture, design and development of scalable servers; Java, Shell scripts, SED AWK and PERL; SQL, UNIX internals including IPC, Signals, Threads and Sockets; Internet protocols TCP/IP, HTTP HTTPS, SMTP, POP, IMAP, LDAP, SSH; version control systems CVS and SVN; and BIRT/JASPER reports.

invent

Hewlett-Packard Company has an opportunity for a

Systems/Software Engineer

Jamestown, RI.
Reference: JAMFMC2

Reqs: BS in CE, CS or rltd & 10 yrs exp as a Sys SW Engr or rltd. Expert knowledge and exp. in C progrmng lang.; Linux Kernel Devlpmt, espec. in file sys & block storage; Devlpmt of high-availability systems; SAN based storage incldg Fibre Channel, iSCSI and ATA over Ethernet. List full name, address & email address on resume.

Please send resume referencing Job# JAMFMC2 to Hewlett-Packard Company, H1-6E-28, 5400 Legacy Drive, Plano, TX 75024. No phone calls please. Must be legally authorized to work in the U.S. without sponsorship. EOE.

invent

Hewlett-Packard Company has an opportunity for the following position in **Palo Alto, CA** and at various unanticipated sites throughout the U.S.

Technology Consultant (Bus Intel Solutns)

Reference: PALLKA2

Req: BS in CS, CE, SW Eng, or rltd; 5 years exp as a Tech Conslt or rltd.; Developing with Informatica, Ab Initio, or DataStage Data Integration tools; Knowledge of Bus. Intel/Data Warehousing solutions; Must have deep technical knowledge in order to implement specific deliverables; Exp with design contribution: List full name, address & email address on resume. Send resume & refer to Job# PALLKA2.

Please send resumes referencing Job# PALLKA2 to Hewlett-Packard Company, H1-6E-28, 5400 Legacy Drive, Plano, TX 75024. No phone calls please. Must be legally authorized to work in the U.S. without sponsorship. EOE.

invent

Hewlett-Packard Company has an opportunity for the following position in **Palo Alto, CA** and at various unanticipated sites throughout the U.S.

Technology Consultant

Reference: PALPKU2

Req: BS in Eng, CS, CE or rltd; 6 years exp as a Tech Conslt or rltd. Exp in/as a Certified Solution Developer (Websphere IIS DataStage Enterprise Edition V7.5); Data Warehouse App Design, Data Warehouse App Devel; Extract, Transform, and Load Design; Bus Intel App Design, Bus Intel App Devel; Working knowledge of Oracle DB, Working knowledge of Teradata DB: List full name, address & email address on resume. Send resume & refer to Job# PALPKU2.

Please send resumes referencing Job# PALPKU2 to Hewlett-Packard Company, H1-6E-28, 5400 Legacy Drive, Plano, TX 75024. No phone calls please. Must be legally authorized to work in the U.S. without sponsorship. EOE.

invent

Hewlett-Packard Company has an opportunity for a

Systems/Software Engineer

Boise, ID
Reference: BOISAG2

Reqs: MS & 1 yr exp & exp w/ application of statistical analysis in engineering environment; applied mathematics; time-series analysis; stochastic process theory; optimizing algorithms; pattern analysis & recognition; wavelets (multiresolution analysis); queuing theory; analytical & simulated modeling of storage systems. List full name, address & email address on resume. Send resume & refer to Job# BOISAG2.

Please send resumes with job number to Hewlett-Packard Company, H1-6E-28, 5400 Legacy Drive, Plano, TX 75024. No phone calls please. Must be legally authorized to work in the U.S. without sponsorship. EOE.

invent

Hewlett-Packard Company has an opportunity for the following position in **Palo Alto, CA** and at various unanticipated sites throughout the U.S.

Technology Consultant

Reference: PALAME

Requires: BS degree in CS or rltd & 8 yrs exp as a Tech Consultant or rltd. DB platforms: Netezza, Teradata, and Oracle; Data integration tools: Ab Initio, Informatica and Data Stage; Change data capture tools such as Golden Gate; Coding in programming langs such as UNIX, PERL and Java; Data modeling with tools such as ERWIN; Front end reporting tools such as Business Objects. List full name, address & email address on resume.

Please send resume referencing Job# PALAME2 to Hewlett-Packard Company, H1-6E-28, 5400 Legacy Drive, Plano, TX 75024. No phone calls please. Must be legally authorized to work in the U.S. without sponsorship. EOE.

invent

Hewlett-Packard Company has an opportunity for a

Systems Software Engineer

Andover, MA
Reference: ANDRNE2

Reqs: MS in CS, IT, Engrg, Math, Physics or rltd & 3 yrs exp as a Systems Software Engineer or rltd. Exp dvlping SW apps; SW build envrnmts, source control mgmt (espec iSVN) and SW installers; Java expertise; Broad tech. Bckgrnd; Exp. with VC++, HTML, Javascript, Css, VBScript, Perl, Visual Studio, Dos batch programming, WSH scripting, Ant, setting up ESX server and usage of VSphere client to manage VMWare Virtual Machines, subversion, SmartSVN, Perforce Administration. List full name, address & email address on resume.

Please send resume referencing Job# ANDRNE2 to Hewlett-Packard Company, H1-6E-28, 5400 Legacy Drive, Plano, TX 75024. No phone calls please. Must be legally authorized to work in the U.S. without sponsorship. EOE.

invent

Hewlett-Packard Company has an opportunity for a

Thermal Engineer

Houston, TX
Reference: HOUVKH2

Reqs: MS in Thermal Mech Engrg, ME or rltd & 1 yr exp as a Thermal Eng or rltd. Exp in CFD mdlng; Air flow/thermal analysis in elect. sys. incl srvr & strg sys dsgn; Chilled water & refrig based cooling soltns. List full name, address & email address on resume.

Please send resume referencing Job# HOUVKH2 to Hewlett-Packard Company, H1-6E-28, 5400 Legacy Drive, Plano, TX 75024. No phone calls please. Must be legally authorized to work in the U.S. without sponsorship. EOE.

invent

Hewlett-Packard Company has an opportunity for a

Software Developer

San Diego, CA
Reference: SANCVA2

eqs: exp w/ Project mgmt skills; Agile methods; Metrics & reporting knowledge; prod dev't life cycle; ITIL; OS knowledge & behavior of operating sys; Sw devl't & QA best pract; a scripting language; & SQL query language. List full name, address & email address on resume.

Please send resumes referencing Job# SANCVA2 to Hewlett-Packard Company, H1-6E-28, 5400 Legacy Drive, Plano, TX 75024. No phone calls. Must be legally authorized to work in the U.S. without sponsorship. EOE.

invent

Hewlett-Packard Company has an opportunity for a

Software QA Engineer

Cupertino, CA
Reference: CUPSVU2

Reqs: MS & 1 yr exp. Exp. w/Automation tools including Quality Center, Quick Test Professional, VB scripting, & Load Runner; Scripting languages Shell & Perl; Programming in Java; Microsoft Exchange & related tech; Jboss app server; Linux & Windows op systems; & DB2, Mysql DB.

List full name, address & email address on resume. Ref Job# CUPSVU2 and send resume to Hewlett-Packard Co., H1-6E-28, 5400 Legacy Dr, Plano, TX 75024. No phone calls pls. Must be legally authorized to work in the U.S. w/o sponsorship. EOE.

invent

Hewlett-Packard Company has an opportunity for a

Software Design Engineer

Bellevue, WA
Reference: BELGWU2

Reqs: BS & 5 yrs exp. Spec. Reqs: 5 yrs exp in SW design & implementation; 3 yrs exp working w/network protocols or network device mgmt; Exp w/MVC UI frameworks such as Spring, JBoss, DOJO; Java, Javascript, CSS, HTML, AJAX, Web Services; & SQL/relational DB w/exp in one or more of: MySQL, MS SQL Server, Oracle. Platform exp on one or more of Linux, Solaris, Windows. Server-side dev't techniques. Exp w/network protocols, design & configuration; & design/code quality.

List full name, address & email address on resume. Ref Job# BELGWU2 & send resumes to Hewlett-Packard Co., H1-6E-28, 5400 Legacy Drive, Plano, TX 75024. No phone calls pls. Must be legally authorized to work in the U.S. w/o sponsorship. EOE.

invent

Hewlett-Packard Company has an opportunity for a

Systems/Software Engineer

Roseville, CA
Reference: ROSFHU2

Reqs: Master's + 2 yrs exp. & exp. with comp & network security; Ethernet network architectures, network switching & mgmt in enterprise & small business environments. List full name, address & email address on resume.

Please send resumes referencing Job# ROSFHU2 to Hewlett-Packard Company, H1-6E-28, 5400 Legacy Drive, Plano, TX 75024. No phone calls please. Must be legally authorized to work in the U.S. without sponsorship. EOE.

invent

Hewlett-Packard Company has an opportunity for a

Software Designer

San Diego, CA
Reference: SANVKO2

Reqs: BS in CS or rltd & 5 yrs exp as a SW Dsgnr or rltd. Expertise in Java prgrmg lang; various web technologies (JavaScript, JSPs, XML, AJAX); App. Srvrs such as Tomcat, WebSphere or WebLogic Test automation eg, JUnit; Performance validation and profiling tools (Jprobe or Optimizelt, LoadRunner, DynaTrace). List full name, address & email address on resume.

Please send resume referencing Job# SANVKO2 to Hewlett-Packard Company, H1-6E-28, 5400 Legacy Drive, Plano, TX 75024. No phone calls please. Must be legally authorized to work in the U.S. without sponsorship. EOE.

ADVERTISER INFORMATION
JUNE 2011

Advertising Personnel

Marian Anderson: Sr. Advertising Coordinator
Email: manderson@computer.org
P: +1 714 821 8380; F: +1 714 821 4010

Sandy Brown: Sr. Business Development Mgr.
Email sbrown@computer.org
P: +1 714 821 8380; F: +1 714 821 4010

Advertising Sales Representatives (display)

Western US/Pacific/Far East:
Eric Kincaid
Email: e.kincaid@computer.org
Phone: +1 214 673 3742
Fax: +1 888 886 8599

Eastern US/Europe/Middle East:
Ann & David Schissler
Email: a.schissler@computer.org, d.schissler@computer.org
Phone: +1 508 394 4026
Fax: +1 508 394 4926

Advertising Sales Representative (Classified Line/Jobs Board)

Greg Barbash
Email: g.barbash@computer.org
Phone: +1 914 944 0940
Fax: +1 508 394 4926

Special Technical Communities

Dejan Milojicic, *HP Labs*
Phil Laplante, *Pennsylvania State University*

STCs offer a new way for members to develop communities focusing on selected technical areas, integrating with contemporary technology in a way that breaks down silos and fosters growth in a dynamic environment.

Businesses must evolve to meet the needs of their customers and an ever-changing, competitive marketplace. Likewise, professional societies must adapt to meet the needs of their members in a dynamic environment.

The IEEE Computer Society (IEEE-CS) has served its membership well over the past 65 years by providing high-quality publications, conferences, standards, technical committees, educational products, and support for local chapters. But these well-defined services have led to a stovepiped structure characterized by rigid control and boundaries in which formal processes and procedures have proliferated and crystallized to the point where it's difficult to respond quickly to the rapidly changing needs of both members and the competitive market.

Clearly, a more flexible structure is needed for the IEEE-CS to meet its members' needs. Special Technical Communities (STCs) represent a new way for members to develop communities focusing on selected technical areas, integrating with contemporary technology in a way that breaks down silos and fosters organic growth. These communities will be able to scale up, scale down, or even shut down when they no longer serve the community's needs. As a result of this new structure, we foresee more synergy across publications, conferences, and standards activities—leading to increased intellectual property creation and enhanced financial sustainability for the IEEE-CS.

CREATING COMMUNITIES

Professional associations are all about creating communities. In the mid-1600s, scientists began creating associations to exchange ideas and practices, and engineers began forming similar groups in the mid-1800s. Today, professional associations are still an important medium of exchange for technologists, researchers, educators, students, and other stakeholders. Within these professional societies, focused communities of interest have evolved and are sustained through meetings, publications, lectures, short courses, standards development, and other activities.

Traditionally, professional associations build communities through letters, printed journals, and face-to-face interaction. Modern communications, however, have changed this process, initially through audio and video conferencing, followed by the use of simple Internet-based services such as e-mail and websites. Soon after e-mail became commonplace, mailing lists, bulletin boards, and basic archives became important media for interaction.

More recently, Web 2.0 and social networking technologies have enhanced remote interaction. Although face-to-face contact is still valued, professionals, especially younger ones, have become used to interacting virtually. Some even question whether face-to-face meetings are necessary if technology can support rich interaction (P. Wiesner et al., "Virtual Communities for the Technical Professional," *Proc. ASEE Ann. Conf.*, 2003; www.asee.org).

With the growth of the IEEE-CS and its stovepiped organizational structure, the formation of new publications, conferences, and other traditional products has become more difficult. For example, even when someone could make a business case to start a new magazine, the process could take several years; discontinuing a publication has always required considerable political willpower. It sometimes requires more than a year to thoroughly review and publish an article. While IEEE-CS is less willing to invest resources when certain areas of interest may only be relatively fleeting, members are less tolerant of delays in creating new publishing outlets and in publishing articles than they once were.

Published by the IEEE Computer Society
0018-9162/11/$26.00 © 2011 IEEE

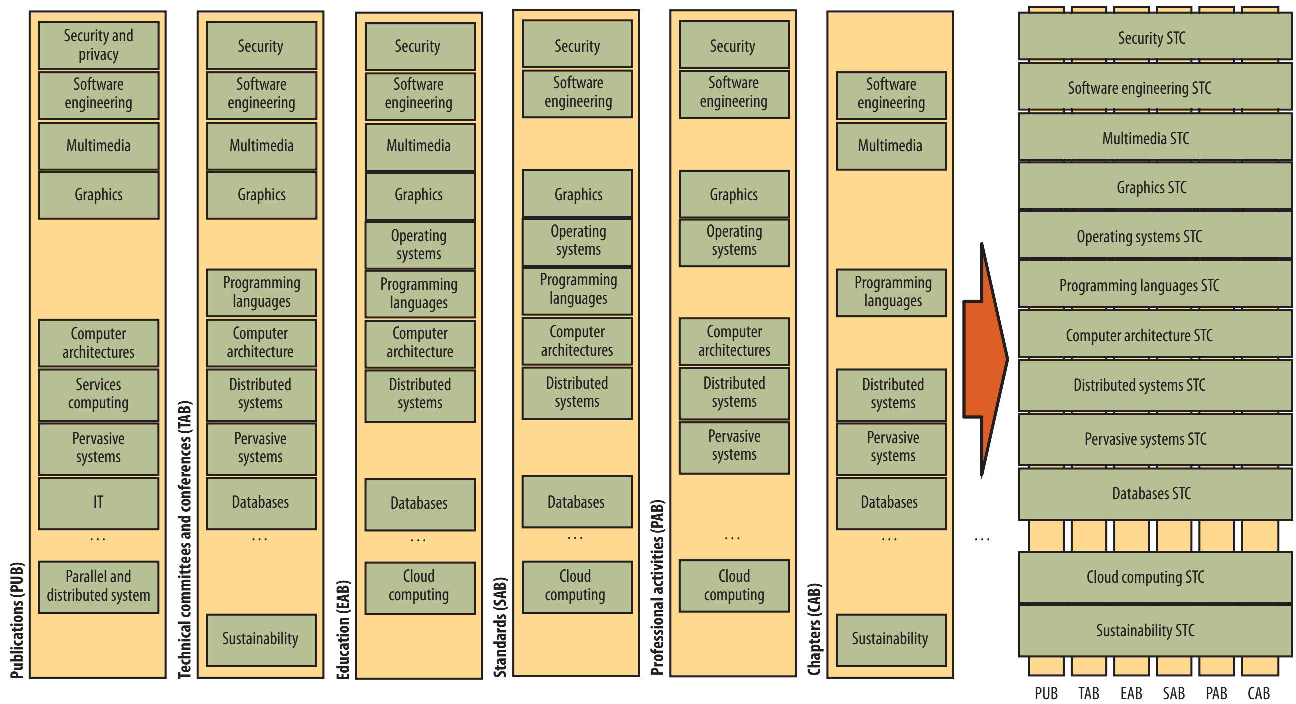

Figure 1. Enhancing organizational structures with technical focus. STCs will enable existing organizational structures (on the left) to better address existing topic areas more collaboratively, while also allowing member interests to coalesce around emerging topics more quickly.

The IEEE-CS has been exploring the use of technology to address these problems for several years. For example, in 2007, Sorel Reisman, the current IEEE-CS president, proposed reducing the organization's stovepiped structure and focusing activities around common technical interests (S. Reisman, "Fast-Tracking Content: The Dramatic Evolution of Intellectual Property," *IT Professional*, Sept/Oct 2007, pp. 58-60). In 2010, the IEEE-CS president introduced instant communities, a self-service based on Liferay technology.

STCs derive from and build on this and other past experience in the IEEE-CS and IEEE. As an example, the software engineering community comprises professionals, educators, researchers, students, and business enterprises. We need to focus this community of interest—including both IEEE-CS members and nonmembers—around relevant technology. As Figure 1 shows, publications, conferences, standards, and so on will result from the collaboration of this community of interest rather than vice versa.

The power of social networking combined with crowdsourcing and technology support, such as self-served mailing lists, wikis, multimedia interaction, digital newsletters, and magazines, could create intellectual property in new ways, reduce peer review time, and improve reviewing quality.

SPECIAL TECHNICAL COMMUNITIES

STCs are intended to reach beyond traditional membership and activities, introducing new revenue-generating opportunities, offering new products and services, and enriching professional activities such as sharing best practices. These activities should also benefit the IEEE-CS by improving both its reputation and financial position. STCs will also strengthen governance by allowing members to feel more involved in decision-making processes through the dynamic organizational structure.

The principles behind the STC concept include the following:

- *Elasticity*. STCs will be easy to create, and will grow, contract, and retire as member needs warrant. They will be inclusive of IEEE-CS and IEEE members as well as nonmembers with shared interests, offering different levels of engagement such as newsfeeds, blogs, information exchanges, reviews, newsletters, and virtual conferences.
- *Self-service*. STCs will require little or no staff support; use up-to-date technology; integrate with social networking tools such as blogs, wikis, alerts, and portals; and allow for personalization and presence on social networking sites such as Facebook, LinkedIn, and Twitter.
- *Customization*. STCs will integrate into existing IEEE-CS

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

CHARACTERISICS OF VIRTUAL COMMUNITIES

The IEEE's Agora project identified the basic characteristics of a virtual community. The italic type in this list describes how STCs would address these issues.

A clearly stated purpose
Technical topic behind each STC

Clearly identified leadership
Identify capable chairs for each STC

Clearly identified opinion leaders
Chair gathers leaders in the area

A pool of dedicated volunteers
Chair identifies STC volunteers

A well-thought-out plan for maintaining the community
Establish an STC operational committee

Availability of seeded and current content
Leverage IEEE-CS CSDL and IEEE Xplore

A well-defined target audience
Identify professionals, academics, and others with related interests

An articulated plan for recruiting members
STCs will reach out opportunistically to the community at large

Regular use of the site by registered users
Focus on ease of use of collaboration tools

intellectual property and contribute to it as well. They will holistically promote education, publications, technical councils, chapters, and standards. STCs will link to other organizations and services relevant to the underlying technology.

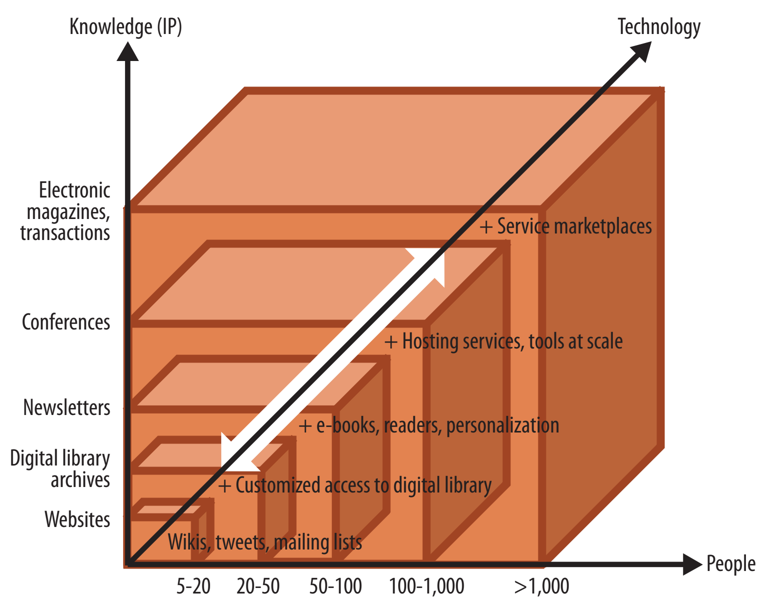

Figure 2. Elasticity of STCs along the dimensions of the number of participants, technology support, and knowledge (intellectual property) produced. As STCs change in size, they will be empowered with more or less technology support and enabled to create IP in a variety of ways. Moving up the dimensions of technology also requires IEEE-CS resources.

STCs will enable IEEE-CS members to keep pace with the world by providing a more dynamic way of creating communities based on new technologies without regard to geographic boundaries. Moreover, because STCs aren't based on a single publication, conference, or standard, their creation will eliminate structural silos. By establishing a Web presence and complementing existing social networking sites, STCs will expose the world to IEEE-CS products, services, and membership benefits, offering members and nonmembers new ways to collaborate.

ORGANIZING AND MANAGING STCS

From 2002 to 2003, the IEEE piloted the use of virtual communities through the Agora project. Several valuable lessons were learned from this experience, as described in the "Characteristics of Virtual Communities" sidebar. Taking into account these lessons and experience from the IEEE-CS's pilot virtual communities, we envision the evolution of STCs from small groups offering lightweight features to much larger communities offering richer feature sets. STCs are more focused communities, some of which will be large, whereas others will be smaller.

The key elasticity concept, enabling flexible growth and scaling down of communities, must be matched by the dynamic management and organization of STCs, avoiding the creation of administrative obstacles. Figure 2 illustrates the envisioned evolution of a new STC. Suppose a small group of IEEE-CS or IEEE members and a few nonmembers gather around an area of interest, for example, open hardware architectures. Numerous existing projects focus on this topic, for example, the Arduino project (www.arduino.cc), which has a well-orga-

Table 1. Current STCs.

Community	Site	Contact
Sustainable computing	http://stc-sustainable-computing.ieee.net	Martin Arlitt and Ishfaq Ahmad
Social networking	http://stc-social-networking.ieee.net	Christian Timmerer and George Thiruvathukal
Systems engineering	http://stc-systems-engineering.ieee.net	Dick Fairley
Gaming	http://stc-dgame.ieee.net	Narisa Chu
Education	http://stc-education.ieee.net	Arnold Pears
Software engineering	http://stc-software-engineering.ieee.net	Paul Croll
Cloud computing	http://stc-cloud-computing.ieee.net	Michael Kozuch and Rick McGeer

nized community. But suppose this group of IEEE-CS members wishes to establish a more formalized community for the purposes of IP capture and broader exposure. A core group of 5 to 20 individuals could create a new STC site based on a simple Web portal service mechanism. For a group of this size, the IEEE-CS might offer some services such as membership discounts, a given number of free downloads from the Computer Society Digital Library (CSDL), or even discounted library subscriptions.

Over time, if the STC grows, so will the scope of its member benefits. The combinations of potential benefits are limitless as the size of a community evolves. For example, if the community grows into hundreds of members, it might be able to support a newsletter. If it grows even more, offering hosted conferences (virtual or face-to-face), electronic magazines, or e-books would be possible. If the community continues to grow, other benefits and offerings could include courses, webinars, books, and discounts on IP bundles. Discounts or even free access could be given to magazines, conferences, standards, EssentialSets and ReadyNotes (annotated collections of papers), and professional certification testing.

All content for the community would be semantically tagged and cross-linked to existing IP in the CSDL, resulting in substantial customization for each STC. The ability to combine peer reviewing with community feedback will be essential to maintain high scientific and technical standards.

At the same time, communities can downsize as interests and technologies change.

The exact levels of offerings for each phase of STC evolution has yet to be finalized. We're developing the mechanisms and guidelines for the creation, growth, and retirement of STCs. In the meantime, Table 1 lists seven STCs currently in a pilot phase.

BUSINESS MODEL

For long-term sustainability, STCs must follow a sound business model. On the revenue side, it's possible that STCs will provide sufficient indirect benefits, such as increased numbers of new and retained members to cover the operational costs. In addition to membership fees, other income streams are possible by providing subscription products and services to an STC's members. It may also be possible to form financial relationships with e-learning site developers, portal operators, vocational training institutes, certification vendors, publishers, universities, technology providers, and technology integrators who could provide services to members and share revenue with the STC and the IEEE-CS.

The Agora team found that it takes about three months to set up a community. During that period, about 25 hours a week of staff time were required. But once a community is set up, only about one hour per week of staff support would be required to maintain and monitor the community as long as the elements necessary for success are in place. Expenditures might include software platform costs and support as well as the costs associated with staff and institutional overhead. Many, but not all of these costs, can be reduced through the use of volunteers.

Once the guidelines for the creation and maintenance of STCs receive approval from IEEE-CS governance bodies, we expect rapid growth both in the number of STCs and in total membership. By the end of 2011, we plan to grow today's seven active STCs with approximately 50 members to 12 STCs with 300 members. By the end of 2012, we anticipate growth to 100 STCs with 1,000 members, and by the end of 2013, we foresee 200 STCs with 10,000 members. This is reasonable and conservative growth based on IEEE Technical Committees membership.

We believe that STCs are the future of the IEEE-CS. However, we recognize that there are challenges ahead, as the Agora team foreshadowed by concluding that "during the 'pioneer phase,' there certainly must be a degree of advocacy to 'sell' the concept of virtual communities." In time, users will respond and adapt to technological opportunities if they receive sufficient benefits, and similarly, STCs will adapt to meet the needs of IEEE-CS members. Perhaps research related to needs might expedite matters. Like any other initiative, virtual communities must compete for relevance.

To meet this challenge, we vigorously advocate the participation of members and nonmembers in the STCs. If you wish to participate in an existing STC, please contact the leads noted in Table 1. If you wish to start a new STC, please contact the authors. C

Acknowledgments

We acknowledge John Walz, Sorel Reisman, and Jim Isaac for contributing many of the ideas presented here and for their encouragement and support for the development of IEEE-CS STCs. We also collectively thank the many other volunteers and IEEE-CS staff who also have been involved in this initiative.

Dejan Milojicic *is a senior researcher and director of the Open Cirrus Cloud Computing testbed at HP Labs. Contact him at dejan@hpl.hp.com.*

Phil Laplante *is a professor of software engineering at Pennsylvania State University. Contact him at plaplante@psu.edu.*

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

Kuck Receives Computer Pioneer Award

Parallel computing researcher David Kuck has been named winner of the 2011 IEEE Computer Society Computer Pioneer Award "for pioneering parallel architectures including the Illiac IV, the Burroughs BSP, and Cedar; and for revolutionary parallel compiler technology including Parafrase and KAP Tools."

Kuck is one of the most influential figures in parallel computing, especially in creating productivity tools for parallel programming. Over the past four decades, he has influenced a wide range of areas including architecture design and evaluation, compiler technology, programming languages, and algorithms. His influence has been both theoretical and practical.

At the University of Illinois, Urbana-Champaign, Kuck created the Center for Supercomputing Research and Development. CSRD was extraordinarily influential in developing parallel computing technology (from hardware to algorithms) in the era of vectorization and SMPs. As founder and director of Kuck and Associates, and later as an Intel Fellow, his work subsequently influenced industry. Every compiler in use today incorporates techniques he pioneered, targeting parallelism in its many forms and managing locality. In this era of multicore and many-core architectures, as well as petascale supercomputers, this work is now more important than it has ever been. As an outgrowth of his compiler work, Kuck initiated efforts that led to the development of OpenMP, the most common solution for incorporating threads into scientific applications.

David Kuck's Parafrase system was the spark that initiated vectorization research.

Kuck is a member of the National Academy of Engineering and is a fellow of the IEEE, ACM, and the American Association for the Advancement of Science. He has received the Charles Babbage Outstanding Scientist Award and the 1993 IEEE Computer Society/ACM Eckert-Mauchly Award.

COMPUTER PIONEER AWARD

The IEEE Computer Society Computer Pioneer Award was established in 1981 to recognize and honor the vision of those people whose efforts resulted in the creation and continued vitality of the computer industry. The award is presented to outstanding individuals whose main contribution to the concepts and development of the computer field was made at least 15 years earlier. The recognition is engraved on a bronze medal specially struck for the Society.

To learn more about Computer Society awards, including the Computer Pioneer Award, visit www.computer.org/awards. C

VMWare's Greene and Rosenblum Win Computer Entrepreneur Awards

DIANE GREENE

Industry innovator Diane Greene, whose efforts contributed to the birth and development of practical virtualization platforms, was recently named winner of the 2011 IEEE Computer Society Computer Entrepreneur Award.

Greene received the award "for creating a virtualization platform that profoundly revolutionized modern computing."

Greene serves on the board of Intuit, the MIT Corporation, and the Peninsula Open Space Trust. She works actively with several private technology companies and is on the Stanford School of Engineering advisory board. Greene was a cofounder and the CEO of VMware from 1998 to 2008, leading the company through an IPO and to a $2 billion capitalization. Prior to VMware, she was the founding CEO of VXtreme, which was sold to Microsoft in 1997 as the basis for its media player. Earlier, Greene held engineering management and development positions at SGI, Tandem, and Sybase.

Greene's academic honors include mechanical engineering, naval architecture, and computer science degrees from the University of Vermont, Massachusetts Institute of Technology, and the University of California, Berkeley, respectively. Greene also won the 1976 women's national double-handed dinghy championship.

Diane Greene led VMWare for a decade.

Mendel Rosenblum is leading an effort to build a new operating system targeted to large-scale shared-memory multiprocessors.

MENDEL ROSENBLUM

VMware cofounder Mendel Rosenblum recently received the 2011 IEEE Computer Society Entrepreneur Award "for creating a virtualization platform that profoundly revolutionized modern computing."

Rosenblum is an associate professor of computer science and electrical engineering at Stanford University. His research interests include system software, distributed systems, and computer architecture. He has published research in the areas of disk storage management, computer simulation techniques, scalable operating system structure, virtualization, computer security, and mobility. Rosenblum received a BA in mathematics from the University of Virginia and an MS and PhD in computer science from the University of California, Berkeley.

As VMware's chief scientist for the company's first decade, Rosenblum helped design and build virtualization technology for commodity computing platforms. He was the 1992 recipient of the National Science Foundation's National Young Investigator Award and the 1994 recipient of an Alfred P. Sloan Foundation research fellowship. He was a cowinner of the 1992 ACM Doctoral Dissertation Award and the 2002 ACM/SIGOPS Mark Weiser Award for creativity and innovation in operating systems research. In 2009, Rosenblum was a cowinner of the ACM System Software Award. He is a member of the National Academy of Engineering and an ACM Fellow.

COMPUTER ENTREPRENEUR AWARD

The IEEE Computer Society established the Computer Entrepreneur Award in 1982 to recognize and honor the technical managers and entrepreneurial leaders who are responsible for the growth of some segment of the computer industry. The efforts must have taken place more than 15 years earlier, and the industry effects must be generally and openly visible.

To learn more about Computer Society awards, including the Computer Entrepreneur Award, visit www.computer.org/awards.

Dozens of New CSDP and CSDA Holders Named

Thousands of software development practitioners have advanced their career opportunities by qualifying as Certified Software Development Professionals or Certified Software Development Associates through IEEE Computer Society certification. In early 2011, 59 more top computing professionals passed the exacting exams for earning CSDP or CSDA status.

The CSDP is the only software development certification that has all of the components of a professional certification, including exam-based testing to demonstrate mastery of a body of knowledge, extensive experience requirements for the performance of the professional work being certified, and continuing professional education standards as measured and relevant to the BOK. The CSDA credential is a software development certification that is intended for recent software engineering graduates or entry-level software development professionals.

The CSDA and CSDP exam specifications were developed through a job analysis process that provided an industry-accepted, systematic procedure for identifying/validating the performance domain of a job and the knowledge and skills that are necessary for a software development professional to perform his or her job.

New CSDP and CSDA holders are named below.

CSDP

A

Alaa Alwani
Rajesh Ananthanarayanan

B

Jacob W. Beningo
Robert Binder
Gregory Blank
Bradley E. Braun
Dwayne Budzak

C

Anthony Candarini
Michael Cox
Maurice Curtin

D

Anastasia R. Davis
Susan Demkowicz
Keshav Deshpande
Alan Durston

F

Richard Fairley
Susan L. Frank

K

Heinz Kabutz
Yegor V. Kosyanchuk

M

Hemant Mahapatra
Betty J. Mills
Arturo Frappe Munoz
Nick T. Mushovic

P

Dan Pilcher
Timothy Procter

S

Yutaka Sato
Thom Schoeffling
Manojkumar Singal
David Southard

T

Leon H. Tabak
Joe Templin
Douglas Thomas
Michael E. Toler
Judd Trayling

V

Christine Vladic

W

Wendy Wallick
Jeffrey Walters
Michael Werling

CSDA

Helene Crowfoot
Alexander L. Gascoigne
Adam Gwin
David A. Harper
Andrew Hay
Qinan Hu
Lu Li
Hrvoje Maric
John Morrison
Andrew Muyanja
Nigel Noronha
Terry Olson
James O. Onyango
Jefferson Sean Orr
Thomas Owens
Daniel Perret
Michael L Richards
Shaopeng Shi
Chelsea Stenner
Tao Sun
Patrick Un
Hairong Wu

COMPUTER SOCIETY CERTIFICATION

In a world where software is pervasive, the need for skilled, competent, software development professionals is greater than ever. IEEE Computer Society certifications benefit graduates by bridging the gap between education and work requirements and verifying students' understanding of fundamental software development practices. Certification also helps established professionals to confirm proficiency in current software development practices and demonstrate a commitment to the profession. Finally, Computer Society certifications support employers in standardizing software development practices and help them protect their investment in a competent and proficient workforce.

Learn more about Computer Society certification programs at www.computer.org/certification.

B. Ward, Editor; bnward@computer.org

CALLS FOR ARTICLES FOR IEEE CS PUBLICATIONS

IEEE Internet Computing plans a May/June 2012 special issue on infrastructures for online social networking services.

The proliferation of rich social media, online communities, and collectively produced knowledge resources has accelerated the convergence of technological and social networks, resulting in a dynamic ecosystem of online social networking services, environments, and applications. OSN sites' success is reshaping the Internet's structure, design, and utility. Moreover, this trend is creating numerous challenges and opportunities for the development, deployment, management, and operation of scalable, secure, interoperable OSN infrastructures that can sustain a cycle of innovative application development, improved end-user experience, high-quality service provision, privacy protection, and healthy market expansion.

IC's guest editors seek recent research results in systems, software, and services that provide novel ubiquitous, scalable, secure, and trustworthy OSN infrastructures.

Articles are due by **1 September**. Visit www.computer.org/portal/web/computingnow/iccfp3 to view the complete call for papers.

IEEE Intelligent Systems plans a January/February 2012 special issue on human-agent-robot teamwork.

HART has become a widely accepted metaphor for describing the nature of multirobot and multiagent cooperation. By virtue of a largely reusable, explicit, formal model of shared intentions, team members try to manage general responsibilities and commitments to each other in a coherent fashion that both enhances performance and facilitates recovery when unanticipated problems arise. For software agents and robots to participate in teamwork alongside people in carrying out complex real-world tasks, they must have some of the capabilities that enable natural and effective teamwork among groups of people.

The guest editors seek innovative contributions to theories, methods, and tools in support of mixed groups of humans, agents, and robots working together in teams. Multidisciplinary research—combining findings from fields such as computer science, AI, cognitive science, anthropology, social psychology, ergonomics, robotics, organizational psychology, and human technology interaction to address the problem of HART—is strongly encouraged.

Articles are due by **1 July**. Visit www.computer.org/portal/web/computingnow/iscfp1 to view the complete call for papers.

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.

CALENDAR

JULY 2011

11-15 July: ICME 2011, IEEE Int'l Conf. on Multimedia and Expo, Barcelona, Spain; www.icme2011.org

18-22 July: COMPSAC 2011, IEEE Int'l Computer Software and Applications Conf., Munich; http://compsac.cs.iastate.edu/

18-22 July: SAINT 2011, Symp. on Applications and the Internet (with COMPSAC), Munich; www.saintconference.org

25-27 July: ARITH 2011, Symp. on Computer Arithmetic, Tubingen, Germany; www.ac.usc.es/arith20

AUGUST 2011

1 Aug: Hot Chips 2011, 23rd Symposium on High-Performance Chips, Palo Alto, California; www.hotchips.org

15-18 Aug: ICGSE 2011, IEEE Int'l Conf. on Global Software Eng., Helsinki, Finland; http://icgse2011.soberit.hut.fi

25-27 Aug: NCA 2011, 10th IEEE Int'l Symp. on Network Computing and Applications, Cambridge, Mass.; www.ieee-nca.org

SEPTEMBER 2011

11-14 Sept: ASAP 2011, 22nd IEEE

Published by the IEEE Computer Society 0018-9162/11/$26.00 © 2011 IEEE

EVENTS IN 2011

July

11-15 ....ICME 2011
18-22 ....COMPSAC 2011
18-22 ....SAINT 2011
25-27 ....ARITH 2011

August

1 .... Hot Chips 2011
15-18 .... ICGSE 2011
25-27 ....NCA 2011

September

11-14 ....ASAP 2011
14-16 ....ECOWS 2011
18-23 ....ITC 2011
25 ....ICSM 2011
26-30 ....Cluster 2011

ASE 2011

The 2011 IEEE/ACM International Conference on Automated Software Engineering brings together researchers and practitioners to share ideas on the foundations, techniques, tools, and applications of automated software engineering. ASE 2011 will include technical papers, experience papers, invited keynotes, tutorials, workshops, tool demonstrations, and a doctoral symposium.

Conference organizers have solicited original technical papers and practical experience reports that address topics including modeling language semantics, open systems development, product line architectures, and specification languages.

ASE 2011 is sponsored by the IEEE Computer Society and the ACM. The conference takes place 6-10 November in Lawrence, Kansas. Visit www.continuinged.ku.edu/programs/ase for complete conference details.

Int'l Conf. on Application-Specific Systems, Architectures, and Processors, Santa Monica, California; http://asap-conference.org

14-16 Sept: ECOWS 2011, 9th IEEE European Conf. on Web Services, Lugano, Switzerland; http://ecows2011.inf.usi.ch

18-23 Sept: ITC 2011, Int'l Test Conf., Anaheim, California; www.itctestweek.org

25 Sept-1 Oct: ICSM 2011, 27th IEEE Int'l Conf. on Software Maintenance, Williamsburg, Virginia; www.cs.wm.edu/icsm2011

26-30 Sept: Cluster 2011, IEEE Int'l Conf. on Cluster Computing, Austin, Texas; www.tacc.utexas.edu/ieee/2011/index.php

OCTOBER 2011

4-7 Oct: LCN 2011, 36th IEEE Conf. on Local Computer Networks, Bonn, Germany; www.ieeelcn.org/index.html

4-7 Oct: SRDS 2011, 30th IEEE Int'l Symp. on Reliable Distributed Systems, Madrid, Spain; http://lsd.ls.fi.upm.es/srds2011

10-14 Oct: PACT 2011, 20th Int'l Conf. on Parallel Architectures and Compilation Techniques, Galveston, Texas; http://pactconf.org

22-25 Oct: FOCS 2011, 52nd IEEE Symp. on Foundations of Computer Science, Palm Springs, California; www.cs.ucr.edu/~marek/FOCS11

NOVEMBER 2011

6-10 Nov: ASE 2011, 26th IEEE/ACM Int'l Conf. on Automated Software Eng., Lawrence, Kansas; www.continuinged.ku.edu/programs/ase

6-13 Nov: ICCV 2011, 13th Int'l Conf. on Computer Vision, Barcelona, Spain; www.iccv2011.org

7-9 Nov: ICTAI 2011, 23rd IEEE Int'l Conf. on Tools with Artificial Intelligence, Boca Raton, Florida; www.cse.fau.edu/ictai2011

12-18 Nov: SC 2011, ACM/IEEE Int'l Conf. for High Performance Computing, Networking, Storage, and Analysis, Seattle; http://sc11.supercomputing.org

Harmonizing Quality Assurance Processes and Product Characteristics

César Pardo, *Kybele Consulting Colombia*
Francisco J. Pino, *University of Cauca, Colombia*
Félix García, *University of Castilla-La Mancha, Spain*
Mario Piattini, *Alarcos Quality Center, Spain*

Adapting software quality assurance processes to product requirements will result in greater product quality and compliance, and thus increased customer satisfaction.

Today, information technology is characterized by rapid innovation and intense competition. Consequently, software organizations must be prepared to develop products that meet quality requirements at a low cost and in a short time.

According to Alfonso Fuggetta, "the quality of a software product heavily depends on the people, organization, and procedures used to create and deliver it" ("Software Process: A Roadmap," *Proc. Conf. Future Software Eng.*, ACM Press, 2000, pp. 25-34). Numerous models have been developed to assure the quality of software processes including CMMI (Capability Maturity Model Integration), the ISO 12207 standard, COBIT (Control Objectives for Information and related Technology), and ITIL (Information Technology Infrastructure Library). Other models aim to guarantee the quality of software products including FURPS (Functionality, Usability, Reliability, Performance, Supportability) and the ISO 9126 and ISO 25000 standards.

Despite the wide range of available models to define processes and develop better products, companies typically implement a process without considering its effect on the quality characteristics of the product they're developing. However, it's impossible to select the proper reference model without understanding how it impacts the product.

THE PRODUCT-PROCESS DILEMMA

Researchers have attempted to ascertain the business value of adopting various CMMI-defined practices, with special emphasis on improving certain characteristics of software product quality. For example, correlating the maturity of the process a company uses with the number of errors in its products detected by customers has made it possible to determine the impact that the number of code lines (software size) has on the number of defects.

Most studies agree that institutionalizing process improvements improves software product quality. However, even companies that use the same development process at the same level of maturity don't achieve the same levels of quality.

What is a major quality requirement in one company may not be so in another. "An organization that produces mission-critical software considers reliability to be the most important factor," Noushin Ashafi observed, "while portability may be a necessity for an organization that produces a software product for a variety of platforms" ("The Impact of Software Process Improvement on Quality: In Theory and Practice," *Information & Management*, Aug. 2003, pp. 677-690). This makes it problematic to extend, say, ISO 9126 to specific software environments.

Because there's no agreed-upon way to validate the effect of process improvements on product quality characteristics, companies don't harmonize the models they use, exponentially increasing development time, cost, and resources required.

AN INTEGRATED APPROACH

The solution to this dilemma is to analyze the various process reference models available and select

Published by the IEEE Computer Society

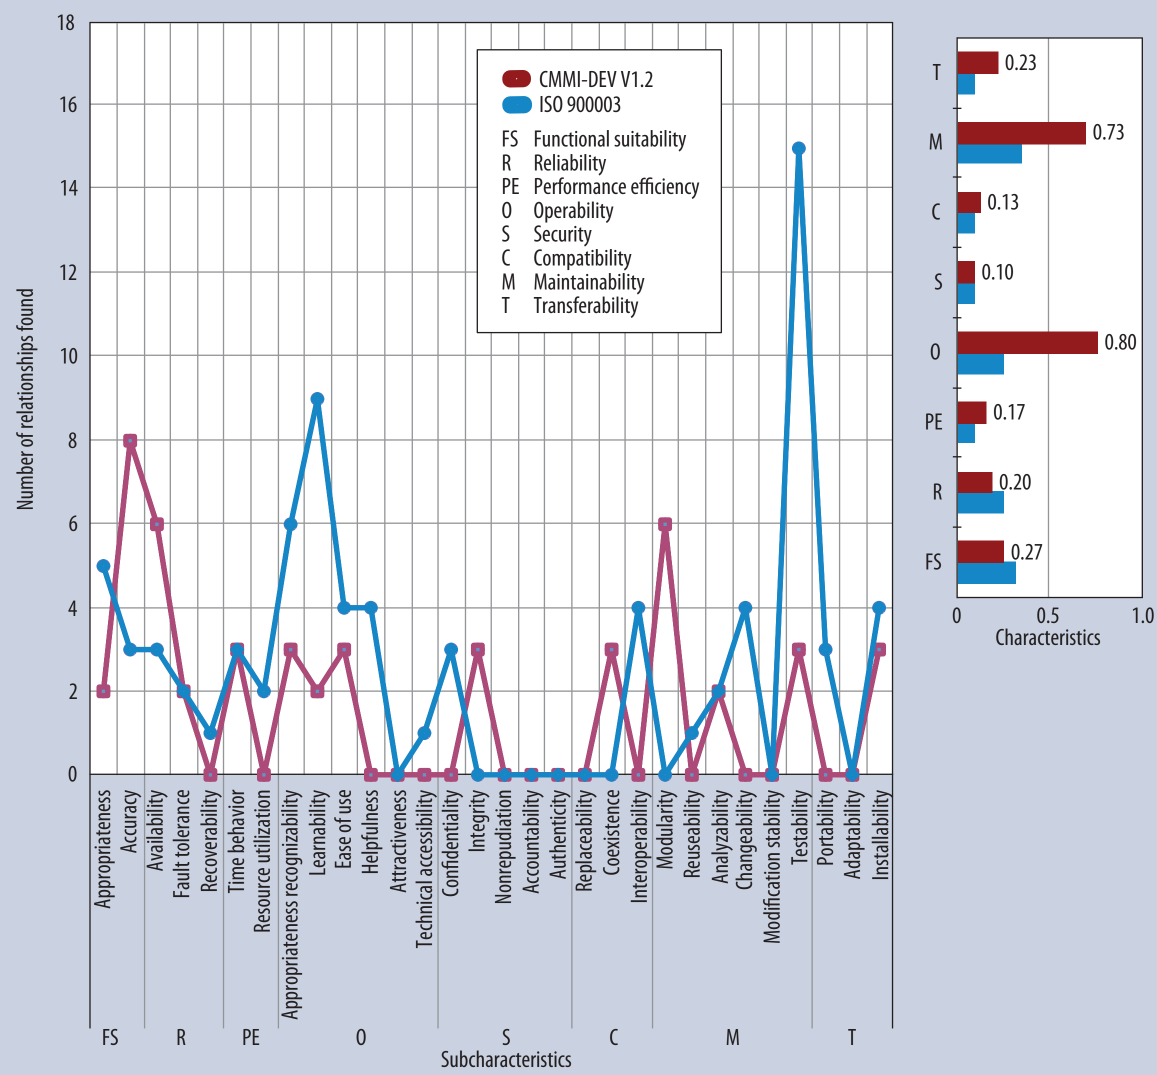

Figure 1. Relationship between ISO 25010 quality characteristics, CMMI-DEV V1.2 specific goals, and ISO 90003 clauses.

those processes most likely to support a product's particular quality requirements. Adapting processes from different models to product requirements will result in greater product quality and compliance, and thus increased customer satisfaction. It will also optimize the resources used for process improvement by linking those resources to the product's quality requirements. Although this integrated approach will likely increase short-term costs, it will ultimately cost less—and provide more value—than implementing separate process and product approaches.

No single process reference model is comprehensive enough to support all of a given product's quality characteristics. For example, specific quality requirements defined in ISO 25010 might best be supported by CMMI-DEV V1.2 or ISO 90003. As Figure 1 shows, CMMI-specific goals (SGs) provide more support for modularity, availability and accuracy, while ISO clauses offer greater support for installability, portability, testability, changeability, interoperability, helpfulness, ease of use, learnability, appropriateness recognizability, and appropriateness.

Special care must be taken when integrating the syntax and semantics of each model's descriptions, which can be at different levels of abstraction and detail. Consider, for example, the quality characteristic of time behavior, which ISO 25010 defines as "the degree to which the software product provides appropriate response and processing times, as well as throughput rates when performing its function, under stated conditions."

CMMI supports this requirement with three SGs. According to SG 3, "The verification methods, procedures, and criteria are used to verify the selected work products and any associated maintenance, training, and support services using the appropriate verification environment. Verification activities should be performed throughout the product life cycle. Practices related to peer reviews as a specific verification method are included in specific goal 2 …"

ISO 90003 has three clauses related to time behavior. According to clause 7.3.6.1, which deals with verification, "Before offering the product for customer acceptance, the organization should validate the operation of the product in accordance with its specified intended use, under conditions similar to the application environment, as specified in the contract. Any differences between the validation environment and the actual application environment, and the risks associated with such differences, should be identified and justified as early …"

CMMI provides greater implicit support for verifying time behavior than ISO, but an organization that carefully integrates CMMI and ISO processes could achieve even better support by tailoring the processes to the specific product requirements.

More work is needed to harmonize two software process engineering approaches that are widely used but rarely integrated in a coherent way: process quality assurance and product quality. To help guide organizations in applying processes from different models to specific quality requirements, researchers must create a formal framework that matches quality characteristics to process descriptions and outlines associated measures to ensure compliance. ■

César Pardo *is a member of Kybele Consulting Colombia, a consulting and research firm in Spinoff, Colombia, focusing on software process improvement, as well as a PhD student in computer science at the University of Castilla-La Mancha. Contact him at cpardo@unicauca.edu.co.*

Francisco J. Pino *is an associate professor at the University of Cauca, Colombia, as well as a PhD in computer science at the University of Castilla-La Mancha and founding partner of Kybele Consulting Colombia. Contact him at fjpino@unicauca.edu.co.*

Félix García *is an associate professor of computer science at the University of Castilla-La Mancha. Contact him at felix.garcia@uclm.es.*

Mario Piattini *is a director at Alarcos Quality Center, a spin-off from the Alarcos Research Group at the University of Castilla-La Mancha. Contact him at mario.piattini@uclm.es.*

Editor: Mike Hinchey, Lero—The Irish Software Engineering Research Centre; mike.hinchey@lero.ie

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

ENTERTAINMENT COMPUTING

The Kinect Digital Out-of-Box Experience

John Solaro, *Microsoft*

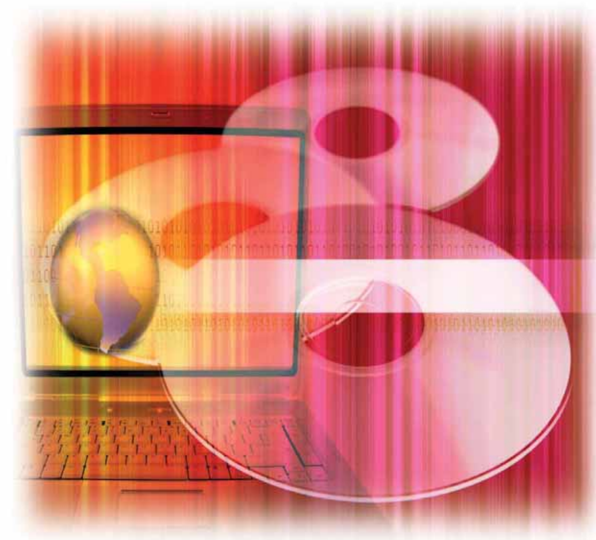

A great OOBE is vital to the success of a consumer electronic device, both for experience and business reasons.

Most modern consumer electronics have one thing in common: you have to set them up in some way before you can use them. Even the most basic gadgets require you to at least turn them on, charge a battery, set up some hardware, or provide basic settings information to let them know you're ready to start using them. Connected software and websites also usually need some sort of account information to get them set up.

As the products we use for fun or function become increasingly powerful, there is often more a user must do to get things working. But the goal should be for technology to get out of the way and let the users enjoy the product they just purchased. A great user experience, after all, starts with an easy setup—which is often much harder to create than you might think.

THE OOBE CHALLENGE

Well before the public knew about Kinect for Xbox 360, then referred to as Microsoft's Project Natal, a small group of Xbox developers started working on creating a simple and fun way for users to set up their Kinect and understand how to use it.

Our focus was on the digital out-of-box experience (OOBE). This is most simply defined as the setup process that happens the moment a user presses the power button on an electronic device for the first time.

A great OOBE is vital to the success of a consumer electronic device, both for experience and business reasons. On the experience side, product teams often forget that the OOBE is the only process or feature that 100 percent of their customers are guaranteed to encounter. It's also the very first experience that users have with a product. Because it sets the tone for all other uses of a product, the first experience needs to be a good one. On the business front, a successful OOBE means fewer returned products or calls to customer support, as well as increased sales based on customers' word-of-mouth recommendations.

Luckily, the leadership working on Kinect realized early on that the OOBE was a critical piece of the overall program. Our team was given plenty of leeway and partnered closely with many others to tackle this important but tough project.

The Kinect OOBE was particularly challenging for several reasons. First, this wasn't just the first time users were experiencing a new product—it was the first time they were experiencing an entirely new technology and way of interacting with a consumer electronics device. Second, the timeframes were aggressive, and the goals were huge. As with other motion-sensing devices, the tie between software, hardware, and the actual physical space in people's homes all play a role in the overall set up and experience, creating many variables outside the product team's control.

Given the nature of the problem we faced with building the Kinect OOBE, it was critical to think through all aspects of the product and bring together representation from design, user research, producers, engineering, game development, brand, marketing, business, product support, and more so they could ask the right questions across the board. How could moments in television commercials reinforce the right way to use Kinect speech commands? How could a video on Xbox.com help educate users so they don't have to call in to customer

Figure 1. Kinect OOBE. (a) The instructions break up the OOBE into sections so that users can better understand them. (b) Users are rewarded when they finish a section. (c) Short animation helps users understand some key points in the initial setup.

support? More than ever, we had to think end to end.

Over time, a core team formed and stayed with the project through its entire life cycle. Many others contributed throughout the process, and our core group eventually collaborated with the proverbial cast of thousands. All the knowledge, learning, facts, and opinions from the various voices ultimately coalesced through iteration and hard work into the OOBE that ultimately shipped.

PROJECT GOALS

Our OOBE's goals were simple and logical, but tough to accomplish:

- customers must be able to easily and correctly set up their Kinect sensor and Xbox 360;
- customers should be playing their first Kinect game within minutes of pressing the power button for the first time;
- we should beat acceptable rates for product returns and customer support calls; and
- customers should understand how to use their Kinect and Xbox 360 and have fun in the process.

All these goals were also focused on keeping Xbox's traditional base of hardcore gamers happy while still welcoming novice or more family-oriented customers to use our ever-broadening entertainment platform.

AN ITERATIVE PROCESS

Kinect's magic was something that no customer had ever experienced before. To most users, it sounded like science fiction. Every user test taught us more about what we didn't know and what we needed to learn. We had to create new concepts to help solve problems we hadn't even imagined going into the process. Everything was new territory, and every fact we know now had to be learned as we went along.

To make the project work even more interesting, other designers and developers were building the new dashboard experience and core Kinect feature set at the same time we were building the OOBE, meaning that the core things we had to teach users about Kinect during the OOBE were constantly changing. An all too familiar quote during the project was, "We're flying a plane while we're building it."

Every time there was new understanding of the technology or the way our customers used it, we iterated on the experience. Our designer, researcher, and producer core team iterated on user flows, wireframes, high-fidelity visual comps, and motion-study animated movies to help tell the story in collaboration with our engineering and business partners. Engineering used these tools to create prototypes of working code that we could test with users to gain new insights and then repeat the process again.

Many tough calls, changes, and long nights were needed to make the project a reality. We ultimately redesigned and rebuilt the OOBE many times to get to the final design.

In the end, the way we completed the project on time and met our goals was much less rocket science than the technology we were helping customers adopt. As with all great experiences, collaboration, user testing, iteration, and an amazing staff were important factors. That, and a lot of humility, so that when it was clear our customers needed a change to the product, we could make it right.

Thanks to the hard work of our team and the feedback from our users, the Kinect OOBE was ultimately a success. There is no doubt in my mind that we could have done even better, but it's great to see that customers can successfully set up their Kinects and Xbox 360s and, most importantly, have fun playing the amazing games while experiencing our system's various entertainment capabilities. The images in Figure 1 show some of the interfaces that our OOBE incorporates to help users get started with their systems.

During a summer internship in college, I worked at a factory outside Chicago where I tried to help make its assembly lines more efficient. In my first week, my new boss had me spend my time working on the assembly line instead of learning the software, getting to know other people, and so on. It was the best way I could have possibly learned how things really worked there.

The Kinect OOBE project was the same way. I learned more about our customers and this product by working on the OOBE than any other way I could imagine.

Building the Kinect OOBE has been the biggest learning experience of my career so far. The story I have told here is admittedly design-focused, given that's the team I work on. Many other people and teams would have different stories to tell and thoughts share about how things went on this project. But I believe that's the nature of any project of this scale and importance.

In my opinion, anyone focused on building great software experiences should work on an OOBE project at some point. The knowledge and understanding gained when you focus on customers using your product for the very first time provides incredible insight into a much broader set of user experiences. C

John Solaro is an executive producer at Microsoft, working on the Interactive Entertainment Business Design Team focused on Xbox. Contact him at jsolaro@microsoft.com.

Editor: Kelvin Sung, Computing and Software Systems, University of Washington, Bothell; ksung@u.washington.edu

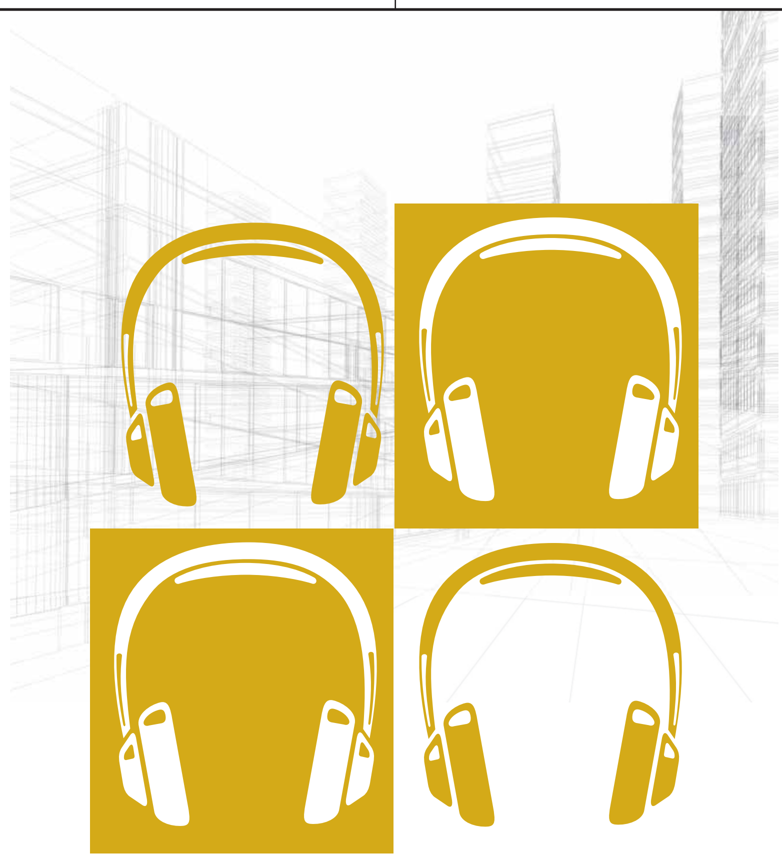

Computer | Previous Page | Contents | Zoom in | Zoom out | Front Cover | Search Issue | Next Page | qMags

Ubiquitous Data Collection: Rethinking Privacy Debates

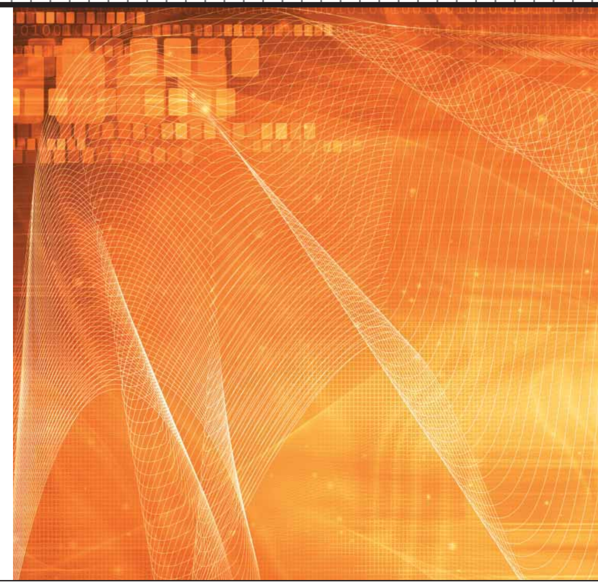

Dan Breznitz, Michael Murphree, and Seymour Goodman
Georgia Institute of Technology

A discussion about the ubiquitous collection, dissemination, and processing of data requires a comprehensive perspective of the risks involved.

Historically, many dystopian novels pictured a future in which an all-encompassing police state uses sophisticated surveillance and technology to gather data on every citizen. Under Big Brother's ever-watchful eye, citizens are reduced to slaves of the state, distrustful of one another and fearful of speaking their minds. Individuals are assumed to have no right to privacy.

THE CURRENT DEBATE

Current debates over online privacy and the right to the security and inviolability of individual data appear to be informed by such visions. Proposed legislation and policy actions from congressional representatives Jackie Speier of California and Joe Barton of Texas emphasize the problems arising from the apparent loss of privacy. Meanwhile, lobbying activity by the Interactive Advertising Bureau and similar groups on behalf of data harvesters attempts to assuage fears of privacy violations.

These debates aren't uniquely American. The European Union has enacted a series of laws designed to preempt certain categories of privacy invasion such as forbidding the storage, transmission, or processing of certain types of data outside sovereign country borders. Cultural traditions in Japan and elsewhere create different types of online privacy and commercial dynamics, which in turn shape the way privacy is interpreted and addressed.

Unlike dystopian predictions of an authoritarian government-created surveillance society, the most active protagonist today is private industry. The fact that private industry does most of the surreptitious gathering, processing, sale, and then use of data may be at least as threatening to a viable democratic society's personal liberty and core values as the traditional fears of a state panopticon.

REDEFINING THE PROBLEM

Regardless of the actor, however, the emphasis in debates about the merits or threats from ubiquitous data collection falls on an individual's right to privacy.

But there's a glaring oversight in arguments that place so much emphasis on a narrow definition of privacy. To illustrate, in the past, an employee who had a dispute with an employer would write a letter to management or the human resources department. The issue would remain between employer and employee and could be settled as such. If a record was kept, it would remain with the employer.

Today, electronic communication and data storage mean that there's not only a permanent record, but one that's relatively open to those technology-savvy enough to access these resources. In the modern world, the record of an employment dispute can result in permanent blacklisting, making it difficult for a job seeker to find work. This is a much greater problem than just the violation of the individual's privacy.

The searchability of massive amounts of data creates another great problem. As so much data is available on the Internet, those interested can use sophisticated ICT tools to coax out relationships in the data. Using the results of these unmonitored and largely unproven analytics, companies—most notably insurers and banks—are beginning to rate potential clients without actually making contact with them. These ratings, based on data that may or may not

Published by the IEEE Computer Society

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

actually predict what it's supposed to, might sharply limit options available to people whose lives are analyzed. Life insurers have begun using social network analysis—such as Facebook friends—to determine a person's likely life expectancy, whether or not these online personalities have any bearing on actual behavior or risk. Thus, the real risk is not the loss of privacy per se, but the ubiquity of data collection, both passive and active, online and offline, combined with powerful ICT tools that allow its manipulation and usage for supposedly predictive purposes. This activity is already occurring, largely without public debate on its merit and influences.

A 2002 film, *Minority Report*, considered a world in which these trends reached their most disturbing conclusion. In the movie, the use of data collection for predictive purposes was designed to prevent crime by anticipating when and where a person would act, allowing for an arrest in advance. Such an extension, although seemingly far-fetched, violates the core Western legal principle of due process under law. However, it's not that great of an extension to see the potential threats to individual security and sovereignty from a society and economy that makes its decisions about individuals based on their online personas and digital signatures.

EXAMINING THE ISSUES

The first issue at stake is the more narrowly defined concern over privacy. This has been extensively researched and is increasingly part of the public understanding and debate. The second issue is the matter of ubiquitous collection of data without our conscious knowledge or approval. This is the area currently receiving much attention in the media and from policymakers in Washington, DC.

The gathering of extensive personal data on the clicks, selections, browsing habits, and preferences of users has become big business. Buying and selling user preferences and interests enabled the creation of a whole new industry. Such innovation is generally to be welcomed, but too often discussions of this data collection simply focus on whether or not it's a threat to individuals' right to privacy. The problem is that such narrow understanding might lead to regulations that both inhibit innovation and don't solve the core issues. There's a need for a systematic examination and categorization of the risks arising from ubiquitous data collection and tracking.

The real risk is not the loss of privacy per se, but the ubiquity of data collection, both passive and active, online and offline.

The third issue—one largely neglected by existing studies, and probably the issue with the most far-reaching long-term consequences—is that the Internet has enabled a whole-scale transformation in data accessibility. Researchers have always collected many types of data about employment, education, housing, or other related subjects. In the past, access to this data was prohibitively expensive, and search was quite cumbersome. Most importantly, the ability to merge different types of data from different sources was severely limited.

As there were no means for the layman or even most companies to comb through the reams of data, collecting it raised few concerns. This data posed little threat to interpersonal trust or personal freedom. Today, however, both companies and determined individuals can easily use other information to search and cross-reference these databases. The sources are endless: self-generated, like Facebook and blogs; ostensibly anonymous, like online message boards; surreptitiously collected through cookies; or even passively collected as we go through our daily lives in what we think of as the separate physical world. All of this creates a pervasive and cumulative digital record of our actions without our even being aware of it.

INTEGRATING THE THREAT

These issues are distinct, but they should be tied together and treated holistically.

The first problem is that, as currently used, the term "privacy" is actually quite amorphous. Attempting to pigeonhole these issues by using a narrow definition or else disregarding them undermines efforts to address the potential problems arising from the ubiquitous collection of data. The issue isn't specifically the loss of individual privacy, since much of the data that is of concern has always been collected. However, it's the qualitative and quantitative difference found in the ubiquitous and often unknown gathering of easily accessible data and its use for other purposes that poses a real threat to individual freedom of expression and, by extension, democracy.

If anyone from employers to vengeful neighbors has access to tools once restricted to a governmental security apparatus, it makes it possible to spy on all citizens with ease. Social capital is the raw material of democratic governance and most effective forms of modern capitalism. Without it, our polity and economy will suffer greatly. The fear of upsetting someone, anyone—since all people now have the search power to act—will erode interpersonal trust and thus social capital. Preventing an erosion of social capital is critical to democracy's long-term viability.

A COMPREHENSIVE PRIVACY FRAMEWORK

We argue that there's a need to develop a new framework for the analysis of the questions surrounding ubiquitous data collection and availability. This *comprehensive privacy framework* expands the field of inquiry and debate beyond a narrow definition of privacy. It includes traditional privacy concerns such as the desire to have personal browsing habits and Internet activities kept out of public view. These are the issues that proposed legislation such as the Do Not Track Me Online Act is currently exploring.

More significantly, the comprehensive privacy framework explicitly addresses the matter of ubiquitous data availability. The fact that the Internet enables massive collation and integration of data for examination and categorization of individuals is not widely appreciated or known.

To the extent that there is awareness, there is no framework for integrating these issues and thus determining what the actual implications are for our society, democracy, and economy. Too much of the existing research merely lists illustrative anecdotes that lack systematic categorization and analysis.

UNANSWERED QUESTIONS

To construct such a framework, future research needs to address four major questions.

What is the actual extent to which individuals have no control over the amount and type of data stored about them in the public or searchable domain? It's crucial to know exactly what types of data are being collected and stored, by whom, and in what ways. This will also reveal whether or not the current legal and regulatory framework addresses the threats that ubiquitous data collection pose to social harmony.

What types of data are currently widely available that should be deleted or have restricted access? Once we have an understanding of the types of data collected and the extent to which individuals can still control or anonymize it, it's important to ask what types of data shouldn't be made available over the Internet and thus begin seeking means to reduce our comprehensive online information footprint.

What can be done to correct errors in online data? By its very nature—and ARPA's original design intent—online data is highly survivable. Once created, it can be replicated and stored infinitely. Errors are similarly stored and replicated. The ubiquity of online data, and the ease of using IT search tools, makes it relatively easy to spot errors. However, how is it possible to ensure that errors have been corrected in all sources? When online data is used for predictive purposes, the persistence of errors is a grave concern.

What alternatives should individuals have to prevent a loss of social capital? Both Congress and the US Department of Commerce are considering proposals that address this question. Should individuals be able to opt out of ubiquitous data gathering? If so, in what ways? How can policymakers fairly balance the interests of individuals who don't want to be spied upon with business interests that rely upon creating, trolling, or preserving online data in perpetuity?

It's time for our privacy debates to reflect the fact that the age of ubiquitous data collection is already here. **C**

Dan Breznitz *is an associate professor in the Sam Nunn School of International Affairs and the College of Management at the Georgia Institute of Technology. Contact him at tbvb@gatech.edu.*

Michael Murphree *is a PhD student in the Sam Nunn School of International Affairs at the Georgia Institute of Technology. Contact him at michael.murphree@gatech.edu.*

Seymour Goodman *is a professor in the Sam Nunn School of International Affairs and the College of Computing at the Georgia Institute of Technology. Contact him at sy.goodman@cc.gatech.edu*

Editor: Jeffrey Voas, National Institute of Standards and Technology; jeffrey.m.voas@gmail.com

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

21st-Century Data Miners Meet 19th-Century Electrical Cables

Cynthia Rudin, Rebecca J. Passonneau, and Axinia Radeva, *Columbia University*

Steve Ierome and Delfina F. Isaac,
Consolidated Edison of New York

Researchers can repurpose even extremely raw historical data for use in prediction.

Electrical grid reliability will be a key issue as peak demand for electricity continues to increase. Grids will need to accommodate a growing population, more high-tech appliances, distributed power generation, and, soon, a large fleet of electric vehicles.

Considering the future of smart grids in the world's major urban centers, it's tempting to picture completely remade grids of shiny new copper wires, each implanted with smart status monitors. This is very unlikely, however, because the existing electrical infrastructure is enormous. New York City alone has more than 94,000 miles of underground cable, enough to wrap around the Earth three and a half times. There is simply too much cable to replace or individually monitor—we're not even close to having truly smart grids.

Are there ways we can make an energy grid smarter without using monitors to assist with reliability? Perhaps there are ways we can learn from reliability failures in the past in order to help maintain reliability and safety in the future. In a sense, can we build a "historically conscious" smarter grid? Our team of scientists at Columbia University and engineers at Consolidated Edison (Con Edison), New York City's power utility, set out to answer a version of this question.

Specifically, we sought to determine whether Con Edison data regarding past failures on the city's low-voltage grid—manhole fires, explosions, smoking manholes, and burnouts—could be used to predict, and thus prevent, future events. Our goal was to create a list of manholes (which serve as access points to the underground electrical grid), ranked from most to least vulnerable, that could support the company's inspection and repair programs, which improve public safety and energy reliability.

PROCESSING RAW DATA

Many electrical grids are very old, and cables from the Thomas Edison era are still reliably carrying current. New York City has the world's oldest grid, and we computed that over 5 percent of Manhattan's low-voltage underground cables were installed before 1930. Con Edison's databanks started in the 1880s, but certainly those historical data weren't created for the purpose of predictive modeling. The historical data are extremely raw and very noisy.

The tasks of matching up failure events to the manholes they refer to and matching manholes to the cables they contain are complicated, mainly because of the various means by which Con Edison has recorded data over the years. Our cable data comes from the company's accounting department, past event records come from the emergency control systems database, and manhole location, inspection, and other data come from different Con Edison sources.

We're trying to predict "serious events," but when viewed through the lens of historical records, it's not always clear whether a past event was serious. Con Edison records events through "trouble tickets." These shorthand notes taken by dispatchers

0018-9162/11/$26.00 © 2011 IEEE

```
FIRE DEPT.REPORTS;CONDITION ORANGE F/O 1655 WEST END AVE.
01/26/00 11:54 MDEPICA DISPATCHED BY 71122
01/26/00 12:19 FERRARO REPORTS: REL. FD. TBL/H M-493784
F/O 1655 WEST END AVE... FOUND RD SOLID COVER 3' OFF AGAINST
VEHICLE...HOLE IS SMOKING...REQS..FLUSH................
CO = 0PPM -> 1655 WEST END AVE...BASEM'T AREA...........DR
01/26/00 14:36 FERRARO REPORTS: IN M-493784 F/O 1655 WEST ST.
HAVE (1)3-500,2-4/0,COPPERED GOING TO BUS COMP. V-72184...DR
01/26/00 18:00 PICA REPORTS: ========= C.F.R.===============
FROM M-493784 F/O 1655 WEST END AVE.. 3-500,2-4/0,AC,4'53'
TO BUS COMPARTMENT V-72184 F/O 1655 WEST END AVE.........
ALL B/O CLEARED......... ALL CO CLEARED...............DR
01/26/00 18:00 MDEPICA COMPLETE BY 23349
***************ELIN REPORT MADE OUT**********MC
07/29/00 00:20 ACT TRBL CHNGD FROM EDSMHF TO EDSMHX BY 71453
08/01/01 09:54 REFERRED TO: CAI ES0012 FYI BY 01585
```

Figure 1. Excerpt of a Con Edison trouble ticket (edited for anonymity) for a manhole explosion event in 2000.

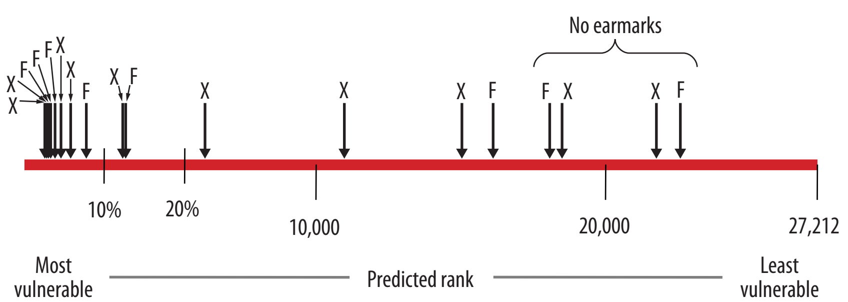

Figure 2. Results from the Bronx blind prediction test. The axis is the predicted rank. Each arrow labeled "X" indicates a 2009 manhole explosion, and each arrow labeled "F" indicates a 2009 manhole fire.

describe the company's responses to events, not the events themselves, and are difficult for a nonexpert to decipher. Figure 1 shows an excerpt of a trouble ticket (edited for anonymity) for a manhole explosion event in 2000.

The text in trouble tickets is very irregular and thus challenging to process in its raw form. There are many spellings of each word—for instance, the term "service box" has at least 38 variations including SB, S, S/B, S.B, S?B, S.B., SBX, S/BX, SB/X, S/XB, /SBX, S.BX, S &BX, S?BX, S BX, S/B/X, S BOX, SVBX, SERV BX, SERV-BOX, SERV/BOX, and SERVICE BOX.

To determine whether an event is serious, we extract information from the ticket text. For example, variations of terms like "smoking" likely connote a serious event, while cable sizes like "2-4/0" or "3-500" or the term "cleared" indicate that an event occurred, since a cable replacement was required.

A NEW FOCUS

With this issue, *Computer* replaces AI Redux with a new column on the use of analytics, data mining, and machine learning to foster discovery in diverse domains. The goal of the Discovery Analytics column is to highlight the pervasive role these technologies now play in science, engineering, humanities, and beyond. Send comments and future article suggestions or proposals to column editor Naren Ramakrishnan, Department of Computer Science, Virginia Tech, Blacksburg, VA; naren@cs.vt.edu.

After processing the tickets, we combine them with the processed cable and manhole location data and several other data sources to obtain an accurate reconstruction of each manhole's decade-long event history, potentially 120-year-old cable history, and inspection results. We then use machine learning algorithms to create a meaningful event prediction model, targeted to predicting failures on individual manholes.

The data processing, and how it's coupled to the statistical model, is the most important step within our knowledge discovery process. Our approach doesn't transform raw data into a form that can be used for many different tasks; generic data transformations wouldn't suffice. Instead, the data processing is specifically geared to our specific modeling task and is driven by the goal of building the predictive model.

Our data processing was done in a transparent way, to ensure that we could explain to Con Edison engineers and managers why a particular manhole was ranked highly, point to the past tickets it was involved in, and describe why we might recommend it for repair or inspection.

PREDICTING SERIOUS EVENTS

We conducted blind tests of the event prediction model in three New York City boroughs: Manhattan, Brooklyn, and the Bronx. For these blind tests, the goal was to predict serious events that happened after the current end date of our database. For instance, when we have data through the end of 2008, we try to predict events in 2009.

Figure 2 shows results from the Bronx test, in which we aimed to predict 2009 events using complete data through 2007 and partial data from 2008. The axis is the predicted rank, from most vulnerable to least vulnerable. Each arrow labeled "X" indicates the predicted rank of a 2009 manhole explosion, and each arrow labeled "F" indicates a 2009 manhole fire. The top 10 percent of manholes (2,721/27,212) on our ranked list contained 44 percent (8/18) of the manholes that experienced a serious event in 2009; the top 20 percent of manholes (5,442/27,212) on the list contained 55 percent (10/18) of the manholes that had a 2009 serious event.

As part of the blind testing, we did case studies of the manholes that experienced a serious event but were at the bottom of the list to understand why or whether our model had failed in those cases. For the case studies, we relied on tools designed to facilitate communication between the Columbia scientists and Con Edison engineers.

One of the tools summarizes all of the raw and processed data that goes into the model regarding a given manhole. This "report card" tool allows us to tell at a glance precisely which trouble tickets, cable records, and other data determine the manhole's ranking. The reports on the bottom four manholes in Figure 2 (labeled "no earmarks") gave no indication that they were vulnerable—these manholes had very few cables and no involvement in past events. Even with a clean and comprehensive database, manhole event prediction can sometimes be a difficult task.

Another tool displays the underground electrical grid's geometry superimposed on Google Earth satellite images, as Figure 3 shows. Each circle in the image is a manhole, and each line between two manholes represents underground low-voltage cables connecting them. This visualization tool enables us to check that our match of cables to manholes is correct and complete and shows the density of trouble tickets in particular neighborhoods.

Figure 3. Image from a visualization tool that displays the underground electrical grid's geometry superimposed on Google Earth satellite images. Each circle is a manhole, and each line between two manholes represents underground low-voltage cables connecting them.

The Columbia-Con Edison manhole event prediction project provides important lessons.

First, researchers can repurpose extremely raw historical data for use in prediction. Databanks similar to Con Edison's are commonly not repurposed, left instead to become "data tombs." But researchers often can analyze and exploit such data to make important contributions—in this case, to devise a better procedure for electrical grid inspection and repair that could improve public safety and energy reliability. The challenge is navigating an ocean of possible data processing tasks, some more rewarding than others, to achieve a more accurate predictive model.

Second, it's possible to maintain future electrical grids using past data. The backbone of our future's smart grids will, for some time, be what exists now. Many power companies face similar challenges, but the usefulness of their historical data will depend on how they collect and store the data, and whether they're willing to expend the resources (and take the risk) to mine it. Our results demonstrate that the investment can be well worth the risk, in that it can help make the grid smarter, safer, and more reliable. ■

***Cynthia Rudin** is an assistant professor in the Operations Research and Statistics group at the MIT Sloan School of Management as well as an adjunct research scientist at Columbia University's Center for Computational Learning Systems. Contact her at rudin@mit.edu.*

***Rebecca J. Passonneau** is a senior research scientist at Columbia University's Center for Computational Learning Systems. Contact her at becky@ccls.columbia.edu.*

***Axinia Radeva** is a staff associate at Columbia University's Center for Computational Learning Systems. Contact her at axinia@ccls.columbia.edu.*

***Steve Ierome** is a distribution engineering secondary system analysis manager at Consolidated Edison of New York. Contact him at ieromes@coned.com.*

***Delfina F. Isaac** is a quality assurance manager at Consolidated Edison of New York. Contact her at isaacd@coned.com.*

Transforming Computer Science Education in High Schools

Jan Cuny, *National Science Foundation*

The CS 10K Project is an ambitious and bold endeavor that represents our best opportunity to transform computing education beginning where we must begin—at the high school level.

Although high school students routinely use technology—social networking, smartphones, searches for information—their interest in courses that could lead to careers in technology is declining.

The 2009 NAEP High School Transcript Study shows that the percentage of US students taking science, technology, engineering, and mathematics courses has increased over the past 20 years for all STEM disciplines except computer science, where participation dropped from 25 percent to 19 percent (http://nces.ed.gov/nationsreportcard/pdf/studies/2011462.pdf). In a 2009 Computer Science Teachers Association survey, high school teachers reported that they were teaching 8 percent fewer CS advanced placement (AP) courses than just two years earlier (http://csta.acm.org/Research/sub/Projects/ResearchFiles/CSTASurvey09CSResults_DCarter.pdf).

These trends are surprising because projected job growth in IT is very strong—much higher than in all other STEM fields combined—and computational competencies are in high demand in many careers.

There's a renewed focus on K-12 STEM education at the national, state, and local levels, but computer science, for the most part, has been left out. How do we change this?

WHAT'S HAPPENING IN US HIGH SCHOOLS?

The term "computer science" has a specific meaning to university faculty and researchers, but I use it here more generally, as it is often used in the high school context, to include the full range of computing and IT-related disciplines. This generation of students can't be just technology consumers. They must also be technology creators, people who can adapt technology to their own ends and who can express themselves computationally. Thus, all students, whether they're future software engineers and IT innovators or biologists, chemists, engineers, journalists, historians, or artists, will need to know foundational computing concepts.

Currently, rigorous computer science is taught in too few schools by too few certified teachers, and underrepresented groups aren't participating. Consider the following:

- In 41 states, computer science doesn't count as a core course for high school students—that is, because it isn't classified as mathematics or science, it doesn't fulfill graduation requirements (http://csta.acm.org).
- Although some schools do offer CS as a college preparatory elective, most offer it only as a career and technical education (CTE) option.
- High school computer science courses often cover only literacy—keyboarding and the use of software packages.
- Only 14 states have adopted significant standards for computer science education in secondary schools.
- In many school districts, the only rigorous computing course that college-bound students can take is AP CS, but it's offered in fewer than 10 percent of US high schools.
- In 2010, only 14,517 high school seniors took the AP CS test, while 194,784 took the AP calculus AB test, 134,871 took the AP biology test, and 109,609 took the statistics test (www.collegeboard.org).
- CS had the worst gender balance of any of the AP tests. In 2010, only 19 percent of the CS test takers were women, while the percentages for calculus, biology, and statistics were 49 percent, 58 percent, and 51 percent, respectively.

0018-9162/11/$26.00 © 2011 IEEE

Published by the IEEE Computer Society

During their high school years, students develop problem-solving skills and make initial postsecondary educational choices. Engagement programs to pique and sustain interest in computing in the early grades will be ineffective if those students don't have opportunities to study computing again during their four years of high school. Likewise, enhancing college computing programs will have little impact unless more students are motivated and prepared to enroll in them.

WHAT SHOULD WE DO?

The CISE Directorate of the National Science Foundation (NSF) intends to serve as a catalyst for an ambitious national effort to strengthen high school computing education. The CS 10K Project will support the development of effective new high school curricula and efforts to prepare 10,000 teachers to teach that material in 10,000 high schools by 2015.

The NSF will support research into how students learn computing concepts, and it will assist in the development and evaluation of a range of curricula, course materials, and approaches to teacher preparation and support.

The new curricula will include introductory (or pre-AP) courses as well as a proposed, entirely new AP CS course that the College Board is developing. The curricula will be rigorous and engaging; in addition to imparting computational thinking skills, they will expose students to the breadth of application and the "magic" of computing. The introductory courses will be designed for both college preparatory and career and technical education students.

The proposed new AP CS course—called CS Principles—will be designed both for students with interests in CS and for those with more general interests in science, engineering, and the humanities (http://csprinciples.org). The College Board will also continue to offer the existing AP CS test.

WHY SO MUCH FOCUS ON AP?

Education in the US is decentralized, with curriculum determined at the state and local school district levels. That makes it difficult to effect the widespread change that's needed.

AP courses are attractive to schools because they increase their academic profile, to students because of their potential to strengthen transcripts, to parents because they can reduce college expenses, and to college admissions officers because they are known quantities.

Although it does include programming, the CS Principles course isn't programming-centric.

The AP imprimatur will thus greatly facilitate adoption of the CS Principles course. More importantly, it's our only single point of national leverage, and it carries with it assurance of fidelity of replication.

SIGNS OF PROGRESS

Curriculum development for the new courses is well under way. With NSF funding, a College Board Commission, working with a distinguished advisory committee, has designed the framework for the CS Principles course. Although it does include programming, the course isn't programming-centric. Instead, it focuses on the underlying principles of computation including problem solving, abstraction, algorithms, data and knowledge creation, and programming.

The course covers the limitations of computation and exposes students to both its breadth of application and related issues of society, culture, and ethics. It's intended to be inspiring, to emphasize creativity, and to be relevant to a diverse group of students. The framework has recently been sent to colleges for vetting, and the resulting comments and suggestions will be incorporated into a final version of the framework this summer.

Progress has also been made in developing curricula that implement the framework. During the 2010-2011 academic year, CS Principles courses were piloted at five universities (AP courses are, after all, college-level courses): Metropolitan State College of Denver; the University of California, Berkeley; the University of North Carolina Charlotte; the University of California, San Diego; and the University of Washington. These courses were different, even in the languages they used: Scratch, BYOB Scratch, Alice, and Processing.

During the 2011-2012 academic year, we expect that additional pilots will be run at 20 to 25 universities and 20 to 40 high schools. SIGCSE, the ACM Education Board and Council, and 100 universities and colleges have endorsed the CS Principles course.

The curriculum for the pre-AP courses will be more flexible, allowing teachers to define their class locally. To assist, the NSF will support the development of a curricular framework, instructional materials, and a set of exemplars. One of those exemplars will be the new Exploring Computer Science (ECS) course, also developed with NSF funding. ECS was piloted this year in more than 20 schools in Los Angeles, San Jose, and Oakland. Its complete curriculum is available on the CSTA website (http://csta.acm.org).

WHAT ARE THE CHALLENGES?

Curriculum isn't our biggest challenge. The biggest challenge will be developing effective teacher preparation and support, and scaling it to reach 10,000 teachers. Few schools today have teachers with any formal

CS training. The computing community must launch an unprecedented effort to prepare teachers, working with in-service as well as pre-service teachers, and in both traditional and alternative certification programs. We'll need to pair face-to-face training with extensive, state-of-the-art online support that includes curricula, instructional materials, assessments, and social networking.

The NSF intends to catalyze this effort through its Computing Education for the 21st Century (CE21) program. To scale the effort to 10,000 schools, though, we'll need more than just NSF support. Success will require establishing a public/private partnership to provide the significant additional funding needed to reach 10,000 teachers. Advocacy groups will also need to address state and local regulations on teacher credentialing, curriculum standards, and the placement of computing within graduation requirements.

The magnitude of this project is daunting. To be successful, we'll need to enlist the entire computing community in the effort—K-12 teachers, university faculty, and undergraduate and graduate students in service learning and outreach programs, as well as professionals serving as citizen scientists.

The CS 10K Project is ambitious and bold, but the situation is dire and not getting any better. CS 10K is our best shot at transforming computing education. Its success will be critical to our nation's ability to innovate and compete in the global 21st-century economy.

Jan Cuny *is a CISE program officer, National Science Foundation. Contact her at jcuny@nsf.gov.*

Editor: Ann E.K. Sobel, Department of Computer Science and Software Engineering, Miami University; sobelae@muohio.edu

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

Continued from page 112

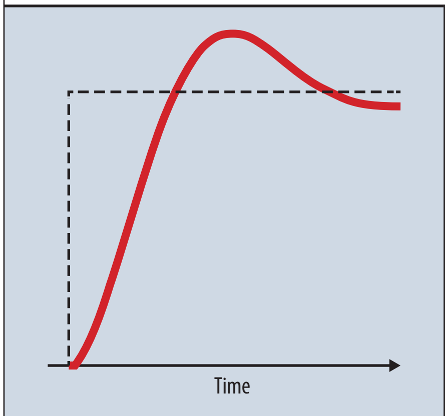

Figure 1. Step response for a soft system.

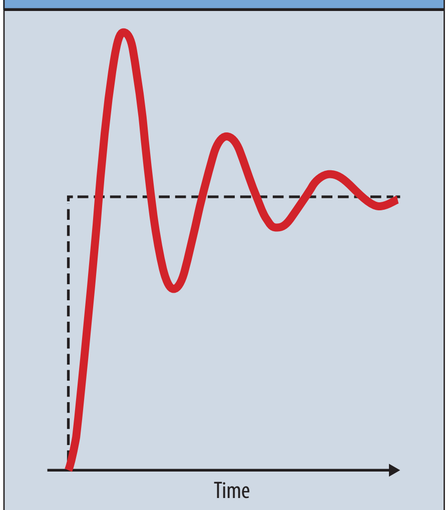

Figure 2. Step response for a hard system.

Of course, in any project, decisions are made to go ahead at a particular time, which usually means deciding to get the new system in place in its entirety by some specific date—in effect, a step function. In big projects, this is often followed by missed targets and cost overruns, in effect, a step response.

One way to moderate step responses is incremental development, that is, to go for a sequence of small steps comprising what is sometimes called a staircase function. This approach would fully develop and install the most important subset of what's needed, then undertake a sequence of projects that add capabilities a few at a time. Later projects would benefit both from the buildup of skills and from the evaluation and experience of preceding projects.

A staged approach wouldn't necessarily mean a succession of systems if the first and succeeding systems were designed to make it easy to add blocks of program code. This used to be done by putting dummy blocks into the system where additions were intended, and I would expect this would still be easy to do with modern development tools.

The developmental increments could be quite small. For example, developers made 6,000 changes to Google's search engine program in 2010, of which 500 were made permanent (tinyurl.com/MagGgTw). Surely, incrementality along these lines could be adopted much more widely, though it would need to be based on a system designed from the start to take small changes safely and easily.

Another approach is to design a system comprising subsystems that interact through a simple data interface (The Profession, June 2006, pp. 98-100). This composite approach is compatible with the incremental approach; it's possible that both are widely used, but much of what I read suggests otherwise (tinyurl.com/WpMyki). A large part of the problem seems to be that failing to continually upgrade an existing system while it's being used forces the huge step of undertaking the development of a completely new system (tinyurl.com/MAgVcPl).

Computing professionals faced with undertaking a very large project are typically hampered by the decisions being made by top-level managers whose immediate subordinates don't understand the difficulties involved. They see digital technology as a universal management panacea. When failure looms or hits, they blame the computer. But companies still undertake big projects, even when doubts are openly expressed (tinyurl.com/EcTFHlCr). Professionals should strive to get the doubts up toward the top and push for the adoption of better approaches.

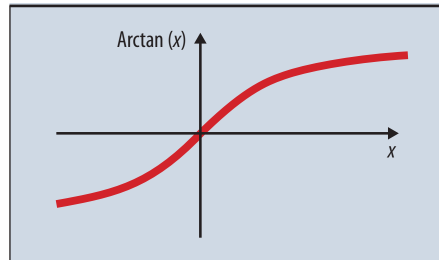

Figure 3. The arctangent function.

SOCIAL RESPONSIBILITY

Beyond the design and implementation of digital systems, the computing professional has a responsibility to consider the effects, present and potential, of such systems on users and on society at large. Because many of these are like step inputs to society, they can have a significant effect.

Consider the so-called flash crash of 2010 (tinyurl.com/WpFlCr). This was a record-breaking stock market plunge apparently caused by the use of computer programs to automatically implement high-frequency trading. "Circuit breakers" have since been installed in stock exchange programs to block rapid changes. However, flash crashes have occurred since then in the currency and commodity markets (tinyurl.com/WSJCmCr).

Clearly, there are social issues here that are relevant to the computing profession. Should automatic trading be legal? If not, how is it to be blocked? Given the ubiquity of computers in financial matters, should all financial transactions be moderated? Are market manipulators using computers to provoke and exploit profitable price fluctuations? How are these to be detected and blocked?

Another computer-related social issue is joblessness, which continues to increase in many countries. The problem is particularly severe for unskilled workers (tinyurl.com/

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

EcDcWkMn). Arguably, this is because technology, particularly digital technology, is being used to an increasing extent for basic manufacturing and clerical tasks. Outsourcing to countries with low-paid workers, often made possible by the Internet, is arguably another factor.

For jobless people to be more employable, they need to acquire more skills. This effectively would require a huge investment in teachers and equipment, but the investment could be made much more effective by using computers to impart the basic skills that will underpin the acquisition of higher-level skills (The Profession, Mar. 2008, pp. 102-104).

But there are more subtle ways in which computers can be exploited for community benefit, ways that would actually use the computer to avoid administrative stepping.

For example, income taxation is often progressive, meaning that the level of income is divided into a few brackets or ranges like a staircase function, and each bracket has a progressively higher tax rate. This often leads to taxpayers going to considerable lengths to get into a lower bracket.

Simple computation could be used to work out a tax rate that moves higher smoothly. An obviously suitable function for this is the arctangent or inverse tangent function of trigonometry (tinyurl.com/WpInTrFn). Subprograms for calculating it are widely available, and it's supported on typical scientific calculators, where it appears as tan^{-1}.

As Figure 3 shows, the horizontal axis would be used for the value of taxable income, and the vertical axis for the taxation rate. The vertical axis can be used directly for the taxation rate because the curve approaches the value of one, that is, 100 percent, asymptotically. With the taxation rate calculation given as $tan^{-1}(ax - b)$, the value of a can be used to adjust the curve's slope and the value of b can be used to adjust the point of zero taxation, that is, to shift the curve's crossover point. A negative taxation rate could be used to calculate a welfare rate.

Smooth changes in income tax rates would take some of the tension out of the process and allow governments to make regular adjustments to suit their budgets. To avoid step changes, smooth taxation would need to be introduced in stages—for example, by having it apply to progressively higher fractions of income during the changeover.

The kind of measures described here to avoid sharp changes in society are speculative, but they could be implemented. Doing so would require the computing profession's leadership and cooperation. Getting at least some of these measures accepted would require the profession's active promotion and support. I hope it happens soon. ■

Neville Holmes is an honorary research associate at the University of Tasmania's School of Computing and Information Systems. Contact him at neville.holmes@utas.edu.au.

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

Computing and the Step Function

Neville Holmes, *University of Tasmania*

Rapid change is dangerous socially and professionally.

One aspect of the engineering education that I began back in the early 1950s seemed very strange at the time. It was the serious consideration of the step function, which was unlike anything treated in high school mathematics because it was discontinuous.

While its relevance to the design of mechanical and electrical machinery soon became obvious, it has recently occurred to me that the step function is just as relevant to the *use* of machinery, in particular, computers.

THE STEP FUNCTION

The step function of my memory is much like what is now called the Heaviside step function (tinyurl.com/WpHvStFn), yielding only two values. Representing it as $y = S(x)$, if the step takes place when $x = x_0$, then one value of y, y', is delivered when $x < x_0$, and the other, y'', when $x > x_0$. For the Heaviside unit step function, x_0, $y' = 0$ and $y'' = 1$. For the perhaps more familiar step function, *signum*, $x_0 = 0$, $y' = -1$, and $y'' = 1$ (tinyurl.com/WpSgnm).

The point of greatest interest is when $x = x_0$. By definition, signum has $y = 0$ when $x = x_0$. In plain language, the signum of a negative value is –1; of a positive value, 1; and of zero, 0. For the Heaviside function, the corresponding value is a matter of convention.

This way of handling critical points is all very well for mathematicians, but for engineers, the application sets the scene. The step function is relevant when a system is required to take a step input and convert it effectively to its output—ideally, a step output.

The engineering problem is that x is usually a value of time, thus the step function requires an instantaneous change from y' to y''. In spatial terms, this is an indefinitely large acceleration. In mechanical terms, it means an indefinitely large force. These are unnatural and, in physical systems at least, neither an input nor an output can be exactly stepped.

The engineering challenge is to take an input that's practically a step function and give an output that can be used as a step function, sometimes after cleaning it up. In digital circuitry, the square waves that encode bits in transmission are ideally a sequence of step functions, and the bit rate is limited by the distortion during transmission.

The output of a system with a step input is called a step response (tinyurl.com/WpStpRsp). Figures 1 and 2 show two different responses to the step function (in dashes). Figure 1 shows the response of a system that absorbs sudden change and moves slowly to the target level, and Figure 2 shows the response of a system that changes fairly quickly once the input steps up.

Both figures show the *overshoot phenomenon*, that is, an output level beyond the final stable level. The significant aspect of this is that the more responsive the system, the greater the overshoot. If the overshoot is too great, the system will fail. This contrast is similar to the difference between dropping a wineglass on a carpet and dropping it on a tiled floor. That's why I've used *hard* and *soft* as labels for the systems in the first two figures.

Socially as much as electronically, changing too rapidly leads to instability and overreaction. The professional problem is that computers are all too often used as agents of rapid change. The professional responsibility is to prevent, or at least mitigate, any ill effects.

TECHNICAL RESPONSIBILITY

Governments and other large organizations seem to have an urge to undertake grand projects. When the project is based on the use of computers, failure seems to occur all too frequently (tinyurl.com/MAgPrOm). This is by no means a recent phenomenon, as I saw only too plainly back in the 1970s.

Continued on page 110

Published by the IEEE Computer Society

CALL FOR PAPERS

ICDCS 2012

The 32nd International Conference on

Distributed Computing Systems

Macau, China
June 18-21, 2012

www.icdcs-2012.org

IEEE computer society
IEEE

Sponsored by

The IEEE Computer Society Technical Committee on Distributed Processing

ORGANIZING & PROGRAM COMMITTEES

General Co-Chairs
Wei Zhao, Univ. of Macau, Macau, China
Ten H. Lai, Ohio State University, USA

Program Co-Chairs
Xavier Defago, JAIST, Japan
Wang-Chien Lee, Penn State University, USA

Program Vice Chairs

Algorithms and Theory
Marcos Aguilera, Microsoft Research, USA
Toshimitsu Masuzawa, Univ. of Osaka, Japan

Cloud Computing Systems
Li Cui, Chinese Academy of Sciences, China
Ling Liu, Georgia Institute of Technology, USA

Data Management and Data Centers
Masaru Kitsuregawa, Univ. of Tokyo, Japan
Jianliang Xu, Baptist Univ., Hong Kong, China

Distributed OS and Middleware
Pascal Felber, Univ. of Neuchatel, Switzerland
Gilles Muller, INRIA/University Pierre and Marie Curie, France

Fault Tolerance and Dependability
Matti Hiltunen, AT&T - Research, USA
Vana Kalogeraki, Athens University of Economics and Business, Greece

Network/Web/P2P Protocols and Applications
John Lui, Chinese Univ., Hong Kong, China
David Yau, Purdue University, USA

Security and Privacy
Tei-Wei Kuo, National Taiwan Univ., Taiwan
Felix Wu, Univ. of California at Davis, USA

Wireless, Mobile, Sensor, and Ubiquitous Computing
Sajal Das, Univ. of Texas at Arlington, USA
Yunhao Liu, HKUST, Hong Kong, China

Program Committee Members
(Please see the conference web page)

Workshops Co-Chairs
Jiannong Cao, Poly. Univ., Hong Kong, China
Min Song, NSF/Old Dominion University, USA

Awards Co-Chairs
Wu-chi Feng, Portland State University, USA
Zhiwei Xu, Chinese Academy of Sci., China

Local Arrangements Co-Chairs
Samson Chang, Univ. of Macau, Macau, China
Biao Chen, Univ. of Macau, Macau, China

Registration Co-Chairs
Weijia Jia, City University, Hong Kong, China
Carrie Stein, Ohio State University, USA

Finance Co-Chairs
Huan Li, Beihang University, China
Dong Xuan, Ohio State University, USA

Publicity Co-Chairs
Ai Chen, Shenzhen Inst. of Adv. Tech., China
Jiming Chen, Zhejiang University, China

Publication Co-Chairs
Min-Te Sun, National Central Univ., Taiwan
Jianjia Wu, Univ. of Macau, Macau, China

International Liaison Co-Chairs
Haruhisa Ichikawa, University of Electro-Communications, Japan
Michel Raynal, IRISA, France
John Stankovic, University of Virginia, USA

TCDP Chair
Jie Wu, Temple University, USA

Steering Committee Chair
Xiaodong Zhang, Ohio State Univ., USA

SCOPE

The conference provides a forum for scientists and engineers in academia, industry and government to present their latest research findings in any aspects of distributed computing. Topics of particular interest include, but are not limited to:

- Algorithms and Theory
- Cloud Computing Systems
- Data Management and Data Centers
- Distributed OS and Middleware
- Fault Tolerance and Dependability
- Network/Web/P2P Protocols and Applications
- Security and Privacy
- Wireless, Mobile, Sensor, and Ubiquitous Computing

WORKSHOPS

Workshops will be held in conjunction with the conference. Workshop proposals should be submitted to Workshops Co-Chairs, Prof. Jiannong Cao (csjcao@comp.polyu.edu.hk) or Prof. Min Song (msong@odu.edu) by **September 1, 2011**. Notification of acceptance will be made by October 1, 2011. Please see the conference web page for details.

PAPER SUBMISSION

Form of Manuscript: All paper submissions should follow the IEEE 8.5" x 11" Two-Column Format. Each submission is limited to 10 pages. If the paper is accepted for publication, up to 2 over-length pages may be purchased for the final camera-ready version.

Electronic Submission: Submissions will be handled via the conference web page.

The proceedings of the conference and the workshops will be published in CDs by the IEEE-CS Conference Publishing Services (CPS).

IMPORTANT DEADLINES

Abstract Registration	**November 1, 2011**
Paper Submission	**November 8, 2011**
Author Notification	February 7, 2012
Final Manuscript	March 14, 2012

For further information, please contact Technical Program Co-Chairs, Prof. Xavier Defago at defago@jaist.ac.jp or Prof. Wang-Chien Lee at wlee@cse.psu.edu

Seapine's Quality-Centric ALM Solutions

Seapine's quality-centric ALM solutions bring extensive enterprise capabilities to your software development organization. Leverage your existing infrastructure investment with solutions that are easy to deploy and manage. Scalable, feature-rich ALM tools integrate seamlessly to provide a quality management system across all phases of your development lifecycle.

Seapine Supports Your Agile Transition

Seapine's ALM tools are cross-platform and built on flexible architectures using open standards to support today's popular development methodologies and best practices. TestTrack 2010.1 supports your transition to agile development while giving your organization the flexibility to manage traditional development processes in parallel.

New agile capabilities enable you to:

- Manage backlog and releases
- Perform sprint planning and tracking
- Conduct user acceptance tests and define doneness criteria
- Support multiple products, projects, and distributed teams
- Generate burn down, velocity, task board, time tracking, live charts, and more!

Learn why thousands of leading enterprises worldwide rely on Seapine's quality-centric ALM solutions to manage their software development lifecycle.

Visit **www.seapine.com/myenterprisealm**
or call **1-888-683-6456** today.

Seapine Software™